Cyclical Returns to Scale and the Slopes of the Phillips Curves*

ChaeWon Baek† Byoungchan Lee‡

Tufts University HKUST

March 26, 2026

Abstract

This paper shows that more procyclical returns to scale in recent decades have flattened the Phillips curve when conventional activity measures, such as the output gap, labor gap, and labor share, are used as forcing variables. In contrast, the marginal cost Phillips curve remains steep. Using a simple, intuitive model with a translog production function, we illustrate a novel channel linking input complementarity and procyclical returns to scale to the identified slopes of the Phillips curves. By utilizing the estimated production functions and quantitative models, we emphasize the importance of our mechanism for rationalizing the US inflation data.

***JEL* Codes**: E30, E23, E52

Keywords: Phillips curve, returns to scale, translog production functions, input complementarity

*We are grateful to Yan Bai, Zhang Chen, David Cook, Russell Cooper, Giancarlo Corsetti, Marc Dordal i Carreras, Yunjong Eo, Yuriy Gorodnichenko, Paymon Khorrami, Moritz Kuhn, Jinil Kim, Jae Won Lee, Younghwan Lee, Francesco Lippi, Yang Lu, Eunseong Ma, Gernot Müller, Woong Yong Park, Benjamin Pugsley, Kwanho Shin, Vincent Sterk, Eric Swanson, Alireza Tahbaz-Salehi, Venky Venkateswaran, Jenny Xu, and the seminar and conference participants at Bank of Korea, Fridays at the Fed (Federal Reserve Bank of Boston), HKUST, HKUST-Jinan Joint Workshop on Macroeconomics, Midwest Macroeconomics Meeting at the Federal Reserve Bank of Kansas City, KIF-KAEA-KAFA Symposium, Korea University, and The 2025 World Congress of the Econometric Society for their helpful comments and suggestions. Musa Orak kindly shared the data for Ohanian et al. (2023) with us, which is greatly appreciated. All errors are our own.

†Chaewon.Baek@tufts.edu, ChaeWon Baek, Department of Economics, Tufts University, Joyce Cummings Center, Room 620G, 177 College Ave, Medford, 02155

‡bclee@ust.hk. Department of Economics, Hong Kong University of Science and Technology, 6068 Lee Shau Kee Business Building, Clear Water Bay, Kowloon, Hong Kong

1 Introduction

Stable inflation dynamics and substantial business cycle fluctuations in the first two decades of the 21st century sparked a reexamination of the relationship between inflation and economic activity, i.e., the Phillips curve. A series of recent studies have documented two divergent empirical results: the Phillips curves based on conventional measures of economic activity, such as the output gap, unemployment gap, and labor shares, have flattened (Stock and Watson, 2007, 2020; Del Negro et al., 2020; Inoue et al., 2024; Smith et al., 2025), and the version based on marginal costs is steep and alive (Gagliardone et al., 2025). We propose a theory that reconciles these two seemingly contradictory empirical patterns.[1] Furthermore, our quantitative analyses predict the growing significance of inflation expectations in understanding recent US inflation data and the Phillips curve, echoing the insights from Coibion and Gorodnichenko (2015), Hazell et al. (2022), and Meeks and Monti (2023).

The commonly used Cobb–Douglas (Galí and Gertler, 1999; Sbordone, 2002; Eichenbaum and Fisher, 2007) and constant elasticity of substitution (CES, Gagnon and Khan, 2005; McAdam and Willman, 2013) production functions in the study of Phillips curves feature constant returns to scale. Switching from those tightly parameterized functional forms to a more flexible translog production function (Christensen et al., 1973; De Loecker and Warzynski, 2012; Hyun et al., 2024) introduces cyclical variations in returns to scale around their long-term value.[2] This novel cyclical fluctuation in returns to scale introduces a wedge between conventional measures of economic activities and marginal costs.[3] Depending on the cyclicality of this wedge, the slopes of the Phillips curves using conventional activity measures as forcing variables could deviate from that of the marginal cost Phillips curve. We show that returns to scale have been more procyclical in recent decades, and this change

[1] To capture variations in marginal costs, Gagliardone et al. (2025) uses average variable costs with explicit consideration of potentially time-varying returns to scale. This cost structure is consistent with our theory, which builds on the microfounded cyclicality of returns to scale.

[2] For studies assuming translog functional forms for the demand side of the economy, see, e.g., Christensen et al. (1975); Bergin and Feenstra (2000); Feenstra (2003); Bilbiie et al. (2012); Lewis and Poilly (2012); Bilbiie et al. (2014); Fujiwara and Matsuyama (2023).

[3] The concept of time-varying returns to scale has recently been explored in relation to the decline in information technology prices (Lashkari et al., 2024), capital misallocation (**?**), and the complementarity between labor and energy (Hyun et al., 2024); these studies employed a nonhomothetic CES production function, an exogenously determined returns-to-scale process, and a translog production function, respectively. In this paper, we utilize a translog function, following Hyun et al. (2024), to investigate the implications of endogenously cyclical returns to scale that arise from the complementarity and substitutability among capital, skilled labor, and unskilled labor.

has resulted in the flattening of the conventional Phillips curves. In contrast, the marginal cost Phillips curve is still steep despite decades-long highly stable inflation before the recent pandemic, exemplified by missing disinflation and reinflation episodes (Coibion and Gorodnichenko, 2015; Ball and Mazumder, 2020).

We illustrate this mechanism using analytical expressions and graphical representations based on a simple model that incorporates a translog production function in an otherwise baseline New Keynesian model (see, e.g., Woodford, 2003; Galí, 2015). We show that the translog form, which is a second-order approximation of a generic production function, is general enough to nest the Cobb–Douglas and CES models within the log-linearization of the equilibrium conditions. Furthermore, relaxing the restrictive Cobb–Douglas and CES functional forms introduces cyclical fluctuations in returns to scale, consistent with the empirical evidence for time-varying returns to scale in Hyun et al. (2024) and **?**.[4] When factor inputs are complementary, the productivity of one input increases with the use of the other, and vice versa. Thus, when inputs are used more than usual, the overall marginal productivity can rise, leading to an increase in returns to scale.

This cyclical fluctuation in returns to scale introduces a novel element in the marginal costs of production in the translog model. In contrast, in the Cobb–Douglas and CES models, the marginal costs consist only of total factor productivity (TFP) and factor prices such as wages. Because conventional measures of economic activity, such as the output gap, labor gap, and labor shares, primarily reflect factor prices and TFP, returns to scale become an omitted variable in conventional Phillips curve regressions.[5]

The Phillips curves are often estimated using instruments for aggregate demand, such as identified monetary policy shocks, to address potential endogeneity due to mismeasurement and price markup shocks (see, e.g., Mavroeidis et al., 2014; Coibion et al., 2018; McLeay and Tenreyro, 2020; Barnichon and Mesters, 2020, 2021; Inoue et al., 2024).[6] However, even

[4] Hyun et al. (2024, section 2.3) documents an empirical pattern supporting the notion of procyclical returns to scale based on industry-level panel data. Using detailed firm-level data, **?**, appendix D shows that more than half of the variance in returns to scale across firms and over time is driven by an autoregressive component, leading to significantly persistent dynamics at business cycle frequencies.

[5] See, e.g., Galí and Gertler (1999); Sbordone (2002); Eichenbaum and Fisher (2007) for empirical studies employing labor shares as a forcing variable in the Phillips curve.

[6] Other intriguing approaches in the literature include introducing alternative measures of inflation and slack (Stock and Watson, 2010, 2020; Ball and Mazumder, 2011, 2019, 2020), using regional variations (Kiley, 2015; Hooper et al., 2020; McLeay and Tenreyro, 2020; Hazell et al., 2022; Fitzgerald et al., 2024), and relying on sectoral and firm-level data (Imbs et al., 2011; Gagliardone et al., 2025). In our quantitative analyses, we take a different route of structural model-based inferences following An and Schorfheide (2007); Schorfheide (2011); Del Negro et al. (2020). This full-information method is complementary to the limited

when using instruments, the estimated slopes of the conventional Phillips curves could still be affected by omitted variable biases if returns to scale are cyclical conditional on monetary policy shocks. In contrast, the marginal cost Philips curve is not subject to this confounding factor because marginal costs already account for variations in returns to scale. Therefore, if the cyclicality of returns to scale and the resulting magnitude of the omitted variable bias change over time, the conventional Phillips curves could flatten without a corresponding change in the marginal cost Phillips curve.[7]

We examine the relevance of this novel mechanism to US inflation dynamics through the lens of medium-scale dynamic stochastic general equilibrium (DSGE) models. We augment the partial equilibrium, four-input, nested CES production model studied in Krusell et al. (2000) and Ohanian et al. (2023) with a normalized translog production function proposed by Hyun et al. (2024). Firms use structure, equipment, skilled labor, and unskilled labor for production. We combine this production block of the economy with standard elements of quantitative New Keynesian general equilibrium models known to be useful in explaining US time series data, such as sticky prices and wages, fixed costs of production, investment adjustment costs, costly capacity utilization, and consumption habits (see, e.g., Christiano et al., 2005; Smets and Wouters, 2007).

We consider three models based on the Cobb–Douglas, nested CES, and translog production functions. We estimate these models separately for two different time periods using an early sample (1966-99) and a late sample (2000-19). To examine structural changes, we focus on production function parameters, the cyclicality of returns to scale, and the slopes of the Phillips curves with different forcing variables. When estimating the models, we employ additional data to the standard aggregate time series in light of the two types of capital (structures and equipment) and labor (skilled and unskilled) in the models. Specifically, we use equipment stock, relative labor hours and wage rates between skilled and unskilled workers.

Our quantitative results are summarized as follows. First, the translog model better matches the US macroeconomic data than the other models. Second, when estimating the unrestricted translog model, we reject the parameter restrictions imposed by the Cobb–

information approaches mentioned above.

[7]This rationalization of the flattening of the Phillips curves for conventional measures of (inverse) slack echoes the findings of Atkeson and Ohanian (2001) that the relationship between inflation and measures of slack may not be stable. This instability could be the case in our framework even when the slope of the marginal cost Phillips curve is stable, as shown in Section 2.

Douglas and nested CES models. Third, by comparing the estimated models using the pre- and post-2000 data, we find that the translog model can jointly replicate the steep marginal cost Phillips curve and the flattening of the conventional Phillips curves. Fourth, the mechanism connecting the cyclical returns to scale to marginal costs and inflation, embedded in the translog model, plays a crucial role in this result. In this analysis, we capture the potential effects of other structural changes on the slopes of the Phillips curves via variation in the Calvo price stickiness parameter. This reduced-form estimate indicates that the contribution of alternative mechanisms to the flattening of the Phillips curves is modest and statistically insignificant. In contrast, the Cobb–Douglas and nested CES models without cyclical returns to scale rely on unrealistically sticky prices and quite flat Phillips curves to match the stable inflation data after 2000. Furthermore, the translog model predicts the increased importance of inflation expectations in understanding inflation data and the Phillips curve. Finally, the historical decomposition of pandemic-era inflation data using the translog model suggests that high inflation in 2021 and 2022 was primarily driven by loose monetary policies and supply-side disturbances. This result emphasizes that monetary policies are still highly relevant to inflation despite the recent flattening of conventional Phillips curves. In contrast, the Cobb–Douglas and CES models do not predict this policy implication. Given their quite flat Phillips curves, these models primarily rely on price markup shocks and shifts in the Phillips curves to explain inflation fluctuations.

