DISCUSSION PAPER SERIES

DP17246

Countercyclical Elasticity of Substitution

Dongya Koh and Raül Santaeulàlia-Llopis

MACROECONOMICS AND GROWTH

CEPR

ISSN 0265-8003

Countercyclical Elasticity of Substitution

Dongya Koh and Raül Santaeulàlia-Llopis

Discussion Paper DP17246
Published 23 April 2022
Submitted 20 April 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Countercyclical Elasticity of Substitution

Abstract

We empirically show that the short-run elasticity of substitution between capital and labor is countercyclical. In recessions, capital and labor are more substitutable than in expansions. This countercyclicality of the elasticity of substitution introduces an asymmetry in an otherwise standard competitive-markets business cycle model that contributes to resolve several labor-market puzzles: the labor productivity puzzle, the Dunlop-Tarshis phenomenon, the hours-productivity puzzle, and the labor share puzzle. Interestingly, the cyclicality of the elasticity of substitution is per se not a source of aggregate fluctuations, but it propagates the effects of other shocks.

JEL Classification: E13, E32

Keywords: Short-run, Elasticity, substitution, Capital, Labor, cycle, Labor market, Labor Share

Dongya Koh - chuhyoung@gmail.com
University of Arkansas

Raül Santaeulàlia-Llopis - loraulet@gmail.com
University of Pennsylvania, Universitat Autònoma de Barcelona, Barcelona School of Economics and CEPR

Acknowledgements
We thank Pol Antras, Cristiano Cantore, Nicolas Coeurdacier, Domenico Ferraro, Nir Jaimovic and Irina Panovska for their comments at different stages of this project. This paper previously circulated under the name “The Shape of the Aggregate Production Function over the Business Cycle and Its Implications for the Labor Market”. Raul Santaeulalia-Llopis thanks the Spanish Ministry of Economy and Competitiveness through the Proyectos I+D+i 2019 Retos Investigacion PID2019-110684RB-I00 Grant and the Severo Ochoa Program for Centers of Excellence in R&D (SEV-2015-0563) for financial support.

Countercyclical Elasticity of Substitution

Dongya Koh
University of Arkansas

Raül Santaeulàlia-Llopis*
University of Pennsylvania
UAB, BSE and CEPR

April 20, 2022

Abstract

We empirically show that the short-run elasticity of substitution between capital and labor (σ_t) is countercyclical. In recessions, capital and labor are more substitutable than in expansions. This countercyclicality of σ_t introduces an asymmetry in an otherwise standard competitive-markets business cycle model that contributes to resolve several labor-market puzzles: the labor productivity puzzle, the Dunlop-Tarshis phenomenon, the hours-productivity puzzle, and the labor share puzzle. Interestingly, the cyclicality of σ_t is *per se* not a source of aggregate fluctuations, but it propagates the effects of other shocks.

JEL Codes: E13, E32
Keywords: Short-Run, Elasticity, Substitution, Capital, Labor, Cycle, Labor Market, Labor Share

*We thank Pol Antras, Cristiano Cantore, Nicolas Coeurdacier, Domenico Ferraro, Nir Jaimovic and Irina Panovska for their comments at different stages of this project. This paper previously circulated under the name "The Shape of the Aggregate Production Function over the Business Cycle and Its Implications for the Labor Market". Raï Santaeulàlia-Llopis thanks the Spanish Ministry of Economy and Competitiveness through the Proyectos I+D+i 2019 Retos Investigacion PID2019-110684RB-I00 Grant and the Severo Ochoa Program for Centers of Excellence in RD (SEV-2015-0563) for financial support.

1 Introduction

Business cycle theory and practice commonly require consistency with a balanced growth path (BGP) (King et al., 1988; Cooley and Prescott, 1995). With investment-specific technological change, BGP restricts the choice of the aggregate production function to be Cobb-Douglas (Uzawa, 1961).[1] That is, the elasticity of substitution between capital and labor must be one. Importantly, the Cobb-Douglas assumption gives rise to a set of well-known labor market puzzles for real business cycle models with competitive markets: the Dunlop-Tarshis phenomenon, labor productivity puzzle, labor share puzzle, and hours-productivity puzzle. In this context, a point that we find worth noting is that, although the BGP restricts the long-run production function, the BGP does not constrain the short-run elasticity of substitution.

In this paper, we empirically show that the short-run elasticity of substitution, σ_t, is countercyclical: capital and labor are more substitutable (σ_t increases) in recessions and more complementary (σ_t decreases) in expansions. We provide two identification strategies to estimate the cyclicality of σ_t: a reduced form strategy (e.g. Antràs, 2004) and a structural estimation that matches labor market comovements and impulse response functions (e.g. Christiano et al., 2005). First, using a theoretical framework that distinguishes short- and long-run production, we conduct a reduced form estimation of the short-run elasticity of substitution from optimal factor input demand equations. Our distinctive feature is that we allow the elasticity to differ between recession and non-recession quarters. We find that the short-run elasticity of substitution is significantly higher in recessions, on average $\sigma_R = 0.642$, than in non-recession quarters, on average $\sigma_{NR} = 0.481$. That is, the elasticity of substitution is countercyclical. We assess our empirical strategy using Monte Carlo experiments on simulated data from a vanilla RBC model with competitive markets. We explore potential sources of estimation bias, including nonlinearities in production due to the countercyclicality of σ_t and additional sample size issues.

However, the reduced-form estimation is not free of caveats. In particular, the optimal demand equations might be omitting relevant variables—or wedges as in Chari et al. (2007)—that directly affect the factor input allocations. For example, a home-production shock—a labor wedge arising from the supply side (e.g. Karabarbounis, 2014)—that changes the incentives to supply labor can affect the estimated value of σ_t. Similarly, capital utilization can propagate the response of factor inputs to productivity shocks (e.g. King and Rebelo, 1999) in a manner that can affect the estimation of σ_t. Our approach to addressing these issues is to a provide a structural estimation of σ_t—treated as a latent variable—using a real business cycle model that incorporates some of the potentially unobserved wedges (e.g., home productivity shocks and endogenous capital

[1] See also the more recent proofs in Schlicht (2006) and Jones and Scrimgeour (2008).

utilization). The structural estimation targets a set of relevant business cycle moments in the data, including impulse responses of the labor market, and serves two purposes. First, by using the model to jointly estimate σ_t and other technological parameters together with potential sources of these unobserved wedges, we partially alleviate the omitted biases of the reduced form estimation. Our structural estimates also deliver a countercyclical σ_t with $\sigma_R = 0.665$ in recessions and $\sigma_R = 0.512$ in non-recession quarters; which are both slightly higher than—but within the confidence intervals of—the reduced-form estimates. Second, the fact that our inferred σ_t is structurally estimated allows us to causally interpret the effects of σ_t on the labor market.

Our main quantitative finding is that the countercyclicality of the elasticity of substitution between capital and labor goes a long way in explaining labor market behavior over the business cycle. Intuitively, when negative productivity shocks are accompanied—as in the data—by an increase in the substitutability between capital and labor, the response of labor becomes less attached to that of capital in recessions than when the elasticity is assumed constant. Hence, a negative productivity shock makes the drop of labor larger—deeper and longer-lasting. The higher responsiveness of labor partially transmits to output, making labor productivity decrease by less in recessions. Since higher responsiveness of output coexists with a lower effect on labor productivity, the cyclical correlation between labor productivity and output decreases and helps resolve the labor productivity puzzle.[2] Further, with a non-unitary elasticity of substitution, the response of wages is a weighted average of labor productivity and productivity shocks where the weight is the elasticity of substitution. In this context, a countercyclical elasticity shifts the weight toward productivity shocks and away from labor productivity in recessions. This disassociation with labor productivity makes wages display a longer-lasting response to negative productivity shocks, which lowers the correlation of wages and output consistently with the Dunlop-Tarshis phenomenon.[3] Ultimately, the behavior of wages and labor productivity determines that of labor share, which, under a countercyclical elasticity, is mildly countercyclical as in the data.[4]

An important aspect of our analysis is that the fluctuations of σ_t have no effects *per se* on equilibrium allocations; we theoretically show that this is the case. That is, σ_t is not a source of aggregate fluctuations. Formally, a high-order Taylor expansion of our production function shows that the partial derivatives of output with respect to σ_t evaluated at steady state disappear.

[2]Labor productivity is highly correlated with output from the mid-1950s to the mid-1980s (Hansen and Wright, 1992; Cooley and Prescott, 1995) but less since the mid-1980s (McGrattan and Prescott, 2012; Galí and van Rens, 2021; Ramey, 2012). In our sample, the cyclical (logged and HP-filtered) correlation is $\rho(lp, y) = 0.15$.

[3]A low (or even negative) cyclical correlation between wages and output is found in household- and establishment-survey data (Brandolini, 1995; Abraham and Haltiwanger, 1995). Using aggregate data, the cyclical correlation between wages and output is $\rho(w, y) = -0.13$ in our sample. Christiano and Eichenbaum (1992) redefine this phenomenon as the correlation between hours and wages, which in our sample is $\rho(eh, w) = -.44$.

[4]See Boldrin and Horvath (1995); Gomme and Greenwood (1995); Ríos-Rull and Santaeulàlia-Llopis (2010). Our sample finds a correlation between labor share and output per capita of $\rho(ls, y) = -0.34$.

However, the cross-partial derivatives between σ_t and other sources of fluctuations (and factor inputs) remain. That is, the effects of σ_t exist only when the economy is already outside of the steady state as captured by the the cross-derivatives (i.e., when other shocks are active). For that reason, a standard orthogonalized variance decomposition is invalid in our framework. We conduct a non-orthogonal *cumulative* variance decomposition in order to assess how σ_t affects the contribution of productivity, investment, government, and home-production shocks. Further, the fact that σ_t works through the cross-derivatives introduces a non-linearity in the model that gives rise to asymmetric cycles.[5,6] The asymmetry arises because the countercyclicality of σ_t shapes the endogenous responses in opposite directions in recessions and expansions—detaching the response of labor to that of capital in recessions and attaching it in expansions.

Our work relates to the extensive work on the estimation of the elasticity of substitution (e.g. Antràs, 2004; León-Ledesma et al., 2010). The elasticity of substitution between capital and labor is a central parameter in the macroeconomic theory (Arrow et al., 1961).[7] Here, it is relevant to note that our work focuses on the cyclicality of the short-run elasticity of substitution assuming a long-run Cobb-Douglas component. Our focus on the short-run elasticity is important because our less-than-one average estimate of the short-run elasticity is not directly comparable to the lower-than-one estimates in the empirical literature that does not distinguish between short- and long-run components of the production function (e.g. Klump and Papageorgiou, 2008; Chirinko et al., 2011; Young, 2013). It is also not directly comparable to the larger-than-one estimates from previous literature that focuses on long-run behavior (Karabarbounis and Neiman, 2014). In this context, and given the implications of treating intellectual property products (IPP) rents as ambiguous income in Koh et al. (2020), we argue that a long-run Cobb-Douglas component is a natural choice to describe long-run labor share behavior.[8] In the pursuit of isolating elasticities across data frequencies, our work is perhaps more closely related to Chirinko and Mallick

[5]Therefore, the first-order linearization of our model would miss the effects of σ_t. We use a third-order perturbation method to solve the model, which we find performs very well in terms of accuracy and computational speed compared with the alternative global methods in our context. High order perturbation methods are a standard solution technique for nonlinear models (Fernández-Villaverde et al., 2011). The method entails a disparity between the risky steady state and the deterministic steady state (Coeurdacier et al., 2011; de Groot, 2013). We also solved stripped model versions with value function iteration and projection methods with Chebyshev polynomials, and our results in terms of accuracy and speed favored the perturbation method. This conforms to the insights in Aruoba et al. (2006) that provide a detailed account of performance across solution techniques.

[6]The idea that the cycle reacts more to negative than to positive productivity shocks has been previously highlighted (Sichel, 1993; Hansen and Prescott, 2005). See also recent work on asymmetric behavior in Aruoba et al. (2017), Chang and Hwang (2015), Ferraro (2018), Hairault et al. (2010) and Ilut et al. (2018).

[7]The elasticity of substitution is a classic theme in production function theory (Hicks, 1932). In the context of the business cycle, Sargent and Wallace (1974) show how restrictions on the shape of the production function help account for the behavior of labor productivity and wages in the Lucas (1970)'s model on capacity and overtime.

[8]See McGrattan and Prescott (2014) and McGrattan (2020) for an assessment of intangible capital over the business cycle. IPP represents a part of aggregate intangible capital (Koh et al., 2020).

(2017) that use a low-frequency bandwidth filtering to estimate the long-run substitution elasticity. Instead, since we focus on estimating the short-run elasticity and its cyclical properties, we purposefully purge long-run information from the time series data. To the best of our knowledge, we are the first to estimate the cyclical properties of aggregate elasticity of substitution between capital and labor. Chirinko (2008) carefully discusses the difficulties in accommodating frameworks that allow for differences in short and long-run elasticities. In this context, we give a short- and long-run meaning to, respectively, the locality and globality of the technology in Jones (2003, 2005). Our functional choice also relates to that of León-Ledesma and Satchi (2019) that set theoretical grounds for this particular form of technology with an endongeous elasticity that is lower in the short run than in the long run. We contribute to this literature by showing empirical evidence of a countercyclical elasticity of substitution and studying its quantitative relevance to the business cycle and labor market.

The rest of the paper is organized as follows. Section 2 provides direct empirical evidence of the countercyclical σ_t and a set of Monte Carlo experiments to assess potential sources of bias. We write our fully-fledged model in Section 3 and discuss the structural estimation of σ_t (and other parameters) in Section 4. Our quantitative results are in Section 5. We discuss the role of σ_t in explaining the business cycle and labor market dynamics in Section 6. Section 7 concludes.

2 Countercyclical Elasticity of Substitution: Evidence from the U.S.

First, we introduce a theoretical framework—i.e., an aggregate production function—that explicitly allows for the elasticity of substitution between capital and labor (σ_t) to move with the cycle in Section 2.1. Second, we use our theoretical framework in order to empirically investigate the cyclical behavior of σ_t using U.S. quarterly data in Section 2.2. We find that σ_t is countercyclical. Third, we embed the countercyclical elasticity into a plain-vanilla business cycle model in order to qualitatively discuss its implications for the labor market in Section 2.3.1. Fourth, we conduct a Monte Carlo exercise to validate our estimation strategy in Section 2.3.2.

2.1 A Theoretical Framework

In order to capture the potentially cyclical behavior of σ_t, we propose an aggregate production function that distinguishes the short- and long-run behavior of factor inputs and production that resembles the local-global technology in Jones (2003, 2005).[9] Our novelty is to explicitly allow

[9]Here, we take the locality of the production function as defining the short-run allocations and the globality of the production function as defining the long-run (or BGP) allocations; see also Acemoglu (2002). In this manner, the long-run component of the production function can be additionally interpreted as 'appropriate' values in the sense of Basu and Weil (1997) and Caballero and Hammour (1998).

for the elasticity of substitution between capital and labor to fluctuate in the short run.

Let us first introduce some notation useful for distinguishing short- and long-run behavior. We assume that each and all of our variables x_t can be traced over time as $x_t = x_0(1+\lambda_x)^t\widehat{x}_t$. There is a long-run component $x_t^* = x_0(1+\lambda_x)^t$ with long-run average growth $\lambda_x = \sum_{t=0}^{T} \frac{\lambda_{x_t}}{T}$ and a short-run component $\ln \widehat{x}_t = \ln\left(\frac{x_t}{x_t^*}\right)$ (log-deviations from the trend). If $\widehat{x}_t = 1$ (i.e., if there are no deviations from trend), then x_t follows a deterministic trend. For simplicity, let's assume that a shock $\widehat{a}_t$ is the only exogenous source generating business cycle fluctuations (i.e., deviations from the trend). Then, we can write the behavior of x_t as:

$$x_t = \begin{cases} \frac{x_t}{x_t^*}x_t^* = \widehat{x}_t x_t^* & \text{if} \quad \widehat{a}_t \neq 1 \\ x_t^* & \text{if} \quad \widehat{a}_t = 1. \end{cases} \tag{1}$$

Hence, we can write an aggregate production function that takes potentially a different shape in the short run, f_t^{SR}, than in the long run, f^{LR}, as follows:

$$y_t = \begin{cases} f^{SR}\left(\frac{k_t}{k_t^*}, \frac{a_t}{a_t^*}\frac{l_t}{l_t^*}, \sigma_t\right) f^{LR}(k_t^*, a_t^* l_t^*) & \text{if} \quad \widehat{a}_t \neq 1 \\ f^{LR}(k_t^*, a_t^* l_t^*) & \text{if} \quad \widehat{a}_t = 1 \end{cases} \tag{2}$$

where y_t is output per capita, k_t capital per capita, l_t aggregate hours per capita, and a_t labor-augmenting technical change. We denote with $y_t^* = f^{LR}(k_t^*, a_t^* l_t^*)$ the long-run production occurring when there are no shocks. In contrast, short-run production $\frac{y_t}{y_t^*} = f_t^{SR}\left(\frac{k_t}{k_t^*}, \frac{a_t}{a_t^*}\frac{l_t}{l_t^*}, \sigma_t\right)$ is used when there are shocks that move technology away from trend. Our differential insight is that short-run production is explicitly shaped by an elasticity of substitution between capital and labor, σ_t, that is allowed to cyclically vary with time. Next, we choose explicit forms for the long- and short-run components of the production function.

First, in order to be consistent with balanced growth, theory imposes strong discipline on the choice of the long-run component of the production function, f^{LR}. In particular, the steady-state growth theorem with investment-specific technical change (ISTC) as an engine of growth (Greenwood et al., 1997, 2000)[10] combined with a strictly positive and constant factor income shares requires the long-run component of production to be Cobb-Douglas (Uzawa, 1961; Schlicht, 2006; Jones and Scrimgeour, 2008). Since our benchmark model incorporates ISTC, we set the long-run component of the production function to be Cobb-Douglas (CD):

$$f^{LR}(k_t^*, a_t^* l_t^*) \;=\; (k_t^*)^{1-\vartheta}(a_t^* l_t^*)^{\vartheta}, \tag{3}$$

[10]See more recent applications in Fisher (2006), Justiniano and Primiceri (2008), Justiniano et al. (2011), Ríos-Rull et al. (2012) and Choi and Ríos-Rull (2020). See a nice summary in Ramey (2016).

which implies a constant long-run elasticity of substitution equal to one.

