Review of Economic Studies (2018) **0**, 1–29 doi:10.1093/restud/rdy002

Advance access publication 22 January 2018

Appropriate Technology and Balanced Growth

MIGUEL A. LEÓN-LEDESMA
University of Kent

and

MATHAN SATCHI
University of Kent

First version received February 2013*; Editorial decision November* 2017*; Accepted January* 2018 (*Eds.*)

We provide a general theoretical characterization of how firms' technology choice on a technology frontier determines the long-run elasticity of substitution between capital and labour. We show that the shape of the frontier determines factor shares and the elasticity of substitution between capital and labour. If there are adjustment costs to technology choice, the short- and long-run elasticities differ, with the long-run always higher. If the technology frontier is log-linear, the production function becomes Cobb–Douglas in the long run but, consistent with empirical evidence, short-run dynamics are characterized by gross complementarity. The approach is easily implementable and yields a powerful way to introduce CES-type production functions in macroeconomic models. We provide an illustration within an estimated dynamic general equilibrium model and show that the use of our production technology provides a good match for the short- and medium-run behaviour of the U.S. labour share.

Key words: Balanced growth, Appropriate technology, Elasticity of substitution

JEL Codes: E25, O33, O40

1. INTRODUCTION

As well known among economists, Uzawa's (1961) steady-state growth theorem states that balanced growth requires either all technical progress to be labour-augmenting or the elasticity of substitution between capital and labour to equal one in the long run.[1] Evidence that factor shares are approximately constant in the long run, such as shown in Figure 1, has led to balanced growth being a standard baseline description of long-run data and also a standard constraint for most solution methods in dynamic macro models. Cobb–Douglas, which imposes a unitary capital-labour elasticity of substitution, sits at odds both with the substantial cyclical fluctuations observed in factor shares and the weight of evidence (reviewed in *e.g.* Chirinko, 2008, and León-Ledesma *et al.*, 2010) which supports a value significantly below unity at

1. See Jones and Scrimgeour (2008) for a useful proof. Here, balanced growth refers to a long-run growth path consistent with Kaldor's facts; in particular, constant factors shares, great ratios, and a constant real interest rate.

The editor in charge of this paper was Francesco Caselli.

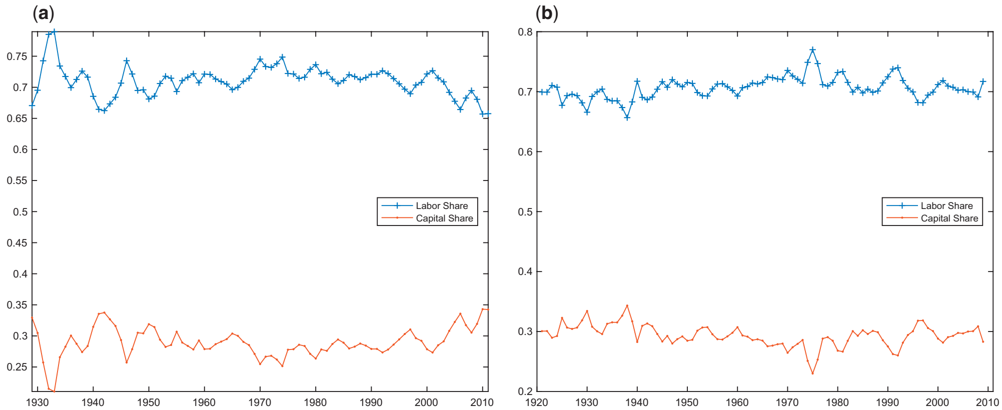

FIGURE 1

Historical capital and labour shares for the U.S. and U.K.

Sources: U.S., Piketty and Saez (2003), U.K., Mitchell (1988) updated by Bank of England. (a) U.S. 1929–2011; (b) U.K. 1921–2012.

standard frequencies. Specifying an elasticity less than unity might be beneficial in modelling fluctuations, but doing so precludes long-run balanced growth unless all technical progress is labour augmenting. However, as discussed below, the assumption that technical progress is indeed *purely* labour augmenting is difficult to justify. In this sense, the theorem constrains modelling practice.

We propose a method of relaxing this constraint by deriving a general production function in which the elasticity of substitution between capital and labour is lower in the short run than in the long run. The special case where the long-run elasticity $\sigma_{LR}=1$ allows flexibility in modelling short-run dynamics (the modeller can choose any short-run elasticity σ_{SR} such that $0<\sigma_{SR}<1$) while retaining general compatibility with balanced growth. A particular focus here is in providing a tractable method to achieve this flexibility. In comparison to alternative approaches, this method can be more easily applied in a wide range of macroeconomic models commonly used for policy making, where the nature of technical progress is not the primary research question.

The framework is based on considering technology choice among firms. The appropriate technology literature (see Atkinson and Stiglitz, 1969, Caselli and Coleman, 2006 and Jones, 2005) analyses this type of technology choice among firms and forms the basis for the current article. Here we particularly build on Caselli and Coleman (2006; henceforth CC). Following an important idea in CC, firms face a standard constant elasticity of substitution production function where output Y is derived from a CES production function

$$Y=\left[A^{\rho}K^{\rho}+B^{\rho}L^{\rho}\right]^{1/\rho}, \tag{1.1}$$

where, in addition to choosing capital K and labour L, the firm also chooses its technology $A\geq 0$ and $B\geq 0$ subject to (A,B) lying within a given technology frontier.[2] Since optimal technology

2. CC draw a world technology frontier. In CC, there are three inputs, capital, unskilled labour, and skilled labour, and the firm chooses the efficiencies (constrained by a technology frontier) of skilled and unskilled labour (which empirical evidence suggests are gross substitutes) rather than of capital and labour (which have short-run gross complementarity) as in our case. Since it is hard to argue that balanced growth applies to these two inputs, this issue (and more generally the time variation of the elasticity of substitution) is not of great relevance in CC.

choices will vary with factor prices, and technology choice influences the quantities of factors employed, technology choice also alters the elasticity of substitution (provided $\rho \neq 0$ as assumed throughout).

Importantly, we assume here that it is costly for firms to adjust their choice of technology. The expressions 'short run' and 'long run' in the article refer respectively to the situations where (following a shock say) no adjustment has occurred and where adjustment is complete. In (1.1) the short-run elasticity σ_{SR} is therefore simply $\sigma_{SR} = \frac{1}{1-\rho}$. Given σ_{SR}, we characterize how the shape of the frontier determines the long-run elasticity σ_{LR}. Its slope in the space of the log efficiencies determines the capital share. A straightforward intuitive explanation of this is given in the article: the capital share is directly related to the slope of the isoquants in this space. Given an interior solution, we show that we always have $\sigma_{LR} > \sigma_{SR}$ and it is the *curvature* of the frontier, in conjunction with σ_{SR} , that determines σ_{LR}. The article provides sufficient conditions for the existence of such an interior solution, the principal one being an upper bound on the curvature or convexity of the frontier.

It follows that a log-linear frontier implies a constant long-run capital share — effectively long-run Cobb–Douglas — and therefore balanced growth. We also provide functional forms that result in a more general long-run CES production function, so we can have a general production function which is CES in both the long- and the short-run limits but has different elasticities at these horizons, with $\sigma_{LR} > 1 > \sigma_{SR}$. Finally, using the key balanced growth case $\sigma_{LR} = 1$, we present an illustrative dynamic macro model that, despite its simplicity, captures well the empirical short- and medium-run behaviour of the labour share.

Empirical Context. Suppose we take compatibility with long-run balanced growth as a model requirement.[3] An alternative approach to reconcile evidence of cyclicality in labour shares with the requirements of balanced growth is to introduce wage and/or price rigidity in a model with a Cobb–Douglas production function. This may generate short-run variations in factor shares if they produce cyclical fluctuations in firm markups. While such rigidities are likely to affect the labour share at business cycle frequencies, at medium-run frequencies we might expect wages and prices to adjust. Following the Comin and Gertler (2006) methodology, we do in fact find significant medium-run fluctuations in factor shares in which rigidities are less likely to explain. As indicated by Beaudry (2005), technical change might play a role in explaining such medium-run phenomena. Furthermore, estimates of capital income built by directly calculating the real user cost as in Klump *et al.* (2007) display considerable business cycle fluctuations that cannot be attributed to changing markups. Bentolila and Saint-Paul (2003) also find evidence that changes in the labour share are significantly driven by technological shifts unrelated to labour market rigidities.

As discussed above, if $\sigma_{SR} = \sigma_{LR} \neq 1$, the alternative offered by Uzawa (1961) is the assumption that technical progress is purely labour-augmenting in the long run. Then permanent labour-augmenting technology shocks will produce short-run fluctuations in factor shares, potentially allowing models to match the data while satisfying balanced growth. However, it is difficult to make a clear theoretical case as to why any permanent technical progress should

3. Some recent literature (see *e.g.* Piketty, 2013 and Karabarbounis and Neiman, 2013) has argued that current trends suggest that the capital share is increasing over time, and so growth is not balanced but best described by $\sigma_{LR} > 1$. This argument is far from settled, since observations on the capital share and the capital-output ratio are disputed due to measurement issues (see *e.g.* Elsby *et al.*, 2013, Bonnet *et al.*, 2004, and Bridgman, 2014). It is important to note that since short-run evidence favours $\sigma_{SR} < 1$, either of these two views of long-run growth suggest $\sigma_{LR} > \sigma_{SR}$ so both allow a useful role for technology choice. Empirical applications which allow $\sigma_{LR} \neq 1$ are a possible topic of future research. The end of Supplementary Appendix B contains a short note on an alternative view put forward by Lawrence (2015), that a falling labour share is best explained by a falling effective capital-labour ratio and capital-labour gross complementarity.

be purely labour augmenting.[4] Empirically, investment-specific technical change (henceforth *IST*) has also taken an important role in macroeconomics in the past two decades (see *e.g.* Greenwood *et al*., 1997 and 2000, and Fisher, 2006, who find *IST* to be one of the key drivers of macroeconomic fluctuations in the U.S. economy). *IST* is clearly not temporary, and relative price data for investment goods, a proxy for *IST*, are clearly trended. Just as with capital-augmenting technical progress, trends in *IST* will not result in balanced growth in conventional models with a CES production function. Nonetheless, our approach provides a justification for the use of CES production functions in modelling short-run dynamics, since any such model can potentially be made compatible with balanced growth by the introduction of technology choice.[5]

Related Theory. This article is most closely related to the "appropriate technology" literature that describes models of technological choice. Prominent examples are CC, and Jones (2005), in turn extended in various interesting ways by Growiec (2008, 2013). In these approaches, the firm typically makes a technological choice by selecting for example a pair (A,B) that represents the efficiency of two inputs of production. The space of available technologies might take the form of a deterministic frontier in (A,B) space (CC), or represent an accumulated stock of arrived individual technologies (A_i,B_i) drawn stochastically from given distributions (Jones, 2005 and Growiec, 2008 and 2013).

In Jones (2005), firms choose the most appropriate of the technologies that have arrived, each of which is Leontief (or CES with a low elasticity of substitution). In the short run, while the firm remains on the current technology, the elasticity of substitution is zero (or low), but in the long run switching to new technologies will cause the elasticity of substitution to increase. The intuition for this is shown in Figure 2, which is very similar to Figure 1 of Jones (2005). The isoquants of the "global" production function that incorporates endogenous technology choice are the convex hull of the isoquants of the individual technologies, and therefore the former should have less curvature than the latter. Jones (2005) shows that if A and B are drawn from independent Pareto distributions, the elasticity of substitution is unity in the long run. Like our present approach, it produces Cobb–Douglas at the firm level in the long run, rather than as a result of aggregation.[6] Growiec (2013), in turn, shows how the use of Weibull distributions can lead to long-run CES.

While these papers provide a rich and elegant description of technical progress, applying these ideas in a conventional macroeconomic modelling framework is difficult. For instance, due to the fact that firms switch to the best available technology, the dynamics in Jones (2005)

4. Theoretical reasons for why technical progress may be purely labour-augmenting are examined in the "induced innovation" strand of the literature, going back to Hicks (1932), Kennedy (1964), Samuelson (1965), Drandakis and Phelps (1966) and Kamien and Schwartz (1968) and including more recently Acemoglu (2002, 2003, 2007) and Zeira (1998) among others. An adequate survey is beyond the scope of this article, but the question of whether the induced innovation literature as a whole produces the outcome of balanced growth without overly-restrictive assumptions on the nature of innovation is not clear (see Acemoglu, 2003, for a useful discussion). One purpose of this article is also to free researchers from formally modelling innovation when this is not essential to the research question. Finally, note that the joint assumptions of a CES production function and purely labour-augmenting technical progress imply long-run cointegration of the log of the capital share and the log of the user cost of capital. See Supplementary Appendix B for a discussion.

