High-order coverage of smoothed Bayesian bootstrap intervals for population quantiles

David M. Kaplan* Lonnie Hofmann†

November 4, 2019

Abstract

Using fractional order statistics, we characterize the high-order frequentist coverage probability of smoothed and unsmoothed Bayesian bootstrap credible intervals for population quantiles. The original (Rubin, 1981) unsmoothed intervals have $O(n^{-1/2})$ coverage error, whereas intervals based on the smoothed Bayesian bootstrap of Banks (1988) have much smaller $O(n^{-3/2}[\log(n)]^3)$ coverage error. No smoothing parameter is required. In special cases, the smoothed intervals are exact. Simulations illustrate our results.

Keywords and phrases: continuity correction, credible intervals, fractional order statistics
MSC 2000 subject classifications: Primary 62G30, secondary 62G15

1 Introduction

We study the frequentist coverage accuracy of certain Bayesian bootstrap credible intervals for population quantiles. Specifically, we consider credible intervals based on the Bayesian bootstrap (BB) of Rubin (1981) and the continuity-corrected Bayesian bootstrap (CCBB) of Banks (1988). With the CCBB, the posterior probability that the quantile of interest is above a given order statistic is identical to the corresponding frequentist coverage probability, even in small samples. Even when the interval endpoint is a "fractional" order statistic (i.e., a weighted average of consecutive order statistics; see (13)), the CCBB and frequentist probabilities remain identical if the underlying distribution is uniform. More generally, as long as the underlying density is twice differentiable, the CCBB and frequentist probabilities differ by merely $O(n^{-3/2}[\log(n)]^3)$ for central (non-extreme) order statistics. Consequently, CCBB credible intervals are also very accurate frequentist confidence intervals. They are

*Department of Economics, University of Missouri. Email: kaplandm@missouri.edu
†Department of Economics, University of Missouri. Email: lrh6df@mail.missouri.edu

also straightforward to compute since the smoothing does not involve any bandwidth or other tuning parameter. Approximate CCBB credible intervals can be computed without simulation, using the beta distribution, with $O(n^{-1})$ error in posterior probability. For the BB without smoothing, the $O(n^{-3/2}[\log(n)]^3)$ coverage error becomes $O(n^{-1/2})$. For comparison, the standard (multinomial) percentile bootstrap and related variants also have $O(n^{-1/2})$ coverage error; e.g., see Hall and Martin (1989).

Despite the high-order accuracy, our results' value is more theoretical than directly practical. For practical frequentist inference, calibrated versions of Hutson's (1999) confidence intervals from Goldman and Kaplan (2017, Thm. 4(iii)) achieve the same $O(n^{-3/2}[\log(n)]^3)$ coverage error as CCBB, without requiring simulation. Similarly, the calibrated versions of Beran and Hall's (1993) confidence intervals from Ho and Lee (2005) have $O(n^{-3/2})$ coverage error and relatively simple computation. Further, Goldman and Kaplan (2017, Thm. 4) show the leading n^{-1} term in the coverage error of Hutson's (1999) intervals to be positive, which in simulations seems helpful especially for more extreme quantiles; see Section 5. For objective Bayesian inference, given the simplicity of BB and CCBB, it is only marginally helpful that the intervals of Goldman and Kaplan (2017) are good approximations of BB and CCBB intervals. However, it is interesting that the Bayesian–frequentist agreement extends to higher-order terms and is even exact for CCBB in special cases.

Further, the CCBB's high-order accuracy here raises the question of whether such accuracy extends to more complex objects of interest and models. Although the approach of Hutson (1999) and Goldman and Kaplan (2017) extends to interquantile ranges and certain other quantile-based objects of interest (Goldman and Kaplan, 2018), the CCBB applies much more broadly. Perhaps relatedly, for quantile regression, Hahn (1997) advocates the BB since it produces a smoother distribution than the multinomial bootstrap, but without any claims of high-order accuracy (only first-order).

Our results can also be summarized as approximations based on the beta distribution CDF. The BB posterior probability of an interval containing the quantile of interest can always be written exactly in terms of the beta CDF. When the interval endpoints are order statistics, the CCBB posterior probability can also be written exactly with the beta CDF. Further, the CCBB posterior probability exactly matches the frequentist coverage probability, whereas the BB does not. When the interval endpoints are between order statistics, CCBB requires simulation, whereas BB does not, but CCBB achieves better coverage accuracy. The CCBB can also be approximated by a beta CDF (without simulation), which leads to theoretical coverage accuracy better than BB but worse than simulated CCBB (if simulation error is negligible), although not for extreme quantiles (see Section 5). The beta CDF for the CCBB approximation is the same beta CDF that facilitates frequentist confidence intervals with higher-order accuracy (Goldman and Kaplan, 2017).

The literature contains a few variations of the Bayesian bootstrap relevant to our paper. Rubin (1981) introduces the technique, which can also be seen as applying an improper Dirichlet process prior to the work of Ferguson (1973). For continuously distributed data, Banks (1988) suggests spreading out the point masses over the intervals between observed sample values, so each distribution function is a linear spline with knots at the sample values. He calls this "histospline smoothing" in allusion to Wahba (1975). This yields a posterior over continuous (instead of discrete) distributions. Meeden (1993) provides an alternative grid-based smoothing procedure that is a proper stepwise Bayesian method, unlike the CCBB. However, setting the grid as the observed sample values yields the CCBB. Meeden (1993) also shows simulation evidence of the benefit of smoothing, but not theoretical justification. We provide such theoretical results.

Motivating our work, Breth (1979) connects nonparametric Bayesian inference on distributions and quantiles with the frequentist sampling distribution of uniform order statistics, but without theoretical results on coverage accuracy. Breth (1979) anticipates the BB as a limiting case of the Dirichlet process model of Ferguson (1973), discussing "the case of a noninformative prior" (p. 641). Breth (1979, p. 642) notes the BB posterior probability of a band for the population CDF can be computed from the frequentist sampling distribution of uniform order statistics for a sample with one fewer observation ("$m = n - 1$"). (The CCBB, which had not yet been proposed, "corrects" this to be the same sample size, $m = n$.) Breth (1979) also mentions credible intervals for quantiles (Section 3) using order statistics as endpoints, but without results on coverage accuracy, which is our main contribution.

Section 2 introduces the credible intervals. Sections 3 and 4 contain theoretical results. Section 5 illustrates the theory with simulations. Section 6 suggests directions for future research. Appendix A contains proofs.

Notation Vectors are bold; e.g., the sample is $\mathbf{Y} = (Y_1, \ldots, Y_n)$. Also, $\mathbb{1}\{\cdot\}$ is the indicator function, $\mathrm{Q}_\tau(Y)$ is the τ-quantile of random variable Y whose CDF is $F_Y(\cdot)$, $\mathrm{P}(\cdot)$ is probability, $\pi_u(\cdot \mid \mathbf{Y})$ is BB posterior probability, and $\pi_s(\cdot \mid \mathbf{Y})$ is CCBB posterior probability. Acronyms used include those for Bayesian bootstrap (BB), confidence interval (CI), continuity-corrected Bayesian bootstrap (CCBB), coverage probability (CP), coverage probability error (CPE), cumulative distribution function (CDF), probability density function (PDF), and probability mass function (PMF). Let $F_\beta(p; a, b)$ and $f_\beta(p; a, b)$ respectively denote the CDF and PDF of a Beta(a, b) distribution evaluated at p.

2 Setup

Let $\mathbf{Y} = (Y_1, \ldots, Y_n)$ denote the data. For most of the paper, we condition on $\mathbf{Y}$ and treat the Y_i as fixed; without conditioning, the Y_i are treated as random variables. The Y_i are iid realizations of random variable Y with continuous CDF $F_Y(\cdot)$, so the Y_i values are unique (with probability 1). Interest is in $Q_\tau(Y)$, the τ-quantile of Y for some $\tau \in (0, 1)$.

Rubin (1981) considers a posterior distribution over possible CDFs $F_Y(\cdot)$. In analogy to the multinomial bootstrap (Efron, 1979), the Bayesian bootstrap restricts attention to discrete distributions with support $\{Y_i\}_{i=1}^n$. Then, a Dirichlet–multinomial model with improper prior is used to compute the posterior. Although not cited by Rubin (1981), this is equivalent to using the nonparametric model of Ferguson (1973, 1974) with improper Dirichlet process prior.

The Bayesian bootstrap posterior over $F_Y(\cdot)$ can be expressed in terms of a probability mass function (PMF). Let the order statistics (i.e., ordered sample values) be

$$Y_{n:1} < Y_{n:2} < \cdots < Y_{n:n}. \tag{1}$$

All PMFs in the support of the posterior can be written in terms of

$$p_k \equiv \mathrm{P}(Y = Y_{n:k}), \quad k = 1, \ldots, n. \tag{2}$$

The Bayesian bootstrap posterior is

$$(p_1, \ldots, p_n) \mid \mathbf{Y} \sim \mathrm{Dir}(1, \ldots, 1), \tag{3}$$

a Dirichlet distribution with all n parameters equal to 1 (i.e., a uniform distribution over the unit $(n-1)$-simplex). Given $(p_1, \ldots, p_n)$, the corresponding CDF is

$$F_Y(y) = \sum_{i=1}^{n} p_i \, \mathbb{1}\{Y_{n:i} \le y\}, \tag{4}$$

where $\mathbb{1}\{\cdot\}$ is the indicator function.