We also estimate production function parameters using industry-level panel data and instrumental variable (IV) methods as described by Hyun et al. (2024). Our estimates using the pre- and post-2000 data align with the structural estimation results based on the translog DSGE models, providing additional empirical evidence of stronger input complementarity at the aggregate level in recent decades than in earlier decades.

The remainder of this paper is organized as follows. Section 2 explains the mechanism behind the flattening of the Phillips curves when conventional measures of activity, such as the output gap, labor gap, and labor shares, are used as forcing variables. Using a simple tractable model with a translog production function, we illustrate the underlying intuition and introduce a novel channel linking input complementarity and procyclical returns to scale to the identified slopes of the Phillips curves. Building on this insight, Section 3 constructs a medium-scale DSGE model with a translog production function for quantitative analyses. Through the lens of the estimated model, we investigate the slopes of the Phillips curves; the relationships among the returns to scale, marginal costs, and inflation; and the inflation

dynamics under the model and during the pandemic period. Section 4 presents empirical evidence for stronger input complementarity in recent decades using industry-level panel data. Section 5 concludes the paper.

2 Illustrative Model

This section illustrates the mechanism connecting input complementarity and cyclical returns to scale with the slopes of the Phillips curves. This illustration is based on a flexible translog production function (Christensen et al., 1973; Hyun et al., 2024) incorporated into an otherwise baseline New Keynesian model (see, e.g., Woodford, 2003; Galí, 2015). This simple tractable model facilitates an intuitive exposition of the novel channel linking input complementarity to the slopes of the Phillips curves.

We describe the model below, focusing on new terms in equilibrium conditions arising from more general production functions than Cobb–Douglas functions. The derivations of standard model elements are kept brief.

Firms. Intermediate goods and their producers are indexed by $i \in [0, 1]$. Firms operate in a monopolistically competitive environment. The final goods are produced by combining intermediate goods in a Dixit and Stiglitz (1977) manner. Price-setting friction of the Calvo (1983) type leads to the following New Keynesian Phillips curve in log-linearization:

$$\hat{\pi}_t = \beta \mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa \hat{mc}_t + \varepsilon^p_t, \tag{2.1}$$

where π, mc, and ε^p represent inflation, real marginal costs, and an exogenous component in price markup, respectively. Hats are used to denote log deviations from the steady-state values of the corresponding variables. The slope of the marginal cost Phillips curve, κ, is given by $\frac{(1-\beta\zeta)(1-\zeta)}{\zeta}$, where β is the discount factor and $1-\zeta$ is the probability of resetting the price of an intermediate good in each period.

Next, we discuss the production function. We drop the firm and time indices, i and t, when they are unnecessary. Suppose that intermediate goods are produced using two types of labor, l_1 and l_2, such that $y = \exp(\varepsilon^a) f(l_1, l_2)$, where ε^a represents an exogenously determined level of TFP. A generic production function, f, increases with each input around the steady state, $(\bar{l}_1, \bar{l}_2)$, where a bar denotes the steady-state value of a variable. We introduce two types of labor because the concept of input complementarity is not natural

Table 1: Comparison of different production functions

	Cobb–Douglas	CES	Translog
Production function	$y = l_1^{\alpha_1} l_2^{\alpha_2}$	$\frac{y}{\bar{y}} = \left[\alpha_1 \left(\frac{l_1}{\bar{l}_1}\right)^{\phi} + \alpha_2 \left(\frac{l_2}{\bar{l}_2}\right)^{\phi}\right]^{1/\phi}$	$y = f(l_1, l_2)$
Log deviation form	$\hat{y} = \alpha_1 \hat{l}_1 + \alpha_2 \hat{l}_2$	$\hat{y} = \alpha_1 \hat{l}_1 + \alpha_2 \hat{l}_2 + 0.5\alpha_1\alpha_2\phi\left[(\hat{l}_1)^2 - 2\hat{l}_1\hat{l}_2 + (\hat{l}_2)^2\right]$	$\hat{y} = \alpha_1 \hat{l}_1 + \alpha_2 \hat{l}_2 + 0.5\beta_{11}(\hat{l}_1)^2 + \beta_{12}\hat{l}_1\hat{l}_2 + 0.5\beta_{22}(\hat{l}_2)^2$
Elast. of output w.r.t. l_j	$\frac{\partial \log y}{\partial \log l_j} = \alpha_j$	$\frac{\partial \log y}{\partial \log l_j} = \alpha_j + \alpha_1\alpha_2\phi(\hat{l}_j - \hat{l}_k),\ j \neq k$	$\frac{\partial \log y}{\partial \log l_j} = \alpha_j + \beta_{jj}\hat{l}_j + \beta_{jk}\hat{l}_k,\ j \neq k$
Returns to scale	$rts = 1$	$rts = 1$	$rts = 1 + (\beta_{11} + \beta_{12})\hat{l}_1 + (\beta_{12} + \beta_{22})\hat{l}_2$

Notes: We ignore the TFP term, ε^a, in this table. The returns to scale are defined as the scale elasticity, i.e., $\frac{\partial \log f(\lambda l_1, \lambda l_2)}{\partial \log \lambda}|_{\lambda=1}$. The CES production function is written in the deviation form following Cantore and Levine (2012). For the CES and translog functions, we ignore the higher-order terms that degenerate when the equilibrium conditions are log-linearized.

when a production function has only a single input. By focusing on a symmetric equilibrium, where $l_1 = l_2$, we write the equilibrium conditions in terms of the aggregate labor, $l = l_1 + l_2$, similar to the baseline New Keynesian models. Note that this modeling assumption is only for illustrative purposes. We examine the input complementarity and substitutability among equipment, skilled labor, and unskilled labor in Section 3 for the quantitative analyses.

We consider a second-order Taylor expansion of f in the logarithm, ignoring the higher-order terms that degenerate when the equilibrium conditions are log-linearized. Suppose that $\alpha_j = \frac{\partial \log y}{\partial \log l_j} = 0.5$ in the steady state for $j = 1, 2$. Similarly, the second-order derivatives are denoted by $\beta_{jk} = \frac{\partial^2 \log y}{\partial \log l_j \partial \log l_k}$ in the steady state for $j, k = 1, 2$. Thus:

$$\hat{y} = \varepsilon^a + \alpha_1 \hat{l}_1 + \alpha_2 \hat{l}_2 + 0.5\beta_{11}(\hat{l}_1)^2 + \beta_{12}\hat{l}_1\hat{l}_2 + 0.5\beta_{22}(\hat{l}_2)^2. \tag{2.2}$$

Clearly, at the first order, f is proportional to a conventional Cobb–Douglas production function, $l_1^{\alpha_1} l_2^{\alpha_2}$, or l because of the symmetry ($l_1 = l_2$). However, the novelty in our theory arises from the second-order properties of the production function, captured by the translog parameters, β_{jk}.

To illustrate the workings of the translog parameters, we compare Equation (2.2) with the Cobb–Douglas and CES production functions. Table 1 shows the key differences among the three functions at the first and second orders, while omitting the TFP and irrelevant higher-order terms. The Cobb–Douglas function is log-linear. Thus, the elasticity of output with respect to each input, α_j, and the returns to scale, $\alpha_1 + \alpha_2$, are constant. For the CES

function, one more parameter, $\phi \leq 1$, governs the coefficients on the second-order terms. At the second order, we can convert the CES function to a translog form with parameter restrictions such that $\beta_{11} = \beta_{22} = \alpha_1\alpha_2\phi$ and $\beta_{12} = -\beta_{11}$. When $\phi < 0$, the two inputs are complementary in the sense that $\beta_{12} > 0$. In this case, the elasticity of the output with respect to l_1 increases in l_2, and vice versa. However, $\frac{\partial \log y}{\partial \log l_1}$ decreases in l_1, balancing the productivity gain from the input complementarity and l_2 and making the returns to scale constant. The translog specification relaxes the restrictions on β_{jk}s. Note that without the restriction that $\beta_{11} = -\beta_{12}$, the input complementarity ($\beta_{12} > 0$) does not necessarily imply that $\frac{\partial \log y}{\partial \log l_1}$ decreases in l_1.

One of the most intriguing implications of this property is that the returns to scale can vary over time under the translog production function. When defined as the scale elasticity, $\frac{\partial \log f(\lambda l_1, \lambda l_2)}{\partial \log \lambda}|_{\lambda=1}$, the returns to scale are given by the sum of the elasticity of output:

$$\begin{aligned} rts_t &= 1 + (\beta_{11} + \beta_{12})\hat{l}_{1t} + (\beta_{12} + \beta_{22})\hat{l}_{2t} \\ &= 1 + (\beta_{11} + 2\beta_{12} + \beta_{22})\hat{l}_t, \end{aligned} \tag{2.3}$$

where the second line uses the symmetry between the two types of labor. For example, suppose that $\beta_{11} = \beta_{22} = 0$, as in the Cobb–Douglas case, the two types of labor are complementary ($\beta_{12} > 0$), and the aggregate labor is procyclical. In this situation, the returns to scale are also procyclical, fluctuating around the long-term value of one (see also Hyun et al., 2024, for empirical evidence for procyclical returns to scale). Thus, the production function (2.2) features constant returns to scale in the long run but allows for short-run variations. In contrast, for the Cobb–Douglas and CES functions, rts_t simplifies to one because the parameter restrictions imply that $\beta_{11} + \beta_{12} = \beta_{12} + \beta_{22} = 0$ in both cases.

