Production technology estimates and balanced growth*

MIGUEL A. LEÓN-LEDESMA,† PETER McADAM,‡§
and ALPO WILLMAN‡†

† *School of Economics, University of Kent,*
Keynes College, CT27NP, Canterbury, UK
(e-mail: m.a.leon-ledesma@kent.ac.uk)
‡*Research Dept., European Central Bank,*
(e-mail: peter.mcadam@ecb.int)
(e-mail: alpo.willman@gmail.com)
§ *Department of Economics, University of Surrey*

Abstract
Capital-labor substitution and TFP estimates are essential features of many economic models. Such models typically embody a balanced growth path. This often leads researchers to estimate models imposing stringent prior choices on technical change. We demonstrate that estimation of the substitution elasticity and TFP growth can be substantially biased if technical progress is thereby mis-specified. We obtain analytical and simulation results in the context of a model consistent with balanced and near-balanced growth (i.e., departures from balanced growth but broadly stable factor shares). Given this evidence, a Constant Elasticity of Substitution production function system is then estimated for the US economy. Results show that the estimated substitution elasticity tends to be significantly lower using a factor-augmenting specification (well below one). We are also able to reject conventional neutrality forms in favor of general factor augmentation with a non-negligible capital-augmenting component. Our work thus provides insights into production and supply-side estimation in balanced-growth frameworks.

JEL classification numbers
C15, C32, E23, O33, O51.

*This paper is a substantially updated and revised version of an earlier draft distributed as "In Dubio Pro CES" (ECB Working Paper No. 1175).We thank two anonymous referees, Beata Javorcik (editor), Youngjae Chang (discussant), Manuel Gómez-Suárez, Rainer Klump, Kangwoo Park (discussant), Efthimios Tsionas (discussant), Anders Warne and participants at numerous seminars for comments. Opinions are the authors' and not necessarily those of the ECB.

Keywords
Balanced Growth, Technical Progress Neutrality, Factor Income share, Constant Elasticity of Substitution, Factor-Augmenting Technical Change.

I Introduction

A balanced growth path (BGP) defines an equilibrium in which macroeconomic variables such as output, consumption, etc., tend to a common growth rate, whilst key underlying ratios (e.g., factor income shares, capital-output ratio, and the real interest rate) are constant, Kaldor (1961). In terms of neoclassical growth theory (Uzawa (1961)), it requires that either, technical progress be labor-augmenting (i.e., Harrod Neutral), or production is Cobb Douglas (i.e., exhibits a unitary elasticity of substitution between input factors).

Although balanced growth looks a reasonable description of many economies and is a common and tractable modelling narrative, these two explanations are widely disputed.[1] For instance, there is now mounting evidence in favor of a below-unity aggregate substitution elasticity (e.g., Chirinko (2008)). Likewise, that all technical change is labor augmenting appears unduly restrictive. Recent theoretical literature (Acemoglu (2007)) also argues that while technical progress is asymptotically labor-augmenting, it may become capital-biased in transition reflecting incentives for factor-saving innovations.[2] Despite these concerns, researchers, guided by tractability and the apparent "stylized facts", invariably impose balanced growth path (BGP) conditions for estimating key supply side parameters such as the elasticity of capital-labor substitution and total factor productivity (TFP).

Arguably, though, the costs of doing so are unknown. To fill this important gap we hence analyze the potential consequences of imposing prior beliefs on the form of technical progress for such estimates. In particular, we study how estimates of the elasticity of substitution and TFP are affected by imposing a priori restrictions on the direction of technical change where the economy may depart to a large or small extent from BGP. To motivate matters, we first use some theory to highlight a set of (potential) pitfalls related to parameter inference and TFP approximations. Then we analyze the practical importance of these biases in a simulation experiment. Finally, in light of our analysis, we estimate a production-technology system of the US economy over 1952-2009 under different technical progress specifications and compare the resulting estimates of the substitution elasticity and TFP. Our reference point is the flexible "factor-augmenting" Constant Elasticity of Substitution (CES) production function.

Following our earlier contribution, León-Ledesma et al. (2010), we exploit Monte-

[1] See Attfield and Temple (2010) for an empirical assessment of the BGP conditions and a discussion of previous studies of the empirical validity of the BGP.

[2] Other perspectives draw on the distributional form of technical change over time, Jones (2005), Growiec (2008), or the endogenous choice of production technology, León-Ledesma and Satchi (2011).

Carlo methods. Compared to that paper, though, the set up and motivation are quite different. León-Ledesma et al. (2010) analyzed means to estimate production-technology parameters (linear, non-linear, single and multi-equation) and showed that the normalized (or indexed) non-linear system estimation allowed for identification of the key technology parameters. Here, we take that approach as given but use it to explore the more applied and more specific topic of econometric mis-specification and the robustness of balanced growth and particular neutrality assumptions.

Our analysis shows that, generally, when the true nature of technical progress is factor-augmenting, imposing Hicks-neutrality leads to biases towards Cobb-Douglas (unit elasticity). Imposing Harrod-neutrality would generally lead to upward biases in the estimated elasticity if the true elasticity is below unity and downward biases if it is above unity. We rationalize these various biases as attempts by the estimator to control for trends in the data (e.g., in capital deepening) otherwise incompatible with the presumed neutrality concept. We also show that TFP growth approximations from CES estimates crucially depend on the elasticity of substitution, which governs the transmission of capital deepening and technical progress components into the evolution of TFP. Hence, biases in the estimated elasticity will be reflected in biases in estimated TFP growth.

When we estimate using US data, many of the previous lessons find an echo in empirical estimates. Although results yield different values for the substitution elasticity for different a priori technical progress restrictions. In all cases, our tests support the general factor-augmenting specification with a capital-labor substitution elasticity well below one. We also find a non-negligible capital-augmenting technical progress component.

The paper is organized as follows. In section II we present some relevant background on the CES production function and in section III discuss the potential biases arising from mis-specification of technical change. In Section IV we present the simulation setup and discuss the results. Section V presents empirical results using US data. Finally, Section VI concludes.

II Background

The CES production function was formally introduced in economics by Arrow et al. (1961) and spawned a vast literature.[3] Following the work of La Grandville (1989) and Klump and de La Grandville (2000), the function is often now expressed in "normalized" (or indexed) form since its parameters then have a direct economic

[3]See Klump et al. (2012) for a survey.

interpretation:

$$Y_t = F\left(\Gamma_t^K K_t,\, \Gamma_t^N N_t\right) = Y_0\left[\pi_0\left(\frac{\Gamma_t^K K_t}{\Gamma_0^K K_0}\right)^{\frac{\sigma-1}{\sigma}} + (1-\pi_0)\left(\frac{\Gamma_t^N N_t}{\Gamma_0^N N_0}\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{1}$$

where the point of time $t = 0$ represents the point of normalization, Y_t represents real output, K_t is the real capital stock and N_t is the labor input.

Terms Γ_t^K and Γ_t^N capture capital and labor-augmenting technical progress. To circumvent problems related to the Diamond-McFadden impossibility theorem, researchers usually assume specific functional forms for technical progress, e.g., $\Gamma_t^K = \Gamma_0^K e^{\gamma_K t}$ and $\Gamma_t^N = \Gamma_0^N e^{\gamma_N t}$ where γ_i denotes growth in technical progress associated to factor i, $t = 1, 2, \ldots, T$ represents a time trend. Technical progress can be Hicks neutral ($\gamma_K = \gamma_N > 0$), Harrod neutral ($\gamma_K = 0$,$\gamma_N > 0$) or, more seldom, Solow-Neutral ($\gamma_K > 0$,$\gamma_N = 0$). A *general* factor-augmenting case ($\gamma_K > 0 \neq \gamma_N > 0$), though, is typically by-passed.

The capital income share at the point of normalization is $\pi_0 = \frac{r_0 K_0}{Y_0}$ (r denotes the real user cost of capital) and the elasticity of substitution between capital and labor inputs is given by the percentage change in factor proportions due to a change in the factor price ratio along an isoquant:

$$\sigma \in [0, \infty) = \frac{d\log(K/N)}{d\log(F_N/F_K)} \tag{2}$$

CES production function (1) nests Cobb Douglas when $\sigma = 1$; the Leontief function (i.e., fixed factor proportions) when $\sigma = 0$; and a linear production function (i.e., perfect factor substitutes) when $\sigma \to \infty$. The higher is σ, the greater the *similarity* between capital and labor: when $\sigma < 1$, factors are gross complements in production and gross substitutes otherwise. It can be shown that with gross substitutes, substitutability between factors allows both the augmentation and bias of technological change to "favor" the same factor.[4] For gross complements, however, a capital-augmenting technological change, to be specific, increases demand for labor (the complementary input) more than it does capital, and vice versa. By contrast, when $\sigma \to 1$ an increase in technology does not produce a bias towards either factor (factor shares will always be constant since any change in factor proportions will be offset by a change in factor prices). Thus, as we shall soon appreciate, the question of whether σ is above or below unity is possibly as important as its numerical value.

[4] In other words, if $\sigma > 1$ and $\gamma_i > \gamma_j$ this implies that $F_i > F_j$ plus that there is a relative rise in the income share of factor i . Hence we can say that technical change related to factor i "favors" factor i in the gross substitutes case.

III Mis-specified technical change: two examples

We now discuss the general issues at stake and analytically derive some potential estimation problems. In Sections III and III, we consider the particular impact of mis-specification of technical progress on the estimation of the elasticity of substitution, and then on TFP estimates and its decompositions.

These examples, note, are meant to be primarily motivational: they usefully highlight many of the issues that will become apparent in both the simulation and data estimation sections.

Mis-specified technical change: parameter inference

The relative capital-to-labor income share, given competitive factor markets and profit maximization, can be expressed as,

$$\Theta_t = \frac{r_t}{w_t}\frac{K_t}{N_t} = \frac{\pi_0}{1-\pi_0}\left(\frac{\Gamma_t^K K_t/K_0}{\Gamma_t^N N_t/N_0}\right)^{\frac{\sigma-1}{\sigma}} \tag{3}$$

Whilst Θ_t is observed, neither the substitution elasticity nor technical change are. For Θ to be constant requires the familiar balanced growth cases of $\sigma = 1$ or Harrod neutrality. But can $d\Theta \approx 0$ (i.e., a *near* balanced growth path) arise when we purposefully depart form these two restrictive assumptions? And what would be the consequences?:

(i). Equation (3) shows that if we assume Hicks neutrality, stable factor shares require $\widehat{\sigma} \to 1$ to offset any trend in capital deepening. Antràs (2004) uses this argument to rationalize Berndt (1976)'s widely-cited finding of Cobb-Douglas for US manufacturing.

(ii). The same is true of Solow neutrality.

(iii). Another possibility, for factor-augmenting technical progress, is that stable factor shares hold if the bias in technical change exactly offsets that of capital deepening. In this case, factor shares are stable independently of the value of the substitution elasticity.

(iv). More intriguingly, however, and independent from the size of σ, Θ would remain broadly constant outside the balanced growth path if r_t "absorbs" some of the trend in capital augmentation. This, though, violates our priors that the real interest rate is stable.[5] However, we can show that this trend absorption need only be modest.[6] If the user cost only partially absorbs the capital-augmenting technical progress, there will also be trends in the factor income shares, but these may be weak when coupled with a moderate pace of capital augmentation.[7] Hence, the broad stability of factor income shares is not a sufficient condition for the correctness of either Cobb-Douglas or Harrod neutrality.

We have seen that the assumption of Hicks neutrality can bias σ towards unity. We can also show that quite generally (although not universally) the Harrod-neutral specification can result in σ estimates that are either upwards or downwards biased when the true DGP contains capital-augmenting technical progress.