Second, while the long-run component of the production function must comply with balanced growth path restrictions, this is not required for the short-run component of the production function. In other words, f_t^{SR} is free to take any shape. In order to allow for cyclical movements of the elasticity of substitution, we choose a flexible nonconstant elasticity of substitution (NCES) technology for short-run production:

$$f^{SR}\left(\frac{k_t}{k_t^*}, \frac{a_t}{a_t^*}\frac{l_t}{l_t^*}, \sigma_t\right) = \left((1-\alpha)\left(\frac{k_t}{k_t^*}\right)^{\frac{\sigma_t-1}{\sigma_t}} + \alpha\left(\frac{a_t}{a_t^*}\frac{l_t}{l_t^*}\right)^{\frac{\sigma_t-1}{\sigma_t}}\right)^{\frac{\sigma_t}{\sigma_t-1}}. \tag{4}$$

where the short-run elasticity of substitution σ_t is allowed to vary over time.

An Aggregate Production Function with Cyclical σ_t. Putting the NCES short-run production (4) together with the BGP-disciplined long-run technology (3), we pose our aggregate production function as:

$$y_t = \begin{cases} \left((1-\alpha)\left(\frac{k_t}{k_t^*}\right)^{\frac{\sigma_t-1}{\sigma_t}} + \alpha\left(\frac{a_t l_t}{a_t^* l_t^*}\right)^{\frac{\sigma_t-1}{\sigma_t}}\right)^{\frac{\sigma_t}{\sigma_t-1}} (k_t^*)^{1-\vartheta}(a_t^* l_t^*)^{\vartheta}, & \text{if} \quad \widehat{a}_t \neq 1 \\ (k_t^*)^{1-\vartheta}(a_t^* l_t^*)^{\vartheta} & \text{if} \quad \widehat{a}_t = 1 \end{cases} \tag{5}$$

In this manner, aggregate production follows a CD technology in the absence of shocks. In contrast, aggregate production follows an NCES-CD technology with an NCES short-run component and a CD long-run component with the presence of shocks. Our specification of the short-run component relates to the normalized constant elasticity of substitution (CES) in de LaGrandville (1989) and Klump and de LaGrandville (2000) that insulates the effects of the elasticity of substitution against the units of the factor inputs. Note that in our case, since we are interested in the cyclical properties of the elasticity of substitution, we purposefully choose the long-run values (i.e., deterministic trends) of output and factor inputs as reference for the normalization. Our choice serves the purpose of making the normalized values of the output and factor inputs equal to (log) deviations from the trend as defined in equation (1). Further, note that if we shut down the cyclicality of the elasticity in the short-run, i.e., set $\sigma_t = \sigma$ for all t, then our aggregate production function is identical to the one proposed in Jones (2003, 2005). Hence, our unique point of departure from previous work is that we allow the elasticity σ_t to vary over time in the short run. We study the empirical plausibility of this assumption next in Section 2.2.

2.2 Empirical Evidence on the Countercyclicality of σ_t

To empirically estimate the short-run elasticity of substitution and its cyclical properties, we focus on the optimal factor input demand equations (e.g. Antràs, 2004). Given long-run values, a representative firm's problem that maximizes profits with our NCES-CD technology and faces competitive markets implies these first-order conditions for labor l_t and capital k_t:

$$FOC(l_t): \qquad d\ln\left(\frac{y_t/y_t^*}{a_t l_t/a_t^* l_t^*}\right) = \sigma_t d\ln\left(\frac{w_t/w_t^*}{a_t/a_t^*}\right) + d\sigma_t \ln\left(\frac{w_t/w_t^*}{a_t/a_t^*}\right) \tag{6}$$

$$FOC(k_t): \qquad d\ln\left(\frac{y_t/y_t^*}{k_t/k_t^*}\right) = \sigma_t d\ln\left(r_t/r_t^*\right) + d\sigma_t \ln\left(r_t/r_t^*\right) \tag{7}$$

$$Ratio: \qquad d\ln\left(\frac{k_t/k_t^*}{a_t l_t/a_t^* l_t^*}\right) = \sigma_t d\ln\left(\frac{w_t/w_t^*}{a_t r_t/a_t^* r_t^*}\right) + d\sigma_t \ln\left(\frac{w_t/w_t^*}{a_t r_t/a_t^* r_t^*}\right) \tag{8}$$

where $d\ln x_t$ is the log-difference of x_t between t and $t-1$. This system of equations is fairly standard with some additional remarks. First, since we are interested in isolating the cyclical behavior of σ_t, we take the first differences of the first-order conditions. This transformation further aims to mitigate the potential non-stationarity of our time series, which can otherwise result in a spurious regression bias of the elasticity; see Antràs (2004).[11] Since we aim to empirically purge our time series of long-run information by taking first differences, our goal is the reverse of that of other studies that explicitly focus on the long-run elasticity using low-frequency data (e.g. Chirinko and Mallick, 2017). Note also that, since the elasticity is assumed nonconstant, the first-difference transformation of the optimal demand equations implies not only the differentiation of factor prices and factor inputs but also an additional term that explicitly takes into account the differentiation of the elasticity, $d\sigma_t$, in each equation (6)-(8) which we need to consider in our estimation.

Second, in order to estimate σ_t we need data on log-deviations from trend of output y_t, labor l_t, capital k_t, wages w_t, return to capital r_t and labor-augmenting technical change a_t.[12] This can be challenging due to the unobservability of a_t. However, a theoretical result emerging from our normalized NCES-CD technology is that we can recover a_t/a^* as a residual using time series of output, labor, capital, and factor shares. We show that this is the case by totally differentiating production around the steady state, which, together with the assumption of competitive markets,

[11] There is a variety of ways in which this issue can be addressed defining what the cycle is through different filtering or detrending techniques. We use the first differences as a benchmark.

[12] Differencing the log-deviations simply means that our FOCs (6)-(8) are in deviations from average growth.

yields:

$$\ln \frac{y_t}{y_t^*} = (1 - ls^*) \ln \frac{k_t}{k_t^*} + ls^* \ln \frac{l_t}{l_t^*} + ls^* \ln \frac{a_t}{a_t^*}, \tag{9}$$

where ls^* is the steady-state (average) labor share of income. The only unknown in this equation is $\ln \frac{a_t}{a_t^*}$. We show full proof of this theoretical result that identifies the labor-augmenting technical change in Appendix A.7. Note that with, for example, additionally unobserved capital-augmenting technical change in production, we would not be able to separately identify $\ln \frac{a_t}{a_t^*}$ and the elasticity σ_t (Diamond et al., 1978). In our model, we embed production with observable capital-augmenting technical change by constructing our series of capital using investment-specific technical change (Greenwood et al., 1997, 2000) identified as the aggregate price of investment relative to consumption (Cummins and Violante, 2002; Fisher, 2006; Ríos-Rull et al., 2012).

Third, in order to estimate the cyclical properties of the elasticity of substitution, we need to choose a parametric specification of σ_t. We assume that σ_t takes two values in the data:[13]

$$\sigma_t = \begin{cases} \sigma_{NR} & \text{during non-recessions} \\ \sigma_R = \sigma_{NR} + \gamma & \text{during recessions,} \end{cases}$$

Then, our econometric specification of the system of optimal demand equations (6)-(8) becomes:

$$d\ln\left(\frac{y_t/y_t^*}{a_t l_t/a_t^* l_t^*}\right) = \sigma_{NR} d\ln\left(\frac{w_t/w_t^*}{a_t/a_t^*}\right) + \mathbf{1}_R \gamma d\ln\left(\frac{w_t/w_t^*}{a_t/a_t^*}\right) + d\sigma_t \ln\left(\frac{w_t/w_t^*}{a_t/a_t^*}\right) + \varepsilon_{1t} \tag{10}$$

$$d\ln\left(\frac{y_t/y_t^*}{k_t/k_t^*}\right) = \sigma_{NR} d\ln(r_t/r_t^*) + \mathbf{1}_R \gamma d\ln(r_t/r_t^*) + d\sigma_t \ln(r_t/r_t^*) + \varepsilon_{2t} \tag{11}$$

$$d\ln\left(\frac{k_t/k_t^*}{a_t l_t/a_t^* l_t^*}\right) = \sigma_{NR} d\ln\left(\frac{w_t/w_t^*}{a_t r_t/a_t^* r_t^*}\right) + \mathbf{1}_R \gamma d\ln\left(\frac{w_t/w_t^*}{a_t r_t/a_t^* r_t^*}\right) + d\sigma_t \ln\left(\frac{w_t/w_t^*}{a_t r_t/a_t^* r_t^*}\right) + \varepsilon_{3t} \tag{12}$$

where $\mathbf{1}_R$ is a recession dummy that takes value one during recession periods and zero otherwise; the estimate σ_{NR} pins down the elasticity in non-recession quarters; $\sigma_{NR} + \gamma$ pins down the elasticity in recession quarters; and ε's are measurement errors.

Our estimation results are in Table 1. We use quarterly U.S. time series data of output y_t, capital k_t, labor l_t, rate of return r_t and wages w_t together with productivity residual a_t—from

[13] We conduct a Monte Carlo experiment on this parametric choice in Section 2.3.2. The Monte Carlo shows that our choice is a compromise between accuracy and feasibility given the small sample of recession quarters available for estimation— forty-one NBER recession quarters in 1947.I-2019.IV.

Figure 1: NBER Recessions and a Proxy, U.S. 1948.I-2019.IV

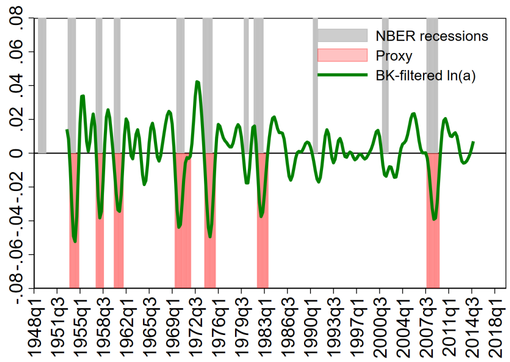

Notes: The figure compares the NBER recession periods (gray shading) with alternative recession periods (red shading) that we proxy from the BK-filtered labor-augmenting technical changes (green line). The proxy method recognizes recessions when $ln(a)$ drops for at least two consecutive quarters by more than 3.4%.

equation (9). We report OLS regression results of each optimal demand equation using the NBER definition of recessions in panel (a) of Table 1. For each of the first order conditions (10)-(12), we set the explicit difference $d\sigma_t$ to zero in our benchmark specification (column 1). Note that if this assumption is not correct, then an omitted variable bias arises in our estimation of σ_{NR} and γ. We control for this potential omitted variable bias in two ways. In a second specification (column 2) we estimate $d\sigma_t$ as a constant $d\sigma$. In a third specification (column 3) we additionally correct for the sign of $d\sigma$: $d\sigma^+ = d\sigma$ if the economy switches from a non-recessions to a recession quarter and $d\sigma^+ = -d\sigma$ if the economy switches from a recession to a non-recession quarter.

Our main finding is that σ_t is countercyclical. First, under our benchmark specification, the average non-recession elasticity, σ_{NR}, is 0.481 across first-order conditions. Results are similar by first-order condition: 0.481 for the labor condition, 0.476 for the capital condition, and 0.485 for the ratio. Controlling for the omitted variable bias with $d\sigma$ or $d\sigma^+$, we find similar average results for σ_{NR} across first-order conditions, respectively, 0.501 and 0.506. Second, we find that our estimate for γ is significant and strictly positive. In our benchmark specification, the average γ across first-order conditions is 0.162. This implies that the average elasticity during recessions, $\sigma_R = \sigma_{NR} + \gamma$, is 0.642. That is, the short-run elasticity of substitution is countercyclical. Specifically, the elasticity of substitution during recessions is 33.7% significantly larger than the elasticity during non-recession periods for all specifications. We also re-conduct our entire analysis using a proxy for NBER recessions defined as periods where a_t drops in more than two consecutive quarters by more than 3.4%; see Figure 1. We construct this ad-hoc definition, which closely overlaps with that of the NBER, to define proxy—NBER-like—recessions when using model-simulated data. Our empirical results with this proxy for NBER recessions also yield a countercyclical elasticity with similar point estimates for σ_{NR} and γ; see panel (b) in Table 1.

Table 1: Estimated σ and its Business Cycle Fluctuations (1948.I-2019.IV)

(a) NBER Recession

	$FOC(l_t)$			$FOC(k_t)$			$FOC\ Ratio$		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
σ_{NR}	0.481***	0.502***	0.507***	0.476***	0.496***	0.501***	0.485***	0.505***	0.511***
	(0.0289)	(0.0301)	(0.0304)	(0.0291)	(0.0304)	(0.0306)	(0.0288)	(0.0300)	(0.0303)
γ	0.161**	0.140**	0.129*	0.174***	0.153**	0.142**	0.152**	0.130**	0.119*
	(0.0496)	(0.0502)	(0.0508)	(0.0502)	(0.0509)	(0.0515)	(0.0492)	(0.0497)	(0.0504)
$d\sigma$		0.0261*			0.0253*			0.0266*	
		(0.0110)			(0.0112)			(0.0109)	
$d\sigma^+$			-0.0810**			-0.0802*			-0.0816**
			(0.0311)			(0.0315)			(0.0308)
Obs.	287	287	287	287	287	287	287	287	287
Adj.R^2	0.648	0.653	0.654	0.643	0.647	0.649	0.652	0.656	0.658

(b) NBER Recession Proxy

	$FOC(l_t)$			$FOC(k_t)$			$FOC\ Ratio$		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
σ_{NR}	0.485***	0.501***	0.493***	0.480***	0.496***	0.488***	0.488***	0.505***	0.497***
	(0.0306)	(0.0310)	(0.0308)	(0.0308)	(0.0312)	(0.0310)	(0.0305)	(0.0309)	(0.0307)
γ	0.127**	0.124**	0.105*	0.139**	0.136**	0.117*	0.119*	0.116*	0.0968*
	(0.0482)	(0.0478)	(0.0491)	(0.0488)	(0.0485)	(0.0498)	(0.0477)	(0.0473)	(0.0486)
$d\sigma$		0.0311**			0.0308**			0.0312**	
		(0.0109)			(0.0110)			(0.0108)	
$d\sigma^+$			-0.141*			-0.135*			-0.144*
			(0.0641)			(0.0647)			(0.0637)
Obs.	287	287	287	287	287	287	287	287	287
Adj.R^2	0.644	0.652	0.648	0.639	0.646	0.642	0.648	0.655	0.651

Notes: Table shows the estimation results of equation (10)-(12) using quarterly data series from 1948.I-2019IV (see Appendix A for detailed construction of data). The upper panel of the table demonstrates the estimation with recession dummies taken from the NBER recessions, while the lower panel takes the recessions from the log deviation of imputed labor-augmenting technical changes. A comparison of two alternative definitions of recessions are shown in Figure 1. Columns (1)-(3) indicate estimating equations (10)-(12), respectively. In practice, we also explicitly take into account the productivity structural breaks in 1974 and 2008, though this does not alter our results. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.3 Monte Carlo Experiment

We conduct a set of Monte Carlo experiments in order to assess the empirical strategy that we proposed to estimate the cyclical behavior of σ_t in Section 2.2. First, we embed a vanilla RBC model with our NCES-CD technology in order to generate model-generated data with cyclical σ_t in Section 2.3.1. Second, we test our empirical strategy by re-conducting our estimation of σ_t on the model simulated data in Section 2.3.2.

2.3.1 A Vanilla RBC Model with Cyclical σ_t

Here, we explore the effects of a countercyclical σ_t on the business cycle in a vanilla RBC model with productivity and investment shocks. We embed our proposed NCES-CD technology into a one-sector model specification of Greenwood et al. (1997, 2000) with productivity shocks and investment shocks as in Fisher (2006). Specifically, a representative household maximizes the present discounted utility:

$$\max_{\{c_t, l_t\}} E_0 \sum_t^{\infty} \beta^t \left[\ln c_t - \kappa \frac{l_t^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}} \right]$$

where c_t denotes consumption and l_t is labor supply. The parameter β is the discount factor, ν is the elasticity of labor supply, and κ is a disutility weight for working. The household's problem is subject to a resource constraint where consumption and investment (i_t) equate production (y_t) and a capital (k_t) law of motion with investment-specific technical change (υ_t) and depreciation rate δ, respectively:

$$\begin{aligned} c_t + i_t &= y_t \\ \upsilon_t \, i_t &= k_{t+1} - (1-\delta) k_t. \end{aligned}$$

Poduction follows the proposed NCES-CD technology:

$$y_t = \begin{cases} \left((1-\alpha) \left(\frac{k_t}{k_t^*} \right)^{\frac{\sigma_t - 1}{\sigma_t}} + \alpha \left(\frac{a_t l_t}{a_t^* l_t^*} \right)^{\frac{\sigma_t - 1}{\sigma_t}} \right)^{\frac{\sigma_t}{\sigma_t - 1}} (k_t^*)^{1-\vartheta} (a_t^* l_t^*)^{\vartheta}, & \text{if} \quad \widehat{a}_t \neq 1 \\ (k_t^*)^{1-\vartheta} (a_t^* l_t^*)^{\vartheta}, & \text{if} \quad \widehat{a}_t = 1. \end{cases} \tag{13}$$

Since the long run is guided by the CD component of technology, the steady-state growth theorem is preserved with ISTC. There are three exogenous shocks to the economy: shocks to productivity a_t, shocks to investment υ_t, and shocks to the elasticity of substitution σ_t. In order to introduce the countercyclicality of σ_t we use a bivariate process between a_t and σ_t:

$$\begin{bmatrix} \ln \widehat{a}_t \\ \ln \sigma_t \end{bmatrix} = \begin{bmatrix} \psi_a & 0 \\ \psi_{\sigma a} & \psi_\sigma \end{bmatrix} \begin{bmatrix} \ln \widehat{a}_{t-1} \\ \ln \sigma_{t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon_{a,t} \\ \varepsilon_{\sigma,t} \end{bmatrix}, \qquad \begin{bmatrix} \varepsilon_{a,t} \\ \varepsilon_{\sigma,t} \end{bmatrix} \sim N(0, \Sigma) \tag{14}$$

where we assume a short-run identification that shocks to σ_t do not affect a_t. We apply a Cholesky on Σ ordering a_t as the first element. The investment shocks υ_t are orthogonal to other shocks and follow an AR(1) process. Parameter values for the simulation of the model are taken from the calibration in our fully-fledged model in Section 3.

