Economics Letters 247 (2025) 112196

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

Testing the convexity hypothesis in nonparametric cost functions☆

Kristiaan Kerstens[a], Shirong Zhao[b],*

[a] Univ. Lille, CNRS, IESEG School of Management, UMR 9221 - LEM - Lille Économie Management, F-59000 Lille, France
[b] School of Finance, Dongbei University of Finance and Economics, Dalian, Liaoning 116025, China,

ARTICLE INFO

Dataset link: http://qed.econ.queensu.ca/jae/2011-v26.2/kumbhakar-tsionas/, https://github.com/srzhao89/kz-cost

JEL classification:
C12
D24

Keywords:
Convexity
Cost function
Convex technology
Nonconvex technology

ABSTRACT

We develop statistical tests to formally test the convexity axiom for cost functions by extending a recently developed statistical tests in nonparametric models of production. The tests are applied to a panel of US electric power generation plants and strong evidence against convexity is found for most years in the resulting cost functions. Our results suggest that empirical researchers need to be more cautious about the often implicit embedding of the convexity assumption for the cost functions.

1. Introduction

A cost function is widely used by empirical researchers as a standard tool to analyze various economic questions of interest, including economies of scale and scope, determining a Lerner index of market power, efficiency, productivity, etc. In fact, various factors may cause technology to be nonconvex (Farrell, 1959). The most important causes of nonconvexity include indivisibilities of inputs and outputs, economies of scale and scope, positive and negative externalities, etc. (see Briec et al. (2022) and Kerstens and Van de Woestyne (2021) for further discussions). While one would expect such nonconvexities to have a potential impact on economic value functions, almost all studies with parametric, semi-parametric or nonparametric specifications of cost functions just impose that technology is convex.

One main reason to test the impact of a convex technology on the cost function is a widely ignored property of the cost function in the outputs. In particular, when technology is convex, then the cost function is convex in the outputs (e.g., Shephard (1970, p. 227)). Thus, by contraposition, when the cost function is nonconvex in the outputs, then the technology is nonconvex.

Knowing whether the cost function exhibits convexity or not in the outputs has important implications for the policy makers and economists, as some researchers (e.g., Kerstens and Van de Woestyne (2021)) have documented significant differences between convex and nonconvex cost function values (between 21% and 38% on average) and an impact on economies of scale assessment. Furthermore, Kerstens and Van de Woestyne (2021, Figures 2–3) illustrate that nonconvex cost functions trace a step function in the outputs (rather than a piece-wise linear function in the convex case) supported by substantially more change points. An earlier study of Balaguer-Coll et al. (2007) also reports a staggering 41.13% difference between convex and nonconvex cost functions at the sample level. However, to the best of our knowledge, in the literature there is so far no statistical test available that can be used to test the convexity of the cost function in the outputs. Thus, no cost function study known to us has ever formally tested whether this convexity assumption can be maintained.

Leveraging on the recent developments of the statistical tests on the convexity for the production technologies (see Kneip et al. (2016), Simar and Wilson (2020a) and Kneip et al. (2022)), this paper proposes a corresponding statistical test to test for the convexity of the cost function in a nonparametric framework (see Ray (2022) for a recent review on nonparametric methods). In particular, it tests whether the cost function is convex or nonconvex in the outputs. To our knowledge, this study represents the first attempt to formally test the convexity assumption for the case of the cost function. Furthermore, we illustrate our statistical methods using panel data from US electric power generation plants covering the years 1986–1998 compiled by Kumbhakar and

☆ We thank a referee and an associate editor of the journal for the most helpful comments we have received. The usual disclaimer applies.
* Corresponding author.
E-mail addresses: k.kerstens@ieseg.fr (K. Kerstens), shironz@163.com (S. Zhao).

https://doi.org/10.1016/j.econlet.2025.112196
Received 14 September 2024; Received in revised form 6 December 2024; Accepted 15 January 2025
Available online 22 January 2025
0165-1765/© 2025 Elsevier B.V. All rights reserved, including those for text and data mining, AI training, and similar technologies.

Tsionas (2011). Anticipating our results, the testing results find strong evidence against convexity of the cost function for most of the years. Hence, empirical researchers should be much more cautious about this often implicitly embedded convexity assumption when estimating cost functions.