5. The introduction of CES production technologies in business cycle analysis has gained relevance in recent years due to an increasing interest in the drivers of factor income shares (see Choi and Ríos-rull, 2009). Cantore *et al*. (2014), for instance, show that the effect of technology shocks on hours worked can solve the technology-hours correlation puzzle when the elasticity of factor substitution differs from one and there are biased technology shocks.

6. The aggregation approach is taken by Houthakker (1955–56). Jones (2005) and Lagos (2006) provide useful discussions of this classic paper. Lagos (2006), in the spirit of Houthakker, derives a Cobb–Douglas form for the aggregate production function by aggregating Leontief production technologies at the firm level using a model with search frictions (assuming an exogenous rental on capital). Since we principally aim at providing a production function, the aims are very different from those here and are primarily directed at accounting for the determinants of observed TFP.

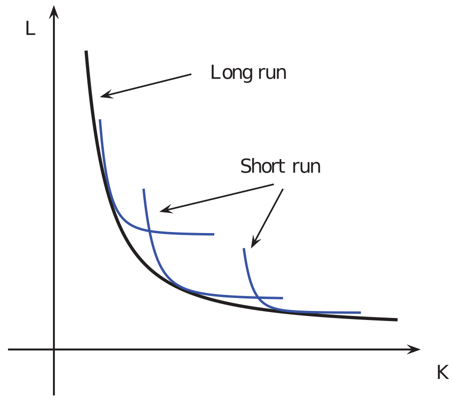

FIGURE 2
Isoquants of the long- and short-run production functions.

have an extreme value property that makes simulation difficult in a conventional forward-looking macroeconomic setting such as a DSGE using the usual solution techniques.[7] CC (see footnote 2) is not explicitly related to balanced growth, but, as we show here, extending the CC framework results in a formulation that is very straightforward to introduce into conventional macroeconomic models.

The rest of the article is organized as follows. The next section contains the key theoretical results. It presents the production technology and its core characteristics, dynamics, and relationship to balanced growth. Section 3 presents an application of this approach to modelling the behaviour of the labour share of income. Section 4 concludes.

2. THE PRODUCTION TECHNOLOGY

Though firm heterogeneity is likely to be of importance when considering technology choice, we leave this for future research to focus on the simplest possible setting. Since constant-returns-to-scale (CRS) is the most important assumption for many of our results, we begin with a generalization of (1.1). Suppose for all firms output Y is given by the following production function

$$Y = F(AK, BL) \equiv F(e^a K, e^b L), \tag{2.2}$$

where $F(.,.)$ is a standard twice continuously differentiable CRS production function with $F_K, F_L > 0$ and $F_{KK}, F_{LL} < 0$.

The efficiencies of capital and labour are A and B respectively, and a and b their logs. As well as choosing K and L, the firm now chooses these efficiencies constrained by a given technology frontier. This represents the choice of relative efficiency of both inputs of production. For instance, capital might make a relatively greater marginal contribution to output in a firm providing web-based customer support service compared to a telephone-based one. A firm may decide to change to a new factory design where the new organization of machinery and operators changes their relative marginal contribution, *e.g.* by changing to web-based provision. This choice of factory, however, is limited by the available designs for new factories determined by the state of knowledge

7. The extreme value property will also have a significant impact on the short-run dynamics of factor shares, increasing the likelihood of sharp adjustments. Here, a standard adjustment cost mechanism results in smoother changes in factor shares.

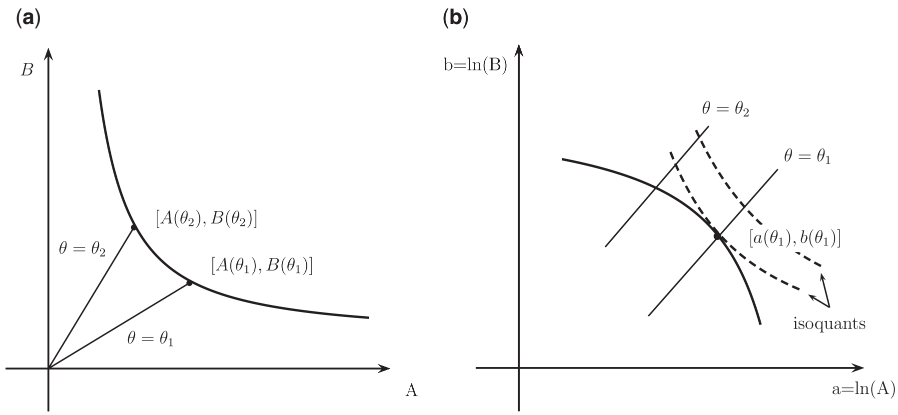

FIGURE 3

The technology frontier. Left panel, the technology frontier in (A,B) space. Right panel, the frontier in log space. (a) The frontier in efficiency space; (b) The frontier in log-efficiency space.

of the economy. In contrast to Jones (2005), we use a continuous technology frontier as in CC. In CC, the technology frontier takes a specific functional form (see footnote 2). To examine balanced growth and other outcomes, we are interested in exploring different possible shapes of the frontier.

We proceed as follows. First, we briefly discuss the frontier and introduce notation and two definitions of equilibrium corresponding to the short and long run (Section 2.1). We then present the main results regarding balanced growth (Section 2.2). Section 2.3 presents more general results in addition to providing the second-order conditions necessary for the balanced growth results. In Section 2.4, we consider the specific case where production takes the CES form in both the short- and long-run limits but with differing elasticities of substitution. Finally, in Section 2.5, we formally introduce adjustment costs in technology choice.

2.1. *The technology frontier*

We can draw the frontier in the space of the efficiencies A and B, or their logs a and b. When we do the latter, we refer to "log-efficiencies" or the "log-frontier". We denote the capital-labour ratio by $k \equiv \frac{K}{L}$ and the ratio of the efficiencies by $\theta \equiv B/A \equiv e^{b-a}$. The quantity $\frac{k}{\theta}$ is then the ratio of capital to labour in efficiency units.

The frontier can only intersect a given ray $B=\theta A$ at one point which we call $(A(\theta;x), B(\theta;x))$, where x is a shift parameter that represents a level of technology. So, we can label any point on the frontier by the ratio of efficiencies θ at that point. The log-efficiencies are $a(\theta;x) \equiv \ln A(\theta;x)$ and $b(\theta;x) \equiv \ln B(\theta;x) = \ln\theta + a(\theta;x)$. The left panel of Figure 3 shows an example of a technology frontier in (A,B) space together with two rays through the origin $B=\theta_1 A$ and $B=\theta_2 A$; the right panel shows the same frontier and "rays" in (a,b) space. We assume the frontier is "smooth" by assuming the function $a(\theta;x)$ — which defines the shape of the frontier — is twice continuously differentiable. We also assume the log-frontier is strictly downward-sloping but not vertical (*i.e.* $a(.)$ is strictly decreasing and $b(.)$ is increasing in θ) so its slope is always defined. This slope is denoted by $s(\theta;x) < 0$.

We assume that technical progress takes the following form: an increase in x, Δx, represents an equi-proportionate shift from each point (A,B) on the frontier to the point $\left(Ae^{\Delta x}, Be^{\Delta x}\right)$. In log-efficiency space, this corresponds to a parallel shift of the log-frontier along 45° lines. An

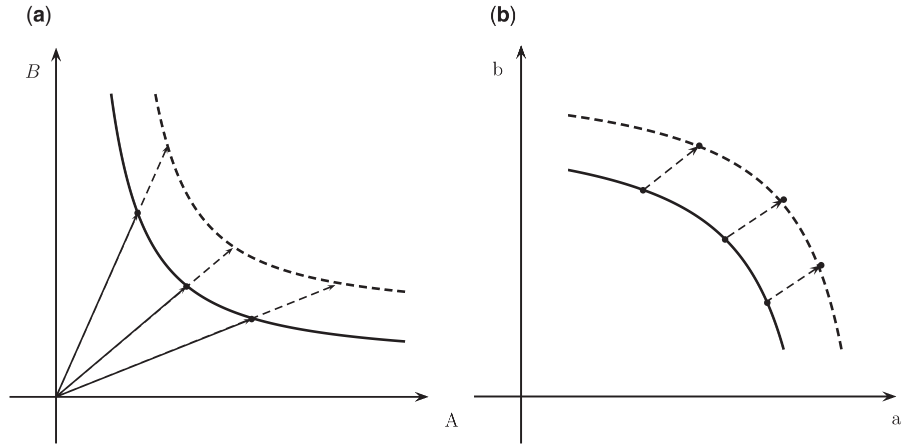

FIGURE 4

Technological progress: an increase in x. Left panel represents the shift in (A,B) space. Right panel represents the shift in log space. (a) In efficiency space; (b) In log-efficiency space.

example is shown in Figure 4. Under this type of expansion, the slope of the frontier will not change as we move along a line of constant θ. Hence $s(\theta;x) \equiv s(\theta)$ is independent of x as are in fact $a'(\theta)$ and $b'(\theta)$. We will see below that this implies that changes in x do not cause firms to change their optimal long-run capital-labour ratio or choice of technology if factor price ratios stay constant, *i.e.* technical progress is Hicks-neutral. Then,

$$s(\theta) \equiv \frac{b'(\theta)}{a'(\theta)} = \frac{1}{\theta a'(\theta)} + 1. \tag{2.3}$$

We assume perfect competition throughout. The firm's problem is to maximize profits $Y - r^k K - wL$ by choosing its quantities of inputs and technology subject to the factor prices of labour and capital, w and r^k respectively. Due to CRS, the solution to this problem determines the capital-labour ratio k in terms of the factor price ratio $\Lambda \equiv \frac{w}{r^k}$. For the moment, we do not explicitly introduce adjustment costs in technology choice. However, it is useful to have two definitions of equilibrium that correspond to the cases of complete (long-run) adjustment in technology choice and to that of no (short-run) adjustment. In the short run, the firm only chooses its optimal inputs within a given factory. In the long run, firms also choose technology θ.

Definition. *Given factor prices r^k and w and a frontier $a(\theta;x)$, an* ***interior long-run equilibrium*** *is a pair (k^*, θ^*), that satisfies the standard first-order conditions $Y_K = r^k$, $Y_L = w$ and $Y_\theta = 0$ and which, if all firms choose $\theta = \theta^*$ and $k = k^*$, is such that no firm can increase its profits by deviating from this choice. Given factor prices r^k and w* ***and*** *current technology (a,b), an* ***interior short-run equilibrium*** *is a capital-labour ratio k^*, that satisfies the standard first-order conditions $Y_K = r^k$ and $Y_L = w$ and is such that, if all firms choose $k = k^*$, no firm can increase its profits by deviating from $k = k^*$ holding (a,b) fixed.*

2.2. *Balanced growth*

Let us now return to Figure 3. Holding K and L fixed, we can draw "efficiency" isoquants in (a,b) space (shown by the dotted lines in the right panel). Let $s^I(a,b)$ represent the the slope of the isoquant through the point (a,b). Note that

$$-s^I(a,b)=\frac{Y_a}{Y_b}=\frac{Ke^aF_1(e^aK,e^bL)}{Le^bF_2(e^aK,e^bL)}=\frac{KY_K}{LY_L}. \tag{2.4}$$

Equation (2.4) applies in and out of equilibrium. The ratio of the marginal gains from increasing capital efficiency to those from increasing labour efficiency must depend on the capital-labour ratio. Beyond that, it depends on the marginal rate of technical substitution between capital and labour, since both efficiency gains and factor quantity increases raise production via efficiency units of the relevant factor. In (an interior long-run) equilibrium, the firm must choose a technology on the frontier where the slope of the frontier equals the slope of the isoquants (*i.e.* $Y_\theta=0$). Since the expression on the right-hand side of (2.4) also equals the ratio of capital to labour income, we have

$$-s(\theta)=\frac{k}{\Lambda}=\frac{\alpha}{1-\alpha}, \tag{2.5}$$

where α is the capital share. Despite its simplicity, equation (2.5) is a very useful one. If the log-frontier is linear, the capital share will remain constant in the long run giving balanced growth, provided we remain at an interior long-run equilibrium.

From equation (2.3), for the capital share α to remain constant in (2.5), we require the following functional form for $a(\theta)$, upto a constant:

$$a(\theta)=x+(\alpha-1)\ln\theta. \tag{2.6}$$

Substituting in (2.2), then gives,

$$Y=F(e^x\theta^{\alpha-1}K,e^x\theta^\alpha L). \tag{2.7}$$

We summarize the above results in the following lemma:

Lemma 1. *Suppose output is given by equation (2.7), where $0<\alpha<1$ is constant, x represents the level of technology, θ technology choice and $F(.,.)$ is a standard CRS production function. If the solutions to the first-order conditions $Y_K=r^k$, $Y_L=w$, and $Y_\theta=0$ maximize firm profits for all x in a given interval, then the capital share is constant and equal to α along the growth path as x increases in that interval.*

Lemma 1 is potentially powerful because it shows that long-run balanced growth does not depend on the shape of $F(.,.)$. If changing technologies is costly in the short and medium run, the shape of $F(.,.)$ will however influence dynamics at these time horizons. Thus for example one can choose a CES production function to model short- and medium-run phenomena without the restrictive assumptions on technology imposed by Uzawa (1961).