Assume the *lhs* of equation (4) below corresponds to the "true" DGP for the observed capital income share and the *rhs* corresponds to the mis-specified Harrod-neutral (h) version:

$$\pi_0 \left(\frac{\Gamma_t^K K_t}{K_0 Y_t} \right)^{\frac{\sigma-1}{\sigma}} = \pi_0 \left(\frac{K_t}{K_0 Y_t} \right)^{\frac{\widehat{\sigma}^h - 1}{\widehat{\sigma}^h}} \tag{4}$$

Taking logs and rearranging,

$$\frac{\sigma - 1}{\sigma} \log \Gamma_t^K = \frac{\widehat{\sigma}^h - \sigma}{\widehat{\sigma}^h \sigma} \log \left(\frac{K_t}{K_0 Y_t} \right) \tag{5}$$

In the true data, $\frac{K_t}{K_0 Y_t} = \left(\Gamma_t^K\right)^{\sigma-1} \left(\frac{r_0}{r_t}\right)^{\sigma}$. Assume $r_t = r_0 \left(\Gamma_t^K\right)^{\alpha}$, $\alpha \in (0,1]$ which implies that the real user cost *partly* absorbs the trend in capital-augmenting technology. It can be shown that with values of $\alpha > \frac{\sigma-1}{\sigma}$, the negative trend in the capital-output ratio corresponds to the positive trend of Γ_t^K. When this condition holds, then in the interval $\alpha \in (0,1]$, $\widehat{\sigma}^h > \sigma$ and with $\sigma > 1$, in turn, $\widehat{\sigma}^h < \sigma$. However, when $\alpha = 0$ and $\sigma > 1$, then the capital-output ratio has a positive trend and $\widehat{\sigma}^h > \sigma > 1$.

[5] However, rather than exhibiting global stability, real interest rates are commonly thought of as regime-wise stationary. Also, depreciation rates (another component of the user cost) have trended upwards over this sample – see Whelan (2002). This is compatible with the commonly-held view that the share of equipment in capital has increased while the share of structures has decreased and hence investment is characterized by shorter mean lives.

[6] Assuming capital augmenting-technical progress is 0.5% annually and even where that is fully absorbed by the real user cost, then the latter would rise from, for instance, 0.05 to 0.064 within 50 years.

[7] Jones (2003) also reports evidence showing capital shares for OECD countries frequently exhibit large variation and medium-run trends. These trends are certainly relevant for typical sample sizes. See also McAdam and Willman (2013).

Mis-specified technical change: TFP calculations

The calculation of TFP is a key application of production function estimates. Predicated on Cobb Douglas, TFP calculations are invariably derived imposing Hicks Neutrality (the "Solow Residual"). However, even if estimates of the *size* of TFP growth are robust to mis-specification, an accurate decomposition of TFP growth offers insights on the mechanisms underlining economic performance and may usefully inform policy.

An exact (or *residual*) method to calculate the contribution of log(TFP) to output is given by,

$$\log\left[\frac{F\left(\Gamma_t^K K_t,\, \Gamma_t^N N_t\right)}{F\left(\Gamma_0^K K_t,\, \Gamma_0^N N_t\right)}\right] = \frac{\sigma}{\sigma-1}\log\left[\frac{\pi_0\left(\frac{\Gamma_t^K K_t}{\Gamma_0^K K_0}\right)^{\frac{\sigma-1}{\sigma}} + (1-\pi_0)\left(\frac{\Gamma_t^N N_t}{\Gamma_0^N N_0}\right)^{\frac{\sigma-1}{\sigma}}}{\pi_0\left(\frac{K_t}{K_0}\right)^{\frac{\sigma-1}{\sigma}} + (1-\pi_0)\left(\frac{N_t}{N_0}\right)^{\frac{\sigma-1}{\sigma}}}\right] \tag{6}$$

For illustrative purposes, it is also useful to present a closed-form approximation for log(TFP) separable from factor inputs. We follow Kmenta (1967) and Klump et al. (2007), by applying an expansion of the normalized log CES production function (1) around $\sigma = 1$:

$$\begin{aligned} y_t &= \pi_0 k_t + a k_t^2 \\ &\quad \underbrace{+\,\pi_0\left[1+\frac{2a}{\pi_0}k_t\right]\gamma_K\cdot\widetilde{t} + (1-\pi_0)\left[1-\frac{2a}{(1-\pi_0)}k_t\right]\gamma_N\cdot\widetilde{t} + a\left[\gamma_K-\gamma_N\right]^2\cdot\widetilde{t}^2}_{\Phi=log\ (TFP)} \end{aligned} \tag{7}$$

where $\widetilde{t} = t - t_0$, $y_t = \log[(Y_t/Y_0)\,/\,(N_t/N_0)]$, $k_t = \log[(K_t/K_0)\,/\,(N_t/N_0)]$, and where $a = \frac{(\sigma-1)\pi_0(1-\pi_0)}{2\sigma}$.

Equation (7) shows that the output-labor ratio can be decomposed into (linear and quadratic) capital deepening and technical change weighted by factor shares and the substitution elasticity – where $sgn\,(a) = sgn\,(\sigma-1)$ and $\lim\limits_{\sigma\in[0,\infty)} a \in \left[-\infty, \frac{1}{2}\pi_0(1-\pi_0)\right)$. In addition, (7) shows that, when $\sigma \neq 1$ and $\gamma_K \neq \gamma_N > 0$, additional (quadratic) curvature is introduced into the production function: ak_t^2 and $a\left[\gamma_K-\gamma_N\right]^2\cdot\widetilde{t}^2$.

The effect of capital deepening on $\log(TFP)$ – given by $2a\widetilde{t}\,(\gamma_K-\gamma_N)$ – switches sign depending on whether factors are gross substitutes or complements. However, although the transmission of individual technology changes to TFP is also a function of σ, generally its sign (and, in particular, the importance of gross substitutes or complements) is ambiguous.[8]

[8] Except in two cases, when $\gamma_K - \gamma_N > 0$:

The effect of σ on TFP through capital deepening can be given an economic interpretation, though. When $\sigma \neq 1$, capital deepening will be biased in favor of one factor of production (changing its income share). Hence, with factor augmenting technical change, an acceleration of capital deepening changes the estimated TFP growth simply because technical progress is biased in favor of one of the factors. If, for instance, $\sigma < 1$, capital deepening would increase the labor share. If $(\gamma_K - \gamma_N) < 0$, capital deepening would lead to an acceleration of the estimated TFP growth.

The expressions for log(TFP) for the restricted neutrality cases are:[9]

$$\text{Harrod} : (1-\pi_0)\left[1 - \frac{2a}{(1-\pi_0)}k_t\right]\gamma_N \cdot \widetilde{t} + a\gamma_N^2 \cdot \widetilde{t}^2 \tag{8}$$

$$\text{Solow} : \pi_0\left[1 + \frac{2a}{\pi_0}k_t\right]\gamma_K \cdot \widetilde{t} + a\gamma_K^2 \cdot \widetilde{t}^2 \tag{9}$$

$$\text{Hicks} : \gamma \cdot \widetilde{t}, \text{ where } \gamma = \gamma_K = \gamma_N \tag{10}$$

The comparisons of (7) with variants (8)-(10) are self evident. For instance, in the Hicks case all improvements in TFP would be attributed to a single factor-neutral component, γ, excluding also any role for capital deepening.

For values of K_t and N_t close to their normalization points, $k_t \approx 0$, one can also obtain two simpler approximation for log(TFP):

$$\Phi^{Simple} = \pi_0\gamma_K \cdot \widetilde{t} + (1-\pi_0)\gamma_N \cdot \widetilde{t} + a\left[\gamma_K - \gamma_N\right]^2 \cdot \widetilde{t}^2 \tag{11}$$

$$\Phi^{LinearWeight} = \pi_0\gamma_K \cdot \widetilde{t} + (1-\pi_0)\gamma_N \cdot \widetilde{t} \tag{12}$$

The first abstracts from capital deepening. This may be considered informative regarding the contribution of capital deepening in TFP estimates based on (6) and (7) - especially so given the rapid capital deepening in the US towards the end of our sample. The second form, which is a simple linear weight of the two constant progress terms, discards all nonlinearities in TFP.

Although all cases coincide at the point of normalization, equation (11) by excluding capital deepening, runs the risk that the nonlinearity in the TFP is not correctly captured. For instance, if the economy is characterized by Harrod neutrality, Φ^{Simple} implies the wrong sign for the quadratic effect term (being positive rather than negative).[10]

$$\frac{\partial\Phi}{\partial\gamma_N|_{\sigma<1,\widetilde{t},k>0}} = (1-\pi_0)\,\widetilde{t}\left[\left\{1 - \frac{k_t\pi_0(\sigma-1)}{\sigma}\right\} - (\sigma-1)(\gamma_K-\gamma_N)\,\widetilde{t}\right] > 0,$$
$$\frac{\partial\Phi}{\partial\gamma_K|_{\sigma>1,\widetilde{t},k>0}} = \widetilde{t}\left[\pi_0\left\{1 + \frac{k_t(1-\pi_0)(\sigma-1)}{\sigma}\right\} + (1-\pi_0)(\sigma-1)(\gamma_K-\gamma_N)\,\widetilde{t}\right] > 0.$$

[9] Individual technical change cannot be identified in the Cobb-Douglas case.

[10] In the Harrod neutral case $k_t = \gamma_N \cdot \widetilde{t}$. Substituting this into (8) results in the following form of

IV The specification bias: simulation evidence

Following our earlier discussion, we now use a simulation exercise for a variety of parameter values of the supply side to quantitatively analyze the bias arising from mis-specified technical progress. We first simulate a consistent DGP for factor inputs, output, and factor payments, and then estimate the relevant parameters using the normalized system approach imposing particular forms of factor neutrality. The simulation follows León-Ledesma et al. (2010), but differs in terms of the stochastic process for factor inputs and, crucially, the way the growth of the capital stock is specified. This will precisely allow us to focus on questions of whether the simulated data is plausible in terms of balanced or near balanced growth trajectories, which is of special relevance in our context.[11]

The normalized system estimator of the parameters consists of the joint estimation of (log-version of) CES function (1) and the first order conditions for K and N. Normalization allows us to fix parameter π_0 to its observed value (capital income share in the baseline period) also simplifying the estimation problem. The 3-equation system of equations is then jointly estimated using a Nonlinear SUR system estimator (which we also use, among several alternative methods, for estimation with US data in section V).[12] In this case, of course, within a system setting, consistent cross-equation parameter restrictions are imposed.

The simulation experiment

We generate data in a consistent way corresponding to a particular evolution of factor inputs, technical progress and output. This Monte Carlo (MC) data is estimated under both correctly specified and mis-specified systems.

We draw M simulated stochastic processes of sample size T for labor, capital, labor- and capital-augmenting technology. Using these, we then derive "potential" or "equilibrium" output (Y_t^*), observed output (Y_t) and real factor payments (w_t and r_t), for a range of parameter values and shock variances. The simulated system is consistent with the normalized approach, so that we ensure our parameters are *deep*, i.e. can be given an economic interpretation and are not a combination of other

the log(TFP): $\pi_0\gamma_K \cdot \widetilde{t} + (1-\pi_0)\,\gamma_N \cdot \widetilde{t} - a\gamma_N^2 \cdot \widetilde{t}^2$ and hence Φ^{Simple} implies the wrong sign for the quadratic term.

[11] We do not focus here on comparison of estimation methods as in León-Ledesma et al. (2010), but on model (mis-) specification and in which direction it affects estimated parameters.

[12] We also considered GMM, 3SLS, and FIML estimators that take into account potential endogeneity bias, but the results remained very similar and are not reported here. In the empirical section, however, we show all these methods.

parameters.