To examine the effects of time-varying returns to scale on marginal costs, we consider the cost minimization problem of an intermediate goods producer:

$$\min_{l_1, l_2} w_1 l_1 + w_2 l_2 - mc\left[\exp(\varepsilon^a) f(l_1, l_2) - y\right],$$

where w_j is the real wage rate for type j labor for $j = 1, 2$. The first-order conditions (FOCs) imply that, in log-linearization, the aggregate marginal costs equal:

$$\hat{mc}_t = \underbrace{\hat{w}_t}_{\text{factor prices}} - \underbrace{\varepsilon^a_t}_{\text{tfp}} - \underbrace{\widehat{rts_t}}_{\text{returns to scale}}, \tag{2.4}$$

where w_t is the aggregate wage rate and $\widehat{rts}_t = (\beta_{11} + 2\beta_{12} + \beta_{22})\hat{l}_t$.[8] The first two terms in Equation (2.4) reflect the average variable costs, $(w_1 l_1 + w_2 l_2)/y$, based on the factor prices and TFP. These terms correspond to the marginal costs in the baseline New Keynesian models with the Cobb–Douglas or CES production functions that assume constant returns to scale. In contrast, the translog specification generates a novel element–the cyclical returns to scale–in the marginal cost. Again, suppose that aggregate labor is procyclical, $\beta_{11} = \beta_{12} = 0$, and $\beta_{12} > 0$. In this case, the marginal productivity gain from the input complementarity, along with a larger-than-usual labor input ($\hat{l} > 0$), decreases the marginal costs in expansions. This mechanism is captured by the procyclical variation in the returns to scale in Equation (2.4).[9]

Households. To further simplify the expression for the marginal costs in Equation (2.4), we specify the household side of the economy. The period utility function of the type j worker is given by $\log c_j - \frac{h_j^{1+1/\sigma_l}}{1+1/\sigma_l}$, where c_j and h_j are the consumption and hours of a type j worker, respectively. The aggregate type j labor input, l_j, equals $0.5h_j$ because the population share of each type is 0.5. σ_l is the Frisch elasticity of the labor supply. Suppose that the labor market is competitive. Using the symmetry, the market clearing condition ($y = c$), and the labor supply schedule ($\text{-}\hat{c}_j + \hat{w}_j = \frac{1}{\sigma_l}\hat{h}_j$ for $j = 1, 2$), it follows that $\hat{w} - \varepsilon^a = (1 + \frac{1}{\sigma_l})\hat{l}$. Thus, Equation (2.4) implies the following:

$$\hat{mc}_t = \left(1 + \frac{1}{\sigma_l}\right)\hat{l}_t - \widehat{rts}_t = \left[1 + \frac{1}{\sigma_l} - (\beta_{11} + 2\beta_{12} + \beta_{22})\right]\hat{l}_t. \tag{2.5}$$

Phillips curves. From the marginal cost Phillips curve (2.1) and the expression for the marginal costs (2.5), we obtain the following labor Phillips curve:

$$\begin{aligned}\hat{\pi}_t &= \beta\mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa\hat{mc}_t + \varepsilon_t^p \\ &= \beta\mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa\left[1 + \frac{1}{\sigma_l} - (\beta_{11} + 2\beta_{12} + \beta_{22})\right]\hat{l}_t + \varepsilon_t^p.\end{aligned}$$

Suppose that $\beta_{11} = \beta_{22} = 0$, as in the Cobb–Douglas model. Then, according to the above

[8] The FOCs are given by $w_j = mc\exp(\varepsilon^a)f_j = mc\exp(\varepsilon^a)\frac{\partial \log f}{\partial \log l_j}\frac{f}{l_j} = mc\frac{y}{l_j}\frac{\partial \log f}{\partial \log l_j}$ for $j = 1, 2$. Because of the symmetry, the aggregate wage rate, w, equals $w_1 = w_2$ and, similarly, $\hat{l} = \hat{l}_1 = \hat{l}_2$. Then, Equation (2.4) follows from the fact that $\hat{y} = \varepsilon^a + \hat{l}$ in log-linearization.

[9] Gagliardone et al. (2025, equation (14)) similarly assumes that the marginal costs depend on the average variable costs and the returns to scale.

equation, stronger input complementarity ($\beta_{12} > 0$) and the resulting procyclical returns to scale yield flatter labor Phillips curves. The following proposition and remark formalize this observation.

Proposition 1 (Phillips curves with cyclical returns to scale)**.** *Assume the model structure discussed in this section. For different measures of economic activity, such as the marginal cost, output gap, labor gap, and labor shares, the following hold:*

$$\begin{aligned}
\textit{marginal cost Phillips curve:} \quad & \hat{\pi}_t = \beta \mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa \hat{mc}_t + \varepsilon_t^p, \\
\textit{output gap Phillips curve:} \quad & \hat{\pi}_t = \beta \mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa(1 + 1/\sigma_l)\textit{output gap}_t + \varepsilon_t^p - \kappa \widehat{rts}_t, \\
\textit{labor gap Phillips curve:} \quad & \hat{\pi}_t = \beta \mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa(1 + 1/\sigma_l)\textit{labor gap}_t + \varepsilon_t^p - \kappa \widehat{rts}_t, \\
\textit{labor share Phillips curve:} \quad & \hat{\pi}_t = \beta \mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa \widehat{\textit{labor share}}_t + \varepsilon_t^p - \kappa \widehat{rts}_t.
\end{aligned}$$

Proof. We use $*$ to denote variables under flexible price equilibrium. Note that $\hat{mc}_t^* = 0$. Thus, Equation (2.5) implies that $\hat{l}_t^* = 0$. Furthermore, output $\text{gap}_t = \hat{y}_t - \hat{y}_t^* = (\varepsilon_t^a + \hat{l}_t) - (\varepsilon_t^a + \hat{l}_t^*) = \hat{l}_t$. Similarly, labor $\text{gap}_t = \hat{l}_t - \hat{l}_t^* = \hat{l}_t$. Then, Equation (2.5) and the marginal cost Phillips curve yield the output gap and labor gap Phillips curves above. Finally, for the labor share Phillips curve, note that $\widehat{\text{labor share}}_t = \hat{w}_t + \hat{l}_t - \hat{y}_t = \hat{w}_t - \varepsilon_t^a = \hat{mc}_t + \widehat{rts}_t$ because of Equation (2.4). □

Remark 1. In practice, when estimating a Phillips curve, IVs that are orthogonal to ε_t^p can be used to address potential endogeneity arising from supply-side disturbances (see Mavroeidis et al., 2014; Coibion et al., 2018; Barnichon and Mesters, 2020, 2021; McLeay and Tenreyro, 2020; Inoue et al., 2024). However, even in this case, if conventional measures of economic activity other than the marginal costs are used as forcing variables, the returns to scale term might behave as an omitted variable in regression equations. If the activity measure is procyclical conditional on the variations captured by the instruments, the procyclical returns to scale could induce a downward omitted variable bias in the estimated slope of the Phillips curves. Finally, if the returns to scale have become more procyclical in recent decades, the omitted variable bias in the identified slope would be more negative, leading to the flattening of the Phillips curves when the conventional activity measures are employed. Note that this prediction holds even when the "true" slope of the marginal cost Phillips curve, κ, remains large (see Gagliardone et al., 2025). Furthermore, this prediction is supported by data; we show stronger input complementarity and more procyclical returns

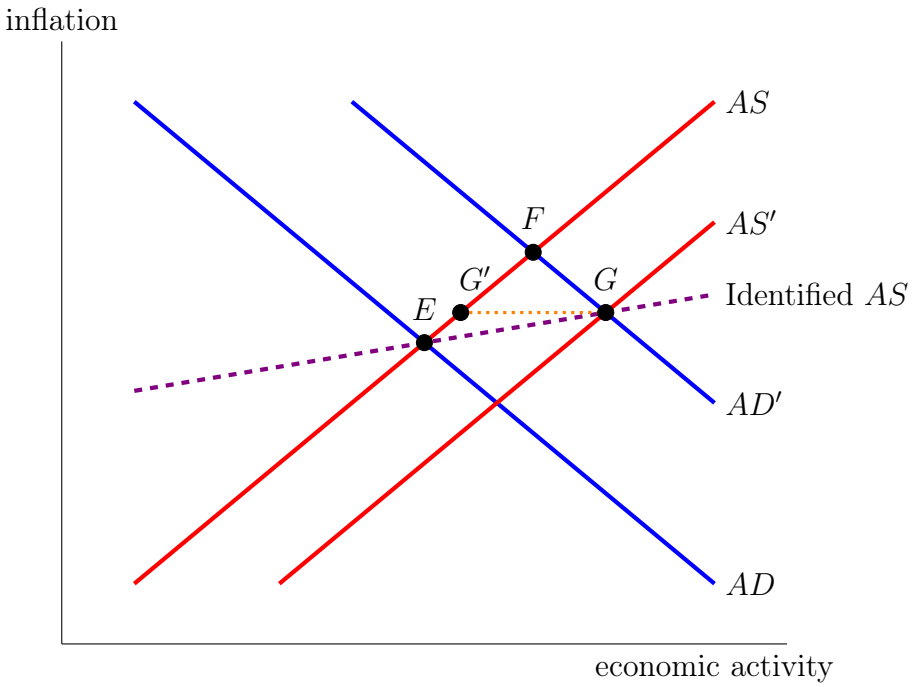


Figure 1: Graphical illustration of the Phillips curve flattening

to scale in recent decades than in earlier decades in Sections 3 and 4.

The downward omitted variable bias in the identified slope coincides with a decrease in the elasticity of marginal costs with respect to the conventional activity measures, such as the output gap, $\frac{\partial \hat{mc}}{\partial \text{output gap}}$. In the translog model, the procyclical returns to scale, represented by a positive value of $\beta_{11} + 2\beta_{12} + \beta_{22}$, decrease the abovementioned elasticity: $\frac{\partial \hat{mc}}{\partial \text{output gap}} = 1 + \frac{1}{\sigma_l} - (\beta_{11} + 2\beta_{12} + \beta_{22})$. Note that the pass-through of the output gap into inflation, $\frac{\partial \hat{\pi}}{\partial \text{output gap}}$, is determined by the slope of the marginal cost Phillips curve, $\frac{\partial \hat{\pi}}{\partial \hat{mc}} = \kappa$, and the elasticity of marginal costs with respect to the output gap, $\frac{\partial \hat{mc}}{\partial \text{output gap}}$. Thus, a small $\frac{\partial \hat{mc}}{\partial \text{output gap}}$ can result in a weak pass-through of the output gap into inflation even when the marginal cost Phillips curve is steep. This theoretical prediction provides an explicit microfoundation for the reduced-form, empirical results that $\frac{\partial \hat{mc}}{\partial \text{output gap}}$ is less than one in recent periods in Gagliardone et al. (2025).

Figure 1 graphically illustrates this mechanism. Assume that $\sigma_l = \infty$ for exposition so that the Phillips curves in Proposition 1 share the same slope, κ. The Phillips curve, conditional on the expected inflation and cost-push shocks, is upward sloping on the economic activity and inflation planes. Thus, the Phillips curve serves as the aggregate supply curve,

labeled AS in Figure 1. Note that estimating the slope of this curve is analogous to identifying the slope of the supply curve in the supply–demand simultaneous equation system. As emphasized by McLeay and Tenreyro (2020), the data consist of the equilibrium pairs of inflation and economic activity levels, such as point E. To properly identify κ, we need to use demand shifters, such as the monetary policy shocks employed in Barnichon and Mesters (2020, 2021), as instruments. The identification assumption is that the AD curve shifts to AD', whereas the AS curve remains the same; therefore, comparing the two equilibrium points (E and F) allows econometricians to estimate κ. However, Proposition 1 implies that the AS curve could shift downward to AS' in response to expansionary monetary policy shocks when conventional measures of economic activity, such as the output gap, labor gap, and labor share, are employed and the returns to scale are procyclical conditional on monetary policy shocks. The increase in $\widehat{rts}_t$ works as an endogenous supply shifter, driving the economy to point G rather than F. Thus, the identified AS curve, denoted by the purple dashed line connecting points E and G in Figure 1, could be flatter than the true underlying AS curve.

