

A FAST METHOD FOR IMPLEMENTING HYPOTHESIS TESTS WITH MULTIPLE SAMPLE SPLITS IN NONPARAMETRIC MODELS OF PRODUCTION

Léopold Simar, Paul W. Wilson

REPRINT | 2025 / 20

ISBA

Voie du Roman Pays 20 - L1.04.01

B-1348 Louvain-la-Neuve

Email : lidam-library@uclouvain.be

<https://uclouvain.be/en/research-institutes/lidam/isba/publication>

A Fast Method for Implementing Hypothesis Tests with Multiple Sample Splits in Nonparametric Models of Production

LÉOPOLD SIMAR

PAUL W. WILSON*

January 2025

Abstract

Kneip et al. (*Journal of Business and Economic Statistics*, 2016) and Daraio et al. (*The Econometrics Journal*, 2018) provide non-parametric tests of (i) convexity versus non-convexity of the production set, (ii) constant versus non-constant returns-to-scale of the frontier, and (iii) separability versus non-separability of the frontier with respect to environmental variables. Among other uses, these tests are essential for deciding which non-parametric efficiency estimator should be used to estimate technical efficiency. Each test requires randomly splitting the sample. Although theory establishes that the tests are valid for *any* random split, results can vary with different splits. This paper provides a computationally efficient method to aggregate test outcomes across multiple sample-splits using a simple decision rule, and we prove that our rule yields tests with asymptotic size no greater than the nominal size. We provide tests using multiple sample-splits to remove the ambiguity resulting from a single sample-split and give extensive Monte Carlo evidence on the size and power of our tests. The computational time required by the new tests is about 0.001 times that required by the bootstrap method proposed by Simar and Wilson (*Journal of Productivity Analysis*, 2020). The reduction in computational burden makes the tests useful in a practical sense for empirical researchers using inexpensive desktop or laptop computers.

Keywords: Hypothesis testing, inference, multiple splits, convexity, returns to scale, separability, DEA, FDH.

JEL classification: C12, C44, C63

*Simar: Institut de Statistique, Biostatistique, et Sciences Actuarielles (ISBA), LIDAM, Université Catholique de Louvain-la-Neuve, Belgium; email: leopold.simar@uclouvain.be.

Wilson: Department of Economics and School of Computing, Clemson University, Clemson, South Carolina, USA; email: pww@clemson.edu.

We have benefited from a discussion with Eugen Pircalabelu. Technical support from the Cyber Infrastructure Technology Integration group at Clemson University is gratefully acknowledged. Any remaining errors are solely our responsibility.

1 Introduction

Non-parametric envelopment estimators are widely used to estimate efficiency of production. Broadly, there are three classes of such estimators: (i) those that impose neither convexity nor constant returns-to-scale (CRS) on the production frontier; (ii) those that impose convexity but not CRS; and (iii) those that impose both convexity and CRS. In addition, conditional versions of the estimators must be used when environmental variables that affect the frontier are present (see Simar and Wilson, 2007, 2011 for details and discussion). Kneip et al. (2016) and Daraio et al. (2018) provide tests of model structure to help the empirical researcher choose an appropriate estimator among the available choices. As discussed below, these tests require randomly splitting the data into two independent subsamples to obtain statistics with non-degenerate limiting distributions under the null, and while the tests are theoretically valid for any random split of the original sample, different results can be obtained with different splits. Simar and Wilson (2020) develop a bootstrap method to remove this ambiguity, but the computational burden of their method makes the approach impractical in many settings. This paper provides a new method to remove the uncertainty introduced by multiple splitting that substantially reduces the computational burden required to implement the tests of Kneip et al. (2016) and Daraio et al. (2018) using multiple sample splits. The new method should be practical for empirical researchers even with samples involving several thousand observations without using large parallel computing systems or other high-throughput systems such as HTCondor.¹

Standard microeconomic theory of the firm defines a production set as the set of feasible combination of input and output quantities for a firm in a given industry. Firms that operate on the frontier of this set are technically efficient, while firms that operate in the interior of the production set are technically *inefficient*. The frontier of the production set serves

¹HTCondor is a high-throughput computing environment that exploits idle computer resources in a network. See <https://htcondor.org/index.html> for details.

as a useful benchmark for gauging the performance of firms in an industry. Unfortunately, however, the frontier is not observed and must be estimated. Similarly, distance between a particular input-output combination and the frontier of the production set is unobserved and must also be estimated.

Several non-parametric estimators of the production set, as well as corresponding estimators of distance to the frontier of the production set, have been proposed in the literature on efficiency and productivity. Deprins et al. (1984) estimate the production set by the free-disposal hull (FDH) of a set of sample observations on input-output combinations. As discussed below in Section 2, this estimator is quite flexible since it does not impose convexity of the production set nor does it require that the frontier display constant returns to scale (CRS). Alternatively, Farrell (1957) estimates the production set by the conical hull of observed input-output pairs, thereby imposing assumptions of (i) convexity of the production set and (ii) constant returns to scale (CRS) on the frontier. Charnes et al. (1978) formalize the corresponding estimator of distance to the frontier in terms of a linear program. Banker et al. (1984) relax the requirement of CRS, but not convexity of the production set, by estimating the production set by the convex hull of the FDH of observed input-output combinations and write the corresponding estimator of distance to the frontier as a linear program. The resulting frontier estimate allows for variable returns to scale (VRS), i.e., for increasing, constant or decreasing returns to scale in various regions. The two estimators based on linear programming are known in the literature as “data envelopment analysis” (DEA) estimators, a term coined by Rhodes (1978) and adopted by both Charnes et al. (1978) and Banker et al. (1984). In the ensuing discussion below, we refer to these as CRS-DEA and VRS-DEA estimators.

The applied, empirical researcher who wishes to avoid restrictive, parametric assumption is faced with the problem of first choosing which estimator to use. A safe choice might be the FDH estimator, since it imposes neither convexity nor CRS. But, as discussed below,

the convergence rate of the FDH estimator is slower than either the VRS-DEA or CRS-DEA estimators, and thus may result in larger estimation error than either of the DEA estimators. On the other hand, one might choose one of the DEA estimators, but this requires convexity of the production set, and also CRS if the CRS-DEA estimator is used. One should test convexity versus non-convexity of the production set before using the VRS-DEA estimators instead of FDH estimators, and one should in addition test CRS versus non-CRS of the frontier before using CRS-DEA estimators instead of VRS-DEA or FDH estimators.²

This paper provides a new method for aggregating information provided by the original tests of Kneip et al. (2016) over multiple random sample-splits, allowing researchers to “integrate out” the uncertainty introduced by random splitting. We provide a simple decision rule and prove that it yields tests with asymptotic less than or equal to the nominal size.³ We provide simulation evidence that suggests our rule yields tests that are conservative (i.e., that have rejection rates less than the nominal test size) when the null is true and that have substantial power for reasonable departures from the null, whereas the bootstrap method of Simar and Wilson (2020) leads to tests that typically over-reject when the null is true. In addition, our new method provides a substantial reduction in computational burden over the bootstrap method of Simar and Wilson (2020); computational time required by the new method is typically about 0.001 times the computational time required by the bootstrap method of Simar and Wilson (2020), making our new method useful to empirical researchers with sample sizes commonly encountered in the literature on efficiency estimation.

Our rule is related to ideas from the statistical literature on multiple hypothesis testing, in particular those of Simes (1986), Sarkar and Chang (1997), Benjamini and Hochberg (1995) and Benjamini and Yekutieli (2001) for controlling false discovery rates. A key difference in

²As discussed in Section 3, Kneip et al. (2022) introduce a test of CRS versus non-CRS for cases where the production set may not be convex.

³Note that one typically chooses a priori the level of significance or nominal test size as the acceptable rate of type-I error. In the tests considered in this paper, the null hypotheses to be tested are composite hypotheses, and so the test size is the supremum of the rejection rates over all possibilities under the null.

our setting, however, is that the null hypothesis (e.g., the frontier is everywhere CRS or the production set is convex) is the same for each test on each random sample-split, whereas in the multiple testing literature, typically a set of *different* null hypotheses are tested.

The remainder of the paper proceeds along the following path. The next section presents a statistical model and describes model features that researchers are often interested in. Section 3 discusses issues surrounding estimation of testing of hypotheses. Estimators of efficiency are discussed in Section 3.1, then the tests of CRS versus non-CRS and convexity versus non-convexity developed by Kneip et al. (2016) are briefly reviewed in Sections 3.2 and 3.3, respectively. Section 4 describes our new method for removing uncertainty arising from multiple sample-splits with the tests of Kneip et al. (2016) and Daraio et al. (2018). Section 5 provides evidence from Monte Carlo experiments on the performance of our new method in terms of realized size and power when testing CRS versus non-CRS and when testing convexity versus non-convexity of the production set, with comparisons to the earlier method of Simar and Wilson (2020).⁴ Summary and conclusions are given in Section 6, and some additional results are given in separate Appendices A–F.

2 The Statistical Model

Let $x, X \in \mathbb{R}_+^p$ denote fixed and stochastic p -vectors of input quantities, and let $y, Y \in \mathbb{R}_+^q$ denote fixed and stochastic q -vectors of output quantities (respectively).⁵ Standard microeconomic theory specifies a production set

$$\Psi := \{(x, y) \mid x \text{ can produce } y\} \quad (2.1)$$

that consists of all feasible combinations of input and output quantities. Standard assumptions on Ψ include (i) Ψ is closed; (ii) production of any non-zero output quantity requires use

⁴Monte Carlo evidence for the Daraio et al. (2018) separability test is omitted to avoid additional computational expense. However, one can expect similar performance and saving of computational burden when our new method is used with the tests of Daraio et al. (2018).

⁵Throughout, we use lower-case letters to denote fixed quantities, and upper-case letters to denote random variables.

of non-zero level of at least one input; and (iii) both inputs and outputs are freely disposable so that for $\tilde{x} \geq x$ and $\tilde{y} \leq y$, if $(x, y) \in \Psi$ then $(\tilde{x}, y) \in \Psi$ and $(x, \tilde{y}) \in \Psi$, where inequalities involving vectors are defined on an element-by-element basis. The third assumption amounts to an assumption of weak monotonicity on the frontier

$$\Psi^\partial = \{(x, y) \mid (x, y) \in \Psi, (\gamma x, \gamma^{-1} y) \notin \Psi \text{ for } \gamma < 1\}. \quad (2.2)$$

Technical efficiency of a point $(x, y) \in \Psi$ is measured in the input or output direction by the distance functions

$$\theta(x, y \mid \Psi) := \inf \{\theta \mid (\theta x, y) \in \Psi\} \quad (2.3)$$

and

$$\lambda(x, y \mid \Psi) := \sup \{\lambda \mid (x, \lambda y) \in \Psi\} \quad (2.4)$$

proposed by Farrell (1957). Alternatively, one can measure technical efficiency using hyperbolic, directional or generalized hyperbolic distance functions proposed by Färe et al. (1985), Chambers et al. (1996) or Wilson (2024). To simplify and to conserve space, we focus our work on the input direction, but one can work in the output or other directions. The particular direction chosen is unlikely to change the main results of this paper.

The input-oriented measure in (2.3) gives the proportion by which input levels can be feasibly, simultaneously reduced for any $(x, y) \in \Psi$ without sacrificing any output. By construction, $\theta(x, y \mid \Psi) \leq 1$ for any $(x, y) \in \Psi$. If $\theta(x, y \mid \Psi) = 1$, then $(x, y) \in \Psi^\partial$, and $\theta(x, y \mid \Psi) < 1$ for any $(x, y) \in \text{int}(\Psi)$ where $\text{int}(\Psi)$ denotes the interior of Ψ .

The production set Ψ , frontier Ψ^∂ and the measures in (2.3)–(2.4) are unobserved in any interesting case. Hence these features must be estimated and inference must be made in order to learn about them. This, in turn, requires a statistical model. We assume a well-defined probability space on which the density $f(x, y)$ of input and output quantities (X, Y) is defined. We also assume that $f(x, y)$ has support on $\mathcal{D} \subset \Psi$, and that $f(x, y)$ is continuously differentiable on $\mathcal{D}$. In addition, we rely on the additional assumptions

of Kneip et al. (2015) as appropriate. These assumptions concern the smoothness of the distance function $\theta(x, y \mid \Psi)$ and the structure of the set $\mathcal{D}$. See Kneip et al. (2015) for details and discussion.

3 Estimation and Hypothesis Testing

3.1 Estimation of Efficiency

Suppose that an independent, identically distributed (iid) sample $\mathcal{X}_n = \{(X_i, Y_i)\}_{i=1}^n$ drawn from the density $f(x, y)$ is available. Several non-parametric estimators of Ψ are available. The least restrictive among the non-parametric estimators of Ψ is the free-disposal hull (FDH)

$$\widehat{\Psi}_{\text{FDH}}(\mathcal{X}_n) := \bigcup_{(X_i, Y_i) \in \mathcal{X}_n} \{(x, y) \in \mathbb{R}_+^{p+q} \mid x \geq X_i, y \leq Y_i\}. \quad (3.1)$$

of the sample observations used by Deprins et al. (1984). Note that $\widehat{\Psi}_{\text{FDH}}(\mathcal{X}_n)$ is not convex when n is finite.

