Bank i Kredyt 46(4), 2015, 299-326

BANK I KREDYT
BANK & CREDIT

The shape of aggregate production functions: evidence from estimates of the World Technology Frontier

Jakub Growiec*, Anna Pajor#, Dorota Górniak‡, Artur Prędki§

Submitted: 16 February 2015. Accepted: 28 April 2015.

Abstract

The article estimates the aggregate production function at the World Technology Frontier on the basis of annual data on inputs and output in 19 highly developed OECD countries in 1970–2004. A comparison of results based on Data Envelopment Analysis and Bayesian Stochastic Frontier Analysis uncovers a number of significant discrepancies between nonparametric estimates of the frontier and parametric (Cobb-Douglas and translog) aggregate production functions in terms of implied technical efficiency levels, partial elasticities, returns to scale, and elasticities of substitution.

Keywords: World Technology Frontier, aggregate production function, partial elasticity, returns to scale, substitutability

JEL: E23, O11, O14, O33, O47

* Narodowy Bank Polski, Economic Institute; Warsaw School of Economics, Department of Quantitative Economics; e-mail: jakub.growiec@nbp.pl.

Cracow University of Economics, Department of Econometrics and Operations Research; e-mail: pajora@uek.krakow.pl.

‡ Alior Bank SA; e-mail: dorota.gorniak@alior.pl.

§ Cracow University of Economics, Department of Econometrics and Operations Research; e-mail: predkia@uek.krakow.pl.

1. Introduction

The aggregate production function, mapping the stocks of aggregate inputs onto the aggregate output, is frequently assumed to take simple forms such as the constant-returns-to-scale (CRS) Cobb-Douglas function. Such postulates can be found not only in the theoretical literature – models of long-run growth, business cycles, etc. – but also in empirical studies. The latter category includes, among others, the growing development accounting literature (see e.g. Hall, Jones 1999; Caselli 2005). The obvious advantage of making such simplifying assumptions lies with their analytical tractability. The CRS Cobb-Douglas function is also in partial agreement with a few broad patterns observed in macro data, such as Kaldor's stylized facts. On the other hand, substantial empirical discrepancies in several other areas have been documented, implying that such overly restrictive functional requirements should in fact be relaxed. In particular, factor shares have been documented to exhibit both short-run fluctuations and prolonged trends (e.g. McAdam, Willman 2013; Karabarbounis, Neiman 2014), whereas elasticities of substitution have been shown to be generally non-unitary (e.g. Duffy, Papageorgiou 2000; Caselli, Coleman 2006; Klump, McAdam, Willman 2007; Chirinko 2008; León-Ledesma, McAdam, Willman 2010).

From a positive viewpoint, there is however no consensus on the preferred functional form of the aggregate (country-level) production function. The principal reason is that estimation of aggregate production functions is notoriously difficult due to multiple empirical issues: measurement uncertainty of input and output aggregates such as GDP, physical capital and human capital, problems with comparability across countries and time, endogeneity of input variables, just to name a few. Yet another important issue, and one that we carefully address in the current paper, is that even though the aggregate production function is a technological concept, one of a technical relationship between inputs and outputs, in reality, country-level productivity may also be affected by non-technological variables such as taxation, presence of various barriers to doing business (corruption, crime, complicated bureaucratic procedures, etc.), sectoral composition of production, labour market institutions, or financial constraints. Hence, to obtain reliable estimates of the technological production function itself, one ought to control for differences in these institutional conditions across countries and time. We tackle this issue by taking the World Technology Frontier approach.

The objective of the current paper is to estimate the aggregate, country-level production function as a relationship between countries' aggregate inputs and their maximum attainable output, computed on the basis of the World Technology Frontier (WTF hereafter) – where the WTF is the best-practice frontier at each moment in time. By doing so, we are able to single out technological aspects of the production processes from their institutional background, at least up to a multiplicative constant. Such estimates of the aggregate production function will be then used as a convenient starting point for further analyses, aimed at deriving this function's crucial characteristics, and discussing which parametric form agrees most with the available empirical evidence. As crucial features of the estimated aggregate production function, we shall investigate its implications for the cross-country distribution of technical inefficiency, the pattern of dependence of its (variable) partial elasticities on factor endowments, (variable) returns-to-scale properties, and its implied (Morishima and Allen-Uzawa) elasticities of substitution.

We estimate the aggregate production function with two alternative methods. First, we apply the nonparametric Data Envelopment Analysis (DEA) approach (for applications in macroeconomics, see e.g. Färe et al. 1994; Kumar, Russell 2002; Henderson, Russell 2005; Jerzmanowski 2007; Badunenko, Henderson, Zelenyuk 2008; Growiec 2012b), augmented with a bootstrap procedure which enables us to adjust for the bias of DEA efficiency estimates as well as to compute standard errors and confidence intervals for these estimates. The key advantage of this first approach is that it does not require one to make any *a priori* assumptions on the functional form of the aggregate production function – and yields testable predictions instead. Unfortunately, since the DEA approach is based on piecewise linear approximations of the true aggregate production function, it is not suited to providing predictions on the function's curvature features such as the elasticities of substitution.

Second, we also apply the Stochastic Frontier Analysis (SFA) methodology (for applications in macroeconomics, see e.g. Koop, Osiewalski, Steel 1999; 2000 and Bos et al. 2010; see Kumbhakar, Lovell 2000 for general reference) which allows us to estimate the production function directly, under certain predefined (parametric) functional specifications. Such parametric models are estimated with Bayesian techniques, particularly well-suited to production function estimation due to their relative robustness under collinearity and measurement error. The advantage of the SFA approach is that it allows to test several parametric specifications directly. It is also useful for drawing precise conclusions on the aggregate production function's elasticities of substitution. SFA allows us to estimate the production function directly, but even if the estimated parametric function is misspecified when taken at face value, sometimes it can still be considered as a reasonable approximation of the true aggregate production function, sufficiently good within some range of input combinations. The translog production function is indeed frequently viewed this way, i.e. as a local second-order Taylor approximation of an arbitrary function.

The current study provides a systematic assessment of the shape of aggregate, country-level production functions based on a comparison of nonparametric and parametric estimates of the WTF. Our baseline interpretation of identified discrepancies between the DEA and SFA is that (bootstrap--augmented) DEA estimates of the frontier may be imprecise due to their potentially low efficiency but should, in principle, capture the true characteristics of the production process without systematic biases; its parametric SFA approximation, on the other hand, is more robust to random noise and the impact of nuisance variables because of being based on a more parsimonious model, but may be fundamentally misspecified. The importance of carrying out such a study follows from the fact that even though investigations based on different *a priori* assumptions on the shape of the aggregate production tend to produce conflicting implications, there is still a clear disconnect between the Cobb-Douglas-, CES-, DEA-, and SFA-based literatures nevertheless. The current paper, along with the monograph by Growiec (2012a) which reproduces some of its results, is thus among the first attempts to bridge this gap.

An earlier draft of the current article, under the same title, has been circulated in 2011 as NBP Working Paper No. 102 (Growiec et al. 2011). Some of the results of (the earlier, significantly longer version of) the current article have been reproduced in the monograph by Growiec (2012a). In particular, the current Table 1 has been reproduced as Table 5.7, Table 2 – split into Tables 5.10 and 5.12, Table 3 – split into Tables 5.11 and 5.13, Table 4 – reproduced as Table 5.16, and Table 5 – as Table 5.17. Given this overlap, it is critical to note that the current article contains original research results whereas the quoted monograph only reproduces them, framing them in a different, somewhat broader context. All this has been acknowledged in Growiec (2012a). As compared to the quoted NBP Working Paper No. 102, the current revision features a streamlined, abridged presentation of the methodology and results as well as contains a range of new interpretations and additional literature references.

The remainder of the paper is structured as follows. Section 2 presents the dataset and reviews the alternative methodologies. Section 3 discusses the properties of the aggregate, country-level production function, inferred from our DEA- and SFA-based estimates of the WTF. Section 4 concludes. For qualifications and extensions, please consult the working paper version of this article, Growiec et al. (2011).

2. Data and methodology

2.1. Data sources and the construction of variables

The macroeconomic dataset used in the current study covers 19 highly developed OECD economies in the period 1970–2004. The output variable Y_{it} (henceforth, i indexes countries and t denotes the year) is GDP and the input variables are the aggregate stocks of physical capital, human capital, subdivided into unskilled and skilled labour, and (for auxiliary purposes only) the "raw" number of employees.

International, annual data on GDP and GDP per worker as well as the total number of workers in 1970–2004 have been taken from the Total Economy Database, developed by the Conference Board and Groningen Growth and Development Data Centre (GGDC). The unit of measurement is the US dollar, converted to constant prices as of year 2008 using updated 2005 EKS PPPs.[1]

Physical capital stocks K_{it} have been constructed using the perpetual inventory method (cf. Caselli 2005). We have used country-level investment shares from the Penn World Table 6.2 (cf. Heston, Summers, Aten 2006). Following Caselli (2005), we also assumed an annual capital depreciation rate of 6%.

Country-level human capital data have been taken from de la Fuente and Doménech (2006). The raw variables provided in this contribution are shares of population aged 25 or above having completed primary, some secondary, secondary, some tertiary, tertiary, or postgraduate education. The considered dataset is of 5-year frequency only and ends in 1995. Nevertheless, the de la Fuente--Doménech dataset has been given priority among all possible education attainment databases due to its presumed superior quality. The original de la Fuente-Doménech data have then been extrapolated forward in the time-series dimension until the year 2000 using Cohen and Soto (2007) schooling data as a predictor for the trends. Furthermore, the human capital data have been interpolated to all intermediate years as well, for human capital variables are, in general, very persistent and not susceptible to business cycle variations.

