

Bayesian and frequentist nonlinear inequality tests

David M. Kaplan* Longhao Zhuo†
Department of Economics, University of Missouri

April 25, 2017

Abstract

Bayesian and frequentist criteria are fundamentally different, but often posterior and sampling distributions are asymptotically equivalent (and normal). We compare Bayesian and frequentist inference on nonlinear inequality restrictions in such cases. To quantify the comparison, we examine the (frequentist) size of a Bayesian hypothesis test (based on a comparable loss function). For finite-dimensional parameters, if the null hypothesis is that the parameter vector lies in a certain half-space, then the Bayesian test has size α ; if the null hypothesis is a subset of a half-space (with strictly smaller volume), then the Bayesian test has size strictly above α ; and in other cases, the Bayesian test's size may be above or below α . For infinite-dimensional parameters, similar results hold. Two examples illustrate our results: inference on stochastic dominance and on curvature of a translog cost function. We hope these results increase awareness of when Bayesian and frequentist inferences may differ significantly in practice, as well as increase intuition about such differences.

JEL classification: C11, C12

Keywords: Bernstein–von Mises theorem, limit experiment, nonstandard inference, stochastic dominance, translog

1 Introduction

Although Bayesian and frequentist properties are fundamentally different, in many cases we can (approximately) achieve both. In other cases, the Bayesian and frequentist summaries of the data differ greatly, and practitioners must carefully consider which to prefer. We provide results on the role of null hypothesis “shape” in determining such differences. We hope to

*Corresponding author. Email: kaplandm@missouri.edu. Mail: Department of Economics, University of Missouri, 118 Professional Bldg, 909 University Ave, Columbia, MO 65211-6040, United States. Many thanks to Tim Armstrong, Jim Berger, Jeremy Fox, Patrik Guggenberger, Zack Miller, Stephen Montgomery-Smith, and Andres Santos for helpful discussion, comments, examples, and references. Thanks also to anonymous reviewers for detailed comments, references, and the generalization beyond normality.

†Email: lz2z4@mail.missouri.edu.

alert practitioners to situations prone to large differences and to foster understanding of why such differences are large.

Economic theory mostly concerns inequalities, often nonlinear.¹ For example, inequality of CDFs characterizes first-order stochastic dominance (SD1), an important concept for welfare analysis. Bayesian and frequentist SD1 inferences may differ greatly, and the direction of the difference partly depends on whether the null hypothesis is dominance or non-dominance. An example of nonlinear inequalities is curvature constraints on production, cost, indirect utility, and other functions. Such constraints usually result from optimization, like utility or profit maximization. Our theoretical results suggest differences between Bayesian and frequentist inference on both SD1 and curvature, even asymptotically. We illustrate these large differences through simulations and example datasets.

Further motivation for studying Bayesian–frequentist differences here is that deriving frequentist tests for general nonlinear inequalities is notoriously difficult; e.g., see Wolak (1991). In contrast, it is (relatively) simple to compute the posterior probability that the parameter satisfies certain nonlinear inequalities, by computing the proportion of draws from the parameter’s posterior in which the constraints are satisfied. Perhaps especially in the absence of a feasible frequentist method, it is helpful to understand if the Bayesian test’s size differs greatly from the nominal α .

Statistically, we consider cases where the sampling distribution of some estimator is asymptotically normal, while the asymptotic posterior is also normal, with the same covariance matrix. (Our formal results relax normality, but it remains the leading case.) For simplicity, we consider the corresponding limit experiment: a single draw from a normal distribution with unknown mean and known covariance. In the limit experiment, an uninformative improper prior captures the assumed asymptotic independence of the prior and posterior.

To quantitatively compare Bayesian and frequentist inference, we characterize the frequentist size of the Bayesian test that rejects the null hypothesis when its posterior probability is below α . In addition to being intuitive and practically salient, there are decision-theoretic reasons to examine this test, as detailed in Appendix B. Although size gives no insight into admissibility, it captures the practical difference between reporting a Bayesian and frequentist result.

Under general conditions, we characterize the role of the shape of the null hypothesis, H_0 , in terms of the Bayesian test’s size. By “the shape of H_0 ,” we mean the shape of the

¹Nonlinear inequalities also come from other sources. For example, $H_0 : \theta_1\theta_2 \geq 0$ can be used to test stability of the sign of a parameter over time (or geography), or whether a treatment attenuates the effect of another regressor; see Kaplan (2015) for details.

parameter subspace where H_0 is satisfied. If H_0 is a half-space, then the Bayesian test has size α exactly. If H_0 is strictly smaller than a half-space, then the Bayesian test’s size is strictly above α . If H_0 is not contained within a half-space (i.e., does not have a supporting hyperplane), then the Bayesian test’s size may be above, equal to, or below α . An immediate corollary of these results is that Bayesian and frequentist tests agree asymptotically for testing a single linear inequality constraint, while frequentist testing is strictly more conservative for testing two or more linear inequality constraints. We provide similar results for infinite-dimensional parameters.

Our results beg the question: if inferences on H_0 can disagree while credible and confidence sets coincide, why not simply report the credible or confidence set?² If interest is primarily in the parameters themselves, then reporting a credible or confidence set may be better than a posterior or p -value for H_0 . However, sometimes interest is in testing implications of economic theory or in specification testing. Other times, inequalities provide economically relevant summaries of a high-dimensional parameter, like whether a certain income distribution stochastically dominates another.

Literature Many papers compare Bayesian and frequentist inference, in a variety of setups. Here, we highlight examples of different types of conclusions: sometimes frequentist inference is more conservative, sometimes Bayesian, sometimes neither.

Some of the literature documents cases where frequentist inference is “too conservative” from a Bayesian perspective. For testing linear inequality constraints of the form $H_0 : \boldsymbol{\theta} \geq \mathbf{0}$ with $\boldsymbol{\theta} \in \mathbb{R}^d$, Kline (2011) finds frequentist testing to be more conservative (e.g., his Figure 1), especially as the dimension d grows; this agrees with our general result. As another example, under set identification, asymptotically, frequentist confidence sets for the true parameter (Imbens and Manski, 2004; Stoye, 2009) are strictly larger than the estimated identified set, whereas Bayesian credible sets are strictly smaller, as shown by Moon and Schorfheide (2012, Cor. 1).³ Our setup is not directly comparable to theirs since a Bayesian credible set cannot be inverted into a test.

Other papers document cases where frequentist inference is “too aggressive” from a Bayesian perspective. Perhaps most famously, in Lindley’s (1957) paradox, the frequentist test rejects while the Bayesian test does not. Berger and Sellke (1987) make a similar argument. In both cases, as noted by Casella and Berger (1987b), the results follow pri-

²Berger (2003) also notes this agreement on credible/confidence sets but disagreement on testing. However, he writes, “The disagreement occurs primarily when testing a ‘precise’ hypothesis” (p. 2), whereas we find disagreements even with inequality hypotheses. Also, Casella and Berger (1987b, p. 344) opine, “Interval estimation is, in our opinion, superior to point null hypothesis testing.”

³There seems to be a typo in the statement of Corollary 1(ii), switching the frequentist and Bayesian sets from their correct places seen in the Supplemental Material proof.

marily from having a large prior probability on a point (or “small interval”) null hypothesis, specifically $P(H_0) = 1/2$. Arguing that $P(H_0) = 1/2$ is “objective,” Berger and Sellke (1987, p. 113) consider $P(H_0) = 0.15$ to be “blatant bias toward H_1 .” Casella and Berger (1987b) disagree, saying $P(H_0) = 1/2$ is “much larger than is reasonable for most problems” (p. 344).

In yet other cases, Bayesian and frequentist inferences are similar or even identical. Casella and Berger (1987a) compare Bayesian and frequentist one-sided testing of a location parameter, given a single draw of X from an otherwise fully known density. They compare the p -value, $p(x)$, to the infimum of the posterior $P(H_0 | x)$ over various classes of priors. In many cases, the infimum is attained by the improper prior of Lebesgue measure on $(-\infty, \infty)$ and equals $p(x)$ (p. 109). Goutis, Casella, and Wells (1996) consider jointly testing multiple one-sided hypotheses. In a single-draw Gaussian shift experiment (similar to this paper), further assuming all components of the vector \mathbf{X} are mutually independent, they consider the Bayesian posterior on H_0 when the (improper, uninformative) prior is adjusted to have $P(H_0) = 1/2$. In this case, the posterior is proportional to one of the frequentist p -values they consider, but it is (weakly) smaller. This complements our setting where we impose neither independence nor $P(H_0) = 1/2$, and we do not restrict the shape of the null hypothesis subspace.

Paper structure and notation Section 2 contains our main results and discussion. Section 3 illustrates our results with stochastic dominance and cost function curvature. Proofs not in the main text are collected in Appendix A. Appendix B provides some decision-theoretic context. Appendix C discusses assumptions for infinite-dimensional parameters. Appendix D contains derivations for the translog cost function example. Acronyms used include those for negative semidefinite (NSD), posterior expected loss (PEL), rejection probability (RP), and first-order stochastic dominance (SD1). Notationally, \subseteq is subset and \subset is proper subset; scalars, (column) vectors, and matrices are respectively formatted as X , \mathbf{X} , and $\underline{\mathbf{X}}$; $0(\cdot)$ denotes the zero function, i.e., $0(t) = 0$ for all t .

2 Results

2.1 Results for one-dimensional parameters

Assumption A1 states the sampling and posterior distributions of the (limit) experiment we consider in the scalar case. The sampling distribution conditions on a fixed parameter θ and treats the data X as random, whereas the posterior distribution conditions on a fixed X and treats θ as random.

Assumption A1. Let $F(\cdot)$ be a continuous CDF with support \mathbb{R} and symmetry $F(-x) = 1 - F(x)$. Let $\theta \in \mathbb{R}$ denote the parameter and $X \in \mathbb{R}$ denote the (lone) observation. The sampling distribution is $X - \theta \mid \theta \sim F$, and the posterior distribution is $\theta - X \mid X \sim F$.