2. Cost function

We first introduce the definitions of technology and cost function. Denoting a vector of p input quantities by $x \in \mathbb{R}_+^p$ and a vector of q output quantities by $y \in \mathbb{R}_+^q$, the production possibility set or technology can be defined as:

$$\Psi := \{(x, y) \in (\mathbb{R}_+^p, \mathbb{R}_+^q) \mid x \text{ can produce at least } y\}. \tag{1}$$

The widely used input distance function (Farrell, 1957) for one evaluated observation (x_0, y_0) is given by:

$$D(x_0, y_0 \mid \Psi) := \inf \left\{\theta \mid (\theta x_0, y_0) \in \Psi\right\}. \tag{2}$$

Switching to a dual representation of technology, the cost function is defined as the minimum expenditure required to produce an output vector y_0, given the production possibility set Ψ and a vector of non-negative input prices $w_0 \in \mathbb{R}_+^p$:

$$C(y_0 \mid w_0, \Psi) := \min \left\{w_0' x : (x, y_0) \in \Psi\right\}. \tag{3}$$

We can then define the cost ratio as a ratio of the minimum cost over the observed cost:

$$CR(x_0, y_0 \mid w_0, \Psi) := \frac{C(y_0 \mid w_0, \Psi)}{c_0}. \tag{4}$$

where $c_0 = w_0' x_0$ is the observed cost. Lemma 3.2 of Simar and Wilson (2020b) shows that:

$$CR(x_0, y_0 \mid w_0, \Psi) = D(c_0, y_0 \mid \Psi_{w_0}), \tag{5}$$

where Ψ_{w_0} is the image of Ψ under $h_{w_0}(x, y) = A_{w_0}[x' \ y']'$, and where:

$$A_{w_0} = \begin{pmatrix} w_0' & 0_q' \\ 0_{p\times q}' & I_q \end{pmatrix}, \tag{6}$$

0_q is a $(q \times 1)$ vector of zeros, $0_{p\times q}$ is a $p \times q$ matrix of zeros, and I_q is a $(q \times q)$ identity matrix. To be more specific, the relation between Ψ_{w_0} and Ψ can be expressed as:

$$\Psi_{w_0} = \{(c, y) \mid (c, y) = h_{w_0}(x, y), \ \forall\, (x, y) \in \Psi\}. \tag{7}$$

Moreover, as h_{w_0} is an affine function, Ψ_{w_0} preserves the convexity if and only if Ψ is convex.

Given a random sample $S_n = \{X_i, Y_i, W_i\}_{i=1}^n$, if the cost function is convex in the outputs, then the cost ratio $CR(x_0, y_0 \mid w_0, \Psi)$ can be estimated using a convex variable returns to scale nonparametric specification of technology as follows[1]:

$$\widehat{\mathrm{CR}}_{\mathrm{C}}(x_0, y_0) = \min_{\theta, s_1, \ldots, s_n} \left\{ \begin{aligned} &\theta \mid y_0 \le \sum_{i=1}^n s_i Y_i, \\ &\theta w_0' x_0 \ge \sum_{i=1}^n s_i w_0' X_i, \ \theta \ge 0, \\ &\sum_{i=1}^n s_i = 1, \ \forall\, s_i \ge 0 \end{aligned} \right\}. \tag{8}$$

If the cost function is nonconvex in the outputs, then it can be estimated using a nonconvex variable returns to scale nonparametric specification of technology as follows[2]:

$$\widehat{\mathrm{CR}}_{\mathrm{NC}}(x_0, y_0) = \min_{\theta, s_1, \ldots, s_n} \left\{ \begin{aligned} &\theta \mid y_0 \le \sum_{i=1}^n s_i Y_i, \\ &\theta w_0' x_0 \ge \sum_{i=1}^n s_i w_0' X_i, \ \theta \ge 0, \\ &\sum_{i=1}^n s_i = 1, \ \forall\, s_i \in \{0, 1\} \end{aligned} \right\}. \tag{9}$$

3. Convexity of the cost function in the outputs: Statistical tests

Starting from Kneip et al. (2016) and Simar and Wilson (2020a), we present the details of extending the convexity test of the production technology into the convexity test of the cost function. Simar and Wilson (2020b) show that under convexity the cost ratio estimated using the convex technology (8) converges at the rate n^{κ_1}, while under nonconvexity the cost ratio estimated using the nonconvex technology (9) converges at the rate n^{κ_2}, where $\kappa_1 = 2/(q+2)$ and $\kappa_2 = 1/(q+1)$.