Human capital aggregates have been constructed from these educational attainment data using the Mincerian exponential formula with a concave exponent, following Hall and Jones (1999), and more directly, Caselli (2005) and Growiec (2012b):

$$H_{it}^{U} = \left(\sum_{j \in S_U} \psi_{j,it} \, e^{\phi(s_{j,it})} \right) L_{it}, \qquad H_{it}^{S} = \left(\sum_{j \in S_S} \psi_{j,it} \, e^{\phi(s_{j,it})} \right) L_{it} \tag{1}$$

[1] The Conference Board and Groningen Growth and Development Centre, Total Economy Database, January 2009 http://www.conference-board.org/data/economydatabase/.

where S_U is the set of groups of people who completed less than 12 years of education (less than elementary, elementary, less than secondary), S_S is the set of groups of people who completed 12 years of education or more (secondary, less than college, college or more), ψ_j captures the share of j-th education group in total working-age population of the given country, s_j represents years of schooling in j-th education group (cf. de la Fuente, Doménech 2006), L_{it} is the total number of workers, and $\phi(s)$ is a concave piecewise linear function:

$$\phi(s) = \begin{cases} 0.134s & s < 4 \\ 0.134 \cdot 4 + 0.101(s-4) & s \in [4, 8) \\ 0.134 \cdot 4 + 0.101 \cdot 4 + 0.068(s-8) & s \geq 8 \end{cases} \tag{2}$$

The overall human capital stock may be computed as the sum of unskilled and skilled labour: $H_{it} = H_{it}^U + H_{it}^S$. The cutoff point of 12 years of schooling, delineating unskilled and skilled labour, seems adequate for the relatively highly developed OECD economies in our sample, though it might be set too high if developed economies were to be considered as well (cf. Caselli, Coleman 2006). We have, however, allowed these two types of labour to be imperfectly substitutable and thus enter the production function separately. The "perfect substitution" case where only total human capital matters for production (and its distribution between unskilled and skilled labour has no impact whatsoever) is an interesting special case of our generalized formulation. The data do not support this assumption.

Should significant outliers be found within our sample, indicating insufficient data quality, the final results are likely to be biased. Additonally, outliers could also indicate an underlying problem of discrepancies bewteen the data-generating process and the estimated specification. Our approach tackles this possibility as well. The same problem could also appear due to business-cycle fluctuations, especially that we only measure the total stocks of physical and human capital in the considered countries, without taking account of their utilization rates which vary significantly across the cycle. Avoiding the impact of short- and medium-term disturbances appears extremely important in an aggregate production function analysis such as ours. Thus, the Hodrick and Prescott (1997) filter with the usual smoothing parameter (λ = 6.25 for annual data) has been applied to all our data to exclude the outliers and high-frequency cyclical variation present in the data.

Unfortunately, when employing the aforementioned panel dataset in parametric analyses such as the SFA, we face a critical problem. Namely, due to the strong multi-collinearity present in the time domain of our smoothed time series, the parametric, Bayesian estimation procedures applied here might lead to uninformative, uninterpretable results. To avoid this unwelcome outcome, we have decided to narrow down the time dimension of the dataset used in our SFA estimations, limiting ourselves to data covering entire decades instead of single years. The presentation of our results takes into account the fact that our DEA results have been obtained for the whole dataset and the SFA results for its subset only. We concentrate on cross-sectional comparisons or on the inferred "time--less" characteristics such as the slope and curvature of the aggregate production function, and do not compare goodness-of-fit statistics if they are computed on the basis of different datasets.

2.2. Methodology

The objective of the current paper is to draw conclusions on the shape of the aggregate, country-level production function, based on two types of estimates of the World Technology Frontier: deterministic DEA-based ones, augmented with the stochastic, nonparametric Simar-Wilson bootstrap, and parametric SFA-based ones, computed using Bayesian procedures.

Data Envelopment Analysis

The idea behind DEA is to construct the best-practice production function as a convex hull of production techniques (input-output configurations) used in countries present in the data, assuming that there is full factor utilization and no measurement error. The production function is then inferred indirectly as a fragment of the boundary of this convex hull for which output is maximized given inputs. More precisely, for each observation $i = 1, 2, ..., I$ and $t = 1, 2, ..., T$, output y_{it} is decomposed as:

$$y_{it} = E_{it} f_t(x_{it}) \tag{3}$$

i.e. into a product of the maximum attainable output given inputs $y^*_{it} \equiv f_t(x_{it})$ and the "efficiency index", i.e. the Shephard distance function $E_{it} \in (0, 1]$. The efficiency index E_{it} measures (vertical) distance to the technology frontier, while the frontier itself is computed nonparametrically as $y^*_{it} = f_t(x_{it})$.

Unless indicated otherwise, the vector of inputs in DEA will be $x_{it} = (K_{it}, H^U_{it}, H^S_{it})$. Formally, the (output-based) deterministic DEA method is a linear programming technique allowing one to find the Shephard distance function E_{jt} for each unit $j = 1, 2, ..., I$ and given $t \in \{1, 2, ..., T\}$ in the sample such that its reciprocal – the Debreu-Farrell efficiency index θ_{jt} – is maximized subject to a series of feasibility constraints (cf. Fried, Lovell, Schmidt 1993):

$$\begin{aligned}
&\max_{\{\theta_{jt}, \lambda_{11}, \ldots, \lambda_{It}\}} \theta_{jt} \\
&\text{s.t. } \theta_{jt}\, y_{jt} \le \sum_{\tau=1}^{t}\sum_{i=1}^{I} \lambda_{i\tau} y_{i\tau} \\
&\sum_{\tau=1}^{t}\sum_{i=1}^{I} \lambda_{i\tau} x_{1,i\tau} \le x_{1,jt} \\
&\sum_{\tau=1}^{t}\sum_{i=1}^{I} \lambda_{i\tau} x_{2,i\tau} \le x_{2,jt} \\
&\quad\vdots \\
&\sum_{\tau=1}^{t}\sum_{i=1}^{I} \lambda_{i\tau} x_{n,i\tau} \le x_{n,jt} \\
&\lambda_{i\tau} \ge 0, \qquad i = 1, 2, \ldots, I,\ \tau = 1, 2, \ldots, t
\end{aligned} \tag{4}$$

where $x_{k,it}$ captures the amount of k-th input available in country i at time t.

It is also additionally assumed that $\sum_{\tau=1}^{t}\sum_{i=1}^{I}\lambda_{i\tau}=1$ in the VRS case (variable returns to scale), or $\sum_{\tau=1}^{t}\sum_{i=1}^{I}\lambda_{i\tau}\leq 1$ in the NIRS case (non-increasing returns to scale). Under the CRS (constant returns to scale) assumption, no further restriction on $\lambda_{i\tau}$'s is necessary.

The Shephard distance function E_{jt} is computed as the reciprocal of the (output-oriented Debreu--Farrell) efficiency index θ_{jt} (that is, $E_{jt}=1/\theta_{jt}$).

Since the data contain a finite number of data points, one for each country and each year, by construction the DEA-based production function is piecewise linear and its vertices are the actually observed efficient input-output configurations (i.e. not dominated by any linear combination of other observed input-output configurations).

As a rule, the WTF is estimated sequentially, so that for computing the WTF in each period t, data from periods $\tau=1,2,...,t$ are used. This corresponds to the assumption that technologies, once developed, remain available for use forever (see e.g. Henderson, Russell 2005).

Advantages and limitations of the deterministic DEA approach

The deterministic DEA is a data-driven approach to deriving the production function from observed input-output pairs. Its unquestionable strength lies in the fact that it does not require any assumptions on the functional form of the aggregate production function (provided that it satisfies the free-disposal property), and provides testable predictions on its shape instead. Indeed, the usual assumption of a Cobb-Douglas aggregate production function may lead to marked biases within growth accounting or levels accounting exercises leading to an overestimation of the role of total factor productivity (TFP), as argued by Caselli (2005) and Jerzmanowski (2007), a feature which is avoided when the DEA approach is adopted. As for the predicted shape of the production function, DEA can only offer its finite-sample, piecewise linear approximation. With sufficiently large data samples, however, certain parametric forms could be tested formally against this approximate DEA-based nonparametric benchmark, such as the Cobb-Douglas or translog.

There are also limitations of the DEA approach. First, its deterministic character makes it silent on the estimation precision of the aggregate production function and of the predicted efficiency levels if inputs and outputs are subject to stochastic shocks. Second, the DEA provides a biased proxy of the actual technological frontier. In fact, even the most efficient units in the sample could possibly operate with some extra efficiency, since they are already aggregates of smaller economic units and must therefore have some internal heterogeneity. Taking account of that, the frontier would be shifted upwards; efficiency is nevertheless normalized to 100% for the most efficient units in the sample. The bootstrap method due to Simar and Wilson (1998; 2000b) helps in this respect by allowing for some corrections in the bias as well as for estimating confidence intervals for the actual efficiency levels and the technological frontier.

Third, the DEA constructs the aggregate production function on the basis of (relatively few) efficient data points. This makes it naturally sensitive to outliers and measurement error. This problem cannot be fully neutralized by bootstrap techniques: while they alleviate sampling bias, they cannot eliminate the bias due to measurement and aggregation errors. To address this problem, our data have been carefully filtered, ensuring that the outlying observations and cyclical fluctuations have been removed.

Simar and Wilson's bootstraps

As mentioned above, our deterministic DEA results have been complemented with Simar and Wilson's (SW) bootstraps. These procedures approximate the sampling distribution of an estimator by repeatedly simulating the Data Generating Process (DGP) under the assumption that the true production function is unknown and consequently the true Shephard distance functions E_{it} (for $i = 1, 2,...,I$ and $t = 1, 2,...,T$) are unknown, too. Simar and Wilson's bootstraps are then used to formulate an approximation of the sampling distribution of the difference $\hat{E}_{it} - E_{it}$, where $\hat{E}_{it}$ is the DEA estimator of E_{it}.