Assumption A1 can be interpreted as a limit experiment where θ is a local mean parameter. Usually F is $N(0, \sigma^2)$, satisfying the continuity, support, and symmetry conditions in A1. For example, if $Y_{ni} \stackrel{iid}{\sim} N(\mu_n, 1)$, $i = 1, \dots, n$, and $\sqrt{n}\mu_n \rightarrow \theta$, then $\sqrt{n}\bar{Y}_n = n^{-1/2} \sum_{i=1}^n Y_{ni} \xrightarrow{d} N(\theta, 1)$; more generally, if $Y_{ni} \stackrel{iid}{\sim} N(m + \mu_n, \sigma^2)$ and $\sqrt{n}\mu_n \rightarrow \theta$, then $\sqrt{n}(\bar{Y}_n - m)/\hat{\sigma} \xrightarrow{d} N(\theta, 1)$ for any consistent estimator $\hat{\sigma}^2 \xrightarrow{p} \sigma^2$. This type of result holds for a wide variety of models, estimators, and sampling assumptions; it is most commonly used for local power analysis but has been used for purposes like ours in papers like Andrews and Soares (2010, eqn. (4.2)). Since θ is the *local* mean parameter, assuming $\theta \in \mathbb{R}$ does not require that the original parameter space (e.g., for $m + \mu_n$ in the example) is \mathbb{R} , but it does exclude boundary points. Results for posterior asymptotic normality date back to Laplace (1820), as cited in Lehmann and Casella (1998, §6.10, p. 515).

Seeing A1 as a limit experiment, implicitly the prior has no asymptotic effect on the posterior, as in the Bernstein–von Mises theorem. In the original parameter space, this is true if the prior density is continuous and positive at the true value (Hirano and Porter, 2009, p. 1696). In the limit experiment, this is equivalent to using an improper uninformative prior. For example, with sampling distribution $X \mid \theta \sim N(\theta, 1)$ and prior $\theta \sim N(m, \tau^2)$, the posterior is

$$\theta \mid X \sim N\left(\frac{\tau^2 X + m}{\tau^2 + 1}, \frac{\tau^2}{\tau^2 + 1}\right),$$

and taking $\tau^2 \rightarrow \infty$ yields the posterior $\theta \mid X \sim N(X, 1)$, satisfying A1.

To formally compare Bayesian and frequentist inference, we use the (frequentist) size of a Bayesian hypothesis test.⁴ Specifically, the Bayesian test rejects the null hypothesis H_0 if and only if the posterior probability of H_0 is below α , i.e., iff $P(H_0 \mid X) \leq \alpha$. In addition to being intuitive, there are decision-theoretic reasons to use this test for comparison; see Appendix B.

Method 1 (Bayesian test). Reject H_0 if $P(H_0 \mid X) \leq \alpha$; otherwise, accept H_0 .

Theorem 1 states our first new result.

Theorem 1. *Let Assumption A1 hold. Consider testing $H_0 : \theta \in \Theta_0$ against $H_1 : \theta \notin \Theta_0$ with the Bayesian test in Method 1, where $\Theta_0 \subset \mathbb{R}$.*

- (i) *If $\Theta_0 = (-\infty, c_0]$, then the Bayesian test has size α , and the type I error rate is exactly α when $\theta = c_0$.*

⁴This is only for comparison; in practice, reporting a posterior probability provides more information.

- (ii) Let $\Theta_0 \subseteq (-\infty, c_0]$ be Borel measurable, and let c_0 belong to the closure of Θ_0 : $c_0 \in \bar{\Theta}_0$. Then, the Bayesian test has size $\geq \alpha$.
- (iii) Continuing from (ii), if additionally the set $\{\theta : \theta \leq c_0, \theta \notin \Theta_0\}$ has positive Lebesgue measure and F has a continuous, strictly positive PDF $f(\cdot)$, then the Bayesian test's RP is $> \alpha$ when $\theta = c_0$, and its size is $> \alpha$.
- (iv) If Θ_0 is not a subset of a half-line, then the Bayesian test's size may be greater than or less than α .

The above results also hold with $[c_0, \infty)$ replacing $(-\infty, c_0]$.

Similar results to (i) are found in the literature, like in Casella and Berger (1987a); special cases of (ii) and (iii) have also been given.

Intuitively, Theorem 1(i) holds by the symmetry of F . Parts (ii) and (iii) hold because when parts of the half-line are carved away to make Θ_0 smaller, the posterior probability of H_0 (at any X) becomes smaller, making the Bayesian test more likely to reject. Part (iv) holds only with a pathological example for size above α ; restricting Θ_0 to be a finite union of intervals, the Bayesian test's size is always below α if Θ_0 is not contained by a half-line. In higher dimensions, however, the result holds without resorting to pathological examples.

Beyond intrinsic interest, the scalar results help prove results in higher dimensions in Section 2.2.

2.2 Results for multi-dimensional parameters

Assumption A2. Let \mathbf{X} and $\boldsymbol{\theta}$ belong to a Banach space of possibly infinite dimension. Let $\phi(\cdot)$ denote a continuous linear functional, with sampling distribution $\phi(\mathbf{X}) - \phi(\boldsymbol{\theta}) \mid \boldsymbol{\theta} \sim F$ and posterior $\phi(\boldsymbol{\theta}) - \phi(\mathbf{X}) \mid \mathbf{X} \sim F$, where F has the properties described in A1.

When the Banach space is \mathbb{R}^d , continuous linear functionals are simply linear combinations $\mathbf{c}'\mathbf{X}$ for some constant vector $\mathbf{c} \in \mathbb{R}^d$. In the leading case of multivariate normality of \mathbf{X} and $\boldsymbol{\theta}$, linear combinations are (scalar) normal random variables, satisfying the assumption. More generally, including infinite-dimensional spaces, if $X(\cdot)$ is a Gaussian process in some Banach space and $\phi(\cdot)$ belongs to the dual of that space, then $\phi(X(\cdot))$ is a scalar normal random variable; e.g., see Definition 2.2.1(ii) in Bogachev (1998, p. 42) and van der Vaart and Wellner (1996, pp. 376–377).

Lower-level conditions sufficient for equivalent, asymptotically multivariate normal sampling and posterior distributions, including for semiparametric models like GMM and quantile regression, are given and discussed in Hahn (1997, Thm. G and footnote 13), Kwan

(1999, Thm. 2), Kim (2002, Prop. 1), and Sims (2010, Sec. III.2), among others. The infinite-dimensional case is discussed more in Appendix C.

Theorem 2 generalizes Theorem 1 to multiple dimensions.

Theorem 2. *Let Assumption A2 hold. Consider testing $H_0 : \boldsymbol{\theta} \in \Theta_0$ against $H_1 : \boldsymbol{\theta} \notin \Theta_0$ with the Bayesian test in Method 1, where Θ_0 is a subset of the Banach space in A2.*

- (i) *If $\Theta_0 = \{\boldsymbol{\theta} : \phi(\boldsymbol{\theta}) \leq c_0\}$, then the Bayesian test has size α , and the type I error rate is exactly α when $\phi(\boldsymbol{\theta}) = c_0$.*
- (ii) *Let $\Theta_0 \subseteq \{\boldsymbol{\theta} : \phi(\boldsymbol{\theta}) \leq c_0\}$ be measurable with respect to the sampling and posterior distributions, and let $c_0 \in \phi(\bar{\Theta}_0)$. Then, the Bayesian test has size $\geq \alpha$.*
- (iii) *Continuing from (ii), additionally assume that the Banach space is \mathbb{R}^d . If (a) the set $\{\boldsymbol{\theta} : \phi(\boldsymbol{\theta}) \leq c_0, \boldsymbol{\theta} \notin \Theta_0\}$ has positive Lebesgue measure, (b) the sampling distribution (given any $\boldsymbol{\theta}$) and posterior distribution (given any \mathbf{X}) have strictly positive PDFs, and (c) $P(\boldsymbol{\theta} \in \Theta_0 \mid \mathbf{X})$ is continuous in \mathbf{X} , then the Bayesian test's RP is $> \alpha$ when $\phi(\boldsymbol{\theta}) = c_0$, and its size is $> \alpha$.*
- (iv) *If Θ_0 is not a subset of any half-space $\{\boldsymbol{\theta} : \phi(\boldsymbol{\theta}) \leq c_0\}$, then the Bayesian test's size may be greater than or less than α .*

Theorem 2(ii) has a geometric interpretation: if Θ_0 has a supporting hyperplane, then the Bayesian test's size is at least α . Theorem 2(iii) gives sufficient conditions for the Bayesian test's RP to be strictly above α for any $\boldsymbol{\theta}$ that is a support point of Θ_0 .

Theorem 2(iii) is difficult to formally extend to infinite-dimensional spaces. The difficulty is in establishing that the (set) difference between the half-space and Θ_0 has strictly positive probability under certain distributions. Still, the intuition is the same as in \mathbb{R}^d : if enough of the half-space is removed, the posterior probability of Θ_0 falls below the probability of the half-space, so the rejection region expands.

Theorem 2(iii) includes a special case explored by Kline (2011). Let $\boldsymbol{\theta} \in \mathbb{R}^d$, with $H_0 : \boldsymbol{\theta} \geq \mathbf{0}$ (elementwise) against $H_1 : \boldsymbol{\theta} \not\geq \mathbf{0}$ (i.e., at least one element $\theta_j < 0$). Kline (2011, p. 3136) explains the possible divergence of Bayesian and frequentist conclusions as the dimension d grows, when the distribution is multivariate normal with identity matrix covariance. He gives the example of observing $\mathbf{X} = \mathbf{0}$, where for large d the Bayesian $P(H_0 \mid \mathbf{X} = \mathbf{0}) \approx 0$ while the frequentist p -value is near one. Inverting his example illustrates Theorem 2(iv). If H_0 and H_1 are switched to get $H_0 : \boldsymbol{\theta} \not\geq \mathbf{0}$ and $H_1 : \boldsymbol{\theta} \geq \mathbf{0}$, then the divergence is in the opposite direction: $P(H_0 \mid \mathbf{X} = \mathbf{0}) \approx 1$, and large $P(H_0 \mid \mathbf{X})$ can occur even when the p -value is near zero. For example, the point $\mathbf{X} = (1.64, 1.64, \dots, 1.64) \in \mathbb{R}^d$

is the corner of the rejection region for the likelihood ratio test with size $\alpha = 0.05$, but the corresponding $P(H_0 | \mathbf{X}) = 0.40$ when $d = 10$, 0.72 when $d = 25$, and 0.99 when $d = 90$.