The null hypothesis is that the cost function is convex in the outputs, while the alternative hypothesis is that it is nonconvex in the outputs. Note that the test statistic proposed by Kneip et al. (2016) involves two independent samples, while researchers typically just have one sample. To be valid, the data are first randomly shuffled and then divided unevenly into two sub-samples with more observations used for the nonconvex estimator due to its slow convergence rate. One sub-sample (n_1) is used for computing convex cost ratios, while the other sub-sample (n_2) is utilized for evaluating nonconvex cost ratios, where $n_1 + n_2 = n$ and $n_1^{\kappa_1} = n_2^{\kappa_2}$. Analogous to (50) in Kneip et al. (2016) for the production technology, under the null hypothesis of convexity of the cost function and provided $q \le 2$, we have:

$$\hat{\tau}_{1,n} = \frac{(\hat{\mu}_{\mathrm{NC},w_x,n_2} - \hat{\mu}_{\mathrm{C},w_x,n_1}) - (\hat{B}_{\mathrm{NC},\kappa_2,n_2} - \hat{B}_{\mathrm{C},\kappa_1,n_1})}{\sqrt{\frac{\hat{\sigma}^2_{\mathrm{NC},n_2}}{n_2} + \frac{\hat{\sigma}^2_{\mathrm{C},n_1}}{n_1}}} \xrightarrow{\mathcal{L}} N(0,1), \tag{10}$$

where $\hat{\mu}_{\mathrm{NC},w_x,n_2}$ and $\hat{\mu}_{\mathrm{C},w_x,n_1}$ are the mean values of the nonconvex and convex cost ratio estimates in the two-subsamples (with sample sizes n_2 and n_1), respectively; $\hat{\sigma}^2_{\mathrm{NC},n_2}$ and $\hat{\sigma}^2_{\mathrm{C},n_1}$ are the corresponding variance estimates; $\hat{B}_{\mathrm{NC},\kappa_2,n_2}$ and $\hat{B}_{\mathrm{C},\kappa_1,n_1}$ are the bias terms estimated using generalized jackknife bootstrap methods discussed in Kneip et al. (2015) and Simar and Wilson (2020b).

If $q \ge 3$, we must use subsets of the subsamples to compute the sample means. Let $\kappa = \kappa_2$, $n_{1,\kappa} = \lfloor n_1^{2\kappa} \rfloor \le n_1$, $n_{2,\kappa} = \lfloor n_2^{2\kappa} \rfloor \le n_2$, where $\lfloor a \rfloor$ is the largest integer that is no more than a. Moreover, let $\hat{\mu}_{\mathrm{NC},w_x,n_{2,\kappa}}$ and $\hat{\mu}_{\mathrm{C},w_x,n_{1,\kappa}}$ be the mean values of the nonconvex and convex cost ratios estimates, with the sizes as $n_{2,\kappa}$ and $n_{1,\kappa}$, randomly selected from the two-subsamples, respectively. Analogous to (52) in Kneip et al. (2016) for the production technology, under the null hypothesis of convexity for the cost function, we have:

$$\hat{\tau}_{2,n} = \frac{(\hat{\mu}_{\mathrm{NC},w_x,n_{2,\kappa}} - \hat{\mu}_{\mathrm{C},w_x,n_{1,\kappa}}) - (\hat{B}_{\mathrm{NC},\kappa_2,n_2} - \hat{B}_{\mathrm{C},\kappa_1,n_1})}{\sqrt{\frac{\hat{\sigma}^2_{\mathrm{NC},n_2}}{n_{2,\kappa}} + \frac{\hat{\sigma}^2_{\mathrm{C},n_1}}{n_{1,\kappa}}}} \xrightarrow{\mathcal{L}} N(0,1). \tag{11}$$

Note that these two tests are one-sided, and rejection of the null hypothesis can be indicated by "large" values of the test statistics $\hat{\tau}_{1,n}$ and $\hat{\tau}_{2,n}$. Furthermore, the tests involve randomly splitting the sample just for a single time. Obviously, different randomly split samples risk yielding different statistics and p-values. To overcome this shortage, Simar and Wilson (2020a, p. 302) suggest repeating the random sample-split for many times (at least 10 times) and recording the test statistics and the corresponding p-values. Under the null hypothesis of convexity,

[1] In the operations research literature often the name data envelopment analysis (DEA) method is employed to designate this type of technologies.

[2] In the operations research literature, nonconvex technologies are sometimes designated as free disposal hull (FDH) methods.

the average of the test statistics is close to 0 and the distribution of these p-values is close to the uniform distribution. Therefore, we have two tests in total, denoted as Test#1 and Test#2. Test#1 relies on the average of the test statistics while Test#2 relies on a Kolmogorov–Smirnov test of uniformity of the distribution of p-values obtained from each test statistic. Moreover, Simar and Wilson (2020a) suggest using the bootstrap methods to make inference, as these random sample-splits are not independent.