Jones (2005) provides a version of (2.7) when $F(.,.)$ is Leontief. In the more general case, we need to establish restrictions on $F(.,.)$ implying that the solution to the first-order conditions does indeed maximize the profits of the firm. This is discussed further below and in Supplementary Appendix A.2. For the solution to the first-order conditions to be unique, we just need that, for any θ, capital and labour are always gross complements in $F(.,.)$. Since this is the empirically preferred assumption for modelling short- and medium-run fluctuations in the labour share, this is unlikely to prove too restrictive. This solution then globally maximizes firm profits if an additional regularity condition holds (closely related to the concept of *strict essentiality* in production functions; see Supplementary Appendix A.2).

2.3. *Some general results*

Without technology choice, if Y is CRS in K and L, and "well-behaved", the standard condition $Y_{KK}<0$ is necessary and sufficient to ensure the appropriate second-order conditions in the firm's problem. Our aim is to develop an equivalent condition with technology choice for a general (non-linear) log-frontier.

Suppose we are at a solution (k^*,θ^*) to the first-order conditions $Y_K=r^k$, $Y_L=w$ and $Y_\theta=0$. Since we assume technology is sticky in the short-run, the short-run elasticity $\sigma_{SR}(k^*,\theta^*)$ between capital and labour is determined purely by the shape of $F(.,.)$ at (k^*,θ^*): it describes how k responds to the factor price ratio $\Lambda\equiv\frac{w}{r^k}$ holding technology (*i.e.* θ) fixed, *i.e.*

$$\sigma_{SR}(k^*,\theta^*)=\left.\frac{\partial\ln k^*}{\partial\ln\Lambda}\right|_{\theta=\theta^*}=\frac{Y_K Y_L}{Y_{KL}Y}. \tag{2.8}$$

The standard expression on the far right-hand side of equation (2.8) only applies because we hold θ constant. This short-run elasticity determines the curvature or convexity of the efficiency isoquants in Figure 3b. In contrast, the long-run elasticity $\sigma_{LR}(k^*,\theta^*)$ gives the response of k^* to a change in Λ when the firm can also optimally choose θ in response to this change.

A measure of the the convexity of the frontier is given by the elasticity $\eta(\theta)$ of the slope of the log-frontier $s(\theta)$ with respect to θ:

$$\eta(\theta)=\frac{\theta s'(\theta)}{s(\theta)}. \tag{2.9}$$

It is straightforward to show that the log-frontier is strictly convex [concave] wherever $\eta(\theta)>0$ [$\eta(\theta)<0$]. We can see from Figure 3b that there is a natural upper bound on the curvature of the log-frontier for the first-order conditions solution to locally maximize firm profits: the log-frontier must at least be less convex than the efficiency isoquants (*i.e.* $Y_{\theta\theta}<0$).[8] If the firm only chose technology, this would also be a sufficient condition. However, since the firm has an additional dimension of optimization in the choice of the capital-labour ratio, in general we need something stricter.

Proposition 2. *The solution to the first-order conditions (k^*,θ^*) locally maximizes firm profits if and only if*

$$\eta(\theta^*)<1-\sigma_{SR}(k^*,\theta^*). \tag{2.10}$$

Proof. See Supplementary Appendix A.1. ||

This is the second-order condition for technology choice. The Appendix also provides a set of sufficient conditions for a solution to the first-order conditions to be unique and globally maximize firm profits, and therefore constitute an interior long-run equilibrium.[9] Assuming condition (2.10) holds, we now derive an expression for the long-run elasticity of substitution, $\sigma_{LR}(k^*,\theta^*)$. Since

8. It can be shown that the log-frontier is less convex than the efficiency isoquants at (k^*,θ^*) iff $\sigma_{SR}(k^*,\theta^*)\eta(\theta^*)<1-\sigma_{SR}(k^*,\theta^*)$, which is always implied by condition (2.10).

9. For example, if condition (2.10) holds for all k and θ, then the marginal rate of technical substitution between labour and capital is strictly monotonic in k (given an optimal choice of θ conditional on k) and so any solution to the first-order conditions is unique. It remains therefore to exclude corner solutions, and a variety of conditions might allow this depending on the shape of the log-frontier and $F(.,.)$. One such set is given in the Supplementary Appendix A.2. This also implies aggregation (see Supplementary Appendix A.4) since all firms will choose the same k_t and θ_t.

$F(.,.)$ is CRS, the partial derivatives $F_1(.,.)$ and $F_2(.,.)$ are homogeneous of degree zero, so we can write

$$\Lambda = \frac{Y_L}{Y_K} = \frac{\theta^* F_2(e^{a(\theta^*)}K, \theta^* e^{a(\theta^*)}L)}{F_1(e^{a(\theta^*)}K, \theta^* e^{a(\theta^*)}L)} = \frac{\theta^* F_2(k^*/\theta^*, 1)}{F_1(k^*/\theta^*, 1)} \equiv \theta^* g\left(\frac{k^*}{\theta^*}\right) \tag{2.11}$$

for some differentiable function $g(.)$. Let ψ be the elasticity of $g(.)$ with respect to k/θ at k^*/θ^*. Let us now take logs and partial derivatives of equation (2.11) with respect to $\ln\Lambda$, holding θ^*constant:

$$1 = \psi \frac{\partial \ln k^*}{\partial \ln \Lambda}. \tag{2.12}$$

It follows from this (and equation (2.8)) that $\psi = 1/\sigma_{SR}$ (dropping arguments for convenience). Noting this, and now taking logs and total derivatives of both first-order conditions (2.11) and (2.5), we have respectively

$$d\ln\theta^* + \frac{d\ln k^* - d\ln\theta^*}{\sigma_{SR}} = d\ln\Lambda \tag{2.13}$$

and

$$\eta(\theta^*) d\ln\theta^* = d\ln k - d\ln\Lambda. \tag{2.14}$$

Therefore, assuming condition (2.10) is satisfied, straightforward algebra gives[10]

$$\sigma_{LR} = \frac{d\ln k^*}{d\ln\Lambda} = \frac{1-\sigma_{SR}-\eta\sigma_{SR}}{1-\sigma_{SR}-\eta} = \sigma_{SR} + \frac{(1-\sigma_{SR})^2}{1-\sigma_{SR}-\eta} \geqslant \sigma_{SR}. \tag{2.15}$$

We can see that in the specific case of a log-linear frontier where $\eta = 0$, as discussed in Section 2.2, the long-run elasticity will be one. If the short-run production function is Cobb–Douglas, then so is the long-run one since then technology choice only amounts to choosing total factor productivity (there being an interior solution if the log-frontier is concave). Otherwise, (2.15) confirms the intuition discussed in the introduction: whenever we have an interior solution with technology choice, the long-run capital-labour elasticity of substitution will always exceed or equal the short-run elasticity. Because technology choice allows the firm to choose a factory design more appropriate to a new capital-labour ratio, it augments the response of the capital-labour ratio to a change in factor price ratios. The more convex the log-frontier is, the greater is the impact of technology choice on the optimal capital-labour ratio. The restriction that capital and labour can be no more than perfect substitutes places an upper bound on the convexity of the frontier;[11] this is the same upper bound given in the second-order condition (2.10) necessary for an interior solution. If the frontier is too convex, the incentive for the firm to deviate will be too strong for an interior solution.

2.4. *Long- and short-run CES production functions*

The previous results were obtained for any generic twice continuously differentiable production function. We now assume that the short-run production function takes the form given in

10. Growiec (January 2017) independently extended a 2015 version of this article to provide this result as well. It was also contained in a December 2016 version of this article.

11. This is perhaps clearer if (2.15) is rearranged as $\eta = (1-\sigma_{SR})\frac{\sigma_{LR}-1}{\sigma_{LR}-\sigma_{SR}}$. The upper bound reached as $\sigma_{LR} \to \infty$ is $1-\sigma_{SR}$.

equation (1.1) which, in the above notation, is

$$Y = \left[\mathrm{e}^{\rho a(\theta;x)} K^{\rho} + \theta^{\rho} \mathrm{e}^{\rho a(\theta;x)} L^{\rho} \right]^{1/\rho}. \tag{2.16}$$

Hence the short-run elasticity of substitution (when the firm chooses only K and L) is given by $\sigma_{SR} = {}^{1}/_{1-\rho}$, and assuming short-run gross complementarity, $\sigma_{SR} < 1$. In Section 2.2 we showed that if the log-frontier was linear, then the long-run production function is Cobb–Douglas. However, suppose we prefer the long-run production function to take a CES form with an elasticity $\sigma_{LR} > \sigma_{SR}$ not necessarily equal to one. For example, we might wish to have a model with $\sigma_{SR} < 1$, in line with evidence on short-run dynamics, but where $\sigma_{LR} > 1$ for reasons described in Piketty (2013) and Karabarbounis and Neiman (2013).

To achieve this, we are interested in the shape of frontier $a(\theta)$ that, given factor prices r^k and w, implies that the firm choosing K, L, and θ in production function (2.16) always chooses the same capital-labour ratio as it would if, instead, it only chose K and L with Y given by a standard CES production technology

$$Y = \begin{cases} e^{x} \left(\alpha K^{\mathcal{R}} + (1-\alpha) L^{\mathcal{R}} \right)^{\frac{1}{\mathcal{R}}} & \text{when } \mathcal{R} \neq 0 \\ e^{x} K^{\alpha} L^{1-\alpha} & \text{when } \mathcal{R} = 0. \end{cases} \tag{2.17}$$

Thus the long-run elasticity of substitution is $\sigma_{LR} = \frac{1}{1-\mathcal{R}}$ where $\rho < \mathcal{R} < 1$. We call the former problem with technology choice P_1 and the latter (standard) problem without it P_2, and the following proposition provides the shape of the frontier for which both have the same solutions. Since $\sigma_{LR} > \sigma_{SR}$, we can see that η is increasing in σ_{LR} in equation (2.19); *i.e.* a higher long-run elasticity corresponds to a more convex frontier.

Proposition 3. *Consider the following function form for the shape of the frontier $a(\theta;x)$:*

$$a(\theta;x,\mathcal{R}) = \begin{cases} x + \frac{1}{\mathcal{R}\zeta} \ln\left(\alpha^{\zeta} + (1-\alpha)^{\zeta} \theta^{-\mathcal{R}\zeta} \right) & \textit{when } \mathcal{R} \neq 0 \\ x + \frac{1}{\rho} [\alpha \ln \alpha + (1-\alpha)\ln(1-\alpha)] - (1-\alpha)\ln\theta & \textit{when } \mathcal{R} = 0 \end{cases} \tag{2.18}$$

where the constant $\zeta \equiv \frac{\rho}{\rho - \mathcal{R}}$.[12] *The functions for the slope and elasticity of the frontier implied by (2.18) are as follows:*

$$s(\theta;\mathcal{R}) = -\left(\frac{\alpha}{1-\alpha} \right)^{\zeta} \theta^{\mathcal{R}\zeta}; \; \eta(\theta;\mathcal{R}) = \mathcal{R}\zeta = (1-\sigma_{SR}) \frac{\sigma_{LR} - 1}{\sigma_{LR} - \sigma_{SR}} < 1 - \sigma_{SR}. \tag{2.19}$$

Then, if P_2 has a unique interior symmetric equilibrium solution,[13] *the unique interior solution to P_1 will result in identical outcomes for $y \equiv {}^{Y}/_{L}$ and $k \equiv {}^{K}/_{L}$ if and only if $a(\theta;x,\mathcal{R})$ takes the form given by equation (2.18).*

Proof. See Supplementary Appendix A.3. ∥

12. Note that, since $\zeta \to 1$ as $\mathcal{R} \to 0$ and that $\frac{\partial \zeta}{\partial \mathcal{R}}|_{\mathcal{R}=0} = \frac{1}{\rho}$, it follows by L'Hôpital's rule that $a(\theta;x,\mathcal{R})$ is continuous in $\mathcal{R}$.