Two remarks are in order. First, the elasticity of substitution is not *per se* a source of aggregate fluctuations. That is, shocks to σ_t *per se* have no effect on the dynamics of labor and output (or on any other model variable for that matter); see the first column in Figure 2.[14] Second, the cyclical σ_t introduces a nonlinearity in production—that is not present in frameworks with constant elasticity of substitution (CES). This nonlinearity generates endogenous asymmetric responses of the labor market to productivity shocks (or, more generally, other sources of fluctuations).[15] In the second and third columns of Figure 2, we show the asymmetric impulse response functions (IRFs) of output and labor market variables to a productivity shock (of the same size, see panel (b)) across different economies that differ in their assumption of σ_t. In the responses to a productivity shock in the CES-CD framework (solid orange line), we use $\sigma = 0.5$ for all periods. Instead, in our NCES-CD economy, productivity shocks affect the dynamics of σ_t as depicted in panel (c). To highlight the asymmetries emerging in the NCES-CD economy, we plot the response to a positive productivity shock (solid blue line) and the response to a negative shock (cyan dashed line) in absolute values.

Intuitively, the asymmetry arises because higher substitutability between capital and labor—a phenomenon that with a countercyclical elasticity occurs in recessions—makes the response of labor less dependent on capital with respect to the CES-CD model and, therefore, more volatile in the NCES-CD model, see panel (c) of Figure 2. The opposite occurs in expansions when the complementarity between labor and capital increases tightening the response of labor to that of capital and making the response of labor less volatile in the NCES-CD model than in the CES-CD model. These effects on labor partly transmit—with a lower than one elasticity—to output, see panel (d). Hence, a negative productivity shock drops labor productivity by a larger amount in the NCES-CD model than in the CES-CD model, whereas the opposite occurs with a positive shock that raises labor productivity less in the NCES-CD model than in the CES-CD model. Finally, we see the differential effects of a productivity shock on wages behavior by subtracting the logged optimal demand of the NCES-CD model (σ_t) from that of the CES-CD model (σ):

$$\ln \frac{\widehat{w}_t(\sigma_t)}{\widehat{w}_t(\sigma)} = \frac{1}{\sigma} \ln \frac{\widehat{lp}_t(\sigma_t)}{\widehat{lp}_t(\sigma)} + \left(\frac{1}{\sigma} - \frac{1}{\sigma_t}\right) \ln \widehat{a}_t + \left(\frac{1}{\sigma_t} - \frac{1}{\sigma}\right) \ln \widehat{lp}_t(\sigma_t) \tag{15}$$

Since we are using the same productivity shock across models, if productivity shocks do not affect the elasticity, i.e., $\sigma_t = \sigma$ for all t, then there is no difference in the response of wages across models. Instead, a countercyclical σ_t implies that, in recessions, the dynamics of wages shifts to follow the dynamics of the productivity shock and away from labor productivity which

[14] We provide a full derivation of this theoretical result in a more general fully-fledged model in Section 3.

[15] We derive this theoretical result in our more general fully-fledged model in Section 3.

Figure 2: Effects of elasticity shocks (σ_t) and productivity shocks (a_t): Illustration from a Vanilla RBC Model

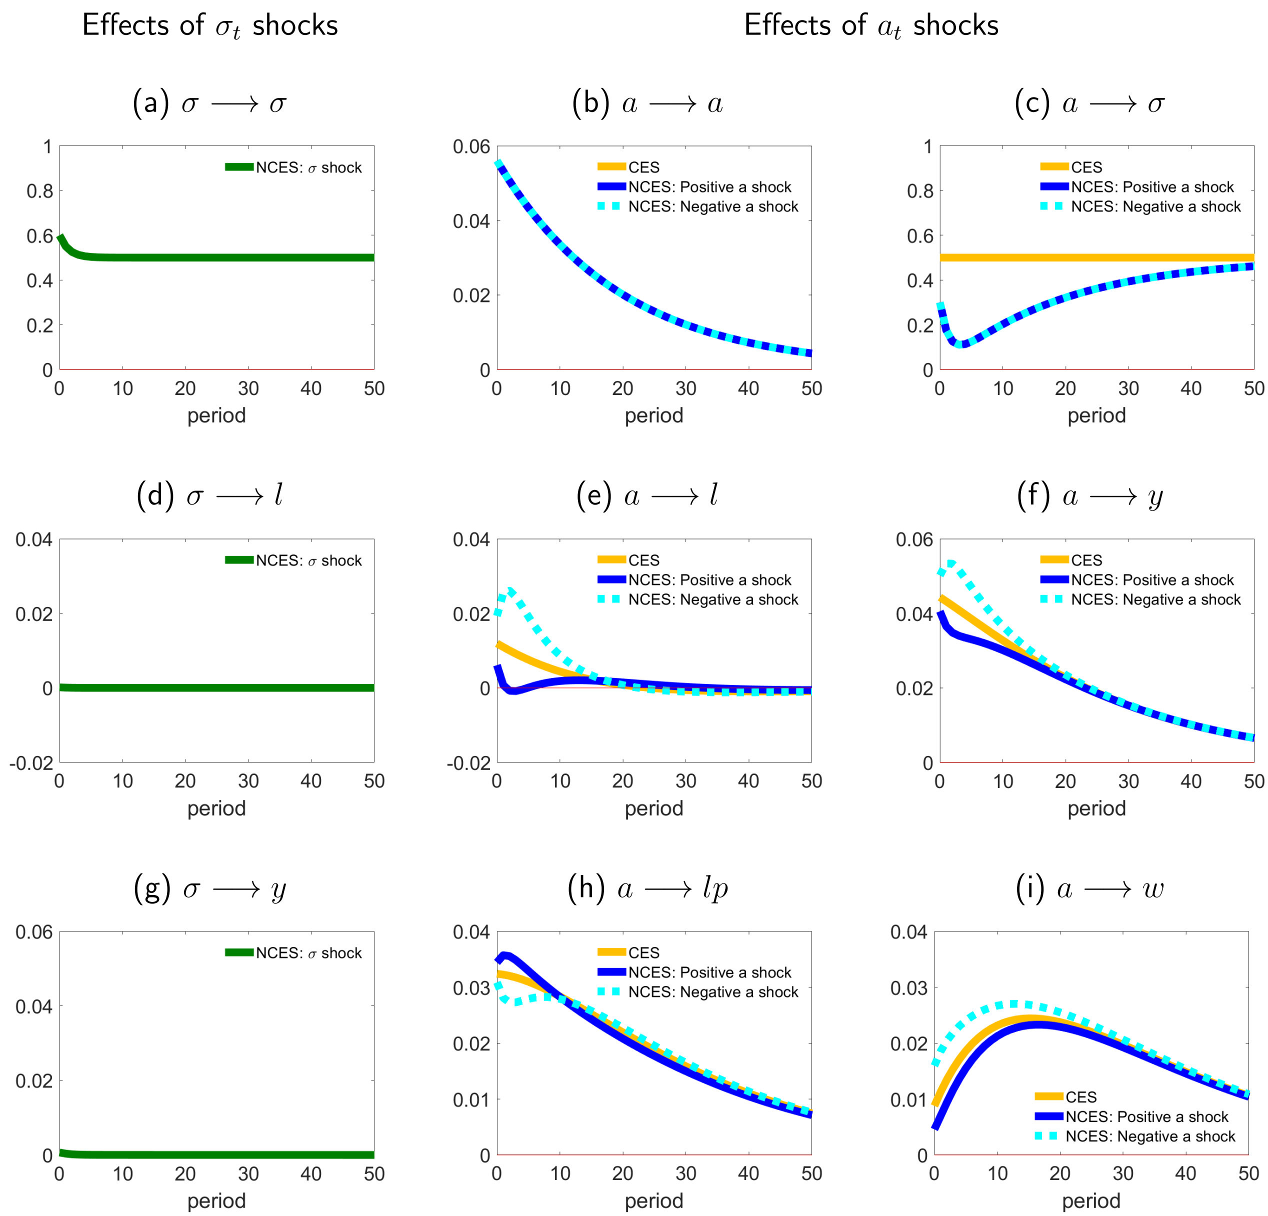

Notes: Panels in the first column show the simulated impulse responses (green line) of model variables for the first 50 periods after one-period elasticity shock σ_t. Panels in the second and third columns show the simulated impulse responses (IRF) to one-period productivity shock a_t. In each panel, the solid blue line shows the IRF to positive productivity shocks from the NCES-CD model; the cyan dashed line shows the *absolute* value of the IRF to negative productivity shocks from the NCES-CD model; the yellow line shows the IRF to the positive productivity shock with CES-CD model.

makes wages show a longer-lasting drop in the NCES-CD model than in the CES-CD model. In contrast, in expansions, the countercyclical σ_t puts more weight on labor productivity and less on the productivity shock making wages have a shorter-lasting increase.

2.3.2 Monte Carlo Results

Here, we assess whether our proposed empirical strategy for the estimation of the cyclical elasticity of substitution in Section 2.2 is warranted. To do so, we apply our econometric specification of the optimal demand equations (10)-(12) on model-simulated data from the vanilla RBC model in Section 2.3.1.[16] The idea is to test, using model-generated data, whether our estimated elasticity $\widehat{\sigma}_t$ recovers the true elasticity σ_t in the model. We are particularly interested in exploring two potential sources of bias in our estimation. First, the cyclicality of σ_t introduces nonlinearities in the responses of the endogenous variables (see our Section 2.3.1), which can affect the estimation of the average short-rung elasticity—and its cyclical properties. Second, we also assess how the small sample size—in terms of availability of recession quarters—in the data affects the accuracy of the estimation.

In panel (a) of Figure 3, we show the true frequency distribution of σ_t separately for recession and nonrecession quarters from a simulated model time series of 10,000 periods.[17] Using this binary specification, the elasticity in nonrecession years is, on average, $\sigma_{NR} = 0.450$, whereas the true elasticity in recession years is, on average, $\sigma_R = 0.863$. That is, the elasticity of substitution is countercyclical—a result that fully emerges from our model specification for productivity shocks and elasticity shocks in (14). We plot these true elasticities (red dots) on the left side of panel (b) in Figure 3. We further zoom in on the nonlinearity of the elasticity of substitution using a 10-quantile partition of the productivity shock $\widehat{a}_t$. The resulting true elasticity (red dots) by deciles are on the right side of panel (b) in Figure 3. Note that due to the nonlinearity the short-run elasticity averages 0.544—a weighted average of recession and nonrecession quarters, whereas the predetermined short-run σ in the model is 0.500 at steady state.

Is our empirical strategy able to recover the true elasticities? First, we find that with our binary specification for σ that allows for the estimated elasticity to differ between nonrecession and recession periods, we can recover estimates of the elasticity that are close to the true values of the elasticity in both nonrecession and recession periods; see our estimates alongside the true values in the left side of panel (b) in Figure 3. Most importantly for us, our estimates can recover the significant countercyclicality of the elasticity generated by the model. This is the case whether we use the FOC of labor (blue dots), FOC of capital (green dots), or their ratio (yellow dots) in the estimation. Second, the estimated elasticities can also capture the true nonlinearity of σ_t when productivity is partitioned in thinner quantiles; see the right side of panel (b) in Figure 3.

[16] In Section 3, we reconduct this experiment using our fully-fledged model with multiple sources of fluctuations.

[17] In the model, we use the proxy NBER-like recessions as defined in Section 2.2. The proportion of quarters in a recession is 15.8%, and this figure in the data is 14.24% (41 quarters out of 288 quarters).

Figure 3: Monte Carlo Experiments

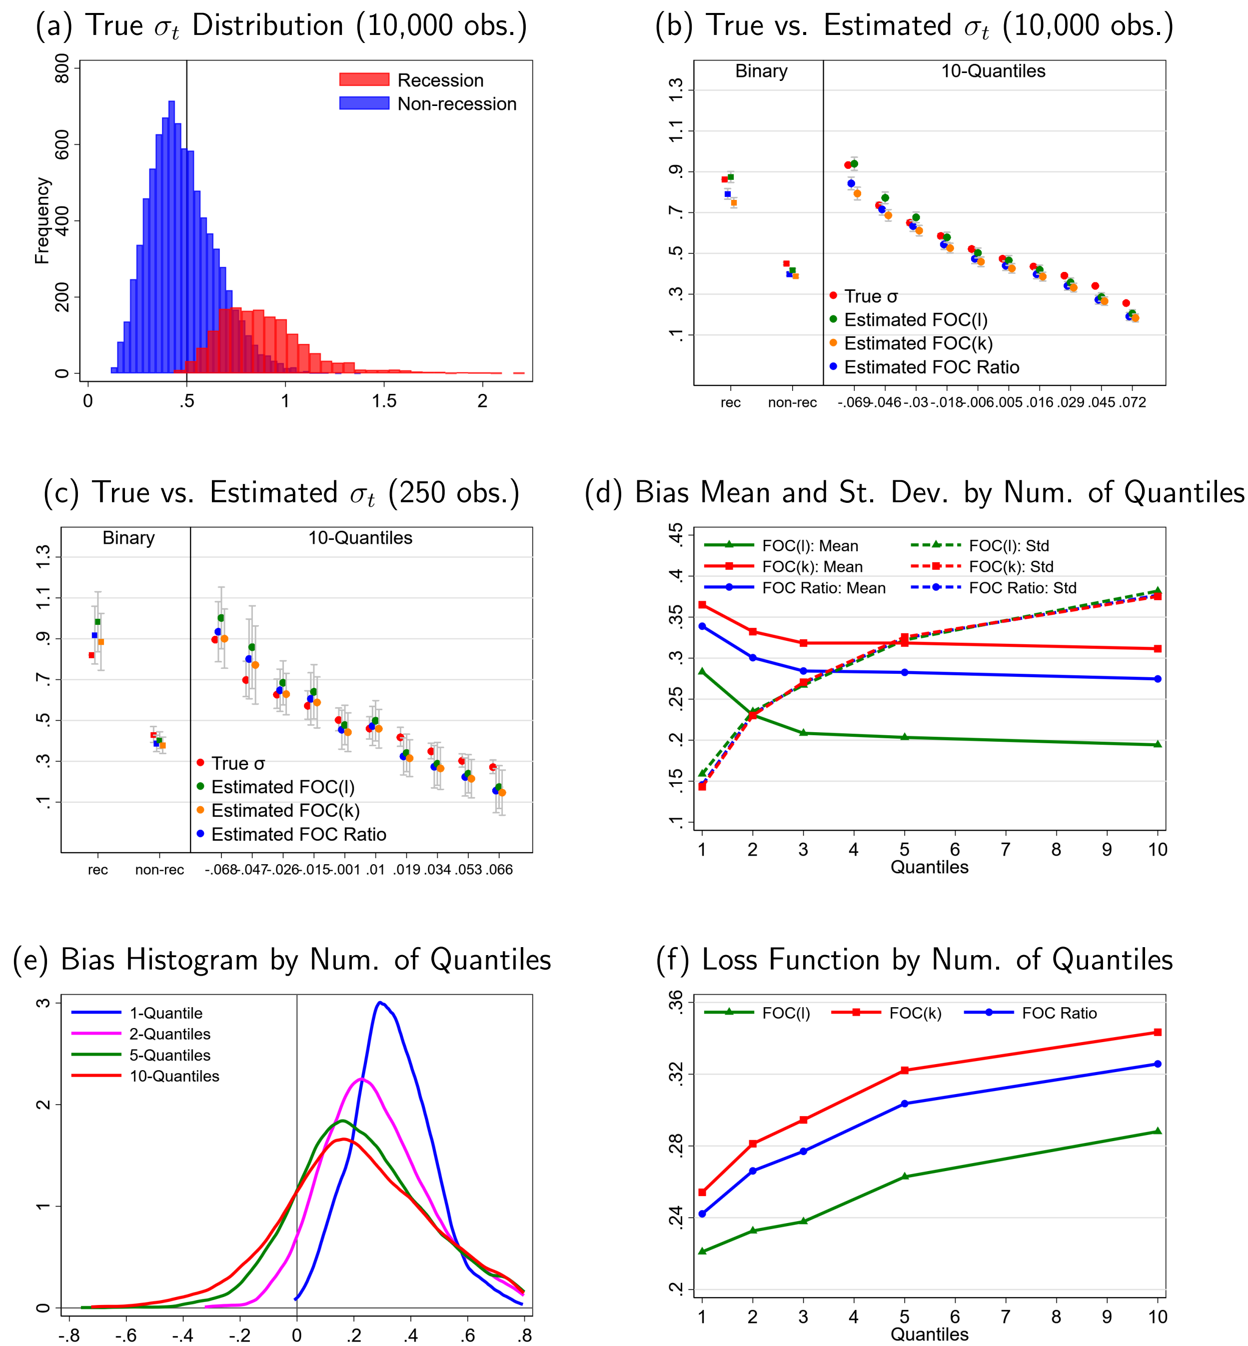

Notes: We generate 10,000 observations to estimate σ using FOCs by the number of quantiles and compute the estimation bias between FOC estimates and the true σ (i.e., 1 - $\widehat{\sigma}/\sigma$). Panel (a) shows the distribution of the estimation bias for the NBER-like recessions and non-recessions. Panel (b) and (c) show the FOC estimates of σ compared to the true σ for binary and 10-quantile cases with 10,000 simulated observations and 250 reduced observations. We conduct 1,000 simulations of 250 observations to compute the mean and standard deviation of the estimation bias, as shown in panel (d). The mean is weighted by the frequency of observations in each quantile. The histogram of estimation bias by the number of quantiles is shown in panel (e). We compute the loss function in panel (f) by weighting the mean and standard deviation of estimation bias in panel (d) by an equal weight of 0.5.

What if we reduce the sample size? In panel (c) of Figure 3, we show how our estimation results change when the sample size of the simulated time series drops to 250 observations. In this scenario, the estimates for the binary specification of σ_t are still able to capture the true differential in the elasticity across nonrecession and recession periods, despite a loss in accuracy. In contrast, when the productivity shock is split into ten quantiles, the loss in accuracy is large with confidence intervals that overlap across adjacent quantiles. Precisely, estimates of σ across contiguous (up to four) quantiles are not significantly different from each other. Hence, dropping the number of observations to a sample size that resembles the actual data availability, we find that estimating the nonlinearity of σ with a fine partition of the productivity shock significantly loses precision, whereas the binary specification still does the job.