In cases where Ψ is convex and Ψ^∂ displays constant returns-to-scale everywhere, Ψ can be estimated by the conical hull of $\widehat{\Psi}_{\text{FDH}}(\mathcal{X}_n)$ given by

$$\widehat{\Psi}_{\text{CRS}}(\mathcal{X}_n) := \{(x, y) \in \mathbb{R}^{p+q} \mid y \leq \mathbf{Y}\boldsymbol{\omega}, x \geq \mathbf{X}\boldsymbol{\omega}, \boldsymbol{\omega} \in \mathbb{R}_+^n\}. \quad (3.2)$$

This CRS-DEA estimator was proposed by Farrell (1957) and written formally as a linear-program by Charnes et al. (1978). Alternatively, when Ψ is convex but Ψ^∂ does not exhibit CRS everywhere, Ψ can be estimated by the convex hull

$$\widehat{\Psi}_{\text{VRS}}(\mathcal{X}_n) := \{(x, y) \in \mathbb{R}^{p+q} \mid y \leq \mathbf{Y}\boldsymbol{\omega}, x \geq \mathbf{X}\boldsymbol{\omega}, \mathbf{i}'_n \boldsymbol{\omega} = 1, \boldsymbol{\omega} \in \mathbb{R}_+^n\}, \quad (3.3)$$

of $\widehat{\Psi}_{\text{FDH}}(\mathcal{X}_n)$, where $\mathbf{X} = (X_1, \dots, X_n)$ and $\mathbf{Y} = (Y_1, \dots, Y_n)$ are $(p \times n)$ and $(q \times n)$ matrices of input and output vectors, respectively; $\mathbf{i}_n$ is an $(n \times 1)$ vector of ones, and $\boldsymbol{\omega}$ is a $(n \times 1)$ vector of weights. This is the variable returns-to-scale (VRS) version of the data

envelopment analysis (DEA) estimator anticipated by Farrell (1957) and written formally as a linear program by Banker et al. (1984).

Substituting one of the non-parametric estimators for Ψ in (2.3) leads to a non-parametric estimator of $\theta(x, y \mid \Psi)$. Let $\mathcal{I}(x, y)$ be the set of indices of observations in $\mathcal{X}_n$ such that $X_i \leq x$ and $Y_i \geq y$. Then substituting $\widehat{\Psi}_{\text{FDH}}(\mathcal{X}_n)$ for Ψ in (2.3) yields the FDH estimator $\widehat{\theta}_{\text{FDH}}(x, y \mid \mathcal{S}_n)$ that can be computed as

$$\widehat{\theta}_{\text{FDH}}(x, y \mid \mathcal{X}_n) = \min_{i \in \mathcal{I}(x, y)} \max_{j=1, \dots, p} \left(\frac{X_{ij}}{x_j} \right) \quad (3.4)$$

where X_{ij} and x_j denote the j th elements of X_i and x (respectively).⁶

Substituting $\widehat{\Psi}_{\text{VRS}}(\mathcal{X}_n)$ in (3.3) for Ψ in (2.3) yields the VRS-DEA estimator

$$\widehat{\theta}_{\text{VRS}}(x, y \mid \mathcal{X}_n) = \min_{\theta, \omega} \{ \theta \mid y \leq \mathbf{Y}\omega, \theta x \geq \mathbf{X}\omega, \mathbf{i}'_n \omega = 1, \omega \in \mathbb{R}_+^n \} \quad (3.5)$$

of $\theta(x, y \mid \Psi)$. Similarly, substituting $\widehat{\Psi}_{\text{CRS}}(\mathcal{X}_n)$ in (3.2) for Ψ in (2.3) yields the CRS-DEA estimator

$$\widehat{\theta}_{\text{CRS}}(x, y \mid \mathcal{X}_n) = \min_{\theta, \omega} \{ \theta \mid y \leq \mathbf{Y}\omega, \theta x \geq \mathbf{X}\omega, \omega \in \mathbb{R}_+^n \} \quad (3.6)$$

of $\theta(x, y \mid \Psi)$. Both $\widehat{\theta}_{\text{VRS}}(x, y \mid \mathcal{X}_n)$ and $\widehat{\theta}_{\text{CRS}}(x, y \mid \mathcal{X}_n)$ can be computed using standard linear-programming methods.

Among the three non-parametric estimators of Ψ described above, $\widehat{\Psi}_{\text{FDH}}(\mathcal{X}_n)$ is the least restrictive, while $\widehat{\Psi}_{\text{CRS}}(\mathcal{X}_n)$ is the most restrictive. By construction, $\widehat{\Psi}_{\text{FDH}}(\mathcal{X}_n) \subseteq \widehat{\Psi}_{\text{VRS}}(\mathcal{X}_n) \subseteq \widehat{\Psi}_{\text{CRS}}(\mathcal{X}_n)$. Consequently, $\widehat{\theta}_{\text{CRS}}(X_i, Y_i \mid \mathcal{X}_n) \leq \widehat{\theta}_{\text{VRS}}(X_i, Y_i \mid \mathcal{X}_n) \leq \widehat{\theta}_{\text{FDH}}(X_i, Y_i \mid \mathcal{X}_n) \leq 1$ for all $(X_i, Y_i) \in \mathcal{X}_n$. The properties of each of the three estimators of $\theta(x, y \mid \Psi)$ are well-known due to results obtained by Park et al. (2000), Kneip et al. (2015) and Daouia et al. (2017) for the FDH estimator; by Kneip et al. (1998) and Kneip et al. (2008, 2015, 2016) for the VRS-DEA estimator; and Park et al. (2010) and Kneip et al. (2015) for the CRS-DEA estimator. Under appropriate assumptions, each of the three estimators is consistent, has a

⁶If $x_j = 0$, then X_{ij}/x_j should be set equal to infinity if $X_{ij} > 0$, or negative infinity if $X_{ij} = 0$.

non-degenerate limiting distribution, converges at rate n^κ and has bias of order $O(n^{-\kappa})$. For the FDH estimator, under either convexity or non-convexity and either CRS or non-CRS, $\kappa = 1/(p + q)$. For the VRS-DEA estimator, provided Ψ is convex, $\kappa = 2/(p + q + 1)$ under VRS or $2/(p + q)$ under CRS. For the CRS-DEA estimator, $\kappa = 2/(p + q)$ provided Ψ is convex and the frontier is CRS.

A problem in any empirical project is to decide which estimator to use. The FDH estimator provides the most flexibility, since it is consistent regardless of whether Ψ is convex or whether Ψ^∂ is CRS. But, its convergence rate is slower than the other two estimators. Under CRS, both the VRS-DEA and CRS-DEA estimators converge at rate $n^{2/(p+q)}$, but the VRS-DEA estimator should be expected to have higher mean-square error than the CRS-DEA estimator in such cases since the VRS-DEA estimator does not impose CRS. The obvious approach is to test for convexity versus non-convexity of Ψ , and for CRS versus non-CRS for the frontier Ψ^∂ . If convexity is rejected, one should use the FDH estimator. If convexity is not rejected, one can consider using either of the DEA estimators, but if CRS is rejected only the VRS-DEA estimator can be considered. Of course, one should remember that failure to reject the null does not, by itself, imply that the null is true. As such, among the three classes of estimators, FDH estimators are the safest choice since they require neither convexity nor CRS.

3.2 The Kneip et al. (2016) Test of CRS versus non-CRS

The Kneip et al. (2016) test of CRS versus non-CRS is based on comparison of means of the VRS-DEA and CRS-DEA estimates. Under the null hypothesis of CRS, the difference between the two means should be “small,” while under the alternative hypothesis the difference should be “large.” Kneip et al. (2015, 2016) prove that under the null hypothesis of CRS, both the VRS-DEA and CRS-DEA estimators converge at rate n^κ where $\kappa = 2/(p+q)$.⁷

⁷Consequently, the rate of convergence depends on the number of dimensions, i.e., the number of inputs and outputs. DEA as well as FDH estimators suffer from the well-known “curse of dimensionality,” meaning

In addition, Kneip et al. (2015) establish that standard CLTs such as the Lindeberg-Feller CLT do not hold for means of non-parametric efficiency estimates if $\kappa \leq 1/2$ due to the bias of non-parametric efficiency estimators. Under the null hypothesis of CRS, $\kappa > 1/2$ if and only if $(p + q) < 4$, and hence standard CLT results apply only in limited cases. The test developed by Kneip et al. (2016) is based on the central limit (CLT) results of Kneip et al. (2015). In addition, Kneip et al. discuss the necessity for the means to be independent in order to obtain a limiting $N(0, 1)$ distribution for the test statistic.

In order to ensure independence between two sample means of efficiency estimates, set $n_1 = \lfloor n/2 \rfloor$ and $n_2 = n - n_1$ where $\lfloor a \rfloor$ denotes the largest integer less than or equal to $a \in \mathbb{R}$. Then randomly split the sample $\mathcal{X}_n$ into two samples $\mathcal{X}_{1,n_1}$, $\mathcal{X}_{2,n_2}$ such that $\mathcal{X}_{1,n_1} \cup \mathcal{X}_{2,n_2} = \mathcal{X}_n$ and $\mathcal{X}_{1,n_1} \cap \mathcal{X}_{2,n_2} = \emptyset$.⁸ Let

$$\hat{\mu}_{\text{VRS},n_1} = n_1^{-1} \sum_{(X_i, Y_i) \in \mathcal{X}_{1,n_1}} \hat{\theta}_{\text{VRS}}(X_i, Y_i \mid \mathcal{X}_{1,n_1}) \quad (3.7)$$

and

$$\hat{\mu}_{\text{CRS},n_2} = n_2^{-1} \sum_{(X_i, Y_i) \in \mathcal{X}_{2,n_2}} \hat{\theta}_{\text{CRS}}(X_i, Y_i \mid \mathcal{X}_{2,n_2}) \quad (3.8)$$

and let

$$\hat{\sigma}_{\text{VRS},n_1}^2 = n_1^{-1} \sum_{(X_i, Y_i) \in \mathcal{X}_{1,n_1}} \left[\hat{\theta}_{\text{VRS}}(X_i, Y_i \mid \mathcal{X}_{1,n_1}) - \hat{\mu}_{\text{VRS},n_1} \right]^2 \quad (3.9)$$

and

$$\hat{\sigma}_{\text{CRS},n_2}^2 = n_2^{-1} \sum_{(X_i, Y_i) \in \mathcal{X}_{2,n_2}} \left[\hat{\theta}_{\text{CRS}}(X_i, Y_i \mid \mathcal{X}_{2,n_2}) - \hat{\mu}_{\text{CRS},n_2} \right]^2. \quad (3.10)$$

In addition, let $\hat{B}_{\text{VRS},\kappa,n_1}$ and $\hat{B}_{\text{CRS},\kappa,n_2}$ denote the generalized jackknife estimates of the bias of $\hat{\mu}_{\text{VRS},n_1}$ and $\hat{\mu}_{\text{CRS},n_2}$ given by Kneip et al. (2016, equations 37 and 38).⁹ Then under the

that DEA and FDH estimators' convergence rates become slower as dimensionality increases. The rate of convergence of an estimator is crucial when using the estimator to make inference.

⁸As a practical matter, the sample $\mathcal{X}_n$ can be randomly split into two subsamples by first randomly shuffling the observations using the algorithm of Durstenfeld (1964) and then giving the first n_1 observations to $\mathcal{X}_{1,n_1}$ and the remaining n_2 observations to $\mathcal{X}_{2,n_2}$.

⁹Note that Kneip et al. (2016, equation 38) contains two typos. The term n_1 in the subscripts should be instead n_2 .

null hypothesis of CRS and the assumptions of Kneip et al. (2015, Theorem 4.2) and Kneip et al. (2016, Theorem 1),

$$\hat{T}_{1,n} := \frac{(\hat{\mu}_{\text{VRS},n_1} - \hat{\mu}_{\text{CRS},n_2}) - (\hat{B}_{\text{VRS},\kappa,n_1} - \hat{B}_{\text{CRS},\kappa,n_2})}{\sqrt{\frac{\hat{\sigma}_{\text{VRS},n_1}^2}{n_1} + \frac{\hat{\sigma}_{\text{CRS},n_2}^2}{n_2}}} \xrightarrow{\mathcal{L}} N(0,1) \quad (3.11)$$

as $n \rightarrow \infty$ provided $(p+q) \leq 5$.

Alternatively, for $(p+q) \geq 5$, set $n_{\ell,\kappa} = \lceil n_\ell^{2\kappa} \rceil$ for $\ell \in \{1, 2\}$ and let $\mathcal{X}_{\ell,n_{\ell,\kappa}}^*$ be a random subset of $n_{\ell,\kappa}$ input-output pairs from $\mathcal{X}_{\ell,n_\ell}$. In addition, let

$$\hat{\mu}_{\text{VRS},n_{1,\kappa}} = n_{1,\kappa}^{-1} \sum_{(X_i, Y_i) \in \mathcal{X}_{1,n_{1,\kappa}}^*} \hat{\theta}_{\text{VRS}}(X_i, Y_i \mid \mathcal{X}_{1,n_1}) \quad (3.12)$$

and

$$\hat{\mu}_{\text{CRS},n_{2,\kappa}} = n_{2,\kappa}^{-1} \sum_{(X_i, Y_i) \in \mathcal{X}_{2,n_{2,\kappa}}^*} \hat{\theta}_{\text{CRS}}(X_i, Y_i \mid \mathcal{X}_{2,n_2}). \quad (3.13)$$

Then under the null hypothesis of CRS and the assumptions of Kneip et al. (2015, Theorem 4.4) and Kneip et al. (2016, Theorem 1),

$$\hat{T}_{2,n} := \frac{(\hat{\mu}_{\text{VRS},n_{1,\kappa}} - \hat{\mu}_{\text{CRS},n_{2,\kappa}}) - (\hat{B}_{\text{VRS},\kappa,n_{1,\kappa}} - \hat{B}_{\text{CRS},\kappa,n_{2,\kappa}})}{\sqrt{\frac{\hat{\sigma}_{\text{VRS},n_{1,\kappa}}^2}{n_{1,\kappa}} + \frac{\hat{\sigma}_{\text{CRS},n_{2,\kappa}}^2}{n_{2,\kappa}}}} \xrightarrow{\mathcal{L}} N(0,1) \quad (3.14)$$

as $n \rightarrow \infty$ provided $(p+q) \geq 5$.