The exact procedure applied here is the homogenous bootstrap described by Simar and Wilson (1998). The procedure is based on the homogeneity assumption (cf. Simar, Wilson 2000a), that random variables $E_{1t},...,E_{It}$ are i.i.d. with an unknown density function g on the support (0, 1] (the output-oriented case). In particular, it means that we assume E_{it} to be independent of the random variables generating observed inputs and output $(\tilde{x}_i^t, \tilde{y}_i^t)$, where $x_i^t = [x'_{11} \ldots x'_{it}]'$, $y_i^t = [y_{11} \ldots y_{it}]'$. Vectors $(\tilde{x}_i^t, \tilde{y}_i^t)$, for $i = 1, 2,...,I$ and $t = 1, 2,...,T$, are assumed to be i.i.d., too. Their realizations are the observed input-output pairs $\{(x_i^t, y_i^t),\ i = 1, 2,...,I,\ t = 1, 2,...,T\}$. We use the procedure boot.sw98 contained in the free software package FEAR (written in R).

As the outcome of the homogenous SW bootstrap we receive, for each unit $i = 1, 2,...,I$ and $t = 1, 2,...,T$, the bootstrap estimate of the Shephard distance function $\hat{E}_{it}$ and a set of bootstrap realizations E_{itb}, b = 1, 2,..., B, where B = 2000 is the number of bootstrap iterations (see Simar, Wilson 1998). Usually, B = 2000 is considered sufficient in the literature. Consequently, we also obtain estimates of the bootstrap bias, variance of $\hat{E}_{it}$, and respective confidence intervals. Estimates $\hat{E}_{it}$ may also be additionally bias-corrected. If the bootstrap procedure is consistent, then asymptotically, these estimates may be used for E_{it}. Some Monte Carlo experiments conducted by Simar and Wilson (1998; 2000a) suggest that this SW bootstrap is indeed consistent. However, no rigorous proof of its consistency exists in the literature so far (cf. Simar, Wilson 2000a). The homogeneity assumption may be relaxed. The inefficiency of a unit would then depend on the observed values of inputs and outputs, i.e. on the pairs $\{(x_i^t, y_i^t)$ ($i = 1, 2,...,I$ and $t = 1, 2,...,T)\}$. Such procedures are called heterogeneous bootstraps and were first proposed in the paper by Simar and Wilson (2000b), where the pairs (x_i^t, y_i^t), were expressed in cylindrical coordinates. In the papers by Kneip, Simar and Wilson (2008; 2009) as well as Park, Jeong and Simar (2009), generalized procedures were proposed, allowing for orthonormal coordinates, with one of them being connected with E_{it} and constant returns to scale. These authors have also proposed formal proofs of consistency of certain bootstrap procedures.

Testing global and local returns to scale

In order to test the extent of returns to scale in the production technology on the basis of DEA-based estimates of the WTF, we have carried out formal tests of global and local returns to scale, first introduced by Löthgren and Tambour (1999) and improved by Simar and Wilson (2002).

As far as the test of global returns to scale is concerned, we use a procedure based on two nested tests proposed by Simar and Wilson (2002). In Test 1, the null hypothesis is tested that the aggregate production function (WTF) exhibits globally constant returns to scale (CRS) against an alternative hypothesis that the technology is characterized by variable returns to scale (VRS). That is:

Test 1:

H_0: technology is globally CRS,

H_1: technology is VRS.

If H_0 is rejected, we shall perform Test 2 with H_0 stating that the technology exhibits globally non-increasing returns to scale (NIRS) against H_1 that the technology is VRS:

Test 2:

H_0: technology is globally NIRS,

H_1: technology is VRS.

Simar and Wilson (2002) discussed various statistics for testing these hypotheses; among these, we have selected the following ratios of means:

$$\hat{S}_t^{CRS} = \frac{\sum_{\tau=1}^{t}\sum_{j=1}^{I}\hat{\theta}_{j\tau}^{CRS}(x_j^\tau, y_j^\tau)}{\sum_{\tau=1}^{t}\sum_{j=1}^{I}\hat{\theta}_{j\tau}^{VRS}(x_j^\tau, y_j^\tau)} \text{ in Test 1}$$

$$\hat{S}_t^{C-NIRS} = \frac{\sum_{\tau=1}^{t}\sum_{j=1}^{I}\hat{\theta}_{j\tau}^{CRS}(x_j^\tau, y_j^\tau)}{\sum_{\tau=1}^{t}\sum_{j=1}^{I}\hat{\theta}_{j\tau}^{NIRS}(x_j^\tau, y_j^\tau)} \text{ in Test 2}$$

where $\hat{\theta}_{j\tau}^{CRS}(x_j^\tau, y_j^\tau)$, $\hat{\theta}_{j\tau}^{VRS}(x_j^\tau, y_j^\tau)$ and $\hat{\theta}_{j\tau}^{NIRS}(x_j^\tau, y_j^\tau)$ are estimators of the (output-oriented) Debreu--Farrell distance function under the assumption of constant, variable, and non-increasing returns to scale, respectively.

By construction $\hat{S}_t^{CRS} \geq 1$ because $\hat{\theta}_{j\tau}^{CRS}(x_j^\tau, y_j^\tau) \geq \hat{\theta}_{j\tau}^{VRS}(x_j^\tau, y_j^\tau)$. The null hypothesis in Test 1 is rejected when $\hat{S}_t^{CRS}$ is significantly greater than 1. The p-value of the null hypothesis is derived by bootstrapping (see Simar, Wilson 2002):[2]

$$p\text{-value} = \sum_{b=1}^{B}\frac{I_{[0,+\infty)}(\hat{S}_t^{CRS,b} - \hat{S}_{obs,t}^{CRS})}{B} \tag{5}$$

where B = 2000 is the number of bootstrap replications, $I_{[0,+\infty)}$ is the indicator function, $\hat{S}_t^{CRS,b}$ is the b-th bootstrap sample, and $\hat{S}_{obs,t}^{CRS}$ is the original observed value. The same methodology is used in Test 2.

In turn, our statistical test of local returns to scale is based on bootstrap confidence intervals proposed by Simar and Wilson (2002). This returns-to-scale test (for each unit $j = 1, 2, \ldots, I$, and $t = 1, 2, \ldots, T$) is performed using the following nested testing procedure:

[2] To test the hypotheses regarding global returns to scale of the technology we use suitably modified codes written by Oleg Badunenko (see http://sites.google.com/site/obadunenko/codes).

Test 1:
$H_0: S_{jt}^{C-NIRS} = 1$ (scale-efficient or increasing returns to scale),
$H_1: S_{jt}^{C-NIRS} > 1$ (decreasing returns to scale).
If H_0 in Test 1 is not rejected, we proceed with the second test:
Test 2:
$H_0: S_{jt}^{CRS} = 1$ (scale-efficient),
$H_1: S_{jt}^{CRS} > 1$ (increasing returns to scale),

where

$$S_{jt}^{CRS} = \frac{\theta_{jt}^{CRS}(x_j^t, y_j^t)}{\theta_{jt}^{VRS}(x_j^t, y_j^t)}, \quad S_{jt}^{C-NIRS} = \frac{\theta_{jt}^{CRS}(x_j^t, y_j^t)}{\theta_{jt}^{NIRS}(x_j^t, y_j^t)}$$

and $\theta_{jt}^{CRS}(x_j^t, y_j^t), \theta_{jt}^{VRS}(x_j^t, y_j^t)$ and $\theta_{jt}^{NIRS}(x_j^t, y_j^t)$ are the output-oriented Debreu-Farrell distance functions under the assumption of constant, variable, and non-increasing returns to scale, respectively.

Let then $\hat{S}_{jt}^{*C-NIRS}(\alpha)$ and $\hat{S}_{jt}^{*CRS}(\alpha)$ denote the lower bound of the bootstrap $(1-\alpha)$-confidence interval for S_{jt}^{C-NIRS} and S_{jt}^{CRS}, respectively. The testing procedure is as follows: (i) if $\hat{S}_{jt}^{*C-NIRS}(\alpha) > 1$, then H_0 in Test 1 is rejected and we conclude that the technology features (locally) decreasing returns to scale; (ii) if $\hat{S}_{jt}^{*C-NIRS}(\alpha) = 1$, then H_0 in Test 1 cannot be rejected and we perform Test 2. If $\hat{S}_{jt}^{*CRS}(\alpha) > 1$, then the hypothesis of scale efficiency is rejected by Test 2 and we conclude that the technology exhibits (locally) increasing returns to scale; (iii) finally, if $\hat{S}_{jt}^{*CRS}(\alpha) = 1$, we conclude that the technology is (locally) scale-efficient.

Stochastic Frontier Analysis

To take a broader picture of the (in)efficiency in aggregate production processes in highly developed OECD countries, the results obtained with the DEA approach have been compared against estimates resulting from stochastic frontier analysis (SFA). In this alternative approach, stochastic disturbances are explicitly taken into account. Potential biases in technical efficiency estimates caused by stochastic variation, outliers and measurement error are thus minimized. Unfortunately, these advantages are only conditional on finding the appropriate parametric representation of the aggregate, WTF-based production function.

In its simplest, log-linear form, the stochastic frontier model for panel data, employed in the current paper, can be written as:

$$y_{it} = x_{it}'\beta + v_{it} - u_{it} \tag{6}$$

where $y_{it} = \ln Y_{it}$ represents the logarithm of output (log GDP), β represents the vector of estimated parameters, $u_{it} \geq 0$ captures technical inefficiency, and v_{it} is the idiosyncratic, symetrically distributed error term. Finally, the vector x_{it} carries information about n factors of production expressed in logarithms, plus a constant term.