We now describe the full DGP for the MC simulations. Capital and labor evolve as non-stationary stochastic processes with drift:

$$K_t = e^{\left(\kappa + \ln K_{t-1} + \varepsilon_t^K\right)}\ ,\ N_t = e^{\left(\eta + \ln N_{t-1} + \varepsilon_t^N\right)} \tag{13}$$

where κ and η are the drift terms. The initial values are $N_0 = 1$, and $K_0 = \pi_0/r_0$, with the real user cost at $r_0 = 0.05$.[13]

The technical progress functions, as described before, are also assumed to be exponential with a deterministic and stochastic component (around a suitable normalization point):

$$\Gamma_t^K = \Gamma_0^K e^{\left(\gamma_K \widetilde{t} + \varepsilon_t^{\Gamma^K}\right)},\ \ \Gamma_t^N = \Gamma_0^N e^{\left(\gamma_N \widetilde{t} + \varepsilon_t^{\Gamma^N}\right)} \tag{14}$$

where Γ_0^K and Γ_0^N are initial values for technology which we also set to unity.

We then obtain equilibrium output from the normalized CES function:

$$Y_t^* = Y_0^* \left[\pi_0 \left(\frac{K_t}{K_0} e^{\left(\gamma_K \widetilde{t} + \varepsilon_t^{\Gamma^K}\right)} \right)^{\frac{\sigma-1}{\sigma}} + (1-\pi_0) \left(\frac{N_t}{N_0} e^{\left(\gamma_N \widetilde{t} + \varepsilon_t^{\Gamma^N}\right)} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{15}$$

with $Y_0^* = 1$. This "equilibrium" output is then used to derive the real factor payments from the FOCs, to which we add a multiplicative shock:

$$r_t = \frac{\partial Y_t^*}{\partial K_t} = \pi_0 \left(\frac{Y_0^*}{K_0} e^{\left(\gamma_K \widetilde{t} + \varepsilon_t^{\Gamma^K}\right)} \right)^{\frac{\sigma-1}{\sigma}} \left(\frac{Y_t^*}{K_t} \right)^{\frac{1}{\sigma}} e^{\varepsilon_t^r} \tag{16}$$

$$w_t = \frac{\partial Y_t^*}{\partial N_t} = (1-\pi_0) \left(\frac{Y_0^*}{N_0} e^{\left(\gamma_N \widetilde{t} + \varepsilon_t^{\Gamma^N}\right)} \right)^{\frac{\sigma-1}{\sigma}} \left(\frac{Y_t^*}{N_t} \right)^{\frac{1}{\sigma}} e^{\varepsilon_t^w} \tag{17}$$

Equations (16) and (17) imply that real factor returns equal their marginal product times a multiplicative shock that temporarily deviate factor payments from equilibrium. All shocks are assumed normally distributed iid: $\varepsilon_t^{\Lambda} \overset{iid}{\sim} N\left(0, \sigma_{\varepsilon^{\Lambda}}\right), \Lambda = [K, N, \Gamma^K, \Gamma^N, r, w]$. These shocks are to be interpreted as unexpected factor market shocks that lead to temporary deviations between marginal products and factor prices. These unobserved shocks do not enter the first order conditions for profit maximization and are hence uncorrelated with factor demands.

[13]For all the experiments we also simulated K_t and N_t such that they displayed deterministic rather than stochastic trends. The main conclusions of the analysis did not change and the results are available on request. Also, initial values for r_0 and K_0 do not affect the results if the system is appropriately normalized.

Because we need to ensure that our artificial data is consistent with national accounts identities, we then obtain the "observed" output series using the identity:

$$Y_t \equiv r_t K_t + w_t N_t \tag{18}$$

Observed and equilibrium outputs differ because of these unobserved factor market shocks introducing a temporary wedge between factor prices and marginal products. We use the "observed" output series for estimation purposes. This ensures that, regardless of the shocks, factor shares sum to unity, which has to be the case in this artificial setting with absent markups.

Hence, the experiment consists of, first, simulating a time series of sample size T for factor inputs, technical progress, and equilibrium output. Second, from these we obtain factor payments and observed output. Finally, we estimate the normalized system, (15)-(17), imposing Hicks-, Harrod- and Solow neutrality in technical progress. We repeat these steps M times and analyze the possible biases arising from mis-specification by looking at the difference between the true and estimated σ.[14]

Table 1 lists the parameters used to generate the simulated series. We fixed the distribution parameter to 0.4.[15] The substitution elasticity is set to a neighborhood around Cobb-Douglas (0.9) and 0.9±0.4 (thus accommodating gross substitute and complements). The labor supply drift (η) is set to 1.5% per year. The values for capital stock drift (κ) will be discussed further below. We use values for technical progress assuming a plausible summation of 2% per year; $\gamma_N = 2\%$ and $\gamma_K = 0\%$ (Harrod-neutral case); $\gamma_N = 0\%$, $\gamma_K = 2\%$ (Solow neutral); and $\gamma_N = \gamma_K = \gamma = 1\%$ (Hicks-neutral). Finally, we have two cases where technical progress is of the general factor augmenting form.

The standard errors of the shocks are chosen so that they also generate series with realistic behavior. We chose a value of 0.05 for the capital and labor stochastic shocks. For the technical-progress parameters, we used a value of 0.01 when the technical progress parameter is set to zero, so that the stochastic component of technical progress does not dominate. When technical progress exceeds zero we used a value of 0.05 so when technical progress is present it is also more volatile.[16] Finally, for shocks

[14] Note, a slightly different way of setting up the Monte Carlo would have been to assume that a constant fraction of output is devoted to investment, with the capital stock then being determined by this investment as well as the assumed rate of depreciation. But this makes controlling and isolating the different growth rates of technical progress (which is a key component of our exercise) somewhat less transparent, and essentially imposes balanced growth from the outset.

[15] In practice, setting different values for π_0 did not affect the results.

[16] For robustness purposes, we also replicated the results assuming no shock when technical progress is zero and also equal shocks for both components. The results were not affected by

to factor payments, we used the standard deviation of the de-trended real wages and the standard deviation of demeaned user cost of capital for the US economy. These take values of 0.05 and 0.1 respectively, reflecting the larger volatility of the real user cost.

We used a sample size $T = 50$ (years).[17] Also, the nonlinear system estimator used requires initial guesses for the parameters, which we set to their true value following Thursby (1980).[18]

The choice of the drift parameter for capital, κ, is important given our emphasis on settings where the economy does not deviate in an evident way from the case of stable factor income shares. Hence, κ is chosen such that we exclude unrealistic income share trends. We can do this by looking again at the expression for the capital-to-labor income share under competitive profit maximization,

$$\Theta_t = \frac{r_t K_t}{w_t N_t} = \frac{\pi_0}{1-\pi_0}\left(\frac{\Gamma_t^K K_t/K_0}{\Gamma_t^N N_t/N_0}\right)^{\frac{\sigma-1}{\sigma}}$$

Accordingly, if $\sigma \neq 1$, capital- and labor-augmenting technical change can lead to ever increasing or decreasing factor shares for given factor proportions. Hence, for given rates of technical progress, to obtain approximately constant shares, we set the drift in K in such a way that we avoid any counter-factual trends in shares.

One simple mechanism to achieve this, following our earlier discussion, is to allow r to absorb some fraction, α, of the trend in capital augmentation (assuming $\Gamma_0^K = \Gamma_0^N = 1$). Hence, we use the following deterministic rule for r:

$$r_t^{det} = r_0 e^{\alpha\left(\gamma_K \cdot \tilde{t}\right)} \tag{19}$$

Now with (19), the FOC of capital results in the following relation for the capital income share,

$$\frac{r_t^{det} K_t}{Y_t} = \pi_0 e^{(1-\alpha)(\sigma-1)\left(\gamma_K \cdot \tilde{t}\right)} \tag{20}$$

Equation (20) shows that, with the constant user cost, i.e. when $\alpha = 0$, the capital augmenting technical change coupled with non-unitary substitution elasticity results in continuously changing factor income shares. However, with $\alpha \to 1$ the

these changes.

[17]Using values of 100 and 30 led to very similar results, although, as expected, the range of estimated values for the parameters increased as we decreased the sample size.

[18]This facilitates comparisons across specifications since we eliminate the effect of arbitrary starting values on results.

larger part of this trend is absorbed by the trend in the user cost. With $\alpha = 1$ factor income shares remain constant independently from the sizes of σ and γ_K. Hence, we can choose α in the unit interval so that factor shares and the real user cost do not display trends that are grossly counter factual.

Once α is chosen, for given technology parameters, we obtain r_t^{det} from (19). Given an exogenous law of motion for N, the CES function and (20) solve for K and Y. Using the value of K from this recursive system, we obtain the *average* rate of growth of K that we then use as the value for κ in our stochastic DGP. This is the value compatible with factor shares and real interest rates that do not display counter-factual trends. Given that parameter α controls the rate of change of r_t^{det}, a sufficiently small value can be set to mimic empirically-relevant paths for r and hence K/Y and Θ. In our experiments, we set $\alpha = 0.5$.

The functional construct of (19) is not without an empirical counterpart. As we know, the real user cost comprises the nominal interest rate (i.e., the risk-free government bond rate or firms' market rates), inflation, capital depreciation, taxes, capital gains etc. All these are time-varying (Figure 5 plots our measure of the user cost series for the US). Thus, if there is technical change which is not solely Harrod neutral alongside approximately constant factor shares, factor payments must be compensating.

Simulation results

Median estimates

Tables 2 to 4 report the Monte Carlo results when the data are generated according to the $\{\gamma_K\,,\ \gamma_N\,\}$ and $\{\sigma\}$ combinations given in Table 1 but then estimated for the respective cases of Hicks-, Harrod- and Solow neutrality. In the tables, we report the median parameter estimates across the 5,000 draws for the substitution elasticity (and its percentiles) and γ_i.

Where the imposed technical change corresponds to the true DGP (labeled "benchmark" in the tables), the parameters are very precisely estimated, reflecting the power of the normalized system. However, in non-benchmark gross complements cases (i.e., the first two columns in each table), systematic upwards bias is always found, i.e.,:

$$\sigma^m - \sigma\,\{0.5,\ 0.9\} > 0$$

The gross-substitute, non-benchmarks cases are less clear cut. Whilst, in all but one case (relating to Harrod neutrality, Table 3) a gross substitutes production function is correctly identified, in all cases but two (corresponding to Hicks neutrality) there is a downward bias:

$$\sigma^m - \sigma\{1.3\} < 0, \text{ with } \sigma^m \approx 1$$

The technical progress parameters also display substantial biases in non-benchmark cases. These biases tend to be upwards for the Solow neutral specification, and downwards for the other two cases.

Distributions

The distribution of the substitution elasticities across the 5,000 draws shed further light on these results (Figures 1, 2 and 3). Regarding the $\sigma = 0.5$ case, we see that the general factor augmenting specification is always tightly distributed around the true value of the substitution elasticity. The Solow neutral specification, though, yields a bimodal distribution for the two cases in which technical progress is net labor-augmenting. To a smaller degree, the Harrod-neutral specification also shows bimodality in two cases. The Hicks neutral specification is almost flat except in the benchmark case of Hicks-neutral augmentation. This is reflected in Table 2, where the 10% and 90% percentiles show a very considerable variation. The distributions also tend to be more skewed when the specified model differs from the true DGP.