13. See Akerlof and Nordhaus (1967), and De La Grandville (2012); see also the conditions for a symmetric equilibrium in the main proposition of CC. If $\mathcal{R} > 0$, a firm choosing either $L = 0$ or $K = 0$ would respectively obtain $Y = e^{x} \alpha^{\frac{1}{\mathcal{R}}} K$ or $Y = e^{x}(1-\alpha)^{\frac{1}{\mathcal{R}}} L$; hence if $\mathcal{R} > 0$, we must have both $e^{x} \alpha^{\frac{1}{\mathcal{R}}} \leq r^k$ and $e^{x}(1-\alpha)^{\frac{1}{\mathcal{R}}} \leq w$. If X grows over time and r^k is constant, then these condition must at some point be violated and a symmetric equilibrium can no longer hold. De La Grandville (2012) shows that if K and L are gross substitutes and there is permanent capital-augmenting technical progress, then r^k is unbounded over time in the long-run growth path. He therefore argues that a CES production function with $\sigma > 1$ is incompatible with competitive equilibrium if there is sustained capital-augmenting progress. This

2.5. *Adjustment costs in technology choice*

We now consider the explicit introduction of adjustment costs as follows. Suppose a change in θ implies a loss of output $\varphi\left(\frac{\theta_t}{\theta_{t-1}}\right)Y$ where $\varphi \geqslant 0$, $\varphi(1)=\varphi'(1)=0$ and $\varphi''(.)>0$. Given a real interest rate r_t, factor prices for capital and labour r_t^k and w_t respectively, the firm chooses θ_t, K_t, and L_t to maximize

$$\sum_{t=0}^{\infty}\left\{\left[\prod_{s=0}^{t}\left(\frac{1}{1+r_s}\right)\right]\left[Y_t\left(1-\varphi\left(\frac{\theta_t}{\theta_{t-1}}\right)\right)-r_t^k K_t - w_t L_t\right]\right\}, \tag{2.20}$$

where Y_t is given by (2.16) and the frontier takes the form (2.18). The transition between the short- and long-run depends on the speed of adjustment and hence how costly it is to change θ.[14] We next consider an empirical application which corresponds to the balanced growth case where the long-run production function is $Y=e^x K^\alpha L^{1-\alpha}$ implying the frontier has the form $a(\theta)=x+\frac{1}{\rho}[\alpha\ln\alpha+(1-\alpha)\ln(1-\alpha)]-(1-\alpha)\ln\theta$.

Before this, however, we describe some broad features of the behaviour of income shares under perfect competition stemming from (2.20), including its implied countercyclicality of the labour share, in line with observed data. This is useful to interpret some of the results from the quantitative model presented in the next section. Writing $\Omega_t = \theta_t/\theta_{t-1}$, the marginal product of capital condition implied by (2.20), given by equation (3.34) below, implies the following two expressions for the capital share:

$$\frac{r_t^k K_t}{\{1-\varphi(\Omega_t)\}Y_t}=\left(\frac{\theta_t^{1-\alpha}Y_t}{X_t K_t}\right)^{-\rho}=\{1-\varphi(\Omega_t)\}^{\frac{1}{1-\rho}}\left\{\frac{r_t^k\theta_t^{1-\alpha}}{X_t}\right\}^{\frac{-\rho}{1-\rho}}. \tag{2.21}$$

In the short-run, because the capital stock is predetermined, we can see from the first quality in equation (2.21), that the capital share will move in the same direction as Y (since $\rho<0$). With a positive TFP shock, we would expect Y/X to rise on impact assuming a positive labour supply response; with other shocks, where X does not rise, we would expect the capital share to be more procyclical still. In the long run, from the second equality, we can see that, if we had a CES production function, *i.e.* if θ was held fixed, this initial rise will be reversed following a permanent positive TFP or IST shock (the latter leading to a permanent fall in r_t^k): the capital share will fall in the long run. With technology choice, θ_t will rise permanently to offset this, causing the capital share to return to its initial value, though the capital share might still fall below its initial value in the medium term. Thus the capital [labour] share initially rises [falls], perhaps falls [rises] above base in the medium term and then returns to base. For this reason, this model might cause an overshooting of the labour share in response to TFP shocks that is in fact found empirically in

is a general property of CES production functions, rather than of their derivation via a model of technology choice. In fact technology choice potentially removes some of the reassurance that La Grandville's argument offers: it is possible to have a frontier such that the labour share falls to zero asymptotically in the long-run growth path, but where σ tends to 1 asymptotically from above in such a way that these conditions are always satisfied. In fact, many simple piecewise polynomial functional forms for the log-frontier satisfy this property.

14. The presence of adjustment costs naturally leads to the question of firm heterogeneity: suppose firms start off with different levels of technology, how will the distribution of technology choice across firms change over time and will it converge? Under constant returns, the addition of capital adjustment costs or some other factor that constrains firm size ought to be useful with this type of heterogeneity. A description of firm entry and exit might also be desirable; for example new (and small) entrants might enter at an "optimal" level of technology, while larger incumbents might face costs in adjusting. Such questions are left for future research.

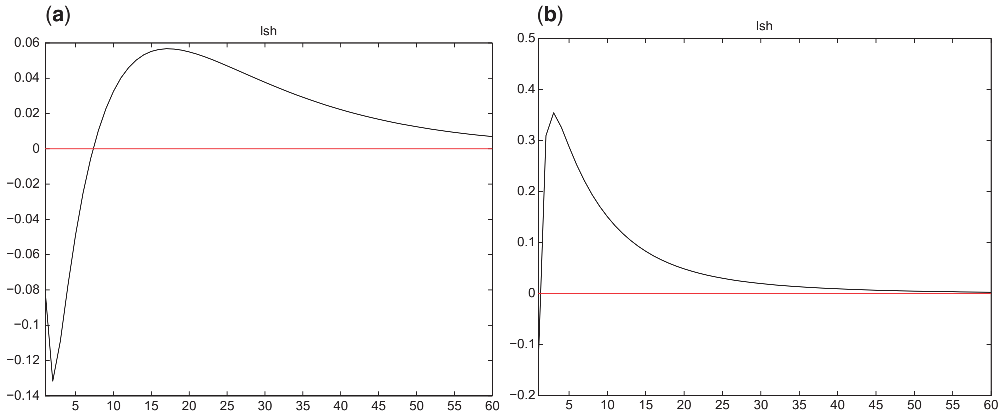

FIGURE 5

Impulse response for the labour share to a 1% X, and Q shocks. $\tau = 20$, $\sigma = 0.2$, $\kappa_i = 0.2$. (a) X shock; (b) Q shock.

U.S. data by Choi and Ríos-rull (2009) and Ríos-Rull and Santaeulàlia-Llopis (2010). Figure 5 at the end of Section 3.2 shows this overshooting response in our estimated model.

Returning briefly to the short-run impact of a TFP shock, note from (2.21) that as θ_t initially adjusts upwards towards its new long-run value, this tends to raise the capital share even further. So though our model is in a sense intermediate between CES and Cobb–Douglas, it is possible that the labour share has a stronger countercyclical response on impact than it would to an equivalent CES model under a TFP shock.[15]

3. THE SHORT- AND MEDIUM-RUN DYNAMICS OF THE LABOUR SHARE

The production technology of Section 2 has a wide range of applications in macroeconomics and growth models where the value of the elasticity of substitution matters for both cyclical and long-run phenomena.[16] In this exercise, we examine how the production function (2.20) performs in the simplest possible model of macroeconomic fluctuations, with a particular view to modelling the labour income share. Business cycle models typically embody a concept of balanced growth in steady state that is consistent with the well-known Kaldor stylized facts. Because of this, we limit ourselves to the specific case of a linear log-frontier so that the long run production function is Cobb–Douglas and the short-run production function is assumed to take the form given in equation (1.1). As a result, Y_t in (2.20) is given by:

$$Y_t = X_t \left((\theta_t^{\alpha-1} K_t)^\rho + (\theta_t^\alpha L_t)^\rho \right)^{\frac{1}{\rho}}, \tag{3.22}$$

15. We find this for the estimated parameters of the technology choice model in Section 3. Comparing the impact of a permanent TFP shock in the model to a CES model with the same σ_{SR}, the labour share falls distinctly more sharply following impact when there is technology choice (contact authors for details). In Section 3, we compare our model with a TFP shock to a CES model with a purely labour-augmenting shock (for compatibility with balanced growth) which alters this comparison.

16. For instance, the analysis of factor-bias in cross-country technology differences (Caselli and Coleman, 2006), the production of "appropriate technologies" and their impact on cross-country income gaps (Basu and Weil, 1998; Acemoglu and Zilibotti, 2001), and the effect of taxes on factor shares (Chirinko, 2002). In business cycle models, Cantore *et al.* (2014) show the important of capital-labour substitution for the response of hours to technology shocks, and Di Pace and Villa (2016) its relevance to match labour market moments in search and matching models.

where $X_t = e^{x_t}$, and, as above, an increase in X_t represents a Hicks-neutral expansion of the technology frontier.

We compare the performance of the model with technology choice (henceforth the TC model) to an equivalent model with the two standard production functions used in modelling macroeconomic fluctuations: Cobb–Douglas and CES. In the CES production function, for compatiblity with balanced growth, we have to replace the Hicks-neutral shock X_t with a labour-augmenting shock Z_t. Since technology choice is absent, the associated adjustment costs disappear so $(1-\varphi)Y_t$ in (2.20) is replaced by $Y_t = \left(\alpha K_t^{\rho} + (1-\alpha)(Z_t L_t)^{\rho}\right)^{\frac{1}{\rho}}$. For Cobb–Douglas, it is replaced by $Y_t = X_t K_t^{\alpha} L_t^{1-\alpha}$.

Given the focus in this exercise on the evolution of the labour share, we introduce investment specific technical change (IST) into the model. Due to the observed decline in the relative price of investment goods, this has been considered an important source of macroeconomic fluctuations (Greenwood *et al*., 1997, 2000 and Fisher, 2006). Permanent investment-specific shocks, however, traditionally have required the restrictive assumption of a Cobb–Douglas functional form for the production function, preventing any fluctuations in factor shares. The combination of investment-specific permanent shocks with the observed volatility of factor income shares[17] (see Growiec *et al*., 2015 for an overview) thus presents a puzzle in macroeconomics. To capture IST, capital accumulation in the model is given by

$$K_{t+1} = Q_t((1-\varphi(\Omega_t))Y_t - C_t) + (1-\delta)K_t, \tag{3.23}$$

where Q_t represents investment specific technical change.

Though counterfactual, in the CES model shocks to Q_t have to be made temporary for compatibility with balanced growth. We can interpret Q_t as a proxy for technical progress in producing investment goods. In a two sector model with consumption and investment goods, with each sector having an identical production structure, we would obtain an equation like (3.23) with Q_t representing Hicks-neutral efficiency in the investment goods sector. In the model with technology choice, an increase in Q_t might represent the impact of an expansion of the technology frontier in the investment goods sector.[18]

Thus we compare a model with technology choice and permanent Hicks-Neutral (X_t) and investment-specific (Q_t) shocks to two equivalent models that do not incorporate technology choice: one with the same shocks and a standard Cobb–Douglas production function, and the other with a CES production function that uses instead permanent labour-augmenting shocks (Z_t) and temporary investment-specific shocks (Q_t) for compatibility with balanced growth. For the purposes of exposition, we describe an overarching model that nests all three of these cases and apply suitable restrictions to examine each case.

To highlight the role of technology choice, these production technologies are compared in an otherwise standard real business cycle model which abstracts from some of the rigidities and frictions emphasized by the business cycle literature. Because these rigidities form important amplification and persistence mechanisms, such an exercise can have limited success in matching the various aspects of the data that we might analyse. Since the impact of such rigidities might be expected to diminish at medium-run frequencies, we also examine the ability of the models

17. As mentioned earlier, factor income shares fluctuations appear to be driven also by factors unrelated to cyclical fluctuations in mark-ups. As we will see below, even at medium-run frequencies when wage and price rigidities are diminished, we observe substantial fluctuations in the labour share.

18. For simplicity, this might be assumed linear with the same slope as the technology frontier in the consumption goods sector. Even then, however, the short-run dynamics in the two sector model would be more complicated due to adjustment costs in technology choice in the investment goods sector.

TABLE 1
U.S. data moments, 1948Q1–2013Q3

	Short-run cycle				Medium-run cycles (CG filter)			
	Std.	Corr.	95% CI		Std.	Corr.	95% CI	
C	0.370	0.799	0.746	0.866	0.803	0.837	0.837	0.877
Inv	2.101	0.884	0.835	0.933	2.070	0.918	0.918	0.937
W	0.441	0.166	0.056	0.364	0.750	0.656	0.656	0.708
L	0.853	0.866	0.827	0.904	0.902	0.574	0.574	0.734
Prod	0.490	0.452	0.401	0.654	0.722	0.664	0.664	0.766
LSH1	0.469	−0.243	−0.398	−0.093	0.335	0.203	−0.025	0.432
LSH2	0.590	−0.416	−0.521	−0.247	0.434	0.181	−0.049	0.412
LSH3	0.481	−0.426	−0.533	−0.319	0.410	0.362	0.181	0.543

Standard deviations relative to output, correlations with output and their 95% confidence intervals.

to match medium-run moments in the data following Comin and Gertler (2006). Given these constraints, we show that the behaviour of the labour share generated by this simple model with technology choice matches reasonably well its observed countercyclicality in the short run and mild procyclicality in the medium run, as well as the shape of its dynamic response conditional on technology shocks.