To further explore the accuracy of our estimation, we generate 1,000-time series (samples) of 250 observations each and compute the mean and standard deviation of the estimation bias on the average elasticity over the business cycle, $\epsilon = 1 - \frac{\widehat{\sigma}}{\sigma}$. The true average σ is computed as the weighted average $\sigma = \sum_q \frac{\mu_q}{\sum_q \mu_q} \sigma_q$ where σ_q is the elasticity for quantile q and μ_q is the weight by the number of observations in quantile q. In the case of our benchmark binary specification, q takes two values, and in the case of the by-decile specification, q takes ten values. The estimated average $\widehat{\sigma}$ is computed analogously. In panel (d) of Figure 3, we find that the average bias ϵ decreases with the number of quantiles of sigma. A single quantile—i.e., a constant specification for σ_t that ignores its cyclicality—underestimates the true average σ by 28-36% across all specifications with a standard deviation across simulations of around 0.15. Using our binary specification for σ that isolates recession and nonrecession periods, we find an average bias of 23-33% with a standard deviation of 0.23 across simulations. Increasing the specification of sigma to 3 and 10 quantiles, we find that the average bias decreases to, respectively, 20-31% and 19-31%. At the same time, the standard deviation associated with adding quantiles to the specification of sigma is 0.27 for 3 quantiles and 0.38 for 10 quantiles. This increasing inaccuracy with the number of quantiles is explained by the associated shrinking sample size (within quintiles), limiting the ability to capture tails of the distribution. In sum, the average bias decreases in a convex fashion, which makes the marginal gains lower by adding quantiles, whereas the inaccuracy of the estimation increases with the addition of quantiles. Clearly, there is a trade-off between increasing the number of quantiles (lower average error) and decreasing it (estimation less dependent on tails and higher accuracy). We can also see this trade-off in panel (e), which shows the histogram of the average bias across specifications. Increasing the number of quantiles reduces the average bias closer to zero but increases the standard deviation across simulations. In this context, we find that our binary specification—our benchmark in the empirical Section 2.2—is a good compromise between the feasibility and accuracy in the estimation of cyclical σ. More formally, the minimization of a simple loss function that puts equal weight on the average and standard deviation of the error ϵ implies our choice of the binary specification; see panel (f) in Figure 3.

Endogeneity, wedges and further measurement issues Another relevant aspect of the estimation of σ_t—or any technological parameter for that matter—is the potential endogeneity of the factor inputs in production (e.g. Antràs, 2004; León-Ledesma et al., 2010). The reason is that the choice of factor inputs, particularly labor, is likely to be correlated with the realization of the productivity shock (i.e., the residual term in the production function estimation). Since we strictly base our estimation on the set of optimal demand equations, the residual term cannot be interpreted as a productivity shock. Further, recall that our NCES-CD production function formulation allows us to identify the productivity shocks independently of the elasticity; see equation (9) in Section 2.2. However, using the optimal demand conditions as we do is far from being free of caveats. Specifically, our estimated model (10)-(12) can be misspecified. That is, our specification of optimal demand might be omitting relevant variables—or wedges as in Chari et al. (2007)—that directly affect the factor input allocations. This is relevant for us because these wedges can prevent the optimal demand conditions that we use for estimation to hold, and hence our estimates of the elasticity of substitution and its cyclical properties can be biased. For example, a positive home-production shock—a labor wedge arising from the supply side (e.g. Karabarbounis, 2014)—that makes labor supply unattractive can affect the estimated value of σ_t. Similarly, capital utilization can propagate the response of factor inputs to productivity shocks (King and Rebelo, 1999), again altering the estimation of the elasticity.[18] Clearly, the reduced form of empirical evidence that we have so far provided (Section 2.2) and the associated Monte Carlo analysis is not ideal for addressing these concerns, in particular, if the wedges are of potentially unobserved nature (e.g., home productivity shocks or capital utilization) which we cannot easily control for or instrument.

In this context, our approach to addressing these issues is to provide a structural estimation for σ_t—an alternative to our reduced-form estimation in Section 2.2—using a fully-fledged business cycle model that incorporates some of the potentially unobserved wedges. The structural estimation serves two purposes. First, using the model to jointly estimate σ_t (and other technological parameters) together with potential sources of these wedges (e.g., home productivity shocks and endogenous capital utilization) in order to match a set of relevant moments in the data, we partially alleviate potentially omitted biases of the reduced form estimation. Second, the fact that our inferred σ_t is structurally estimated allows us to causally assess the effects of

[18] Further, if firms make investment choices (as opposed to households), the presence of capital adjustment costs—or any other form of intertemporal wedges—that slow down the response of capital to a productivity shock can reduce the estimated value of the elasticity of substitution (Caballero, 1994; Chirinko and Mallick, 2017). Our Appendix I extends our model by adding capital adjustment costs.

σ_t on the labor market, which is our ultimate goal.

3 A Fully-Fledged Real Business Cycle Model

We pose a competitive-market business cycle model (e.g. Cooley and Prescott, 1995; Ríos-Rull et al., 2012) with a set of standard exogenous sources of aggregate fluctuations: productivity shocks, investment shocks, government shocks, and home-production shocks. The key differential ingredient of our exercise is that we introduce a production function section that allows the short-run elasticity of substitution between aggregate capital and labor to move with the cycle. The economy is populated by a continuum of identical agents that receive utility from market consumption, home-produced consumption, and disutility from working.[19]

$$\max_{\{c_t,x_t,i_t,h_t,e_t,u_t\}} E_0 \sum_t^{\infty} \beta^t N_t \left[u(c_t) - v(h_t)e_t + m(x_t)\right] \tag{16}$$

where c_t is market consumption, h_t is the hours worked conditional on working (i.e., the intensive margin of labor supply), e_t is the fraction of days of work (i.e., the extensive margin of labor supply), $x_t = b_t(1-e_t)$ is home-produced consumption where b_t is a home-productivity shock, E_0 is the conditional expectation operator, $\beta \in (0,1)$ is a discount factor, and population grows at a constant rate λ_n. The social planner solves (16) subject to the resource constraint

$$c_t + i_t + g_t = y_t \tag{17}$$

where i_t is real investment, g_t is government expenditure, and y_t is output per capita. The law of motion of capital in efficiency units, k_t, is

$$v_t\, i_t = (1+\lambda_n)k_{t+1} - (1-\delta(u_t))k_t \tag{18}$$

where $v_t i_t$ is investment in efficiency units and v_t is investment-specific technical change.[20] We add endogenous capital utilization, a propagation mechanism emphasized in King and Rebelo (1999), with depreciation defined as an increasing function of utilization (Ríos-Rull et al., 2012).

[19] The incorporation of home-produced consumption (Benhabib et al., 1991; Greenwood et al., 1995; Chang and Schorfheide, 2003) is a feature that is directly related to the separate and explicit treatment of the intensive and extensive margins of labor supply (Cho and Cooley, 1994).

[20] As in Fisher (2006), v_t is the inverse of the relative price of investment in terms of consumption, $v_t = \frac{1}{p_t}$.

Technology. Output per capita, y_t, is produced using the aggregate production function,

$$y_t = \begin{cases} \left((1-\alpha)\left(\frac{u_t k_t}{u_t^* k_t^*}\right)^{\frac{\sigma_t - 1}{\sigma_t}} + \alpha\left(\frac{a_t e_t h_t}{a_t^* e_t^* h_t^*}\right)^{\frac{\sigma_t - 1}{\sigma_t}}\right)^{\frac{\sigma_t}{\sigma_t - 1}} (u_t^* k_t^*)^{1-\vartheta}(a_t^* e_t^* h_t^*)^{\vartheta}, & \text{if} \quad \widehat{s}_t \neq 1 \\ (u_t^* k_t^*)^{1-\vartheta}(a_t^* e_t^* h_t^*)^{\vartheta}, & \text{if} \quad \widehat{s}_t = 1 \end{cases} \tag{19}$$

where $\widehat{s}_t$ is the set of shocks $\widehat{s}_t = \{\widehat{a}_t, \widehat{\sigma}_t, \widehat{v}_t, \widehat{g}_t, \widehat{b}_t\}$. Under $\widehat{s}_t = 1$ there are no shocks—hence no aggregate fluctuations in this model, and we are back to a steady state.

Unobservable shocks. There are four shocks that we consider unobservable: $\{\widehat{a}_t, \widehat{\sigma}_t, \widehat{v}_t, \widehat{b}_t\}$. We model productivity shocks $\widehat{a}_t$ and shocks to the elasticity $\widehat{\sigma}_t$ with a joint bivariate process,

$$\begin{bmatrix} \ln \widehat{a}_t \\ \ln \sigma_t \end{bmatrix} = \begin{bmatrix} \psi_{1a} & 0 \\ \psi_{1\sigma a} & \psi_{1\sigma} \end{bmatrix} \begin{bmatrix} \ln \widehat{a}_{t-1} \\ \ln \sigma_{t-1} \end{bmatrix} + \begin{bmatrix} \psi_{2a} & 0 \\ \psi_{2\sigma a} & \psi_{2\sigma} \end{bmatrix} \begin{bmatrix} \ln \widehat{a}_{t-1} \\ \ln \sigma_{t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon_{a,t} \\ \varepsilon_{\sigma,t} \end{bmatrix} \tag{20}$$

where note that σ_t is affected by lags from $\widehat{a}_t$.[21] Further, $\epsilon_{a,t}$ and $\epsilon_{\sigma,t}$ draw from a normal joint distribution with a zero mean and variance-covariance matrix $\Sigma_{a,\sigma}$. We orthogonalize a la Cholesky assuming that innovations to a_t affect the elasticity but not the other way round. Hence, the effects of a_t on σ_t occur through the mixed persistence parameters $(\psi_{1a\sigma}, \psi_{2a\sigma})$ and the off-diagonal element of the orthogonalized variance-covariance matrix, $\upsilon_{a,\sigma}$. As discussed below, the fact that σ_t does not affect a_t ensures that σ_t is *per se* not a source of fluctuations.

We treat investment shocks, v_t, as unobservable. We assume that v_t is exogenously given and follows the AR(2) process with a linear trend λ_v and the initial value $v_0 = 1$, then

$$\ln \widehat{v}_t = \psi_{1v} \ln \widehat{v}_{t-1} + \psi_{2v} \ln \widehat{v}_{t-2} + \varepsilon_{v,t} \quad \text{with} \quad \varepsilon_{v,t} \sim iid(0, \upsilon_v^2) \tag{21}$$

Finally, the unobservable home-productivity shock follows,

$$\ln \widehat{b}_t = \psi_{1b} \ln \widehat{b}_{t-1} + \psi_{2b} \ln \widehat{b}_{t-2} + \varepsilon_{b,t} \quad \text{with} \quad \varepsilon_{b,t} \sim iid(0, \upsilon_b^2) \tag{22}$$

Observable shocks. We assume that g_t follows a time-varying fraction of total output, $g_t = (1 - 1/\tau_t)y_t$ and τ_t. This government spending follows,

$$\ln \tau_t = (1 - \psi_g) \ln \tau_t^* + \psi_g \ln \tau_{t-1} + \varepsilon_{g,t} \quad \text{with} \quad \varepsilon_{g,t} \sim iid(0, \upsilon_g^2) \tag{23}$$

[21] The AR(2) allows us to capture potential humps/bumps in IRFs to own innovations (Ríos-Rull et al., 2012).

Definition 1 [Equilibrium] *An equilibrium for this economy is the sequence of optimal allocations, $\{c_t, x_t, i_t, h_t, e_t, u_t\}_{t=0}^{\infty}$, that solve (16) subject to the resource constraint (17), the aggregate capital law of motion (18), the NCES-CD production function (19), the productivity and elasticity shocks (20), the investment shock process (21), home-productivity shocks (22), the process of government expenditure (23), and initial values of shocks and aggregate capital.*[22]

Definition 2 [Steady-State Equilibrium] *A steady-state equilibrium for this economy is an equilibrium (as defined above) with no shocks, that is, under $\widehat{s}_t = 1$.*

We highlight two theoretical results.

Result 1. Cyclical σ_t is NOT a source of aggregate fluctuations *per se*, but propagates the effects of other sources of aggregate fluctuations. To see this, write compactly the production function (19) dropping time subscripts as,

$$F(\mathbf{x}, \sigma) = f^{SR}(\widehat{\mathbf{x}}, \sigma)\, f^{LR}(\mathbf{x}^*, 1), \tag{24}$$

where $\mathbf{x}$ is an m-dimensional vector of elements including all the factor inputs of production and labor-augmenting technical change; $f^{SR}(\widehat{\mathbf{x}}, \sigma)$ is the short-run component of the production function; and $f^{LR}(\mathbf{x}^*, 1)$ is the long-run component of the production function with input factors evaluated at the steady state and with an elasticity of substitution equal to one. We consider sources of fluctuations the exogenous variables that move the endogenous variables in $\mathbf{x}$ away from steady state. The following theorem shows that innovations to σ_t are not a source of fluctuations, but they affect the propagation of other sources of fluctuations. This result is general to any business cycle model with a short- and long-run decomposition of production as that follows (24).

Theorem 1. *In an nth-order approximation around the steady state of f defined in equation* (24) *where σ does not impact exogenous sources of fluctuations:*

(a) Shocks to the elasticity of substitution, σ, do not have any impact on f through the partial derivatives $\frac{\partial^n f}{\partial \sigma^n}\Big|_*$ of any order $n > 0$. That is, shocks to σ are not a source of fluctuations.

(b) If $n > 1$ then σ-shocks propagate the effects of other sources of fluctuations on f through the higher-order cross partial derivatives of f between σ and $\widehat{\mathbf{x}}$.

Proof. See Online Appendix D. □

[22] See the entire set of model equilibrium conditions and the model stationarization in Appendix E.

Intuitively, part (a) of Theorem 1 hinges on that, by construction, the short-run component of production does not deviate from the steady state in response to σ_t. Since σ_t appears solely as the exponent of the short-run production (and factor inputs), its effect vanishes when the base is one, which occurs at the steady state. Hence, the ability of σ_t *per se* to generate aggregate fluctuations (i.e., deviations from the steady state) is nil. However, σ_t can propagate the effects of actual sources of fluctuations (e.g., a productivity shock). Precisely, if the economy is not at the steady state (i.e., $\widehat{x} \neq 1$) due to a productivity (or other) shock (e.g., $\widehat{a} \neq 1$), then σ_t can propagate the effects of that shock. In part (b) of Theorem 1, we show that this propagation power of σ_t occurs through the higher-order cross partial derivatives.

Result 2. Cyclical σ_t generates asymmetric business cycles. An nth order (with $n > 1$) approximation of production—as implemented in Theorem 1—shows that our NCES-CD technology introduces potential nonlinearities due to σ_t. A second-order approximation of production suffices to highlight this effect:

$$P_2(\widehat{\mathbf{x}}, \sigma) = f|_* + \sum_{i}^{m} f_{\widehat{x}_i}|_* (\widehat{x}_i - \widehat{x}_i^*) + \frac{1}{2!}\left[\sum_{i=1}^{m}\sum_{j=1}^{m} f_{x_j x_i}\big|_* (x_i - x_i^*)(x_j - x_j^*) + 2\sum_{i=1}^{m} f_{x_i \sigma}|_* (x_i - x_i^*)(\sigma - \sigma^*)\right]$$

where we denote $f_{x_j x_i} = \frac{\partial^2 f}{\partial x_i \partial x_j}$ for all pairs (i, j) as the elements in the submatrix of the Hessian of f that includes all but its last column and row, the term $f_{x_i \sigma}$ captures the cross partial derivatives between $\widehat{\mathbf{x}}$ and σ (i.e., the off-diagonal elements in the last row and column of the Hessian). Note that in deriving $P_2(\widehat{\mathbf{x}}, \sigma)$ we used Theorem 1(a) setting the first and the second-order partial derivative of f with respect to σ to zero (i.e., the diagonal bottom-right element of the Hessian). Hence, deviations of σ from steady state introduce a nonlinearity in production through the nonzero cross-derivative $f_{x_i \sigma}$. The nonzero cross-derivative generates an asymmetry in the response of the endogenous model variables to sources of fluctuations. In particular, if a negative productivity shock a_t is accompanied by an increase in σ_t (i.e. higher substitution between capital and labor in recessions), then the response of labor to a_t is likely to be more elastic (in absolute terms) than that of a positive productivity shock to a_t which reduces σ_t (i.e. higher complementarity between capital and labor in expansions).

Since our production is not linear in σ_t, a solution method based on a first-order approximation would ignore the potential role of σ for the business cycle. For this reason, we need to use higher-than-first-order perturbation methods(or other global methods, for that matter) to solve our model. In addition, we have to solve this model many times for structural estimation. We find that a third-order perturbation method is well suited for our purposes, displaying a good balance

Table 2: Calibrated Parameters

Parameters	Value	Target	Value
β	.9944	δ_0	.013
δ_1	.0034	i/y	.28
ϑ,α	.66	labor share	.66
κ_1	59.86	h^*	.31
λ_k	.0111	λ_v	.0064
λ_a	.0014	λ_y	.0047

Notes: The choice of calibrated parameters is discussed in section 4. Each of these calibrated parameters targets a long-run moment in quarterly terms.

between speed and accuracy. [23]

4 Estimation

We pose a calibration strategy that identifies some model parameters so that steady-state allocations match the long-run behavior of U.S. data. Then, we structurally estimate the rest of the model parameters using standard simulated method of moments (SMM) techniques.

Calibrated Parameters For consistency with balanced growth path, we assume log-utility in consumption, $u(c)$, and CRRA functions for $m(x_t)$ and $v(h)$. A social planner maximizes

$$\max_{\{c_t,k_{t+1},h_t,e_t,u_t\}} E_0 \sum_t^{\infty} \beta^t N_t \left[\ln c_t - \kappa_1 \frac{h_t^{1+\frac{1}{\nu_1}}}{1+\frac{1}{\nu_1}} e_t + \kappa_2 \frac{(b_t(1-e_t))^{1+\frac{1}{\nu_2}}}{1+\frac{1}{\nu_2}} \right] \tag{25}$$

where $\nu_1 > 0$ is the elasticity of intensive margin of labor supply, $\nu_e = -\nu_2 \frac{1-e}{e} > 0$ is the elasticity of extensive margin of labor supply, and κ_1 and κ_2 are the utility weight parameters. To choose the elasticity of labor supply of the intensive margin, we use the micro-estimate of .72 that takes into account a second earner (Heathcote et al., 2014). Then, from the first-order condition of hours per worker, we obtain κ_1=59.86 by targeting average long-run hours per worker to h_0=.31. Using a value for the long-run (average) labor share ϑ=.66, the steady-state private consumption-to-output ratio c_0/y_0=.39, and a long-run employment per capita e_0=.46.