The CCBB of Banks (1988) has a similar posterior, except probabilities are spread over intervals between order statistics. The n order statistics define $n+1$ intervals. Let $Y_{n:0}$ and $Y_{n:n+1}$ denote the lower and upper bounds of the support of Y; these bounds only affect inference on extreme quantiles, which we do not consider. Instead of (2),

$$p_k \equiv \mathrm{P}(Y_{n:k-1} < Y \le Y_{n:k}), \quad k = 1, \ldots, n+1,$$

and Y is uniformly distributed (flat PDF) within each interval. Similar to (3), the posterior distribution of these interval probabilities is

$$(p_1, \dots, p_{n+1}) \mid \mathbf{Y} \sim \mathrm{Dir}(1, \dots, 1), \tag{5}$$

a Dirichlet distribution with all $n+1$ parameters equal to 1 (i.e., a uniform distribution over the unit n-simplex). Instead of (4), the continuity-corrected (histospline-smoothed) CDF corresponding to $(p_1, \dots, p_{n+1})$ is

$$F_Y(y) = \sum_{i=1}^{n+1} \left[p_i \, \mathbb{1}\{Y_{n:i} \le y\} + p_i \frac{y - Y_{n:i-1}}{Y_{n:i} - Y_{n:i-1}} \, \mathbb{1}\{Y_{n:i} > y > Y_{n:i-1}\} \right]. \tag{6}$$

If $Y_{n:1} \le y \le Y_{n:n}$, then the bounds do not affect $F_Y(y)$. For intuition, evaluating (6) at an order statistic yields

$$F_Y(Y_{n:k}) = \sum_{i=1}^{k} p_i. \tag{7}$$

The CDF in (6) linearly interpolates between such points.

Given (5), distributions of various sums of the p_i are given by Wilks (1962, pp. 236–238). The CCBB posterior distribution of $(p_1, \dots, p_{n+1})$ is the same as the frequentist sampling distribution of standard uniform order statistic "spacings,"

$$(U_{n:1}, U_{n:2} - U_{n:1}, \dots, U_{n:n} - U_{n:n-1}, 1 - U_{n:n}) \sim \mathrm{Dir}(1, \dots, 1), \tag{8}$$

where the $U_{n:k}$ are order statistics based on $U_i \overset{iid}{\sim} \mathrm{Unif}(0,1)$, $i = 1, \dots, n$. This makes sampling from the posterior easy. Given (3), (5), and (8), results from Wilks (1962) give

$$\begin{aligned} \text{BB: } p_1 + \cdots + p_k \mid \mathbf{Y} &\sim \mathrm{Beta}(k, n-k), \\ \text{CCBB: } p_1 + \cdots + p_k \mid \mathbf{Y} &\sim \mathrm{Beta}(k, n+1-k). \end{aligned} \tag{9}$$

Already, (5) and (8) suggest a close connection between the CCBB and frequentist methods. If $F_Y(\cdot)$ is continuous, then the probability integral transform (Fisher, 1932; Neyman, 1937; Pearson, 1933) combined with (8) gives the joint frequentist sampling distribution

$$(F_Y(Y_{n:1}), F_Y(Y_{n:2}) - F_Y(Y_{n:1}), \dots, F_Y(Y_{n:n}) - F_Y(Y_{n:n-1}), 1 - F_Y(Y_{n:n})) \sim \mathrm{Dir}(1, \dots, 1). \tag{10}$$

From the frequentist view, $F_Y(\cdot)$ is fixed, whereas the $Y_{n:k}$ are random. However, (10) is also the CCBB posterior distribution conditional on $\mathbf{Y}$, treating the $Y_{n:k}$ as fixed and $F_Y(\cdot)$ as random. One implication is that uniform bands for $F_Y(\cdot)$ can be constructed in finite samples with the same exact $1 - \alpha$ coverage probability and posterior probability; e.g., see Breth (1979, p. 642) and Goldman and Kaplan (2016, Prop. 10).

Section 3 uses posterior probabilities like $\pi(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y})$. Here, the order statistic $Y_{n:k}$ is fixed (conditioned on), whereas the population quantile $\mathrm{Q}_\tau(Y)$ is random. With continuous $F_Y(\cdot)$, the condition $Y_{n:k} < \mathrm{Q}_\tau(Y)$ is equivalent to $F_Y(Y_{n:k}) < \tau$, and $F_Y(Y_{n:k}) = \sum_{i=1}^k p_i$, whose distribution is (9). We treat the complications from interpolation between order statistics in Section 4.

3 Results for order statistic endpoints

We first state assumptions for analyzing intervals with order statistic endpoints. Section 4 considers more general intervals whose endpoints are fractional order statistics, $Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k})$ for some $\epsilon \in [0, 1)$; this section considers the special case $\epsilon = 0$.

Assumption A1. The Y_i are iid realizations of random variable Y with CDF $F_Y(\cdot)$, and the Y_i are unique over $i = 1, \ldots, n$: $Y_i \neq Y_k$ for any $i \neq k$. Interest is in the τ-quantile, $\mathrm{Q}_\tau(Y) \equiv \inf\{y : F_Y(y) \geq \tau\}$.

Assumption A1 implies the BB and CCBB Dirichlet posteriors in (3) and (5). The uniqueness assumption ensures the CCBB posterior includes only continuous CDFs. If $F_Y(\cdot)$ is continuous, then $(Y_1, \ldots, Y_n)$ are unique with probability 1. Although continuity of $F_Y(\cdot)$ is needed for comparison with frequentist coverage probability, it is not needed (and hence not in A1) for our results expressing BB and CCBB posterior probabilities in terms of the beta CDF.

Some additional notation is helpful. Denote the unsmoothed BB (subscript u) and smoothed CCBB (subscript s) posterior probabilities given data $\mathbf{Y}$ by

$$\pi_u(\cdot \mid \mathbf{Y}), \quad \pi_s(\cdot \mid \mathbf{Y}), \tag{11}$$

based on (3)–(6). Denote the CDF of a Beta(a, b) distribution by $F_\beta(\cdot; a, b)$.

We first characterize the BB and CCBB posterior probabilities of one-sided credible intervals whose lower endpoint is an order statistic. For either $\pi = \pi_s$ or $\pi = \pi_u$, such probabilities are

$$\pi(\mathrm{Q}_\tau(Y) \in [Y_{n:k}, \infty) \mid \mathbf{Y}) = \pi(Y_{n:k} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y}).$$

Although generally such intervals cannot exactly achieve a prespecified nominal level, these results already show the very close connection between the CCBB posterior and frequentist sampling probabilities, as well as the benefit of smoothing. Also, in practice, it may be preferable to have, e.g., an exact 93% interval than an approximate 95% interval.

The unsmoothed part of Theorem 1 is similar to a result from Efron (1982, p. 82, eqn. (10.23)). The proof uses the equivalence $Y_{n:k} < \mathrm{Q}_\tau(Y) \iff \sum_{i=1}^k p_i < \tau$ and (9).

Theorem 1. *Under Assumption A1, given* (3)–(6) *and* (11)*, for* $k = 1, \ldots, n$,

$$\begin{aligned}
\pi_u(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) &= F_\beta(\tau; k, n-k), \\
\pi_u(Y_{n:k} \le \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) &= F_\beta(\tau; k-1, n+1-k), \\
\pi_s(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) = \pi_s(Y_{n:k} \le \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) &= F_\beta(\tau; k, n+1-k).
\end{aligned}$$

If additionally $F_Y(\cdot)$ *is continuous, then* $\pi_s(Y_{n:k} \le \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) = \mathrm{P}(Y_{n:k} \le \mathrm{Q}_\tau(Y))$*, i.e., the CCBB posterior probability equals the frequentist coverage probability.*

Theorem 1 shows the BB and CCBB posterior probabilities of certain credible intervals can be computed analytically, without sampling from the posterior. These extend readily to lower one-sided and two-sided intervals since

$$\begin{aligned}
\pi(Y_{n:k} \ge \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) &= 1 - \pi(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}), \\
\pi(Y_{n:k} \le \mathrm{Q}_\tau(Y) \le Y_{n:r} \mid \mathbf{Y}) &= \pi(Y_{n:k} \le \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) - \pi(Y_{n:r} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}).
\end{aligned} \tag{12}$$

Given (12), the CCBB–frequentist equality in Theorem 1 also holds for two-sided intervals where both endpoints are order statistics. For any interval with order statistic endpoints, the exact, finite-sample coverage probability and CCBB posterior probability are identical and easily calculated with the beta CDF. The BB posterior probability always differs, although the difference's magnitude is small with large n, as formalized in Corollary 1.1.