3.1. *The data*

Table 1 presents the behaviour of some key macroeconomic aggregates using quarterly data for the U.S. for the 1948Q1:2013:Q3 period. The way we filter the data follows Comin and Gertler (2006). They distinguish between short- and medium-run frequencies of the data to capture medium-run cycles. The *medium-run cycle* is obtained using a band pass filter that includes frequencies between 2 and 200 quarters, *i.e.* it filters the data using a very smooth non-linear trend. The medium-run cycle is made up of a *short-run component* (frequencies between 2 and 32 quarters) and a *medium-run component* (frequencies between 32 and 200 quarters). Since most of the data are non-stationary in levels, we applied the filters to the growth rate of the series and reconstructed the filtered levels using their cumulative sum.[19] Note that, although the filter includes frequencies of 200 quarters, in the time domain, this translated into cycles of around 10–12 years as in Comin and Gertler (2006). In the results below, we present the standard deviation of the variables relative to output, the correlation with output, and the 95% GMM confidence intervals for this correlation.

The construction of the data follows standard procedures in the literature (see Supplementary Appendix C for details and sources). Output (Y) was measured as output of the non-farm business sector over civilian non-institutionalized population, consumption (C) is real non-durable and services consumption over civilian non-institutionalized population, investment (Inv) is real private fixed investment plus durable consumption over civilian non-institutionalized population, wages (W) are compensation per hour in the non-farm business sector, and hours worked (L) are measured as all hours in the non-farm business sector over civilian non-institutionalized population. Labor productivity ($Prod$) is measured as Y/L. The labour share measure is important for our exercise. However, measuring the labour share of income is complicated by problems related to how certain categories of income should be imputed to labour and capital owners. Supplementary Appendix C contains a more thorough discussion of the measures of the labour share we used and their construction. Following Gomme and Rupert (2004), we present three

19. For stationary variables this procedure and filtering directly the level series yielded virtually the same results.

TABLE 2
Other countries, labour share moments

	Short-run cycle				Medium-run cycles (CG filter)			
	Std.	Corr.	95% CI		Std.	Corr.	95% CI	
Australia	0.729	−0.382	−0.542	−0.240	0.331	−0.245	−0.432	−0.058
Canada	0.894	−0.555	−0.735	−0.374	0.625	−0.430	−0.644	−0.215
Netherlands	1.148	−0.668	−0.815	−0.521	0.272	−0.463	−0.713	−0.213
Spain	0.672	−0.232	−0.612	0.148	0.308	−0.313	−0.690	0.064
U.K.	0.999	−0.507	−0.663	−0.351	0.539	0.277	0.099	0.454

Sample dates: Australia 1960Q1–2010Q4, Canada 1981Q1–2010Q4, Netherlands 1988Q1–2013Q3, Spain 1995Q1–2010Q4, and U.K. 1960Q1–2010Q4.
Standard deviations relative to output, correlations with output and their 95% confidence intervals.

different measures. The first ($LSH1$) is the labour share of income in the non-farm business sector as reported by the Bureau of labour Statistics. The second ($LSH2$) is the labour share of the domestic corporate non-financial business sector, which is calculated as one minus corporate profits and interests net of indirect taxes over value added. The third measure ($LSH3$) also follows Gomme and Rupert (2004) and calculates the labour share as unambiguous labour income over unambiguous capital income plus unambiguous labour income. We also obtained quarterly data for the labour share for Australia, Canada, The Netherlands, Spain, and the U.K. reported in Table 2. The countries were chosen on the basis of data availability, and Supplementary Appendix C gives details of the sources and data construction.

The behaviour of consumption, investment, and labour market variables is standard, and the short- medium-run split displays the same behaviour as that reported in Comin and Gertler (2006). The standard deviations of all variables, except for investment, increase in the medium run relative to output. Correlations with output also increase with the exception of hours worked, which become less procyclical in the medium-run. The labour share displays a clear countercyclical behaviour at short-run frequencies, with a standard deviation that is about 50% that of output for the U.S., and even larger for the other countries in the sample. In the long run, if balanced growth holds, factor income shares should display no variation or correlation with output. We observe that, at medium-run frequencies, the standard deviation of the labour share falls, and its correlation with output becomes positive, although not significant for two out of the three measures. For the rest of the countries, the standard deviations of the labour share also fall when compared to short-run frequencies, and the medium-run countercyclical behaviour becomes milder. In the case of the U.K., the medium-run correlation with output is positive, but only marginally significant. Overall, hence, we observe that the business cycle countercyclical behaviour of the labour share tends to fade at medium-run frequencies, as does its volatility, indicating a process of convergence towards balanced growth in the long run. In the medium run, the labour share becomes mildly procyclical in some cases. We use these results as a benchmark for the model developed next.

3.2. *The model*

We use a standard RBC model with optimizing representative households and firms. Households maximize their lifetime utility defined over their stream of consumption and leisure, and firms maximize profits. We use a decentralized version of the model where households own the capital and rent it to firms. Note that by including a labour-augmenting efficiency term in the production function (3.22), we get:

$$Y_t = X_t \left((\theta_t^{\alpha-1} K_t)^\rho + (\theta_t^\alpha Z_t L_t)^\rho \right)^{\frac{1}{\rho}}. \tag{3.24}$$

This gives a general production function that nests the three models described at the beginning of Section 3: our model with technology choice, and equivalent models using respectively a CES and Cobb–Douglas production function. Each model has an investment specific shock and one further technology shock. In our technology choice model, we *eliminate the labour-augmenting shock* Z_t in (3.24) by setting $Z_t=1$. For Cobb–Douglas, we also set $Z_t=1$ and adjustment costs to zero. For the CES model, we eliminate the Hicks-neutral shock by setting $X_t=1$, retaining only a permanent labour-augmenting Z_t for compatibility with balanced growth. For the latter, adjustment costs are set to a very large value, so the production function is effectively CES in both the short- and the long-run at the frequencies considered. Since this simplifies the exposition, we describe this general model below.[20]

For the sake of brevity, we skip some of the detail for the explanation of the standard parts of the model. Households choose consumption (C_t), hours worked (L_t), capital stock (K_{t+1}), and one-period non-contingent bonds (B_{t+1}) to maximize their expected lifetime utility $U(\cdot)$:

$$\underset{C_t,L_t,K_{t+1},B_{t+1}}{\text{Max}} E_0 \sum_{t=0}^{\infty} \beta^t U(C_t, L_t),$$

subject to the budget constraint

$$C_t + I_t + B_{t+1} = r_t^k K_t + w_t L_t + (1+r_t) B_t, \tag{3.25}$$

and the law of motion for capital

$$K_{t+1} = (1-\delta) K_t + Q_t I_t. \tag{3.26}$$

I_t is investment in new capital stock, r_t^k is the rental rate of capital, r_t is the interest rate on one-period bonds, w_t are wages, and δ is the rate of depreciation of capital. Investment-specific technical change enters the capital accumulation equation by increasing the productivity of new investment goods. In this model, Q_t is also the inverse of the price of investment relative to consumption goods. We specify a utility function separable in consumption and labour:

$$U_t = \log C_t - \upsilon \frac{L_t^{1+\mu}}{1+\mu}, \tag{3.27}$$

where μ is the inverse of the Frisch elasticity, and υ is a shift parameter.

The firm's problem is to choose K_t, L_t, and θ_t to maximize (2.20) subject to the technology constraint given by the production function (3.24) and the adjustment costs to a change in technology, $\varphi\left(\frac{\theta_t}{\theta_{t-1}}\right)$. The law of motion for technology shocks is given by:

$$d\log Z_t = (1-\kappa_Z)\nu_Z + \kappa_Z d\log Z_{t-1} + (1-\kappa_Z)\epsilon_Z, \tag{3.28}$$

$$d\log X_t = (1-\kappa_X)\nu_X + \kappa_X d\log X_{t-1} + (1-\kappa_X)\epsilon_X, \tag{3.29}$$

$$d\log Q_t = (1-\kappa_Q)\nu_Q + \kappa_Q d\log Q_{t-1} + (1-\kappa_Q)\epsilon_Q, \tag{3.30}$$

so that technological progress is specified as (permanent) rate of growth shocks with drifts ν_i and persistence parameters κ_i for $i=Z,X,Q$. The innovations ϵ_i are zero mean normally distributed

20. In Supplementary Appendix E, we compare the performance of several other models nested by this general model featuring different combinations of shocks. In two of them, we introduce labour-augmenting shocks in our technology choice model. There, Z_t is not seen as a shock to the frontier, but as a shock that affects labour productivity independently of how the firm organizes production. This is just implemented to check the robustness of alternative specifications of the shocks.

with covariance matrix Σ. This specification nests the pure random walk when $\nu_i = 0$, $\kappa_i = 0$, and all the off-diagonal elements of Σ are zero.

Defining $\Omega_t = \frac{\theta_t}{\theta_{t-1}}$ and dropping the expectations operator from forward-looking variables for notation convenience, the first-order conditions for households and firms yield:

$$\frac{C_{t+1}}{C_t} = \beta(1+r_t), \tag{3.31}$$

$$w_t = \upsilon L_t^{\mu} C_t, \tag{3.32}$$

$$1+r_t = \frac{(1-\delta)Q_t}{Q_{t+1}} + r_{t+1}^k Q_t, \tag{3.33}$$

$$\{1-\varphi(\Omega_t)\}\left(\theta_t^{\alpha-1} X_t\right)^{\rho}\left(\frac{Y_t}{K_t}\right)^{1-\rho} = r_t^k, \tag{3.34}$$

$$\{1-\varphi(\Omega_t)\}\left(\theta_t^{\alpha} X_t Z_t\right)^{\rho}\left(\frac{Y_t}{L_t}\right)^{1-\rho} = w_t, \tag{3.35}$$

$$\alpha\{1-\varphi(\Omega_t)\} - \frac{r_t^K K_t}{Y_t} - \left\{\Omega_t \varphi'(\Omega_t) - \frac{\Omega_{t+1}}{1+r_t}\varphi'\left(\Omega_{t+1}\right)\frac{Y_{t+1}}{Y_t}\right\} = 0. \tag{3.36}$$

The capital accumulation equation is given by (3.23) introduced earlier. Equations (3.31)–(3.35) are standard in RBC models: a consumption Euler equation (3.31), labour supply (3.32), an arbitrage condition in capital markets (3.33), and the two factor demand equations (3.34)–(3.35). The final equation (3.36) is the first-order condition for θ_t with adjustment costs. Note that, if there were no adjustment costs, the equation would reduce to $\alpha = \frac{r_t^k K_t}{Y_t}$, *i.e.* the Cobb–Douglas outcome.

An equilibrium in this context is a set of decision rules $\mathcal{D}_t = \mathcal{D}(K_t, Z_t, X_t, Q_t)$ for $\mathcal{D}_t = \left\{C_t, L_t, K_{t+1}, B_{t+1}, \theta_t\right\}$ such that (3.22), (3.23) and (3.31)–(3.36) are satisfied. The model is then appropriately stationarized by dividing all trended variables $K_t, C_t, Y_t, w_t, \theta_t$ by their stochastic trends. The trend for K_t, C_t, Y_t, w_t is defined by $\bar{Y}_t = Z_t X_t^{\frac{1}{1-\alpha}} (Q_{t-1})^{\frac{\alpha}{1-\alpha}}$, whereas the trend for θ_t is given by $\bar{\Theta}_t = (X_t Q_{t-1})^{\frac{1}{1-\alpha}}$.

The functional form for the technology adjustment costs is assumed to be a symmetric exponential function[21] $\varphi(\theta_t/\theta_{t-1}) = 1 - \mathrm{e}^{-\frac{1}{2}\tau(\theta_t/\theta_{t-1}-1)^2}$. Parameter τ determines the speed of adjustment. The model then nests a standard RBC with Cobb–Douglas when $\tau = 0$, and an RBC model with CES production function as $\tau \to \infty$. Note that, in the latter case, only the Z_t process can be allowed to contain stochastic or deterministic trends. Given the observed decline in the relative price of investment, forcing Q_t to be temporary is counterfactual.

Before entering into the parameterization/estimation of the model, and a detailed analysis of the transmission of shocks in both the data and the model, it is worth analysing the transmission of the two key shocks, Q_t and X_t, to the labour share. Figure 5 presents the impulse response of the labour share to a 1% shock to X_t and Q_t. The parameters take standard values in the literature (Table 3). We also used a short run σ of 0.2, a rate of growth persistence of shocks of 0.2 and a τ value of 20.