Theorem 2(iii) and the example in Kline (2011) are partly a result of the prior $P(H_0)$ being small when Θ_0 is small. That is, the prior over the *parameter* is always the same (regardless of Θ_0), so the implicit prior $P(H_0)$ shrinks when Θ_0 shrinks. (Technically, the limit experiment’s prior is an improper constant prior, so $P(H_0)$ is not well-defined, but the qualitative idea remains.) Unless Θ_0 is a half-space, this differs from Berger and Sellke (1987) and others who only consider “objective” priors with $P(H_0) = 0.5$. Whether placing a prior on the null (like $P(H_0) = 0.5$) or on the parameter is more appropriate depends on the empirical setting; e.g., do we have prior reason to suspect SD1? Often it is easier computationally not to set a specific $P(H_0)$; for example, one may use the same Bayesian bootstrap posterior advocated by Chamberlain and Imbens (2003) to easily compute probabilities of many different hypotheses. However, for hypotheses like SD1, this can lead to a very small “ $P(H_0)$ ” and consequently very large rejection probabilities (and size distortion).

Although the implicit $P(H_0)$ partially explains Theorem 2(iii), the shape of Θ_0 still plays an important role. For example, in \mathbb{R}^2 , let $\boldsymbol{\theta} = (\theta_1, \theta_2)$ and $H_0 : \theta_1\theta_2 \geq 0$, so Θ_0 comprises the first and third quadrants (and thus is not contained in any half-space). This Θ_0 is the same “size” as the half-space $\{\boldsymbol{\theta} : \theta_1 \geq 0\}$. However, with bivariate normal sampling and posterior distributions, Theorem 2(i) implies the Bayesian test of the half-space has exact size α , whereas the size of the Bayesian test of $H_0 : \theta_1\theta_2 \geq 0$ may be strictly above or below α , depending on the correlation. For example, let $\mathbf{X} = (X_1, X_2)$ have a bivariate normal sampling distribution with $\text{Corr}(X_1, X_2) = -1$. Then the test is equivalent to a scalar test where H_0 is a finite, closed interval, in which case the Bayesian test’s size strictly exceeds α by Theorem 1(iii). The same holds for other negative correlations, but size dips below α as the correlation gets closer to zero (and is strictly below α at zero correlation and positive correlations); see Kaplan (2015, §3.2) for details. This example also shows that, unlike for the one-dimensional Theorem 1(iv), there are interesting, non-pathological examples for Theorem 2(iv) where size is strictly above α .

Theorem 2 includes linear inequalities in \mathbb{R}^d as a special case. Part (i) states that for a single linear inequality $H_0 : \mathbf{c}'\boldsymbol{\theta} \leq c_0$, the Bayesian test has size α . Part (iii) states that for multiple linear inequalities, the Bayesian test’s size is strictly above α , and its RP is strictly above α at every boundary point of Θ_0 .

Whether Θ_0 is treated as H_0 or H_1 affects the size of the Bayesian test: if Θ_0 satisfies Theorem 2(iii), then its complement does not. Combining (iii) and (iv), this means the Bayesian test of $H_0 : \boldsymbol{\theta} \in \Theta_0$ may have size strictly above α while the Bayesian test of $H_0 : \boldsymbol{\theta} \in \Theta_0^c$ has size strictly below α , or vice-versa. This is indeed the case for SD1, as seen

in Section 3.1.

Many nonlinear inequalities could be recast as linear inequalities, but often with unnecessary additional approximation error. Beyond cases where reparameterization is possible, one could argue that the delta method usually implies asymptotic normality of functions of estimators, so nonlinear equalities may be written as linear inequalities for (nonlinear) functions of parameters. Although true, the delta method adds a layer of approximation error between the theoretical results and realistic finite-sample settings; we remove this layer by treating nonlinear inequalities directly.

Similarly, interpreting $n \rightarrow \infty$ literally would unnecessarily restrict the class of possible shapes of Θ_0 . For example, looking at a shrinking $n^{-1/2}$ neighborhood of a point in Θ , we could never have a disk-shaped Θ_0 for the local parameter. (Or, we would need a *sequence* of Θ_0 with radius proportional to $n^{-1/2}$.) However, $n \rightarrow \infty$ is only a tool for approximation; considering unrestricted Θ_0 should more accurately reflect finite-sample settings.

3 Examples

3.1 Example: first-order stochastic dominance

The example of testing first-order stochastic dominance (SD1) illustrates some of the results from Section 2. Let $X_i \stackrel{iid}{\sim} F_X(\cdot)$, $Y_i \stackrel{iid}{\sim} F_Y(\cdot)$, and $F_0(\cdot)$ is non-random, where all distributions are continuous. One-sample SD1 is $F_X(\cdot) \leq F_0(\cdot)$; two-sample SD1 is $F_X(\cdot) \leq F_Y(\cdot)$.

Theorem 2 implies the Bayesian test's asymptotic size is $\geq \alpha$ when the null hypothesis is SD1, while its size may be $< \alpha$ if the null is non-SD1. Consider the one-sample setup with $X_i \stackrel{iid}{\sim} F(\cdot)$. Implicitly taking $F(\cdot)$ or $F_0(\cdot)$ to be a drifting sequence, let $\sqrt{n}(F(\cdot) - F_0(\cdot)) \rightarrow \theta(\cdot)$, the local parameter. SD1 of F over F_0 is $F(\cdot) \leq F_0(\cdot)$, or equivalently $\theta(\cdot) \leq 0(\cdot)$. Since $\sqrt{n}(\hat{F}(\cdot) - F(\cdot)) \rightsquigarrow B(F(\cdot))$ for standard Brownian bridge $B(\cdot)$,

$$X(\cdot) \equiv \sqrt{n}(\hat{F}(\cdot) - F_0(\cdot)) = \sqrt{n}(\hat{F}(\cdot) - F(\cdot)) + \sqrt{n}(F(\cdot) - F_0(\cdot)) \rightsquigarrow B(F(\cdot)) + \theta(\cdot),$$

so the limit experiment is $X(\cdot) - \theta(\cdot) \mid \theta(\cdot) \sim B(F(\cdot))$. Note $B(F(\cdot))$ is a mean-zero Gaussian process with covariance function $\text{Cov}(t_1, t_2) = F(t_1)[1 - F(t_2)]$ for $t_1 \leq t_2$. Although $F(\cdot)$ is unknown, $\hat{F}(\cdot) \xrightarrow{a.s.} F(\cdot)$ uniformly by the Glivenko–Cantelli theorem, so asymptotically the covariance is known while $\theta(\cdot)$ remains unknown. Letting $\phi(\theta(\cdot)) = \theta(x)$ for some given $x \in \mathbb{R}$, $\Theta_0 \equiv \{\theta(\cdot) : \theta(\cdot) \leq 0(\cdot)\} \subset \{\theta(\cdot) : \phi(\theta(\cdot)) \leq 0\}$, satisfying the condition of Theorem 2(ii). The complement Θ_0^c satisfies Theorem 2(iv) instead. When the null hypothesis is SD1, i.e., $H_0 : \theta(\cdot) \in \Theta_0$, the intuition is similar to that for tests of $H_0 : \boldsymbol{\theta} \leq \mathbf{0}$ with $\boldsymbol{\theta} \in \mathbb{R}^d$ for large $d < \infty$ as in Kline (2011), where the Bayesian test has higher RP.

When the null is non-SD1, the intuition is reversed, as are the results below.

In Table 1, we compare frequentist p -values with the Bayesian posterior probability of H_0 in a particular dataset. For a given h , we set $X_i = i/(n + 1) + hn^{-1/2}$ for $i = 1, \dots, n$. This sample is compared with either the standard uniform distribution or a second sample with $Y_i = i/n$ for $i = 1, \dots, n - 1$. When the null is SD1, i.e., testing $H_0 : F_X(\cdot) \leq F_Y(\cdot)$ against $H_1 : F_X(\cdot) \not\leq F_Y(\cdot)$, we use the KS p -values from `ks.test` in R (R Core Team, 2017). When the null is non-SD1, i.e., testing $H_0 : F_X(\cdot) \not\leq F_Y(\cdot)$ against $H_1 : F_X(\cdot) \leq F_Y(\cdot)$, we use the two-sample p -value from Davidson and Duclos (2013) and a one-sample p -value from an intersection–union test based on Goldman and Kaplan (2016). For the Bayesian posterior probabilities, the Bayesian bootstrap variant of Banks (1988) is used. Details may be seen in the provided code.

Table 1: Frequentist p -values and Bayesian posterior probabilities of H_0 .