The $\sigma = 0.9$ case is interesting given its proximity to Cobb-Douglas, and thus the heightened relevance of the issues raised in Section III. Note that the densities are now more symmetric and display limited dispersion, and in several cases there is a clear bias towards Cobb-Douglas at the median. Consistent with the $\sigma = 0.5$ case above, most median estimates exhibit upward biases. As discussed earlier, a unitary substitution elasticity is a strong attractor: pulling estimates to the log-linear form captures the broadly balanced growth characteristics of the simulated data minimizing the cost of the imprecise technical change component. Recalling approximation (7), $\widehat{\sigma} \to 1$, neutralizes the effect of quadratic curvature in capital deepening and technical bias, and minimizes the weight given to the individual technical progress components. Furthermore, bi- or multi-modality is more severe than in the $\sigma = 0.5$ (or indeed $\sigma = 1.3$) case, even so for the cases where both forms of technical change are permitted; thus, even the factor-augmenting specification shows a (second) peak around unity in all cases.

For $\sigma = 1.3$ the distributions are, by contrast, much flatter. The factor augmenting specification, despite capturing very well the true values of σ, also tend to display a small local maximum around a value of one. Overall there is always one specification for which there is a clear bias towards unity. The only case of an upwards bias happens with the Harrod-neutral specification for the pure capital augmenting case.

Our simulation exercises were necessarily stylized. In particular, we analyzed an environment of balanced or near balanced growth. This has several advantages. First, it corresponds to situation common to many developed countries (over reasonably-sized samples). Second, it places our exercises within a familiar context, making the interpretation and motivation of results more transparent. However, third, it in fact makes for a particularly challenging exercise since estimates – framed in the neighborhood of a balanced growth path – may degenerate to unitary elasticities and overlook or strongly bias the nature of technical change. Our next step is to analyze how these potential biases affect estimates of the supply-side parameters and estimates of TFP growth for the US economy.

V Supply Side Estimation of the US Economy

Data

We use the U.S. annual national income and product accounts (NIPA) data released by the Bureau of Economy analysis (BEA) for the private non-residential sector over the period 1952 to 2009. Output (at current and constant prices) is evaluated at factor cost, i.e. net of indirect taxes minus subsidies. Hence, current price private non-residential output equals gross domestic product minus taxes on production and imports less subsidies, general government value added and gross housing value added. In calculating the (chained dollars) constant price output the constant price gross domestic product is scaled down in proportion of to the base year's (2005) indirect tax content, of which constant price general government and gross housing value added are subtracted.

Employment is defined as the sum of self-employed persons and the private sector full-time equivalent employees (both from NIPA tables). NIPA tables do not report the income of proprietors (self-employed) divided into labor and capital income. Therefore, in calculating labor income we follow a common practice (e.g., Klump et al. (2007)), use the private sector compensation of employees as a shadow price of labor of self-employed workers. Accordingly, total labor income equals the private sector compensation of employees scaled up by the labor share of self-employed workers.

For capital we use the quantity index of net stock of non-residential private capital from the BEA fixed asset tables. Capital income and the implied measure of the user cost are calculated from the accounting identity of non-residential private sector conditional for an assumed 10% markup (which is a common benchmark in macroeconomic models, e.g., Clarida et al. (1999)).[19]

Figure 4 presents some variables of interest. Against the balanced-growth path hypothesis the capital-output ratio appears to show a trend over the sample period. An ADF test does not reject the null of non stationarity of the capital-output ratio. This trend expresses itself also in a trend difference between average labor productivity (output-labor ratio) and capital intensity (capital-labor ratio). The share of labor income shows sizeable annual variation. Although a sort of inverted U (or double U) trend profile cannot be observed, an ADF test rejects the null of non stationarity of the labor income share. Over the whole sample, the trends of real wages and labor productivity are quite close to each other although, most of the time, the real wage index exceeds the labor productivity index. The real user cost looks stationary until early 1990s but thereafter is shows a clear upward trend reflecting the return of the labor (and capital) income share back to the level where it was in the early part of the sample period. Hence, in terms of an ADF test the real user cost is non-stationary. We also discover that, in line with our discussion in section 4.1, the actual data evolution of the real user cost contributes towards retaining the stationarity of factor income shares.

Specification

Given the practical existence of a markup over factor costs in the data, the estimated model includes an extra parameter $\mu = 0.1$. This captures an average markup which, consistent with our data construction, we restrict to 10%.

Also, with real data, to diminish the size of stochastic component in the point of normalization we prefer to define the normalization point in terms of sample averages (geometric averages for growing variables and arithmetic ones otherwise). The non-linearity of the CES function, in turn, implies that the sample average of production need not exactly coincide with the level of production implied by the production function with sample averages of the right hand variables. Following Klump et al. (2007), we therefore introduce an additional parameter ζ whose expected value is around unity. Hence, we define $Y_0 = \zeta\bar{Y},\ K_0 = \bar{K},\ N_0 = \bar{N}$; $t_0 = \bar{t}$ and $\pi_0 = \overline{\pi}$ where the bar refers to the appropriate type of sample average. The estimated system, allowing for factor augmentation, is then,

[19]The benefit of this approach is that we do not have to explicitly calculate the user cost, which has long been recognized as being a complex exercise and with scope for large measurement error. E.g., Jorgensen and Yun (1991). However, for robustness, we also used a user cost calculation and let the average markup to be freely estimated. This did not change substantially the results.

$$\log r = \log\left(\frac{\bar{\pi}}{1+\mu}\frac{\zeta\bar{Y}}{\bar{K}}\right) + \frac{1}{\sigma}\log\left[\frac{Y/(\zeta\bar{Y})}{K/\bar{K}}\right] + \frac{\sigma-1}{\sigma}\gamma_K(t-\bar{t}) \tag{21}$$

$$\log(w) = \log\left(\frac{(1-\bar{\pi})}{1+\mu}\frac{\zeta\bar{Y}}{\bar{N}}\right) + \frac{1}{\sigma}\log\left(\frac{Y/(\zeta\bar{Y})}{N/\bar{N}}\right) + \frac{\sigma-1}{\sigma}\gamma_N(t-\bar{t}) \tag{22}$$

$$\log\left(\frac{Y}{\bar{Y}}\right) = \log\zeta + \frac{\sigma}{\sigma-1}\log\left[\bar{\pi}\left(\frac{e^{\gamma_K(t-\bar{t})}K}{\bar{K}}\right)^{\frac{\sigma-1}{\sigma}} + (1-\bar{\pi})\left(\frac{e^{\gamma_N(t-\bar{t})}N}{\bar{N}}\right)^{\frac{\sigma-1}{\sigma}}\right] \tag{23}$$

For the estimation of the system we fix parameter $\bar{\pi}$ to its sample average, which is one of the empirical advantages of normalization. We also obtained the results estimating $\bar{\pi}$ freely, but it made minimal difference to the other relevant parameters.

The system is then estimated using a variety of methods to account for cross-equation error correlation and regressor endogeneity. We used Nonlinear Seemingly Unrelated Regression (NLSUR) methods, Nonlinear 3-Stage Least Squares (NL3SLS), Fully Information Maximum Likelihood (FIML), and Generalized Method of Moments (GMM) methods. All of these four estimations are implemented accounting for cross-equation parameter restrictions.

Estimation results

The results of the four estimation methods for the factor augmenting specification of the system are reported in Table 5.[20] Table 6 reports the results of the Hicks-, Harrod-, and Solow-neutral specifications for the case of the NLSUR estimator. We report only this case to save space as the rest of the estimation methods encountered essentially the same patterns. Table 5 also reports p-values for tests of the null hypothesis of a unitary σ. The following rows display p-values for Wald tests of restrictions on technical progress to statistically discriminate between the different nested specifications. We also report ADF and Phillips-Perron (PP) unit root residual tests.[21] Given that we do not know the distribution of the statistic under the null, we use bootstrapped p-values following and Chang and Park (2003). For the

[20]Note we conducted a number of robustness and sensitivity exercises. Initial conditions of all parameters were varied around plausible supports with practically no impact on final results in every case. Plus, for the HAC standard errors we tried both Bartlett and Quadratic kernel options and various choices for bandwidth selections, again with negligible difference on results. Details available.

[21]The PP tests are robust to serial correlation by using a heteroskedasticity- and autocorrelation-consistent covariance matrix estimator. We used the Z_ρ version of this test. We also used KPSS residual tests for stationarity. The results (not reported) supported stationarity in the majority of the cases.

instruments-based estimators, we used first lags of the log of the user cost and real wage, normalized employment, capital stock and log-output, and the time trend.

The results in Table 5 show similar results for the estimated value of σ that ranges from 0.4 (FIML) to 0.7 (3SLS). Manifestly, these estimates are well below and significantly different from unity. Estimates of technical progress coefficients are very stable across estimation methods. Labor-augmenting technical progress is estimated to be around 2% per year, whereas capital-augmenting technical progress is 0.4% per year in most of the cases.[22] However, we can appreciate the large value for the Solow-augmenting specification: since capital attracts a below half weight in capital share the value of γ_K must be suitably high to match movements in TFP. Overall, technical progress is net labor-saving, but with non-negligible capital-augmenting technical progress. The scale parameter, ζ, is practically indistinguishable from unity as expected. In all cases, the null of non-stationarity for the residuals of each equation is rejected according to the bootstrapped p-values.

Regarding other specifications, we see that the σ estimates are substantially different from those obtained with general factor augmentation. The point estimate of σ with Hicks and Solow neutrality is indistinguishable from one. The Harrod-neutral specification also yields a higher estimate for σ, although still significantly below unity. These findings are consistent with those from the simulation experiment and our previous analytical results.

The Hicks specification biases the estimate of the substitution elasticity towards one. The Solow neutral specification also leads to a sharp bias towards Cobb-Douglas. Again, looking back at the results in Table 4 this is consistent with our simulations, which showed that the more the DGP deviates from Solow neutrality, the stronger the bias towards unity. In the case of the Harrod-neutral specification which, together with Hicks-neutral, is most commonly used for estimation, we observe that the results are biased upwards. This bias is consistent with that found in the simulation experi-

[22]One of our referees suggested that we consider the capital input series used by the BLS in its multifactor productivity release, rather than the BEA non-residential private capital stock. Doing so, generates relatively similar results: we find an elasticity of 0.71 rather than 0.69, and a labour augmenting term of 0.02 which is as before. The difference lies in the rate of growth of capital augmentation. We find this to be 0.004; under this alternative capital stock definition now we find it to be -0.011. The difference can be related to the greater weight of capital equipment (particularly IT equipment) and smaller weight of long-lived structures in the BLS series compared to ours. This resulted in (relative to our series) a faster growth rate of capital stock, especially in the later part of the sample and an associated slower development of capital rental price. A negative rate of technical progress therefore is by no means un-interpretable in the context of this alternative series. Given the reduction in user cost, capital becomes - to use the language of the directed technical change literature - the less scarce factor and accordingly, with such factor abundance, leaves firms with little incentive to focus technical improvements on the abundant factor.

ment with positive values for capital-augmenting technical progress. As discussed in section III, the Harrod-neutral specification could results in upward biases if the true σ is below one.

Finally, the Wald tests for the restrictions implied by specific forms of factor augmentation, always reject the restrictions in favor of the general factor augmenting specification. Hence, our results support the use of a more general specification for technical progress and confirm our claim that mis-specification of technical progress can lead to important biases in the estimated substitution elasticity.

Figure 5 plots the model residuals for the four specifications for the NLSUR estimator. For the user cost, the four models yield similar fit except towards the end of the sample where both the factor augmenting and the Harrod-neutral specifications capture better the increase in the user cost. Importantly, the fit for output appears to be almost identical for the four specifications. The main difference emerges in the way the models fit wages, with the factor augmenting specification displaying larger fluctuations.[23] Of course, even if the three models yield similar fit for variables such as output, the implications of the different estimates of the substitution elasticity and technical progress to explain the evolution of factor shares are still different. As we will now see this is also true for TFP growth.