We can observe that the X_t shock leads to an initial fall and then an overshooting after around 15 quarters akin to the empirical overshooting to TFP shocks in U.S. data found in

21. Other functional forms (*e.g.* quadratic) do not change the results quantitatively.

TABLE 3
Calibrated parameters and priors

α	0.33
μ	0.33
β	0.99
δ	0.025
ρ	$-9/-2.3$
κ_Q	0.266
ν_Q	0.0018
$stdv(\epsilon_Q)$	0.0067
Priors	
$stdv(\epsilon_i)$	InvGamma(0.001, 1)
ν_i	Normal(0.005, 0.01)
κ_i	Beta(0.1, 0.1)
τ	Gamma(10, 5)

Choi and Ríos-rull (2009) and Ríos-Rull and Santaeulàlia-Llopis (2010). The intuition for this is discussed in Section 2.5. The impact of the Q_t can be analysed through its impact on the rental rate of capital r_t^k as shown in equation (2.21). Initially, the falling value of capital can cause a compensating increase in r_t^k just as an increase in depreciation would (see also equation (3.33)). As capital prices continue to fall, so must the rental rate, however, and the shock acts much like a capital-augmenting process which increases the labour share with $\rho < 0$.

3.3. *Calibration and estimation*

To simulate the model, we obtain parameter values by a combination of calibration and estimation. We calibrate those parameters for which we can obtain an observable steady state condition or use information from previous studies, and estimate the rest of the parameters. Table 3 presents the calibration values in the first eight rows. We used a standard value for the steady state capital share of 0.33. β is set to 0.99 as we are matching quarterly data, whereas the depreciation rate is a standard 2.5% per quarter. Parameter μ is set to 0.33, which is consistent with macroeconomic estimates of a Frisch elasticity between 2 and 4 (see Peterman, 2016 and Chetty, 2012). For the parameters driving the law of motion of investment-specific technical change, we estimated equation (3.30) using data for the relative price of investment goods for the 1948Q1-2013:Q3 period. The data were obtained using the implicit deflator for fixed investment and durable goods over the price deflator for non-durables and services consumption. All data were obtained from the BEA. We estimated a drift coefficient of 0.0018 per quarter, a persistence of 0.266, and a standard deviation of the residual of 0.0067.

The values used for the *short-run* elasticity of substitution ranged from 0.1 to 0.3 (implying a ρ coefficient between -9 and -2.3). Time series estimates of the elasticity of substitution for the U.S. range from 0.4 to 0.7 (see León-Ledesma *et al*., 2010 and 2015). In our context, these estimated elasticities would be capturing the average value during the adjustment towards unity in the long run. Hence, the short-run elasticity must be below these benchmark values. To test this effect, we simulated the model with a short-run elasticity of 0.2 and an adjustment speed coefficient $\tau = 20$.[22] We then simulated time-series for the data and estimated the elasticity of substitution using an OLS regression on the log first-order condition for labour. The estimated

22. A τ of 20 is a common value obtained from the estimates discussed below, and implies that a 1% change in θ_t incurs a reasonable output cost of 0.1%.

TABLE 4
Posterior estimates of parameters for technology choice model

Parameter	Posterior mean	90% credible set
τ	19.51	[15.15, 23.90]
κ_X	0.081	[0.068, 0.093]
ν_X	0.003	[0.002, 0.004]
$stdv(\epsilon_X)$	0.014	[0.013, 0.015]

value for the elasticity of substitution was 0.6, which is comfortably in the range of estimates from previous studies. Hence, low values in the range of 0.1–0.3 for the short-run elasticity of substitution are consistent with the estimates offered in the literature.

The rest of the parameters were estimated using Bayesian likelihood methods based on the state-space representation of the model and now standard in the DSGE literature. The bottom part of Table 3 presents the prior distributions used to obtain posterior modes using MCMC methods. The priors are standard in the estimated DSGE literature. The standard deviation of shocks (other than ϵ_Q) follows an Inverse Gamma distribution as they are bounded below by zero and unbounded above. Drifts follow a normal distribution, and persistence coefficients follow beta distributions as they are restricted to the open unit interval. The prior for the adjustment speed τ is drawn from a gamma distribution as it excludes negative draws.

We estimated different versions of the model, allowing for a variety of shocks (see Section 3.4 below and Supplementary Appendix E). All the models contained combinations of two technology shocks and hence we only used a maximum of two observables given that non-singularity requires the same number of observables and shocks.[23] The observables used were the first difference of the log of labour productivity ($d\log(Prod_t)$) and hours per-capita (L_t). We also used alternative variables as observables such as the growth rate of output per capita, consumption growth, investment growth, and the labour share. However, the results remained robust to the choice of observable variables. Given that we estimated a large number of models, we only report those for a model with technology choice and X_t and Q_t shocks in Table 4 where we report posterior means and the 90% credible set for the estimated parameters.[24] For the rest of the models, it suffices to mention that, in the overwhelming majority of cases, posteriors appeared to be far from the priors, indicating that data are adding relevant information in the estimation process. Standard deviations of shocks vary by specification, but generally match well the volatility of output. Drift parameters are also consistent with the average rate of growth of per capita variables in the economy. Finally, the estimated posterior mode for τ was generally found to be close to 20.

For the RBC model with Cobb–Douglas, the estimation follows the same procedure, but τ is restricted to be zero. For the CES model, as discussed above, investment-specific technical change (Q_t) is specified as a stationary AR(1) process, with persistence and standard deviations estimated as above. Again as described above, we use labour-augmenting technology shocks (Z_t) instead of X_t for the CES model. The drift and persistence of Z_t are similarly estimated using Bayesian likelihood methods with above priors; we obtain posterior means $\kappa_Z = 0.21$ and $\nu_Z = 0.005$. Also for this model, we calibrated the elasticity of substitution to 0.5, a higher value than the technology choice model, as there is no adjustment towards Cobb–Douglas.

Finally, following a shock in the technology choice model, we would expect the initial response to be as if the elasticity of substitution is $\sigma = 0.2$, but the long-run response to be as if it were equal to one. Using the estimated posterior mode for τ, we can see this in Figure 6, where we

23. We do not add ad hoc measurement errors.

24. A complete set of results, data, and codes are available upon request.

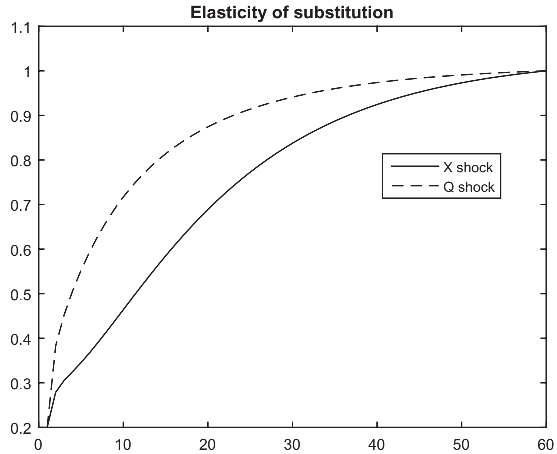

FIGURE 6

Impulse response for an implied elasticity of substitution to a 1% X, and Q shocks using posterior means of parameters.

plot the evolution of an "implied elasticity" after a X_t and a Q_t shock. The construction of this implied elasticity of substitution is explained in Supplementary Appendix D.

3.4. *Model and data moments*

Synthetic data for the macroeconomic variables considered were generated using calibrated values and posterior means of the estimated parameters. We simulated 2,000 data points and kept the last 261 to have the same sample size as in the data. We then applied the same data filters, so we can consistently compare the short- and medium-run moments with those in the data. We compare a large number of models with Cobb–Douglas, CES, and technology choice with different combinations of shocks. Details of this comparison are available in Supplementary Appendix E. As mentioned earlier, we focus on three models, each of which has two technology shocks:

(1) RBC: a standard RBC model with Cobb–Douglas and both permanent Hicks-neutral (X_t) and investment-specific shocks (Q_t).
(2) CES: a model with a standard CES production function with permanent labour-augmenting (Z_t) and temporary investment-specific shocks (Q_t).
(3) TC: a model with technology choice and permanent Hicks-neutral (X_t) and investment-specific shocks (Z_t).

It is obvious from the outset that the RBC model is unable to generate any dynamics in the labour share. On the other hand, because IST shocks can only be temporary, the CES model is inconsistent with the observed trends in the relative price of investment. These are clear a priori disadvantages of these models. Nevertheless, we can compare them in terms of other data moments to see whether the introduction of technology choice comes at the cost of missing important features of the data relative to other models. We present the results for the short-run moments in Table 5 and for the medium-run moments in Table 6. For ease of comparison, we report again the data moments. For the labour share, we used the simple average of the three measures for the U.S. economy. Nevertheless we will refer back to Table 2 for the rest of countries when necessary.

TABLE 5

Short-run theoretical and data moments: relative standard deviations to and correlations with output

	RBC		CES		TC		Data	
	Std.	Corr.	Std.	Corr.	Std.	Corr.	Std.	Corr.
C	0.641	0.775	0.789	0.003	0.650	0.967	0.370	0.799
Inv	2.790	0.900	4.270	0.866	2.173	0.976	2.101	0.884
W	0.692	0.898	0.632	0.184	0.651	0.968	0.441	0.166
L	0.486	0.779	0.631	0.547	0.313	0.499	0.853	0.866
Prod	0.692	0.898	0.841	0.779	0.886	0.952	0.490	0.452
LSH	na	na	0.371	−0.933	0.207	−0.669	0.510	−0.362

TABLE 6

Medium-run theoretical and data moments: relative standard deviations to and correlations with output

	RBC		CES		TC		Data	
	Std.	Corr.	Std.	Corr.	Std.	Corr.	Std.	Corr.
C	0.932	0.953	0.867	0.937	0.811	0.967	0.803	0.837
Inv	1.583	0.844	1.707	0.881	1.732	0.940	2.070	0.918
W	0.940	0.974	0.876	0.950	0.812	0.969	0.750	0.656
L	0.229	0.368	0.114	0.480	0.323	0.668	0.902	0.574
Prod	0.940	0.974	0.951	0.995	0.820	0.956	0.722	0.664
LSH	na	na	0.154	−0.473	0.108	*0.056*	0.390	*0.248*

Italics indicate not significantly different from zero using GMM 95% confidence intervals.

For the TC models, we used a short-run elasticity of substitution of 0.2 ($\rho = -4$). Italics denote statistically insignificant correlations using 95% GMM confidence intervals.

In the short run, the RBC model with Cobb–Douglas reflects well-known results in the literature. Consumption appears to be more volatile than in the data, and labour market moments perform poorly, with an excess volatility of wages relative to hours worked. The CES model does a better job at matching the correlation of wages with output. However, the CES model generates twice the volatility of investment relative to the data, and an almost zero consumption correlation. The TC model is able to generate lower consumption volatility and an investment volatility closer to the data. However, it is not able to match well the labour market moments, just as in standard business cycle models without rigidities.[25] Finally, regarding the short-run moments of the labour share, the CES model generates an excessive countercyclical behaviour, although it matches well its volatility. The TC model does a better job at generating a reasonable countercyclical behaviour, especially when compared to the moments for other countries in Table 2. Overall, in the short run, the RBC and TC models perform better in terms of matching data moments, with the TC model performing best in terms of the behaviour of the labour share.

It is for the medium run where the performance of the model with technology choice (TC) is best relative to the data and other models. The model outperforms the basic RBC and CES in most counts. Importantly, the short-run countercyclical behaviour of the labour share now becomes slightly procyclical, but not significant, in line with the data. There is also a substantial fall in its volatility relative to output. However, this volatility still appears lower than in the data. As

25. We introduced indivisible labour in several versions of the model following Hansen (1985). Indivisible labour increases the volatility of hours and reduces the volatility of wages as expected. But, in the models with choice of technology, the short- and medium-run behaviour of the labour share remains consistent with the data. This is a promising avenue for future research as it would help matching not only the labour share, but the *joint* cyclical behaviour of its components.

the long-run elasticity of substitution approaches unity, both its volatility and countercyclicality fall when we look at medium-run frequencies.

It is also worth mentioning that a model comparison based on posterior odds ratios between the models with choice of techniques and a standard RBC with Cobb–Douglas or CES only, favour dramatically the former. In most cases, if we assign equal prior model probabilities, the posterior probability of the models with choice of techniques is always higher than 0.99.

In conclusion, the TC model does the best job at fitting labour share moments, and only slightly worse at fitting some of the labour market moments only in the short run. The lack of success at reproducing labour market moments is, however, common to most standard macroeconomic models without frictions. The introduction of other mechanisms such as indivisible labour or search and matching frictions (see Di Pace and Villa 2016) should improve all models' moments, but is beyond the scope of our illustration. Our conclusions are also supported by the fuller comparison of models provided in Supplementary Appendix E.