According to Simar and Wilson (2020b), when there is more than one input the nonparametric estimators of cost efficiency are shown to have faster rates of convergence than the corresponding estimators of technical efficiency. In addition, the convergence rates for nonparametric estimators of cost efficiency depend solely on q, whereas those for technical efficiency depend on both p and q. These distinctions extend to the differences between the proposed tests for cost functions and the existing tests for production technologies: the proposed tests for cost functions rely only on q, while existing production technology tests require both p and q. However, it should be noted that the proposed tests for cost functions necessitate data on both input and output quantities as well as on input prices, while the existing tests for production technologies require only input and output quantity information.

4. Empirical illustration

We apply our tests to the publicly available data set of privately investor-owned US electric power generation plants compiled by Kumbhakar and Tsionas (2011). The sample contains 1065 plant-year observations from 1986 to 1998.[3] Plants are assumed to use labor (X_1), fuel (X_2), and capital (X_3) to generate the single output electricity (Y). Output is measured as the net steam electric power generation in megawatt-hours: it is the amount of power produced using fossil-fuel fired boilers to produce steam for turbine generators during a given period of time. The corresponding prices for the inputs are also available in this data set. More details on inputs, output and input prices as well as descriptive statistics are found in Kumbhakar and Tsionas (2011, p. 278–279). For each year, we have 82 observations except for the last year in which there are only 81 observations.

From the engineering literature, we know that the economic dispatch problem seeks an optimal power schedule giving the minimum fuel cost for a set of online thermal units while matching the demands and a variety of constraints (including valve loading effects, transmission losses, multiple fuels, etc.). This is a complex non-smooth and nonconvex optimization problem (see Yang et al. (2013)). The resulting cost function is also non-smooth and nonconvex (see Alawode et al. (2018)). In the economics literature, electricity generation is almost universally modeled using convex cost functions. An exception is the study by Kerstens and Van de Woestyne (2021) who report significant differences between convex and nonconvex cost estimates for Chilean hydro-power plants.

Repeating the random sample-split 100 times, the results in Table 1 illustrate that for each year the mean cost ratio estimates of $\hat{\mu}_{\mathrm{NC},w_x}$ are much larger than those of $\hat{\mu}_{\mathrm{C},w_x}$: this indicates that the cost function is likely nonconvex in this single output. Indeed, the two tests of convexity extended from Kneip et al. (2016) and Simar and Wilson (2020a) show that the p-values are smaller than 10% for 16 out of 26 cases in total. More specifically, there are 12 cases in which the p-values are smaller than 5%. Moreover, over a span of 13 years, both tests reject convexity in 8 years, and neither test rejects convexity in the remaining 5 years.

Given the general reasons for nonconvexities in economics, the arguments from engineering in favor of nonconvexities in electricity generation, and the fact that it is implausible that technologies switch status in terms of convexity over time, we find strong evidence that these cost functions for US electric utilities are nonconvex in the outputs. Once ample evidence is found against convexity, researchers should stop using convex estimators and instead use nonconvex estimators. Indeed, nonconvex estimators are always consistent and can approximate convex estimators in case of a true convex world, while convex estimators are only consistent under convexity and remain biased in case of a true nonconvex world (see, e.g., Kneip et al. (2016)).

Table 1
Mean cost ratio estimates under convex and nonconvex technologies: Tests of convexity using uneven splits with 100 sample-splits.

Year	$\hat{\mu}_{\mathrm{C},w_x}$	$\hat{\mu}_{\mathrm{NC},w_x}$	Test#1		Test#2	
			Statistics	p-values	Statistics	p-values
1986	0.726	0.860	2.041	0.151	0.627	0.232
1987	0.739	0.871	2.055	0.277	0.658	0.249
1988	0.711	0.866	2.906	0.006	0.752	0.016
1989	0.720	0.872	2.967	0.010	0.787	0.020
1990	0.695	0.844	1.974	0.024	0.680	0.010
1991	0.721	0.849	1.928	0.017	0.563	0.064
1992	0.725	0.844	1.747	0.117	0.559	0.174
1993	0.710	0.852	2.283	0.000	0.665	0.003
1994	0.708	0.839	1.806	0.008	0.535	0.052
1995	0.725	0.846	1.756	0.013	0.536	0.058
1996	0.725	0.829	1.256	0.186	0.410	0.334
1997	0.702	0.813	0.983	0.375	0.350	0.475
1998	0.717	0.846	1.709	0.013	0.531	0.055

Note: In each year, $\hat{\mu}_{\mathrm{C},w_x}$ and $\hat{\mu}_{\mathrm{NC},w_x}$ are computed using all the observations in that year. We use 100 sample-splits to compute these two statistics, 1000 bootstraps to compute the sampling distribution of these two tests under the null, and 100 bootstraps to compute the bias term. For more details about these two tests, see Simar and Wilson (2020a).