Theorem 1 can be used to compute the difference between BB and CCBB posterior probabilities of intervals. Specifically, for intervals of the form $[Y_{n:k}, \infty)$ or $(Y_{n:k}, \infty)$, Corollary 1.1 characterizes the difference between BB and CCBB posterior probabilities that $\mathrm{Q}_\tau(Y)$ lies in the interval. Let $f_\beta(\cdot; a, b)$ be the Beta(a, b) PDF.

Corollary 1.1. *Given Theorem 1,*

$$\begin{aligned}
\pi_u(Y_{n:k} \le \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) - \pi_s(Y_{n:k} \le \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) &= \frac{1-\tau}{n} f_\beta(\tau; k, n+1-k), \\
\pi_s(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) - \pi_u(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) &= (\tau/n) f_\beta(\tau; k, n+1-k).
\end{aligned}$$

If $c_1 n \le k \le c_2 n$ *for constants* $0 < c_1 < c_2 < 1$ *as* $n \to \infty$*, then both differences are* $O(n^{-1/2})$*. More generally, if* $c_1 n^r \le k \le c_2 n^r$ *or* $c_1 n^r \le n - k \le c_2 n^r$ *for* $0 < r < 1$*, then both differences are* $O(n^{-r/2})$*. If instead* k *or* $n - k$ *is fixed as* $n \to \infty$*, then the differences are* $O(1)$*.*

Corollary 1.1 shows the benefit of smoothing. For intervals with an order statistic endpoint, the CCBB posterior probability exactly matches the coverage probability (even in finite samples), whereas the BB differs by $O(n^{-1/2})$ or possibly more. If τ and the nominal credible/confidence level are fixed, then asymptotically the endpoint will be a "central order statistic" satisfying the condition $c_1 n \leq k \leq c_2 n$, in which case the BB differs from the CCBB (and frequentist coverage probability) by $O(n^{-1/2})$. If interest is in more extreme quantiles, modeled asymptotically with $\tau \to 0$ or $\tau \to 1$ as $n \to \infty$, then the endpoint of interest would be an "intermediate" or "extreme" order statistic, with the larger $O(n^{-r/2})$ or even $O(1)$ difference between BB and CCBB posterior probabilities. However, these bounds may be too pessimistic since they are computed using the mode of the beta distribution. For the purpose of credible intervals of reasonable nominal levels, it may be possible to derive tighter bounds for the intermediate and extreme order statistics; this is left to future work.

4 Results for fractional order statistics

We now consider "fractional" order statistic endpoints. This is useful when a certain nominal level (like 95%) is desired but not attainable at an integer order statistic. For integer k and interpolation weight $\epsilon \in [0, 1)$, the $(k + \epsilon)$th fractional order statistic is

$$Y_{n:k+\epsilon} \equiv Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}). \tag{13}$$

Figure 1 illustrates the following intuition for the CCBB distribution of $\mathrm{Q}_\tau(Y)$. Given the piecewise linear smoothing of CCBB, the τ-quantile in a given posterior draw is also linearly interpolated. Given a draw of $(p_1, \ldots, p_{n+1})$,

$$\mathrm{Q}_\tau(Y) = Y_{n:c} + \frac{(\tau - \sum_{i=1}^{c} p_i)}{p_{c+1}}(Y_{n:c+1} - Y_{n:c}), \quad c \equiv \max\{d : \sum_{i=1}^{d} p_i \leq \tau\}, \tag{14}$$

i.e., c is the largest integer such that $\sum_{i=1}^{c} p_i \leq \tau$. That is, since the CCBB CDF is linearly interpolated between the points

$$F(Y_{n:c}) = \sum_{i=1}^{c} p_i \text{ and } F(Y_{n:c+1}) = \sum_{i=1}^{c+1} p_i,$$

the CCBB inverse CDF (quantile function) $F_Y^{-1}(\cdot)$ is linearly interpolated between

$$Y_{n:c} = F_Y^{-1}(\sum_{i=1}^{c} p_i) \text{ and } Y_{n:c+1} = F_Y^{-1}(\sum_{i=1}^{c+1} p_i).$$

By inspection, (14) has the same form as (13), with $k = c$ and $\epsilon = (\tau - \sum_{i=1}^{c} p_i)/p_{c+1}$. Equation (14) can also be interpreted as moving along the straight line in the quantile function with slope $(Y_{n:c+1} - Y_{n:c})/p_{c+1}$, starting at value $Y_{n:c}$ and moving to the right by

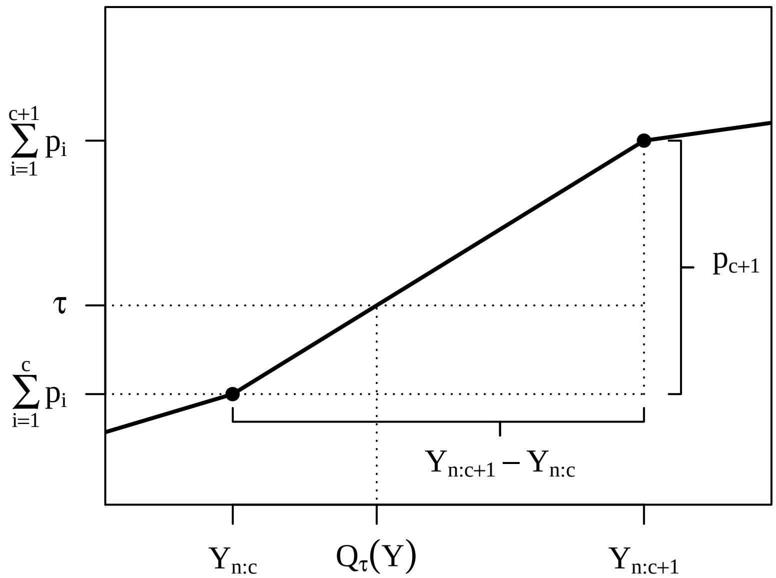

Figure 1: Illustration of (14).

$\tau - \sum_{i=1}^{c} p_i$.

Since the practical value of Theorem 2 is to assess credible/confidence intervals (CIs), we focus specifically on fractional order statistics that would be CI endpoints given τ. For central quantiles (asymptotically fixed τ), Lemma 3 of Goldman and Kaplan (2017) says a $1-\alpha$ one-sided CI lower endpoint has

$$u \equiv (k+\epsilon)/(n+1) = \tau - n^{-1/2} z_{1-\alpha}\sqrt{\tau(1-\tau)} + O(n^{-1}), \tag{15}$$

where $z_{1-\alpha}$ is the $(1-\alpha)$-quantile of the standard normal distribution. Thus, Theorem 2 focuses on $u = \tau - dn^{-1/2}$.

Assumption A2 is required to compare the CCBB posterior probability with high-order coverage probability. It is the same as Assumption A2 of Goldman and Kaplan (2017). It requires enough smoothness to analytically derive the higher-order terms.

Assumption A2. The CDF $F_Y(\cdot)$ in A1 satisfies: i) $F_Y'(\mathrm{Q}_\tau(Y)) > 0$; ii) $F_Y'''(\cdot)$ is continuous in a neighborhood of $\mathrm{Q}_\tau(Y)$.

Theorem 2 shows that the CCBB posterior probability of fractional order statistic intervals is identical to the frequentist coverage probability in certain special cases, and it matches up to a $O(n^{-3/2}[\log(n)]^3)$ difference more generally. Theorem 2 also provides an analytic approximation to the posterior probability based on the beta distribution, up to the same error term. The BB posterior in Theorem 2 follows immediately from Theorem 1 since the BB distributions have no probability strictly between order statistics. As in Section 3, the results for a lower endpoint CI translate to an upper endpoint and two-sided CIs with (12).

Theorem 2. *Let Assumption A1 and* (3)*,* (5)*, and* (13) *hold in each of the following.*

(i) *If* $(k+\epsilon)/(n+1) = \tau - dn^{-1/2}$*, then*

$$\begin{aligned}\pi_u(Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < Q_\tau(Y) \mid \mathbf{Y}) &= F_\beta(\tau; k, n-k),\\ \pi_s(Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < Q_\tau(Y) \mid \mathbf{Y}) &= P(U_{n:k} + \epsilon(U_{n:k+1} - U_{n:k}) < \tau)\\ &= F_\beta(\tau; k+\epsilon, n+1-k-\epsilon)\\ &\quad + n^{-1}\frac{\epsilon(1-\epsilon)}{\tau(1-\tau)}\frac{d}{\sqrt{\tau(1-\tau)}}\phi\Big(d/\sqrt{\tau(1-\tau)}\Big)\\ &\quad + O(n^{-3/2}[\log(n)]^3),\end{aligned}$$

where $\phi(\cdot)$ *is the standard normal PDF and* (8) *is the joint distribution of the* $U_{n:k}$*.*

(ii) *If* $Y \sim$ *Unif*(a,b)*, or if* $F_Y(\cdot)$ *is continuous and* $\epsilon = 0$*, then the CCBB and frequentist probabilities are identical:*

$$\pi_s(Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < Q_\tau(Y) \mid \mathbf{Y}) = P(Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < Q_\tau(Y)).$$

(iii) *If Assumption A2 holds, then*

$$\begin{aligned}\pi_s(Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < Q_\tau(Y) \mid \mathbf{Y}) &= P(Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < Q_\tau(Y))\\ &\quad + O(n^{-3/2}[\log(n)]^3).\end{aligned}$$

Theorem 2 shows the close connection between the CCBB and confidence intervals from Hutson (1999) and Goldman and Kaplan (2017). Hutson (1999) selects $k+\epsilon$ to set $F_\beta(\tau; k+\epsilon, n+1-k-\epsilon) = 1-\alpha$. This is the same beta CDF in Theorem 2(i). Goldman and Kaplan (2017, Thm. 4) show this leaves an error in coverage probability equal to the n^{-1} term and remainder in Theorem 2(i), and they suggest adjusting α to remove the n^{-1} term. Here, the CCBB's posterior probability already captures the n^{-1} term, as well as the beta CDF term. Section 5 illustrates cases where this improved theoretical accuracy translates to better simulation performance, and cases where it does not.