TFP estimates

We obtained estimates of TFP *growth* arising from (6) and the simplified approximations (11) and (12).[24] Figure 6 plots the NLSUR estimates of TFP separately for each specification (alongside capital deepening). Results using the other estimators yielded similar conclusions. The Hicks-neutral specification, necessarily yields constant growth of TFP and, hence, is not plotted separately. The rest of specifications will always yield increasing or decreasing TFP *growth* (except when linear weight, (12), is used). This can be seen in expressions (7) and (11), whose rate of growth is going to be trended owing to the quadratic component. Whether the trend is positive or negative depends on parameter "a", whose sign is a function of whether $\sigma \gtrless 1$ (except in the Hicks case when the trend is zero).

The simple form excluding capital deepening applies wrong trends to the growth rate in TFP in the context of factor-augmenting and Harrod-neutral specifications.

[23]Interestingly, this is a result that Fisher et al. (1977) also obtained in a simulation experiment analyzing production function aggregation. Despite many specifications providing a good fit for output, wages proved much more sensitive to the estimated values of σ.

[24]The Kmenta approximations (7)-(10) and the exact residual method (6) yield practically identical TFP and are not reported for brevity.

Under the Solow neutral specification, however, it works quite satisfactorily. We may conclude that the inclusion of capital deepening is important to capture correctly nonlinearities in TFP growth rates. It is interesting to see that especially our favored factor-augmenting case implies an acceleration in TFP growth from the second half of the 1990s until the mid-2000s.[25] This is compatible with the then observed acceleration of productivity growth (e.g., Basu et al. (2003) and Jorgenson (2001)). The TFP growth spike at the end of the sample simply reflects the rapid cyclical drop in employment due to the financial crisis. Both the factor augmenting and Harrod neutral specifications display very similar TFP growth patterns. However, because of the lower estimate of σ, the residual based estimate in the factor-augmenting case displays more pronounced fluctuations and a sharper trend increase. From our perspective of specification bias, it is worth noting that the differences in annualized TFP growth towards the end of the sample are substantial.

Progress: what have we learnt?

Pulling together the salient points arising from the analytical, simulation, and empirical estimates, we can extract a series of important lessons about estimation and analysis of supply-side systems:

Implications of a priori choices on the nature of technical change

Estimation of the substitution elasticity can be substantially biased if the form of technical progress is mis-specified. For some parameter values, when factor shares are relatively constant, there could be an inherent bias towards Cobb-Douglas, but this is not the only possible direction of bias.

Our empirical results show that the estimated substitution elasticity tends to be significantly lower using a factor augmenting specification and is well below one. We were able to reject Hicks-, Harrod- and Solow-neutral specifications in favor of general factor augmentation with a non-negligible capital-augmenting component.

Beware Cobb-Douglas

Situations of near balanced growth may lead to estimation erroneously favoring the unitary elasticity case. This is clear in some cases such as Hicks Neutrality where

[25] This is consistent with the idea that investment in IT led to an economy-wide productivity increase. In our model, however, we do not separate types of capital and so cannot infer anything about the specific source of this acceleration. However, as far as this capital deepening is related with investment in new technologies, our results seem to support the contention that there was a productivity acceleration in the US from the mid-1990s until the early 2000s.

a unitary bias shrinks the importance of trended capital deepening. Similarly, when seen through the lens of the augmented Kmenta approximation, a unitary elasticity shrinks the impact of quadratic curvature in capital deepening and biased technical change. Furthermore, the MC distributions tended to show a separate mode for the unitary elasticity case, particularly if initial conditions were set within that neighborhood.

There is no simple solution to degenerate Cobb-Douglas estimates, other than some of the practices followed here: discriminating on the basis of global statistical criterion among competing specifications; varying initial conditions and checking for local maxima; inspecting the great ratios to check for stationarity; and hints in the data for the potential presence of capital-augmenting or non-constant technical progress components (e.g., see the discussion in Klump et al. (2007)).

Aggregate studies favoring Cobb-Douglas, though, are far rarer than its theoretical dominance might suggest.[26] But there is still a tendency in the literature to report high near-unity substitution elasticities and neglect the role of biases in technical change. Given how useful the analysis of biased technical change has proved (Acemoglu (2009)) in accounting for growth experiences, this is clearly an error of non-trivial proportion.

The fit of the production function vs. the fit of factor returns

Our empirical results implicitly make an important, even startling, point. The quite similar production-function residuals suggest that the goodness of fit of production functions appears relatively robust to mis-specified technical neutrality assumptions (an early indication of this was given by Willman (2002)). The reason is that mis-specification of technical change under a CES production function implies compensating bias in the estimate of the elasticity of substitution.

However, an important qualification (echoing that of Fisher et al. (1977)) is that using an "incorrect" production function may simply shift estimation failures elsewhere. In our case, this arose most clearly in factor returns equations where there is considerable variation in the fit across specifications.[27]

[26]See, for instance, Table 1 of León-Ledesma et al. (2010).

[27]This is what Christoffel et al. (2011) report for their macro-econometric forecasting and simulation model, the NAWM which employs an aggregate Cobb-Douglas production function: good forecasting performance for many real variables (including the output gap) but large, persistent errors in forecasting real wages and the labor share.

TFP growth

The dispersion of TFP estimates mirrors that of the real wage. Monitoring the level and sources of TFP growth is a key application of the production function literature and a key input into policy debates. Recalling Figure 4, we see an acceleration in US labor productivity from the mid-1990s until the mid-2000s driven by capital deepening in combination with technical change. And yet (Figure 6) these patterns are obscured under Harrod- and Solow-neutral specification – and disappear under Hicks-neutrality –.

There is an important lesson to be drawn here. Given the discussions in Sections II and V, we know that whether the substitution elasticity is above or below unity matters for the transmission of capital deepening and factor-augmenting technical change for TFP's evolution. Getting the substitution elasticity right is hence necessary to correctly estimate TFP growth.

VI Conclusions

Balanced growth requires stringent conditions on the structural parameters driving the production function and factor payments. Given that, we studied the effect of imposing specific forms of technical progress neutrality for estimates of key supply side parameters, such as the substitution elasticity.

Specifically, we studied how estimates of the elasticity of factor substitution and TFP growth are affected by imposing mis-specified a priori restrictions on the factor saving nature of technical change in a context where an economy may depart from a BGP. We showed analytically that, when the true nature of technical progress is factor-augmenting, imposing Hicks-neutrality leads to biases towards Cobb-Douglas. Imposing Harrod-neutrality would generally lead to upward biases in the estimated elasticity if the true elasticity is below unity and downward biases if above unity. Because TFP growth approximations from CES production function estimates depend on the substitution elasticity, these biases will also be reflected in biases in estimated TFP growth. We carried out an extensive simulation exercise that supports these conclusions and showed that the biases can be substantial in terms of magnitude.

We then estimated a CES supply side system for the US economy and found that many of the previous lessons found an echo in empirical estimates. Furthermore, we could reject the Hicks-, Harrod- and Solow-neutral specifications in favor of a general factor augmenting one. We found that capital-augmenting technical progress is non-negligible (0.4% per year). Importantly, the substitution elasticity is found to be substantially below one, emphatically rejecting Cobb-Douglas. We also provide evidence that the implied TFP growth estimates for the various specifications used is substantially different. Our work thus provides insights into production and supply-side estimation and design in balanced-growth based macroeconomic frameworks.

References

Acemoglu, D. (2007). Equilibrium bias in technology. *Econometrica*, 175:1371–1410.

Acemoglu, D. (2009). *Introduction to Modern Economic Growth.* MIT Press.

Antràs, P. (2004). Is the US Aggregate Production Function Cobb-Douglas? New Estimates of the Elasticity of Substitution. *Contributions to Macroeconomics*, 4:Article 4.

Arrow, K. J., Chenery, H. B., Minhas, B. S., and Solow, R. M. (1961). Capital-labor substitution and economic efficiency. *Review of Economics and Statistics*, 43(3):225–250.

Attfield, C. and Temple, J. R. W. (2010). Balanced growth and the great ratios: New evidence for the US and UK. *Journal of Macroeconomics*, 32:937–956.

Basu, S., Fernald, J. G., Oulton, N., and Srinivasan, S. (2003). The case of the missing productivity growth, or does information technology explain why productivity accelerated in the US and not in the U.K.? *NBER Macroeconomics Annual*, 18:9–63.

Berndt, E. R. (1976). Reconciling alternative estimates of the elasticity of substitution. *Review of Economics and Statistics*, 58(1):59–68.

Chang, Y. and Park, J. Y. (2003). A sieve bootstrap for the test of a unit root. *Journal of Time Series Analysis*, 24:379–400.

Chirinko, R. S. (2008). σ: The Long and Short of It. *Journal of Macroeconomics*, 30(2):671–686.

Christoffel, K., Coenen, G., and Warne, A. (2011). Forecasting with DSGE Models. In Clements, M. and Hendry, D., editors, *Oxford Handbook on Economic Forecasting.* Oxford University Press. Chapter 4, pp. 89-127.

Clarida, R., Gali, J., and Gertler, M. (1999). The science of monetary policy: A new keynesian perspective. *Journal of Economic Literature*, 37(4):1661–1707.

Fisher, F. M., Solow, R. M., and Kearl, J. M. (1977). Aggregate Production Functions: Some CES Experiments. *Review of Economic Studies*, 44(4):305–320.

Growiec, J. (2008). Production Functions and Distributions of Unit Factor Productivities: Uncovering the Link. *Economics Letters*, 101(1):87–90.

Jones, C. I. (2003). Growth, Capital Shares, and a New Perspective on Production Functions. Mimeo, University of California Berkeley.

Jones, C. I. (2005). The shape of production functions and the direction of technical change. *Quarterly Journal of Economics*, 120(2):517–549.

Jorgensen, D. and Yun, K.-Y. (1991). *Tax Reform and the Cost of Capital*. Oxford University Press.

Jorgenson, D. W. (2001). Information technology and the US economy. *American Economic Review*, 91:1–32.

Kaldor, N. (1961). Capital accumulation and economic growth. In Lutz, F. A. and Hague, D. C., editors, *The Theory of Capital*, pages 177–222. New York: St. Martin's Press.

Klump, R. and de La Grandville, O. (2000). Economic growth and the elasticity of substitution: Two theorems and some suggestions. *American Economic Review*, 90(1):282–291.

Klump, R., McAdam, P., and Willman, A. (2007). Factor Substitution and Factor Augmenting Technical Progress in the US. *Review of Economics and Statistics*, 89(1):183–92.

Klump, R., McAdam, P., and Willman, A. (2012). The Normalized CES Production Function: Theory And Empirics. *Journal of Economic Surveys*, 26(5):769–799.

Kmenta, J. (1967). On Estimation of the CES Production Function. *International Economic Review*, 8(2):180–189.

La Grandville, Olivier de. (1989). In Quest of the Slutsky Diamond. *American Economic Review*, 79(3):468–481.

León-Ledesma, M. A., McAdam, P., and Willman, A. (2010). Identifying the Elasticity of Substitution with Biased Technical Change. *American Economic Review*, 100(4):1330–1357.

León-Ledesma, M. A. and Satchi, M. (2011). The Choice of CES Production Techniques and Balanced Growth. Studies in Economics 1113, Department of Economics, University of Kent.

McAdam, P. and Willman, A. (2013). Medium Run Redux. *Macroeconomic Dynamics*, (forthcoming).

Thursby, J. (1980). Alternative CES estimation techniques. *Review of Economics and Statistics*, 62(2):259–299.