Further, using the average growth rate of the inverse of the relative price of investment in terms of consumption (Fisher, 2006), which we directly retrieve from the data using a quarterly quality-adjusted price index for investment in equipment and a price index for consumption

[23] Our model contains five shocks, and four of them exhibit AR(2) dynamics. This implies a total of ten state variables—including the endogenous capital stock. Aruoba et al. (2006) provide a comprehensive analysis of perturbation methods and additional methods with a standard RBC model of three state variables, capital, productivity, and investment shocks. They find that higher-order perturbation methods (at least second-order) display a superior performance over alternative methods in terms of accuracy, speed, and the amount of coding.

(see Appendix A), we obtain λ_v=0.0064, which together with λ_y=.0047, and $\vartheta = .66$ implies λ_a=.0014. Further, note that with ISTC capital (in efficiency units) and output do not grow at the same rate; specifically, $(1+\lambda_k) = (1+\lambda_y)(1+\lambda_v)$ and λ_k=.0111.

Our capital depreciation positively depends on capital utilization rates, with

$$\delta(u_t) = \delta_0 + \delta_1(u_t^{1+1/\zeta} - 1), \tag{26}$$

as in Ríos-Rull et al. (2012). This form has a property that $\delta(u_t)|_{u_t=u^*} = \delta_0$ at steady state when we normalize $u^* = 1$. We choose $\delta_0 = .0013$, which is the average depreciation in Cummins and Violante (2002) for capital series that are adjusted for investment-specific technical change. Given capital growth λ_k =.0111, population growth λ_n =.0025, depreciation δ_0 =.0013 and $\frac{i_0}{y_0}$ =.28, the quarterly ratio $\frac{k_0}{y_0}$ =10.5 can be pinned down using the law of motion of capital (in efficiency units) along balanced growth. We have normalized $v_0 = 1$. Further, the rate of return (net of depreciation) is $r_0 = (1-\vartheta)\frac{y_0}{k_0}$ =.0324. To identify β we use the consumption Euler equation along balanced growth, $(1+\lambda_n)(1+\lambda_k) = \beta\,(r_0 + 1 - \delta_0)$, which yields a β=.9944. Given $\widehat{ls}_t = 1$ on average, $ls_0 = \frac{w_0 e_0 h_0}{y_0} = \alpha$. This implies that α is equal to long-run labor share, i.e. $\alpha = \vartheta$ =.66. Finally, we borrow the estimates for g_t process ψ_g =.95 and υ_g =.007 from Ríos-Rull et al. (2012). Table 2 shows the calibrated values described above.

Estimation Method We use SMM to estimate the remaining set of parameters. First, we choose a set of targeted moments, ϕ^d, which consists of output per capita $var(y)$ and labor share $var(ls)$; the correlation between wages and output per capita, $\rho(w, y)$; the correlation between labor productivity and output per capita, $\rho(lp, y)$; the correlation between labor share and output per capita, $\rho(ls, y)$; and the correlation between employment per capita and average hours per worker, $\rho(e, h)$. We also target the long-run (steady-state) average of employment per capita, e_0. In addition, we target the data IRFs computed by the estimated trivariate AR process of productivity, labor share and output per capita. We choose to target a set of dynamic multipliers of $\Theta_{z,z} = \{\theta_{z,z}^{(s_1)}, \theta_{z,z}^{(s_2)}, \dots, \theta_{z,z}^{(s_n)}\}$ and $\Theta_{z,ls} = \{\theta_{z,ls}^{(s_1)}, \theta_{z,ls}^{(s_2)}, \dots, \theta_{z,ls}^{(s_n)}\}$ where $\theta_{u,x}^{s} = \frac{\partial x_{t+s}}{\partial u_t}$ is the response of x at period $t+s$ to a shock u at t. The periods of dynamic multipliers that we choose to match are associated with the following eight periods, $\{s_1 = 1, s_2 = 2, s_3 = 5, s_4 = 10, s_5 = 20, s_6 = 30, s_7 = 40, s_8 = 50\}$, balancing initial responses and potential long-lasting impacts.

Our set of estimated parameters, γ, contains the short-run elasticity of substitution σ at steady state (1 parameter); the joint dynamics of productivity shocks a_t and shocks to the elasticity σ_t (9 parameters) (see equation (20)); the properties of investment shocks, v_t (3 parameters, equation (21)); home-productivity shocks, b_t (3 parameters, equation (22)); the elasticity of the extensive margin of labor supply and its disutility weight, respectively, ν_2 and κ_2 (2 parameters); and the

Table 3: Structurally Estimated Parameters

Parameter	Value	95% C.I.	Parameter	Value	95% C.I.
Elasticity:			Shocks to b:		
σ	0.5981	[0.5943,0.6019]	ψ_{1b}	1.0370	[-0.2809,2.3550]
			ψ_{2b}	-0.0619	[-1.2392,1.1153]
Shocks to a and σ:			υ_b	0.0031	[-0.0003,0.0065]
ψ_{1a}	0.9536	[0.9518,0.9554]			
ψ_{2a}	-0.0595	[-0.0618,-0.0571]	Shocks to υ:		
υ_a	0.0090	[0.0088,0.0092]	$\psi_{1\upsilon}$	1.6763	[1.6762,1.6763]
$\psi_{1\sigma}$	0.9830	[0.0394,1.9265]	$\psi_{2\upsilon}$	-0.6808	[-0.6809,-0.6808]
$\psi_{2\sigma}$	-0.2049	[-0.7763,0.3666]	υ_υ	0.0035	[0.0034,0.0036]
υ_σ	0.0145	[-0.0208,0.0499]			
$\psi_{1a\sigma}$	-6.8144	[-8.2596,-5.3691]	Home-production utility:		
$\psi_{2a\sigma}$	2.9812	[-2.4067,8.3691]	ν_2	-5.5979	[-6.6821,-4.5138]
$\upsilon_{a\sigma}$	-0.4827	[-1.4951,0.5298]	κ_2	1.9096	[1.8738,1.9455]
Elasticity w.r.t utilization:					
ζ	0.1192	[0.0926,0.1457]			

Notes: The variance of estimates are computed by $V = \frac{1}{T}\left(DW^{-1}D'\right)^{-1}$ where $D = \frac{\partial(\phi^d - \phi^m(\gamma))}{\partial\gamma}$ and W is the weighting matrix that we specified in section 4. The standard errors are the square root of the diagonals of V.

elasticity of capital depreciation with respect to utilization, ζ (1 parameter).

Given a guess for γ, we solve and simulate the model economy to obtain a set of model-generated moments, $\phi^m(\gamma)$, associated with the targeting moments.[24] We iterate until we find γ that minimizes the distance between the model moments $\phi^m(\gamma)$ and the data moments ϕ^d:

$$J = \min_{\gamma}\ \left[\phi^d - \phi^m(\gamma)\right]' W^{-1} \left[\phi^d - \phi^m(\gamma)\right]$$

where we use an estimated variance-covariance matrix of ϕ^d (Ω^d) adjusted by the number of simulations (N) and samples (T) as an optimal weighting $W = (1 - T/N)\Omega^d$.[25]

Estimation Results Our estimation results are in Table 3. The short-run elasticity of substitution at steady state is σ=.5981 with a 95% confidence interval [.5943,.6019].[26] This shows that the short-run component of the production function is significantly more complementary than Cobb-Douglas. In our setting, business cycle moments are, ceteris paribus, highly sensitive to σ, posing an ideal ground for its identification. First, the variance of output and labor share bound σ between .55 and .79 (see panel (a) in Figure 4). Second, the targeted labor market comovements

[24]We simulate the model economy 10,000 periods and drop the first 1,000.

[25]We estimate the full variance-covariance matrix using bootstrap methods (Lee and Ingram, 1991; Bloom, 2009). This implies a total of 23×23 elements. Due to the high degree of nonlinearity of our model, we proceed to incorporate the off-diagonal for efficiency reasons (Ruge-Murcia, 2012).

[26]The confidence intervals are computed using the standard errors as the square root of the diagonal of V where $V = \frac{1}{T}\left(DW^{-1}D'\right)^{-1}$ where W is the weighting matrix and $D = \frac{\partial(\phi^d - \phi^m(\gamma))}{\partial\gamma}$ is the response of the loss function to changes in the set of estimated parameters.

Figure 4: Identification of σ: Sensitivity of Targeted Moments

(a) Ratio of Model to Data Variances (b) Correlations with Output

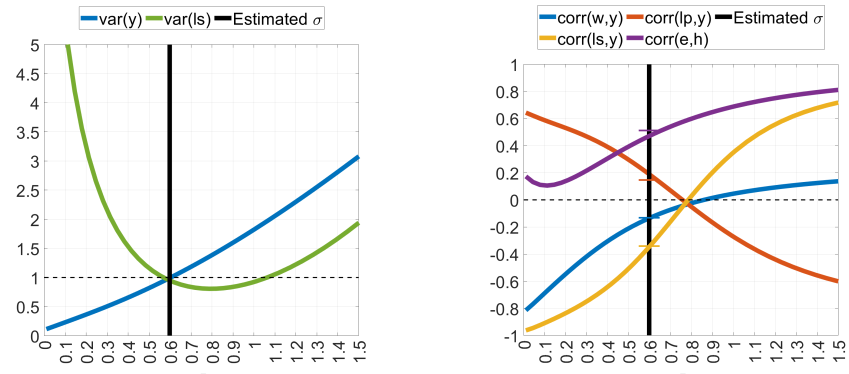

Notes: Panel (a) shows the behavior of model HP-filtered variances of output and labor share (relative to data counterparts) in response to *ceteris paribus* changes in the average short-run elasticity of substitution $\sigma \in (0,1.5]$. The relative variance $\widetilde{\upsilon_x}$ of a variable x is defined as the model-generated variance of that variable divided by its data counterpart; hence, a value of $\widetilde{\upsilon}_x$ =1.00 implies that the model exactly matches the data. Panel (b) shows the HP-filtered correlations of model labor market variables with output where tick marks indicate the value of the correlations in the data.

are strictly monotonic in σ. Hence, there are respective unique values of σ that pin down each of these correlations and these are close to our estimated σ (see panel (b) in Figure 4).

To be consistent with the empirical literature on the dynamics of labor share (Ríos-Rull and Santaeulàlia-Llopis, 2010), we use the IRF of labor share in response to neutral productivity shocks z_t (i.e., model-generated total factor productivity) to discipline the properties of a_t shocks captured by ψ_{1a}, ψ_{2a} and υ_a.[27] The top panel of Figure 5 shows, ceteris paribus, the effects of ψ_{1a} (left panel), ψ_{2a} (center panel) and ν_a (right panel) on the IRF of labor share to productivity shocks. The initial drop of labor share in response to productivity (to -.0022 log points in the data) pushes for estimated values close to one for ψ_{1a} and to zero for ψ_{2a}, while favoring values of υ_a around .0070. The peak of the labor share response (to .0021 log points at the 26th quarter) is informative about the value of ψ_{a2} close to -.1; see also the bottom panels of Figure 5 that shows the associated counter sets of the IRFs under study. The period of overshooting around the 6th quarter pushes ψ_{1a} and ψ_{2a} somewhat below but close to, respectively, one and zero. In all, our algorithm settles for ψ_{1a} =.9536, ψ_{2a} = -.0595 and υ_a.0090.

The model also replicates untargeted IRFs to neutral productivity shocks, in particular, the

[27] As in Ríos-Rull and Santaeulàlia-Llopis (2010), neutral productivity is model-generated total factor productivity computed under the assumption of no investment-specific technical change and no capital utilization. To construct this residual z, we apply algorithms described in Appendix F to the model. The fact that we use model-generated series of output, labor, investment, and factor shares to recover z implies that the productivity z is endogenous. We discuss the construction of neutral productivity shocks in Appendix A.6. Our approach to match IRFs is standard (e.g. Christiano et al., 2005; Altig et al., 2011).

Figure 5: Identification of ψ_{1a}, ψ_{2a} and υ_a: IRFs of Labor Share to TFP shocks

(a) $\psi_{a1}: \quad z \rightarrow ls$ (b) $\psi_{a2}: \quad z \rightarrow ls$ (c) $\upsilon_a: \quad z \rightarrow ls$

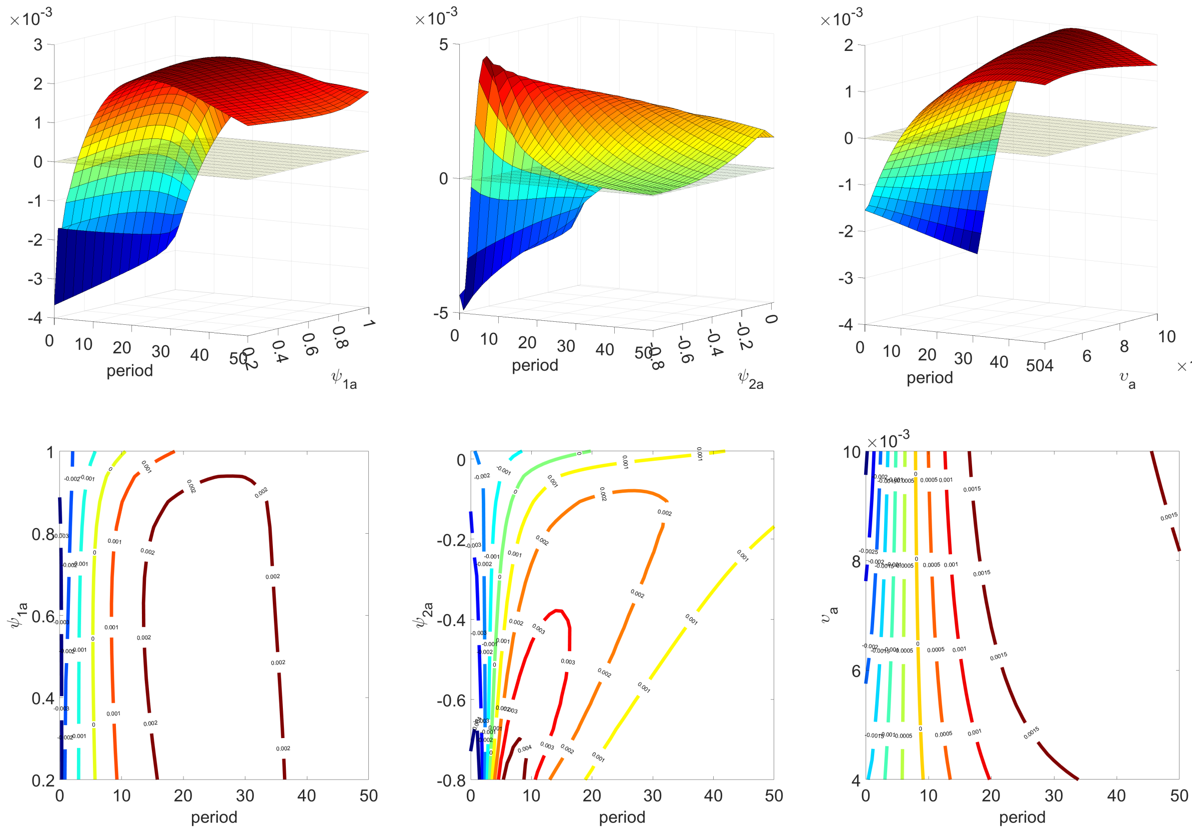

Notes: The top panels show the model IRF of labor share in logs in response to one standard deviation TFP shocks z_t for *ceteris paribus* changes in the properties of the productivity shocks a_t: ψ_{a1} in top-left panel (a), ψ_{a2} in top-center panel (b) and υ_a in top-right-panel (c). The bottom panels show the associated contour sets of the IRFs of interest.

overshooting response of labor share. To compute these IRFs we recover model-generated total factor productivity z_t from simulated data. Then, we estimate a trivariate process of total factor productivity z_t, labor share, and output (in log deviations from trend) identified in that order with a Cholesky decomposition. The resulting model IRFs are in Figure 6. In all cases, the model IRFs closely follow their data counterparts. Our main success is that labor share displays the *overshooting* property of the same size and length that we observe in the data. First, in response to a productivity innovation, labor share initially drops to -.0026 points. Second, labor share rapidly increases after the initial drop to reach a maximum of .0019 log points above the mean before the 30th quarter. Third, this response of labor share persists in keeping itself alive at .0017 log points above the mean for 50 quarters after the productivity innovation occurred.[28]

[28]Recent literature on labor share considers matching this *overshooting* response of labor share endogenously in settings that move away from frictionless RBC models through search frictions (Choi and Ríos-Rull, 2009; Moscarini and Postel-Vinay, 2008; Justiniano and Michelacci, 2012; Yashiv, 2012), optimal contracts (De Graeve et al., 2010), Schumpeterian models (Francois and Lloyd-Ellis, 2009), reallocation shocks (Lettau and Ludvigson, 2013), firm entry costs (Shao and Silos, 2014), countercyclical markups (Colciago and Rossi, 2015), and search

Figure 6: Data vs. Model IRFs to TFP Shocks

(a) $z \to z$ (targeted) (b) $z \to ls$ (targeted) (c) $z \to y$ (non-targeted)

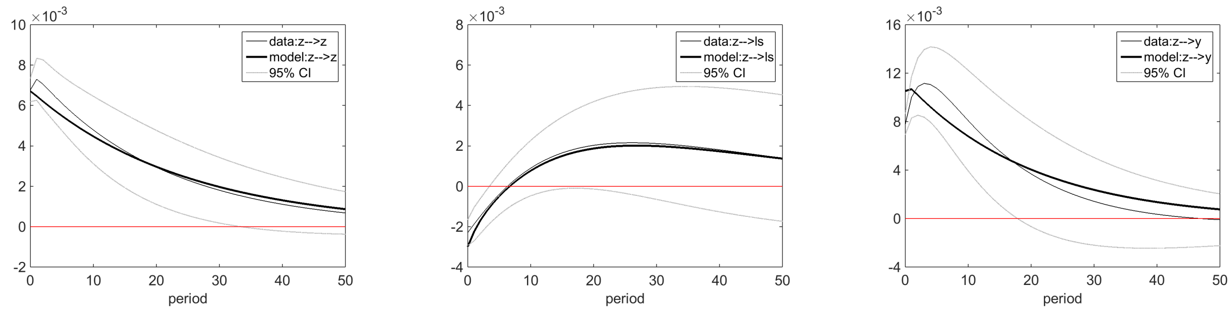

Notes: Data and model IRFs of TFP (top panel), labor share (center panel) and output (bottom panel) in logs in response to one standard deviation TFP shocks. See a discussion of these results in Section 4.