Note that in the $(\hat{mc}, \hat{\pi})$ pair, the equilibrium moves from point E to G' because $\widehat{rts}$ does not appear in the marginal cost Phillips curve. However, for a conventional forcing variable, e.g., the labor gap, the new equilibrium is represented by point G. That is, the labor gap increases more than the marginal costs do, and the difference between the two, the distance between points G' and G, reflects the increase in returns to scale (see also the first equality in Equation (2.5)). If this mechanism has become more effective in recent decades, the AS curve would shift further downward. Then, econometricians can observe the flattening of the conventional Phillips curves even when the underlying slope, κ, remains largely unchanged.

Remark 2. A specific supply-side effect of demand disturbances propagating through the cyclical returns to scale is central to the above mechanism for the flattening of the conventional Phillips curves without an accompanying decrease in the slope of the marginal cost Phillips curve. Relatedly, aggregate demand might affect the supply block of the economy through other channels, such as endogenous R&D (Comin and Gertler, 2006; Anzoategui et al., 2019; Bianchi et al., 2019; Okada, 2022; Ma and Zimmermann, 2023; Antolin-Diaz and Surico, 2024; Cloyne et al., 2024), misallocation (Basu and Fernald, 2002; Eisfeldt and Rampini, 2006; Hsieh and Klenow, 2009; Beaudry and Portier, 2014; Alam, 2020; Baqaee et al., 2024), firm dynamics (Smirnyagin, 2023; Fujiwara and Matsuyama, 2023), search effort (Qiu and Ríos-Rull, 2022; Bai et al., 2024b), increasing returns to scale (Hall, 1990; Basu and

Fernald, 1997), and capacity utilization (Basu, 1996; Boehm and Pandalai-Nayar, 2022). If the effects of these mechanisms have recently become more procyclical, the identified Phillips curves could be flatter than previous curves for similar reasons.

A key distinction between our mechanism and other prominent supply-side channels of monetary policy is that the cyclical returns to scale do not affect TFP at the first order; the TFP in our model (ε_t^a) is exogenous and equal to the Solow residual at the first order. In contrast, several other supply-side mechanisms described above could directly affect measured aggregate TFP (see also Ma and Zimmermann, 2023; Jordà et al., 2023).

Given this consideration, we empirically investigate how TFP responds to monetary policy shocks in the pre- and post-2000 samples (see Appendix C.5 for details). We document that the level of TFP responded procyclically to monetary policy shocks in the early sample. However, in the later sample, the impulse responses are close to zero and statistically insignificant. Thus, we find no significant evidence of increasingly procyclical TFP responses to monetary policy shocks in recent periods. To be clear, we do not argue that those other supply-side mechanisms are weak. Rather, our findings suggest they have not strengthened enough in the past two to three decades to drive the observed flattening of the Phillips curves. Note that our empirical results on TFP do not contradict our proposed explanation on the basis of the procyclical returns to scale because the returns to scale can endogenously vary over time without affecting ε_t^a in our framework.

Remark 3. As an alternative to Equation (2.2), we can assume a normalized translog form proposed by Hyun et al. (2024):

$$y_{it} = \underbrace{\exp(\varepsilon_t^a) l_{1,it}^{\alpha_1} l_{2,it}^{\alpha_2}}_{\text{Cobb–Douglas}} \times \underbrace{\exp\left(s.o.t._{it}\right)}_{\text{Translog terms}},$$

$$\text{where} \quad s.o.t._{it} = \beta_{11}\hat{l}_{1,it}\hat{l}_{1t} + \beta_{12}\hat{l}_{1,it}\hat{l}_{2t} + \beta_{21}\hat{l}_{2,it}\hat{l}_{1t} + \beta_{22}\hat{l}_{2,it}\hat{l}_{2t}, \tag{2.6}$$

and $\beta_{12} = \beta_{21}$. This production function for firm i has the following properties: First, the second-order term, $s.o.t._{,it}$, degenerates in the steady state because the log deviation terms are zero, $\hat{l}_{j,it} = \hat{l}_{jt} = 0$ for $j = 1, 2$. Thus, the steady state of the model is identical to that of the Cobb–Douglas model and, in particular, does not depend on the translog parameters, $\{\beta_{jk} : j, k = 1, 2\}$. This property allows us to focus on the business cycle implications of input complementarity and cyclical returns to scale without altering the long-term predictions of the model. Second, the production functions (2.2) and (2.6) yield

the same aggregate equilibrium conditions in log-linearization. Furthermore, the production function in a deviation form from the steady state is given by:

$$\frac{y_{it}}{\bar{y}} = \exp(\varepsilon_t^a) \left(\frac{l_{1,it}}{\bar{l}_1}\right)^{\alpha_{1t}} \left(\frac{l_{2,it}}{\bar{l}_2}\right)^{\alpha_{2t}},$$

where $\alpha_{1t} = \alpha_1 + \beta_{11}\hat{l}_{1t} + \beta_{12}\hat{l}_{2t}$, $\alpha_{2t} = \alpha_2 + \beta_{21}\hat{l}_{1t} + \beta_{22}\hat{l}_{2t}$, and a bar denotes the steady-state values. Thus, the translog structure induces a potentially time-varying elasticity of output with respect to each input, α_{1t} and α_{2t}, leading to time-varying returns to scale, $rts_t = \alpha_{1t} + \alpha_{2t}$. Finally, the production function (2.6) does not allow each firm to choose its returns to scale, reflected by the fact that α_{1t} and α_{2t} are exogenous from firm i's perspective. This property simplifies the second-order conditions for firm i's cost minimization problem, similar to the Cobb–Douglas model (see Appendix B.1.2).[10]

Remark 4. Other models with nonconstant returns to scale may also exhibit procyclical returns to scale, potentially explaining the flattening of the conventional Phillips curves. However, in these models, procyclical returns to scale are typically accompanied by corresponding procyclical fluctuations in the measured TFP. In contrast, the translog framework described above allows for procyclical returns to scale without affecting the cyclicality of the measured TFP. Because the TFP has not responded to monetary policy shocks more procyclically in recent periods (Remark 2), our proposed explanation emphasizes the cyclical returns to scale that are microfounded by the translog model.

To elaborate on the above point, consider the following production function:

$$y_{it} = \exp(\varepsilon_t^a) l_{it}^{\alpha_{3t}} l_t^{\alpha_{4t}} \exp\left(\alpha_{5t}\hat{l}_{it}\right), \tag{2.7}$$

where we impose the symmetry between the two types of labor for simplicity and write the production function in terms of firm i's labor input, l_{it}, and the aggregate labor, l_t. Because $l_{it}^{\alpha_{3t}} \exp\left(\alpha_{5t}\hat{l}_{it}\right)$ is proportional to $l_{it}^{\alpha_{3t}+\alpha} \exp\left((\alpha_{5t} - \alpha)\hat{l}_{it}\right)$ for any constant α, we assume that the steady-state value of α_{5t}, denoted by $\bar{\alpha}_5$, equals zero.

This model features the internal increasing returns to scale à la Benhabib and Farmer

[10] Furthermore, because rts_t is not a variable that pertains to individual firms, the omitted variable in the Phillips curve regressions due to rts_t could be (partially) absorbed by time (or more granular) fixed effects when disaggregated data are utilized for identification. Relatedly, Gagliardone et al. (2025) controls for the returns to scale (i) by using industry-by-time fixed effects at the granular industry level (table 2, columns (c)-(d)) and (ii) by explicitly modeling and estimating translog production functions for each sector, following the estimation methods in Lenzu et al. (2024) (table A.3).

(1994) when $\alpha_{3t} = \alpha_3 > 1$ and $\alpha_{4t} = \alpha_{5t} = 0$. The external increasing returns to scale occur when, e.g., $\alpha_{3t} = 1$, $\alpha_{4t} = \alpha_4 > 0$, and $\alpha_{5t} = 0$, as in Baxter and King (1991, 1993). Finally, the second-order term in the normalized translog form in Equation (2.6) leads to $\alpha_{3t} = \alpha_1 + \alpha_2$, $\alpha_{4t} = 0$, and $\alpha_{5t} = (\beta_{11} + \beta_{12} + \beta_{21} + \beta_{22})\hat{l}_t$ under the symmetry between the two types of labor ($\hat{l}_{it} = \hat{l}_{1,it} = \hat{l}_{2,it}$ and $\hat{l}_t = \hat{l}_{1t} = \hat{l}_{2t}$).

For exposition, we focus on the relationship between real marginal costs, mc, and labor shares, wl/y. It follows from firm i's cost minimization problem that:

$$\frac{w_t l_{it}}{y_{it}} = mc_{it}(\alpha_{3t} + \alpha_{5t}).$$

Thus, the cyclical returns to scale in Proposition 1, $\widehat{rts}_t$, could result from time variations in α_{3t} and α_{5t}. Note that the conventional internal increasing returns to scale, $\alpha_{3t} = \bar{\alpha}_3 > 1$, and the external returns to scale parameter, α_{4t}, do not yield a wedge between $\hat{mc}$ and $\widehat{wl/y}$. That is, the increasing but time-invariant returns to scale cannot explain the seemingly contradictory empirical results regarding the slopes of the conventional and marginal cost Phillips curves.

Suppose that $\alpha_{4t} = 0$ to further clarify the differences between the cyclical returns to scale arising from the (normalized) translog production functions (α_{5t}) and from a Cobb–Douglas production function with time-varying elasticities of output with respect to each input (α_{3t}). At the first order, the production function is given by:

$$\hat{y}_{it} = \underbrace{(\varepsilon^a + \bar{\alpha}_3 \log(\bar{l})\hat{\alpha}_{3t})}_{\text{measured TFP}} + \underbrace{\bar{\alpha}_3 \hat{l}_{it}}_{\text{factor inputs}}.$$

Thus, unlike α_{5t}, α_{3t} affects output and the cyclicality of the measured TFP at the first order. If $\bar{l} > 1$, α_{3t} contributes to the measured TFP procyclically. Because we find no significant evidence of increasingly procyclical TFP responses to monetary policy shocks in recent periods (Remark 2 and Appendix C.5), more procyclical α_{3t} is less likely to be the major driver of the flattening of the conventional Phillips curves. Thus, we focus on α_{5t} in this paper, microfounded by the (normalized) translog production function.[11]

[11] The procyclical component in the measured TFP due to α_{3t} does not disappear when we use the cost-based Solow residual as a measure of the level of productivity, following Hall (1990). Since labor is the only input in Equation (2.7), the corresponding cost share equals one. Thus, the cost-based Solow residual, $\Delta \log y - \Delta \log l$, is given by $\Delta\varepsilon^a + (\bar{\alpha}_3 - 1)\Delta \log l + \log(\bar{l})\Delta\alpha_3$ at the first order, where Δ denotes first differences. Note that a variation in α_{3t} induces a cyclical component in the measured TFP, even

3 Quantitative Analyses

This section constructs a medium-scale DSGE model with a translog production function for quantitative analyses. We estimate the model parameters using Bayesian methods and show that the translog model fits the US macroeconomic data better than the Cobb–Douglas and nested CES models do. By comparing the estimated models using pre- and post-2000 data, we find the following: (i) The translog model successfully replicates both a steep marginal cost Phillips curve and the flattening of the conventional Phillips curves. (ii) In doing so, our mechanism plays a crucial role by connecting the cyclical returns to scale, marginal costs, and inflation. (iii) The translog model also predicts the increased importance of the expected inflation in the US inflation dynamics in recent periods. (iv) The historical decomposition of the pandemic-era inflation data implies that high inflation in 2021 and 2022 was primarily driven by loose monetary policy as well as supply-side shocks.