Uzawa, H. (1961). Neutral inventions and the stability of growth equilibrium. *Review of Economic Studies*, 28:117–124.

Willman, A. (2002). Euro area Production Function and Potential Output: A Supply Side System Approach. Working Paper No. 153, European Central Bank.

Tables and Figures

Table 1

Parameter values for the Monte Carlo

Parameter	Description	Values
π_0	Distribution parameter	0.4
σ	Substitution elasticity	0.5, 0.9, 1.3
γ_K	K-Augmenting Technical Progress*	0.00, 0.005, 0.01, 0.015, 0.02
γ_N	N-Augmenting Technical Progress*	0.02, 0.015, 0.01, 0.005, 0.00
η	Labor growth rate	0.015
κ	Capital growth rate	See text
$Y_0^* = N_0$	Normalization values for Y and N	1
K_0	Normalization value for K	π_0/r_0
r_0	Normalization value for the user cost	0.05
α	Capital Trend Absorption in r	0.5
$\sigma_{\varepsilon_t^N}$, $\sigma_{\varepsilon_t^K}$	Std. Error, N and K DGP Shock	0.05
$\sigma_{\varepsilon_t^{\Gamma^{K,N}}}$	Std. Error, N and K-Augmenting	0.01 for $\gamma_{K,N} = 0$;
	technical progress shock	0.05 for $\gamma_{K,N} \neq 0$
$\sigma_{\varepsilon_t^w}$	Std. Error, Real Wage shock	0.05
$\sigma_{\varepsilon_t^r}$	Std. Error, Real Interest Rate shock	0.1
T	Sample Size	50
M	Monte Carlo Draws	5,000

Notes: "*" $\gamma_N + \gamma_K = 0.02$.

Table 2

Monte Carlo results. Hicks-neutral specification

	$\sigma = 0.5$	$\sigma = 0.9$	$\sigma = 1.3$
		$\gamma_K = 0.00$, $\gamma_N = 0.02$	
σ^m	0.840	0.976	1.061
10% : 90%	0.656 : 1.316	0.894 : 1.099	0.897 : 1.1946
γ^m	0.012	0.012	0.012
		$\gamma_K = 0.005$, $\gamma_N = 0.015$	
σ^m	0.647	0.942	1.173
10% : 90%	0.484 : 0.848	0.840 : 1.065	1.008 : 1.360
γ^m	0.011	0.011	0.011
	Benchmark	$\gamma_K = \gamma_N = 0.01$	
σ^m	0.511	0.909	1.300
10% : 90%	0.465 : 0.585	0.806 : 1.042	1.091 : 1.549
γ^m	0.010	0.010	0.010
		$\gamma_K = 0.015$, $\gamma_N = 0.005$	
σ^m	0.614	0.901	1.424
10% : 90%	0.375 : 0.823	0.786 : 1.058	1.050 : 1.902
γ^m	0.008	0.009	0.009
		$\gamma_K = 0.02$, $\gamma_N = 0.00$	
σ^m	0.849	0.945	1.444
10% : 90%	0.491 : 1.460	0.782 : 1.097	0.909 : 2.400
γ^m	0.008	0.008	0.008

Notes: Superscript m denotes median values.

Table 3

Monte Carlo results. Harrod-neutral specification

	$\sigma = 0.5$	$\sigma = 0.9$	$\sigma = 1.3$
	Benchmark	$\gamma_K = 0.00$, $\gamma_N = 0.02$	
σ^m	0.511	0.900	1.288
10% : 90%	0.477 : 0.553	0.813 : 1.009	1.097 : 1.555
γ^m	0.020	0.020	0.020
		$\gamma_K = 0.005$, $\gamma_N = 0.015$	
σ^m	0.572	0.912	1.260
10% : 90%	0.432 : 0.722	0.810 : 1.020	1.026 : 1.666
γ^m	0.017	0.019	0.018
		$\gamma_K = \gamma_N = 0.01$	
σ^m	0.709	0.947	1.161
10% : 90%	0.426 : 0.904	0.815 : 1.068	0.917 : 1.829
γ^m	0.016	0.017	0.016
		$\gamma_K = 0.015$, $\gamma_N = 0.005$	
σ^m	0.850	0.987	1.023
10% : 90%	0.552 : 1.225	0.857 : 1.112	0.812 : 1.888
γ^m	0.015	0.015	0.015
		$\gamma_K = 0.02$, $\gamma_N = 0.00$	
σ^m	0.970	1.013	0.940
10% : 90%	0.750 : 1.997	0.922 : 1.154	0.756 : 1.776
γ^m	0.013	0.013	0.013

Table 4

Monte Carlo results. Solow-neutral specification

	$\sigma = 0.5$	$\sigma = 0.9$	$\sigma = 1.3$
		$\gamma_K = 0.00$, $\gamma_N = 0.02$	
σ^m	0.805	1.005	1.004
10% : 90%	0.710 : 0.966	0.962 : 1.061	0.929 : 1.062
γ^m	0.022	0.030	0.030
		$\gamma_K = 0.005$, $\gamma_N = 0.015$	
σ^m	0.799	0.994	1.039
10% : 90%	0.676 : 0.958	0.946 : 1.059	0.945 : 1.107
γ^m	0.026	0.028	0.027
		$\gamma_K = \gamma_N = 0.01$	
σ^m	0.750	0.977	1.089
10% : 90%	0.569 : 0.929	0.917 : 1.047	0.990 : 1.167
γ^m	0.025	0.025	0.024
		$\gamma_K = 0.015$, $\gamma_N = 0.005$	
σ^m	0.621	0.950	1.174
10% : 90%	0.490 : 0.768	0.876 : 1.042	1.057 : 1.271
γ^m	0.020	0.023	0.022
	Benchmark	$\gamma_K = 0.02$, $\gamma_N = 0.00$	
σ^m	0.517	0.914	1.307
10% : 90%	0.467 : 0.576	0.822 : 1.022	1.179 : 1.462
γ^m	0.020	0.020	0.020

Table 5

Estimates of Factor-Augmenting Production Technology System, 1952-2009

	NLSUR	**FIML**	**GMM**	**3SLS**
ζ	1.001***	0.999***	1.003***	0.999***
σ	0.694***	0.439***	0.720***	0.721***
γ_K	0.004***	0.005***	0.004***	0.002**
γ_N	0.020***	0.020***	0.020***	0.020***
		Tests & Restrictions		
$\sigma = 1$	[0.000]	[0.000]	[0.000]	[0.000]
$Hicks : \gamma_K = \gamma_N$	[0.000]	[0.000]	[0.000]	[0.000]
$Harrod : \gamma_N = 0$	[0.000]	[0.000]	[0.000]	[0.030]
$Solow : \gamma_K = 0$	[0.000]	[0.000]	[0.000]	[0.000]
$J - test$	–	–	[0.239]	[0.499]
ADF_r, PP_r	[0.005], [0.001]	[0.006], [0.000]	[0.006], [0.001]	[0.004], [0.001]
ADF_w, PP_w	[0.006], [0.003]	[0.013], [0.017]	[0.008], [0.002]	[0.007], [0.002]
ADF_Y, PP_Y	[0.009], [0.031]	[0.010], [0.030]	[0.010], [0.022]	[0.012], [0.024]

Notes: ***, ** and * respectively indicate the 1%, 5%, and 10% level of significance using robust standard errors. "–" denotes not applicable. The p-values for the residual ADF and Phillips-Perron (PP) tests were obtained from 2,500 bootstrap draws.

Table 6

Estimates by Neutrality Assumption, 1952-2009

	Factor-Aug.	**Hicks**	**Harrod**	**Solow**
ζ	1.001***	1.001***	1.001***	1.000***
σ	0.694***	0.997***	0.841***	1.004***
γ	–	0.017***	–	–
γ_K	0.004***	–	–	0.087***
γ_N	0.020***	–	0.021***	–
		Tests & Restrictions		
$\sigma = 1$	[0.000]	[0.258]	[0.000]	[0.186]
ADF_r, PP_r	[0.005], [0.001]	[0.005], [0.002]	[0.005], [0.000]	[0.004], [0.002]
ADF_w, PP_w	[0.006], [0.003]	[0.010], [0.037]	[0.004], [0.001]	[0.004], [0.002]
ADF_Y, PP_Y	[0.009], [0.031]	[0.013], [0.002]	[0.008], [0.032]	[0.014], [0.027]

Notes: All estimations reported using NLSUR. See also notes to Table 5.