H_0	n	h	X (non)SD1 Unif(0, 1)		X (non)SD1 Y	
			frequentist	Bayesian	frequentist	Bayesian
SD1	100	0.0	0.981	0.009	0.990	0.010
SD1	1000	0.0	0.998	0.000	0.999	0.000
non-SD1	100	0.0	0.630	0.991	1.000	0.988
non-SD1	100	0.5	0.157	0.526	0.020	0.688
non-SD1	100	0.9	0.035	0.165	0.015	0.356
non-SD1	1000	0.0	0.632	0.998	1.000	0.998
non-SD1	1000	0.5	0.159	0.587	0.015	0.729
non-SD1	1000	0.9	0.036	0.175	0.010	0.410

Table 1 illustrates how the shape of H_0 affects differences between Bayesian and frequentist inference, consistent with our theoretical results. When the null is SD1, the subspace of distribution functions satisfying H_0 has a very sharp “corner” at $F_0(\cdot)$, or equivalently at $\theta(\cdot) = 0(\cdot)$. Consequently, when $\hat{F}_X(\cdot) \approx F_0(\cdot)$, or when $\hat{F}_X(\cdot) \approx \hat{F}_Y(\cdot)$, the Bayesian posterior places nearly zero probability on H_0 . In the limit experiment, when $X(\cdot) = 0(\cdot)$, the posterior of $\theta(\cdot)$ is $B(F(\cdot))$, so $P(\theta(\cdot) \leq 0(\cdot) \mid X(\cdot) = 0(\cdot)) = P(B(F(\cdot)) \leq 0(\cdot)) = P(B(\cdot) \leq 0(\cdot)) = 0$; Table 1 shows zero up to a few decimal places already at $n = 1000$. In stark contrast, the frequentist p -value is near one when the estimated $\hat{F}_X(\cdot)$ is near $F_0(\cdot)$ or $\hat{F}_Y(\cdot)$. These results are qualitatively similar to those for the one-sample, finite-dimensional example in Kline (2011, §4).

Table 1 also shows that when H_0 is non-SD1, the results reverse. The previous arguments still apply regarding when the posterior of SD1 is very small, but now SD1 is H_1 . The frequentist test is more skeptical about H_0 than the Bayesian test, for both one-sample

and two-sample inference: the frequentist p -values are always significantly lower than the Bayesian posterior probabilities of H_0 across a range of h .⁵ Larger n only magnifies the difference.

Table 2: Frequentist and Bayesian rejection probabilities, $\alpha = 0.1$, 1000 replications each.

H_0	n	h	X over Unif(0, 1)		X over Y	
			frequentist	Bayesian	frequentist	Bayesian
SD1	100	0.0	0.098	0.980	0.080	0.975
SD1	1000	0.0	0.103	1.000	0.094	1.000
non-SD1	100	0.0	0.000	0.000	0.002	0.000
non-SD1	100	0.9	0.349	0.185	0.281	0.040
non-SD1	100	1.3	0.683	0.566	0.475	0.195
non-SD1	1000	0.0	0.000	0.000	0.000	0.000
non-SD1	1000	0.9	0.295	0.128	0.278	0.023
non-SD1	1000	1.3	0.674	0.515	0.521	0.163

In Table 2, we compare rejection probabilities of the Bayesian and frequentist tests. The DGPs have $X_i \stackrel{iid}{\sim} \text{Unif}(hn^{-1/2}, 1 + hn^{-1/2})$ for $i = 1, \dots, n$, $Y_i \stackrel{iid}{\sim} \text{Unif}(0, 1)$ for $i = 1, \dots, n$ (for two-sample inference), and $F_0(\cdot)$ is the Unif(0, 1) CDF. The hypotheses, methods, and notation are the same as for Table 1.

Table 2 shows the same patterns as Table 1. When H_0 is SD1 and $h = 0$, the Bayesian type I error rate is nearly 100%, whereas the frequentist tests control size at $\alpha = 0.1$. When H_0 is non-SD1, although no tests reject when $h = 0$, the frequentist tests have much steeper power curves as h increases.⁶ As in Table 1, the differences do not diminish with larger n .

3.2 Example: curvature constraints

One common nonlinear inequality hypothesis in economics is a “curvature” constraint like concavity. Such constraints come from economic theory, often the second-order condition of an optimization problem like utility maximization or cost minimization. As noted by O’Donnell and Coelli (2005), the Bayesian approach is appealing for imposing or testing curvature constraints due to its (relative) simplicity. However, according to Theorem 2,

⁵For two-sample, non-SD1 testing, we also tried an intersection–union max- t test similar to Kaur, Prakasa Rao, and Singh (1994); the p -values are larger but still consistently below the Bayesian posterior probabilities, with p -values of 0.717, 0.263, and 0.114 with $n = 100$ and $h \in \{0, 0.5, 0.9\}$ (respectively), and 0.718, 0.244, and 0.109 with $n = 1000$ and $h \in \{0, 0.5, 0.9\}$ (respectively).

⁶Although less powerful than the test of Davidson and Duclos (2013), the intersection–union t -test (not shown in table) still has better power than the Bayesian test, with rejection probabilities 0.089 and 0.294 with $n = 100$ and $h \in \{0.9, 1.3\}$ (respectively), and 0.084 and 0.294 (again) with $n = 1000$ and $h \in \{0.9, 1.3\}$ (respectively).

since curvature is usually satisfied in a parameter subspace much smaller than a half-space, standard Bayesian inference may be much less favorable toward the curvature hypothesis than frequentist inference would be; i.e., the size of the Bayesian test in Method 1 may be well above α . For example, in Table 3 below, the Bayesian test rejects well over 50% of the time given parameter values just inside the boundary of a curvature constraint. Our simulation example concerns concavity of cost in input prices.

Our example uses a cost function with the “translog” functional form (Christensen, Jorgenson, and Lau, 1973). This has been a popular way to parameterize cost, indirect utility, and production functions, among others. The translog is more flexible than many traditional functional forms, allowing violation of certain implications of economic theory, such as curvature, without reducing such constraints to the value of a single parameter. Since Lau (1978), there has been continued interest in methods to impose curvature constraints during estimation, as well as methods to test such constraints. Although “flexible,” the translog is still parametric, so violation of curvature constraints may come from misspecification (of the functional form) rather than violation of economic theory.⁷

Our example concerns concavity of a translog cost function in input prices.⁸ With output y , input prices $\mathbf{w} = (w_1, w_2, w_3)$, and total cost $C(y, \mathbf{w})$, the translog model is

$$\begin{aligned} \ln(C(y, \mathbf{w})) = & a_0 + a_y \ln(y) + (1/2)a_{yy}[\ln(y)]^2 + \sum_{k=1}^3 a_{yk} \ln(y) \ln(w_k) \\ & + \sum_{k=1}^3 b_k \ln(w_k) + (1/2) \sum_{k=1}^3 \sum_{m=1}^3 b_{km} \ln(w_k) \ln(w_m). \end{aligned} \quad (1)$$

Standard economic assumptions imply that $C(y, \mathbf{w})$ is concave in \mathbf{w} (as in Kreps, 1990, §7.3), which corresponds to the Hessian matrix (of C with respect to \mathbf{w}) being negative semidefinite (NSD), which in turn corresponds to all the Hessian’s principal minors of order p (for all $p = 1, 2, 3$) having the same sign as $(-1)^p$ or zero.

For simplicity, we consider local concavity at the point $(1, 1, 1, 1)$:

$$H_0 : \underline{\mathbf{H}} \equiv \left. \frac{\partial^2 C(y, \mathbf{w})}{\partial \mathbf{w} \partial \mathbf{w}'} \right|_{(y, \mathbf{w})=(1,1,1,1)} \text{ is NSD.} \quad (2)$$

This is necessary but not sufficient for global concavity; rejecting local concavity implies rejection of global concavity. In Appendix D, we show that even this weaker constraint

⁷With a nonparametric model, one may more plausibly test the theory itself, although there are always other assumptions that may be violated; see Dette, Hoderlein, and Neumeyer (2016) for nonparametrically testing negative semidefiniteness of the Slutsky substitution matrix.

⁸The “translog” example on page 346 of Dufour (1989) is even simpler but appears to ignore the fact that second derivatives are not invariant to log transformations.

corresponds to a set of parameter values much smaller than a half-space, so Theorem 2(iii) applies.

Our simulation DGP is as follows. To impose homogeneity of degree one in input prices, we use the normalized model (with error term ϵ added)

$$\begin{aligned} \ln(C/w_3) = & a_0 + a_y \ln(y) + (1/2)a_{yy}[\ln(y)]^2 + \sum_{k=1}^2 a_{yk} \ln(y) \ln(w_k/w_3) \\ & + \sum_{k=1}^2 b_k \ln(w_k/w_3) + (1/2) \sum_{k=1}^2 \sum_{m=1}^2 b_{km} \ln(w_k/w_3) \ln(w_m/w_3) + \epsilon \end{aligned} \quad (3)$$

for both data generation and inference.⁹ The parameter values are $b_1 = b_2 = 1/3$, $b_{11} = b_{22} = 2/9 - \delta$ (more on δ below), and $b_{12} = -1/9$ to make some of the inequality constraints in H_0 close to binding, as well as $a_0 = 1$, $a_y = 1$, $a_{yy} = 0$, $a_{yk} = 0$. The other parameter values follow from imposing symmetry ($b_{km} = b_{mk}$) and homogeneity. When $\delta = 0$, \mathbf{H} is a matrix of zeros, on the boundary of being NSD in that each principal minor equals zero (and none are strictly negative). When $\delta > 0$, all principal minors are strictly negative (other than $\det(\mathbf{H}) = 0$, which is always true under homogeneity). We set $\delta = 0.001$. In each simulation replication, an iid sample is drawn, where $\ln(y)$ and all $\ln(w_k)$ are $N(0, \sigma = 0.1)$, $\epsilon \sim N(0, \sigma_\epsilon)$, and all variables are mutually independent. There are $n = 100$ observations per sample, 500 simulation replications, and 200 posterior draws per replication. The local monotonicity constraints $b_1, b_2, b_3 \geq 0$ were satisfied in 100.0% of replications overall.

The posterior probability of H_0 is computed by a nonparametric Bayesian method with improper Dirichlet process prior, i.e., the Bayesian bootstrap of Rubin (1981) based on Ferguson (1973) and more recently advocated in economics by Chamberlain and Imbens (2003). To accommodate numerical imprecision, we deem an inequality satisfied if it is within 10^{-7} . The simulated type I error rate is the proportion of simulated samples for which the posterior probability of H_0 was below α .