5. Conclusions

Convexity of the cost function is implicitly assumed by almost all the researchers when studying various economic questions of interest. However, whether the cost function is convex in the outputs or not is ultimately an empirical question that needs to be tested using some econometric tools, which are absent in the current literature. Building on the recently developed statistical tests on the convexity for the production technology (Kneip et al., 2016; Simar and Wilson, 2020a; Kneip et al., 2022), this paper proposes the corresponding statistical tests on the convexity for the cost function. Consequently, we provide the tools for empirical researchers that can be used in future tests of convexity for the cost function in other sectors.

Acknowledgments

Shirong Zhao acknowledges the support from Liaoning Social Science Foundation (L22CJY011).

Data availability

Data are available on http://qed.econ.queensu.ca/jae/2011-v26.2/kumbhakar-tsionas/. Code is available on https://github.com/srzhao89/kz-cost.

References

Alawode, K., Jubril, A., Kehinde, L., Ogunbona, P., 2018. Semidefinite programming solution of economic dispatch problem with non-smooth, non-convex cost functions. Electr. Power Syst. Res. 164, 178–187.
Balaguer-Coll, M.T., Prior, D., Tortosa-Ausina, E., 2007. On the determinants of local government performance: A two-stage nonparametric approach. Eur. Econ. Rev. 51 (2), 425–451.

[3] Kumbhakar and Tsionas (2011) mention a larger data set: the article also contains the year 1999. The data on the Journal of Applied Econometrics data archive is more limited: see http://qed.econ.queensu.ca/jae/2011-v26.2/kumbhakar-tsionas/.

Briec, W., Kerstens, K., Van de Woestyne, I., 2022. Nonconvexity in production and cost functions: An exploratory and selective review. In: Ray, S., Chambers, R., Kumbhakar, S. (Eds.), Handbook of Production Economics. Vol. 2, Springer, Singapore, pp. 721–754.

Farrell, M.J., 1957. The measurement of productive efficiency. J. Roy. Statist. Soc. Ser. A 120, 253–281.

Farrell, M.J., 1959. The convexity assumption in the theory of competitive markets. J. Polit. Econ. 67 (4), 377–391.

Kerstens, K., Van de Woestyne, I., 2021. Cost functions are nonconvex in the outputs when the technology is nonconvex: Convexification is not harmless. Ann. Oper. Res. 305 (1), 81–106.

Kneip, A., Simar, L., Wilson, P.W., 2015. When bias kills the variance: Central limit theorems for DEA and FDH efficiency scores. Econometric Theory 31, 394–422.

Kneip, A., Simar, L., Wilson, P.W., 2016. Testing hypotheses in nonparametric models of production. J. Bus. Econom. Statist. 34, 435–456.

Kneip, A., Simar, L., Wilson, P.W., 2022. Conical FDH Estimators of General Technologies, with Applications to Returns to Scale and Malmquist Productivity Indices. LIDAM Discussion Paper ISBA 2022/24, Université catholique de Louvain, Institute of Statistics, Biostatistics and Actuarial Sciences.

Kumbhakar, S.C., Tsionas, E.G., 2011. Stochastic error specification in primal and dual production systems. J. Appl. Econometrics 26 (2), 270–297.

Ray, S., 2022. Data Envelopment Analysis: A nonparametric method of production analysis. In: Ray, S., Chambers, R., Kumbhakar, S. (Eds.), Handbook of Production Economics. Vol. 1, Springer, Singapore, pp. 409–470.

Shephard, R.W., 1970. Theory of Cost and Production Functions. Princeton University Press, Princeton.

Simar, L., Wilson, P.W., 2020a. Hypothesis testing in nonparametric models of production using multiple sample splits. J. Prod. Anal. 53, 287–303.

Simar, L., Wilson, P.W., 2020b. Technical, allocative and overall efficiency: Estimation and inference. European J. Oper. Res. 282 (3), 1164–1176.

Yang, L., Fraga, E.S., Papageorgiou, L.G., 2013. Mathematical programming formulations for non-smooth and non-convex electricity dispatch problems. Electr. Power Syst. Res. 95, 302–308.