Given this notation, the case $x_{it} = (1, \ln K_{it}, \ln H^U_{it}, \ln H^S_{it})$ represents our benchmark Cobb--Douglas specification with physical capital, unskilled labour and skilled labour as inputs. However, we shall also extend this vector to accommodate cross-terms as in:

$$x_{it} = (1, \ln K_{it}, \ln H^U_{it}, \ln H^S_{it}, \ln^2 K_{it}, \ln^2 H^U_{it}, \ln^2 H^S_{it}, \ldots \\ \ldots \ln K_{it} \ln H^U_{it}, \ln K_{it} \ln H^S_{it}, \ln H^U_{it} \ln H^S_{it}) \tag{7}$$

in which case equation (6) becomes the translog production function.

Constant returns to scale are either tested or directly imposed, wherever necessary, by writing down the production function in its intensive form. We shall do this in some of our estimated specifications, along with introducing certain regularity conditions which serve as a source of prior information and depend on the specification of the frontier. These regularity conditions enter the analysis via restricting average input elasticities to be non-negative:

$$EL_K = \frac{\partial y_i}{\partial \ln K_i} \geq 0, \; EL_{H^U} = \frac{\partial y_i}{\partial \ln H^U_i} \geq 0, \; EL_{H^S} = \frac{\partial y_i}{\partial \ln H^S_i} \geq 0$$

The sum of these three partial elasticities represents the measure of average returns to scale (scale elasticity).

Another issue is that applying the Stochastic Frontier methodology to panel data requires one to keep track of technological progress, which can strongly affect production capabilities. Obviously, it is “unfair” to evaluate the efficiency of observations from the past against a frontier estimated with a dataset including more recent data as well, since at earlier times, production processes could not enjoy the possibilities offered by technologies developed later on. To address this issue, we employ Battese and Coelli’s (1992; 1995) decomposition of the inefficiency term u_{it}. It takes the following form:

$$u_{it} = u_i \cdot z_t \tag{8}$$

where the fixed effect u_i is an exponentially distributed random variable,[3] whereas $z_t = \exp[-\eta(t-T)]$, where a positive (or negative) η indicates decreasing (or increasing, respectively) inefficiency over time.

Hence, the Battese-Coelli methodology urges the modeller to assume that the random part of u_{it} is time-invariant, and its temporal evolution is described by a deterministic function z_t with an estimated parameter η. This rigidity is, however, partly overcome when the WTF is estimated sequentially, so that for each period t, data from periods $\tau = 1, 2, \ldots, t$ are used. In such case, temporal shifts in u_{it} appear not only due to changes in z_t, but also due to the consecutive re-estimations of the WTF. The inefficiency term u_{it}, the Debreu-Farrell efficiency measure θ_{it} and the Shephard distance measure E_{it} are interrelated via the equality $\theta_{it} = 1/E_{it} = \exp(-u_{it})$.

[3] Robustness tests have been done upon the alternative assumption of a half-normal distribution, hardly affecting any of the results.

Estimating the WTF sequentially allows the fixed effect u_i to be reassessed in every period. As a result, we dispose of the uneasy assumption of a unique pattern of convergence to the WTF across all countries and years (e.g. Kumbhakar, Wang 2005) assume u_i to be a function of capital per worker in the initial period). On the other hand, we do not risk overparametrization of our model, which would have likely happened, had we assumed the parameters in β to be time-dependent (e.g. following linear trends as in Koop, Osiewalski, Steel 1999; Makieła 2009). Such an approach would be inadequate for a time horizon comparable to the one employed in our study.

Bayesian estimation framework

In the current study, all structural parameters of the production function $(y_{it} \mid u_{it}, \Theta) \sim N(x_{it}' \beta - u_{it}, \sigma^2)$, contained in the vector β, as well the variance of disturbances v_{it} and u_{it}, the mean of the inefficiency term u_{it}, denoted by φ^{-1}, and the pace of technological progress η, will be estimated in a Bayesian procedure. The first step of this procedure consists in making assumptions on the considered shapes of parameter distributions and endowing them with appropriate priors. The vector β is assumed to take the multivariate normal distribution (possibly truncated, to depict the regularity conditions), $\beta \sim N(\mu, \Sigma)$. The prior distribution of σ^{-2} is taken close to the "usual" flat prior, as in Koop, Osiewalski and Steel (1999). v_{it}'s are treated as independent normal variables with zero mean, unknown variance and with no autocorrelation over time (for all t, v_{it} is independent of $v_{i,t-1}$). The analysis starts with an assumption that u_{it}'s are independent exponentially distributed variables with mean φ^{-1} and no autocorrelation. In this case, $\varphi^{-1} \sim \mathrm{Exp}(-\ln r^*)$, which implies that prior median efficiency is equal to r^*. According to the findings presented in the literature, r^* should take the values from the interval [0.5, 0.9] (see Marzec, Osiewalski 2008; Makieła 2009). Having found that the final results are insensitive to any value choice out of the aforementioned interval, the prior efficiency median was set to 0.75.

The complexity of stochastic frontier models makes numerical integration methods inevitable. In the current study, as in most recent Bayesian literature, this procedure is based upon Gibbs sampling, which involves taking sequential random draws from the full conditional posterior distribution (cf. e.g. Koop, Steel, Osiewalski 1995). Under very mild assumptions (see Tierney 1994), these draws converge to the distribution of draws from the joint posterior. In the current research, implemented in WinBUGS, the characteristics of joint posterior distribution have been calculated on the basis of 300 000 burn-in-draws and 300 000 accepted (final) draws for different starting points. To evaluate the convergence of the Markov Chain Monte Carlo (MCMC) estimation procedure, the following tests were done: (a) assessment of the history plot (which plots the estimated parameter value against the iteration number); (b) autocorrelation tests: high autocorrelation might imply slow exploration of the entire posterior distribution; (c) evaluation of posterior kernel density plots.

Advantages and limitations of the SFA approach

A large amount of work has been devoted in the literature to the development of Bayesian methods suitable for making inference in stochastic frontier models. Some of the important advantages of this approach include: (i) the possibility of exact inference on technical efficiency in the presence of idiosyncratic disturbances, (ii) the possibility of using prior knowledge on the shape of aggregate production functions, and (iii) relatively easy incorporation of ideas and restrictions such as regularity conditions, or the optimal treatment of parameter and model uncertainty.

Although applications of Bayesian approaches to SFA are widespread in the empirical literature, some competing methods, such as the aforementioned deterministic DEA, or estimating CES production functions (cf. Duffy, Papageorgiou 2000; León-Ledesma, McAdam, Willman 2010), have also been strongly advocated. Undoubtedly, SFA makes it possible to account for the stochastic disturbances and measurement error to which the DEA method seems quite sensitive (cf. Koop, Steel 2001). Moreover, unlike the CES approach, it also allows one to take the WTF perspective. However, when choosing SFA (based upon either classical or Bayesian econometrics), the researcher has to make far more *a priori* assumptions: in particular, the parametrically specified translog production function might exhibit nonmonotonicities. The utmost objective of comparing DEA to SFA is thus to make these assumptions empirically testable, and hence to (partially) bridge the gap between the disconnected DEA and SFA literatures.

3. Results

Our estimates of the WTF provide testable implications on the following properties of the aggregate (country-level) production function:

1. Implied technical efficiency levels. How far is a given country in a given year from the WTF if the latter takes the given functional form? Is there any congruence of these distance measures across different functional specifications?

2. Partial elasticities. Are partial elasticities constant (as in the Cobb-Douglas specification)? If not, are they systematically related to inputs? If so, what is the pattern of dependence? Do we observe meaningful shifts in partial elasticities across time, from which we could infer that technical change favors some factors at the expense of others? Do the observed regularities agree or disagree with the hypothesis of skill-biased technical change?

3. Returns to scale. Viewed globally, can returns to scale be diagnozed as constant or variable? For each given country and year, are they constant, decreasing or increasing?

4. Elasticities of substitution. Are Morishima and Allen-Uzawa (two-factor) elasticities of substitution constant across countries and time (as they are in the Cobb-Douglas and CES specifications)? If not, can we observe indications of greater complementarity or substitutability of certain inputs in certain countries? Is the elasticity of substitution generally above or below one? Do the observed regularities agree or disagree with the hypothesis of capital-skill complementarity?

The first broad finding here is that the CRS Cobb-Douglas specification is the one which most frequently performs badly in our tests. We are however not able to offer an alternative parametric form of the function that would be in good agreement with nonparametric (bias-corrected) DEA results. In particular, our SFA-based estimates of translog production functions indicate visible departures of this particular functional specification from the nonparametric DEA results, too: the discrepancy pertains to implied efficiency levels, identified partial elasticities, and returns-to-scale properties. On the other hand, the same translog estimations provide a strong argument why the CRS Cobb-Douglas is too simple a specification to match the complex patterns present in the data. In fact, partial elasticities vary substantially across countries, are heavily correlated with factor endowments, and a number of Morishima (and Allen-Uzawa) elasticities of substitution are far away from unity.

The available evidence on constant vs. variable returns to scale is ambiguous. In a test of global constancy of returns to scale, the null of constant returns to scale can be rejected against the alternative of variable returns to scale with 99% confidence. In a series of DEA-based tests of local returns to scale (in a given country and year), however, the null of their constancy is relatively rarely rejected (although some countries do exhibit decreasing, rather than constant, returns to scale throughout the whole considered period). Unlike the DEA, the translog specification diagnozes a sharp correlation between returns to scale and the size of the economy.