3.1 Model

We augment the partial equilibrium, four-input, nested CES model of production described in Krusell et al. (2000) and Ohanian et al. (2023) with a translog production function. We combine this novel framework for the supply block of the economy with a quantitative New Keynesian general equilibrium model. Firms use structure, equipment, skilled labor, and unskilled labor in production. Following Christiano et al. (2005), Smets and Wouters (2007), and Justiniano et al. (2010), the model features standard frictions known to be useful for explaining the US aggregate data, such as sticky prices and wages, fixed costs of production, investment adjustment costs, costly capacity utilization, and external consumption habits. The standard structure in the other parts of the model allows us to emphasize the input complementarity and substitution structure, the workings of the translog production function, the resulting cyclical returns to scale, and their contribution to the slopes of the Phillips curves.

We employ a translog production function for our baseline quantitative analyses for several reasons. First, the translog framework naturally provides a microfoundation for cyclical returns to scale, as illustrated in Section 2. Second, translog parameters have a struc-

when the returns to scale are constant, $\bar{\alpha}_3 = 1$ (Fernald, 2014b), or after controlling for the effects of conventional increasing returns to scale, $\bar{\alpha}_3 > 1$ (Basu and Fernald, 1997; Basu et al., 2006). In contrast, α_{5t}, representing our mechanism based on translog production functions, does not influence the measured TFP at the first order.

tural interpretation as a semi-elasticity of output elasticity with respect to each input (e.g., $\beta_{12} = \frac{\partial}{\partial \log l_2}\left(\frac{\partial \log y}{\partial \log l_1}\right)$ in Equation (2.2)), informing on the degree of input complementarity. Third, we show in Section 3.2.3 that the data used for estimation provide sufficient information about the degree of input complementarity, resulting in reasonably precise translog parameter estimates. This is the case despite the fact that a multi-factor translog production function introduces several additional parameters to estimate compared with the standard Cobb–Douglas function. Finally, as a robustness check, we further examine two simpler models in which a smaller number of parameters disciplines the cyclicality of returns to scale. Specifically, we consider (i) a model where the returns to scale co-move with the output gap, with its sensitivity governed by a single parameter, and (ii) a two-factor model à la Smets and Wouters (2007) with three translog parameters. The cyclicality of the returns to scale and its implications for the slopes of the Phillips curves largely remain similar, although the translog model is statistically significantly preferred over the simpler output gap-based model (See Appendices C.6 and C.7).

Below, we briefly introduce the translog model with a focus on the novel elements. See Appendix B for additional details.

3.1.1 Firms

The final goods are produced by combining intermediate goods, indexed by $i \in [0, 1]$, via a Kimball (1995) aggregator.[12] Intermediate goods producers can randomly reset their prices with a probability of $1 - \zeta_p$ in each period. As a result, the model features a New Keynesian price Phillips curve, which is examined in detail in Section 3.3.

The production function for the intermediate good i is given by:

$$\begin{aligned} Y_t(i) &= \exp(\varepsilon_t^a)[K_{st}(i)]^{\alpha_{ks}}[K_{et}(i)]^{\alpha_{ke}}[\gamma^t L_{st}(i)]^{\alpha_s}[\gamma^t L_{ut}(i)]^{\alpha_u} \times \exp(s.o.t._{it}) - \gamma^t \upsilon, \qquad (3.1) \\ s.o.t._{it} &= \left(\hat{k}_{i,et}, \hat{l}_{i,st}, \hat{l}_{i,ut}\right) \boldsymbol{\beta} \left(\hat{k}_{et}, \hat{l}_{st}, \hat{l}_{ut}\right)', \end{aligned}$$

where $K_{st}(i)$, $K_{et}(i)$, $L_{st}(i)$, and $L_{ut}(i)$ are structures, capital equipment services, skilled

[12] There are several modeling strategies for real rigidities, such as quasi-kinked demand curves (Kimball, 1995), firm-specific capital (Eichenbaum and Fisher, 2007), industry-specific factor markets (Woodford, 2003, 2005), and roundabout production networks (Basu, 1995). We employ the quasi-kinked demand curves following Smets and Wouters (2007). The effects of other potential mechanisms may be incorporated into the estimated slope of the price Phillips curve via the Calvo price-stickiness parameter in a reduced-form manner.

labor hours, and unskilled labor hours, respectively. γ is the (gross) growth rate on the balanced growth path (BGP). υ represents the fixed cost in production. The profit is zero on the BGP, $(\Phi - 1)\bar{Y}_t = \gamma^t \upsilon$, where Φ and $\bar{Y}_t$ are the gross price markup and output on the BGP, respectively. The aggregate productivity, ε_t^a, follows an exogenous process. We assume that $\alpha_{ks} + \alpha_{ke} + \alpha_s + \alpha_u = 1$.

Similar to Equation (2.6) in the illustrative model, we assume a normalized translog structure for the second-order terms following Hyun et al. (2024). We use bars and hats to denote the BGP values and log deviations from these values, respectively; $\hat{k}_{i,et} = \log\left(\frac{K_{et}(i)}{\bar{K}_{et}}\right)$, $\hat{l}_{i,st} = \log\left(\frac{L_{st}(i)}{\bar{L}_{st}}\right)$, $\hat{l}_{i,ut} = \log\left(\frac{L_{ut}(i)}{\bar{L}_{ut}}\right)$, $\hat{k}_{et} = \log\left(\frac{K_{et}}{\bar{K}_{et}}\right)$, $\hat{l}_{st} = \log\left(\frac{L_{st}}{\bar{L}_{st}}\right)$, and $\hat{l}_{ut} = \log\left(\frac{L_{ut}}{\bar{L}_{ut}}\right)$, where K_{et}, L_{st}, and L_{ut} are aggregate equipment services, skilled labor hours, and unskilled labor hours, respectively. $\boldsymbol{\beta}$ is a symmetric matrix consisting of the six translog parameters:

$$\boldsymbol{\beta} = \begin{pmatrix} \beta_{kk} & \beta_{ks} & \beta_{ku} \\ \beta_{ks} & \beta_{ss} & \beta_{su} \\ \beta_{ku} & \beta_{su} & \beta_{uu} \end{pmatrix}.$$

We assume that structures appear only in the first-order term in the production function following Krusell et al. (2000) and Ohanian et al. (2023). However, arbitrary degrees of input complementarity and substitutability among the other three factors, represented by $\boldsymbol{\beta}$, are allowed by adopting a translog functional form.

Firms use four inputs, $K_{st}(i)$, $K_{et}(i)$, $L_{st}(i)$, and $L_{ut}(i)$, to minimize the production cost, $R_{st}K_{st}(i) + R_{et}K_{et}(i) + W_{st}L_{st}(i) + W_{ut}L_{ut}(i)$, where R_{st}, R_{et}, W_{st}, and W_{ut} are the nominal rental rate of structures, rental rate of equipment services, skilled wage rate, and unskilled wage rate, respectively. $MC_t(i)$ is the Lagrange multiplier associated with the constraint $Y_t(i) \geq Y$, representing nominal marginal costs.

The FOCs for this cost-minimization problem are as follows:

$$K_{st}(i): \quad \frac{R_{st}K_{st}(i)}{P_t\tilde{Y}_t(i)} = mc_t(i)\alpha_{ks}, \tag{3.2}$$

$$K_{et}(i): \quad \frac{R_{et}K_{et}(i)}{P_t\tilde{Y}_t(i)} = mc_t(i)\alpha_{ke,t}, \quad \alpha_{ke,t} \equiv \left(\alpha_{ke} + \beta_{kk}\hat{k}_{et} + \beta_{ks}\hat{l}_{st} + \beta_{ku}\hat{l}_{ut}\right), \tag{3.3}$$

$$L_{st}(i): \quad \frac{W_{st}L_{st}(i)}{P_t\tilde{Y}_t(i)} = mc_t(i)\alpha_{st}, \quad \alpha_{st} \equiv \left(\alpha_s + \beta_{ks}\hat{k}_{et} + \beta_{ss}\hat{l}_{st} + \beta_{su}\hat{l}_{ut}\right), \tag{3.4}$$

$$L_{ut}(i): \quad \frac{W_{ut}L_{ut}(i)}{P_t\tilde{Y}_t(i)} = mc_t(i)\alpha_{ut}, \quad \alpha_{ut} \equiv \left(\alpha_u + \beta_{ku}\hat{k}_{et} + \beta_{su}\hat{l}_{st} + \beta_{uu}\hat{l}_{ut}\right), \tag{3.5}$$

where $\tilde{Y}_t(i) = Y_t(i) + \gamma^t \upsilon$, P_t is the aggregate price level, and $mc_t(i)$ represents real marginal costs. Thus, the factor income shares, adjusted by fixed costs in production, are proportional to marginal costs times the elasticity of $\tilde{Y}_t(i)$ with respect to each input. Note that $\alpha_{ke,t}$, α_{st}, and α_{ut} are time-varying and endogenously determined. Furthermore, although the second-order term in the production function (3.1) vanishes in log linearization, the translog parameters appear in the log-linearized FOCs and influence firms' factor input decisions.

3.1.2 Households

A representative household consists of a continuum of members indexed by $j \in [0, 1]$. In each period, $\chi_t \in (0, 1)$ fraction of members are skilled workers, and the remaining $1 - \chi_t$ fraction of workers are unskilled, where χ_t follows an exogenous process. We assume that members with $j \in [0, 1 - \chi_t)$ and $j \in [1 - \chi_t, 1]$ are unskilled and skilled, respectively.

The period utility function of member j is given by $\frac{A_{1j}}{1-\sigma_c}\big(C_t(j) - hC_{t-1}\big)^{1-\sigma_c} \exp\Big(A_{2j}\frac{\sigma_c - 1}{1+\sigma_l} H_t(j)^{1+\sigma_l}\Big)$, where $C_t(j)$ is consumption, hC_{t-1} represents an external consumption habit, and $H_t(j)$ is labor hours. The preference parameters, A_{1j} and A_{2j}, vary across worker types. For skilled workers, $A_{1j} = A_{2j} = 1$. For unskilled workers, $A_{1j} = A_1$ and $A_{2j} = A_2$, where the values of A_1 and A_2 are chosen to maintain the "representativeness" of agents such that consumption and labor hours are the same across the worker types on the BGP. This property allows us to deviate minimally from the representative agent framework. However, $C_t(j)$ and $H_t(j)$ can vary across types in the short term.