To describe the estimated joint dynamics between a_t and σ_t, we report the effects of productivity on the elasticity of substitution in Figure 7. After a productivity innovation a_t, the elasticity of substitution σ_t drops to -.01 log points; it keeps sharply dropping to reach a minimum at -.11 log points in the 5th quarter, after which it slowly converges to mean from below in a concave fashion. Note that the effects of a_t on σ_t generate a stronger response of σ_t than its own innovations; indeed, 98% of the fluctuations in σ_t are generated by productivity innovations a_t. The estimated effects of productivity a_t on the elasticity σ_t will turn out to be crucial to understanding labor market behavior, an issue that we discuss in detail in Section 6.

Regarding investment shocks, our algorithm settles for ψ_{1v} =1.6763 and ψ_{2v} =-.6808, which is a highly persistent estimation with $|\psi_{1v} + \psi_{2v}|$ =.9955 and υ_v =.0035.[29] Note that confidence intervals on the properties of home-productivity shocks suggest that these are not significant. At the same time, home-productivity shocks are particularly relevant for the targeted correlation of employment per capita and hours per worker; see our Appendix H. In terms of the extensive margin of labor supply, note that the first-order condition of employment per capita sets the relationship between the elasticity of the extensive margin of labor supply ν_2 and the utility weight κ_2, that

and matching theory (Mangin and Sedlacek, 2018). Even though some of these models are successful generating the overshooting effect of labor share in response to productivity shocks, the size of overshooting remains less than 1/10 of the actual overshoot in the data with the notable exception of Choi and Ríos-Rull (2020) that introduce a bias in productivity shocks towards new plants in the context of a putty-clay framework with competitive wage setting. More recently, Cantore et al. (2021) also document the response of the labor share to a monetary shock and discuss the ability of New-Keynesian models in replicating it.

[29] Overall, the structurally estimated values are largely consistent with the data estimates. We discuss differences between our benchmark results and those attained under the assumption that $\widehat{\upsilon}_t$ is observable—i.e., entirely identified by the relative price of investment—in our Appendix G.

Figure 7: The IRF Dynamics of the Elasticity of Substitution, σ_t

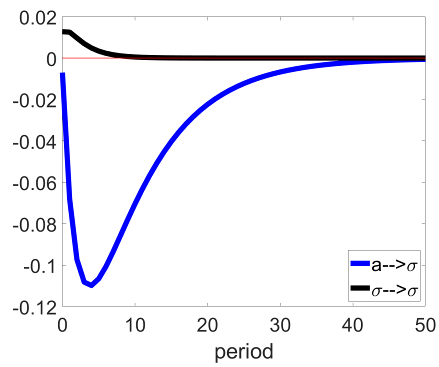

Notes: Model IRFs of the elasticity of substitution σ_t in level in response to one standard deviation productivity shocks a_t and its own shocks. See the identification assumptions used to generate these IRFs from our model in Section 3 and a discussion of these results in Section 4.

is,

$$(1-e) = \left(\frac{1}{b}\right)^{\nu_2+1} \left(\frac{1}{\kappa_2}\right)^{\frac{1}{\nu_2}} \left(\frac{1}{\nu_1} v(h)\right)^{\nu_2}. \tag{27}$$

The fact that our choice of κ_2 depends on the value of ν_2, which, in turn, governs the cyclical behavior of e, explains why we add both κ_2 and ν_2 to the set of parameters to be structurally estimated. We find ν_2 and κ_2 useful in delivering the long-run employment per capita e_0 and the cyclical correlation between employment per capita and hours per worker $\rho(e, h)$. The effects of parameter combinations of κ_2 and ν_2 on the long-run employment per capita e_0 are in panel (a) of Figure 8 and on the cyclical correlation $\rho(e, h)$ in panel (b) of Figure 8. The flat planes in each of those figures is the actual targeted value, $e_0 = .46$ and $\rho(e, h) = .51$. Regarding e_0, we find that κ_2 must be between 1.5 and 2, while the whole spectrum depicted for ν_2 serves the purpose of getting e_0. Regarding $\rho(e, h)$, κ_2 must be between 1.5 and 2.5 and ν_2 must lie between -4 and -8. That is, there is a bounded region for combinations of κ_2 and ν_2 that delivers both moments e_0 and $\rho(e, h)$ at the same time. Within this region, our algorithm picks a ν_2 =-5.5979. This implies an elasticity of the extensive margin of labor supply of $\nu_e = -\nu_2 \frac{1-e}{e}$ =6.5714 with an implied 95% confidence interval of [5.2988,7.8442]. The utility weight associated with not working in the market, κ_2, is 1.9096.[30]

Finally, the elasticity of depreciation with respect to utilization is ζ =0.1192. This result falls in the lower end of the range for values of ζ explored in King and Rebelo (1999) and Ríos-Rull

[30] A larger elasticity of labor supply in the extensive margin than in intensive margin is a common result in the literature; see Keane and Rogerson (2012) for a comprehensive discussion. Chetty et al. (2012) show different values for this elasticity settling for a benchmark extensive margin elasticity of 2.77.

Figure 8: Identification of κ_2 and ν_2: Sensitivity of Targeted Moments

(a) Employment Rate: e_0 (b) Cyclical correlation: $\rho(e,h)$

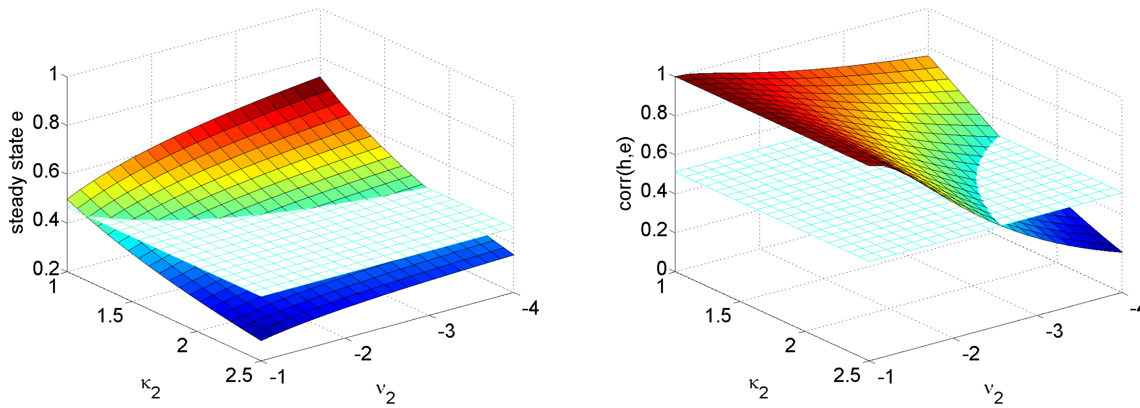

Notes: Panel (a) shows the combinations of κ_2 and ν_2 that deliver the steady-state employment per capita $e_0 =$ 0.46 (i.e., the flat plane in the figure). Panel (b) shows the combinations of κ_2 and ν_2 that deliver the correlation of employment per capita and hours per worker $\rho(e,h) =$ 0.51 (i.e., the flat plane in the figure).

et al. (2012). We find that ζ has strong implications for our selection of targeted business cycle moments, in particular for correlations of labor market variables with output (see our Appendix H).

5 Quantitative Results

Our NCES-CD model generates the size of the aggregate fluctuations and the comovements of labor market variables with output and among themselves that we observe in the data. Table 4 shows a battery of standard business cycle moments of the U.S. 1948.I-2019.IV and the model counterparts.[31] Our model fully captures the aggregate fluctuations of output, y, and also its persistence.[32] Regarding the behavior of the labor input, our model accounts for almost all fluctuations in hours per capita eh, 3.32×100/3.64 = 91%. We can decompose the variance of hours per capita in its two components, employment per capita e (i.e., the extensive margin of labor supply) and average hours worked per worker h (i.e., the intensive margin of labor supply) using $var(eh) = var(e) + var(h) + 2cov(e,h)$. Splitting the contribution of the covariance equally between the e and h implies that the model e accounts for (3.04+.125)/3.32 = 95% of the total fluctuations of hours per capita, and h accounts for the remaining (.03+.125)/3.32 = 5%. These figures are very close to those in the data where e accounts for 81% and h for 19% of the variance in hours per capita. Regarding wages and labor productivity, our model generates

[31] These variables are logged (except for the interest rate) and HP-filtered.

[32] The fact that consumption is negatively correlated with output is a known feature of models with investment-specific technical change shocks (Justiniano and Primiceri, 2008; Guerrieri et al., 2010) and/or government shocks (King and Rebelo, 1999). In Appendix I, we show how we can get consumption right without distorting our labor market results: shutting down government shocks, adding adjustment costs, and/or removing wealth effects with the preferences in Greenwood et al. (1988) help make consumption procyclical.

Table 4: Business Cycle Behavior: NCES-CD Model vs. U.S.1948.I–2019.IV

	U.S. Data			NCES(SR)-CD(LR)		
	υ_x	$\rho(y,x)$	$\rho(x,x')$	υ_x	$\rho(y,x)$	$\rho(x,x')$
Output:						
y	3.31	1.00	0.87	3.28	1.00	0.72
Labor Market:						
eh	3.64	0.88	0.91	3.32	0.91	0.72
e	2.51	0.82	0.93	3.04	0.91	0.72
h	0.28	0.73	0.82	0.03	0.51	0.73
w	0.84	-0.13	0.75	0.63	-0.14	0.75
lp	0.81	0.15	0.73	0.58	0.19	0.70
ls	0.55	-0.34	0.75	0.52	-0.35	0.80
Elasticity:						
w/r	10.70	-0.75	0.82	9.32	-0.77	0.73
eh/k	3.39	0.90	0.91	3.50	0.88	0.72
σ	-	-	-	2.21	-0.48	0.90
Cons./Inv.:						
c	2.15	0.93	0.88	0.72	-0.26	0.73
i	32.29	0.95	0.86	18.24	0.92	0.72
R	0.01	0.80	0.82	0.00	0.84	0.72
Shocks:						
a^*	-	-	-	1.42	0.79	0.71
v^*	1.00	0.00	0.92	0.98	0.24	0.92
g^*	-	-	-	0.84	0.34	0.71
b^*	-	-	-	0.18	-0.12	0.75
TFP Residual:						
z^*	0.91	0.70	0.75	0.74	0.83	0.71

Notes: y denotes output per capita, c indicates consumption per capita, i indicates quality-adjusted investment per capita, R denotes the rate of return, $eh = H/N$ is hours per capita, e is employment per capita, h is average hours per worker, w is real wage, lp indicates labor productivity, ls is labor share, eh/k is the factor labor-capital input ratio, w/r is a factor price ratio. The data series of factor prices is constructed as $w = ls_t \frac{y_t}{eh_t}$ and $r = (1 - ls_t)\frac{y_t}{k_t}$. See Online Appendix A for the data definitions and variable construction. The statistic υ_x refers to the variance of the time series x, $\rho(x,y)$ refers to the correlation of x with output per capita, and $\rho(x,x')$ refers to the autocorrelation of x. For the computations of these statistics all time series have been logged (except the rate of return) and HP-filtered. The data moments of z^* and v^* are computed under the assumption of full utilization (see Appendix A). In Appendix B.1, we provide confidence intervals using block bootstrap with replacement on the time series data (Lee and Ingram, 1991; Bloom, 2009).

75% of the fluctuations in wages and 72% of the fluctuations in (average) labor productivity. Further, the model fully accounts for 95% of the fluctuations in labor share. The autocorrelation coefficients of w, lp, and ls are also very similar between the model and data (slightly below .80).

The comovements of the labor market variables are in Table 5. We find that all comovements between output, hours per capita, employment per capita, wages, labor productivity, and labor share are quantitatively consistent with the data. That is, our competitive-markets model with

Table 5: Labor Market Comovements: NCES-CD Model vs. U.S. 1948.I–2019.IV

	U.S. Data							NCES-CD Model						
	y	eh	e	h	w	lp	ls	y	eh	e	h	w	lp	ls
y	1.00	0.88	0.82	0.73	-0.13	0.15	-0.34	1.00	0.91	0.91	0.51	-0.14	0.19	-0.35
eh		1.00	0.97	0.70	-0.44	-0.33	-0.14		1.00	1.00	0.54	-0.37	-0.23	-0.17
e			1.00	0.51	-0.41	-0.39	-0.03			1.00	0.47	-0.38	-0.23	-0.17
h				1.00	-0.36	-0.02	-0.42				1.00	-0.12	-0.07	-0.06
w					1.00	0.67	0.42					1.00	0.57	0.50
lp						1.00	-0.39						1.00	-0.43
ls							1.00							1.00

Notes: See footnote of Table 4. In Appendix B.1, we provide confidence intervals using block bootstrap with replacement on the time series data (Lee and Ingram, 1991; Bloom, 2009).

countercyclical elasticity helps resolve the four labor market puzzles discussed in our introduction:

(a) The Dunlop-Tarshis phenomenon. We obtain a correlation between wages and output of $\rho(w, y) = -.14$, which is almost on target, $\rho(w, y) = -.13$.

(b) The labor productivity puzzle. The low correlation between labor productivity and output that we obtain, $\rho(lp, y) = .19$, is clearly close to its data counterpart, $\rho(lp, y) = .15$.

(c) The labor share puzzle. The countercyclical pattern of labor share almost matches the negative correlation coefficient of $\rho(ls, y) = -.35$. Further, we are getting the comovement of labor share and output for the right reasons as each of these four elements explaining $\rho(ls, y)$, which can be written as a function of $\rho(w, y)$, $\rho(lp, y)$, $\rho(w, lp)$ and $\frac{var(lp)}{var(w)}$, behave as their corresponding data counterparts. Note that the first two elements are the Dunlop-Tarshis phenomenon and the labor productivity puzzle that we just resolved. Likewise, $\rho(w, lp) = .57$ in our model is similar to that statistic in the data, .67. Our model $\frac{var(lp)}{var(w)} = .92$ is also close to that number in the data, .96 (Table 4). In other words, we are getting not only $\rho(ls, y)$ right, but also the cyclical joint behavior of w, lp, and y that explains $\rho(ls, y)$.

(d) The hours-productivity puzzle. The model comovement of hours per capita with productivity, measured either as wages or labor productivity, behaves as the data. Our model delivers $\rho(w, eh) = -.37$ and $\rho(lp, eh) = -.23$, and these are, respectively, $-.44$ and $-.33$ in the data.

Note that our model successfully matches a larger set of non-targeted (in estimation) comovements across labor market variables in Table 5, including (d). The behavior of hours per capita and labor share in the model also implies the low correlation between them, $\rho(ls, eh) = -.17$, that we observe in the data, -.14. Decomposing hours per capita in terms of employment per capita e and hours per worker h, we find that the comovement of e with w, lp and ls in the model $\rho(w, e) = -.38$, $\rho(lp, e) = -.23$ and $\rho(ls, e) = -.17$ are consistent with their data counterparts $\rho(w, e) = -.41$, $\rho(lp, e) = -.39$ and $\rho(ls, e) = .03$. While the comovement of hours per worker

h and labor productivity $\rho(lp, h)$ =-.07 is close to target -.02, we have some difficulty in entirely matching the comovement of wages and hours per worker with a model correlation of $\rho(w, h) =$ –.12, which is more negative in the data -.36 and, hence, in matching the comovement of labor share and hours per worker with a model correlation of $\rho(ls, h) = -.06$, which is –.42 in the data. Finally, the correlation of wages and labor share $\rho(w, ls) = .50$, and the correlation of labor productivity and labor share $\rho(lp, ls) = -.43$ are also consistent with their data counterparts, respectively $\rho(w, ls) = .42$ and $\rho(lp, ls) = -.39$.[33]

6 What Explains Our Results? The Role of σ_t

We show the effects of σ_t on business cycle moments in Section 6.1 and further explore the effects of σ_t on labor market dynamics in Section 6.2.

Table 6: The Effects of σ_t: A Non-orthogonal *Cumulative* Variance Decomposition

(a) NCES-CD Model

Activated Shocks:	Productivity Shocks $\{a,\sigma\}$			+ Investment Shocks $\{a,v,\sigma\}$			+ Government Shocks $\{a,v,g,\sigma\}$			+ Home Prod. Shocks $\{a,v,g,b,\sigma\}$		
Moments:	υ_x	$\rho(y,x)$	$\rho(x,x')$	υ_x	$\rho(y,x)$	$\rho(x,x')$	υ_x	$\rho(y,x)$	$\rho(x,x')$	υ_x	$\rho(y,x)$	$\rho(x,x')$
Output:												
y	2.21	1.00	0.72	2.80	1.00	0.73	3.23	1.00	0.72	3.28	1.00	0.72
Labor Market:												
eh	1.08	0.94	0.72	2.38	0.91	0.72	3.22	0.91	0.71	3.32	0.91	0.72
e	0.95	0.94	0.72	2.10	0.91	0.72	2.85	0.91	0.71	3.04	0.91	0.72
h	0.00	0.95	0.72	0.01	0.91	0.72	0.01	0.91	0.71	0.03	0.51	0.73
w	0.09	0.70	0.86	0.39	0.10	0.77	0.61	-0.12	0.75	0.63	-0.14	0.75
lp	0.37	0.83	0.74	0.50	0.38	0.71	0.57	0.21	0.70	0.58	0.19	0.70
ls	0.23	-0.63	0.72	0.45	-0.31	0.80	0.52	-0.34	0.80	0.52	-0.35	0.80

(b) CES-CD Model

Activated Shocks:	Productivity Shocks $\{a\}$			+ Investment Shocks $\{a,v\}$			+ Government Shocks $\{a,v,g\}$			+ Home Prod. Shocks $\{a,v,g,b\}$		
Moments:	υ_x	$\rho(y,x)$	$\rho(x,x')$	υ_x	$\rho(y,x)$	$\rho(x,x')$	υ_x	$\rho(y,x)$	$\rho(x,x')$	υ_x	$\rho(y,x)$	$\rho(x,x')$
Output:												
y	2.01	1.00	0.71	2.42	1.00	0.72	2.87	1.00	0.73	2.91	1.00	0.73
Labor Market:												
eh	0.71	0.99	0.71	1.50	0.92	0.73	2.37	0.91	0.73	2.46	0.91	0.73
e	0.63	0.99	0.71	1.32	0.92	0.73	2.09	0.91	0.73	2.26	0.90	0.73
h	0.00	0.99	0.71	0.01	0.92	0.73	0.01	0.91	0.73	0.02	0.47	0.74
w	0.07	0.66	0.89	0.25	0.01	0.77	0.46	-0.26	0.74	0.49	-0.28	0.74
lp	0.36	0.98	0.74	0.43	0.67	0.74	0.50	0.42	0.73	0.51	0.39	0.73
ls	0.17	-0.98	0.71	0.20	-0.96	0.71	0.24	-0.97	0.71	0.24	-0.97	0.71

Notes: See footnote of Table 4. In the CES-CD case we set $\sigma_t = \overline{\sigma} = .5981\ \forall t$.