3.1. Implied efficiency levels

Table 1 presents a comparison of seven different characterizations of the World Technology Frontier in the year 2000, computed on the basis of data for 1970–2000. In consecutive columns, we document Debreu-Farrell efficiency measures θ_i (such that potential output of country i at WTF is $Y_i^* = \theta_i Y_i$) computed according to the following methodologies: (a) bias-corrected DEA, (b) SFA with the Cobb-Douglas specification, (c) SFA translog. The variants differ also in the choice of inputs and assumptions on returns to scale. In the Table, we also report correlations between efficiency indexes computed on the basis of each specification. It turns out that in the cross-sectional dimension, DEA-based and SFA-based predictions on technical efficiency are quite strongly positively correlated.[4] Broadly the same group of countries is found to be closest to the frontier in all considered cases: Ireland, UK, and USA, and broadly the same group of countries lags behind: Finland, Greece, Japan, and Switzerland.

We do find some meaningful differences, however. Firstly, in the case of CRS Cobb-Douglas (and only in that case), all countries are found to be close to the frontier (less than 10% inefficiency) and there is little variation across countries. Given such comparison, this result suggests potential difficulties in identifying the inefficiency distribution u_{it} under this functional specification.[5] Secondly, the correlation between the CRS translog and unrestricted translog results is very strong, suggesting that there are only minor departures from CRS in the translog case. Thirdly, correlations are generally stronger within DEA and SFA estimates than between these two groups, suggesting that functional specification of the aggregate production function is quantitatively more important than the choice of its inputs.

Furthermore, treating the bootstrap-augmented DEA ($\theta_{DEA}(K, H^U, H^S)$) and the unrestricted translog SFA ($\theta_{TL}(K, H^U, H^S)$) as our two "benchmarks", representing the most general, nesting specifications in each of the two approaches, we have also computed the RMSE distance measures, quantifying the differences in predicted Debreu-Farrell technical inefficiency measures obtained from alternative methodologies. The results are in line with expectations: the distances are largest between the two general methodologies (DEA/SFA), and within these methodologies, the distances are the larger, the "simpler" is the measure in question (with the unrestricted Cobb-Douglas being an exception).

The conclusion from these findings is that the parametric functional forms used in our SFA analyses, especially the CRS Cobb-Douglas, are likely to be misspecified. They also constitute suggestive

4 We do not compare our DEA and SFA results across the time-series dimension because as opposed to DEA, our SFA results are based on data of decadal frequency only.

5 Direct evidence against the CRS Cobb-Douglas specification will be provided in section 3.4 where we estimate a non--unitary and non-constant elasticity of substitution of the underlying production function.

evidence that allowing for imperfect substitutability between unskilled and skilled labour helps obtain significantly different (and thus certainly better, since this step allows for more generality) results, supporting the related findings by Growiec (2012b; 2013).

On the other hand, the discrepancy between our DEA and SFA results could also indicate that the former method provides a relatively rough approximation of the WTF due to, e.g. sharp underrepresentation of certain input-output mixes in our dataset. Hence, the intepretation of these discrepancies in terms of Cobb-Douglas or translog production function misspecification must be treated with caution.

In sum, despite several important differences listed above, the ranking of countries in terms of their technical efficiency is similar under all functional specifications of the WTF. Hence, according to this test, the translog production function and the nonparametric frontier seem to identify a similar location of the WTF. Let us now assess its slope and curvature properties.

3.2. Partial elasticities

We shall now check if the partial elasticities of the aggregate production function tend to vary across countries and time if they are not restricted against such behavior. To this end, we have computed the partial elasticities of the aggregate production function both with DEA and SFA-translog. In related studies, Bernanke and Gürkaynak (2001) and Gollin (2002) have documented substantial differences in average levels of capital and labour income shares across countries. McAdam and Willman (2013) as well as Karabarbounis and Neiman (2014) have shown that factor shares also exhibit substantial short--run variation and prolonged trends in the time-series dimension. Our current exercise, documenting the variability of estimated partial elasticities, is complementary to these studies: partial elasticities and factor shares coincide under constant returns to scale and perfect competition, but not in general.

In the DEA approach, partial elasticities have been computed on the basis of the solution to each unit's optimal program. Knowing its maximum attainable output given inputs as well as the neighbouring efficient units, we have identified each of its partial elasticities on the basis of the local slope of the (piecewise linear) production function, projected along the axis associated with the respective factor of production. For fully efficient units, located in vertices of the technology set, so that their left-sided and right-sided partial elasticities do not coincide, we report only the right-sided partial elasticities here (i.e. percentage changes in output given a 1% increase in the respective input, holding everything else constant).[6]

First four (DEA-based) columns of Tables 2–3 document a negative correlation between the estimated partial elasticities and the scale of the economy. Hence, they confirm that the nonparametric WTF has more curvature than the Cobb-Douglas production function for which this correlation is zero. Another finding is that while the average partial elasticities are generally in line with the ones present in the established literature, they tend to vary largely across countries and exhibit visible trends across time. The DEA production function specification implies a consistent falling trend in the partial elasticity of unskilled labour, a moderately increasing trend in the physical capital elasticity, and an essentially flat trend in skilled labour elasticity. The finding that some of the reported right-sided

[6] Left-sided partial elasticities as well as partial elasticities based on the Simar-Wilson bootstrap are available from the authors upon request.

partial elasticities are very low or even zero, is an artifact of the construction of the DEA frontier as a convex hull of observed input-output pairs, with zero slope imposed on the function to the right of the highest efficient unit. Also by construction, left-sided partial elasticities must be greater or equal to the right-sided ones here. Hence, on the basis of right-sided partial elasticities reported here, one cannot make any inference regarding returns to scale. This will be done separately, using a different methodology.

Under the SFA translog specification, even though partial elasticities are by definition dependent on factor endowments, it remains uncertain if this would generate substantial variability of partial elasticities across countries and time. This would be the case only if second-order terms are quantitatively important. The latter four columns of Tables 2–3 show that it is indeed the case: translog-based partial elasticities vary strongly across countries and time. Reassuringly, similar patterns are observed both in the CRS and the VRS case.

In sum, despite visible differences (the correlations between partial elasticities computed for each of the 19 countries and 25 years in question according to the DEA- and SFA-translog specifications range from $E_K = 0.5234$, through $E_{H^S} = 0.0954$, to $E_{H^U} = -0.2717$), partial elasticities presented above share a few common properties. First, they vary largely across countries and time. Second, they are generally negatively correlated with output per worker (apart from the skilled labour elasticity under the translog specification), indicating that the frontier production function has more curvature than suggested by the Cobb-Douglas production function. Both these findings provide evidence against the latter functional specification.

Third, we also find that the unskilled labour elasticity is robustly falling over time, in line with the concept of skill-biased technical change: the more productive and more technologically advanced is the economy, the less it relies on unskilled labour for production. Fourth, this fall is counteracted by respective increases in the skilled labour elasticity (in the translog specification), and also partially by increases in the physical capital elasticity (in the DEA case). Both these trends are in line with the skill-biased technical change hypothesis, too, although the latter is conditional on some degree of capital-skill complementarity (cf. Henderson 2009). As we shall see shortly, the analysis of Allen-Uzawa and Morishima elasticities of substitution provides partial support for such complementarity. Fifth, we find a marked difference between partial elasticities estimated on the basis of DEA and SFA: in the former case, partial elasticities are much closer to the benchmark values found in other (not WTF-based) literature[7] than in the latter case. The average capital elasticity is around 0.4 in DEA as compared to 0.6 in SFA translog, and the skilled labour elasticity is around 0.4 in DEA as compared to 0.25 in SFA translog. This could be suggestive either of some production function misspecification issues inherent in the parametric WTF estimations, or of a bias induced when partial elasticities are identified with factor shares.

[7] The capital elasticity is usually taken to be close to 1/3, and the labour (or human capital) elasticity – close to 2/3 (e.g. Kydland, Prescott 1982), in line with average factor shares observed in the USA.

3.3. Returns to scale

Apart from the location and curvature issues discussed above, our WTF estimates also provide interesting conclusions on local and global returns to scale. Indeed, one advantage of the methods used in the current analysis is that they do not require the researcher to impose *a priori* restrictions on whether returns to scale are decreasing, increasing, or constant. Instead, this property is obtained as a result and can be statistically tested against the null of constant returns.

First, we have carried out the Simar and Wilson (2002) statistical test of global returns to scale. For each year 1980–2004, we reject the null of constant returns to scale for the alternative of variable returns to scale at 1% significance.

In tests of local returns to scale, carried out for each unit in the sample separately, on the other hand, we have found locally constant returns to scale in 66.5% of all cases. Decreasing returns have been identified in 29.1% of cases, and in some countries, such as Japan, France, Italy and the Netherlands, they have been found for all or almost all considered years. We have also observed a tendency of decreasing returns becoming more widespread in more recent years. On the other hand, increasing returns are found rarely (in 4.4% of all considered cases), and Finland in 1981–1989 is the only case when increasing returns were found in more than two consecutive years.

In sum, DEA-based returns-to-scale tests provide mixed evidence on this property. On the one hand, the aggregate production function is often locally indistinguishable from CRS; on the other hand, it is also robustly identified as globally VRS. A tentative conclusion would be that even though there exist both upward and downward departures from CRS (which are not systematically related to the inputs or output), returns to scale are also close to constant (perhaps marginally decreasing) on average.

Some inference on returns to scale can also be done using our SFA results. As shown in Table 4, results of estimations of the Cobb-Douglas and the translog production functions without the CRS restriction lead to a conclusion that returns to scale are generally country-specific, yet globally close to constant (again, perhaps marginally decreasing). When computed for the entire sample of countries, the scale elasticity is slightly below unity but not distinguishable from unity in the statistical sense. Country-specific translog production function estimates indicate, however, that returns to scale depend on the size of the economy. Unlike in the DEA case, they are decidedly increasing in the US and decreasing in smaller economies such as Norway and Ireland. Yet, this result might also reflect the misspecification of the estimated translog production function, so it should be treated with care. The relationship between the estimated scale elasticities and *per capita* variables is generally very weak and inconclusive.