A family head solves the resource allocation and portfolio choice problems by deciding the consumption $(C_t(j))$ of the family members, family-level risk-free nominal bonds (B_t), structure investment (I_{st}), equipment investment (I_{et}), structure stock $(\tilde{K}_{st})$, equipment stock $(\tilde{K}_{et})$, and equipment utilization rate (Z_t) to maximize the following utilitarian welfare function:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \int_0^1 \left[A_{1j} \frac{1}{1-\sigma_c} \big(C_t(j) - hC_{t-1}\big)^{1-\sigma_c} \exp\left(A_{2j} \frac{\sigma_c - 1}{1+\sigma_l} H_t(j)^{1+\sigma_l} \right) \right] dj$$

subject to the budget constraint as follows:

$$\int C_t(j)\, dj + I_{st} + I_{et} + \frac{B_t}{\exp(\varepsilon_t^b) R_t P_t} + T_t \tag{3.6}$$

$$\leq \frac{B_{t-1}}{P_t} + \frac{W_{st}\int_{1-\chi_t}^{1} H_t(j)\,dj}{P_t} + \frac{W_{ut}\int_{0}^{1-\chi_t} H_t(j)\,dj}{P_t} + \frac{R_{st}\tilde{K}_{s,t-1}}{P_t} + \frac{R_{et}Z_t\tilde{K}_{e,t-1}}{P_t} - a(Z_t)\tilde{K}_{e,t-1} + \frac{\Pi_t^w}{P_t} + \frac{\Pi_t}{P_t}$$

and the capital accumulation equations for structure and equipment: $\tilde{K}_{st} = (1-\delta_s)\tilde{K}_{s,t-1} + \left[1 - S_s\left(\frac{I_{st}}{I_{s,t-1}}\right)\right] I_{st}$ and $\tilde{K}_{et} = (1-\delta_e)\tilde{K}_{e,t-1} + \exp(\varepsilon_t^i)\left[1 - S_e\left(\frac{I_{et}}{I_{e,t-1}}\right)\right] I_{et}$. R_t is the (gross) nominal returns to bonds. ε_t^b denotes an exogenous premium in returns to bonds, reflecting a structural shock to the demand for safe and liquid assets (Fisher, 2015). T_t represents lump-sum taxes. Π_t^w and Π_t are the profits of labor unions and firms paid out as dividends. $a(\cdot)$ reflects costly capacity utilization. We assume that $a(\bar{Z}) = 0$ and $\bar{Z} = 1$, where $\bar{Z}$ is the steady-state level of utilization. Following Smets and Wouters (2007), we reparametrize $\frac{a'(\bar{Z})}{a''(\bar{Z})}$ as $\frac{1-\psi}{\psi}$. The amount of structure and equipment services households can rent to firms is given by $K_{st} = \tilde{K}_{s,t-1}$ and $K_{et} = Z_t\tilde{K}_{e,t-1}$, respectively. $S_s(\cdot)$ and $S_e(\cdot)$ are the investment adjustment costs for structure and equipment such that $S_s(\gamma) = S_e(\gamma) = 0$, $S_s'(\gamma) = S_e'(\gamma) = 0$, $S_s''(\gamma) = \varphi_s > 0$, and $S_e''(\gamma) = \varphi_e > 0$. ε_t^i is an exogenously determined (equipment) investment-specific productivity.

3.1.3 Other Components of the Model

Households relegate labor supply decisions to labor market institutions. The skilled labor supply from households is differentiated and channeled to skilled labor unions. These labor unions set nominal wages for each differentiated skilled labor service subject to Calvo-type frictions, where the resetting probability in each period is denoted by $1 - \zeta_s$. The differentiated labor services are combined into composite skilled labor services using a Kimball aggregator and sold to intermediate goods producers. This structure yields a New Keynesian skilled wage Phillips curve. We introduce a similar labor market structure for unskilled labor. With a probability of ζ_u, the unskilled labor unions cannot adjust wages in each period. The model similarly features a New Keynesian unskilled wage Phillips curve. The skilled and unskilled wage markup shocks have independent ARMA(1,1) processes, following Smets and Wouters (2007).

The total supply of unskilled labor hours to intermediate goods producers is given by $\int_0^{1-\chi_t} H_t(j)\,dj$. Thus, the market for unskilled labor clears when $L_{ut} = \int L_{ut}(i)di = \int_0^{1-\chi_t} H_t(j)\,dj$. Similarly, we have $L_{st} = \int L_{st}(i)di = \int_{1-\chi_t}^{1} H_t(j)\,dj$.

The central bank's monetary policy rule is given by: $\frac{R_t}{\bar{R}} = \left(\frac{R_{t-1}}{\bar{R}}\right)^{\rho}\left[\left(\frac{\Pi_t}{\bar{\Pi}}\right)^{r_\pi}\left(\frac{Y_t}{Y_t^*}\right)^{r_y}\right]^{1-\rho}\left(\frac{Y_t/Y_{t-1}}{Y_t^*/Y_{t-1}^*}\right)^{r_{\Delta y}}\exp(\varepsilon_t^r)$, where $\bar{R}$ is the steady state nominal (gross) rate and Y_t^* is the natural

level of output in an economy with flexible prices and wages and without markup shocks to prices and wages. ρ determines the degree of interest rate smoothing. r_π, r_y, and $r_{\Delta y}$ reflect the responsiveness of the policy rate to each forcing variable.

The aggregate resource constraint in this economy is standard: $C_t+I_t+G_t+a(Z_t)\tilde{K}_{e,t-1} = Y_t$, where aggregate consumption and investment are given by $C_t = \int C_t(j)\,dj$ and $I_t = I_{st} + I_{et}$, respectively. Government spending, G_t, is exogenously determined.

3.1.4 Cobb–Douglas and Nested CES Models

The translog model nests similar models on the basis of the Cobb–Douglas and nested CES production functions in log linearization of the corresponding equilibrium conditions. Clearly, the Cobb–Douglas model can be obtained by imposing the restriction that the translog parameters are zero, $\boldsymbol{\beta} = 0$.

Next, we consider a nested CES production function employed in Krusell et al. (2000) and Ohanian et al. (2023). In a deviation form based on dimensionless parameters such as those in Cantore and Levine (2012); Cantore et al. (2014, 2015), we can write the nested CES production function as follows:

$$\frac{Y_t(i)+\gamma^t\upsilon}{\bar{Y}_t+\gamma^t\upsilon} = \exp(\varepsilon_t^a)\left[\frac{K_{st}(i)}{\bar{K}_{st}}\right]^{\alpha_{ks}}\left\{\tilde{\alpha}_u\left[\frac{L_{ut}(i)}{\bar{L}_{ut}}\right]^{\phi_{ux}}+\tilde{\alpha}_x\left(\tilde{\alpha}_k\left[\frac{K_{et}(i)}{\bar{K}_{et}}\right]^{\phi_{ks}}+\tilde{\alpha}_s\left[\frac{L_{st}(i)}{\bar{L}_{st}}\right]^{\phi_{ks}}\right)^{\frac{\phi_{ux}}{\phi_{ks}}}\right\}^{\frac{1-\alpha_{ks}}{\phi_{ux}}},$$

where $\tilde{\alpha}_u+\tilde{\alpha}_x = 1$, $\tilde{\alpha}_k+\tilde{\alpha}_s = 1$, and $-\infty \le \phi_{ux}, \phi_{ks} \le 1$. Similar to the results shown in Table 1, this nested CES function can be converted to a translog form. The steady-state income shares in the translog function are obtained as follows: $\alpha_u = (1-\alpha_{ks})\tilde{\alpha}_u$, $\alpha_{ke} = (1-\alpha_{ks})\tilde{\alpha}_x\tilde{\alpha}_k$, and $\alpha_s = (1-\alpha_{ks})\tilde{\alpha}_x\tilde{\alpha}_s$. See Appendix B.4 for similar results for the translog parameters.

Finally, all three models share the same log-linearized aggregate production functions around the BGP, $\hat{Y}_t = \Phi(\varepsilon_t^a + \alpha_{ks}\hat{K}_{st} + \alpha_{ke}\hat{K}_{et} + \alpha_s\hat{L}_{st} + \alpha_u\hat{L}_{ut})$.

3.2 Bayesian Estimation

We estimate the three models with different production functions using pre- and post-2000 US time series data and Bayesian methods. A comparison of the model fits reveals that the translog model better matches the data than the Cobb–Douglas and CES models do. Fur-

thermore, the estimated production function parameters reflect potential structural changes in production technology between the two sample periods, which could further influence the cyclicality of returns to scale and the identified slopes of the Phillips curves.

3.2.1 Data, Priors, and Methods

Data. We use eight quarterly variables and three annual variables for estimation. The quarterly variables include skilled workers' employment shares and the seven observables employed by Smets and Wouters (2007), such as growth rates of per capita real GDP, consumption, and investment; labor hours; the growth rate of the real wage rate; price inflation; and the nominal interest rate. We replace the federal funds rate with the shadow rate of Wu and Xia (2016) from 2009:q1 to 2015:q4 because of the binding zero lower bound.

Using the Current Population Survey (CPS), we compute skilled employment shares. Following Krusell et al. (2000) and Ohanian et al. (2023), college graduates are considered skilled workers. Because of the data limitation in the early period, this variable is only available for the first quarters until 1975. Since then, it has been available every quarter. We remove the seasonal variation using X-13 ARIMA.

In light of two types of capital (structures and equipment) and labor (skilled and unskilled), we borrow three related annual variables from Ohanian et al. (2023) to discipline the model parameters. These variables are the growth rate of per capita equipment stock (including intellectual property products), the changes in relative labor hours, and the changes in the relative wage rates between skilled and unskilled workers. The capital stock is constructed using perpetual inventory methods. The two labor market variables are calculated using the CPS data. These annual variables are assumed to be observable in the fourth quarter of each year.

Our sample begins in 1966:q1 as in Smets and Wouters (2007) and ends in 2019:q4 before the recent pandemic. We divide the data into early (1966-99) and late (2000-19) samples and estimate each model for each sample period. In total, we have six estimated models for the three production functions and the two sample periods.

State-space system. The state equation is constructed using the Sims (2002) method for each parameter value. The measurement equation consists of available observations among the eleven variables described above each quarter, leading to a mixed-frequency state-space model (see, e.g., Schorfheide and Song, 2015, 2025).

Real GDP, consumption, investment, price inflation, and the nominal interest rate are mapped to Y_t, C_t, I_t, P_t/P_{t-1}, and R_t in the model, respectively. The aggregate labor hours are matched with $H_t \equiv \int_0^1 H_t(j)\,dj = L_{st} + L_{ut}$. The aggregate wage rate is given by the total wage bills divided by the aggregate hours, $W_t \equiv \frac{W_{st}L_{st}+W_{ut}L_{ut}}{P_t H_t}$. The skilled employment share corresponds to χ_t. We assume that $\log(\chi_t) = \log(\bar{\chi}) + \varepsilon_t^{\chi}$, where ε_t^{χ} follows a mean-zero AR(1) process. We set $\bar{\chi}$ at 0.21 for the pre-2000 sample and 0.325 for the post-2000 sample, which are the values of χ_t around the midpoint of each sample.