[33] Further, the cyclical correlations of e and h with output in the model, respectively .91 and .51, are similar to their data counterparts, respectively .82 and .73. The correlations of e and h with eh are also very close between the model, respectively 1.00 and .54, and the data, respectively .97 and .70. Further, the model also accounts for the comovement between e and h with a correlation coefficient of .47 in the model and .51 in the data.

6.1 The effects of σ_t on business cycle moments

Before getting into the effects of σ_t, we note that the structurally estimated elasticity shows a significant countercyclical behavior with a correlation with output of -.48; see Table 4. Note that the countercyclical property of our structural estimation of σ_t coincides with the reduced-form estimation from our empirical strategy in Section 2.2. Indeed, applying our empirical estimation strategy from Section 2.2 on the simulated series of our fully-fledged model, we find that the structurally estimated elasticity is $\sigma_{NR} =$ 0.512 in non-recession quarters and $\sigma_R =$0.665 in recession quarters, which are significantly different from each other. Further, these model estimates are within the confidence intervals of the data counterparts in Section 2.2. We also note that σ_t displays substantial fluctuations with a (logged and HP-filtered) variance approximately two-thirds that of output.

6.1.1 A non-orthogonalized variance decomposition

Since σ_t does not generate fluctuations *per se* (see theoretical result 1 in Section 3) and affects the equilibrium allocations solely through the cross-derivatives (see theoretical result 2 in Section 3), a standard variance decomposition of the orthogonalized shocks does not deliver the total variance in the NCES-CD model. In fact, the sum of the fluctuations generated by the orthogonalized shocks in the NCES-CD model differ from the actual variance of the NCES-CD model—and are formally identical to the total variance of the CES-CD model that ignores the amplifying effect of σ_t. For this reason, in order to assess the actual effects of σ_t on business cycle moments, we follow a non-orthogonal strategy that we apply separately to the NCES-CD model and the CES-CD model in, respectively, panel (a) and panel (b) of Table 6. First, we explore the effects of productivity shocks a_t without and with σ_t, and then cumulatively add investment shocks v_t, government shocks g_t, and finally home-productivity shocks b_t.

Productivity shocks a_t. Productivity shocks are the most significant contributors to the cyclical variance of output, 2.21/3.28=67%, and hours per capita, 1.08/3.32=33%, in the NCES-CD model; see column $\{a, \sigma\}$ in panel (a) of Table 6. This is also the case in the CES-CD model but to a lesser extent with a_t shocks explaining 2.01/3.28=61% of the cyclical variance of output .71/3.32=21% of the variance of hours per capita; see column $\{a\}$ in panel (b) of Table 6. That is, the cyclical elasticity σ_t propagates the effects of productivity shocks boosting the contribution of these shocks to the cyclicality of hours per capita by 57%. Further, we find that a_t shocks are far from resolving the labor market puzzles in the CES-CD case but move the labor market in the right direction in the NCES-CD case. In particular, with a constant elasticity a_t shocks make labor productivity strongly comove with output, $\rho(lp, y) =$ .98, generating a too countercyclical labor share, $\rho(ls, y)$ = -.98. Instead, in the NCES-CD case, a_t shocks start to move the correlation of labor productivity and output away from 1, $\rho(lp, y)$ = .83, also generating a substantially lower countercyclicality of the labor share $\rho(ls, y)$ = -.63 that is closer to the data.

Investment shocks v_t. The cyclical σ_t also propagates the effects of v_t shocks. With a CES-CD technology, v_t shocks contribute to generate (2.42-2.01)/2.91=14% of output fluctuations and (1.50-0.71)/2.46=32% of the fluctuations in hours per capita; see column $\{a, v\}$ in panel (b) of Table 6. With NCES-CD the contribution to the fluctuations of output and hours per capita are, respectively, (2.80-2.21)/3.28=18% and (2.38-1.08)/3.32=39%; see column $\{a, v, \sigma\}$ in panel (a) of Table 6. Hence, whereas technological shocks $\{a, v\}$ account for 75% of output fluctuations and 53% of hours per capita fluctuations in the CES-CD model, these figures are, respectively, 85% and 72% in the NCES-CD model. Further, we find that v_t shocks substantially drop the correlation of wages and output in both models. Distinctively across models, v_t shocks reduce the correlation of labor productivity and output to .38 in the NCES-CD case. This reduces the countercyclicality of labor share to $\rho(ls, y)$ =–.31, moving the NCES-CD model closer to the data. In contrast, $\rho(ls, y)$ barely changes when v_t shocks are introduced in the CES-CD model.

Government shocks g_t. The effect of government shocks on model behavior is similar to that of investment shocks but less sizeable, particularly in the NCES-CD model. In the CES-CD model, government shocks contribute the fluctuations of output and hours per capita by, respectively, (2.87-2.42)/2.91=15% and (2.37-1.50)/2.46=35%; see column $\{a, v, g\}$ in panel (b) of Table 6. In the NCES-CD model these figures are, respectively, (3.23-2.80)/3.28=13% for output (3.22-2.38)/3.32=25% for hours per capita; see column $\{a, v, g, \sigma\}$ in panel (a) of Table 6. In terms of the labor market co-movements, government shocks help reduce the correlation of output and labor productivity from 0.67 to 0.42 in the CES-CD model and from 0.38 to 0.21 in the NCES-CD model. The rest of co-movents remains largely unaltered by government shocks.

Home-productivity shocks b_t. Home-productivity shocks do not alter the size of fluctuations or labor market behavior. This occurs in both the NCES-CD model (column $\{a, v, g, b, \sigma\}$ in panel (a) of Table 6) and the CES-CD model (column $\{a, v, g, b\}$ in panel (b) of Table 6). However, b_t shocks are key in getting the correlation between employment and hours per worker which is $\rho(e, h)$ = .38 in the CES-CD model and .47 in the NCES-CD model, closer to the data .51.

6.1.2 NCES-CD model vs CES-CD model.

The NCES-CD model outperforms the CES-CD model. To see this compare column $\{a, v, g, b, \sigma\}$ in panel (a) of Table 6 with column $\{a, v, g, b\}$ in panel (b) of Table 6. First, the CES-CD model falls short in generating the business cycle fluctuations of the NCES-CD model and, hence, the data. The variance of output and hours per capita in the CES-CD model is, respectively, 89% and 74% that of the NCES-CD model. Likewise, the variance of wages, labor productivity, and labor share in the CES-CD model is, respectively, 78%, 88% and 46% that of the NCES-model. Second, wages and labor productivity move too strongly with output in the CES-CD case. The correlation of wages with output goes from $\rho(w, y)$ =–.14 in the NCES-CD case to –.28 in the CES-CD, while it is –.13 in the data. At the same time, the correlation of labor productivity with output raises from $\rho(lp, y) =$ .19 in the NCES-CD case to .39 in the CES-CD, while it is .15 in the data. With the CES-CD technology, labor share is almost perfectly negatively correlated with output, $\rho(ls, y) =$ –.97, while this figure is -.35 in the NCES-CD model and –.34 in the data. That is, a constant elasticity of substitution revives the labor productivity and the labor share puzzles. Further, whereas the hours-productivity correlation is negative in our NCES-CD model, –.23, and in the data, –0.33, the two variables are uncorrelated in the CES-CD model.[34]

6.2 The effects of σ_t on labor market dynamics

Since cyclical σ_t introduces a non-linearity in the model, we plot the labor market responses to both a positive productivity shock (solid blue line) and a negative productivity shock (dashed cyan line). We plot the response to the negative productivity shock in absolute terms. We also plot the (symmetric) responses to a productivity shock in the CES-CD model (solid orange line).

A negative productivity shock a_t generates a long-lasting hump-shaped dynamic response of output (in absolute terms) in the NCES-CD model compared with CES-CD (panel (a), Figure 9). Since the size of productivity shocks a_t is the same across models, the differential response of output is due to the dynamic responses of factor inputs in the NCES-CD model compared with the CES-CD model. After the initial impact of a negative productivity shock that drops hours to -3.0%, the dynamic response of hours (in absolute terms) shows a hump shape that reaches a peak of -4.2% after three quarters (panel (b), Figure 9). In the same fashion, capital services keep dropping from -0.9% at impact to -1.7% after twelve quarters to slowly converge to the mean from below (panel (c), Figure 9). These long-lasting, dynamic responses of factor inputs in the NCES-CD model are due to the rise in the substitutability between factor inputs that detaches the behavior of hours and capital, as shown by the ratio of hours to capital that is more responsive in the NCES-CD model than in the CES-CD model (panel (d), Figure 9). The mechanism is straightforward. While capital services respond immediately to productivity shocks through capital utilization u_t (panel (e), Figure 9), the dynamics of capital services are yet slow.

[34] One may argue that we might be forcing it to perform below its possibilities without re-estimating the CES-CD model. To address this issue, we conduct a full re-estimation of the CES-CD model with exactly the same target moments (and weighting matrix) as those in the NCES-CD model; see our Appendix J.

Figure 9: The Effects of σ_t on Labor Market Dynamics: NCES-CD Model vs. CES-CD Model

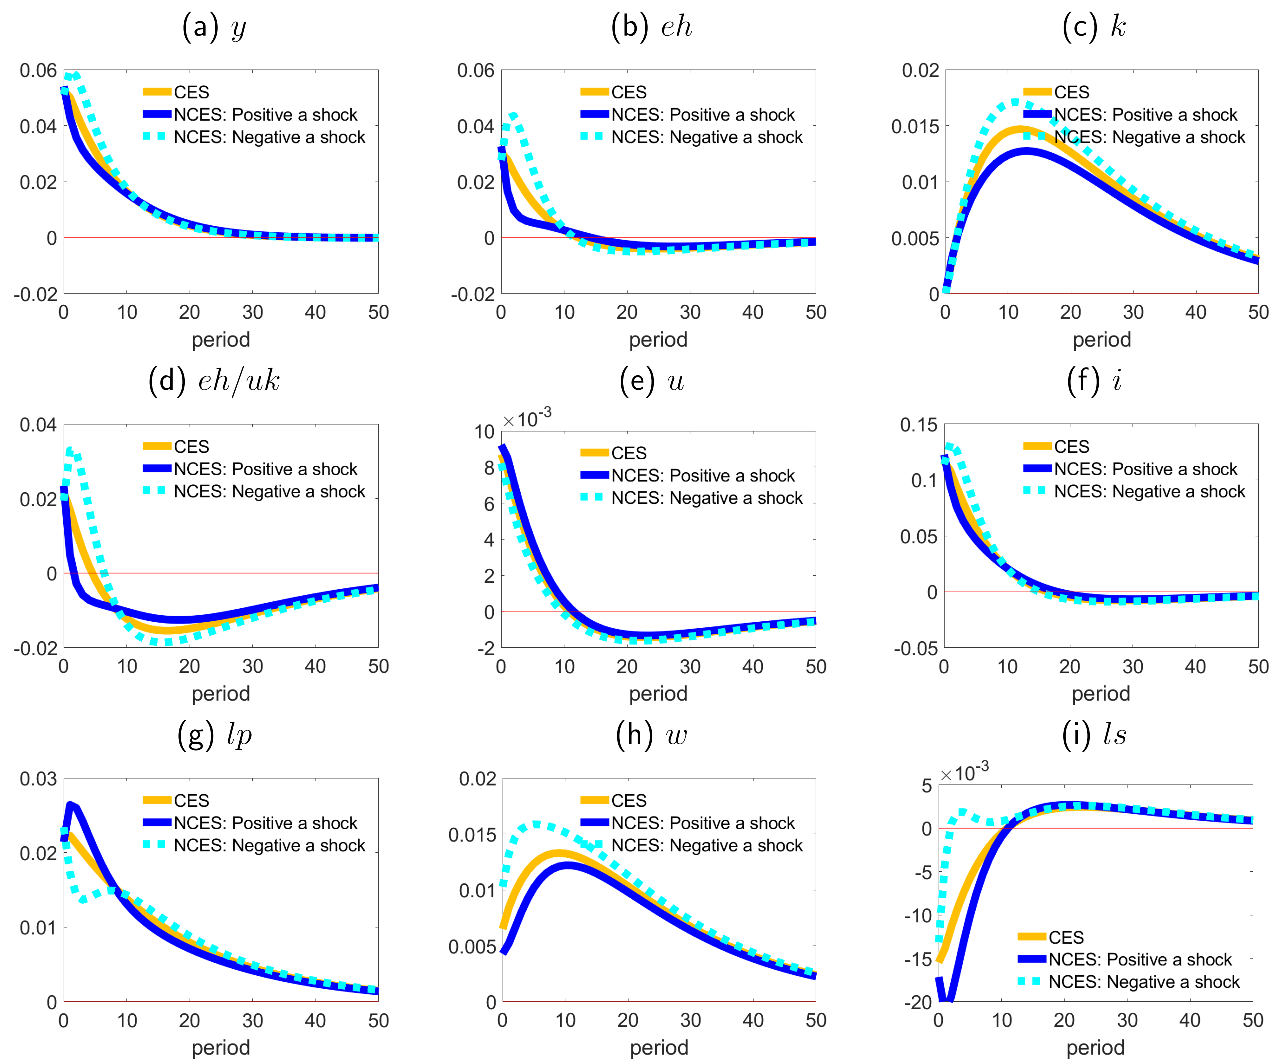

Notes: Simulated impulse responses of model variables for the first 50 periods after one-period productivity shock a_t. In each panel, the solid blue line shows the IRF to positive productivity shocks from the NCES-CD model; the blue dashed line shows the absolute value of the IRF to negative productivity shocks from the NCES-CD model; the yellow line shows the IRF to the positive productivity shock with CES. We plot simulated IRFs instead of orthogonalized IRFs because the nonlinear effect of σ_t cannot be orthogonalized from the productivity shocks.

It takes time to build new capital k_t since it is bounded by investment response in the previous period and the higher depreciation rate due to higher utilization. However, with the higher degree of substitutability between capital and labor that follows a negative productivity shock, hours per capita become less dependent on capital and respond more strongly to the negative productivity shock. Last, the stronger response of factor inputs encourages investment (panel (f), Figure 9) which helps explain the long-lasting, dynamic response of capital services in the NCES model. The opposite behavior occurs after a positive productivity shock which increases the degree of complementarity between capital and labor, causing the response of hours to inherit part of the less responsive behavior of physical capital. Hours per capita increase by 3.0% at prompt in both the CES and the NCES models but drop at a much faster rate in the NCES case reaching 1.0%

Figure 10: IRFs of Labor Share to TFP Shocks: NCES-CD Model vs. CES-CD Model

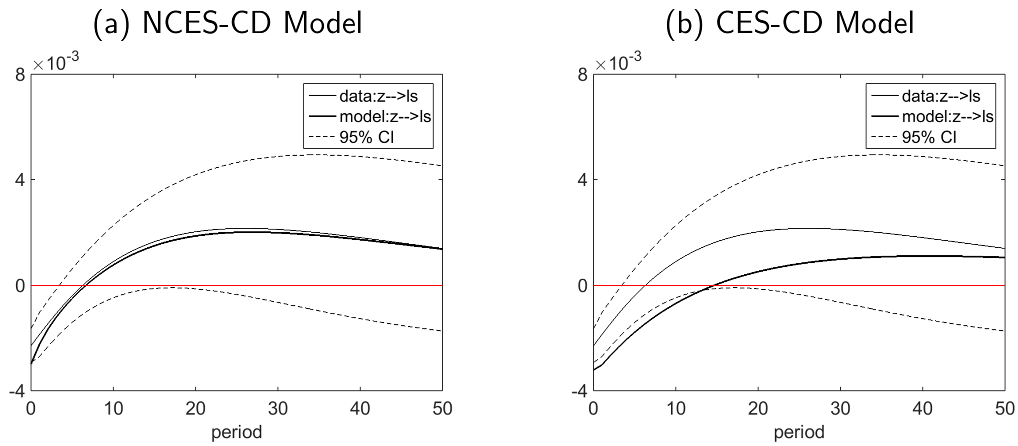

Notes: IRFs of labor share in logs in response to one standard deviation TFP shocks. To be consistent with the empirical evidence (Appendix B.1), we identify TFP shocks as model-generated Solow residuals z_t. Panel (a) refers to the IRFs from our NCES-CD model, while panel (b) refers to the CES-CD model.

after two quarters. Analogously, after a positive productivity shock, capital services increase by roughly 0.9% at prompt in both models, but the CES model exhibits long-lasting dynamics where capital services reach a peak of 1.5% in ten quarters, whereas the NCES model reaches a peak of 1.3% after four quarters.

Clearly, negative productivity shocks have larger effects on output and labor than productivity shocks with a countercyclical elasticity. Hence, given that positive and negative productivity shocks occur equally often in the model, the bigger effects of the negative productivity shocks dominate the cycle. This explains the larger aggregate fluctuations—particularly, the larger cyclical variance of output and hours—in the NCES model than in the CES model described in the previous section.