In sum, the parametric and nonparametric approaches both tend to invalidate the assumption of constancy of global returns to scale, although on average, returns to scale might be in fact approximately constant. At the level of individual observations, there is however little congruence between results obtained with either method.

3.4. Morishima and Allen-Uzawa elasticities of substitution

Another important characteristic of the shape of a production function is its elasticity of substitution. In the two-input world, this quantity is uniquely defined and interpreted as local curvature of the isoquant (contour line of the production function), i.e. percentage change in the marginal rate of substitution between inputs given a 1% change in their relative price. The elasticity of substitution is an important measure of flexibility of production processes or the ease with which the inputs can be substituted. However, since our results described above (as well as the respective findings due to e.g. Caselli, Coleman 2006; Growiec 2012b) provide evidence against homogeneity of human capital, we are considering three-input production functions here, for which the elasticity of substitution is no longer a unique concept.

The two most frequently mentioned concepts of elasticity of substitution for n-input functions are the Allen-Uzawa and the Morishima elasticity (cf. Blackorby, Russell 1989). The first one is defined as (cf. Hoff 2004):

$$\sigma_{ij}^{A} = \frac{\sum_{k=1}^{n} X_k F_{X_k}}{X_i X_j} \frac{H_{ij}}{|H|}, \qquad i \neq j \tag{9}$$

for any two inputs $X_i, X_j \in \{K, H^U, H^S\}$, with $|\mathrm{H}|$ being the determinant of the bordered Hessian matrix:

$$H = \begin{bmatrix} 0 & F_K & F_{H^U} & F_{H^S} \\ F_K & F_{KK} & F_{KH^U} & F_{KH^S} \\ F_{H^U} & F_{KH^U} & F_{H^U H^U} & F_{H^U H^S} \\ F_{H^S} & F_{KH^S} & F_{H^U H^S} & F_{H^S H^S} \end{bmatrix} \tag{10}$$

and H_{ij} being the cofactor of (i, j)-th element in the H matrix.

The Allen-Uzawa elasticity of substitution is symmetric and simplifies to the unique elasticity of substitution in the two-input case. Unfortunately, as forcefully argued by Blackorby and Russell (1989), it does not measure the curvature of the underlying production function or the ease of input substitution appropriately, nor does it provide information about the comparative statics of income shares.

These two important critisms do not apply to the Morishima elasticity of substitution, which it thus a more theoretically sound concept of elasticity of substitution. The Morishima elasticity of substitution is defined as:

$$\sigma_{ij}^{M} = \frac{F_{X_j}}{X_i} \frac{H_{ij}}{|H|} - \frac{F_{X_j}}{X_j} \frac{H_{ij}}{|H|}, \qquad i \neq j \tag{11}$$

and thus $\sigma_{ij}^{M} \neq \sigma_{ji}^{M}$, signifying that the current measure is not symmetric. In the current study, we compute both measures.

Unfortunately, it is not possible to compute meaningful estimates of the elasticity of substitution for the DEA-based WTF, because – by construction – the production function is then piecewise linear, implying that the elasticity of substitution must be locally equal to either zero or infinity. Hence, we have computed these estimates only for our translog specification. Even these results should be interpreted with care, though, because the translog production function, as a log-quadratic function, might be a bad approximation of the true production function for atypical units, and sometimes even exhibits non-monotonicities. Country-specific results are presented in Table 5.

Our translog-based estimates of Morishima and Allen-Uzawa elasticities of substitution imply the following regularities:

1. According to Allen-Uzawa elasticities of substitution, capital and unskilled labour, as well as capital and skilled labour, are gross substitutes on average. Skilled and unskilled labour are generally complementary. There is substantial variation in these elasticities of substitution across countries. For all pairs of factors, substitutability does not exhibit any clear time trend. In some countries, the trend is increasing, whereas in others it is decreasing.

2. According to Morishima elasticities of substitution, when capital price increases, it can be relatively easily substituted with unskilled labour, but not with skilled labour. When unskilled labour wage increases, some of it can be substituted with capital, somewhat more easily than with skilled labour. When skilled labour wage increases, it can be relatively easily substituted with capital, easier than with unskilled labour.

3. Neither definition of the elasticity of substitution supports its constancy across countries and time (required in the CES case). None of the computed values of elasticity is close to unity on average (as required in the Cobb-Douglas specification).

Thus, our analysis provides partial evidence for the disputed concept of capital-skill complementarity. Using the translog specification instead of the CES, and basing our discussion on Morishima elasticities of substitution, we can provide complementary information to what is usually said in the related literature. In particular, we observe a one-sided relationship here: capital-skill complementarity is observed on average only when their relative price changes due to changes in capital price, not the skilled wage.

Unfortunately, our three-factor approach does not allow us to draw implications on whether the estimated elasticity of substitution between capital and labour is lower or greater than unity, i.e. if capital and labour are gross complements or substitutes. Studies surveyed by León-Ledesma, McAdam and Willman (2010), based on time-series data, tend to imply that they are gross complements, whereas cross-country studies are less conclusive and sometimes suggest the opposite (Duffy, Papageorgiou 2000; Karabarbounis, Neiman 2014). See also Chirinko (2008).

4. Conclusions

Summing up, the objective of the current paper has been to investigate the shape of the aggregate (country-level) production function based on the estimates of the World Technology Frontier (WTF). Using annual data on inputs and output in 19 highly developed OECD countries in 1970–2004, we have estimated the WTF both non-parametrically and parametrically (using the bias-corrected DEA and SFA approach, respectively) and then used these estimates to assess several properties of the aggregate production function.

We have obtained the following principal results:

– the CRS Cobb-Douglas production function fails to reproduce the important properties of our data (inferred inefficiency levels, estimated partial elasticities, elasticities of substitution),

– the (non-parametric) bootstrap-augmented DEA frontier is also markedly different from the unrestricted Cobb-Douglas and the translog, even though the latter offers much more flexibility and can be fitted to the data relatively well,

– regardless of the approach taken, the ranking of countries with respect to their technical efficiency is relatively stable,

– partial elasticities of the aggregate production function are correlated with inputs both in the DEA and in the translog case, and they vary substiantially across countries and time, providing strong evidence against the Cobb-Douglas specification, and also providing partial support for the skill-biased technical change hypothesis,

– tests of returns to scale based on the DEA, Cobb-Douglas and translog representations of the frontier provide mixed evidence on this property: returns to scale seem to be globally variable, but close to constant on average,

– unskilled and skilled labour are not perfectly substitutable,

– elasticities of substitution vary largely across countries and time, and there are some (partial) indications of capital-skill complementarity.

The results presented here constitute one of the first attempts so far to put together nonparametric (DEA) and parametric (SFA) estimates of the WTF, with the objective of drawing quantitative implications for the shape of aggregate production functions. Hence, even if our results are not entirely conclusive, at least they could be considered a useful benchmark for further investigations.

There is indeed a number of issues related to the current study that might be addressed in further research. First, one could re-estimate the WTF parametrically, using alternative estimation methodologies (e.g. maximum likelihood methods, different assumptions regarding the distribution of technical inefficiency and technical change, etc.), or nonparametrically, using alternative bootstraps. Such exercises could assess the extent to which our results are sensitive to the choice of estimation method. Second, one could also allow for other parametric forms of the aggregate production function apart from the Cobb-Douglas and the translog, in particular the CES production function. Third, one could look for links between the current study and the literature on misallocation (see e.g. Hsieh, Klenow 2009) that provides more direct explanations of cross-country differences in technical efficiency, albeit (again) generally based on the Cobb-Douglas production function specification.

References

Badunenko O., Henderson D.J., Zelenyuk V. (2008), Technological change and transition: relative contributions to worldwide growth during the 1990's, *Oxford Bulletin of Economics and Statistics*, 70(4), 461–492.

Battese G., Coelli T. (1992), Frontier production functions, technical efficiency and panel data: with application to paddy farmers in India, *Journal of Productivity Analysis*, 3(1–2), 153–169.

Battese G., Coelli T. (1995), A model for technical inefficiency effects in a stochastic frontier production function for panel data, *Empirical Economics*, 20, 325–332.

Bernanke B.S., Gürkaynak R.S. (2001), Is growth exogenous? Taking Mankiw, Romer and Weil seriously, in: B.S. Bernanke, K. Rogoff (eds.), *NBER Macroeconomics Annual 2001*, MIT Press, Cambridge.

Blackorby C., Russell R.R. (1989), Will the real elasticity of substitution please stand up? (A comparison of the Allen/Uzawa and Morishima elasticities), *American Economic Review,* 79(4), 882–888.

Bos J.W.B., Economidou C., Koetter M., Kolari J.W. (2010), Do all countries grow alike?, *Journal of Development Economics*, 91(1), 113–127.

Caselli F. (2005), Accounting for cross-country income differences, in: P. Aghion, S. Durlauf (eds.), *Handbook of economic growth*, Elsevier, Amsterdam.

Caselli F., Coleman W.J. (2006), The World Technology Frontier, *American Economic Review,* 96(3), 499–522.

Chirinko R.S. (2008), σ: the long and short of it, *Journal of Macroeconomics*, 30, 671–686.

Cohen D., Soto M. (2007), Growth and human capital: good data, good results, *Journal of Economic Growth*, 12(1), 51–76.

Duffy J., Papageorgiou C. (2000), The specification of the aggregate production function: a cross--country empirical investigation, *Journal of Economic Growth*, 5, 83–116.