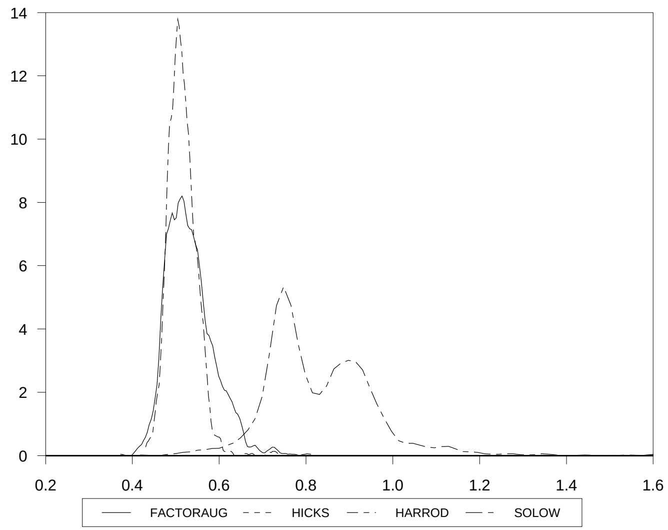

(a) $\gamma_k = 0.00$, $\gamma_N = 0.02$

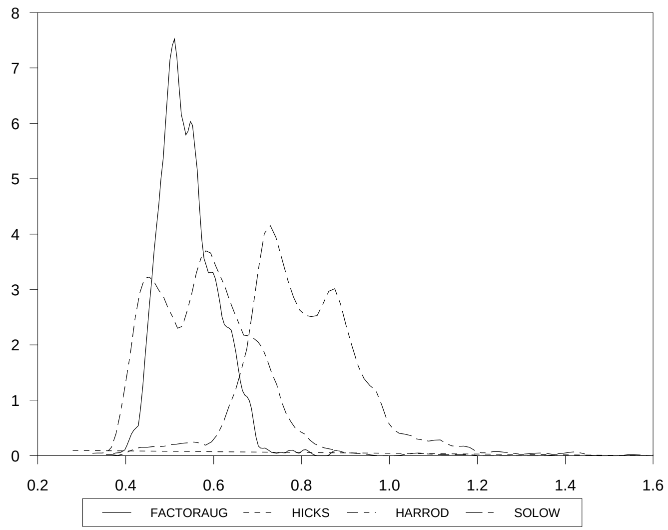

(b) $\gamma_k = 0.05$, $\gamma_N = 0.015$

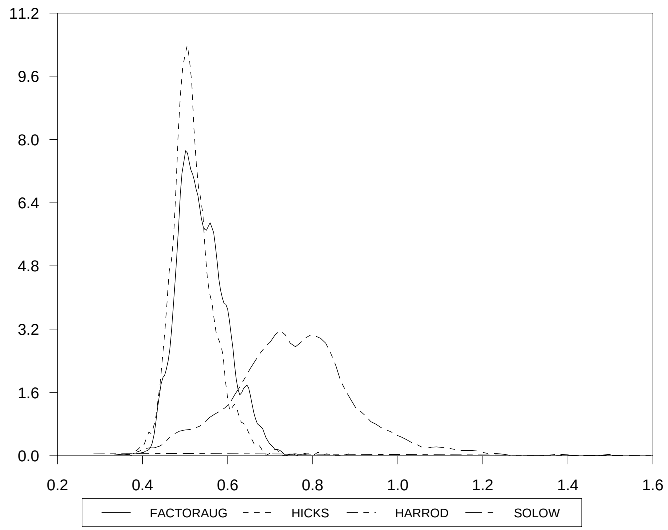

(c) $\gamma_k = 0.01$, $\gamma_N = 0.01$

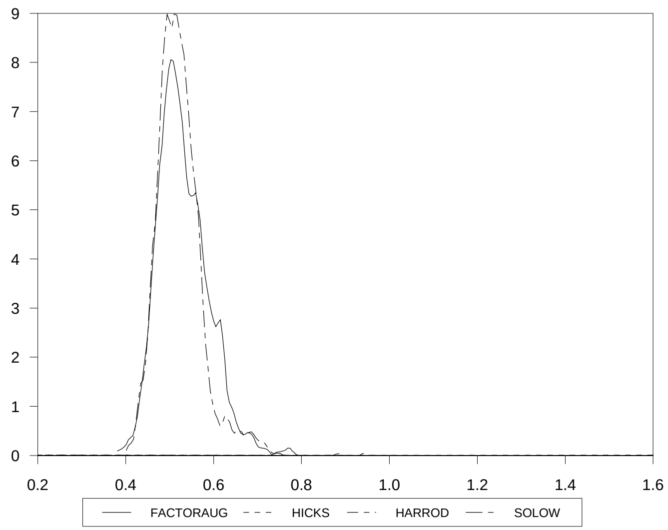

(d) $\gamma_k = 0.02$, $\gamma_N = 0.00$

Figure 1. Distribution of estimated σ. True $\sigma = 0.5$

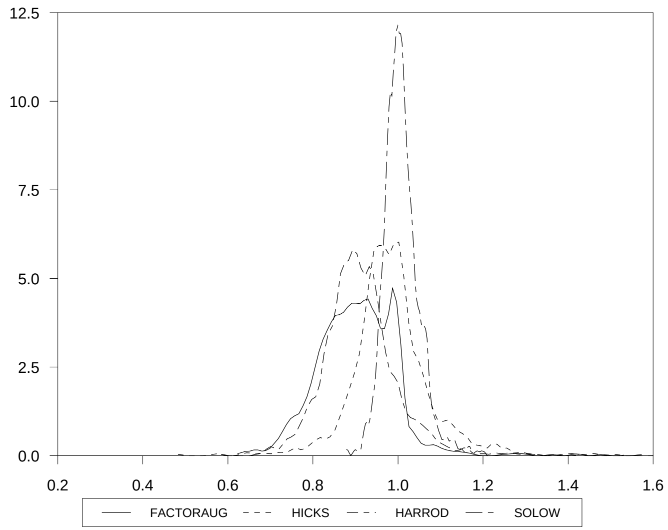

(a) $\gamma_k = 0.00$, $\gamma_N = 0.02$

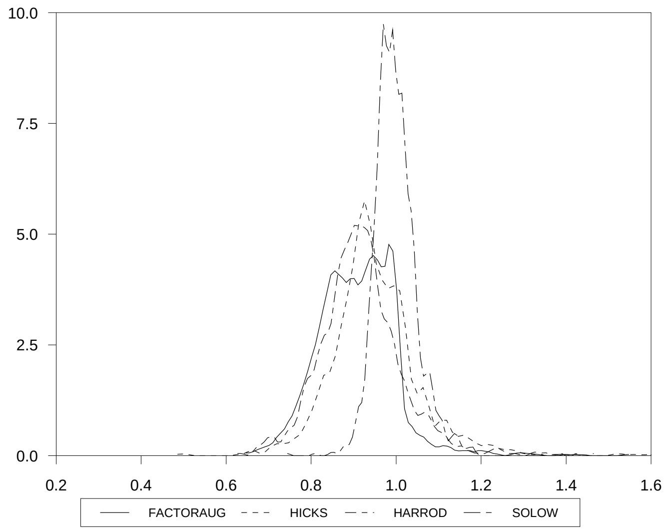

(b) $\gamma_k = 0.05$, $\gamma_N = 0.015$

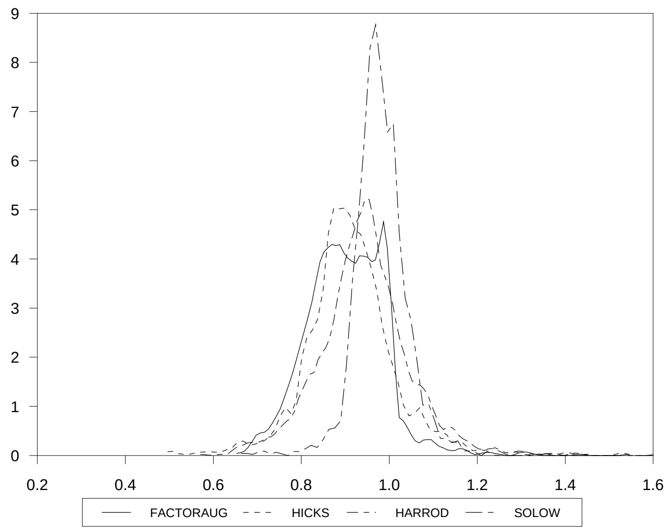

(c) $\gamma_k = 0.01$, $\gamma_N = 0.01$

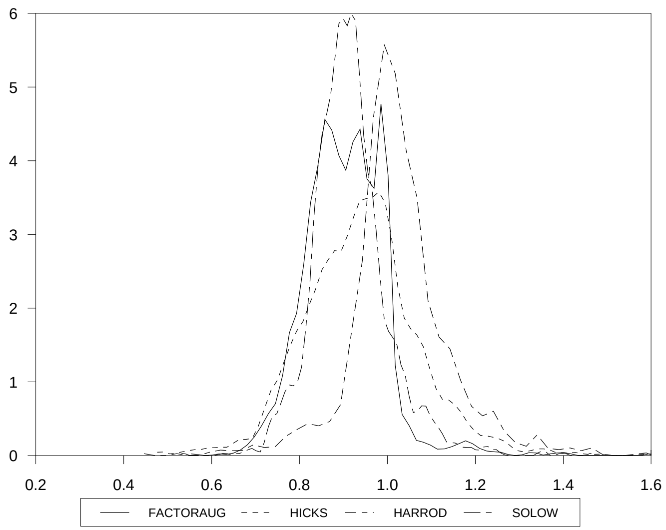

(d) $\gamma_k = 0.02$, $\gamma_N = 0.00$