As discussed by Kneip et al. (2016, Section 3.2), one might consider computing estimates $\hat{\theta}_{\text{VRS}}(X_i, Y_i \mid \mathcal{X}_n)$ and $\hat{\theta}_{\text{CRS}}(X_i, Y_i \mid \mathcal{X}_n)$ for $i = 1, \dots, n$ using the full sample $\mathcal{X}_n$, and then computing the sample means

$$\hat{\mu}_{\text{VRS},n}^{\text{full}} = n^{-1} \sum_{i=1}^n \hat{\theta}_{\text{VRS}}(X_i, Y_i \mid \mathcal{X}_n) \quad (3.15)$$

and

$$\hat{\mu}_{\text{CRS},n}^{\text{full}} = n^{-1} \sum_{i=1}^n \hat{\theta}_{\text{CRS}}(X_i, Y_i \mid \mathcal{X}_n) \quad (3.16)$$

computed using all of the n observations in $\mathcal{X}_n$. One might then build a statistic analogous to $\hat{T}_{1,n}$ defined in (3.11) based on the difference $\hat{\mu}_{\text{VRS},n}^{\text{full}} - \hat{\mu}_{\text{CRS},n}^{\text{full}}$. But, the resulting

statistic would have a degenerate limiting distribution under the null since the variance of $n^a (\hat{\mu}_{\text{VRS},n}^{\text{full}} - \hat{\mu}_{\text{CRS},n}^{\text{full}})$ collapses to zero as $n \rightarrow \infty$ for any $a \leq 1/2$ due to the result in Kneip et al. (2016, equation 22), making inference impossible. See Kneip et al. (2016, p. 440) for additional discussion.

As discussed by Kneip et al. (2016), depending on the value of $(p + q)$, either $\hat{T}_{1,n}$ or $\hat{T}_{2,n}$ (but not both, except when $p + q = 5$) can be used to test the null hypothesis of constant returns to scale, with critical values obtained from the standard normal distribution. For $j \in \{1, 2\}$ and one sample-split, the null hypothesis of CRS is rejected if $\hat{p} = 1 - \Phi(\hat{T}_{j,n}) < \alpha$ is less than, say, .1, .05, or .01. For $(p + q) = 5$, one should expect tests based on $\hat{T}_{2,n}$ to have better size properties than tests based on $\hat{T}_{1,n}$ due to the order of negligible remainder terms that are ignored in either case. See Kneip et al. (2015) for details, and Kneip et al. (2021) for additional discussion in a similar setting.¹⁰

3.3 The Kneip et al. (2016) Test of Convexity versus Non-Convexity

The vast majority of non-parametric studies of efficiency in production settings rely on DEA estimators and thus implicitly, if not explicitly, assume that the underlying production set is convex. Statistical consistency of DEA estimators requires convexity of the production set, while consistency of FDH estimators does not depend on convexity. Convexity of the production set is typically assumed without testing. However, Koopmans, 1957 and McFadden, 1978) note that the widespread use of convexity assumptions in the theory of the firm is due to nothing other than the resulting mathematical convenience. Others, including Farrell (1959), Blinder (1981), Malinvaud (1985); Ramey (1991), Varian (1992), Mas-Colell et al. (1995) and Hall (2000) discuss why convexity of the production set may not be an

¹⁰ The statistics described here are based on comparisons of means of efficiency estimates, one which imposes the null and the other which does not. Alternatively, one can build statistics based on the integrated difference between two frontier estimates, one which imposes the null and one which does not. However, simulation evidence does not indicate any advantages of these alternative statistics in terms of size or power of the resulting tests. See the separate Appendix A for details and discussion.

appropriate assumption. Whether a production set is convex or not convex is necessarily an empirical question in any realistic problem. Convexity is a strong assumption, and should be tested whenever DEA estimators are used since their statistical consistency depends on convexity of the production set. By contrast, FDH estimators do not require convexity, but their convergence rate is slower than that of DEA estimators for given values of $(p + q)$.¹¹

The test of convexity versus non-convexity of Ψ proposed by Kneip et al. (2016) is similar in spirit to the test of CRS versus non-CRS discussed above in Section 3.2. The test is based on the difference between a sample mean of FDH estimates and a sample mean of VRS-DEA estimates. Under the null hypothesis of convexity of Ψ , the difference between the two sample means should be “small,” while under the alternative hypothesis that Ψ is not convex, the difference should be large. The FDH efficiency estimator does not depend on convexity and is consistent in either case, but the VRS-DEA estimator is inconsistent if Ψ is not convex. Under either convexity or non-convexity, $\hat{\Psi}_{\text{FDH}}(\mathcal{X}_n) \subseteq \Psi$. Under convexity, $\hat{\Psi}_{\text{VRS}}(\mathcal{X}_n) \subseteq \Psi$, but under the alternative hypothesis, $\hat{\Psi}_{\text{VRS}}(\mathcal{X}_n) \not\subseteq \Psi$.

As with the test statistics in (3.11) and (3.14) presented in Section 3.2, independence between the sample means of FDH and VRS-DEA estimates is needed to obtain a non-degenerate limiting distribution under the null. Kneip et al. (2016) use a larger number of observations for the FDH estimates than for the VRS-DEA estimates due to the slower convergence rate of the FDH estimator by setting $n_1 + n_2 = n$ and $n_1^{2/(p+q+1)} = n_2^{1/(p+q)}$, and then solving for n_1 and n_2 as described by Kneip et al. ((2016), Section 3.3). Then let $\mathcal{X}_{1,n_1}$ and $\mathcal{X}_{2,n_2}$ denote mutually exclusive, collectively exhaustive sub-samples of sizes n_1 and n_2 of the full sample $\mathcal{X}_n$.

Let

$$\hat{\mu}_{\text{FDH},n_2} = n_2^{-1} \sum_{(X_i, Y_i) \in \mathcal{X}_{2,n_2}} \hat{\theta}_{\text{FDH}}(X_i, Y_i \mid \mathcal{X}_{2,n_2}) \quad (3.17)$$

¹¹A number of recent empirical studies have tested and rejected convexity in a variety of applications. See, for example, Apon et al. (2015), O’Loughlin and Wilson (2021), Wilson and Zhao (2022) and Simar and Wilson (2024).

denote the sample mean of FDH efficiency estimates computed from the sub-sample $\mathcal{X}_2$ obtained by randomly splitting the full sample $\mathcal{X}_n$ as described in Section 3.2. Let

$$\hat{\mu}_{\text{FDH},n_2,\kappa} = n_{2,\kappa}^{-1} \sum_{(X_i, Y_i) \in \mathcal{X}_{2,n_2,\kappa}^*} \hat{\theta}_{\text{FDH}}(X_i, Y_i \mid \mathcal{X}_{2,n_2}) \quad (3.18)$$

denote the sample mean of the same FDH estimates, but over only $n_{2,\kappa}$ of the n_2 observations in $\mathcal{X}_{2,\kappa}$. In both (3.17) and (3.18), the efficiency estimates are computed using all of the observations in $\mathcal{X}_{2,n_2}$. Finally, let $\hat{B}_{\text{FDH},\kappa,n_2}$ denote the generalized jackknife estimate of the bias of $\hat{\mu}_{\text{FDH},n_2}$, computed as described by Kneip et al. (2016), and let

$$\hat{\sigma}_{\text{FDH},n_2}^2 = n_2^{-1} \sum_{(X_i, Y_i) \in \mathcal{X}_{2,n_2}} \left[\hat{\theta}_{\text{FDH}}(X_i, Y_i \mid \mathcal{X}_{2,n_2}) - \hat{\mu}_{\text{FDH},n_2} \right]^2. \quad (3.19)$$

Note that the convergence rate of the FDH estimator dominates the rate of the VRS-DEA estimator. As shown by Kneip et al. (2016), under appropriate assumptions including the null hypothesis of convexity of Ψ , Kneip et al. (2015, Theorem 4.2) can be used to establish that

$$\hat{T}_{3,n} := \frac{(\hat{\mu}_{\text{FDH},n_2} - \hat{\mu}_{\text{VRS},n_1}) - (\hat{B}_{\text{FDH},\kappa,n_2} - \hat{B}_{\text{VRS},\kappa,n_1})}{\sqrt{\frac{\hat{\sigma}_{\text{FDH},n_2}^2}{n_2} + \frac{\hat{\sigma}_{\text{VRS},n_1}^2}{n_1}}} \xrightarrow{\mathcal{L}} N(0, 1) \quad (3.20)$$

as $n \rightarrow \infty$ provided $(p + q) \leq 3$. Alternatively, if $(p + q) \geq 3$, then

$$\hat{T}_{4,n} := \frac{(\hat{\mu}_{\text{FDH},n_{2,\kappa}} - \hat{\mu}_{\text{VRS},n_{1,\kappa}}) - (\hat{B}_{\text{FDH},\kappa,n_2} - \hat{B}_{\text{VRS},\kappa,n_1})}{\sqrt{\frac{\hat{\sigma}_{\text{FDH},n_2}^2}{n_{2,\kappa}} + \frac{\hat{\sigma}_{\text{VRS},n_1}^2}{n_{1,\kappa}}}} \xrightarrow{\mathcal{L}} N(0, 1) \quad (3.21)$$

as $n \rightarrow \infty$.

Similar to the test of CRS versus non-CRS, either $\hat{T}_{3,n}$ or $\hat{T}_{4,n}$ in cases where $(p + q) = 3$. For reasons similar to those discussed at the end of Section 3.2, one should expect $\hat{T}_{4,n}$ to yield better size properties than $\hat{T}_{3,n}$ when $(p + q) = 3$.

Analogous to the discussion near the end of Section 3.2, one might consider computing sample mean

$$\hat{\mu}_{\text{FDH},n}^{\text{full}} = n^{-1} \sum_{(X_i, Y_i) \in \mathcal{X}_n} \hat{\theta}_{\text{FDH}}(X_i, Y_i \mid \mathcal{X}_n) \quad (3.22)$$

using the full set of observations in $\mathcal{X}_n$ and use this with (3.15) to construct a test statistic based on the difference $\hat{\mu}_{\text{VRS},n}^{\text{full}} - \hat{\mu}_{\text{FDH},n}^{\text{full}}$. However, as discussed by Kneip et al. (2016, Section 3.3), the resulting test statistic has zero asymptotic variance, and hence a degenerate distribution under the null, for reasons similar to those given near the end of Section 3.2, thereby rendering inference impossible. See Kneip et al. (2016, Section 3.3) for additional discussion.

As noted earlier, tests based on the statistics in (3.11)–(3.14) and (3.20)–(3.21) are valid for a single split of the original sample, but two researchers working independently may arrive at different conclusions unless they use the same random number generator (RNG) and the same seed value to split the sample. The bootstrap method described by Simar and Wilson (2020) provides a way to remove this ambiguity, but at the cost of a large, perhaps prohibitive, computational burden. The next section describes our new method that substantially reduces this computational expense.¹²

4 Removing Uncertainty Introduced by Multiple Sample-Splits

4.1 A New Decision Rule

As discussed above, the tests of CRS versus non-CRS and of convexity versus non-convexity of the production set require randomly splitting the sample $\mathcal{X}_n$ into two sub-samples so that the test statistics $\hat{T}_{\ell,n}$, $\ell \in \{1, \dots, 4\}$ have non-degenerate limiting distributions under the null. The asymptotic normality obtained for each of the statistics makes it easy to compute the corresponding p -value $\hat{p}_{\ell,n} = 1 - \Phi(\hat{T}_{\ell,n})$, where $\Phi(\cdot)$ denotes the standard normal distribution function. By the probability integral transform, necessarily $\hat{p}_{\ell,n}$ is asymptotically distributed uniform on $(0, 1)$ under the particular null hypothesis that

¹²Similar to the observation in footnote 10, one can instead construct statistics based on the integrated difference between two frontier estimates, one which imposes convexity and the other which does not. See the separate Appendix A for discussion and details.

is tested. For a single sample-split, the null hypothesis being testing is rejected when the realized p -value is small, e.g., less than .1, .05 or .01. Tests based on a single sample-split are thus asymptotically valid, but the random splitting introduces some noise and arbitrariness. For example, two researchers might draw different conclusions when using the same data if they use different RNGs or different seeds when using the same RNG to split the sample. A dishonest researcher might be tempted to try different random splits until the desired result is attained.

To remove the uncertainty introduced by the random splitting of the sample, suppose the sample is split S times. In order to simplify notation in the discussion that follows, let $s = 1, \dots, S$ index the sample-splits, and let $\hat{\tau}_s$ denote the appropriate statistic (i.e., one of $\hat{T}_{\ell,n}$, $\ell \in \{1, \dots, 4\}$) for testing the given null hypothesis H_0 computed from the s th sample-split of a sample of size n . Denote the corresponding p -value by $\hat{p}_s$. Each of the S p -values is asymptotically distributed uniform on $(0, 1)$ under H_0 , but the S p -values cannot be independent, nor can the test statistics $\hat{\tau}_s$ be independent. The dependence arises from the fact that when the sample $\mathcal{X}_n$ is split into two sub-samples of size n (assuming n is even for simplicity), and then split again, the two subsamples resulting from the second split will contain about half of the observations in each of the two sub-samples resulting from the first split.