The countercyclicality of σ_t also has direct implication for labor market dynamics. First, after a negative productivity shock, the longer-lasting drop in hours after a recession is not accompanied by such a drop in output. This leads to a less pronounced decline in labor productivity after a recession in the NCES-CD model than the one in the CES model (panel (g), Figure 9). Note that this less pronounced decline in productivity coexists with a larger drop in output, which helps resolve the labor productivity puzzle by lowering the correlation between labor productivity and output in the NCES model. It also helps resolve the hours-productivity puzzle by lowering the correlation between labor productivity and hours per capita. Second, since the dynamic response of productivity shocks a_t to its own innovations is larger than the response of labor productivity, more substitutability between labor and capital in recessions—which shifts the wage response towards a_t and away from lp_t (see equation (15)), implies a longer-lasting response of wages in the NCES model than in the CES model (see panel (h), Figure 9). This implies that while the response of wages is smaller than that of output, the duration of the response of wages exceeds that of output. Specifically, after a negative productivity shock, wages keep lowering even when output is already reverting to the mean from below and drives the correlation of wages and output toward zero in the NCES model. This helps explain the Dunlop-Tarshis phenomenon.

Finally, combining the response of wages and labor productivity, we find a larger overshooting of labor share in the NCES case than in the CES case, a feature that is more salient in recessions (panel (i), Figure 9). These dynamics lower the correlation of labor share and output in the NCES model. Instead, in the CES model, labor share and output are almost perfectly negatively correlated. Last, to compare our results to those in Ríos-Rull and Santaeulàlia-Llopis (2010), we also show the role of σ_t in the response of labor share to model-generated TFP shocks in Figure 10. Again, σ_t plays a crucial role in generating the overshooting response of labor share.

7 Conclusion

We show that the elasticity of substitution between capital and hours, σ_t, is countercyclical. That is, it is easier to substitute capital and labor in recessions than in expansions. The countercyclicality of σ_t generates an asymmetry over the business cycle. Since recessions are accompanied by higher substitutability between capital and labor, labor is more responsive—less dependent on capital adjustments—than when the elasticity is assumed constant. The opposite occurs in expansions because the complementarity in capital and labor increases and, hence, the response of labor becomes more attached to the response of capital. We quantify that these effects are larger in recessions than in expansions and, hence, recession effects dominate the cycle.

The countercyclicality of the elasticity of substitution helps resolve several labor market puzzles at once. First, a negative productivity shock, which raises substitutability between capital and hours, implies a longer-lasting drop of hours that only partially transmits to output. The larger drop in hours than in output implies that labor productivity experiences a lower decline with countercyclical σ_t than with a constant elasticity. This helps explain the labor-productivity puzzle. Second, a countercyclical elasticity shifts the response of wages to that of productivity shocks and away from labor productivity which—making wages display a longer-lasting response to negative productivity shocks than that of output—lowers the correlation of wages and output. This helps explains the Dunlop-Tarshis phenomenon. Finally, the labor share inherits the behavior of wages and labor productivity. Under a countercyclical σ_t, we find that labor share is mildly countercyclical as in the data.

Our results suggest that a good theory of the countercyclicality of σ_t is potentially a good theory for the labor market. For future research, it seems quite natural to think about theories that endogenize σ_t from a technological perspective in a manner that the countercyclical property of σ_t is preserved. An alternative reading of our results is that our σ_t might reflect labor market (or other) frictions. That is, σ_t might encapsulate all sorts of frictions that we have decided to intentionally ignore in our analysis. In this context, our successful results in replicating cyclical labor market behavior suggest that one external validation exercise for noncompetitive models (e.g., Neo-Keynesian and Diamond-Mortensen-Pissarides settings) is the assessment of whether those models can replicate the countercyclicality of σ_t—for example, by applying our reduced-form estimation of cyclical σ_t to model-simulated data, as we have done with our model.

References

Abraham, K. G. and Haltiwanger, J. C. (1995). Real Wages and the Business Cycle. *Journal of Economic Literature*, 33(3):1215–1264.

Acemoglu, D. (2002). Directed Technical Change. *Review of Economic Studies*, 69(4):781–809.

Altig, D., Christiano, L. J., Eichenbaum, M., and Linde, J. (2011). Firm-Specific Capital, Nominal Rigidities and the Business Cycle. *Review of Economic Dynamics*, 14(2):225–247.

Antràs, P. (2004). Is the US Aggregate Production Function Cobb-Douglas? New Estimates of the Elasticity of Substitution. *Contributions to Macroeconomics*, 4(1).

Arrow, K. J., Chenery, H. B., Minhas, B. S., and Solow, R. M. (1961). Capital-Labor Substitution and Economic Efficiency. *Review of Economics and Statistics*, 43(3):225–250.

Aruoba, S. B., Bocola, L., and Schorfheide, F. (2017). Assessing DSGE Model Nonlinearities. *Journal of Economic Dynamics and Control*, 83:34–54.

Aruoba, S. B., Fernández-Villaverde, J., and Rubio-Ramírez, J. F. (2006). Comparing Solution Methods for Dynamic Equilibrium Economies. *Journal of Economic Dynamics and Control*, 30(12):2477–2508.

Basu, S. and Weil, D. N. (1997). Appropriate Technology and Growth. *The Quarterly Journal of Economics*, 113(4):1025–1054.

Benhabib, J., Rogerson, R., and Wright, R. (1991). Homework in Macroeconomics: Household Production and Aggregate Fluctuations. *Journal of Political Economy*, 99(6):1166–1187.

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica*, 77(3):623–685.

Boldrin, M. and Horvath, M. (1995). Labor Contracts and Business Cycles. *Journal of Political Economy*, 103(5):972–1004.

Brandolini, A. (1995). In Search of a Stylised Fact: Do Real Wages Exhibit a Consistent Pattern of Cyclical Variability? *Journal of Economic Surveys*, 9(2):103–63.

Caballero, R. J. (1994). Small Sample Bias and Adjustment Costs. *The Review of Economics and Statistics*, 76(1):52–58.

Caballero, R. J. and Hammour, M. L. (1998). Jobless Growth: Appropriability, Factor Substitution, and Unemployment. *Carnegie-Rochester Conference Series on Public Policy*, 48:51–94.

Cantore, C., Ferroni, F., and León-Ledesma, M. (2021). The Missing Link: Monetary Policy and the Labor Share. *Journal of the European Economic Association*, 19(3):1592–1620.

Chang, Y. and Hwang, S. (2015). Asymmetric Phase Shifts in U.S. Industrial Production Cycles. *The Review of Economics and Statistics*, 97(1):116–133.

Chang, Y. and Schorfheide, F. (2003). Labor-Supply Shifts and Economic Fluctuations. *Journal of Monetary Economics*, 50:1751–1768.

Chari, V. V., Kehoe, P. J., and McGrattan, E. (2007). Business Cycle Accounting. *Econometrica*, 75(3):781–836.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2012). Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities. *NBER Macroeconomics Annual*, 27(1):1–56.

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

Chirinko, R. S., Fazzari, S. M., and Meyer, A. P. (2011). A New Approach to Estimating Production Function Parameters: The Elusive Capital-Labor Substitution Elasticity. *Journal of Business and Economic Statistics*, 29(4):587–594.

Chirinko, R. S. and Mallick, D. (2017). The substitution elasticity, factor shares, and the low-frequency panel model. *American Economic Journal: Macroeconomics*, 9(4):225–253.

Cho, J.-O. and Cooley, T. F. (1994). Employment and Hours Over the Business Cycle. *Journal of Economic Dynamics and Control*, 18:411–432.

Choi, S. and Ríos-Rull, J. V. (2009). Understanding the Dynamics of Labor Share: The Role of Noncompetitive Factor Prices. *Annals of Economics and Statistics*, 95/96:251–278.

Choi, S. and Ríos-Rull, J.-V. (2020). Labor Share and Productivity Dynamics. *The Economic Journal*, 131:2856–2886.

Christiano, L. J. and Eichenbaum, M. (1992). Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations. *American Economic Review*, 82(3):430–450.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, 113(1):1–45.

Coeurdacier, N., Rey, H., and Winant, P. (2011). The Risky Steady-State. *American Economic Review: Papers & Proceedings*, 101(3):398–401.

Colciago, A. and Rossi, L. (2015). Firm Entry, Endogenous Markups and the Dynamics of the Labor Share of Income. *Macroeconomic Dynamics*, 19(6):1309–1331.

Cooley, T. F. and Prescott, E. C. (1995). Economic Growth and Business Cycles. In Cooley, T. F., editor, *Frontiers of Business Cycle Research*. pp. 1-38. Princeton University Press, Princeton, NJ.

Cummins, J. and Violante, G. L. (2002). Investment-Specific Technical Change in the US (1947-2000): Measurement and Macroeconomic Consequences. *Review of Economic Dynamics*, 5:243–284.

De Graeve, F., Dossche, M., Emiris, M., Sneessens, H., and Wouters, R. (2010). Risk Premiums and Macroeconomic Dynamics in a Heterogeneous Agent Model. *Journal of Economic Dynamics and Control*, 34(9):1680–1699.

de Groot, O. (2013). Computing the risky steady state of DSGE models. *Economics Letters*, 120(3):566–569.

de LaGrandville, O. (1989). In Quest of the Slutsky Diamond. *American Economic Review*, 79:468–1481.

Diamond, P. A., McFadden, D., and Rodriguez, M. (1978). Measurement of the Elasticity of Substitution and Bias of Technical Change. In *Production Economics*, pages 125–147. Amsterdam: North-Holland.

Fernández-Villaverde, J., Guerrón-Quintana, P., Rubio-Ramírez, J. F., and Uribe, M. (2011). Risk Matters: The Real Effects of Volatility Shocks. *American Economic Review*, 101(6):2530–2561.

Ferraro, D. (2018). The Asymmetric Cyclical Behavior of the U.S. Labor Market. *Review of Economic Dynamics*, 30:145–162.

Fisher, J. D. (2006). The Dynamic Effects of Neutral and Investment-Specific Technology Shocks. *Journal of Political Economy*, 114(3):413–451.

Francois, P. and Lloyd-Ellis, H. (2009). Schumpeterian Business Cycles with Pro-Cyclical R&D. *Review of Economic Dynamics*, 12(4):567–591.

Galí, J. and van Rens, T. (2021). The Vanishing Procyclicality of Labor Productivity. *The Economic Journal*, 131(633):302–326.

Gomme, P. and Greenwood, J. (1995). On the Cyclical Allocation of Risk. *Journal of Economic Dynamics and Control*, 19(1-2):91–124.

Greenwood, J., Hercowitz, Z., and Huffman, G. (1988). Investment, Capacity Utilization, and the Real Business Cycle. *American Economic Review*, 78(3):402–418.

Greenwood, J., Hercowitz, Z., and Krusell, P. (1997). Long-Run Implications of Investment-Specific Technological Change. *American Economic Review*, 83:342–362.

Greenwood, J., Hercowitz, Z., and Krusell, P. (2000). The Role of Investment-Specific Technological Change in the Business Cycle. *European Economic Review*, 44:91–115.

Greenwood, J., Rogerson, R., and Wright, R. (1995). Household Production in Real Business Cycle Theory. In Cooley, T. F., editor, *Frontiers of Business Cycle Research*. pp. 157–174. Princeton University Press, Princeton, NJ.

Guerrieri, L., Henderson, D., and Kim, J. (2010). Interpreting Investment-Specific Technology Shocks. International Finance Discussion Papers, Board of Governors of the Federal Reserve System, Number 1000.

Hairault, J.-O., Langot, F., and Osotimehin, S. (2010). Matching Frictions, Unemployment Dynamics and the Cost of Business Cycles. *Review of Economic Dynamics*, 13(4):759–779.

Hansen, G. D. and Prescott, E. C. (2005). Capacity Constraints, Asymmetries, and the Business Cycle. *Review of Economic Dynamics*, 8(4):850–865.

Hansen, G. D. and Wright, R. (1992). The Labor Market in Real Business Cycle Theory. *Quarterly Review, Federal Reserve Bank of Minneapolis*, pages 2–12.

Heathcote, J., Storesletten, K., and Violante, G. L. (2014). Consumption and Labor Supply with Partial Insurance: An Analytical Framework. *American Economic Review*, 104(7):2075–2126.

Hicks, J. (1932). *The Theory of Wages*. MacMillan, London.

Ilut, C., Kehrig, M., and Schneider, M. (2018). Slow to Hire , Quick to Fire : Employment Dynamics with Asymmetric Responses to News. *Journal of Political Economy*, 126(5):2011–2071.

Jones, C. I. (2003). Growth, Capital Shares, and a New Perspective on Production Functions. Stanford Working Paper.

Jones, C. I. (2005). The Shape of Production Functions and the Direction of Technical Change. *Quarterly Journal of Economics*, 120(2):517–549.

Jones, C. I. and Scrimgeour, D. (2008). A New Proof of Uzawaís Steady-State Growth Theorem. *Review of Economics and Statistics*, 90(1):180–2.

Justiniano, A. and Michelacci, C. (2012). The Cyclical Behavior of Equilibrium Unemployment and Vacancies in the US and Europe. In *NBER-ISOM Macro Annual Vol. 8*. National Bureau of Economic Research, Inc.

Justiniano, A. and Primiceri, G. E. (2008). The Time-Varying Volatility of Macroeconomic Fluctuations. *American Economic Review*, 98(3):604–641.

Justiniano, A., Primiceri, G. E., and Tambalotti, A. (2011). Investment Shocks and Business Cycles. *Review of Economic Dynamics*, 14(1):101–121.

Karabarbounis, L. (2014). The Labor Wedge: MRS vs. MPN. *Review of Economic Dynamics*, 17(2):206–223.

Karabarbounis, L. and Neiman, B. (2014). The Global Decline of the Labor Share. *Quarterly Journal of Economics*, 129(1):61–103.

Keane, M. and Rogerson, R. (2012). Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom. *Journal of Economic Literature*, 50(2):464–76.

King, R. G., Plosser, C. I., and Rebelo, S. (1988). Production, Growth and Business Cycles: I. The Basic Neoclassical Model. *Journal of Monetary Economics*, 21:195–232.

King, R. G. and Rebelo, S. (1999). Resuscitating Real Business Cycles. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*. chap. 14, pp. 927-1007. Elsevier Science Publications, Amsterdam.

Klump, R. and de LaGrandville, O. (2000). Economic Growth and the Elasticity of Substitution: Two Theorems and Some Suggestions. *American Economic Review*, 90:282–291.

Klump, R. and Papageorgiou, C. (2008). The CES Production Function in the Theory and Empirics of Economic Growth. *Journal of Macroeconomics*, 30(2):599–600.

Koh, D., Santaeulàlia-Llopis, R., and Zheng, Y. (2020). Labor Share Decline and Intellectual Property Products Capital. *Econometrica*, 88(6):2609–2628.

Lee, B. and Ingram, B. F. (1991). Simulation Estimation of Time-Series Models. *Journal of Econometrics*, 47:197–205.

León-Ledesma, M. A., McAdam, P., and Willman, A. (2010). Identifying the Elasticity of Substitution with Biased Technical Change. *American Econmic Review*, 100(4):1330–1367.

León-Ledesma, M. A. and Satchi, M. (2019). Appropriate technology and balanced growth. *Review of Economic Studies*, 86(2):807–835.

Lettau, M. and Ludvigson, S. C. (2013). Shocks and Crashes. In *NBER Macroeconomics Annual 2013, Volume 28*, NBER Chapters. National Bureau of Economic Research, Inc.

Lucas, R. E. (1970). Capacity, Overtime, and Empirical Production Functions. *The American Economic Review Papers and Proceedings*, 60(2):23–27.

Mangin, S. and Sedlacek, P. (2018). Unemployment and the Labor Share. *Journal of Monetary Economics*, 94:41–59.

McGrattan, E. (2020). Intangible Capital and Measured Productivity. *Review of Economic Dynamics*, 37(1):S147–S166.

McGrattan, E. R. and Prescott, E. C. (2012). The Labor Productivity Puzzle. In Ohanian, L., Taylor, J., and Wright, I., editors, *Government Policies and the Delayed Economic Recovery*. Hoover Institution Press, pp. 115-154.

McGrattan, E. R. and Prescott, E. C. (2014). A Reassessment of Real Business Cycle Theory. *American Economic Review*, 104(5):177–82.

Moscarini, G. and Postel-Vinay, F. (2008). The Timing of Labor Market Expansions: New Facts and New Hypothesis. In *NBER Macroeconomics Annual 2008*.

Ramey, V. A. (2012). The Impact of Hours Measures on the Trend and Cycle Behavior of U.S. Labor Productivity. UC San Diego Working Paper.

Ramey, V. A. (2016). Chapter 2 - Macroeconomic Shocks and Their Propagation. In *Handbook of Macroeconomics Vol.2*, pages 71–162. Elsevier.

Ríos-Rull, J. V. and Santaeulàlia-Llopis, R. (2010). Redistributive Shocks and Productivity Shocks. *Journal of Monetary Economics*, 57(8):931–948.

Ríos-Rull, J. V., Schorfheide, F., Fuentes-Albero, C., Kryshko, M., and Santaeulàlia-Llopis, R. (2012). Methods versus Substance: Measuring the Effects of Technology Shocks. *Journal of Monetary Economics*, 59(8):826–846.

Ruge-Murcia, F. (2012). Estimating Nonlinear DSGE Models by the Simulated Method of Moments: With an Application to Business Cycles. *Journal of Economic Dynamics and Control*, 36(6):914–938.

Sargent, T. J. and Wallace, N. (1974). The Elasticity of Substitution and Cyclical Behavior of Productivity, Wages, and Labor's Share. *The American Economic Review Papers and Proceedings*, 64(2):257–263.

Schlicht, E. (2006). A Variant of Uzawa's Theorem. *Economics Bulletin*, 5(6):1–5.

Shao, E. and Silos, P. (2014). Accounting for the Cyclical Dynamics of Income Shares. *Economic Inquiry*, 52:778–795.

Sichel, D. E. (1993). Business Cycle Asymmetry: A Deeper Look. *Economic Inquiry*, 31(2):224–236.

Uzawa, H. (1961). Neutral Inventions and the Stability of Growth Equilibrium. *Review of Economic Studies*, 28(2):117–124.

Yashiv, E. (2012). Frictions and the Joint Behavior of Hiring and Investment. IZA Discussion Papers 6636, Institute for the Study of Labor (IZA).

Young, A. T. (2013). U.S. Elasticities of Substitution and Factor Augmentation At the Industry Level. *Macroeconomic Dynamics*, 17(04):861–897.