Färe R., Grosskopf S., Noriss M., Zhang Z. (1994), Productivity growth, technical progress, and efficiency change in industrialized countries, *American Economic Review*, 84(1), 66–83.

Fried H.O., Lovell C.A.K., Schmidt S.S., eds. (1993), *The measurement of productive efficiency: techniques and applications*, Oxford University Press, New York, Oxford.

de la Fuente A., Doménech R. (2006), Human capital in growth regressions: How much difference does data quality make?, *Journal of the European Economic Association*, 4(1), 1–36.

Gollin D. (2002), Getting income shares right, *Journal of Political Economy*, 110(2), 458–474.

Growiec J. (2012a), *Zagregowana funkcja produkcji w ekonomii wzrostu gospodarczego i konwergencji*, Oficyna Wydawnicza SGH, Warszawa.

Growiec J. (2012b), The World Technology Frontier: What can we learn from the US States?, *Oxford Bulletin of Economics and Statistics*, 74(6), 777–807.

Growiec J. (2013), On the measurement of technological progress across countries, *Bank i Kredyt,* 44(5), 467–504.

Growiec J., Pajor A., Pelle D., Prędki A. (2011), *The shape of aggregate production functions: evidence from estimates of the World Technology Frontier*, NBP Working Paper, 102, Narodowy Bank Polski, Warszawa.

Hall R.E., Jones C.I. (1999), Why do some countries produce so much more output per worker than others?, *Quarterly Journal of Economics*, 114(1), 83–116.

Henderson D.J. (2009), A non-parametric examination of capital-skill complementarity, *Oxford Bulletin of Economics and Statistics*, 71(4), 519–538.

Henderson D.J., Russell R.R. (2005), Human capital and convergence: a production-frontier approach, *International Economic Review*, 46(4), 1167–1205.

Heston A., Summers R., Aten B. (2006), *Penn World Table version 6.2*, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.

Hodrick R.J., Prescott E.C. (1997), Postwar U.S. business cycles: an empirical investigation, *Journal of Money, Credit and Banking*, 29(1), 1–16.

Hoff A. (2004), The linear approximation of the CES function with *n* input variables, *Marine Resource Economics*, 19, 295–306.

Hsieh C.-T., Klenow P. (2009), Misallocation and manufacturing TFP in China and India, *Quarterly Journal of Economics*, 124, 1403–1448.

Jerzmanowski M. (2007), Total factor productivity differences: appropriate technology vs. efficiency, *European Economic Review*, 51, 2080–2110.

Karabarbounis L., Neiman B. (2014), The global decline of the labour share, *Quarterly Journal of Economics*, 129(1), 61–103.

Klump R., McAdam P., Willman A. (2007), Factor substitution and factor-augmenting technical progress in the US: a normalized supply-side system approach, *Review of Economics and Statistics*, 89, 183–192.

Kneip A., Simar L., Wilson P.W. (2008), Asymptotics and consistent bootstraps of DEA estimators in non-parametric frontier models, *Econometric Theory*, 24, 1663–1697.

Kneip A., Simar L., Wilson P.W. (2009), A computationally efficient, consistent bootstrap for inference with non-parametric DEA estimators, Discussion Paper, 0903, Institut de Statistique, Université catholique de Louvain.

Koop G., Osiewalski J., Steel M.F.J. (1999), The components of output growth: a stochastic frontier analysis, *Oxford Bulletin of Economics and Statistics*, 61(4), 455–487.

Koop G., Osiewalski J., Steel M.F.J. (2000), Measuring the sources of output growth in a panel of countries, *Journal of Business and Economic Statistics*, 18(3), 284–299.

Koop G., Steel M.F.J., Osiewalski J. (1995), Posterior analysis of stochastic frontier models using Gibbs sampling, *Computational Statistics*, 10, 353–373.

Koop G., Steel M.F.J. (2001), Bayesian analysis of stochastic frontier models, in: B. Baltagi (ed.), *A companion to theoretical econometrics*, Blackwell, Oxford.

Kumar S., Russell R.R. (2002), Technological change, technological catch-up, and capital deepening: relative contributions to growth and convergence, *American Economic Review*, 92(3), 527–548.

Kumbhakar S.C., Lovell C.A.K. (2000), *Stochastic frontier analysis,* Cambridge University Press, Cambridge.

Kumbhakar S.C., Wang H.-J. (2005), Estimation of growth convergence using a stochastic production frontier approach, *Economics Letters*, 88(3), 300–305.

Kydland F.E., Prescott E.C. (1982), Time to build and aggregate fluctuations, *Econometrica*, 50(6), 1345–1370.

León-Ledesma M., McAdam P., Willman A. (2010), Identifying the elasticity of substitution with biased technical change, *American Economic Review*, 100(4), 1330–1357.

Löthgren M., Tambour M. (1999), Testing scale efficiency in DEA models: a bootstrapping approach, *Applied Economics*, 31, 1231–1237.

McAdam P., Willman A. (2013), Medium run redux, *Macroeconomic Dynamics*, 17(4), 695–727.

Makieła K. (2009), Economic growth decomposition. An empirical analysis using Bayesian Frontier Approach, *Central European Journal of Economic Modelling and Econometrics*, 1, 333–369.

Marzec J., Osiewalski J. (2008), Bayesian inference on technology and cost efficiency of bank branches, *Bank i Kredyt*, 39, 29–43.

Park B.U., Jeong S.O., Simar L. (2009), Asymptotic distribution of Conical-Hull estimators of directional edges, Discussion Paper, 0907, Institut de Statistique, Université catholique de Louvain.

Simar L., Wilson P.W. (1998), Sensitivity of efficiency scores: How to bootstrap in nonparametric frontier models, *Management Sciences*, 44(1), 49–61.

Simar L., Wilson P.W. (2000a), Statistical inference in nonparametric frontier models: state of the art, *Journal of Productivity Analysis*, 13, 49–78.

Simar L., Wilson P.W. (2000b), A general methodology for bootstrapping in non-parametric frontier models, *Journal of Applied Statistics*, 27(6), 779–802.
Simar L., Wilson P.W. (2002), Non-parametric tests of returns to scale, *European Journal of Operational Research*, 139(1), 115–132.
Tierney L. (1994), Markov chains for exploring posterior distributions, *Annals of Statistics*, 22, 1701–1762.

Acknowledgements

The authors thank Michał Gradzewicz, Kamil Makieła, Jacek Osiewalski, anonymous Reviewers, and seminar participants at Narodowy Bank Polski, Warsaw School of Economics, University of Trento, and Warsaw University, as well as the European Workshop on Efficiency and Productivity Analysis (EWEPA 2011) and the EcoMod 2011 conference, for their useful comments and suggestions.

Financial support from the Economic Research Committee, Narodowy Bank Polski, is gratefully acknowledged. The views presented here belong to the authors and have not been endorsed by Narodowy Bank Polski. All remaining errors are authors' responsibility.

Appendix

Table 1

Technical efficiency in 2000 – comparison of alternative measurements

	DEA (K, L) [CRS]	DEA (K, H) [VRS]	DEA (K, Hu, Hs) [VRS]	SFA-CD (K, Hu, Hs) [VRS]	SFA-CD (K, Hu, Hs) [CRS]	SFA--translog (K, Hu, Hs) [CRS]	SFA--translog (K, Hu, Hs) [VRS]	Mean
Australia	1.2021	1.2381	1.2421	1.2029	1.0438	1.0979	1.0810	1.1583
Austria	1.1628	1.2407	1.2397	1.1959	1.0395	1.0869	1.0857	1.1502
Belgium	1.0734	1.1972	1.1750	1.1506	1.0328	1.0655	1.0580	1.1075
Canada	1.2075	1.2700	1.2685	1.1599	1.0320	1.0772	1.0590	1.1535
Denmark	1.2168	1.2248	1.2447	1.2376	1.0466	1.1070	1.1296	1.1724
Finland	1.2469	1.3526	1.3589	1.3015	1.0634	1.1446	1.1688	1.2338
France	1.1406	1.2041	1.2097	1.1552	1.0421	1.0884	1.0876	1.1326
Greece	1.3207	1.3532	1.3419	1.2452	1.0553	1.1263	1.1231	1.2237
Ireland	1.0133	1.0635	1.0835	1.1174	1.0166	1.0324	1.0922	1.0598
Italy	1.1824	1.2216	1.0809	1.1355	1.0424	1.0818	1.0870	1.1188
Japan	1.4641	1.4780	1.3949	1.2724	1.0728	1.1686	1.1991	1.2928
Netherlands	1.2522	1.2683	1.2088	1.2641	1.0637	1.1250	1.1016	1.1834
Norway	1.0377	1.0574	1.1603	1.2700	1.0563	1.1038	1.1409	1.1180
Portugal	1.2806	1.2104	1.1743	1.0205	1.0035	1.0118	1.0198	1.1030
Spain	1.1920	1.1944	1.1203	1.0346	1.0122	1.0243	1.0221	1.0857
Sweden	1.1964	1.1791	1.1470	1.2281	1.0476	1.1052	1.1073	1.1444
Switzerland	1.4088	1.4592	1.4726	1.3335	1.0688	1.1369	1.1547	1.2907
UK	1.0140	1.0263	1.0243	1.0396	1.0109	1.0235	1.0229	1.0231
USA	1.0104	1.1210	1.0887	1.0152	1.0051	1.0181	1.0355	1.0420
Corr. with DEA	0.8222	0.9110	1.0000	0.7615	0.7421	0.7908	0.7189	
RMSE Dev. / DEA	0.0748	0.0520	0.0000	0.0822	0.1993	0.1518	0.1467	
Corr. with SFA-TL	0.5566	0.6202	0.7189	0.9014	0.9027	0.9139	1.0000	
RMSE Dev. / SFA-TL	0.1428	0.1662	0.1467	0.1009	0.0625	0.0217	0.0000	

Source: reproduced as Table 5.7 in Growiec (2012a).