Corollary 2.1 characterizes the difference between the BB and CCBB probabilities with fractional central order statistics, which is $O(n^{-1/2})$ as in Corollary 1.1.

Corollary 2.1. *The difference between the BB and CCBB posterior probabilities in Theorem 2(i) is* $O(n^{-1/2})$ *if* $c_1 n \le k \le c_2 n$*.*

5 Simulations

To illustrate our results, we simulate the coverage probability (CP) of four types of one-sided intervals: BB credible intervals; CCBB credible intervals; "HGK" confidence intervals proposed by Hutson (1999) and studied by Goldman and Kaplan (2017); and "GKc," the calibration proposed by Goldman and Kaplan (2017). Code in R (R Core Team, 2017) to replicate the results is on the second author's website.

Table 1 shows the fractional order statistics chosen by each method in each case. That is, each interval has the form $[Y_{n:k+\epsilon}, \infty)$, with $Y_{n:k+\epsilon} = Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k})$ as in (13), and Table 1 shows the $k + \epsilon$ chosen by each method. For each method, this depends only on $1 - \alpha$ (the nominal level), n, and τ.

Table 1: Fractional order statistic indices $(k + \epsilon)$ for endpoints

n	τ	BB	CCBB	HGK	GKc
18	0.2000	2.00	1.66	1.53	1.69
18	0.2500	2.00	2.24	2.16	2.25
18	0.4978	6.00	5.99	6.00	6.00
18	0.9000	14.00	14.58	14.44	14.77
18	0.9799	17.00	17.00	17.00	17.00
18	0.9900	17.00	17.77	17.49	18.00

Nominal level 0.95; 300,000 posterior draws.

The BB $k+\epsilon$ sets $\epsilon = 0$ and sets k as the largest integer such that $F_\beta(\tau; k-1, n+1-k) \geq 1 - \alpha$, using Theorem 1. This provides the shortest interval with at least $1 - \alpha$ posterior probability: per Theorems 1 and 2, for $\epsilon > 0$, the BB interval $[Y_{n:k+\epsilon}, \infty)$ has the same posterior probability as $[Y_{n:k+1}, \infty)$, but the latter is strictly shorter.

The HGK and GKc $k+\epsilon$ are also based on a beta CDF. HGK sets $1-\alpha = F_\beta(k+\epsilon, n+1-k-\epsilon)$ exactly, per (7) and (8) in Hutson (1999) or (3) and (5) in Goldman and Kaplan (2017). GKc uses the same formula after replacing α with $\alpha + n^{-1}z_{1-\alpha}\phi(z_{1-\alpha})\epsilon(1-\epsilon)/[\tau(1-\tau)]$, per (8) in Goldman and Kaplan (2017). If the resulting k is larger than the HGK k, then GKc sets $\epsilon = 0$.

The CCBB $k+\epsilon$ requires simulation. Per Theorem 2(ii), it sets $1-\alpha = \mathrm{P}(U_{n:k}+\epsilon(U_{n:k+1} - U_{n:k})$, with probability simulated using (8) (see code for details). On a (sub)standard personal computer, it takes a few seconds to compute the CCBB $k+\epsilon$ for all of Table 1. (The Dirichlet draws depend only on n, not τ or α.)

Table 1 indirectly illustrates Theorems 1 and 2. BB is clearly different; besides always having $\epsilon = 0$ (unlike the others), its values are set apart from the other three except in special cases. As Theorem 2 suggests, the other three methods are relatively similar, especially

CCBB and GKc. In special cases, all three are identical (up to CCBB simulation error). In other cases, CCBB is between HGK and GKc, usually closer to GKc.

The different τ are chosen to illustrate different types of situations. Some are chosen so that HGK has $\epsilon = 0$, in which case exact finite-sample CP is attainable (regardless of the underlying distribution). Up to simulation error, the CCBB also has exact CP in these cases, per Theorem 2(ii). In other cases, we choose τ such that HGK has ϵ closer to 0.5 in order to maximize its coverage error, the leading term of which is proportional to $\epsilon(1-\epsilon)$ (Goldman and Kaplan, 2017, Thm. 4(i)). The leading coverage error term is also inversely proportional to $\tau(1-\tau)$, meaning coverage error is larger (all else equal) for τ closer to zero or one, so some such τ are also considered. One particularly extreme τ is chosen to test the limits of our central quantile results; in this case $n\tau > n-1$.

Besides different τ, different distributions F_Y are chosen to illustrate our results. Theorem 2(ii) says CCBB should provide exact finite-sample CP when F_Y is uniform. For comparison, normal and exponential distributions are also used. Although no method's $k+\epsilon$ depends on F_Y, the CP does depend on F_Y if $\epsilon > 0$.

CP is simulated as follows. Samples are drawn iid from the various F_Y distributions shown in Table 2 along with sample size n and quantile index τ. To isolate the effect of F_Y most clearly, the non-uniform samples are based on the standard uniform U_i, with $Y_i = F_Y^{-1}(U_i)$. The nominal (confidence or credibility) level is 0.95.

Table 2 illustrates our theoretical results in the following ways. Note that the CP does not depend on F_Y for intervals with $\epsilon = 0$, which is always true of BB and in special cases true of CCBB, HGK, and GKc (when Table 1 shows an integer value of $k+\epsilon$). As a general summary of Theorem 2, Corollary 2.1, and Goldman and Kaplan's (2017) Theorem 4, for central quantiles (fixed τ as $n \to \infty$), BB intervals have CP differ from the nominal level by $O(n^{-1/2})$, HGK by $O(n^{-1})$, and GKc and CCBB by $O(n^{-3/2}[\log(n)]^3)$.

First, per Theorem 2(ii), CCBB has exact CP when $\epsilon = 0$, which is also true of HGK and GKc (Goldman and Kaplan, 2017, p. 333). Table 2 shows that CP is indeed the nominal 0.95 (up to simulation error) for CCBB, HGK, and GKc when $\epsilon = 0$ (i.e., when the value in Table 1 is an integer). This is true regardless of F_Y. Somewhat coincidentally, CP is also 0.95 in these cases for BB, even though the BB posterior of the interval is strictly above 0.95, per Corollary 1.1.

Second, also per Theorem 2(ii), CCBB has exact CP when F_Y is uniform. This is true in Table 2 (up to simulation error). This is true even for the most challenging $\tau = 0.99$.

Third, CCBB is notably more accurate than BB. For all but the largest (most extreme) τ, CCBB has CP between 0.945 and 0.956, whereas BB's CP varies between 0.900 and 0.972. With $\tau = 0.99$, CCBB is exact when F_Y is uniform, but when F_Y is normal or exponential,

Table 2: Simulated coverage probability

n	τ	F_Y	BB	CCBB	HGK	GKc
18	0.200	Unif$(0,1)$	0.900	0.950	0.961	0.946
18	0.200	N$(0,1)$	0.900	0.956	0.965	0.953
18	0.200	Exp(1)	0.900	0.947	0.959	0.944
18	0.250	Unif$(0,1)$	0.960	0.950	0.954	0.949
18	0.250	N$(0,1)$	0.960	0.951	0.954	0.950
18	0.250	Exp(1)	0.960	0.949	0.953	0.948
18	0.498	Unif$(0,1)$	0.950	0.951	0.950	0.950
18	0.498	N$(0,1)$	0.950	0.951	0.950	0.950
18	0.498	Exp(1)	0.950	0.951	0.950	0.950
18	0.900	Unif$(0,1)$	0.972	0.951	0.959	0.934
18	0.900	N$(0,1)$	0.972	0.947	0.956	0.929
18	0.900	Exp(1)	0.972	0.945	0.954	0.927
18	0.980	Unif$(0,1)$	0.950	0.950	0.950	0.950
18	0.980	N$(0,1)$	0.950	0.950	0.950	0.950
18	0.980	Exp(1)	0.950	0.950	0.950	0.950
18	0.990	Unif$(0,1)$	0.986	0.950	0.974	0.835
18	0.990	N$(0,1)$	0.986	0.898	0.954	0.835
18	0.990	Exp(1)	0.986	0.890	0.946	0.835

Nominal level 0.95; 300,000 replications.