The dependence means that the test statistics $\{\hat{\tau}_s\}_{s=1}^S$ cannot simply be averaged, nor is it obvious how one can combine the information in the set $\{\hat{p}_s\}_{s=1}^S$ of p -values to make inference. Moreover, the dependence across sample-splits is likely heavy, and is of complicated form. If the p -values were independent, then one could adapt the modified Bonferroni procedure for independent multiple tests introduced by Simes (1986) by rejecting the null hypothesis if for any $s = 1, \dots, S$, $\hat{p}_{(s)} \leq s\alpha/S$ for a test of size α , where $\hat{p}_{(1)}, \dots, \hat{p}_{(S)}$ denote the ordered p -values from the S independent samples. But in the present context, this cannot yield valid inference due to the dependence across sample-splits.

Clearly, the null hypothesis H_0 should be rejected if “too many” of the p -values are “small.” A useful rule for rejecting (or not rejecting) H_0 for testing at some predefined level α (e.g., $\alpha = 0.1, 0.05$ or 0.01) is to reject H_0 if more than $\alpha \times 100$ -percent of the S sample-splits yield p -values less than some predetermined level γ . The following proposition provides a way to choose γ .

Proposition 4.1. *For a test of H_0 given the set $\{\hat{p}_s\}_{s=1}^S$ of p -values resulting from S splits of the sample $\mathcal{X}_n$, let $\beta = \#\{\hat{p}_s < \gamma\}_{s=1}^S$ and reject H_0 if $\beta > \alpha S$ for $\alpha \in (0, 1)$. If $\gamma = \alpha^2$, the rate of type-I error using this rule is less than or equal to α as $n \rightarrow \infty$.*

Proof. Since $\hat{p}_s \mid H_0 \xrightarrow{\mathcal{L}} \text{uniform}(0, 1)$, $E(\mathbb{1}(\hat{p}_s < \gamma) \mid H_0) \rightarrow \gamma$ as $n \rightarrow \infty$ where $\mathbb{1}(\cdot)$ denotes the indicator function. Then $E(\beta \mid H_0) = S \times E(\mathbb{1}(\hat{p}_s < \gamma) \mid H_0) \rightarrow S\gamma$ as $n \rightarrow \infty$. Applying the stated rule, and using Markov’s rule,

$$\begin{aligned} \Pr(\text{reject } H_0 \mid H_0) &= \Pr(\beta \geq \alpha S \mid H_0) \\ &\leq \frac{E(\beta \mid H_0)}{\alpha S} \\ &\rightarrow \frac{\gamma}{\alpha} \end{aligned} \tag{4.1}$$

as $n \rightarrow \infty$. Setting $\gamma = \alpha^2$ yields the result. ■

Note that the conditioning on the null hypothesis H_0 in the proof amounts to conditioning on the event “ H_0 is true.” The Proposition ensures that rejecting H_0 if the number of the S p -values corresponding to the S sample-splits that are smaller than α^2 is bigger than αS yields a test with asymptotic size less than or equal to α .¹³

¹³Recall that the rate of type-I error for a given test is the probability of rejecting H_0 when H_0 is true, and this is called the *size* of the test. The probability of rejecting H_0 when H_0 is false is called the *power* of the test. The rate of type-II error is the probability of failing to reject H_0 when H_0 is false. There is a fundamental trade-off between the type-I and type-II errors in the sense that one cannot simultaneously reduce the probabilities of both.

4.2 Other Decision Rules

Various proposals for merging p -values from different tests have been proposed in the statistical literature. Examples include Hommel (1983), Simes (1986), Benjamini and Hochberg (1995), Vovk and Wang (2020), DiCiccio et al. (2020), Benjamini and Yekutieli (2001) and Vovk et al. (2022). Many of the methods that have been proposed result in tests with low power (at least in our setting) or require one or more tuning parameters that have a large influence on test outcomes and which must (somehow) be optimized. For example, DiCiccio et al. (2020, Procedure 2.1) involves choosing a nominal test size α , fixing $r \in (0, 1]$ and then rejecting the null hypothesis if the proportion of p -values less than or equal to $r\alpha$. But, the choice of r is critical. If r is too small, the test will over-reject, and if it is too large, the test will under-reject. No guidance is given for how the tuning parameter r might be chosen. See the separate Appendix D for additional discussion and an illustration.

Vovk and Wang (2020) consider various rules for combining information from multiple tests using merging functions. We examine 13 of the rules proposed by Vovk and Wang (2020) in the separate Appendix E. Our simulations suggest that with $S = 1,000$ sample-splits, eleven of the rules have little or no power in many cases while two severely over-reject. Reducing the number of sample-splits to only 10 improves the power of some of the rules, but with only 10 splits, there remains substantial ambiguity due to the random splitting as discussed below in Section 5.2.¹⁴ See the Separate Appendix E for additional discussion and details.

In a different setting, where S different hypotheses H_s are tested, Benjamini and Yekutieli (2001) examine properties of the procedure proposed earlier by Benjamini and Hochberg

¹⁴In Section 5.2, we show that using only 10 splits with the rule of Proposition 4.1 leaves substantial ambiguity due to the random sample-splitting. While we do not examine this question explicitly using the Vovk and Wang (2020), we see no reason to expect a different outcome.

(1995) for the case of multiple, dependent tests. Benjamini and Hochberg define

$$\widehat{k} := \max_{s=1, \dots, S} \left\{ s \mid \widehat{p}_{(s)} \leq \frac{s}{S} \alpha \right\} \quad (4.2)$$

and reject the corresponding hypotheses $H_{(1)}, \dots, H_{(\widehat{k})}$. Under quite general assumptions on the structure of the dependency among the tests, Benjamini and Yekutieli (2001) prove that this procedure controls the false discovery rate at a level less than or equal to α .

In the context of each of the tests considered in this paper, there is only one hypothesis to be tested. Nonetheless, the Benjamini and Hochberg rule provides a potentially useful way to “tune” the rejection rule for controlling type-I error by adopting the following rule: compute $\widehat{k}$ using (4.2), and reject the null hypothesis if $\widehat{k} > \alpha S$, i.e., if $\widehat{k}$ is too large—larger than $(100 \times \alpha)$ -percent of the S p -values obtained from S sample-splits. While we have no formal proof that this rule controls the rejection rate to result in tests of size α , the results of Benjamini and Yekutieli (2001) suggest that the rule should result in a rate of type-I error less than or equal to α . We discuss in the separate Appendix C results from extensive Monte Carlo experiments that seem to confirm our assertion. The Monte Carlo evidence reveals that tests based on the rule in Proposition 4.1 and on the modified Benjamini and Hochberg (1995) rule are typically conservative, with realized sizes smaller than the nominal size α , but with power that increases rapidly with increasing departures from the null hypothesis. The conservative nature of the new tests contrast with the simulation results of Simar and Wilson (2020), where the tests often have realized sizes larger than α . The Monte Carlo evidence also suggests that the performance of tests based on the rule in Proposition 4.1 in terms of size and power is often better than tests based on the modified Benjamini and Hochberg (1995) rule, but the differences are not large. Moreover, we provide theoretical justification for the new rule in the proof of Proposition 4.1.

4.3 The Simar and Wilson (2020) Bootstrap Method

Before turning to the Monte Carlo evidence, it is useful to recall the bootstrap method proposed by Simar and Wilson (2020). The solution offered by Simar and Wilson (2020) involves computing the sample mean

$$\bar{\hat{\tau}} = S^{-1} \sum_{s=1}^S \hat{\tau}_s \quad (4.3)$$

of the test statistics obtained from the S splits and using the corresponding p -values to compute the statistic

$$\hat{K} = \sup_{u \in [0,1]} \left(\hat{F}_{\hat{p}}(u) - u \right) \quad (4.4)$$

similar to the Kolmogorov-Smirnov statistics for a one-sample test, where $\hat{F}_{\hat{p}}(\cdot)$ denotes the empirical distribution function of the $\{\hat{p}_s\}_{s=1}^S$.¹⁵ Then, after the initial S splits, a bootstrap with N_B replications is conducted to obtain approximations of the distributions of the statistics $\bar{\hat{\tau}}$ and $\hat{K}$. On each replication, a bootstrap sample of size n is generated under the conditions of the null, and then this bootstrap sample is split S times. The resulting test statistics and p -values are combined to yield two bootstrap values $\bar{\hat{\tau}}^*$ and $\hat{K}^*$. After N_B replications, the bootstrap values are used to compute critical values for $\bar{\hat{\tau}}$ and $\hat{K}$, leading to two different tests of the null hypothesis.

Note that the bootstrap is used to estimate tail probabilities, and a large number of replications is needed to do this with a useful degree of accuracy. Simar and Wilson (2020) recommend setting $N_B = 1,000$ based on experience with simulations. Simar and Wilson (2020) use only $S = 10$ splits in their simulations, and note after performing several experiments with $S = 100$ that the increase produced little difference in terms of the achieved sizes or power of the tests considered. But while this is the case in a Monte Carlo framework, in an applied setting, 10 splits are insufficient to remove the ambiguity resulting from the random

¹⁵Simar and Wilson (2020) write their statistic in terms of the absolute value of the difference $\hat{F}_{\hat{p}}(u) - u$, as in the usual Kolmogorov-Smirnov statistic, but only positive values of the difference identify the alternative from the null hypothesis.

splitting. In other words, one may obtain one result with 10 splits, and a different result with another 10 splits. Even with 100 splits, there may be some variation. Recent empirical papers by O’Loughlin and Wilson (2021) Wilson and Zhao (2022) and Simar and Wilson (2024) use 1,000 sample splits in their tests. In any case, with N_B replications, S splits and N_J jackknife replications (for estimation of bias terms), the method of Simar and Wilson (2020) requires computation of $SnN_j(N_B + 1)$ efficiency estimates. By contrast, the new procedure described above requires computation of only SnN_j efficiency estimates. Hence, one should expect the new procedure to require roughly 0.001 times the time required by the Simar and Wilson (2020) bootstrap method with S splits and 1,000 bootstrap replications.

5 Monte Carlo Evidence

5.1 Simulation Framework

We employ a DGP similar to those used by Kneip et al. (2016) and Simar and Wilson (2020) to examine the performance in terms of size and power of tests CRS versus non-CRS and of convexity versus non-convexity of the production set Ψ and to facilitate comparison with earlier work. Specifically, we consider the effects of increasing dimensionality by letting $q = 1$, $p \in \{1, 2, 3, 4, 5\}$, and we consider the effects of increasing sample size by letting $n \in \{50, 100, 150, 200, 1,000\}$ (we report results only for $n = 100$ and 1,000 in the main paper to conserve space; full sets of results for the five different sample sizes are shown in the separate Appendix B). As discussed in Kneip et al. (2016), our methods can be used with similar results in situations where there is more than one output. Here, we use $q = 1$ to simplify simulation of data. In all of the theoretical results on properties of DEA and FDH estimators, including Korostelev et al. (1995a, 1995b), Park et al. (2000), Park et al. (2010), Kneip et al. (1998), Kneip et al. (2008, 2011 2015 2016 2021 2022), Gijbels et al. (1999) and Wilson (2011), the dimensionality ($p + q$) rather than the ratio p/q affects properties of the estimators. Hence we do not expect a loss of generality from simulating only one output.

We model the technology by

$$Y = g \left(\prod_{j=1}^p \left(\tilde{X}^j \right)^{1/p} \right), \quad (5.1)$$

where $\tilde{X}^j$ is the efficient level of the j th input and the function $g(\cdot): \mathbb{R}_+^1 \mapsto \mathbb{R}_+^1$ is either convex, linear or concave. When $g(\cdot)$ is linear and passes through the origin, Ψ^∂ is CRS; otherwise, it is not. When $g(\cdot)$ is convex (from above) or linear, the production set Ψ is convex; otherwise, Ψ is not convex. To describe the function $g(\cdot)$, consider the transformation $(x, y) \mapsto (s, t)$ such that

$$(s, t) = \left(\sqrt{2}(x + y - 1), \sqrt{2}(y - x) \right). \quad (5.2)$$

In (s, t) -space, the coordinate system relative to that in (x, y) -space has been rotated through a clockwise angle of $\pi/4$ radians and then shifted by a distance of $-\sqrt{2}$ to the left. Alternatively, one can view the coordinate system in (x, y) -space relative to that in (s, t) -space as having been rotated through a counter-clockwise angle of $\pi/4$ radians and shifted along a 45-degree ray from the origin in the positive orthant of (x, y) -space by a distance of $\sqrt{2}$. In (s, t) -space, the function $g(\cdot)$ is transformed to a function $\psi(\cdot): \mathbb{R}^1 \mapsto \mathbb{R}^1$ such that

$$t = \psi(s) := \text{sgn}(\delta) \left[c \left(a^2 + \delta^2 s^2 \right)^{1/2} - d \right] \quad (5.3)$$

where $a = 0.5$, $c = 0.75$, $d = 0.375$, $\delta \in \mathbb{R}$ and $\text{sgn}(\cdot)$ is the sign function such that

$$\text{sgn}(\delta) := \begin{cases} -1 & \text{if } \delta < 0; \\ 0 & \text{if } \delta = 0; \\ 1 & \text{if } \delta > 0. \end{cases} \quad (5.4)$$

The parameter δ is used to control the shape and curvature of the function. Note that the transformation from (x, y) -space to (s, t) -space in (5.2) is easily inverted; given a point (s, t) ,

$$(x, y) = \left(\frac{s - t}{2\sqrt{2}} + \frac{1}{2}, \frac{s + t}{2\sqrt{2}} + \frac{1}{2} \right). \quad (5.5)$$

In experiments where we test CRS versus non-CRS, we consider the effects of increasing departures from the null by letting $\delta \in \mathcal{A} = \{0.0, -0.1, -0.2, -0.3, -0.4, -0.6, -1.4\}$. In

experiments where convexity of Ψ is tested, we augment $\mathcal{A}$ with $\mathcal{B} = \{0.1, 0.2, 0.3, 0.4, 0.6, 0.8, 1.0, 1.2\}$, and consider $\delta \in (\mathcal{A} \cup \mathcal{B})$. The left panel in Figure 1 shows various shapes of the function $\psi(s)$ in (s, t) -space corresponding to the different values of δ listed above. For $\delta = 0$, $\psi(s)$ is a flat, horizontal line in (s, t) -space so that $t = 0 \forall s$. In (x, y) -space, the corresponding function $g(x)$ coincides with the 45-degree ray from the origin when $\delta = 0$. The negative values of δ produce curves that are convex from above, while the positive values produce the curves that are concave from above in the right panel of Figure 1. In both cases, curvature increases as $|\delta|$ increases. In the left-hand panel of Figure 1, in (s, t) -space, the rhombus with corners at $(-\sqrt{2}, 0)$, $(0, \sqrt{2})$, $(\sqrt{2}, 0)$, and $(0, -\sqrt{2})$ formed by the dashed lines corresponds to the square with corners at $(0, 0)$, $(0, 2)$, $(2, 2)$ and $(2, 0)$ in (x, y) -space depicted in the right-hand panel

Data are generated by first computing, for a given, fixed value of δ in the set given above, the corresponding value of s , denoted $s_{\max}(\delta)$, where the curve in (5.3) intersects the line $t = \sqrt{2} - s$ or $t = -\sqrt{2} + s$ (depending on whether δ is positive or negative), and then generating uniform random numbers S_i on the interval $(0.7 \max(-s_{\max}(\delta), -0.9\sqrt{2}), 0.7 \min(s_{\max}(\delta), 0.9\sqrt{2}))$. The factor 0.7 used in the endpoints of this interval avoids generating observations with output levels very close to zero which can cause numerical problems. Plugging the S_i into (5.3) for s gives corresponding values T_i ; pairs (S_i, T_i) are then transformed to pairs (V_i, Y_i) using the inverse transformation in (5.5).