Table 2
Partial elasticities estimated from the DEA (piecewise linear) and the SFA-translog production function. Cross-country averages

	DEA				Translog			
	E_K	E_{HU}	E_{HS}	scale	E_K	E_{HU}	E_{HS}	scale
Australia	0.35	0.16	0.48	0.99	0.61	0.06	0.28	0.95
Austria	0.39	0.21	0.41	1.01	0.65	0.02	0.22	0.89
Belgium	0.32	0.31	0.37	1.00	0.66	0.09	0.18	0.93
Canada	0.32	0.21	0.65	0.90	0.51	0.02	0.40	0.93
Denmark	0.45	0.20	0.43	1.03	0.66	-0.04	0.21	0.83
Finland	0.46	0.24	0.32	1.02	0.66	0.00	0.18	0.84
France	0.21	0.25	0.51	0.98	0.60	0.21	0.29	1.10
Greece	0.47	0.30	0.24	1.01	0.67	0.09	0.16	0.92
Ireland	0.26	0.25	0.38	0.37	0.72	-0.06	0.12	0.78
Italy	0.57	0.00	0.23	0.53	0.57	0.33	0.23	1.13
Japan	0.00	0.00	0.75	0.75	0.53	0.28	0.37	1.17
Netherlands	0.13	0.35	0.55	0.94	0.56	0.14	0.26	0.96
Norway	0.00	0.20	0.32	0.40	0.50	0.01	0.28	0.80
Portugal	0.54	0.25	0.15	0.94	0.68	0.15	0.09	0.91
Spain	0.34	0.25	0.23	0.65	0.63	0.26	0.18	1.07
Sweden	0.55	0.18	0.43	1.03	0.64	-0.01	0.24	0.87
Switzerland	0.30	0.73	0.57	1.05	0.49	0.00	0.34	0.83
UK	0.65	0.06	0.07	0.70	0.72	0.16	0.23	1.11
USA	0.00	0.00	0.00	0.00	0.45	0.20	0.50	1.15
Mean	0.42	0.27	0.40	0.85	0.60	0.10	0.25	0.96
Corr. with *Y/L*	-0.19	-0.22	-0.15	-0.39	-0.55	-0.06	0.58	0.05

Notes: means have been computed excluding zeros. Translog parameters, computed for the 1980–2000 dataset, have been assumed constant over time.

Source: reproduced as Tables 5.10 and 5.12 in Growiec (2012a).

Table 3
Partial elasticities estimated from the DEA (piecewise linear) and the SFA-translog production function. Annual averages

	DEA				Translog			
	E_K	E_{HU}	E_{HS}	scale	E_K	E_{HU}	E_{HS}	scale
1980	0.30	0.43	0.39	0.90	0.63	0.13	0.20	0.96
1981	0.31	0.41	0.40	0.90	0.63	0.13	0.20	0.96
1982	0.30	0.39	0.40	0.88	0.63	0.12	0.20	0.96
1983	0.28	0.42	0.39	0.89	0.63	0.12	0.21	0.96
1984	0.35	0.38	0.41	0.91	0.63	0.12	0.21	0.96
1985	0.33	0.43	0.41	0.93	0.63	0.11	0.21	0.96
1986	0.42	0.33	0.41	0.85	0.63	0.11	0.22	0.96
1987	0.36	0.33	0.42	0.83	0.63	0.11	0.22	0.96
1988	0.35	0.33	0.40	0.82	0.62	0.11	0.23	0.96
1989	0.33	0.32	0.42	0.82	0.62	0.11	0.24	0.96
1990	0.31	0.29	0.43	0.80	0.61	0.11	0.24	0.96
1991	0.31	0.29	0.41	0.83	0.61	0.11	0.24	0.96
1992	0.39	0.30	0.43	0.91	0.61	0.10	0.25	0.96
1993	0.45	0.30	0.41	0.93	0.62	0.10	0.25	0.96
1994	0.47	0.31	0.39	0.94	0.62	0.09	0.25	0.96
1995	0.46	0.32	0.40	0.92	0.61	0.09	0.26	0.96
1996	0.46	0.15	0.43	0.82	0.61	0.09	0.26	0.96
1997	0.50	0.14	0.42	0.80	0.60	0.09	0.27	0.96
1998	0.55	0.12	0.38	0.79	0.59	0.09	0.28	0.95
1999	0.53	0.12	0.40	0.81	0.58	0.08	0.29	0.95
2000	0.50	0.16	0.38	0.81	0.57	0.08	0.30	0.95
2001	0.49	0.17	0.37	0.81	0.56	0.08	0.30	0.95
2002	0.51	0.14	0.38	0.79	0.56	0.08	0.31	0.95
2003	0.51	0.12	0.36	0.77	0.55	0.08	0.32	0.94
2004	0.57	0.11	0.34	0.76	0.54	0.08	0.32	0.94
Mean	0.42	0.27	0.40	0.85	0.60	0.10	0.25	0.96
Corr. with Y/L	-0.19	-0.22	-0.15	-0.39	-0.55	-0.06	0.58	0.05

Notes: means have been computed excluding zeros. Translog parameters, computed for the 1980–2000 dataset, have been assumed constant over time.

Source: reproduced as Tables 5.11 and 5.13 in Growiec (2012a).

Table 4
Returns to scale – evidence from stochastic frontier estimates

	Mean scale elasticity	**1970**	**2000**
SFA-CD (K, Hu, Hs)	0.960		
Australia	0.947	0.939	0.953
Austria	0.889	0.896	0.881
Belgium	0.929	0.924	0.933
Canada	0.939	0.962	0.897
Denmark	0.835	0.840	0.827
Finland	0.840	0.827	0.842
France	1.104	1.111	1.087
Greece	0.911	0.890	0.929
Ireland	0.780	0.771	0.782
Italy	1.119	1.096	1.139
Japan	1.177	1.180	1.160
Netherlands	0.956	0.939	0.957
Norway	0.804	0.821	0.763
Portugal	0.898	0.868	0.926
Spain	1.057	1.023	1.089
Sweden	0.879	0.893	0.854
Switzerland	0.830	0.834	0.822
UK	1.114	1.125	1.091
USA	1.164	1.199	1.124
Translog (K, Hu, Hs) mean	0.956	0.955	0.950
Corr. with K/L	0.041	-0.065	-0.028
Corr. with H^U/L	-0.017	-0.227	0.136
Corr. with H^S/L	0.212	0.214	0.078
Corr. with Y/L	0.028	0.203	-0.059
Corr. with L	0.709	0.759	0.638
Corr. with Y	0.654	0.708	0.582

Note: scale elasticities above (below) 1 indicate IRS (respectively, DRS).

Source: reproduced as Table 5.16 in Growiec (2012a).

Table 5
Morishima and Allen-Uzawa elasticities of substitution, inferred from the translog production function. Cross-country averages

	Morishima EoS						**Allen-Uzawa EoS**		
	$E(K, H^U)$	$E(K, H^S)$	$E(H^U, K)$	$E(H^U, H^S)$	$E(H^S, K)$	$E(H^S, H^U)$	$E(K, H^U)$	$E(K, H^S)$	$E(H^U, K)$
Australia	-16.37	-7.19	0.93	-0.93	0.29	0.61	1.46	0.46	-4.92
Austria	-2.37	-8.57	0.73	-0.39	0.41	0.06	1.02	0.57	-3.45
Belgium	-8.06	-6.61	0.69	-0.11	0.33	0.07	0.99	0.47	-8.21
Canada	2.46	-12.51	0.89	-0.89	0.30	0.32	1.65	0.57	-2.69
Denmark	6.13	-8.21	0.74	-0.31	0.45	-0.14	0.94	0.58	-2.22
Finland	0.13	-8.12	0.67	-0.15	0.41	-0.03	0.87	0.54	-3.31
France	16.80	-32.20	-0.26	1.29	1.02	-2.10	-0.51	1.92	8.88
Greece	-5.30	-7.61	0.58	1.05	0.38	-0.52	0.81	0.53	-12.18
Ireland	3.62	-4.33	0.57	-0.03	0.27	-0.01	0.65	0.31	-2.89
Italy	-9.71	-30.44	0.22	-0.45	0.60	0.33	0.45	1.20	3.41
Japan	12.73	-36.51	-0.07	0.85	0.76	-2.05	-0.17	1.73	4.13
Netherlands	-63.48	23.46	2.09	-2.51	-0.70	1.93	3.56	-1.18	-24.26
Norway	2.32	-20.30	0.75	-0.29	0.34	-0.03	1.20	0.56	-2.69
Portugal	0.46	-16.35	-0.11	3.87	1.17	-2.27	-0.14	1.57	-22.14
Spain	-7.57	-16.32	0.33	-0.83	0.51	0.60	0.56	0.89	5.84
Sweden	3.58	-9.17	0.76	-0.41	0.44	-0.10	1.05	0.61	-2.84
Switzerland	2.16	-19.60	0.79	-0.57	0.36	-0.02	1.35	0.62	-1.95
UK	-190.52	66.19	5.84	-12.44	-3.00	20.51	9.90	-5.31	-50.22
USA	431.60	-65.92	-1.33	4.15	1.22	-21.34	-3.62	3.26	10.12
Mean	9.40	-11.59	0.78	-0.48	0.29	-0.22	1.16	0.52	-5.87

Note: translog parameters, computed for the 1980–2000 dataset, have been assumed constant over time.

Source: reproduced as Table 5.17 in Growiec (2012a).