Table 3 shows the type I error rate of the Bayesian bootstrap test of (2) given our DGP. The values of α and σ_ϵ are varied as shown in the table. As a sanity check, when $\sigma_\epsilon = 0$, the RP is zero since the constraints are satisfied by construction. As σ_ϵ increases, as suggested by Theorem 2, the RP increases well above α , even over 50%.¹⁰ Although the Bayesian test's size distortion with the null of local NSD is clearly bad from a frequentist perspective, it reflects the Bayesian method's need for great evidence to conclude in favor of local NSD,

⁹Alternatively, cost share equations may be used. Shephard's lemma implies that the demand for input k is $x_k = \partial C / \partial w_k$. The cost share for input k is then $s_k = x_k w_k / C = (\partial C / \partial w_k)(w_k / C) = \partial \ln(C) / \partial \ln(w_k) \equiv r_k = b_k + a_{yk} \ln(y) + \sum_{j=1}^3 b_{jk} \ln(w_j)$.

¹⁰The results with $\delta = 0.01$ and $\sigma_\epsilon \in [0, 1]$ are similar to Table 3; with $\delta = 0$, RP jumps to over 80% even with $\sigma_\epsilon = 0.001$.

Table 3: Simulated type I error rate of the Bayesian bootstrap test of local NSD.

σ_ϵ	Type I error rate	
	$\alpha = 0.05$	$\alpha = 0.10$
0.00	0.000	0.000
0.10	0.090	0.172
0.20	0.360	0.546
0.30	0.574	0.756
0.40	0.660	0.810
0.50	0.734	0.882

which may be reasonable since the translog form does not come from economic theory and since only a small part of the parameter space satisfies local NSD. Either way, it is helpful to understand the divergent behavior of Bayesian and frequentist inference in this situation.

4 Conclusion

We have explored the difference between Bayesian and frequentist inference on general nonlinear inequality constraints, providing formal results on the role of the shape of the null hypothesis parameter subspace. The examples of first-order stochastic dominance and local curvature constraints illustrate our results; either the Bayesian or frequentist test may have higher rejection probability depending how the labels H_0 and H_1 are assigned. We have separate work in progress detailing nonparametric Bayesian inference for first-order and higher-order stochastic dominance. Investigation of approaches like Müller and Norets (2016) applied to nonlinear inequality testing remains for future work. It would also be valuable to extend this paper’s analysis to allow priors with $P(H_0) = 1/2$ or other values.

References

- Andrews, D. W. K. and G. Soares (2010). Inference for parameters defined by moment inequalities using generalized moment selection. *Econometrica* 78(1), 119–157.
- Banks, D. L. (1988). Histospline smoothing the Bayesian bootstrap. *Biometrika* 75(4), 673–684.
- Berger, J. O. (2003). Could Fisher, Jeffreys and Neyman have agreed on testing? *Statistical Science* 18(1), 1–32.
- Berger, J. O. and T. Sellke (1987). Testing a point null hypothesis: The irreconcilability of p values and evidence. *Journal of the American Statistical Association* 82(397), 112–122.
- Bogachev, V. I. (1998). *Gaussian Measures*, Volume 62 of *Mathematical Surveys and Monographs*. American Mathematical Society.

- Casella, G. and R. L. Berger (1987a). Reconciling Bayesian and frequentist evidence in the one-sided testing problem. *Journal of the American Statistical Association* 82(397), 106–111.
- Casella, G. and R. L. Berger (1987b). Testing precise hypotheses: Comment. *Statistical Science* 2(3), 344–347.
- Chamberlain, G. and G. W. Imbens (2003). Nonparametric applications of Bayesian inference. *Journal of Business & Economic Statistics* 21(1), 12–18.
- Christensen, L. R., D. W. Jorgenson, and L. J. Lau (1973). Transcendental logarithmic production frontiers. *Review of Economics and Statistics* 55(1), 28–45.
- Davidson, R. and J.-Y. Duclos (2013). Testing for restricted stochastic dominance. *Econometric Reviews* 32(1), 84–125.
- Dette, H., S. Hoderlein, and N. Neumeyer (2016). Testing multivariate economic restrictions using quantiles: The example of Slutsky negative semidefiniteness. *Journal of Econometrics* 191(1), 129–144.
- Dufour, J.-M. (1989). Nonlinear hypotheses, inequality restrictions, and non-nested hypotheses: Exact simultaneous tests in linear regressions. *Econometrica* 57(2), 335–355.
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Annals of Statistics* 1(2), 209–230.
- Freedman, D. (1999). On the Bernstein–von Mises theorem with infinite-dimensional parameters. *Annals of Statistics* 27(4), 1119–1140.
- Goldman, M. and D. M. Kaplan (2016). Comparing distributions by multiple testing across quantiles. Working paper, available at <http://faculty.missouri.edu/~kaplandm>.
- Goutis, C., G. Casella, and M. T. Wells (1996). Assessing evidence in multiple hypotheses. *Journal of the American Statistical Association* 91(435), 1268–1277.
- Hahn, J. (1997). Bayesian bootstrap of the quantile regression estimator: A large sample study. *International Economic Review* 38(4), 795–808.
- Hirano, K. and J. R. Porter (2009). Asymptotics for statistical treatment rules. *Econometrica* 77(5), 1683–1701.
- Imbens, G. W. and C. F. Manski (2004). Confidence intervals for partially identified parameters. *Econometrica* 72(6), 1845–1857.
- Kaplan, D. M. (2015). Bayesian and frequentist tests of sign equality and other nonlinear inequalities. Working paper, available at <http://faculty.missouri.edu/~kaplandm>.
- Kaur, A., B. L. S. Prakasa Rao, and H. Singh (1994). Testing for second-order stochastic dominance of two distributions. *Econometric Theory* 10(5), 849–866.
- Kim, J.-Y. (2002). Limited information likelihood and Bayesian analysis. *Journal of Econometrics* 107(1), 175–193.
- Kline, B. (2011). The Bayesian and frequentist approaches to testing a one-sided hypothesis about a multivariate mean. *Journal of Statistical Planning and Inference* 141(9), 3131–3141.
- Kreps, D. M. (1990). *A Course in Microeconomic Theory*. Princeton University Press.
- Kwan, Y. K. (1999). Asymptotic Bayesian analysis based on a limited information estimator. *Journal of Econometrics* 88(1), 99–121.
- Laplace, P.-S. (1820). *Théorie Analytique des Probabilités* (3rd ed.). Paris: V. Courcier.
- Lau, L. J. (1978). Testing and imposing monotonicity, convexity, and quasi-convexity constraints. In M. Fuss and D. McFadden (Eds.), *Production Economics: A Dual Approach*

- to *Theory and Applications*, Volume 1 of *Contributions to Economic Analysis*, Chapter A.4, pp. 409–453. Amsterdam: North-Holland.
- Lehmann, E. L. and G. Casella (1998). *Theory of Point Estimation* (2nd ed.). New York: Springer.
- Lehmann, E. L. and J. P. Romano (2005). *Testing Statistical Hypotheses* (3rd ed.). Springer Texts in Statistics. Springer.
- Lindley, D. V. (1957). A statistical paradox. *Biometrika* 44(1–2), 187–192.
- Lo, A. Y. (1983). Weak convergence for Dirichlet processes. *Sankhyā: The Indian Journal of Statistics, Series A* 45(1), 105–111.
- Lo, A. Y. (1987). A large sample study of the Bayesian bootstrap. *Annals of Statistics* 15(1), 360–375.
- Moon, H. R. and F. Schorfheide (2012). Bayesian and frequentist inference in partially identified models. *Econometrica* 80(2), 755–782.
- Müller, U. K. and A. Norets (2016). Credibility of confidence sets in nonstandard econometric problems. *Econometrica* 84(6), 2183–2213.
- O’Donnell, C. J. and T. J. Coelli (2005). A Bayesian approach to imposing curvature on distance functions. *Journal of Econometrics* 126(2), 493–523.
- R Core Team (2017). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Rubin, D. B. (1981). The Bayesian bootstrap. *Annals of Statistics* 9(1), 130–134.
- Sims, C. A. (2010). Understanding non-Bayesians. Unpublished book chapter, available at <http://sims.princeton.edu/yftp/UndrstndgNnBsns/GewekeBookChapter.pdf>.
- Stoye, J. (2009). More on confidence intervals for partially identified parameters. *Econometrica* 77(4), 1299–1315.
- van der Vaart, A. W. and J. A. Wellner (1996). *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Series in Statistics. New York: Springer.
- Wolak, F. A. (1991). The local nature of hypothesis tests involving inequality constraints in nonlinear models. *Econometrica* 59(4), 981–995.

A Mathematical proofs

A.1 Proof of Theorem 1

Proof. For (i), the Bayesian test rejects iff

$$\alpha \geq \text{P}(\theta \leq c_0 \mid X) = \text{P}(\theta - X \leq c_0 - X \mid X) \equiv F(c_0 - X).$$

Given any $\theta \leq c_0$ (so H_0 holds), the frequentist rejection probability (RP) is

$$\begin{aligned} \text{P}(F(c_0 - X) \leq \alpha \mid \theta) &= \text{P}(\overbrace{1 - F(X - c_0)}^{\text{by A1 symmetry}} \leq \alpha \mid \theta) \\ &= \text{P}(F(X - c_0) \geq 1 - \alpha \mid \theta) \end{aligned}$$

$$\begin{aligned} & \overbrace{\leq \text{P}(F(X - \theta) \geq 1 - \alpha \mid \theta)}^{\text{since } \theta \leq c_0 \text{ under } H_0} \\ & = \alpha \end{aligned}$$

since $F(X - \theta) \mid \theta \sim \text{Unif}(0, 1)$ (the ‘‘probability integral transform’’). If $\theta = c_0$, then the \leq becomes $=$.