If $p = 1$, then set $X_i = V_i \theta_i^{-1}$ where θ_i is a draw from the distribution with density

$$f(t \mid \lambda) = \begin{cases} \lambda t^{-2} e^{-\lambda(t^{-1}-1)} & \forall t \in (0, 1], \\ 0 & \text{otherwise,} \end{cases} \quad (5.6)$$

while setting $\lambda = 2$.¹⁶ Alternatively, if $p > 1$, then a pair (S_i, T_i) is generated and transformed to a pair (V_i, Y_i) as before. To split V_i into p efficient input levels $\tilde{X}_{ij}$, $j = 1, \dots, p$ that will be projected away from the frontier to obtain inefficient “observed” input levels, first

¹⁶Note that for θ_i drawn from the density in (5.6), $1/\theta_i$ has the shifted exponential density $f(t \mid \lambda) = \lambda e^{\lambda(t-1)} \forall t \in [1, \infty)$ and $E(1/\theta_i) = 1.5$.

generate a $(p \times 1)$ vector U_i (with j th element U_{ij}) of uniform deviates on $(0, 1)$. Writing $V_i = \prod_{j=1}^p (\tilde{X}_{ij})^{1/p}$, taking logs and rearranging terms yields $p \log V_i = \sum_{j=1}^p \log \tilde{X}_{ij}$. Then set

$$X_i = \theta_i^{-1} \exp \left(\frac{U_i}{\sum_{j=1}^p U_{ij}} p \log V_i \right), \quad (5.7)$$

where θ_i is a draw from the density in (5.6), again with $\lambda = 2$. For $q = 1$ and any $p = 1, 2, \dots$, the procedure described above yields a simulated sample $\mathcal{X}_n = \{(X_i, Y_i)\}_{i=1}^n$.

In all of our experiments, we use 1,000 Monte Carlo trials. On each trial, we generate a sample of size n , and then for each test considered we record whether the test rejects the null hypothesis of convexity at nominal test sizes $\alpha \in \{0.1, 0.05, 0.01\}$. For each experiment and each test size, we report the proportion among the 1,000 trials where the null is rejected, providing Monte Carlo estimates of rejection rates.

5.2 Performance of Tests Based on Proposition 4.1

Table 1 reports Monte Carlo estimates of rejection rates from experiments with tests of CRS versus non-CRS using the statistics $\hat{T}_{1,n}$ and $\hat{T}_{2,n}$ of Kneip et al. (2016) shown in (3.11) and (3.14) and the simple rule given by Proposition 4.1. In each experiment, we use $S = 1,000$ sample-splits. In the table, “Case I” identifies cases where the statistics based on means over n efficiency estimates are used, while “Case II” indicates those where the statistics based on means over a sub-sample of efficiency estimates are used.

Overall, the results in Table 1 reveal there is a price to pay for increasing dimensionality with a given sample size. As dimensionality increases for a given sample size, the power of the tests diminish for moderate departures from the null. This should not be surprising given that the convergence rates of the VRS-DEA and CRS-DEA estimators depend on the dimensionality of problem. In every case except the one where $n = 100$ and $p = 3$, $q = 1$ in Table 1, the tests are conservative, with rejection rates smaller than the nominal test size.

Monte Carlo estimates of rejection rates obtained from experiments with tests of con-

vexity versus non-convexity of the production set using the statistics $\widehat{T}_{3,n}$ and $\widehat{T}_{4,n}$ of Kneip et al., 2016 shown in (3.20) and (3.21) and the rule in Proposition 4.1 are shown in Table 2. Whereas the Monte Carlo experiments of Kneip et al. (2016) and Simar and Wilson (2020) included only two cases where the null is true, corresponding to $\delta = -1.4$ and $\delta = 0.0$ in our framework, we report in Table 2 results for a range of negative values of δ to give an idea of how our tests perform under varying degrees of convexity of Ψ . In Table 2, “Case I” again identifies cases where the statistics based on means over n efficiency estimates are used, while “Case II” indicates those where statistics based on means of sub-samples of efficiency estimates are used.

The results in Table 2 reveal that the tests of convexity of the production set using the rule in Proposition 4.1 are conservative, with realized size or rejection rates smaller than the nominal test size in almost every case. The exceptions occur at all three levels where $p = q = 1$ and $n = 100$, and where $\alpha = 0.01$, $p = 2$, $q = 1$ and $n = 100$, where the tests over-reject when $\delta = 0.0$. Overall, power is seen to decrease with increasing dimensionality, and to increase with increasing departure from the null, as one should expect.

Note that in any testing situation, there is typically a tradeoff between size and power. In classical hypothesis testing, it is customary to choose a priori a small size for a test and then accept the power that comes with that choice. By design, null hypotheses are constructed and tested where one hopes to minimize type-I error. A conservative test, i.e., one that results in realized size smaller than the nominal size that is chosen a priori, is typically preferred to an aggressive test that results in realized size greater than the nominal size. In view of this, the results in Tables 1–2 are an improvement over the results in Kneip et al. (2016, Tables 3 and 5) based on tests with a single sample-split, resulting in over-rejection in many cases, even when $n = 10,000$.

5.3 How Many Times Should the Sample be Split?

As noted earlier, in all of our Monte Carlo experiments discussed so far, we use $S = 1,000$ sample-splits. But is this excessive? If the sample size n is even and the sample is split into two subsamples each of size $n/2$, then there are $\binom{n}{n/2}$ possible sample splits. With only $n = 50$ observations, there are more than 1.26×10^{24} possible sample-splits. With $n = 100$ observations, there are more than 1.0089×10^{29} possible splits. So one can never hope to examine all possible splits, and the empirical researcher must use a sample of the possible splits. But how big should this sample of splits be in order to eliminate (most) of the ambiguity or randomness introduced by the random splitting?

To gain some insight into this question, we conducted a limited set of additional experiments. In these experiments, we use the same framework to simulate data as that used for the experiments reported in Table 1. We consider $\delta \in \{0.0, -0.1, -0.2, -0.3, -0.4, -0.6\}$, $n \in \{100, 200, 500\}$, $p \in \{1, 2, 3\}$ and $q = 1$. In each experiment, the values of δ , n and p are fixed. On each trial, for a given sample, we examine the outcomes (i.e., either rejection or non-rejection of the null hypothesis) of ten separate, independent tests of the null hypothesis versus the alternative hypothesis with each test using S different sample-splits. On each trial, this amounts to simulating the outcomes of ten different empirical researchers working independently with the same sample.

We consider $S \in \{1, 5, 10, 50, 100, 200, 1,000, 5,000, 10,000\}$. For a given trial and number S of sample-splits, we record whether the ten independent tests lead to the same conclusion (i.e., reject or failure to reject) for tests with nominal sizes $\alpha \in \{0.10, 0.05, 0.01\}$. For $S = 1$, we use the Kneip et al. (2016) test (relying on asymptotic normality) on each of ten different sample-splits. For $S \geq 1$, we obtain the p -values from the appropriate Kneip et al. (2016) test on each split in each of 10 sets of S splits, and then use the rule in Proposition 4.1 with each of the ten groups of S p -values, recording whether each of the ten tests rejects or fails to reject the null, and then whether the tests agree. Note that with

$S = 10,000$, we require $10 \times S = 100,000$ sample splits on each trial for these experiments.

Table 3 shows, for $n = 100$ and each value of δ , p , α and S , the number of trials among 1,000 Monte Carlo trials where the ten tests agree (i.e., all indicate rejection of the null or all indicate failure to reject the null). The maximum possible value in each cell is of course 1,000, corresponding to the number of Monte Carlo trials. This value is achieved in only two cases, where $S = 5,000$ or $10,000$. Overall, the results indicate that with only one sample split, there is often low probability that ten researchers working independently with the same sample would arrive at the same conclusions. For example, when $\delta = 0.0$ (so the null is true), $\alpha = 0.1$ and $p = q = 1$, only 163 of 1,000 trials yield agreement among the ten tests. So the probability that ten researchers working independently would arrive at the same conclusion using the same data is estimated to be only 0.163.

Overall, the results in Table 3 suggest that for fixed values of δ and S , (i) agreement varies, but not systematically, as the test size is reduced, (ii) agreement tends to decrease as the number $(p + q)$ of dimensions increases when either the null is true or the departure from the null is small (e.g., $\delta = -0.1$ or -0.2), (iii) agreement tends to increase as the number $(p + q)$ of dimensions increases when the departure from the null is large (e.g., $\delta \leq -0.3$), and (iv) agreement tends to increase with the number S of sample-splits. The results in Table 3 show that among the cases considered, with $S = 1,000$ sample splits the probability that ten different researchers working independently on the same sample have probability of 0.55 or greater of reaching the same conclusion. But, the results also indicate that using more than 1,000 splits should be expected to further reduce the ambiguity. Of course, empirical researchers are confronted with a trade-off between reducing ambiguity and time spent waiting for results. For a given number of dimensions $(p + q)$ and a given sample size n , the computational cost of our new method increases approximately linearly with S , and so researchers might consider using more than 1,000 splits unless their sample-size is very large. Inverting the famous quote often attributed to Mies van der Rohe, clearly *more is*

less, i.e., more splits lead to less disagreement across tests using the same data but different splits.

Note that Simar and Wilson (2020) report estimates of rejection rates from experiments using only 10 sample-splits due to the computational burden of their bootstrap procedure. As noted earlier in Section 4, several recent empirical papers, including O’Loughlin and Wilson (2021), Wilson and Zhao (2022) and Simar and Wilson (2024) use the Simar and Wilson (2020) method with 1,000 sample splits to test model features, but doing so required use of a massively parallel computer system. The simulations for the results in Tables 1–2 were run on Intel Xeon ES-2665 processors. In Table 1 where results for tests of CRS versus non-CRS are given, the simulations with $S = 1,000$ sample splits, $n = 1,000$ and $p = 1, \dots, 5$ used on average 21.66, 34.46, 47.53, 63.56 and 84.18 minutes, respectively. In Table 2 where results for tests of convexity versus non-convexity are given, the simulations with $S = 1,000$ sample splits, $n = 1,000$ and $p = 1, \dots, 5$ used on average 4.02, 4.98, 5.11, 6.40 and 6.21 minutes, respectively. (tests of convexity require less time than tests of CRS because FDH estimates are used instead of CRS-DEA estimates, and the FDH estimates can be computed faster than CRS-DEA estimates). Since computation times are linear in S , one could easily use more than 1,000 splits with 1,000 observations and up to six dimensions. For example, with $(p + q = 6)$ and $n = 1,000$, testing CRS versus non-CRS using $S = 5,000$ sample-splits should be expected to require about seven hours on a single Xeon E-2665 processor, and testing convexity versus non-convexity with $S = 10,000$ splits should require about one hour. Moreover, the latest computers use even faster processors; for example, a single core on the Apple M3 Max processor is 1.5–3.0 times faster than a single core on the Intel Xeon ES-2665 processor. In addition, researchers faced with five or six dimensions can likely use dimension-reduction methods as discussed by Wilson (2018), which would make testing even faster. Our new testing method based on the decision rule in Proposition 4.1 makes testing feasible for any researcher using an ordinary, desktop or laptop computer of recent

vintage without parallel computation with sample sizes similar to those seen in the empirical literature, which typically involve well less than 1,000 observations.

5.4 Performance of Tests Based on Earlier Methods

In order to facilitate comparison of our new method based on Proposition 4.1 with the methods of Kneip et al. (2016) and Simar and Wilson (2020), we report in Table 4 Monte Carlo results for tests of convexity versus non-convexity using just one sample split and relying on asymptotic normality of our test statistics. In addition, we report in Tables 5–6 results obtained using the bootstrap method of Simar and Wilson (2020). For the bootstrap results, we use $N_B = 1,000$ bootstrap replications as recommended by Simar and Wilson (2020), but due to the resulting computational burden we use only $S = 100$ sample-splits in each experiment.