CCBB has CP below 0.90, whereas BB has 0.986 CP. One possible interpretation of these results is that our central quantile asymptotic results in Theorem 2 are not as accurate for such an extreme quantile index as $\tau = 0.99$ with $n = 18$. Although the technical characterization of extreme quantiles and order statistics is asymptotic and does not directly apply to finite samples, the fact that $n\tau = 17.82 > n - 1$ implies an extreme asymptotic approximation is probably more appropriate. Studying CCBB accuracy within an extreme quantile framework would be valuable in future work.

Fourth, CCBB is generally more accurate than HGK, as well as GKc, although there are important caveats. It is somewhat unfair to compare them when F_Y is uniform since the uniform distribution was picked precisely because CCBB is exact in that case. Nonetheless, CCBB usually has CP closer to nominal with the other F_Y, too. For all but the largest (most extreme) τ, for normal and exponential F_Y, CCBB has CP in the range $[0.945, 0.956]$, while the HGK range is $[0.950, 0.965]$, and GKc is $[0.927, 0.953]$. (It may be possible to improve GKc's finite-sample performance by iteratively recomputing the adjusted α and ϵ, although theoretical accuracy would not improve.) As noted by Goldman and Kaplan (2017), though, HGK almost never undercovers since the leading n^{-1} term in its coverage error is always positive (regardless of F_Y, regardless of two-sided, upper one-sided, or lower one-sided intervals), whereas CCBB and GKc can have CP above or below the nominal level. This is particularly notable with the most extreme $\tau = 0.99$, where CCBB significantly undercovers (for the non-uniform F_Y) and GKc undercovers even more, whereas HGK has CP in the range $[0.946, 0.974]$ among the three F_Y. Again, the $\tau = 0.99$ results are partly explained by the central quantile asymptotic framework breaking down at such an extreme quantile.

6 Conclusion

For population quantiles of a continuous distribution, we have established the high-order accuracy of credible intervals from the smoothed Bayesian bootstrap of Banks (1988), with exact finite-sample coverage probability in special cases. For intervals with fractional order statistic endpoints, posterior and coverage probabilities agree up to $O(n^{-3/2}[\log(n)]^3)$. In contrast, without smoothing, the Bayesian bootstrap of Rubin (1981) yields interval probabilities that differ by $O(n^{-1/2})$ from both CCBB and frequentist probabilities.

Future work can try to extend these results to more general settings. For example, one could relax the iid assumption by allowing for sampling weights using the weighted gamma process prior of Lo (1993). Other possible extensions include conditional quantiles, quantile differences, and interquantile ranges, as studied from the frequentist order statistic perspective by Goldman and Kaplan (2017, 2018), or extreme quantiles and quantile regression.

Additionally, for more extreme quantiles, future work could investigate a frequentist method similar to CCBB but achieving exact coverage at a different (non-uniform) reference distribution, perhaps using an estimate of the true distribution. Such a method would not improve upon the CCBB's higher-order coverage accuracy for central quantiles, but it might significantly improve performance with more extreme quantiles.

References

Abramowitz, M., Stegun, I. A. (Eds.), 1972. Handbook of Mathematical Functions With Formulas, Graphs, and Mathematical Tables. Vol. 55 of Applied Mathematics Series. U.S. Department of Commerce, 10th printing.
URL https://eric.ed.gov/?id=ED250164

Banks, D. L., 1988. Histospline smoothing the Bayesian bootstrap. Biometrika 75 (4), 673–684.
URL https://www.jstor.org/stable/2336308

Beran, R., Hall, P., 1993. Interpolated nonparametric prediction intervals and confidence intervals. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 55 (3), 643–652.
URL https://www.jstor.org/stable/2345876

Breth, M., 1979. Nonparametric Bayesian interval estimation. Biometrika 66 (3), 641–644.
URL https://www.jstor.org/stable/2335186

David, H. A., Nagaraja, H. N., 2003. Order Statistics, 3rd Edition. Wiley, New York.
URL https://doi.org/10.1002/0471722162

Efron, B., 1979. Bootstrap methods: Another look at the jackknife. Annals of Statistics 7 (1), 1–26.
URL https://www.jstor.org/stable/2958830

Efron, B., 1982. The Jacknife, the Bootstrap, and Other Resampling Plans. Siam.
URL https://doi.org/10.1137/1.9781611970319

Ferguson, T. S., 1973. A Bayesian analysis of some nonparametric problems. Annals of Statistics 1 (2), 209–230.
URL https://doi.org/10.1214/aos/1176342360

Ferguson, T. S., 1974. Prior distributions on spaces of probability measures. Annals of Statistics 2 (4), 615–629.
URL https://www.jstor.org/stable/2958401

Fisher, R. A., 1932. Statistical Methods for Research Workers, 4th Edition. Oliver and Boyd, Edinburgh.

Goldman, M., Kaplan, D. M., 2016. Evenly sensitive KS-type inference on distributions, working paper available at https://faculty.missouri.edu/~kaplandm.

Goldman, M., Kaplan, D. M., 2017. Fractional order statistic approximation for nonparametric conditional quantile inference. Journal of Econometrics 196 (2), 331–346.
URL https://doi.org/10.1016/j.jeconom.2016.09.015

Goldman, M., Kaplan, D. M., 2018. Non-parametric inference on conditional quantile differ-

ences and linear combinations, using L-statistics. Econometrics Journal 21 (2), 136–169.
URL https://doi.org/10.1111/ectj.12095

Hahn, J., 1997. Bayesian bootstrap of the quantile regression estimator: A large sample study. International Economic Review 38 (4), 795–808.
URL https://www.jstor.org/stable/2527216

Hall, P., Martin, M. A., 1989. A note on the accuracy of bootstrap percentile method confidence intervals for a quantile. Statistics & Probability Letters 8 (3), 197–200.
URL https://doi.org/10.1016/0167-7152(89)90121-1

Ho, Y. H. S., Lee, S. M. S., 2005. Calibrated interpolated confidence intervals for population quantiles. Biometrika 92 (1), 234–241.
URL http://biomet.oxfordjournals.org/content/92/1/234.abstract

Hutson, A. D., 1999. Calculating nonparametric confidence intervals for quantiles using fractional order statistics. Journal of Applied Statistics 26 (3), 343–353.
URL https://doi.org/10.1080/02664769922458

Lo, A. Y., 1993. A Bayesian method for weighted sampling. Annals of Statistics 21 (4), 2138–2148.
URL https://www.jstor.org/stable/2242333

Meeden, G., 1993. Noninformative nonparametric Bayesian estimation of quantiles. Statistics & Probability Letters 16 (2), 103–109.
URL https://doi.org/10.1016/0167-7152(93)90153-A

Neyman, J., 1937. »Smooth test» for goodness of fit. Skandinavisk Aktuarietidskrift 20 (3–4), 149–199.
URL https://doi.org/10.1080/03461238.1937.10404821

Pearson, K., 1933. On a method of determining whether a sample of size n supposed to have been drawn from a parent population having a known probability integral has probably been drawn at random. Biometrika 25, 379–410.
URL https://doi.org/10.1093/biomet/25.3-4.379

R Core Team, 2017. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.
URL https://www.R-project.org/

Robbins, H., 1955. A remark on Stirling's formula. The American Mathematical Monthly 62 (1), 26–29.
URL https://doi.org/10.2307/2308012

Rubin, D. B., 1981. The Bayesian bootstrap. Annals of Statistics 9 (1), 130–134.
URL http://projecteuclid.org/euclid.aos/1176345338

Wahba, G., 1975. Optimal convergence properties of variable knot, kernel, and orthogonal series methods for density estimation. Annals of Statistics 3 (1), 15–29.
URL https://www.jstor.org/stable/2958077

Wilks, S. S., 1962. Mathematical Statistics. Wiley, New York.