The models discussed in this section lack the variety of shocks and rigidities standard in medium-scale quantitative DSGE models. A legitimate question, thus, is whether our results would survive the introduction of price and wage rigidities (among others) present in most of these models. Let us take a standard New Keynesian (NK) model where prices (and possibly wages) are sticky, and imperfect competition implies a markup over marginal costs. An expansionary demand shock (*e.g.* expansionary monetary policy) increases optimal prices, but since prices are sticky, they remain below optimal. As a consequence, the price markup over costs will fall, increasing the labour share. In a wide variety of existing NK models, the labour share is procyclical conditional on demand shocks.[26] Thus, in such models where dynamics are predominantly driven by demand shocks, a procyclical labour share may be generated, which is counterfactual given our and much other evidence.[27]

Note, however, that a body of evidence questions the validity of a wide variety of such models in terms of the dynamics of the labour share. For example, Cantore *et al*. (2017) show that the labour share robustly increases after a contractionary MP shock in the U.S., U.K., Canada, Australia, and the Euro Area. Hall (2012) uses advertising evidence to suggest that markups are procyclical, and, further, Nekarda and Ramey (2013), surveying a wide range of methods and updating existing evidence, find that markups are in fact procyclical and slightly procyclical or acyclical conditional on demand shocks.

3.5. *Dynamic transmission in the data and in the model*

The performance of the model to reproduce the behaviour of the labour share can also be assessed in terms of its ability to match the dynamic transmission of shocks observed in the data. In order to do so, we first identify the effect of technology shocks on the labour share using a structural VAR (SVAR) on the data. We then compare this response with that of model-generated data. To do so, we take the following steps:

(1) Using a SVAR, identify investment-specific and "neutral" technology shocks in the data, and analyse the impulse responses of the labour share to these shocks.

26. In a New Keynesian model, the markup increases conditional on a technology shock so this in fact helps generate more countercyclicality of the labour share for supply shocks.

27. Use of this production technology in comparison to Cobb Douglas will mitigate this issue in the following sense. Put simply, if μ is the markup and s is the labour share, under Cobb–Douglas $(1+\mu)s$ will be constant and so a countercyclical μ implies a procyclical s. An appropriate modification of equation (2.21) in Section (2.5), shows that here $(1+\mu)s$ is countercyclical, so at least some countercyclicality of μ can be absorbed without a procyclical s. Of course, how quantitatively significant this is is another matter. This is left for future research.

(2) Generate simulated data from the theory model driven by investment-specific and "neutral" stochastic shocks.
(3) Apply the same SVAR to the artificial data and obtain impulse responses.
(4) Compare impulse responses from data and simulated data.

By "neutral" in this context we mean shocks that do not affect the relative price of investment goods in the long run (*i.e.* X_t in our model), following naming conventions in the literature. The reason we compare model and data this way follows Chirinko (2008) who express concerns about the ability of SVARs with long-run restrictions to identify model shocks. Comparing data and model with the *same* (potentially mispecified) metric, we ensure we are carrying out an appropriate model evaluation.

For the identification of the two shocks, we follow Fisher (2006). Intuitively, the identification strategy is that only investment-specific technology shocks can have permanent effects on the price of investment relative to consumption and that both investment-specific and neutral shocks can have a long-run impact on labour productivity. This intuition can then be used to construct a SVAR with long-run restrictions. The basic information set in Fisher (2006) consists of $[\Delta Prod \quad \Delta Prel \quad \ln h]'$. $\Delta Prod$ is the rate of growth of labour productivity, $\Delta Prel$ is the inverse rate of growth of the relative price of investment, and $\ln h$ is the log of hours worked. We use this information set and add the (log) labour share ordered last.[28] The SVAR can be represented as a structural Vector Moving Average (VMA) if it satisfies invertibility and stability conditions. Thus, the structural VMA in our case is:

$$\begin{pmatrix} \Delta Prod_t \\ \Delta Prel_t \\ \ln h_t \\ LSH_t \end{pmatrix} = \begin{pmatrix} \mathbb{C}_{1,1}(L) & \mathbb{C}_{1,2}(L) & \mathbb{C}_{1,3}(L) & \mathbb{C}_{1,4}(L) \\ \mathbb{C}_{2,1}(L) & \mathbb{C}_{2,2}(L) & \mathbb{C}_{2,3}(L) & \mathbb{C}_{2,4}(L) \\ \mathbb{C}_{3,1}(L) & \mathbb{C}_{3,2}(L) & \mathbb{C}_{3,3}(L) & \mathbb{C}_{3,4}(L) \\ \mathbb{C}_{4,1}(L) & \mathbb{C}_{4,2}(L) & \mathbb{C}_{4,3}(L) & \mathbb{C}_{4,4}(L) \end{pmatrix} \begin{pmatrix} \epsilon_{\Delta prod,t} \\ \epsilon_{\Delta prel,t} \\ \epsilon_{h,t} \\ \epsilon_{lsh,t} \end{pmatrix}, \tag{3.37}$$

where $\epsilon_{\Delta prod,t}$, $\epsilon_{\Delta prel,t}$, $\epsilon_{h,t}$, $\epsilon_{lsh,t}$ are contemporaneously correlated shocks with variance covariance matrix $\mathcal{V}$. Since these shocks are correlated, they cannot be interpreted as structural innovations. The problem is to identify structural shocks $[\varepsilon_{N,t}, \varepsilon_{IST,t}, \varepsilon_{h,t}, \varepsilon_{lsh,t}]'$. Here, $\varepsilon_{N,t}$ and $\varepsilon_{IST,t}$ are the neutral and investment specific structural innovations of interest to us. The variance of ε_t is normalized to 1 so that $\mathrm{E}(\varepsilon_t \varepsilon_t') = I$. To transform ϵ_t into orthogonal innovations, we pre-multiply times matrix $\mathbb{D}^{-1}$ such that $\varepsilon_t = \mathbb{D}^{-1}\epsilon_t$. $\mathbb{D}$ is an invertible 4×4 matrix such that $\mathbb{D}\mathbb{D}' = \mathcal{V}$.

Defining the vector of observables as $\mathcal{Y}_t$, then the VMA model (3.37) can be written in terms of structural shocks as $\mathcal{Y}_t = \mathbb{G}\varepsilon_t$, where $\mathbb{G} = \mathbb{C}\mathbb{D}$ and the elements of $\mathbb{C}$ are $\mathbb{C}_{i,j}(L)$ in (3.37). The restriction that $\mathbb{D}\mathbb{D}' = \mathcal{V}$ gives us 10 equations, but $\mathcal{V}$ has 16 elements. Thus, to identify the orthogonal shocks we need six restrictions. We use restrictions on the long-run impact matrix derived from theory. Define $\mathbb{G}(1)$ as the long-run cumulative impact matrix such that $\mathbb{G}(1)_{i,j} = \sum_{s=1}^{\infty} \mathbb{G}_{i,j}^{(s)}$. $\mathbb{G}(1)_{i,j}$ is the cumulative impact on variable i of shock j. By imposing restrictions on these elements, we can identify the orthogonal shocks ε_t.

The first set of restrictions we use come from Fisher (2006). It implies that only IST shocks can have a cumulative impact on $\Delta Prel_t$ (*i.e.* a permanent *level* effect on the relative price of investment). This implies that $\mathbb{G}(1)_{1,2} = \mathbb{G}(1)_{1,3} = \mathbb{G}(1)_{1,4} = 0$. The second set of restrictions come from the well known Galí (1999) identification assumption that only productivity shocks can have long-run effects on labour productivity. In our case this implies that both IST and neutral

28. Ordering the labour share last or before (log) hours does not affect the results.

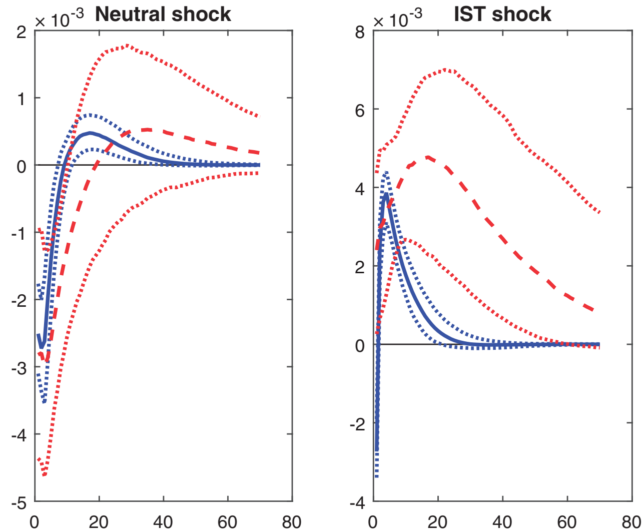

FIGURE 7

Impulse responses of the labour share to a 1% Neutral (left), and IST (right), shock from the structural VAR. Dashed lines are the IRFs obtained using actual data, solid lines the IRFs using model-simulated data. Solid and dashed lines are the median IRF and the dotted lines the 68% credible sets.

shocks can have effects on labour productivity, but not the other two shocks: $\mathbb{G}(1)_{2,3} = \mathbb{G}(1)_{2,4} = 0$. This gives us five restrictions. The final restriction comes from the chosen VAR ordering, but it is not important in our context as we are only interested in the effect of neutral and IST shocks on the labour share since these are the shocks present in our theory model. Since we ordered the labour share last, we impose the restriction that $\mathbb{G}(1)_{3,4} = 0$, that is, that labour share shocks cannot have an effect on the cumulative impulse response of hours.[29] Note that, since both hours and the labour share enter in levels, none of the shocks can have an impact on their level in the long run.

The data used for the vector of observables are the same used for the analysis of the data moments and estimation of the DSGE model. The results reported below use the $LSH2$ measure of the labour share but, as explained below, we also used the other measures. The VAR is estimated using Bayesian methods with a standard Minnesota prior. We used a lag length of 2 for the VAR. We then draw 1,000 times from the posterior distribution of the parameters and obtain impulse responses to the neutral and IST technology innovations and plot the median and the 68% credible sets. We repeat this with the simulated data for the model as explained above. We used the TC model with permanent X_t and Q_t shocks reported above. We simulate 2,000 time series for the variables of interest and drop the first 1,738 so we are left with the same number of observations as for the actual data. The parameters used for the simulation come from the estimation/calibration in the previous subsection.

Figure 7 displays the resulting impulse responses for the data (in red) and the model-generated data (in blue). For the neutral shock, the pattern and the shape of the response of the labour share in both datasets look remarkably similar, save for two aspects. First, the IRF for the data seems to peak later than in the model generated data. Second, the model generated data appears to

29. If we ordered hours last, the restriction implies that hours shocks cannot have a long-run effect on the cumulative labour share.

display less persistence. Nevertheless, the credible sets overlap for most of the periods. For the IST shock, the result is not as satisfactory. However, save for the impact response, both data and model generate IRFs displaying an initial increase of the labour share. The model, however, cannot generate the persistence present in the data IRF. The simple structure of the theory model is not well suited to generate the amplification and persistence we observe in the data. However, this is a well known common problem of DSGE models, particularly those that do not include standard short-run rigidities. With that in mind, the IRF comparison shows that the model does a reasonably good job at reproducing the dynamic transmission of technology shocks for the labour share. More complex structural models including technology choice are left for future research.

We finally carried out some robustness exercises that we only comment briefly for reasons of space. We first experimented with the three measures of the labour share. With the exception of the first measure ($LSH1$), for which the response to IST shocks leads to a fall with large uncertainty bands, the other measures give similar patterns, especially regarding neutral shocks. We also used the Fernald (2014) measure of utilization adjusted TFP instead of labour productivity, leading to results that were remarkably close to the original ones. We also carried out a sample split before and after 1983, a period around which important changes in the U.S. macroeconomy took place. The results for the post-1983 sample are close to those for the whole sample. For the pre-1983 sample, the neutral shocks display a similar pattern, but the IST shocks show the opposite response, albeit with a large degree of uncertainty. Finally, since the identification of the orthogonal innovations does not require the use of hours as an observable, we dropped it for both the actual and simulated data. The results in this case still show a close link between theory generated and actual data IRFs, but the pattern for the neutral shocks does not display the hump-shaped response we see for the original set of observables.[30]

4. CONCLUSIONS

We argue that modelling firms' technology choice on a technology frontier presents at least two distinct advantages for macroeconomic modelling. The first is that the shape of a technology frontier determines the long-run elasticity of substitution between capital and labour. Thus, the frontier determines jointly how the capital/labour share and the long-run elasticity of substitution evolve along the long-run growth path. We provide a theoretical characterization of this process for any generic well behaved production function.