For (ii), because $\Theta_0 \subseteq (-\infty, c_0]$, then for any X ,

$$\text{P}(\theta \in \Theta_0 \mid X) \leq \text{P}(\theta \leq c_0 \mid X).$$

Consequently, the rejection region for $H_0 : \theta \in \Theta_0$ is at least as big as the rejection region for $H_0 : \theta \leq c_0$: for some $r \in \mathbb{R}$,

$$\mathcal{R}_1 \subseteq \mathcal{R}_2, \mathcal{R}_1 \equiv \{X : \text{P}(\theta \leq c_0 \mid X) \leq \alpha\} = [r, \infty), \mathcal{R}_2 \equiv \{X : \text{P}(\theta \in \Theta_0 \mid X) \leq \alpha\}. \quad (4)$$

Given any $\theta \in \Theta_0$, the probability that X falls in the new, larger rejection region (\mathcal{R}_2) is at least as big as the probability that X falls in the old, smaller rejection region (\mathcal{R}_1) from (i). In particular, when $\theta = c_0$, the RP was exactly α in (i). Since the new rejection region is weakly larger, the new RP when $\theta = c_0$ must be $\geq \alpha$. If $c_0 \in \Theta_0$, then the proof is complete. Otherwise, with $\mathcal{R}_1 = [r, \infty)$ from (4),

$$\begin{aligned} \sup_{\theta \in \Theta_0} \text{P}(X \in \mathcal{R}_2 \mid \theta) & \geq \overbrace{\lim_{\theta \rightarrow c_0^-} \text{P}(X \in \mathcal{R}_2 \mid \theta)}^{\text{since } c_0 \in \bar{\Theta}_0} \geq \overbrace{\lim_{\theta \rightarrow c_0^-} \text{P}(X \in \mathcal{R}_1 \mid \theta)}^{\text{by (4)}} \\ & = \lim_{\theta \rightarrow c_0^-} \text{P}(X \in [r, \infty) \mid \theta) = \lim_{\theta \rightarrow c_0^-} \text{P}(X - \theta \geq r - \theta \mid \theta) = \lim_{\theta \rightarrow c_0^-} 1 - F(r - \theta) \overbrace{= 1 - F(r - c_0)}^{\text{by continuity of } F} \\ & = \text{P}(X \in \mathcal{R}_1 \mid \theta = c_0) \overbrace{= \alpha}^{\text{by Part (i)}}. \end{aligned}$$

For (iii), let $\Delta \equiv \{\theta : \theta \leq c_0, \theta \notin \Theta_0\}$. Given any X ,

$$\text{P}(\theta \in \Delta \mid X) = \text{P}(\theta - X \in \Delta - X \mid X) = \int_{\mathbb{R}} \mathbf{1}\{(t + X) \in \Delta\} f(t) dt > 0$$

since by assumption $f(t) > 0$ for all t and Δ has positive Lebesgue measure. Then, given any X ,

$$\text{P}(\theta \in \Theta_0 \mid X) = \text{P}(\theta \leq c_0 \mid X) - \overbrace{\text{P}(\theta \in \Delta \mid X)}^{>0} < \text{P}(\theta \leq c_0 \mid X).$$

In particular, this is true at $X = r$, where r is from $\mathcal{R}_1 = [r, \infty)$ in (4):

$$\mathbb{P}(\theta \in \Theta_0 \mid X = r) = \overbrace{\mathbb{P}(\theta \leq c_0 \mid X = r)}{=\alpha} - \overbrace{\mathbb{P}(\theta \in \Delta \mid X = r)}{>0} < \alpha.$$

Since $\mathbb{P}(\theta \in \Theta_0 \mid X = r) = \int_{\mathbb{R}} \mathbf{1}\{t \in \Theta_0\} f(t - X) dt$ is continuous in X , there exists some $\epsilon > 0$ such that $\mathbb{P}(\theta \in \Theta_0 \mid X = r - \epsilon) < \alpha$, too. Thus, the rejection region must be a superset of $[r - \epsilon, \infty)$: not only is it strictly larger than $[r, \infty)$, but the newly added part has positive Lebesgue measure. Consequently,

$$\begin{aligned} \mathbb{P}(\{X : \mathbb{P}(\theta \in \Theta_0 \mid X) \leq \alpha\} \mid \theta = c_0) &\geq \mathbb{P}(X \geq r - \epsilon \mid \theta = c_0) \\ &= \overbrace{\mathbb{P}(X \geq r \mid \theta = c_0)}{=\alpha} + \overbrace{\mathbb{P}(r - \epsilon \leq X < r \mid \theta = c_0)}{>0} > \alpha. \end{aligned}$$

The continuity of F then completes the proof for when $c_0 \in \bar{\Theta}_0$ but $c_0 \notin \Theta_0$: size is

$$\begin{aligned} \sup_{\theta \in \Theta_0} \text{RP}(\theta) &\geq \sup_{\theta \in \Theta_0} \mathbb{P}(X \geq r - \epsilon \mid \theta) \geq \lim_{\theta \rightarrow c_0^-} \mathbb{P}(X \geq r - \epsilon \mid \theta) = \lim_{\theta \rightarrow c_0^-} 1 - F(r - \epsilon - \theta) \\ &= 1 - F(r - \epsilon - c_0) = \overbrace{1 - F(r - c_0)}{=\alpha} + \overbrace{[F(r - c_0) - F(r - c_0 - \epsilon)]}_{>0} > \alpha. \end{aligned}$$

For (iv), two examples suffice. First, consider $\Theta_0 = \mathbb{Z}$, the integers, which are not contained in any half-line. Given any X , $\mathbb{P}(\theta \in \mathbb{Z} \mid X) = 0$ since F is continuous and \mathbb{Z} is countable. Thus, the Bayesian test always rejects, so its size equals one ($> \alpha$). Second, consider $\Theta_0 = \mathbb{R} \setminus \{0\}$. Given any X , $\mathbb{P}(\theta \in \Theta_0 \mid X) = 1 - \mathbb{P}(\theta = 0 \mid X) = 1$ since F is continuous. Thus, the Bayesian test always accepts, so its size equals zero ($< \alpha$).

□

A.2 Proof of Theorem 2

Proof. For (i), the argument parallels that for Theorem 1(i). The Bayesian test rejects iff

$$\alpha \geq \mathbb{P}(\phi(\boldsymbol{\theta}) \leq c_0 \mid \mathbf{X}) = \mathbb{P}(\phi(\boldsymbol{\theta}) - \phi(\mathbf{X}) \leq c_0 - \phi(\mathbf{X}) \mid \mathbf{X}) \equiv F(c_0 - \phi(\mathbf{X})).$$

Given any $\boldsymbol{\theta}$ such that $\phi(\boldsymbol{\theta}) \leq c_0$ (so H_0 holds), the RP is

$$\begin{aligned} \mathbb{P}(F(c_0 - \phi(\mathbf{X})) \leq \alpha \mid \boldsymbol{\theta}) &= \mathbb{P}\left(\overbrace{1 - F(\phi(\mathbf{X}) - c_0)}^{\text{by symmetry}} \leq \alpha \mid \boldsymbol{\theta}\right) \\ &= \mathbb{P}(F(\phi(\mathbf{X}) - c_0) \geq 1 - \alpha \mid \boldsymbol{\theta}) \\ &= \overbrace{\mathbb{P}(F(\phi(\mathbf{X}) - \phi(\boldsymbol{\theta})) \geq 1 - \alpha \mid \boldsymbol{\theta})}^{\text{since } \phi(\boldsymbol{\theta}) \leq c_0 \text{ under } H_0} \end{aligned}$$

$$= \alpha$$

since $F(\phi(\mathbf{X}) - \phi(\boldsymbol{\theta})) \mid \boldsymbol{\theta} \sim \text{Unif}(0, 1)$. If $\phi(\boldsymbol{\theta}) = c_0$, then the \leq becomes $=$.

For (ii), the argument parallels that for Theorem 1(ii). Because $\Theta_0 \subseteq \{\boldsymbol{\theta} : \phi(\boldsymbol{\theta}) \leq c_0\}$, then for any \mathbf{X} ,

$$\mathbb{P}(\boldsymbol{\theta} \in \Theta_0 \mid \mathbf{X}) \leq \mathbb{P}(\phi(\boldsymbol{\theta}) \leq c_0 \mid \mathbf{X}).$$

Consequently, the rejection region for $H_0 : \boldsymbol{\theta} \in \Theta_0$ is at least as big as the rejection region for $H_0 : \phi(\boldsymbol{\theta}) \leq c_0$: for some $r \in \mathbb{R}$,

$$\begin{aligned} \mathcal{R}_1 &\subseteq \mathcal{R}_2, \quad \mathcal{R}_1 \equiv \{\mathbf{X} : \mathbb{P}(\phi(\boldsymbol{\theta}) \leq c_0 \mid \mathbf{X}) \leq \alpha\} = \{\mathbf{X} : \phi(\mathbf{X}) \geq r\}, \\ \mathcal{R}_2 &\equiv \{\mathbf{X} : \mathbb{P}(\boldsymbol{\theta} \in \Theta_0 \mid \mathbf{X}) \leq \alpha\}. \end{aligned} \tag{5}$$