A Proofs

A.1 Proof of Theorem 1

Proof. First,

$$Y_{n:k} < \mathrm{Q}_\tau(Y) \iff \sum_{i=1}^{k} p_i < \tau.$$

Thus,

$$\pi_u(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) = \pi_u(\sum_{i=1}^{k} p_i < \tau \mid \mathbf{Y}).$$

By (9), $\sum_{i=1}^{k} p_i \mid \mathbf{Y} \sim \mathrm{Beta}(k, n-k)$ for BB. Since the beta distribution is continuous,

$$\pi_u(\sum_{i=1}^{k} p_i < \tau \mid \mathbf{Y}) = \pi_u(\sum_{i=1}^{k} p_i \leq \tau \mid \mathbf{Y}) = F_\beta(\tau; k, n-k). \tag{16}$$

Since the BB posterior's support includes only discrete distributions with support $\{Y_i\}_{i=1}^n$,

$$Y_{n:k} \leq \mathrm{Q}_\tau(Y) \iff Y_{n:k-1} < \mathrm{Q}_\tau(Y).$$

Combining this with (16),

$$\pi_u(Y_{n:k} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) = \pi_u(Y_{n:k-1} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) = F_\beta(\tau; k-1, n-(k-1)).$$

For the CCBB, the same steps hold as for (16), but with $\sum_{i=1}^{k} p_i \mid \mathbf{Y} \sim \mathrm{Beta}(k, n+1-k)$ from (9), so

$$\pi_s(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) = \pi_s(\sum_{i=1}^{k} p_i < \tau \mid \mathbf{Y}) = \pi_s(\sum_{i=1}^{k} p_i \leq \tau \mid \mathbf{Y}) = F_\beta(\tau; k, n+1-k).$$

Since the CCBB posterior of $\mathrm{Q}_\tau(Y)$ is continuous, $\pi_s(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) = \pi_s(Y_{n:k} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y})$.

With continuous $F_Y(\cdot)$, using Wilks (1962, **8.7.4**),

$$\mathrm{P}(Y_{n:k} \leq \mathrm{Q}_\tau(Y)) = \mathrm{P}(F_Y(Y_{n:k}) \leq \tau) = F_\beta(\tau; k, n+1-k) = \pi_s(Y_{n:k} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y}). \quad \square$$

A.2 Proof of Corollary 1.1

Proof. The connection between beta and binomial distributions is used below. The PDF of a Beta(a, b) distribution is (Abramowitz and Stegun, 1972, **6.2.2**, **26.1.33**)

$$f_\beta(p; a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} p^{a-1}(1-p)^{b-1},$$

where $\Gamma(\cdot)$ is the gamma function. If m is an integer, then $\Gamma(m) = (m-1)!$ (Abramowitz and Stegun, 1972, **6.1.6**). Thus, if a and b are integers,

$$f_\beta(p; a, b) = \frac{(a+b-1)!}{(a-1)!(b-1)!} p^{a-1}(1-p)^{b-1}.$$

The PMF of a Binomial(m, p) distribution is (Abramowitz and Stegun, 1972, **26.1.20**)

$$f_B(d; m, p) = \binom{m}{d} p^d (1-p)^{m-d} = \frac{m!}{d!(m-d)!} p^d (1-p)^{m-d}.$$

With $d = k-1$, $m = n-1$, and $p = \tau$,

$$f_B(k-1; n-1, \tau) = \frac{(n-1)!}{(k-1)!(n-k)!} \tau^{k-1}(1-\tau)^{n-k}.$$

With $p = \tau$, $a = k$, and $b = n+1-k$,

$$f_\beta(\tau; k, n+1-k) = \frac{(n)!}{(k-1)!(n-k)!} \tau^{k-1}(1-\tau)^{n-k} = n f_B(k-1; n-1, \tau). \tag{17}$$

The beta and binomial CDFs can both be written in terms of the regularized incomplete beta function I. The beta CDF is $F_\beta(p; a, b) = I_p(a, b)$ (David and Nagaraja, 2003, eqn. (1.3.2)). Applying **26.5.24** in Abramowitz and Stegun (1972),

$$1 - F_B(d; m, p) = \overbrace{\sum_{s=d+1}^{m} f_B(s; m, p) = I_p(d+1, m-d)}^{\text{from } \mathbf{26.5.24}} = F_\beta(p; d+1, m-d). \tag{18}$$

From Theorem 1, applying (18) with $p = \tau$ and either $(d = k-2, m = n-1)$ or $(d = k-1, m = n)$,

$$\begin{aligned}
&\pi_u(Y_{n:k} \le Q_\tau(Y) \mid \mathbf{Y}) - \pi_s(Y_{n:k} \le Q_\tau(Y) \mid \mathbf{Y}) \\
&= F_\beta(\tau; k-1, n+1-k) - F_\beta(\tau; k, n+1-k) \\
&= [1 - F_B(k-2; n-1, \tau)] - [1 - F_B(k-1; n, \tau)] \\
&= F_B(k-1; n, \tau) - F_B(k-2; n-1, \tau).
\end{aligned}$$

The two binomial CDFs may be expressed in terms of $B_i \stackrel{iid}{\sim} \text{Bernoulli}(\tau)$:

$$F_B(k-2; n-1, \tau) = \mathrm{P}(\sum_{i=1}^{n-1} B_i \leq k-2), \quad F_B(k-1; n, \tau) = \mathrm{P}(\sum_{i=1}^{n} B_i \leq k-1).$$

Since $\sum_{i=1}^{n} B_i = B_n + \sum_{i=1}^{n-1} B_i$ and $B_n \leq 1$,

$$\sum_{i=1}^{n-1} B_i \leq k-2 \implies B_n + \sum_{i=1}^{n-1} B_i \leq 1 + k - 2 \implies \sum_{i=1}^{n} B_i \leq k-1.$$

Thus, the CDF difference equals the probability that the latter inequality holds but the former does not, which only occurs when $\sum_{i=1}^{n-1} B_i = k-1$ and $B_n = 0$:

$$\begin{aligned}
F_B(k-1; n, \tau) - F_B(k-2; n-1, \tau) &= \mathrm{P}(\sum_{i=1}^{n} B_i \leq k-1 \text{ and } \sum_{i=1}^{n-1} B_i > k-2) \\
&= \mathrm{P}(\overbrace{\sum_{i=1}^{n-1} B_i = k-1 \text{ and } B_n = 0}^{\text{use } \{B_i\}_{i=1}^{n-1} \perp\!\!\!\perp B_n}) \\
&= \mathrm{P}(\sum_{i=1}^{n-1} B_i = k-1)\, \mathrm{P}(B_n = 0) \\
&= \overbrace{f_B(k-1; n-1, \tau)}^{\text{apply (17)}}(1-\tau) \\
&= n^{-1} f_\beta(\tau; k, n+1-k)(1-\tau).
\end{aligned}$$

The other posterior difference is derived using the same steps. From Theorem 1, applying (18) with $p = \tau$ and either $(d = k-1, m = n)$ or $(d = k-1, m = n-1)$, and again using the B_i defined earlier,

$$\begin{aligned}
&\pi_s(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) - \pi_u(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) \\
&= F_\beta(\tau; k, n-k+1) - F_\beta(\tau; k, n-k) \\
&= [1 - F_B(k-1; n, \tau)] - [1 - F_B(k-1; n-1, \tau)] \\
&= F_B(k-1; n-1, \tau) - F_B(k-1; n, \tau) \\
&= \mathrm{P}(\sum_{i=1}^{n-1} B_i \leq k-1) - \mathrm{P}(\sum_{i=1}^{n} B_i \leq k-1) \\
&= \mathrm{P}(\sum_{i=1}^{n-1} B_i = k-1 \text{ and } B_n = 1)
\end{aligned}$$

$$
\begin{aligned}
&= \mathrm{P}(\sum_{i=1}^{n-1} B_i = k-1)\,\mathrm{P}(B_n = 1) \\
&= f_B(k-1; n-1, \tau)\tau \\
&= \frac{\tau}{n} f_\beta(\tau; k, n+1-k).
\end{aligned}
$$

The mode of a Beta(a, b) distribution is $(a-1)/(a+b-2)$ for $a, b > 1$. Thus for $1 < k < n$, the mode of a Beta$(k, n+1-k)$ distribution is $(k-1)/(n-1)$, i.e.,

$$
\arg\sup_{0 \le t \le 1} f_\beta(t; k, n+1-k) = (k-1)/(n-1).
$$

Thus, the maximum value of the beta PDF is $f_\beta((k-1)/(n-1); k, n+1-k)$. Applying the Stirling-type approximation

$$
m! = \sqrt{2\pi} m^{m+(1/2)} e^{-m} e^{r_m}, \quad \frac{1}{12m+1} < r_m < \frac{1}{12m} \tag{19}
$$

of Robbins (1955) and evaluating (17) at the mode $\tau = (k-1)/(n-1)$,

$$
\begin{aligned}
&f_\beta((k-1)/(n-1); k, n+1-k) \\
&= \frac{n!}{(k-1)!(n-k)!}[(k-1)/(n-1)]^{k-1}[\overbrace{1-(k-1)/(n-1)}^{=[n-1-(k-1)]/(n-1)}]^{n-k} \\
&= \frac{n[(n-1)!]}{(k-1)!(n-k)!}[(k-1)/(n-1)]^{k-1}[(n-k)/(n-1)]^{n-k} \\
&= n\frac{(n-1)!}{(k-1)!(n-k)!}(k-1)^{k-1}(n-k)^{n-k}[1/(n-1)]^{(k-1)+(n-k)} \\
&= n\frac{(n-1)!}{(n-1)^{n-1}}\frac{(k-1)^{k-1}}{(k-1)!}\frac{(n-k)^{n-k}}{(n-k)!} \\
&= n\frac{\sqrt{2\pi}(n-1)^{1/2}e^{-(n-1)}e^{r_{n-1}}}{2\pi(k-1)^{1/2}(n-k)^{1/2}e^{-(k-1)}e^{-(n-k)}e^{r_{k-1}}e^{r_{n-k}}} \\
&= n(2\pi)^{-1/2}\sqrt{\frac{n-1}{(k-1)(n-k)}}\underbrace{\exp\{\overbrace{-(n-1)+(k-1)+(n-k)}^{=0}\}}_{=1}\exp\{r_{n-1}-r_{k-1}-r_{n-k}\}.
\end{aligned}
$$