The results in Table 4 are similar to those appearing in Kneip et al. (2016, Table 5). Here, we consider more cases where the null hypothesis is true than considered by Kneip et al. (2016). Similar to the results obtained by Kneip et al. (2016), the results in Table 4 reveal that the tests substantially over-reject $\delta = 0.0$, and even in some cases when $\delta = -1.4$. This contrasts with our new method which typically rejects at rates smaller than the nominal test size. Moreover, the results in Table 4 suggest that the large rejection rates reported by Kneip et al. (2016) for $\delta = 0.0$ are not merely an artifact of CRS and the resulting change in the convergence rate of the VRS-DEA estimator (which has no effect on the rate of the test statistics, which are tied to the slower rates of the FDH estimator). In Table 4, the estimated rejection rates are still well-above the nominal test sizes for values of δ between 0.0 and -1.4 in many cases.

Recall that the bootstrap method of Simar and Wilson (2020) leads to two tests. Table 5 shows results for the test based on averaging test statistics across splits as in (4.3), and Table 6 gives results for the test based on the Kolmogorov-Smirnov statistic in (4.4). Both

tables reveal substantial over-rejection in most cases where $-0.6 \leq \delta \leq 0.0$, even with 1,000 observations. In fact, the over-rejection with the bootstrap is in many cases worse than in Table 4 where only one split is used. For example, with $n = 1,000$, $\delta = 0.0$, $(p + q) = 6$ and nominal test size 0.1, the rejection rates are 0.173, 0.393 and 0.359 in Tables 4–6, respectively. With only two dimensions (i.e., with $p = q = 1$, the corresponding rejection rates are 0.186, 0.232 and 0.215. By contrast, in Table 2, the rejection rates are 0.000 and 0.001 for $(p + q) = 5$ and 1, and in Table A.2 the corresponding rates are both 0.000.

It is well-known (e.g., see Hall, 1992) that in regression problems where data show evidence of ill-behaved higher moments (e.g., when data are heavily skewed), bootstrap methods may provide more reliable inference than classical methods that rely on asymptotic normality. But our situation is rather different from the regression setting. In the simulations we have conducted, the bootstrap method of Simar and Wilson (2020) yields tests that substantially over-reject, leading to rates of type-I errors larger than nominal test sizes, whereas our new method provides conservative tests. In almost any testing scenario, one should hope to minimize type-I error (i.e., rejection of the null when the null is true) while accepting the resulting level of power (i.e., ability to reject the null when false). This, combined with the fact that the computational burden of the Simar and Wilson (2020) bootstrap renders the method impractical except in very simple, limited conditions, makes our new method based on Proposition 4.1 especially attractive to empirical researchers.

There is, however, a price to pay for the better size properties of our new method. For small departures from the null, our new tests have less power than either the Kneip et al. (2016) tests or the Simar and Wilson (2020) tests. For example, with $n = 1,000$, $p = q = 1$ and nominal test size 0.1, the rejection rates for $\delta = 0.01$, 0.02 and 0.03 are 0.006, 0.647 and 1.000 in Table 2 and 0.000, 0.019 and 0.918 in Table A.2 using our new method. In Table 4 where only one split is used, the corresponding rates are 0.216, 0.376 and 0.668. With $\delta = 0.02$, the Kneip et al. (2016) test has less power than our new method with the

originals statistic, and when $\delta = 0.03$, the Kneip et al. (2016) test has less power than our new method with either the original or alternative statistic. The two bootstrap tests yield corresponding rates of 0.296, 0.847 and 1.000 (Test #1, based on (4.3)), and 0.288, 0.791 and 1.000 (Test #2, based on (4.4)). For $\delta = 0.01$ or 0.02 , the bootstrap tests have greater power than either of our new tests, but the new tests have the same power as the bootstrap tests when $\delta = 0.03$.

Careful inspection of Figure 1 reveals that with $\delta = 0.01$ or 0.02 , the departure from the null is quite small. While the two curves lying just above the 45-degree line in the right panel of Figure 1 are concave from above, their curvature is very slight. The Kneip et al. (2016) and Simar and Wilson (2020) tests have greater power than our new tests for these two cases in part because they yield much larger rejection rates than the new tests when $\delta = 0.0$. But when $\delta = 0.03$, the new tests have power similar to the bootstrap tests.

As noted above, Simar and Wilson (2020) used only 10 sample-splits in their Monte Carlo trials, and here, in Tables 5–6, we use only 100 splits. The rejection rates shown by Simar and Wilson (2020, Tables 5–6) for $\delta = 0.0$ are in many cases smaller than the corresponding values in Tables 5–6. However, as demonstrated above in Section 5.2, 10 splits or even 100 splits are not enough to stabilize the results. But the computational burden imposed by the bootstrap makes it costly to use more splits. In the experiment with $\delta = 0.0$, $S = 100$, $p = q = 1$ and $n = 1,000$ that produced the results in Tables 5–6, the time required for each Monte Carlo trial running on Intel Xeon ES-2665 processors ranged from 526.1 to 589.0 minutes, with a mean time of 573.4 minutes. For the same case in Table 2, but with $S = 1,000$ splits instead of only 100 splits, the time required for each Monte Carlo trial ranges from 5.659 to 5.905 minutes, with a mean time of 5.781 minutes. The same experiment was run with only 100 sample-splits, requiring only 0.5215 to 0.6125 minutes for each trial, with a mean time of 0.5721 minutes. This confirms the statement made in Section 1—the time required by our new method is only about 0.001 times the time required by the Simar and Wilson (2020)

method.

Analogous to Tables 4–6, we also report in the separate Appendix F. Monte Carlo results for tests of CRS versus non-CRS using just one sample split with reliance on asymptotic normality of the test statistics in (3.11)–(3.14) (see Table F.4) as well as results obtained using the bootstrap method of Simar and Wilson (2020) (see Tables F.5–F.6). Testing CRS versus non-CRS requires considerably more time than testing convexity versus non-convexity of Ψ . This is due to the fact that when testing CRS, the CRS-DEA estimator replaces the FDH estimator used in the convexity tests, and the FDH estimates can be computed faster than the CRS-DEA estimates. The experiment in Table 1 with 1,000 sample splits, $\delta = 0.0$, $p = q = 1$ and $n = 1,000$ required 22.039 to 26.292 minutes for each Monte Carlo trial, with a mean time of 23.805 minutes. The corresponding experiment in Table A.1 required 26.234 to 36.548 minutes for each trial, with a mean time of 29.506 minutes. By contrast, the corresponding experiment in Tables F.5–F.6 using the Simar and Wilson (2020) bootstrap test, with $n = 1,000$, $N_B = 1,000$ bootstrap replications but only $S = 100$ sample-splits, required 1605.370 to 3611.190 minutes for each trial, with a mean time of 2824.835 minutes. With $n = 1,000$, $(p+q) \geq 3$ and $n = 1,000$, each trial requires more than 72 hours (with only $S = 100$ splits). Consequently, for the cases with $n = 1,000$, only results from experiments where $p = q = 1$ are reported in Tables F.5–F.6 due to the excessive computational burden.¹⁷

Rejection rates obtained using only one sample-split in Table F.4 are comparable to those obtained by Kneip et al. (2016, Table 3). Rejection rates obtained using the Simar

¹⁷Comparing the test statistics in (3.11) and (3.14) reveals that when the sample $\mathcal{X}_n$ is split evenly, the difference in computation time for the test of CRS versus non-CRS versus the test of convexity versus non-convexity using our new approach based on (4.2) is approximately $S \frac{n}{2} (t_{\text{CRS},n,p+q} - t_{\text{FDH},n,p+q})$, where $t_{\text{CRS},n,p+q}$ and $t_{\text{FDH},n,p+q}$ are the expected times to compute the CRS-DEA and FDH efficiency estimates (respectively) for one observation in a sample of size n with p inputs and q outputs. Using the Simar and Wilson (2020) bootstrap method, the difference in computation time for the tests of CRS versus non-CRS versus the test of convexity versus non-convexity is approximately $(N_B + 1) S \frac{n}{2} (t_{\text{CRS},n,p+q} - t_{\text{FDH},n,p+q})$. Even if the difference in time required to compute CRS-DEA and FDH efficiency estimates for a single observation is a fraction of a second, the difference in time required by the test of CRS versus non-CRS versus the test of convexity versus non-convexity can be quite large depending on the method used, the sample size, and the dimensionality (in terms of $p + q$) of the problem.

and Wilson (2020) bootstrap method with $n = 100$ and with $n = 1,000$ for the case with $p = q = 1$ in Tables F.5–F.6 are similar to those obtained by Simar and Wilson (2020) using only 10 observations. The bootstrap tests as well as the test with a single sample-split over-reject, but less severely than corresponding tests of convexity versus non-convexity, whereas our new method yields conservative tests.

6 Summary and Conclusions

Testing model structure is necessary in order to choose an appropriate estimator with which to estimate efficiency. The tests of CRS versus non-CRS and of convexity versus non-convexity developed by Kneip et al. (2016) require randomly splitting the observed sample into two independent sub-samples. Although the tests are theoretically valid for any random split, results can vary across different splits. Aggregating information from multiple sample-splits is complicated by the fact that the test statistics obtained from each split cannot be independent. Simar and Wilson (2020) propose a bootstrap method to deal with this, but the computational burden makes their method impractical with more than a few hundred observations. However, Monte Carlo evidence reveals that the resulting tests tend to over-reject when the null is true. In addition, as discussed in Section 5.2, a large number of sample-splits—1,000 or more—are needed to remove the ambiguity resulting from different sample-splits, and this further exacerbates the computational burden of the Simar and Wilson (2020) method.

We propose a simple rule in Proposition 4.1 with a proof that can be used with multiple splitting without resorting to bootstrap methods. Our new tests are conservative, but retain power comparable with the Simar and Wilson (2020) method except when the departure from the null is slight. Our new method requires approximately 0.1-percent of the computational time required by the Simar and Wilson (2020) method. Given that minimizing type-I errors (at the expense of type-II errors) is typically the goal in hypothesis testing, we believe that

the substantial reduction in computational burden afforded by our new method is worth the price paid in terms of under-rejection when the null is true and loss of power when departures from the null are slight. Our method makes testing feasible for empirical researchers who do not have access to high-performance computing machinery, even when they have samples with several thousands of observations.

References

- Apon, A. W., L. B. Ngo, M. E. Payne, and P. W. Wilson (2015), Assessing the effect of high performance computing capabilities on academic research output, *Empirical Economics* 48, 283–312.
- Banker, R. D., A. Charnes, and W. W. Cooper (1984), Some models for estimating technical and scale inefficiencies in data envelopment analysis, *Management Science* 30, 1078–1092.
- Benjamini, Y. and Y. Hochberg (1995), Controlling the false discovery rate: A practical and powerful approach to multiple testing, *Journal of the Royal Statistical Society. Series B (Methodological)* 57, 289–300.
- Benjamini, Y. and D. Yekutieli (2001), The control of false discovery rate in multiple testing under dependency, *Annals of Statistics* 29, 1165–1188.
- Blinder, A. S. (1981), Retail inventory behavior and business fluctuations, *Brookings Papers on Economic Activity* 2, 443–505.
- Chambers, R. G., Y. Chung, and R. Färe (1996), Benefit and distance functions, *Journal of Economic Theory* 70, 407–419.
- Charnes, A., W. W. Cooper, and E. Rhodes (1978), Measuring the efficiency of decision making units, *European Journal of Operational Research* 2, 429–444.
- Daouia, A., L. Simar, and P. W. Wilson (2017), Measuring firm performance using nonparametric quantile-type distances, *Econometric Reviews* 36, 156–181.
- Daraio, C., L. Simar, and P. W. Wilson (2018), Central limit theorems for conditional efficiency measures and tests of the ‘separability condition’ in non-parametric, two-stage models of production, *The Econometrics Journal* 21, 170–191.
- Deprins, D., L. Simar, and H. Tulkens (1984), Measuring labor inefficiency in post offices, in M. M. P. Pestieau and H. Tulkens, eds., *The Performance of Public Enterprises: Concepts and Measurements*, Amsterdam: North-Holland, pp. 243–267.
- DiCiccio, C. J., T. J. DiCiccio, and J. P. Romano (2020), Exact tests via multiple data splitting, *Statistics and Probability Letters* 166, 108865.
- Durstenfeld, R. (1964), Algorithm 235: Random permutation, *Communications of the ACM* 7, 420.