The three annual variables are observed (with measurement errors) in the fourth quarter. The growth rate of per capita equipment stock corresponds to $\log(\tilde{K}_{et}/\tilde{K}_{e,t-4}) + \nu_{kt}$. ν_{kt} represents a measurement error in the empirical capital stock variable computed using perpetual inventory methods. ν_{kt} has an independently and identically distributed normal density with a mean of zero and a standard deviation of $\sigma_{\nu,k}$. The relative annual labor hours, denoted by $rell_t^{ann}$, equal $\frac{L_{st}+L_{s,t-1}+L_{s,t-2}+L_{s,t-3}}{L_{ut}+L_{u,t-1}+L_{u,t-2}+L_{u,t-3}}$. The annual change in this variable, $\log(rell_t^{ann}/rell_{t-4}^{ann})$, is observed with measurement errors, $\nu_{ht} \sim N(0, \sigma_{\nu,h}^2)$, because the labor market variables in the CPS are known to be subject to nonnegligible measurement errors (see, e.g., Bound and Krueger, 1991). Similarly, the relative annual wage rate, $relw_t^{ann}$, is given by $\frac{W_{st}L_{st}+\ldots+W_{s,t-3}L_{s,t-3}}{L_{st}+\ldots+L_{s,t-3}} / \frac{W_{ut}L_{ut}+\ldots+W_{u,t-3}L_{u,t-3}}{L_{ut}+\ldots+L_{u,t-3}}$. The measurement of its annual change, $\log(relw_t^{ann}/relw_{t-4}^{ann})$, involves measurement errors denoted by $\nu_{wt} \sim N(0, \sigma_{\nu,w}^2)$.

The model includes nine structural shocks and three measurement errors. The structural shocks include shocks to TFP (ε_t^a), investment-specific productivity (ε_t^i), monetary policy (ε_t^r), risk premium (ε_t^b), government spending (ε_t^g), price markup (ε_t^p), skilled wage markup (ε_t^s), unskilled wage markup (ε_t^u), and the employment share of skilled workers (ε_t^{χ}).

Priors. We use standard priors for conventional parameters, similar to Smets and Wouters (2007) and Justiniano et al. (2010). For the other parameters, we set the following priors. We assume that $\alpha_{ks} \sim N(0.1, 0.005^2)$ given the results in Krusell et al. (2000) and Ohanian et al. (2023). Combined with the prior for $\alpha_{ke} \sim N(0.25, 0.02^2)$, this assumption implies that the capital income share a priori ranges from 0.3-0.4 within the two standard deviation intervals. Given the steady increase in the skilled labor supply (Goldin and Katz, 2009), the mean of α_s is assumed to be 0.2 and 0.3 for the early and late samples, respectively.

The translog parameters, e.g., β_{kk}, have normal priors with a mean of zero and a standard deviation of 0.15. This prior is rather loose and less informative, which is illustrated by the fact that previous estimates based on nested CES functions are within a one-standard deviation interval. For example, when $\phi_{ux} = 0.401$, $\phi_{ks} = -0.495$, and $\alpha_{ks} = 0.117$, following

Krusell et al. (2000) in combination with $\alpha_{ke} = 0.25$ and $\alpha_s = 0.2$, the corresponding translog parameters are less than 0.09 in absolute value (see Equation (B.19) in Appendix B.1.2).

Shocks to χ_t, ε_t^{χ}, has an AR(1) process: $\varepsilon_t^{\chi} = \rho_{\chi}\varepsilon_{t-1}^{\chi} + \eta_t^{\chi}$. Because this variable is slow-moving and steadily increases throughout the sample period, we assume that ρ_{χ} has a beta distribution with a mean of 0.9. Finally, the standard deviations of the measurement errors, $\sigma_{\nu,k}$, $\sigma_{\nu,h}$, and $\sigma_{\nu,w}$, have inverse gamma distributions with a mean of 0.15 and a standard deviation of 0.03. Thus, a priori, measurement errors are assumed not to be excessively large.

We fix δ_s and δ_e at $0.026/4$ and $0.16/4$, respectively, according to Fernald (2014a, p. 18) and Ohanian et al. (2023, fig 3.3(b)). The steady-state share of government spending (0.18), steady-state wage markups (1.5), and Kimball curvature parameters (10) are set at the same values as those in Smets and Wouters (2007).

Methods. The posterior distributions are computed using a Metropolis–Hastings algorithm with a chain length of 200,000. The likelihood of the state-space system is calculated using a Kalman filter. The acceptance rates of the chain for the six cases range from 22%-30%. We summarize the posterior distributions using the mode and credible intervals. See Appendix C.3 for the complete list of the estimated parameters, priors, and posteriors for the six cases.

3.2.2 Model Comparison

This section shows that the estimated translog model better matches the data than the two alternative models do. Specifically, we test and reject the restrictions imposed on the translog parameters by the Cobb–Douglas and nested CES forms, compare the size of measurement errors across the models, and assess the log-likelihood in the posterior modes and marginal data densities (MDDs).[13] The results favor the translog model.

The production function (3.1) implies that the aggregate returns to scale are influenced by the fixed cost and the translog structure. Let lowercase letters denote the detrended real variables around the BGP, e.g., $y_t = \frac{Y_t}{\gamma^t}$, $k_{et} = \frac{K_{et}}{\gamma^t}$, $l_{st} = L_{st}$, and $l_{ut} = L_{ut}$. In log-linearization:

$$\widehat{rts}_t = \underbrace{-\frac{\Phi - 1}{\Phi}\hat{y}_t}_{\widehat{rts}^{fc}} + \underbrace{\left(\beta_k.\hat{k}_{et} + \beta_s.\hat{l}_{st} + \beta_u.\hat{l}_{ut}\right)}_{\widehat{rts}^{tl}}, \tag{3.7}$$

[13] The MDD avoids overfitting due to extra degrees of freedom by systematically accounting for the number of parameters in the model, similar to its asymptotic approximation, i.e., the Bayesian information criterion.

Table 2: Model comparison

	(1)	(2)	(3)	(4)	(5)	(6)
	Early sample (1966-99)			Late sample (2000-19)		
	Cobb–Douglas	Nested CES	Translog	Cobb–Douglas	Nested CES	Translog
Panel A. Returns to scale parameters						
$\beta_{k\cdot}$	0	0	-0.27	0	0	-0.21
	-	-	(-0.63, -0.13)	-	-	(-0.27, -0.19)
$\beta_{s\cdot}$	0	0	-0.13	0	0	0.23
	-	-	(-0.23, -0.06)	-	-	(0.15, 0.33)
$\beta_{u\cdot}$	0	0	-0.37	0	0	-0.28
	-	-	(-0.69, -0.11)	-	-	(-0.48, -0.16)
Panel B. Measurement errors						
σ_{ν,k_e}	3.79	3.83	3.87	0.23	1.69	0.37
	(3.19, 4.58)	(3.23, 4.67)	(3.22, 4.68)	(0.19, 0.38)	(1.25, 2.34)	(0.30, 0.50)
$\sigma_{\nu,h}$	3.13	0.14	0.14	0.14	0.14	0.14
	(2.66, 3.83)	(0.10, 0.21)	(0.10, 0.21)	(0.10, 0.21)	(0.10, 0.22)	(0.12, 0.16)
$\sigma_{\nu,w}$	0.14	1.44	1.21	2.55	1.46	1.28
	(0.10, 0.22)	(1.20, 1.84)	(1.04, 1.49)	(2.11, 3.15)	(1.18, 1.86)	(1.12, 1.66)
Panel C. Log-likelihood in the posterior mode and marginal data densities						
log-likelihood	-96.10	-87.20	-70.27	71.88	85.13	114.42
log MDD	-347.75	-332.69	-311.52	-193.08	-165.21	-173.65

Notes: Panel A shows the estimated returns to scale parameters ($\beta_{k\cdot}$, $\beta_{s\cdot}$, and $\beta_{u\cdot}$) in the posterior mode and the equal-tailed 95% credible intervals. Columns (1)-(3) and (4)-(6) present the results from the three different models estimated using the early and late samples, respectively. Panel B regards the magnitude of the measurement errors. Panel C illustrates the log-likelihood in the posterior mode and the MDD. The MDD is computed using the Sims et al. (2008) algorithm.

where $\beta_{k\cdot} = \beta_{kk}+\beta_{ks}+\beta_{ku}$, $\beta_{s\cdot} = \beta_{ks}+\beta_{ss}+\beta_{su}$, and $\beta_{u\cdot} = \beta_{ku}+\beta_{su}+\beta_{uu}$. The countercyclical variation in $\widehat{rts}_t$ due to the fixed costs in production, captured by $\widehat{rts}^{fc}$, is shared among the three models.[14] In contrast, the cyclical variation originating from the second-order properties of the production function, denoted by $\widehat{rts}^{tl}$, is unique to the translog model. For the Cobb–Douglas model, $\beta_{k\cdot} = \beta_{s\cdot} = \beta_{u\cdot} = 0$ because all six translog parameters are zero. Like $\beta_{11}+\beta_{12} = \beta_{12}+\beta_{22} = 0$ for the two-input CES model in Table 1, the translog form of the nested CES model also satisfies $\beta_{k\cdot} = \beta_{s\cdot} = \beta_{u\cdot} = 0$, implying that $\widehat{rts}^{tl}_t = 0$.[15] We test the validity of these restrictions by estimating the translog model.

Panel A in Table 2 shows the results. The table illustrates the estimated returns to scale parameters in the posterior mode and the equal-tailed 95% credible intervals. Columns (1)-(3) and (4)-(6) present the results from the three different models estimated using the early and late samples, respectively. Per the discussion above, the returns to scale parameters,

[14] The fixed costs are linked to the gross price markup, Φ, through the zero-profit condition on the BGP.

[15] Choosing six translog parameters subject to these three restrictions and the additional restriction $\beta_{ku} = -\tilde{\alpha}_k\beta_{uu}$, arising from the specific nesting order of the two CES aggregators, is equivalent to determining the two CES parameters, ϕ_{ux} and ϕ_{ks}. See Appendix B.4.

$\beta_{k\cdot}$, $\beta_{s\cdot}$, and $\beta_{u\cdot}$, are zero in Columns (1), (2), (4), and (5). Note that none of the six 95% credible intervals in Panel A Columns (3) and (6) include zero, rejecting the prediction of the time-invariant $\widehat{rts}_t^{tl}$ under the Cobb–Douglas and CES models.

The results shown in Panel B imply that rationalizing the data for equipment stock, relative labor hours, and relative wage rates jointly is challenging in a general equilibrium framework. There is no case in Panel B where the magnitude of all three measurement errors is small. Given that, the translog model performs no worse than the other two models do. For example, the translog model better explains the post-2000 relative wage data ($\sigma_{\nu,w}$) than the Cobb–Douglas model does. Compared with the translog model, the nested CES model is less successful at matching the capital stock data (σ_{ν,k_e}) in the recent sample.