The second advantage is that technology choice naturally leads to a situation where the elasticity of substitution between capital and labour is larger in the long run than in the short run. Our framework allows for a long-run elasticity that can be unity (Cobb–Douglas) or any other value (larger than the short-run elasticity). If there are adjustment costs to technology choice, then the short- and long-run elasticities will differ after an exogenous shock. A particular focus of our framework is to provide a tractable and easily implementable resolution to the "balanced growth conundrum" created by the balance growth path theorem without requiring explicit models of R&D. If balanced growth is a good description of the long-run growth path, this prevents the inclusion of certain kinds of *permanent* technical progress in macroeconomic models when, in accordance with empirical evidence, the elasticity of capital-labour substitution is below one.

30. Note that the VAR identification scheme is only able to distinguish between neutral and IST shocks. However, the theory model can have two types of neutral shocks (X_t and Z_t). Identifying separately these two neutral shocks would require additional restrictions as the model would be under-identified. A combination of long-run and sign restrictions has the potential of dealing with this problem by selecting among IRFs from the under-identified model. We leave this for future research as it departs from the main focus of this article. However, provisional estimates give promising results in support of our model with technology choice.

Using the above framework, we show that, if the technology frontier is log-linear, this leads to a class of production functions that are consistent with balanced growth even in the presence of permanent investment specific or other kinds of biased technical progress, but where short-run dynamics are characterized by gross complementarity.

As an application, we present a stochastic general equilibrium business cycle model with the production technology and estimate it using U.S. data for the 1948:Q1–2013:Q3 period. We show that the model does a good job at matching the behaviour of the labour share of income at short- and medium-run frequencies: the labour share is countercyclical and volatile in the short run, and almost acyclical and smoother in the medium run. The model also performs well in terms of data moments and statistical behaviour against a standard RBC model with Cobb–Douglas, and an RBC model with short- and long-run CES only. It is also capable of reproducing the overshooting of the labour share in reaction to a technology innovation obtained from structural VAR estimates. Extensions of this approach for further research could consider its introduction in multi-sector growth models, the estimation of technology frontiers, a more detailed specification of the labour market, and a richer set of non-technology shocks.

Acknowledgments. We are grateful to Daron Acemoglu, Timo Boppart, Alan Carruth, David de la Croix, Julio Dávila, Olivier de La Grandville, Jakub Growiec, Dirk Krueger, John Leahy, Peter McAdam, Keisuke Otsu, John Peirson, Xavier Raurich, José-Víctor Ríos-Rull, Katsuyuki Shibayama, José I. Silva, seminar participants at the universities of Barcelona, Maastricht, CORE (U. Catholique de Louvain), Paris School of Economics, the 1st China Meeting of the Econometric Society, the XXXIIX Meeting of the Spanish Economic Association, and the 31st Meeting of the European Economic Association for helpful comments and suggestions. We are grateful to Filippo Ferroni for making available codes for Bayesian VAR estimation. We thank especially Francesco Caselli for, among other very useful feedback, painstakingly outlining a very valuable restructuring of the article, and Jonathan Temple for very helpful discussions. We also thank the referees for excellent suggestions. The usual disclaimers apply.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

ACEMOGLU, D. (2002), "Directed Technical Change", *The Review of Economic Studies*, **69**, 781–809.

ACEMOGLU, D. (2003), "Labor-and Capital-augmenting Technical Change", *Journal of the European Economic Association*, **1**, 1–37.

—— (2007), "Equilibrium Bias of Technology", *Econometrica*, **75**, 1371–1409.

ACEMOGLU, D. and GUERRIERI, V. (2008), "Capital Deepening and Nonbalanced Economic Growth", *Journal of political Economy*, **116**, 467–498.

ACEMOGLU, D. and ZILIBOTTI, F. (2001), "Productivity differences", *The Quarterly Journal of Economics*, **116**, 563–606.

AKERLOF, G. and NORDHAUS, W. D. (1967), "Balanced Growth-a Razor's Edge?", *International Economic Review*, **8**, 343–348.

ATKINSON, A. B. and STIGLITZ, J. E. (1969), "A New View of Technological Change", *The Economic Journal*, **79**, 573–578.

BARRO, R. E. and SALA-I MARTÍN, X. (2003), *Economic Growth*, 2nd edn. (MIT Press).

BASU, S. and WEIL, D. N. (1998), "Appropriate Technology and Growth", *The Quarterly Journal of Economics*, **113**, 1025–1054.

BEAUDRY, P. (2005), "Innis Lecture: Explorations in Medium-run Macroeconomics", *Canadian Journal of Economics*, **38**, 1136–1159.

BONNET, O., BONO, P.-H., CHAPELLE, G., *et al.* (2014), "Does Housing Capital Contribute to Inequality? A Comment on Thomas Piketty's Capital in the 21st Century" (Sciences Po Departement of Economics Discussion Papers 2014-07).

BRIDGMAN, B. (2014), "Is Labor's Loss Capital's Gain? Gross Versus Net Labor Shares", (Working Paper No. 0114, Bureau of Economic Analysis (BEA)).

CANTORE, C., FERRONI, F. and LEÓN-LEDESMA, M. A. (2017), "The Missing Link: Monetary Policy and the Labour Share" (Technical report).

CANTORE, C., LEÓN-LEDESMA, M. A., MCADAM, P., *et al.* (2014), "Shocking Stuff: Technology, Hours, and Factor Substitution", *Journal of the European Economic Association*, **12**, 108–128.

CASELLI, F. and COLEMAN, W. J. (2006), "The World Technology Frontier", *The American Economic Review*, **96**, 499–522.

CHETTY, R. (2012), "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labour Supply", *Econometrica*, **80**, 969–1018.

CHIRINKO, R. S. (2002), "Corporate Taxation, Capital Formation, and the Substitution Elasticity between Labour and Capital", *National Tax Journal*, **55**, 339–355.

—— (2008), "σ: The Long and Short of It", *Journal of Macroeconomics*, **30**, 671–686.

CHOI, S. and RÍOS-RULL, J.-V. (2009), "Understanding the Dynamics of the Labour Share: The Role of Noncompetitive Factor Prices", *Annals of Economics and Statistics/Annales d'Économie et de Statistique*, **95/96**, 251–277.

COMIN, D. and GERTLER, M. (2006), "Medium-Term Business Cycles", *The American Economic Review*, **96**, 523–551.

DE LA GRANDVILLE, O. (1989), "In Quest of the Slutsky Diamond", *The American Economic Review*, **79**, 468–481.

—— (2008), "Economic Growth: a Unified Approach" (Cambridge University Press).

—— (2012), "Why Is σ Systematically Observed as Smaller than 1? And Why is it of Fundamental Importance?" (Technical report).

DI PACE, F. and VILLA, S. (2016), "Factor Complementarity and Labour Market Dynamics", *European Economic Review*, **82**, 70–112.

DRANDAKIS, E. M. and PHELPS, E. S. (1966), "A Model of Induced Invention, Growth and Distribution", *The Economic Journal*, **76**, 823–840.

ELSBY, M. W., HOBIJN, B. and ŞAHIN, A. (2013), "The Decline of the US labour Share", *Brookings Papers on Economic Activity*, **2013**, 1–63.

FERNALD, J. (2014), "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity" (Working Paper No. 2012–19, FRBSF).

FISHER, J. D. (2006), "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks", *Journal of Political Economy*, **114**, 413–451.

GALÍ, J. (1999), "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?", *American Economic Review*, **89**, 249–271.

GOLLIN, D. (2002), "Getting Income Shares Right" *Journal of political Economy*, **110**, 458–474.

GOMME, P. and RUPERT, P. (2004), "Measuring Labour's Share of Income" (Policy Discussion Papers, Federal Reserve Bank of Cleveland, November).

GREENWOOD, J., HERCOWITZ, Z. and KRUSELL, P. (1997), "Long-run Implications of Investment-Specific Technological Change", *The American Economic Review*, **87**, 342–362.

—— (2000), "The Role of Investment-Specific Technological Change in the Business Cycle", *European Economic Review*, **44**, 91–115.

GROWIEC, J. (2008), "A New Class of Production Functions and an Argument Against Purely Labour-Augmenting Technical Change", *International Journal of Economic Theory*, **4**, 483–502.

—— (2013), "A Microfoundation for Normalized Ces Production Functions with Factor-Augmenting Technical Change", *Journal of Economic Dynamics and Control*, **37**, 2336–2350.

—— (2017), "Factor-Specific Technology Choice", (NBP Working Papers 265, Narodowy Bank Polski, Economic Research Department).

GROWIEC, J., MCADAM, P. and MUCK, J. (2015), "Endogenous Labour Share Cycles: Theory and Evidence" (Working Paper Series 1765, European Central Bank).

HALL, R. E. (2012), "The Cyclical Response of Advertising Refutes Counter-Cyclical Profit Margins in Favor of Product-Market Frictions" (NBER Working Papers 18370, National Bureau of Economic Research).

HANSEN, G. D. (1985), "Indivisible Labour and the Business Cycle", *Journal of Monetary Economics*, **16**, 309–327.

HICKS, J. (1932), *The Theory of Wages* (London: Macmillan).

HOUTHAKKER, H. S. (1955), "The Pareto Distribution and the Cobb–Douglas Production Function in Activity Analysis", *The Review of Economic Studies*, **23**, 27–31.

JONES, C. I. (2005). "The Shape of Production Functions and the Direction of Technical Change", *The Quarterly Journal of Economics*, **120**, 517–549.

JONES, C. I. and SCRIMGEOUR, D. (2008), "A New Proof of Uzawa's Steady-State Growth Theorem", *The Review of Economics and Statistics*, **90**, 180–182.

JUSTINIANO, A., PRIMICERI, G. E. and TAMBALOTTI, A. (2011), "Investment Shocks and the Relative Price of Investment", *Review of Economic Dynamics*, **14**, 102–121.

KAMIEN, M. I. and SCHWARTZ, N. L. (1968), "Optimal Induced Technical Change", *Econometrica*, **36**, 1–17.

KARABARBOUNIS, L. and NEIMAN, B. (2013), "The Global Decline of the Labour Share", *The Quarterly Journal of Economics*, **129**, 61–103.

KENNEDY, C. (1964), "Induced Bias in Innovation and the Theory of Distribution", *The Economic Journal*, **74**, 541–547.

KLUMP, R. and DE LA GRANDVILLE, O. (2000), "Economic Growth and the Elasticity of Substitution: Two Theorems and Some Suggestions", *The American Economic Review*, **90**, 282–291.

KLUMP, R., MCADAM, P. and WILLMAN, A. (2007), "Factor Substitution and Factor-Augmenting Technical Progress in the United States: A Normalized Supply-Side System Approach", *The Review of Economics and Statistics*, **89**, 183–192.

KLUMP, R. and SAAM, M. (2008), "Calibration of Normalised Ces Production Functions in Dynamic Models", *Economics Letters*, **99**, 256–259.

LAGOS, R. (2006), "A model of TFP", *The Review of Economic Studies*, **73**, 983–1007.

LAWRENCE, R. Z. (2015), "Recent Declines in Labour's Share in us Income: A Preliminary Neoclassical Account" (NBER Working Paper No. 21296. National Bureau of Economic Research).

LEÓN-LEDESMA, M. A., MCADAM, P. and WILLMAN, A. (2010), "Identifying the Elasticity of Substitution with Biased Technical Change", *The American Economic Review*, **100**, 1330–1357.
—— (2015), "Production Technology Estimates and Balanced Growth", *Oxford Bulletin of Economics and Statistics*, **77**, 40–65.
LEÓN-LEDESMA, M. A. and SATCHI, M. (2011), "The Choice of CES Production Techniques and Balanced Growth" (Studies in Economics 1113, School of Economics, University of Kent).
MITCHELL, B. R. (1988), *British Historical Statistics* (Cambridge: Cambridge University Press).
NEKARDA, C. J. and RAMEY, V. A. (2013), "The Cyclical Behaviour of the Price-Cost Markup" (NBER Working Papers 19099. National Bureau of Economic Research).
PETERMAN, W. B. (2016), "Reconciling Micro and Macro Estimates of the Frisch Labour Supply Elasticity", *Economic Inquiry*, **54**, 100–120.
PIKETTY, T. (2013), *Capital in the Twenty-First Century* (Cambridge: Harvard University Press).
PIKETTY, T. and SAEZ, E. (2003), "Income Inequality in the United States", 1913–1998, *The Quarterly Journal of Economics*, **118**, 1–41.
RÍOS-RULL, J.-V. and SANTAEULÀLIA-LLOPIS, R. (2010), "Redistributive Shocks and Productivity Shocks", *Journal of Monetary Economics*, **57**, 931–948.
SAMUELSON, P. A. (1965), "A Theory of Induced Innovation Along Kennedy-Weisäcker Lines", *The Review of Economics and Statistics*, **47**, 343–356.
UZAWA, H. (1961), "Neutral Inventions and the Stability of Growth Equilibrium", *The Review of Economic Studies*, **28**, 117–124.
ZEIRA, J. (1998), "Workers, Machines, and Economic Growth", *The Quarterly Journal of Economics*, **113**, 1091–1117.