- Färe, R., S. Grosskopf, and C. A. K. Lovell (1985), *The Measurement of Efficiency of Production*, Boston: Kluwer-Nijhoff Publishing.
- Farrell, M. J. (1957), The measurement of productive efficiency, *Journal of the Royal Statistical Society A* 120, 253–281.
- (1959), The convexity assumption in the theory of competitive markets, *Journal of Political Economy* 67, 377–391.
- Gijbels, I., E. Mammen, B. U. Park, and L. Simar (1999), On estimation of monotone and concave frontier functions, *Journal of the American Statistical Association* 94, 220–228.
- Hall, G. J. (2000), Non-convex costs and capital utilization: A study of production scheduling at automobile assembly plants, *Journal of Monetary Economics* 45, 681–716.
- Hall, P. (1992), *The Bootstrap and Edgeworth Expansion*, New York: Springer-Verlag.
- Hommel, G. (1983), Tests of the overall hypothesis for arbitrary dependence structures, *Biometrical Journal* 25, 423–430.
- Kneip, A., B. Park, and L. Simar (1998), A note on the convergence of nonparametric DEA efficiency measures, *Econometric Theory* 14, 783–793.
- Kneip, A., L. Simar, and P. W. Wilson (2008), Asymptotics and consistent bootstraps for DEA estimators in non-parametric frontier models, *Econometric Theory* 24, 1663–1697.
- (2011), A computationally efficient, consistent bootstrap for inference with non-parametric DEA estimators, *Computational Economics* 38, 483–515.
- (2015), When bias kills the variance: Central limit theorems for DEA and FDH efficiency scores, *Econometric Theory* 31, 394–422.
- (2016), Testing hypotheses in nonparametric models of production, *Journal of Business and Economic Statistics* 34, 435–456.
- (2021), Inference in dynamic, nonparametric models of production: Central limit theorems for Malmquist indices, *Econometric Theory* 37, 537–572.
- (2022), Conical FDH estimators of general technologies, with applications to returns to scale and Malmquist productivity indices. LIDAM Discussion Paper ISBA 2022/24; available at <http://hdl.handle.net/2078.1/264666>.
- Koopmans, T. C. (1957), *Three Essays on the State of Economic Science*, New York: McGraw-Hill, Inc.
- Korostelev, A., L. Simar, and A. B. Tsybakov (1995a), Efficient estimation of monotone boundaries, *The Annals of Statistics* 23, 476–489.
- (1995b), On estimation of monotone and convex boundaries, *Publications de l’Institut de Statistique de l’Université de Paris XXXIX* 1, 3–18.
- Malinvaud, E. (1985), *Lectures on Microeconomic Theory*, Amsterdam: Elsevier Science Publishers B.V.

- Mas-Colell, A., M. D. Whiston, and J. R. Green (1995), *Microeconomic Theory*, New York: Oxford University Press, Inc.
- McFadden, D. (1978), Duality of production, cost, and profit functions, in M. Fuss and D. McFadden, eds., *Production Economics: A Dual Approach to Theory and Applications, Volume 1*, Amsterdam: North-Holland Publishing Company, pp. 3–109.
- O’Loughlin, C. T. and P. W. Wilson (2021), Benchmarking the performance of U.S. municipalities, *Empirical Economics* 60, 2665–2700.
- Park, B. U., S.-O. Jeong, and L. Simar (2010), Asymptotic distribution of conical-hull estimators of directional edges, *Annals of Statistics* 38, 1320–1340.
- Park, B. U., L. Simar, and C. Weiner (2000), FDH efficiency scores from a stochastic point of view, *Econometric Theory* 16, 855–877.
- Ramey, V. A. (1991), Nonconvex costs and the behavior of inventories, *Journal of Political Economy* 99, 306–334.
- Rhodes, E. (1978), *Data Envelopment Analysis and Related Approaches for Measuring the Efficiency of Decision Making Units with an Application to Program Follow Through in U.S. Education*, Ph.D. thesis, Carnegie-Mellon University, School of Urban and Public Affairs, Pittsburgh, PA.
- Sarkar, S. K. and C.-K. Chang (1997), The Simes method for multiple hypothesis testing with positively dependent test statistics, *Journal of the American Statistical Association* 92, 1601–1608.
- Simar, L. and P. W. Wilson (2007), Estimation and inference in two-stage, semi-parametric models of productive efficiency, *Journal of Econometrics* 136, 31–64.
- (2011), Two-Stage DEA: Caveat emptor, *Journal of Productivity Analysis* 36, 205–218.
- (2020), Hypothesis testing in nonparametric models of production using multiple sample splits, *Journal of Productivity Analysis* 53, 287–303.
- (2024), New tools for evaluating the performance of health-care providers using DEA and FDH estimators, in S. Grosskopf, V. Valdmanis, and V. Zelenyuk, eds., *The Cambridge Handbook of Healthcare: Productivity, Efficiency, Effectiveness*, chapter 12, Cambridge, UK: Cambridge University Press, pp. 351–403.
- Simes, R. J. (1986), An improved Bonferroni procedure for multiple tests of significance, *Biometrika* 73, 751–754.
- Varian, H. R. (1992), *Microeconomic Analysis*, New York: W. W. Norton & Co. Third edition.
- Vovk, V., B. Wang, and R. Wang (2022), Admissible ways of merging p -values under arbitrary dependence, *The Annals of Statistics* 50, 351–375.
- Vovk, V. and R. Wang (2020), Combining p -values via averaging, *Biometrika* 107, 791–808.
- Wilson, P. W. (2011), Asymptotic properties of some non-parametric hyperbolic efficiency estimators, in I. Van Keilegom and P. W. Wilson, eds., *Exploring Research Frontiers in Contemporary Statistics and Econometrics*, Berlin: Springer-Verlag, pp. 115–150.

- (2018), Dimension reduction in nonparametric models of production, *European Journal of Operational Research* 267, 349–367.
 - (2024), A generalized hyperbolic distance function for benchmarking performance: Estimation and inference. Unpublished working paper, Department of Economics, Clemson University, Clemson, South Carolina 29634 USA.
- Wilson, P. W. and S. Zhao (2022), Investigating the performance of Chinese banks over 2007–2014, *Annals of Operations Research* Online publication available at <https://doi.org/10.1007/s10479-022-04925-8>.

Statements and Declarations

No funds, grants, or other financial support were received during the preparation of this manuscript. The authors have no relevant financial or non-financial interests or other competing interests to disclose.

Table 1: Rejection Rates for Returns-to-Scale Test, Proposition 4.1 Rule, Original Statistics, 1,000 Splits

n	δ	Case I						Case II					
		$p = 1, q = 1$		$p = 2, q = 1$		$p = 3, q = 1$		$p = 4, q = 1$		$p = 5, q = 1$			
		.10	.05	.01	.10	.05	.01	.10	.05	.01	.10	.05	.01
100	0.0	0.011	0.004	0.004	0.066	0.048	0.044	0.141	0.128	0.141	0.007	0.005	0.003
	-0.1	0.016	0.009	0.005	0.075	0.057	0.051	0.149	0.131	0.151	0.007	0.006	0.004
	-0.2	0.062	0.033	0.027	0.145	0.122	0.088	0.194	0.180	0.196	0.010	0.006	0.004
	-0.3	0.365	0.274	0.176	0.343	0.285	0.235	0.343	0.322	0.320	0.031	0.027	0.017
	-0.4	0.837	0.766	0.618	0.684	0.619	0.530	0.602	0.570	0.550	0.099	0.078	0.056
	-0.6	0.998	0.994	0.978	0.981	0.958	0.912	0.930	0.916	0.885	0.434	0.371	0.290
	-0.8	1.000	1.000	1.000	0.998	0.998	0.994	0.996	0.992	0.990	0.799	0.731	0.646
	-1.0	1.000	1.000	1.000	1.000	1.000	0.998	1.000	1.000	0.998	0.952	0.927	0.873
	-1.2	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.991	0.981	0.958
	-1.4	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.997	0.995	0.989
1000	0.0	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	-0.1	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	-0.2	0.459	0.198	0.036	0.064	0.014	0.004	0.022	0.007	0.004	0.000	0.000	0.000
	-0.3	1.000	1.000	1.000	0.997	0.980	0.837	0.884	0.777	0.505	0.000	0.000	0.000
	-0.4	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.999	0.190	0.077	0.020
	-0.6	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.999	0.993	0.904
	-0.8	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.965
	-1.0	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.952
	-1.2	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	-1.4	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000

NOTE: The frontier is CRS for $\delta = 0$, and non-CRS for $\delta < 0$.

Table 2: Rejection Rates for Convexity Test, Proposition 4.1 Rule, Original Statistics, 1,000 Splits

n	δ	Case I						Case II					
		$p = 1, q = 1$			$p = 2, q = 1$			$p = 3, q = 1$			$p = 4, q = 1$		
		.10	.05	.01	.10	.05	.01	.10	.05	.01	.10	.05	.01
100	-1.4	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	-0.6	0.006	0.009	0.034	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	-0.4	0.032	0.038	0.080	0.000	0.000	0.001	0.000	0.000	0.000	0.000	0.000	0.000
	-0.3	0.059	0.071	0.121	0.001	0.001	0.003	0.000	0.000	0.000	0.000	0.000	0.001
	-0.2	0.108	0.128	0.212	0.002	0.004	0.007	0.000	0.000	0.000	0.000	0.000	0.005
	-0.1	0.172	0.194	0.299	0.005	0.007	0.013	0.000	0.000	0.000	0.000	0.000	0.007
	0.0	0.205	0.227	0.341	0.008	0.010	0.014	0.000	0.000	0.000	0.000	0.000	0.008
	0.1	0.211	0.248	0.371	0.015	0.022	0.022	0.001	0.001	0.006	0.000	0.001	0.007
	0.2	0.365	0.410	0.532	0.025	0.023	0.044	0.004	0.002	0.012	0.001	0.002	0.009
	0.3	0.657	0.701	0.784	0.067	0.066	0.099	0.012	0.017	0.023	0.001	0.002	0.022
	0.4	0.925	0.934	0.954	0.231	0.224	0.252	0.039	0.048	0.080	0.006	0.007	0.042
	0.6	0.998	0.997	0.998	0.740	0.752	0.751	0.311	0.313	0.359	0.053	0.067	0.170
	0.8	1.000	1.000	1.000	0.972	0.976	0.981	0.741	0.755	0.785	0.322	0.340	0.484
	1.0	1.000	1.000	1.000	0.999	0.999	0.997	0.948	0.958	0.966	0.715	0.735	0.813
	1.2	1.000	1.000	1.000	1.000	1.000	1.000	0.991	0.993	0.994	0.913	0.923	0.949
	1.4	1.000	1.000	1.000	1.000	1.000	1.000	0.999	0.998	1.000	0.974	0.976	0.980
	1.6	1.000	1.000	1.000	1.000	1.000	1.000	0.998	0.999	1.000	0.987	0.990	0.990
1000	-1.4	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	-0.6	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	-0.4	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	-0.3	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	-0.2	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	-0.1	0.001	0.000	0.001	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	0.0	0.001	0.001	0.002	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	0.1	0.006	0.004	0.007	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	0.2	0.646	0.593	0.531	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	0.3	1.000	1.000	1.000	0.046	0.015	0.012	0.000	0.000	0.000	0.000	0.000	0.000
	0.4	1.000	1.000	1.000	0.969	0.932	0.763	0.050	0.018	0.010	0.000	0.001	0.000
	0.6	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.997	0.938	0.935	0.822	0.174
	0.8	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.999	0.889
	1.0	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.999
	1.2	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	1.4	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	1.6	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000

NOTE: The production set is convex for $\delta \leq 0$, and non-convex for $\delta > 0$.

Table 3: Number of Trials among 1,000 where 10 Tests of Constant versus Non-Constant Returns to Scale using S Sample-Splits Agree, $n = 100$

δ	S	— $p = 1, q = 1$ —			— $p = 2, q = 1$ —			— $p = 3, q = 1$ —		
		.10	.05	.01	.10	.05	.01	.10	.05	.01
0.0	1	163	348	733	84	221	600	82	208	541
	5	271	615	956	122	396	865	87	320	796
	10	705	407	912	473	201	764	383	141	647
	50	880	797	638	648	537	369	547	379	220
	100	937	946	940	765	765	727	641	607	526
	200	964	965	961	838	842	780	728	702	569
	1000	988	990	995	928	927	917	876	868	779
	5000	995	997	998	971	968	966	944	936	894
	10000	994	998	999	978	981	972	960	953	928
−0.1	1	147	316	714	78	212	585	79	206	534
	5	252	574	948	108	391	861	83	313	795
	10	682	367	899	448	194	759	372	132	644
	50	847	762	593	633	510	360	537	369	215
	100	923	925	920	752	744	713	645	599	520
	200	955	953	945	826	818	779	729	706	567
	1000	984	986	989	920	917	903	870	865	771
	5000	990	997	998	965	968	969	946	927	893
	10000	995	999	998	975	977	974	958	953	927
−0.2	1	71	208	604	56	168	533	59	179	504
	5	134	431	903	70	328	818	71	267	765
	10	468	220	812	361	141	702	302	121	611
	50	655	547	413	511	395	275	457	310	179
	100	767	790	801	635	641	622	599	543	475
	200	838	871	852	740	729	676	693	654	531
	1000	940	951	956	874	883	850	844	822	746
	5000	973	980	979	948	940	937	932	918	882
	10000	983	985	989	964	961	962	947	945	913
−0.3	1	21	78	391	33	94	396	42	119	434
	5	27	200	780	34	181	721	47	192	698
	10	185	72	608	194	53	563	206	79	527
	50	289	201	170	319	215	133	339	229	125
	100	387	396	460	446	415	419	436	413	358
	200	538	529	513	556	533	476	576	520	392
	1000	776	778	746	793	777	693	794	777	634
	5000	900	891	874	893	906	854	905	897	833
	10000	924	920	903	929	935	896	931	924	895

NOTE: The frontier is CRS for $\delta = 0$, and non-CRS for $\delta < 0$.

Table 3: Number of Trials among 1,000 where 10 Tests of Constant versus Non-Constant Returns to Scale using S Sample-Splits Agree, $n = 100$ (continued)

δ	S	— $p = 1, q = 1$ —			— $p = 2, q = 1$ —			— $p = 3, q = 1$ —		
		.10	.05	.01	.10	.05	.01	.10	.05	.01
−0.4	1	7	21	185	12	52	261	20	67	300
	5	32	50	528	22	69	560	24	103	568
	10	52	36	319	75	40	368	98	40	369
	50	358	249	65	282	190	87	266	183	82
	100	515	357	186	434	329	228	409	327	224
	200	658	498	262	557	467	294	552	486	299
	1000	845	759	550	775	736	566	782	732	572
	5000	933	901	775	902	868	807	885	869	779
	10000	953	924	846	928	904	860	916	906	844
−0.6	1	68	22	33	22	12	95	18	20	130
	5	340	91	139	137	40	268	96	27	303
	10	421	388	66	170	165	120	118	121	152
	50	917	835	447	680	576	227	572	490	213
	100	969	908	553	810	682	317	702	601	271
	200	981	955	741	897	804	505	792	728	467
	1000	996	987	933	961	915	785	914	880	751
	5000	1000	994	975	978	959	910	964	951	887
	10000	1000	998	978	982	969	929	972	966	925

NOTE: The frontier is CRS for $\delta = 0$, and non-CRS for $\delta < 0$.