Given any $\boldsymbol{\theta} \in \Theta_0$, the probability that \mathbf{X} falls in the new, larger rejection region (\mathcal{R}_2) is at least as big as the probability that \mathbf{X} falls in the old, smaller rejection region (\mathcal{R}_1) from (i). In particular, when $\phi(\boldsymbol{\theta}) = c_0$, the RP was exactly α in (i). Since the new rejection region is weakly larger, the new RP when $\phi(\boldsymbol{\theta}) = c_0$ must be $\geq \alpha$. If $c_0 \in \phi(\Theta_0)$, then the proof is complete. Otherwise, with $\mathcal{R}_1, \mathcal{R}_2$ from (5), and $\boldsymbol{\theta}^*$ any value such that $\phi(\boldsymbol{\theta}^*) = c_0$ (with the limit formed by a sequence of $\boldsymbol{\theta}$ within Θ_0),

$$\begin{aligned} \sup_{\boldsymbol{\theta} \in \Theta_0} \mathbb{P}(\mathbf{X} \in \mathcal{R}_2 \mid \boldsymbol{\theta}) &\stackrel{\text{since } c_0 \in \phi(\bar{\Theta}_0)}{\geq} \lim_{\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}^*} \mathbb{P}(\mathbf{X} \in \mathcal{R}_2 \mid \boldsymbol{\theta}) \stackrel{\text{by (5)}}{\geq} \lim_{\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}^*} \mathbb{P}(\mathbf{X} \in \mathcal{R}_1 \mid \boldsymbol{\theta}) \\ &= \lim_{\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}^*} \mathbb{P}(\phi(\mathbf{X}) \geq r \mid \boldsymbol{\theta}) = \lim_{\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}^*} \mathbb{P}(\phi(\mathbf{X}) - \phi(\boldsymbol{\theta}) \geq r - \phi(\boldsymbol{\theta}) \mid \boldsymbol{\theta}) = \lim_{\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}^*} 1 - F(r - \phi(\boldsymbol{\theta})) \\ &\stackrel{\text{by continuity of } F, \phi}{=} 1 - F(r - c_0) \stackrel{\text{by (i)}}{=} \alpha. \end{aligned} \tag{6}$$

For (iii), the argument is similar to Theorem 1(iii). Let $\Delta \equiv \{\boldsymbol{\theta} : \phi(\boldsymbol{\theta}) \leq c_0, \boldsymbol{\theta} \notin \Theta_0\}$. Given the stated assumption that the posterior distribution of $\boldsymbol{\theta}$ has a strictly positive PDF for any \mathbf{X} , and the assumption that Δ has positive Lebesgue measure, then $\mathbb{P}(\boldsymbol{\theta} \in \Delta \mid \mathbf{X}) > 0$ for any \mathbf{X} . Using (5), let \mathbf{X}^* be any value such that $\phi(\mathbf{X}^*) = r$. Then,

$$\mathbb{P}(\boldsymbol{\theta} \in \Theta_0 \mid \mathbf{X}^*) = \overbrace{\mathbb{P}(\boldsymbol{\theta} \in \mathcal{R}_1 \mid \mathbf{X}^*)}^{=\alpha \text{ by (i)}} - \overbrace{\mathbb{P}(\boldsymbol{\theta} \in \Delta \mid \mathbf{X}^*)}^{>0} < \alpha.$$

By the assumption that $\mathbb{P}(\boldsymbol{\theta} \in \Theta_0 \mid \mathbf{X})$ is continuous in \mathbf{X} , there is some ϵ -ball \mathcal{B} around \mathbf{X}^* for which $\mathbb{P}(\boldsymbol{\theta} \in \Theta_0 \mid \mathbf{X}) < \alpha$, too. The ball \mathcal{B} has positive Lebesgue measure, as does the part of it lying outside \mathcal{R}_1 , $\mathcal{B} \cap \mathcal{R}_1^c$, since \mathbf{X}^* is on the boundary of \mathcal{R}_1 . Since the sampling distribution of \mathbf{X} given any $\boldsymbol{\theta}$ has a strictly positive PDF, $\mathbb{P}(\mathbf{X} \in (\mathcal{B} \cap \mathcal{R}_1^c) \mid \boldsymbol{\theta}) > 0$ for any

θ . Size is thus

$$\begin{aligned} \sup_{\theta \in \Theta_0} \mathbb{P}(\mathbf{X} \in \mathcal{R}_2 \mid \theta) &\geq \sup_{\theta \in \Theta_0} \left[\mathbb{P}(\mathbf{X} \in \mathcal{R}_1 \mid \theta) + \mathbb{P}(\mathbf{X} \in (\mathcal{B} \cap \mathcal{R}_1^c) \mid \theta) \right] \\ &\geq \underbrace{\mathbb{P}(\mathbf{X} \in \mathcal{R}_1 \mid \theta^*)}_{=\alpha \text{ by (6)}} + \underbrace{\mathbb{P}(\mathbf{X} \in (\mathcal{B} \cap \mathcal{R}_1^c) \mid \theta^*)}_{>0} > \alpha. \end{aligned}$$

For (iv), two examples suffice. These can be essentially the same as the examples for Theorem 1(iv). First, consider $H_0 : \phi(\theta) \neq 0$. As in the scalar case, given any \mathbf{X} , $\mathbb{P}(\phi(\theta) \neq 0 \mid \mathbf{X}) = 1$ since F is continuous, so the Bayesian test never rejects and its size is zero. Second, consider $H_0 : \phi(\theta) \in \mathbb{Z}$ (the integers). This H_0 has zero posterior probability given any \mathbf{X} , so the Bayesian test always rejects and has size equal to one. □

B Decision-theoretic context

The Bayesian test examined in this paper is a generalized Bayes decision rule that minimizes posterior expected loss (PEL) for the loss function taking value $1 - \alpha$ for type I error, α for type II error, and zero otherwise. Let $\mathbb{P}(\cdot \mid \mathbf{X})$ denote the posterior probability given observed data \mathbf{X} . The PEL for the decision to reject H_0 is $(1 - \alpha) \mathbb{P}(H_0 \mid \mathbf{X})$, i.e., the type I error loss times the posterior probability that rejecting H_0 is a type I error. Similarly, the PEL of accepting H_0 is $\alpha[1 - \mathbb{P}(H_0 \mid \mathbf{X})]$, the type II error loss times the probability that accepting H_0 is a type II error. PEL is thus minimized by rejecting H_0 if $\mathbb{P}(H_0 \mid \mathbf{X}) \leq \alpha$ and accepting H_0 otherwise. This is implicitly the same Bayesian test considered by Casella and Berger (1987a), although they compare $\mathbb{P}(H_0 \mid \mathbf{X})$ to a frequentist p -value rather than discuss “testing” explicitly.

Our results compare α to the Bayesian test’s RP at certain parameter values as well as the Bayesian test’s size (i.e., supremum of type I error rate over parameter values satisfying H_0). Primarily, we compare with α (instead of some other value) since it is the most relevant comparison in practice: can the α in Method 1 also be interpreted as the size of the test?

Secondarily, we compare the Bayesian test’s size to α for decision-theoretic reasons. Specifically, if an unbiased frequentist test with size α exists, then it is the minimax risk decision rule given the same loss function used to derive the Bayesian decision rule. Even without unbiasedness, this is approximately true given the values of α most common in practice. Assume $\theta \in \Theta$, $H_0 : \theta \in \Theta_0 \subset \Theta$, $H_1 : \theta \notin \Theta_0$. The minimax risk decision rule

minimizes

$$\begin{aligned} & \max \left\{ (1 - \alpha) \sup_{\boldsymbol{\theta} \in \Theta_0} P_{\boldsymbol{\theta}}(\text{reject}), \alpha \sup_{\boldsymbol{\theta} \in \Theta_1} P_{\boldsymbol{\theta}}(\text{accept}) \right\} \\ & = \max \left\{ (1 - \alpha) \sup_{\boldsymbol{\theta} \in \Theta_0} P_{\boldsymbol{\theta}}(\text{reject}), \alpha \left[1 - \inf_{\boldsymbol{\theta} \in \Theta_1} P_{\boldsymbol{\theta}}(\text{reject}) \right] \right\}, \end{aligned} \quad (7)$$

where $P_{\boldsymbol{\theta}}(\cdot)$ is the probability under $\boldsymbol{\theta}$. Having (exact) size α means $\sup_{\boldsymbol{\theta} \in \Theta_0} P_{\boldsymbol{\theta}}(\text{reject}) = \alpha$. Unbiasedness means $\sup_{\boldsymbol{\theta} \in \Theta_0} P_{\boldsymbol{\theta}}(\text{reject}) \leq \inf_{\boldsymbol{\theta} \in \Theta_1} P_{\boldsymbol{\theta}}(\text{reject})$; by continuity (in $\boldsymbol{\theta}$) of the power function, $\inf_{\boldsymbol{\theta} \in \Theta_1} P_{\boldsymbol{\theta}}(\text{reject}) = \alpha$, so the maximum risk in (7) is $\alpha(1 - \alpha)$. Increasing the nominal level above α increases the first term inside the max in (7) above $\alpha(1 - \alpha)$, and decreasing the level below α increases the second term above $\alpha(1 - \alpha)$; thus, the unbiased test with size α is the minimax risk decision rule. See also Lehmann and Romano (2005, Problem 1.10) on unbiased tests as minimax risk decision rules.

Absent an unbiased test, the minimax-risk-optimal size of a given test is above α , but the magnitude of the difference is very small for conventional α . Consider varying the size of a test, $\gamma_0 = \sup_{\boldsymbol{\theta} \in \Theta_0} P_{\boldsymbol{\theta}}(\text{reject})$. As a function of γ_0 , let $\gamma_1(\gamma_0) = \inf_{\boldsymbol{\theta} \in \Theta_1} P_{\boldsymbol{\theta}}(\text{reject})$, so the risk in (7) is $\max\{(1 - \alpha)\gamma_0, \alpha(1 - \gamma_1(\gamma_0))\}$. In the extreme, $\gamma_1(\gamma_0) = 0$ for all γ_0 , and the maximum risk is α for any test with $\gamma_0 \leq \alpha/(1 - \alpha)$, while maximum risk is larger than α if $\gamma_0 > \alpha/(1 - \alpha)$. If instead $\gamma_1(\gamma_0)$ is strictly increasing in γ_0 but $\gamma_1(\alpha) < \alpha$, then minimax risk is achieved at some $\gamma_0 \in (\alpha, \alpha/(1 - \alpha))$. For example, rounding to two significant digits, if $\alpha = 0.05$, then $\gamma_0 \in (0.050, 0.053)$, or if $\alpha = 0.1$, then $\gamma_0 \in (0.10, 0.11)$. Such small divergence of γ_0 from α is almost imperceptible in practice.