For the approximation error, since $r_{k-1}, r_{n-k} > 0$,

$$
\exp\{r_{n-1} - r_{k-1} - r_{n-k}\} \le \exp\{r_{n-1}\} \le \exp\{1/[12(n-1)]\} = 1 + O(n^{-1}).
$$

Thus,

$$f_\beta((k-1)/(n-1); k, n+1-k) \le n(2\pi)^{-1/2}\sqrt{\frac{n-1}{(k-1)(n-k)}}[1+O(n^{-1})]. \tag{20}$$

If $c_1 n \le k \le c_2 n$ for $0 < c_1 < c_2 < 1$ as $n \to \infty$, i.e., if k is of order n, then the upper bound rate of (20) simplifies further. Specifically, this implies $k-1 \ge c_1 n - 1$ and $n-k \ge n - c_2 n = n(1-c_2)$, so

$$\sqrt{\frac{n-1}{(k-1)(n-k)}} \le \sqrt{\frac{n}{(c_1 n-1)n(1-c_2)}} = [(c_1 n - 1)(1-c_2)]^{-1/2} = O(n^{-1/2}). \tag{21}$$

Replacing $n(2\pi)^{-1/2}$ with $O(n)$ in (20), the posterior difference is

$$n^{-1}\overbrace{f_\beta((k-1)/(n-1); k, n+1-k)}^{\text{use (20)}} = n^{-1}O(n)O(n^{-1/2})[1+O(n^{-1})] = O(n^{-1/2}).$$

If k is not required to have rate n, then the bound on the difference is larger. If k or $n-k$ is fixed as $n \to \infty$ (i.e., $Y_{n:k}$ is an "extreme order statistic"), then (21) becomes $O(1)$ instead of $O(n^{-1/2})$, so the overall posterior difference becomes $O(1)$ instead of $O(n^{-1/2})$. If $c_1 n^r \le k \le c_2 n^r$ for $0 < r < 1$, then (21) becomes

$$\sqrt{\frac{n-1}{(k-1)(n-k)}} \le \sqrt{\frac{n}{(c_1 n^r - 1)(n - c_2 n^r)}} = O(n^{-r/2}).$$

The same holds if instead $c_1 n^r \le n-k \le c_2 n^r$. In either case, the overall posterior difference also becomes $O(n^{-r/2})$. □

A.3 Proof of Theorem 2

Proof. For Theorem 2(i), the BB result follows immediately from Theorem 1 since

$$\pi_u(Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) = \pi_u(Y_{n:k} < \mathrm{Q}_\tau(Y) \mid \mathbf{Y}).$$

For the CCBB result in Theorem 2(i), first we show that (conditional on $\mathbf{Y}$)

$$\overbrace{Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < \mathrm{Q}_\tau(Y)}^{\text{event } A} \iff \overbrace{\sum_{i=1}^{k} p_i + \epsilon p_{k+1} < \tau}^{\text{event } B}. \tag{22}$$

Write the posterior $\mathrm{Q}_\tau(Y)$ with $c \equiv \max\{d : \sum_{i=1}^d p_i \le \tau\}$ as in (14). We show (22) holds for $c < k$, $c > k$, and $c = k$. If $c < k$, implying $\sum_{i=1}^{k} p_i > \tau$, then neither A nor B occurs:

$$\text{for } A, \quad \overbrace{\mathrm{Q}_\tau(Y) < \underbrace{Y_{n:c+1}}_{}}^{\text{by (14)}} \underbrace{\le Y_{n:k}}_{\text{by } c<k} \overbrace{\le Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k})}^{\text{by } \epsilon\ge 0};$$

$$\text{for } B, \quad \overbrace{\sum_{i=1}^{k} p_i + \epsilon p_{k+1} \ge \sum_{i=1}^{k} p_i}^{\text{by } \epsilon\ge 0} \underbrace{\vphantom{\sum} > \tau}_{\text{implied by } c<k} .$$

If $c > k$, implying $c \ge k+1$ and thus $\sum_{i=1}^{k+1} p_i \le \tau$, then both A and B occur:

$$\text{for } A, \quad \overbrace{\mathrm{Q}_\tau(Y) \ge Y_{n:c}}^{\text{by (14)}} \underbrace{\ge Y_{n:k+1}}_{\text{by } c>k} \overbrace{> Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k})}^{\text{by } \epsilon<1};$$

$$\text{for } B, \quad \overbrace{\sum_{i=1}^{k} p_i + \epsilon p_{k+1} < \sum_{i=1}^{k+1} p_i}^{\text{by } \epsilon<1} \underbrace{\vphantom{\sum} \le \tau}_{\text{implied by } c>k} .$$

If $c = k$, then event A in (22) becomes

$$\overbrace{Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < \mathrm{Q}_\tau(Y)}^{\text{event } A \text{ in (22)}} \underbrace{= Y_{n:k} + \frac{(\tau - \sum_{i=1}^{k} p_i)}{p_{k+1}}(Y_{n:k+1} - Y_{n:k})}_{\text{from (14) with } c=k}$$

$$\iff \epsilon < \frac{\tau - \sum_{i=1}^{k} p_i}{p_{k+1}}$$

$$\iff \underbrace{\sum_{i=1}^{k} p_i + \epsilon p_{k+1} < \tau}_{\text{event } B \text{ in (22)}} .$$

Since $A \iff B$ in (22), the corresponding posterior probabilities are equal:

$$\pi_s\{\overbrace{Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < \mathrm{Q}_\tau(Y)}^{A} \mid \mathbf{Y}\} = \pi_s\{\overbrace{\sum_{i=1}^{k} p_i + \epsilon p_{k+1} < \tau}^{B} \mid \mathbf{Y}\}. \tag{23}$$

The right-hand side of (23) can then be approximated by Theorem 2(i) of Goldman and Kaplan (2017). Their theorem concerns the distribution of (linear combinations of) fractional order statistics. Recall from (5) and (8) that the joint posterior distribution of $(p_1, \ldots, p_{n+1})$ is the same as the joint sampling distribution of standard uniform order statistic spacings $(U_{n:1}, U_{n:2} - U_{n:1}, \ldots, 1 - U_{n:n})$ given iid sampling. The sum of the first k spacings is $U_{n:k}$, and the $(k+1)$th spacing is $U_{n:k+1} - U_{n:k}$, so the posterior distribution of $\sum_{i=1}^{k} p_i + \epsilon p_{k+1}$ equals the sampling distribution of $U_{n:k+\epsilon} \equiv U_{n:k} + \epsilon(U_{n:k+1} - U_{n:k})$, i.e., the fractional $(k+\epsilon)$th order statistic. That is, for any value $0 \le r \le 1$,

$$\pi_s(\sum_{i=1}^{k} p_i + \epsilon p_{k+1} \le r \mid \mathbf{Y}) = \mathrm{P}\{U_{n:k} + \epsilon(U_{n:k+1} - U_{n:k}) \le r\}. \tag{24}$$

Theorem 2(i) of Goldman and Kaplan (2017) approximates the right-hand side of (24).

The following shows specifically how to apply Theorem 2(i) of Goldman and Kaplan (2017) to $U_{n:k+\epsilon}$. Their Assumption 1 is iid sampling, as in our case. In their notation, $J = 1$, $\psi_1 = 1$, and $F(\cdot)$ is the standard uniform CDF: for $0 \le t \le 1$, $F(t) = t$, $F^{-1}(t) = t$, and PDF $f(t) = 1$. Thus, their Assumption 2 is satisfied for any $0 < u < 1$ since $f(u) > 0$ and $f''(u) = 0$ is continuous in u. Their u_1 and ϵ_1 are the same as u and ϵ here. So far, this makes their general fractional L-statistic L^L equal to our $U_{n:k+\epsilon}$. Also, their $\mathbb{X}_0$ equals our u, and their $\mathcal{V}_{\boldsymbol{\psi}}$ equals our $u(1-u)$. Their L^I follows a $\mathrm{Beta}((n+1)u, (n+1)(1-u)) = \mathrm{Beta}(k+\epsilon, n+1-k-\epsilon)$ distribution. Thus, setting their $\mathbb{X}_0 + n^{-1/2}K$ equal to our τ, their Theorem 2(i) says in their notation that

$$\mathrm{P}(L^L < \mathbb{X}_0 + n^{-1/2}K) - \mathrm{P}(L^I < \mathbb{X}_0 + n^{-1/2}K) = O(n^{-1}),$$

which in our notation is

$$\mathrm{P}(U_{n:k+\epsilon} < \tau) - F_\beta(k+\epsilon, n+1-k-\epsilon) = O(n^{-1}). \tag{25}$$