Figure 2. Distribution of estimated σ. True $\sigma = 0.9$

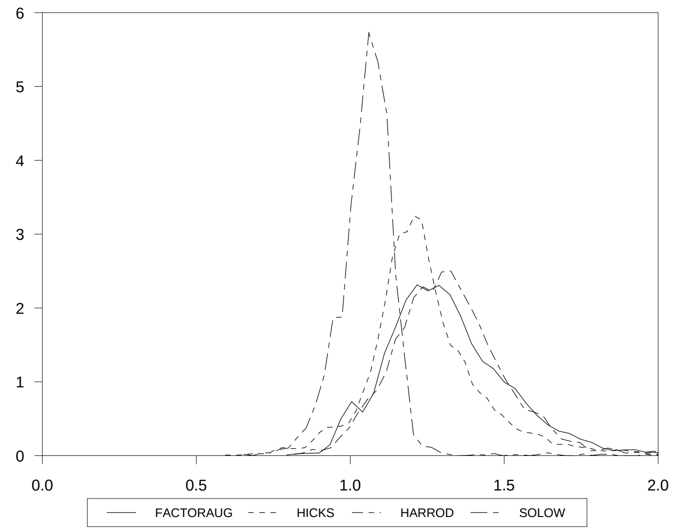

(a) $\gamma_k = 0.00$, $\gamma_N = 0.02$

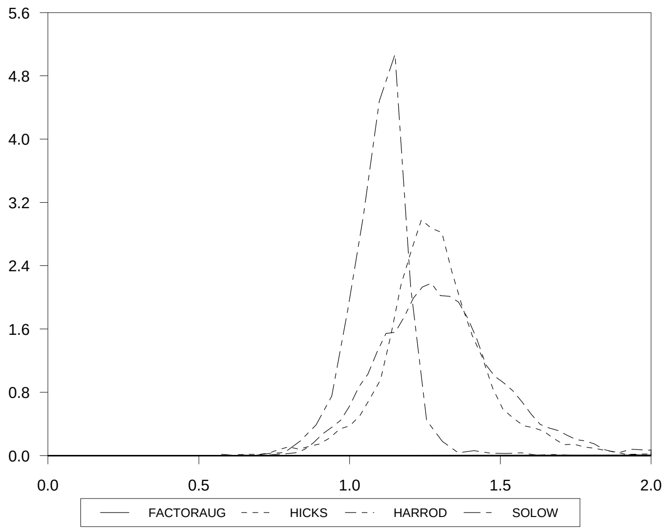

(b) $\gamma_k = 0.05$, $\gamma_N = 0.015$

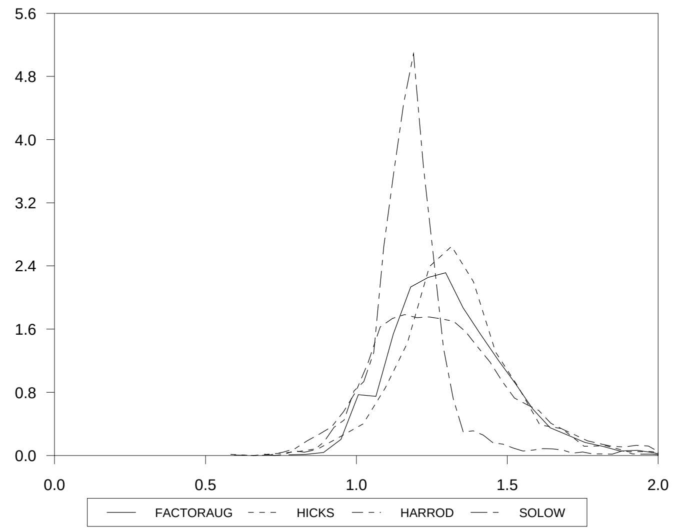

(c) $\gamma_k = 0.01$, $\gamma_N = 0.01$

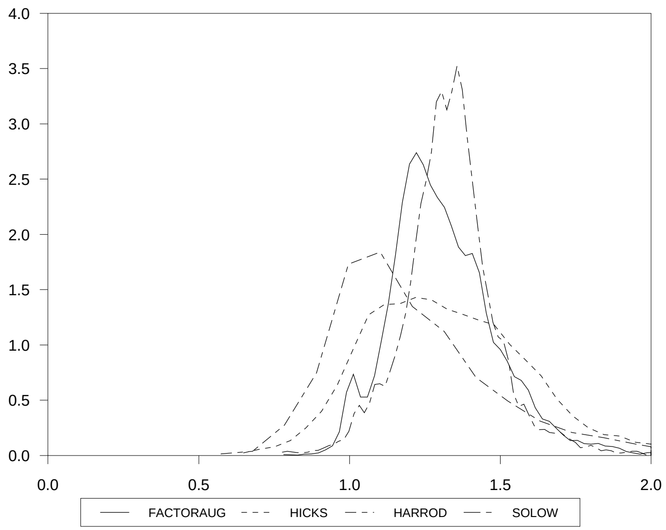

(d) $\gamma_k = 0.02$, $\gamma_N = 0.00$

Figure 3. Distribution of estimated σ. True $\sigma = 1.3$

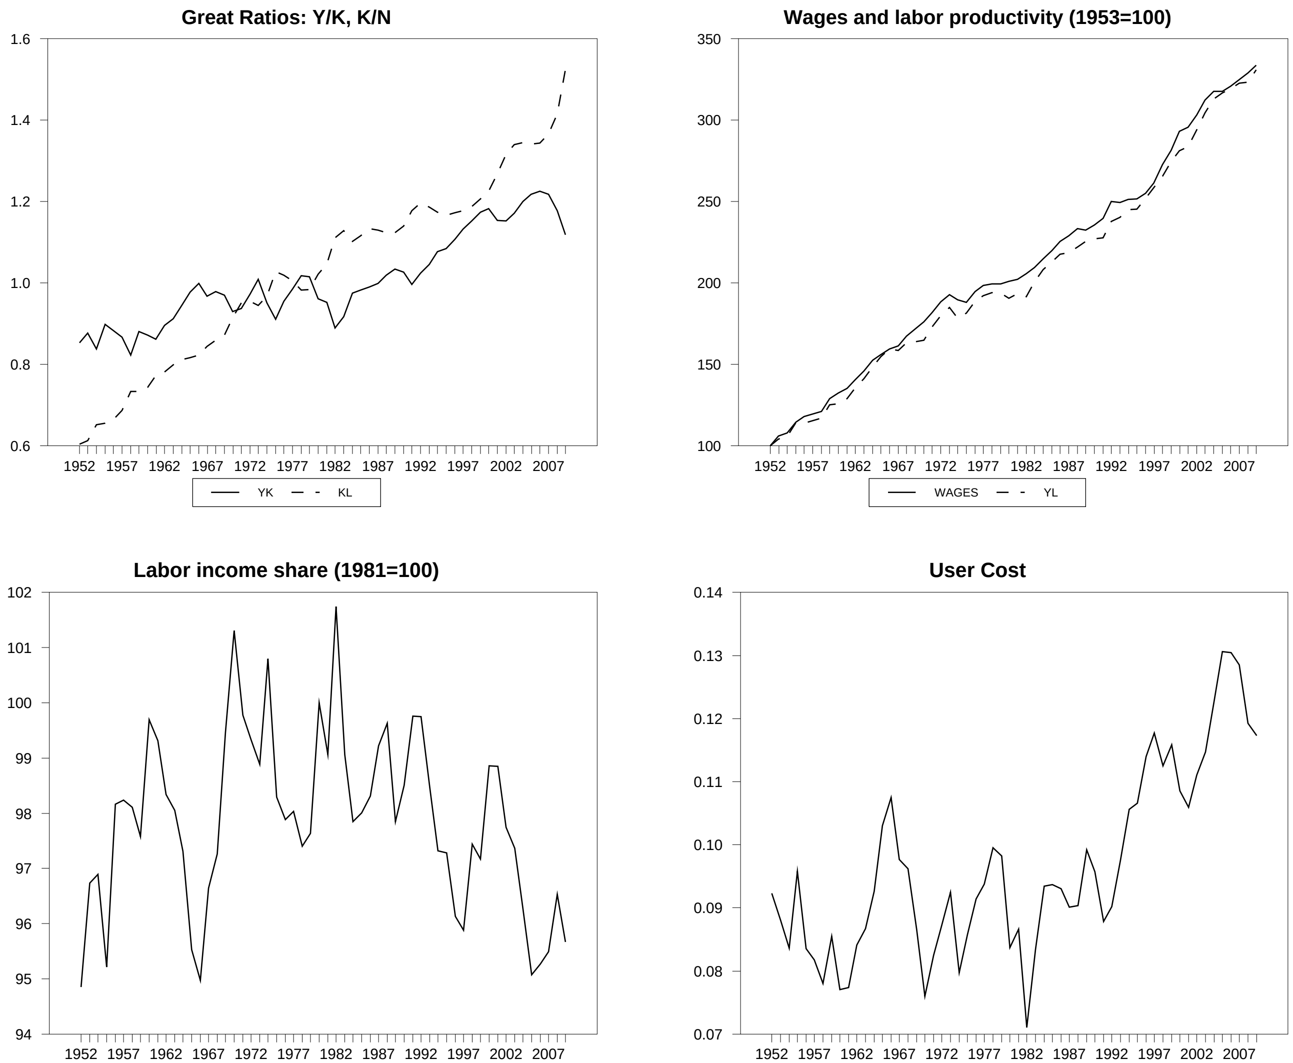

Figure 4. Key variables for the US economy

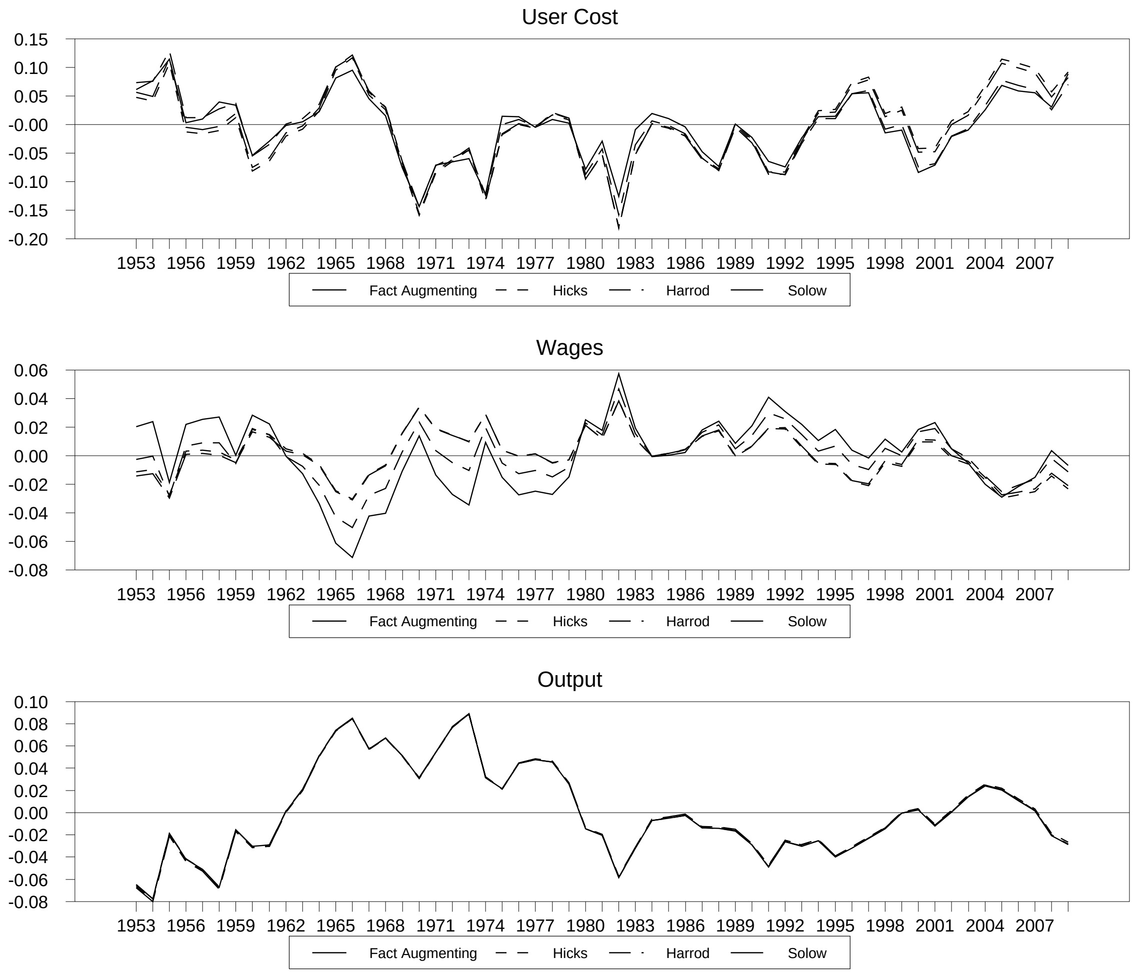

Figure 5. Residuals for the user cost, Wage and Output equations: four specifications (NLSUR)

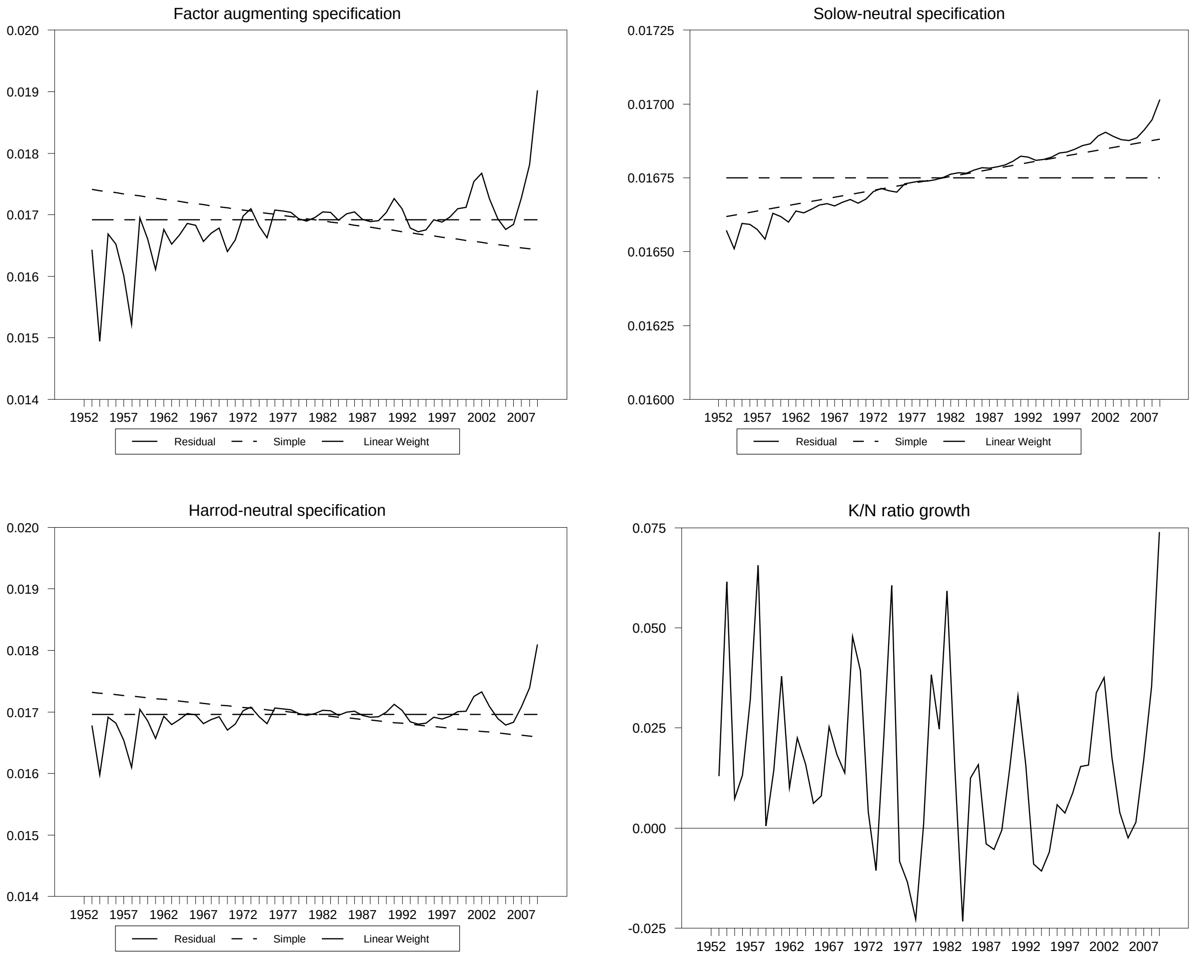

Figure 6. Total Factor Productivity and K/N Ratio Growth (NLSUR)