Ideally, a single decision rule minimizes both the maximum risk in (7) and the PEL. However, if the Bayesian test's size is significantly above or below α , this is not possible. In such cases, it may help to use both Bayesian and frequentist inference and to carefully consider the differences in optimality criteria.

C Discussion of assumptions with function spaces

For estimators of functions, it is common to have a (frequentist) Gaussian process limit with sample paths continuous with respect to the covariance semimetric; e.g., see van der Vaart and Wellner (1996). A natural question is whether the (asymptotic, limit experiment) sampling and posterior distributions are ever equivalent in the sense of

$$X(\cdot) - \theta(\cdot) \mid \theta(\cdot) \sim \mathbb{G}, \quad \theta(\cdot) - X(\cdot) \mid X(\cdot) \sim \mathbb{G},$$

where \mathbb{G} is a mean-zero Gaussian process with known covariance function.

Unfortunately, as discussed by Freedman (1999) and others, such a Bernstein–von Mises result is rare with infinite-dimensional spaces. As explained by Hirano and Porter (2009, p. 1696), in finite dimensions the prior often behaves locally like Lebesgue measure (if its PDF is continuous and positive at the true parameter value), but in infinite-dimensional Banach spaces there is no analog of Lebesgue measure, let alone one that most priors would satisfy.

However, the important special case of inference on a continuous CDF can satisfy the necessary assumptions. Assume iid sampling. On the frequentist side,

$$\sqrt{n}(\hat{F}(\cdot) - F(\cdot)) \rightsquigarrow B(F(\cdot)), \quad (8)$$

an F -Brownian bridge (where $B(\cdot)$ is a standard Brownian bridge), with \rightsquigarrow denoting weak convergence in $\ell^\infty(\bar{\mathbb{R}})$; e.g., see van der Vaart and Wellner (1996, Ex. 2.1.3). For weak convergence under sequences $F_n(\cdot) \rightarrow F(\cdot)$, see Sections 2.8.3 and 3.11 and especially Theorem 3.10.12 in van der Vaart and Wellner (1996). For a nonparametric Bayesian method using the Dirichlet process of Ferguson (1973), Lo (1983, Thm. 2.1) shows that a centered (at $\hat{F}(\cdot)$) and \sqrt{n} -scaled version of the posterior converges to the same F -Brownian bridge if the prior dominates $F(\cdot)$. Even with an improper prior, i.e., using the Bayesian bootstrap of Rubin (1981), Lo (1987, Thm. 2.1) shows that the centered and scaled posterior converges as in (8).

For our purpose of approximating the finite-sample difference between Bayesian and frequentist testing, considering a fixed DGP and drifting centering parameter can be just as helpful as considering a fixed centering parameter and drifting DGP. In the finite-dimensional case, the limit experiment can come from

$$\begin{aligned} \sqrt{n}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}) &\xrightarrow{d} N(\mathbf{0}, \boldsymbol{\Sigma}), \\ \mathbf{X} = \sqrt{n}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_{0,n}) &= \sqrt{n}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}) + \overbrace{\sqrt{n}(\boldsymbol{\mu} - \boldsymbol{\mu}_{0,n})}^{\rightarrow \boldsymbol{\theta}} \xrightarrow{d} N(\boldsymbol{\theta}, \boldsymbol{\Sigma}), \end{aligned} \quad (9)$$

so \mathbf{X} is a test statistic based on the observed (computed) $\hat{\boldsymbol{\mu}}$ and centered at $\boldsymbol{\mu}_{0,n}$. This does not have a literal meaning like “we must change $\boldsymbol{\mu}_0$ if our sample size increases,” just as a drifting DGP does not mean literally that “the population distribution changes as we collect more data”; rather, it is simply a way to capture the idea of $\boldsymbol{\mu}_0$ being “close to” the true $\boldsymbol{\mu}$ in the asymptotics. For the posterior, letting $\boldsymbol{\theta} = \sqrt{n}(\boldsymbol{\mu} - \boldsymbol{\mu}_{0,n})$ and again $\mathbf{X} = \sqrt{n}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_{0,n})$,

$$\begin{aligned} \sqrt{n}(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}) &\xrightarrow{d} N(\mathbf{0}, \boldsymbol{\Sigma}), \\ \boldsymbol{\theta} = \sqrt{n}(\boldsymbol{\mu} - \boldsymbol{\mu}_{0,n}) &= \mathbf{X} + \overbrace{\sqrt{n}(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}})}^{\xrightarrow{d} N(\mathbf{0}, \boldsymbol{\Sigma})} \xrightarrow{d} N(\mathbf{X}, \boldsymbol{\Sigma}). \end{aligned} \quad (10)$$

For the infinite-dimensional case, for general rate of convergence n^r , a similar setup to (9) is

$$n^r(\hat{\mu}(\cdot) - \mu(\cdot)) \rightsquigarrow \mathbb{G}(\cdot), \quad (11)$$

$$X(\cdot) = n^r(\hat{\mu}(\cdot) - \mu_{0,n}(\cdot)) = n^r(\hat{\mu}(\cdot) - \mu(\cdot)) + \overbrace{n^r(\mu(\cdot) - \mu_{0,n}(\cdot))}^{\rightarrow \theta(\cdot)} \rightsquigarrow \mathbb{G}(\cdot) + \theta(\cdot),$$

where $\mathbb{G}(\cdot)$ is a mean-zero Gaussian process and $\theta(\cdot)$ is the (non-random) local mean parameter. For the posterior, with $\theta(\cdot) = n^r(\mu(\cdot) - \mu_{0,n}(\cdot))$ and $X(\cdot) = n^r(\hat{\mu}(\cdot) - \mu_{0,n}(\cdot))$ like in (10),

$$\begin{aligned} n^r(\mu(\cdot) - \hat{\mu}(\cdot)) &\rightsquigarrow \mathbb{G}(\cdot), \\ \theta(\cdot) = n^r(\mu(\cdot) - \mu_{0,n}(\cdot)) &= X(\cdot) + n^r(\mu(\cdot) - \hat{\mu}(\cdot)) \rightsquigarrow \mathbb{G}(\cdot) + X(\cdot). \end{aligned} \quad (12)$$

D Derivation of translog constraints

The Hessian is a nonlinear function of the translog parameters, and it depends on (y, \mathbf{w}) . Letting¹¹

$$r_k \equiv \frac{\partial \ln(C)}{\partial \ln(w_k)} = a_{yk} \ln(y) + b_k + \sum_{j=1}^3 b_{jk} \ln(w_j), \quad (13)$$

a general element of \mathbf{H} is

$$\begin{aligned} H_{mk} &= \frac{\partial^2 C}{\partial w_m \partial w_k} = \frac{\partial}{\partial w_m} \frac{\partial C}{\partial w_k} = \frac{\partial}{\partial w_m} (r_k C / w_k) = \frac{\partial r_k}{\partial w_m} (C / w_k) + \frac{\partial C}{\partial w_m} (r_k / w_k) + \frac{\partial w_k^{-1}}{\partial w_m} (r_k C) \\ &= (b_{mk} / w_m) (C / w_k) + r_m (C / w_m) (r_k / w_k) - \mathbf{1}\{k = m\} w_k^{-2} (r_k C) \\ &= C \frac{b_{mk} + r_m r_k - \mathbf{1}\{k = m\} r_k}{w_m w_k}. \end{aligned}$$

Since each element is proportional to $C > 0$, the value of C does not affect whether or not \mathbf{H} is NSD: \mathbf{H} is NSD iff \mathbf{H}/C is NSD. This may be helpful if the translog parameters are estimated from cost share equations and C is not directly observed.

The local NSD condition in (2) corresponds to a set of parameter values much smaller than a half-space. A necessary (but not sufficient) condition for NSD is that all the principal minors of order $p = 1$ are non-positive, i.e., that $H_{11} \leq 0$, $H_{22} \leq 0$, and $H_{33} \leq 0$. In terms of the parameters, using (13), $H_{11} \leq 0$ iff $b_{11} + r_1^2 - r_1 \leq 0$, i.e., $b_{11} \leq r_1(1 - r_1)$. With $(y, \mathbf{w}) = (1, 1, 1, 1)$, $r_k = b_k$, so $H_{11} \leq 0$ iff $b_{11} \leq b_1(1 - b_1)$. After imposing symmetry ($b_{mk} = b_{km}$) and homogeneity of degree one in input prices ($b_{m1} + b_{m2} + b_{m3} = 0$, $m = 1, 2, 3$), all b_{mk} can be written in terms of b_{11} , b_{12} , and b_{22} : $b_{21} = b_{12}$, $b_{13} = -b_{11} - b_{12}$, etc. Also

¹¹Some notation is from O'Donnell and Coelli (2005).

from homogeneity, $b_1 + b_2 + b_3 = 1$, and from monotonicity, $b_k = r_k \geq 0$, so $0 \leq b_1 \leq 1$. Thus, $b_1(1 - b_1) \in [0, 0.25]$, so $b_{11} \leq b_1(1 - b_1)$ is larger than the half-space defined by $b_{11} \leq 0$ but smaller than the half-space defined by $b_{11} \leq 0.25$. A similar argument for $H_{22} \leq 0$ at $(1, 1, 1, 1)$ yields $b_{22} \leq b_2(1 - b_2) \leq 0.25$. From the constraints on H_{11} and H_{22} alone, Θ_0 is a subset of the “quarter-space” defined by $b_{11} \leq 0.25$ and $b_{22} \leq 0.25$. Adding the constraints for the other principal minors of \mathbf{H} makes Θ_0 even smaller.

Since the local concavity H_0 in (2) corresponds to a subset of a *quarter*-space in the parameter space, Theorem 2(iii) suggests that we expect the Bayesian test’s size to exceed α . The results in Table 3 show this to be the case here.