Theorem 2(i) of Goldman and Kaplan (2017) further provides an analytic expression for the n^{-1} term. Given $\mathbb{X}_0 = u$ and $\mathbb{X}_0 + n^{-1/2}K = \tau$, solving for their K yields $K = \sqrt{n}(\tau - u)$. Given $u = \tau - dn^{-1/2}$, $K = \sqrt{n}(\tau - u) = d$. Plugging into their Theorem 2(i) yields, first in their notation and then translated into ours with $K = d$, $\mathcal{V}_{\boldsymbol{\psi}} = u(1-u)$, $\psi_1 = 1$, and $\epsilon_1 = \epsilon$,

$$\begin{aligned}
&\mathrm{P}(U_{n:k+\epsilon} < \tau) - F_\beta(k+\epsilon, n+1-k-\epsilon) \\
&= n^{-1}\frac{K\exp\{-K^2/(2\mathcal{V}_{\boldsymbol{\psi}})\}}{\sqrt{2\pi\mathcal{V}_{\boldsymbol{\psi}}^3}} \overbrace{\psi_1}^{=1} \overbrace{\epsilon_1(1-\epsilon_1)}^{=\epsilon(1-\epsilon)} \overbrace{[f(F^{-1}(u))]^{-2}}^{=1} + O(n^{-3/2}[\log(n)]^3) \\
&= n^{-1}\frac{d\exp\{-d^2/[2u(1-u)]\}}{\sqrt{2\pi[u(1-u)]^3}}\epsilon(1-\epsilon) + O(n^{-3/2}[\log(n)]^3)
\end{aligned}$$

$$= n^{-1}\frac{\epsilon(1-\epsilon)}{u(1-u)}\frac{d}{\sqrt{u(1-u)}}\frac{1}{\sqrt{2\pi}}\exp\{-(1/2)[d/\sqrt{u(1-u)}]^2\} + O(n^{-3/2}[\log(n)]^3)$$

$$= n^{-1}\frac{\epsilon(1-\epsilon)}{\tau(1-\tau)}\frac{d}{\sqrt{\tau(1-\tau)}}\phi\Big(d/\sqrt{\tau(1-\tau)}\Big) + O(n^{-3/2}[\log(n)]^3), \tag{26}$$

using the fact that $u = \tau + O(n^{-1/2})$ and writing $\phi(\cdot)$ as the standard normal PDF.

For Theorem 2(ii), the $\epsilon = 0$ result follows immediately from Theorem 1 and the result from Wilks (1962) that $F_Y(Y_{n:k}) \stackrel{d}{=} U_{n:k} \sim \text{Beta}(k, n+1-k)$ if $F_Y(\cdot)$ is continuous, so

$$\text{P}(Y_{n:k} < \text{Q}_\tau(Y)) = \text{P}\{F_Y(Y_{n:k}) < F_Y(\text{Q}_\tau(Y))\} = \text{P}\{U_{n:k} < \tau\} = F_\beta(\tau; k, n+1-k).$$

The other result comes from (24), which directly gives the equivalence for $\text{Unif}(0,1)$ that can be extended to $\text{Unif}(a,b)$ as follows. Let $V_i \stackrel{iid}{\sim} \text{Unif}(a,b)$, with order statistics $V_{n:k}$ and $\text{Q}_\tau(V) = a + (b-a)\tau$. Then, $(V_i - a)/(b-a) \stackrel{iid}{\sim} \text{Unif}(0,1)$, and

$$\begin{aligned}
&\text{P}\{V_{n:k} + \epsilon(V_{n:k+1} - V_{n:k}) < a + (b-a)\tau\} \\
&= \text{P}\{V_{n:k} - a + \epsilon[(V_{n:k+1} - a) - (V_{n:k} - a)] < (b-a)\tau\} \\
&= \text{P}\left\{\frac{V_{n:k} - a}{b-a} + \epsilon\left[\frac{V_{n:k+1} - a}{b-a} - \frac{V_{n:k} - a}{b-a}\right] < \tau\right\} \\
&= \text{P}\{U_{n:k} + \epsilon(U_{n:k+1} - U_{n:k}) < \tau\}.
\end{aligned}$$

For Theorem 2(iii), plugging $d = z_{1-\alpha}\sqrt{\tau(1-\tau)} + O(n^{-1/2})$ from (15) into (26),

$$\begin{aligned}
&\pi_s\{Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < \text{Q}_\tau(Y) \mid \mathbf{Y}\} - F_\beta(k+\epsilon, n+1-k-\epsilon) \\
&= n^{-1}\frac{\epsilon(1-\epsilon)}{\tau(1-\tau)}\frac{d}{\sqrt{\tau(1-\tau)}}\phi\Big(d/\sqrt{\tau(1-\tau)}\Big) + O(n^{-3/2}[\log(n)]^3) \\
&= n^{-1}\frac{\epsilon(1-\epsilon)z_{1-\alpha}}{\tau(1-\tau)}\phi(z_{1-\alpha}) + O(n^{-3/2}[\log(n)]^3),
\end{aligned}$$

for any $0 < \alpha < 1$. This matches the coverage probability in Theorem 4(i) from Goldman and Kaplan (2017), noting that their p is our τ. That is, their Theorem 4(i) says that if A1 and A2 and (15) hold, then

$$\begin{aligned}
&\text{P}\{Y_{n:k} + \epsilon(Y_{n:k+1} - Y_{n:k}) < \text{Q}_\tau(Y)\} - F_\beta(k+\epsilon, n+1-k-\epsilon) \\
&\qquad = n^{-1}\frac{\epsilon(1-\epsilon)z_{1-\alpha}}{\tau(1-\tau)}\phi(z_{1-\alpha}) + O(n^{-3/2}[\log(n)]^3).
\end{aligned}$$

More directly, the n^{-1} term in their Theorem 4(i) does not depend on the underlying distribution (other than satisfying A2). Since by Theorem 2(ii) the CCBB posterior probability is equivalent to the coverage probability when the underlying distribution is uniform, it is also equal to the coverage probability given any other distribution satisfying A2 up to the $O(n^{-3/2}[\log(n)]^3)$ remainder. □

A.4 Proof of Corollary 2.1

Proof. For $\epsilon = 0$, Corollary 1.1 gives the $O(n^{-1/2})$ difference. The difference cannot be smaller over the larger range of $\epsilon \in [0, 1)$; below we show it is not bigger, either.

For $0 < \epsilon < 1$, the CCBB probability is bounded by

$$\pi_s(Y_{n:k+1} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) < \pi_s(Y_{n:k+\epsilon} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) < \pi_s(Y_{n:k} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y}).$$

From Theorems 1 and 2, this implies

$$F_\beta(\tau; k+1, n-k) < \pi_s(Y_{n:k+\epsilon} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) < F_\beta(\tau; k, n+1-k). \tag{27}$$

The difference between the unsmoothed BB probability $F_\beta(\tau; k, n-k)$ and the bounds in (27) can be computed using Corollary 1.1. The upper bound difference comes directly from Corollary 1.1:

$$F_\beta(\tau; k+1, n-k) - F_\beta(\tau; k, n-k) = \frac{\tau}{n} f_\beta(\tau; k, n+1-k). \tag{28}$$

Replacing k with $k+1$ in Theorem 1 yields

$$\pi_u(Y_{n:k+1} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) = F_\beta(\tau; k, n-k), \quad \pi_s(Y_{n:k+1} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) = F_\beta(\tau; k+1, n-k). \tag{29}$$

Further applying Corollary 1.1 after replacing k with $k+1$,

$$F_\beta(\tau; k, n-k) - F_\beta(\tau; k+1, n-k) = \frac{1-\tau}{n} f_\beta(\tau; k+1, n-k). \tag{30}$$

This is the difference between the BB probability and the lower bound in (27).

As ϵ goes from 0 to 1, the CCBB probability goes from the upper bound to the lower bound in (27). Thus, the BB–CCBB difference goes from

$$-\frac{\tau}{n} f_\beta(\tau; k, n+1-k) \text{ to } \frac{1-\tau}{n} f_\beta(\tau; k+1, n-k). \tag{31}$$

That is,

$$-\frac{\tau}{n} f_\beta(\tau; k, n+1-k) \leq \pi_u(Y_{n:k+\epsilon} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) - \pi_s(Y_{n:k+\epsilon} \leq \mathrm{Q}_\tau(Y) \mid \mathbf{Y}) \leq \frac{1-\tau}{n} f_\beta(\tau; k+1, n-k). \tag{32}$$

Since the lower bound is negative and the upper bound positive, and the difference varies continuously with ϵ, there is some magic ϵ for which the BB and CCBB actually agree. However, since ϵ is determined by the nominal level and sample size (i.e., not chosen independently), in general the difference is larger. As stated in Corollary 1.1, with $c_1 n \leq k \leq c_2 n$, the bounds in (32) are $O(n^{-1/2})$. $\square$