Table 4: Rejection Rates for Convexity Test, Asymptotic Normality, Original Statistics, 1 Split

n	δ	Case I						Case II					
		$p = 1, q = 1$			$p = 2, q = 1$			$p = 3, q = 1$			$p = 4, q = 1$		
		.10	.05	.01	.10	.05	.01	.10	.05	.01	.10	.05	.01
100	-1.4	0.119	0.075	0.025	0.095	0.054	0.014	0.113	0.064	0.015	0.112	0.069	0.023
	-0.6	0.176	0.110	0.045	0.134	0.074	0.031	0.146	0.087	0.027	0.148	0.088	0.026
	-0.4	0.203	0.127	0.056	0.150	0.092	0.035	0.158	0.106	0.037	0.165	0.103	0.033
	-0.3	0.222	0.140	0.064	0.165	0.103	0.045	0.167	0.109	0.041	0.175	0.109	0.034
	-0.2	0.241	0.161	0.075	0.175	0.111	0.047	0.179	0.114	0.041	0.182	0.115	0.041
	-0.1	0.253	0.176	0.084	0.186	0.113	0.054	0.185	0.118	0.044	0.184	0.120	0.040
	0.0	0.254	0.186	0.084	0.197	0.117	0.049	0.176	0.124	0.042	0.171	0.111	0.036
	0.1	0.259	0.195	0.088	0.203	0.120	0.052	0.176	0.127	0.044	0.171	0.112	0.037
	0.2	0.286	0.215	0.103	0.212	0.138	0.052	0.186	0.130	0.049	0.179	0.120	0.038
	0.3	0.358	0.270	0.151	0.234	0.158	0.064	0.200	0.138	0.054	0.195	0.132	0.040
	0.4	0.466	0.370	0.223	0.276	0.184	0.083	0.221	0.147	0.068	0.211	0.142	0.053
	0.6	0.663	0.585	0.419	0.363	0.269	0.128	0.278	0.191	0.093	0.239	0.177	0.069
	0.8	0.800	0.735	0.599	0.481	0.369	0.205	0.372	0.275	0.137	0.292	0.214	0.100
	1.0	0.890	0.847	0.717	0.563	0.470	0.298	0.444	0.355	0.208	0.350	0.260	0.141
	1.2	0.920	0.899	0.818	0.650	0.569	0.386	0.508	0.415	0.261	0.412	0.307	0.162
	1.4	0.946	0.926	0.882	0.713	0.631	0.449	0.560	0.463	0.313	0.457	0.347	0.199
	1.6	0.937	0.912	0.849	0.725	0.639	0.471	0.575	0.495	0.326	0.465	0.383	0.221
1000	-1.4	0.104	0.051	0.018	0.120	0.055	0.007	0.112	0.068	0.018	0.136	0.076	0.019
	-0.6	0.132	0.066	0.020	0.138	0.083	0.015	0.139	0.075	0.025	0.155	0.091	0.019
	-0.4	0.144	0.075	0.020	0.157	0.098	0.017	0.146	0.087	0.024	0.165	0.101	0.021
	-0.3	0.153	0.085	0.022	0.166	0.097	0.018	0.154	0.093	0.028	0.180	0.107	0.025
	-0.2	0.165	0.097	0.026	0.177	0.102	0.023	0.161	0.101	0.030	0.191	0.115	0.028
	-0.1	0.185	0.103	0.028	0.185	0.112	0.026	0.168	0.105	0.030	0.197	0.120	0.030
	0.0	0.186	0.126	0.032	0.202	0.117	0.031	0.176	0.103	0.035	0.193	0.119	0.031
	0.1	0.216	0.138	0.038	0.210	0.122	0.031	0.182	0.107	0.037	0.196	0.122	0.032
	0.2	0.376	0.273	0.131	0.240	0.146	0.039	0.197	0.116	0.040	0.209	0.129	0.037
	0.3	0.668	0.567	0.388	0.319	0.226	0.073	0.242	0.142	0.052	0.235	0.148	0.042
	0.4	0.922	0.878	0.759	0.461	0.352	0.176	0.309	0.206	0.079	0.279	0.189	0.059
	0.6	0.997	0.997	0.988	0.728	0.618	0.426	0.528	0.390	0.190	0.430	0.320	0.136
	0.8	1.000	1.000	0.999	0.903	0.844	0.673	0.703	0.589	0.362	0.609	0.463	0.258
	1.0	1.000	1.000	1.000	0.970	0.944	0.842	0.822	0.719	0.520	0.720	0.612	0.379
	1.2	1.000	1.000	1.000	0.983	0.971	0.928	0.899	0.820	0.634	0.812	0.706	0.498
	1.4	1.000	1.000	1.000	0.995	0.982	0.961	0.930	0.883	0.734	0.859	0.788	0.582
	1.6	1.000	1.000	1.000	0.992	0.981	0.952	0.921	0.876	0.725	0.865	0.773	0.592

NOTE: The production set is convex for $\delta \leq 0$, and non-convex for $\delta > 0$.

Table 5: Rejection Rates for Convexity Test, Simar and Wilson (2000) Bootstrap Test #1, Original Statistics, 100 Splits, 1,000 Bootstrap Replications

n	δ	Case I						Case II					
		$p = 1, q = 1$			$p = 2, q = 1$			$p = 3, q = 1$			$p = 4, q = 1$		
		.10	.05	.01	.10	.05	.01	.10	.05	.01	.10	.05	.01
100	-1.4	0.037	0.017	0.004	0.037	0.010	0.003	0.030	0.012	0.000	0.035	0.014	0.002
	-0.6	0.140	0.078	0.025	0.137	0.065	0.016	0.110	0.060	0.011	0.103	0.061	0.010
	-0.4	0.169	0.098	0.027	0.204	0.117	0.028	0.165	0.100	0.026	0.155	0.086	0.022
	-0.3	0.202	0.110	0.035	0.252	0.135	0.034	0.202	0.124	0.039	0.183	0.105	0.029
	-0.2	0.226	0.132	0.040	0.297	0.182	0.048	0.240	0.146	0.047	0.217	0.131	0.042
	-0.1	0.255	0.149	0.046	0.332	0.211	0.062	0.272	0.172	0.057	0.243	0.157	0.052
	0.0	0.259	0.158	0.042	0.332	0.229	0.078	0.320	0.212	0.069	0.246	0.153	0.046
	0.1	0.269	0.170	0.050	0.349	0.243	0.083	0.328	0.218	0.072	0.250	0.163	0.048
	0.2	0.339	0.223	0.077	0.389	0.280	0.116	0.368	0.254	0.104	0.278	0.187	0.066
	0.3	0.530	0.387	0.167	0.514	0.381	0.203	0.467	0.333	0.145	0.340	0.239	0.094
	0.4	0.784	0.667	0.437	0.729	0.590	0.352	0.616	0.499	0.262	0.475	0.339	0.163
	0.6	0.986	0.970	0.907	0.959	0.930	0.812	0.884	0.823	0.645	0.788	0.698	0.450
	0.8	1.000	1.000	0.993	0.999	0.997	0.981	0.989	0.977	0.935	0.960	0.927	0.817
	1.0	1.000	1.000	1.000	1.000	1.000	0.999	0.999	0.996	0.991	0.998	0.986	0.967
	1.2	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.997	0.999	0.999	0.993
	1.4	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.999	0.999
	1.6	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.999
1000	-1.4	0.053	0.015	0.002	0.058	0.031	0.008	0.083	0.038	0.007	0.081	0.034	0.008
	-0.6	0.126	0.064	0.012	0.161	0.081	0.026	0.204	0.109	0.034	0.198	0.123	0.024
	-0.4	0.156	0.083	0.022	0.201	0.113	0.031	0.260	0.153	0.047	0.274	0.175	0.049
	-0.3	0.171	0.095	0.022	0.231	0.128	0.037	0.300	0.192	0.064	0.320	0.209	0.071
	-0.2	0.188	0.106	0.024	0.282	0.157	0.048	0.346	0.236	0.081	0.379	0.258	0.097
	-0.1	0.204	0.122	0.027	0.332	0.201	0.068	0.405	0.281	0.108	0.428	0.305	0.122
	0.0	0.232	0.121	0.032	0.339	0.230	0.086	0.426	0.291	0.113	0.451	0.316	0.126
	0.1	0.296	0.175	0.052	0.383	0.254	0.094	0.463	0.318	0.125	0.481	0.338	0.144
	0.2	0.847	0.751	0.486	0.594	0.447	0.211	0.606	0.459	0.223	0.599	0.446	0.216
	0.3	1.000	0.999	0.997	0.965	0.931	0.786	0.903	0.843	0.639	0.849	0.765	0.523
	0.4	1.000	1.000	1.000	1.000	1.000	1.000	0.999	0.993	0.985	0.994	0.990	0.948
	0.6	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	0.8	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	1.0	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	1.2	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	1.4	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	1.6	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000

NOTE: The production set is convex for $\delta \leq 0$, and non-convex for $\delta > 0$.

Table 6: Rejection Rates for Convexity Test, Simar and Wilson (2000) Bootstrap Test #2, Original Statistics, 100 Splits 1,000 Bootstrap Replications

n	δ	Case I						Case II					
		$p = 1, q = 1$			$p = 2, q = 1$			$p = 3, q = 1$			$p = 4, q = 1$		
		.10	.05	.01	.10	.05	.01	.10	.05	.01	.10	.05	.01
100	-1.4	0.046	0.021	0.003	0.039	0.014	0.003	0.040	0.014	0.000	0.051	0.019	0.003
	-0.6	0.138	0.084	0.019	0.136	0.073	0.015	0.113	0.054	0.006	0.112	0.059	0.011
	-0.4	0.164	0.100	0.028	0.199	0.112	0.029	0.161	0.089	0.018	0.154	0.090	0.019
	-0.3	0.185	0.107	0.031	0.241	0.143	0.037	0.199	0.109	0.030	0.189	0.104	0.027
	-0.2	0.212	0.121	0.036	0.282	0.171	0.044	0.234	0.128	0.038	0.208	0.130	0.035
	-0.1	0.235	0.143	0.045	0.313	0.205	0.053	0.259	0.155	0.047	0.228	0.143	0.042
	0.0	0.251	0.157	0.040	0.323	0.211	0.077	0.291	0.176	0.062	0.227	0.135	0.042
	0.1	0.259	0.162	0.044	0.332	0.224	0.085	0.305	0.188	0.070	0.230	0.151	0.042
	0.2	0.340	0.204	0.077	0.384	0.264	0.102	0.337	0.223	0.093	0.266	0.164	0.055
	0.3	0.489	0.361	0.152	0.486	0.359	0.162	0.427	0.306	0.131	0.319	0.214	0.084
	0.4	0.747	0.615	0.374	0.682	0.538	0.316	0.565	0.447	0.233	0.433	0.305	0.122
	0.6	0.976	0.954	0.855	0.947	0.916	0.771	0.866	0.785	0.588	0.745	0.626	0.382
	0.8	0.999	0.997	0.986	0.998	0.995	0.972	0.987	0.971	0.895	0.942	0.903	0.751
	1.0	1.000	1.000	0.998	1.000	1.000	0.998	0.997	0.996	0.987	0.992	0.986	0.944
	1.2	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.999	0.997	1.000	0.998	0.985
	1.4	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.999	1.000	1.000	0.998
	1.6	1.000	1.000	0.999	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.998
1000	-1.4	0.052	0.018	0.003	0.068	0.026	0.008	0.084	0.039	0.008	0.086	0.036	0.007
	-0.6	0.128	0.064	0.013	0.151	0.079	0.022	0.180	0.109	0.019	0.181	0.105	0.023
	-0.4	0.161	0.084	0.020	0.184	0.096	0.030	0.241	0.149	0.040	0.235	0.144	0.045
	-0.3	0.175	0.094	0.022	0.209	0.113	0.037	0.273	0.181	0.053	0.289	0.178	0.063
	-0.2	0.179	0.100	0.024	0.256	0.138	0.041	0.321	0.206	0.070	0.341	0.221	0.079
	-0.1	0.193	0.110	0.029	0.299	0.174	0.050	0.364	0.235	0.091	0.390	0.253	0.100
	0.0	0.215	0.118	0.029	0.314	0.199	0.066	0.374	0.260	0.098	0.400	0.261	0.106
	0.1	0.288	0.170	0.052	0.347	0.226	0.078	0.408	0.281	0.113	0.425	0.280	0.121
	0.2	0.791	0.677	0.395	0.527	0.377	0.164	0.542	0.386	0.171	0.532	0.386	0.170
	0.3	1.000	1.000	0.993	0.945	0.879	0.682	0.845	0.740	0.497	0.799	0.681	0.408
	0.4	1.000	1.000	1.000	1.000	1.000	0.999	0.995	0.989	0.959	0.984	0.964	0.878
	0.6	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	0.8	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	1.0	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	1.2	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	1.4	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	1.6	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000

NOTE: The production set is convex for $\delta \leq 0$, and non-convex for $\delta > 0$.

Figure 1: Function $g(\cdot)$ used to Simulate Data