For the pre-2000 data, the translog model is consistently favored according to both the log-likelihood in the posterior mode and the MDD (Panel C). For the post-2000 data, the two criteria alternatively select the translog and nested CES models. Thus, given the results in Panels A-C, we conclude that the translog model, which predicts the time-varying returns to scale, matches the data better than the two alternative models do, at least in the posterior modes.

3.2.3 Production Function Parameters

Next, we discuss the production function parameters. Table 3 shows that the steady-state factor income shares (Panel A) and price markups (Panel C) are largely comparable across the three models in a given sample period. Between the two sample periods, the skilled labor share (α_s) increases, and the price markup (Φ) decreases. Consistent with the results in Smets and Wouters (2007, table 1A), Φ is relatively large, and capital income shares are relatively small compared with conventional calibration in all six cases.

The data used for estimation provide sufficient information about the production function parameters. The posterior distributions of the translog parameters are centered around non-zero values and are significantly less dispersed than their priors (Panel B). The lengths of the 95% credible intervals are considerably shorter than those of the similar intervals a priori, which are approximately 0.6 (see Figure C.1 in Appendix C.3 for the shape of the posterior distributions compared with the priors in each sample period). Furthermore, the Hessian of the posterior density, evaluated at the mode, has full rank, implying that our model passes the diagnostic test for parameter identifiability in a DSGE model (Canova and Sala, 2009).

The translog model rejects the Cobb–Douglas and nested CES specifications, which is

Table 3: Production function parameters

	(1)	(2)	(3)	(4)	(5)	(6)
	Early sample (1966-99)			Late sample (2000-19)		
	Cobb–Douglas	Nested CES	Translog	Cobb–Douglas	Nested CES	Translog
Panel A. Factor income shares in the steady state						
α_{ks}	0.10	0.10	0.10	0.10	0.09	0.10
	(0.09, 0.11)	(0.09, 0.10)	(0.09, 0.10)	(0.09, 0.11)	(0.08, 0.10)	(0.10, 0.11)
α_{ke}	0.19	0.18	0.20	0.14	0.12	0.09
	(0.16, 0.22)	(0.15, 0.21)	(0.16, 0.22)	(0.11, 0.17)	(0.10, 0.15)	(0.08, 0.10)
α_s	0.09	0.15	0.12	0.20	0.24	0.24
	(0.06, 0.18)	(0.12, 0.19)	(0.10, 0.14)	(0.13, 0.24)	(0.23, 0.31)	(0.20, 0.26)
α_u	0.62	0.58	0.58	0.57	0.54	0.57
	(0.52, 0.66)	(0.53, 0.63)	(0.56, 0.62)	(0.53, 0.63)	(0.46, 0.56)	(0.55, 0.61)
Panel B. Translog and CES parameters						
β_{kk}	0	0.11	-0.09	0	-0.04	0.04
	-	(0.09, 0.13)	(-0.42, -0.03)	-	(-0.14, -0.03)	(0.01, 0.04)
β_{ks}	0	-0.03	-0.06	0	0.08	0.00
	-	(-0.04, -0.01)	(-0.12, -0.02)	-	(0.07, 0.19)	(-0.01, 0.04)
β_{ku}	0	-0.08	-0.12	0	-0.04	-0.25
	-	(-0.10, -0.07)	(-0.23, 0.07)	-	(-0.05, -0.03)	(-0.32, -0.23)
β_{ss}	0	0.09	0.08	0	0.00	0.13
	-	(0.07, 0.12)	(0.07, 0.10)	-	(-0.09, 0.02)	(0.04, 0.18)
β_{su}	0	-0.07	-0.16	0	-0.07	0.10
	-	(-0.09, -0.05)	(-0.20, -0.10)	-	(-0.14, -0.07)	(0.04, 0.20)
β_{uu}	0	0.15	-0.10	0	0.11	-0.13
	-	(0.12, 0.18)	(-0.33, 0.02)	-	(0.10, 0.19)	(-0.33, 0.00)
ϕ_{ux}	0	0.74	-	0	0.52	-
	-	(0.60, 0.84)	-	-	(0.47, 0.86)	-
ϕ_{ks}	0	0.78	-	0	-0.65	-
	-	(0.64, 0.92)	-	-	(-1.95, -0.57)	-
Panel C. Gross price markup in the steady state						
Φ	1.65	1.69	1.63	1.36	1.36	1.39
	(1.53, 1.77)	(1.54, 1.81)	(1.49, 1.77)	(1.23, 1.47)	(1.19, 1.50)	(1.27, 1.50)

Notes: Columns (1)-(3) and (4)-(6) illustrate the estimation results for the Cobb–Douglas, nested CES, and translog models using the early and late samples, respectively. Panel A shows the factor income shares in the steady state in the posterior mode and the equal-tailed 95% credible intervals. Panels B and C cover the translog and CES parameters and the steady-state price markups, respectively.

consistent with the model comparison results in Table 2. Most of the 95% credible intervals for the translog parameters in Columns (3) and (6) in Panel B do not include zero, rejecting the Cobb–Douglas function. In Columns (2) and (5), we show the implied translog parameters by ϕ_{ux}, ϕ_{ks} and the nested CES function. Note that the CES-based β_{kk} and β_{uu} (and β_{su} for the late sample) have different signs from the corresponding translog estimates in Columns (3) and (6). Furthermore, the magnitudes of β_{ku} and β_{ss} in Column (5) substantially differ from those in Column (6). Thus, despite being tractable, the nested CES function might not be flexible enough to capture several aspects of the input complementarity and

substitution patterns in the data.

Finally, by comparing Columns (3) and (6) in Panel B, we find that β_{kk}, β_{ks}, β_{ku}, and β_{su}, in particular, changed substantially. An increase in β_{kk} probably reflects an increase in intellectual property products capital (Koh et al., 2020) and an innate positive externality in "ideas" (Romer, 1990). β_{ks} is closely related to the degree of the capital-skill complementarity (Griliches, 1969; Autor et al., 2003; Chirinko, 2008; Chirinko and Mallick, 2017; Berlingieri et al., 2024), which also increases in the nested CES model, as shown in Columns (2) and (5). The capital deepening and concurrent automation of low- and medium-skill tasks might explain the decrease in β_{ku} (Goos and Manning, 2007; Michaels et al., 2014). A negative β_{su} for the pre-2000 data seems consistent with the conventional view on the substitutability between skilled and unskilled labor (Autor et al., 2008). However, when estimated using post-2000 data, β_{su} becomes positive (Column (6)).[16] Furthermore, this pattern does not occur in the nested CES model; β_{su} in Columns (2) and (5) are similar to each other. In Section 4, using industry-level panel data and IVs, we present empirical results that are consistent with the translog model above.

3.3 Slopes of the Phillips Curves

This section presents the main quantitative exercise: the analysis of the slopes of the Phillips curves. We show that the translog model can jointly replicate the steep marginal cost Phillips curve and the flattening of conventional Phillips curves in recent decades. In contrast, the Cobb–Douglas and nested CES models rely on unrealistically sticky prices and quite flat Phillips curves to match the inflation data after 2000.

The New Keynesian price Phillips curve in our model features real marginal costs as a forcing variable:

$$\hat{\pi}_t = \tilde{\beta}\gamma \mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa \hat{mc}_t + \varepsilon_t^p, \tag{3.8}$$

where $\tilde{\beta} = \beta\gamma^{-\sigma_c}$ is the discount factor in the detrended economy. The slope of this marginal cost Phillips curve, κ, is determined by $\frac{(1-\zeta_p\tilde{\beta}\gamma)(1-\zeta_p)}{\zeta_p}\frac{1}{1+\theta_p(\Phi-1)}$. Relative to the illustrative model in Section 2, κ includes a term representing strategic complementarity in price setting, depending on the curvature of the Kimball aggregator ($\theta_p = 10$ following Smets and Wouters

[16] See also Havranek et al. (2024, table 5) for a wide range of estimates in the literature, including the case of gross complementarity between skilled and unskilled labor.

2007) and the gross price markup (Φ).

We compute the model-predicted slopes of the Phillips curves with different forcing variables as follows. Let x_t denote a measure of economic activity, such as the real marginal cost, output gap, labor gap, and labor shares. Then, the Phillips curve based on x_t is given by:

$$\hat{\pi}_t = \beta_x \mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa_x x_t + error_{x,t},$$

where κ_x represents the slope of this Phillips curve. We project $\hat{\pi}_t$ on $\mathbb{E}_t[\hat{\pi}_{t+1}]$ and x_t and obtain κ_x from the projection coefficient:

$$\begin{pmatrix} * \\ \kappa_x \end{pmatrix} = \left[\text{var}_{c,mp} \begin{pmatrix} \mathbb{E}_t[\hat{\pi}_{t+1}] \\ x_t \end{pmatrix} \right]^{-1} \text{cov}_{c,mp} \left[\begin{pmatrix} \mathbb{E}_t[\hat{\pi}_{t+1}] \\ x_t \end{pmatrix}, \hat{\pi}_t \right], \tag{3.9}$$

where subscripts c and mp indicate that we use the cyclical variation with the periodicity from 6-32 quarters driven by monetary policy shocks. We compute the (co)variances in Equation (3.9) by integrating the frequency-domain representation of the model (while only allowing monetary policy shocks) over the business cycle frequencies. Equivalently, we can obtain κ_x by applying the ideal bandpass filter to the infinite time series of $\hat{\pi}_t$, $\mathbb{E}_t[\hat{\pi}_{t+1}]$, and x_t and use monetary policy shock series as IVs.[17] Thus, our formula corresponds to the population version of empirical practices of estimating Phillips curves using detrended data by filters (Mavroeidis et al., 2014; Stock and Watson, 2020) and the identified monetary policy shocks as instruments (Barnichon and Mesters, 2020, 2021). However, our results are robust to not using filters (see Appendix C.4).

We consider the following forcing variables: marginal costs ($\hat{mc}_t$), the output gap ($\hat{y}_t - \hat{y}^*_t$), the labor gap ($\hat{h}_t - \hat{h}^*_t$), labor shares ($\widehat{w_t h_t / y_t}$) following Galí and Gertler (1999); Sbordone (2002); Eichenbaum and Fisher (2007), and fixed cost adjusted labor shares ($\widehat{w_t h_t / (y_t + \upsilon)}$). The last variable is included because it equals $\hat{mc}_t$ in the Cobb–Douglas model (see Equations (3.4) and (3.5)). Note that when $x_t = \hat{mc}_t$, κ_x coincides with the underlying slope of the New Keynesian price Phillips curve, κ, in Equation (3.8).

[17] When $x = mc$, this method correctly identifies κ in Equation (3.8). Let $\mathcal{B}(L)$ denote the ideal bandpass filter, where L is a lag operator. Because $\mathcal{B}(L)$ is a linear, time-invariant filter, applying it to Equation (3.8) leads to $\{\mathcal{B}(L)\hat{\pi}_t\} = \tilde{\beta}\gamma\{\mathcal{B}(L)\mathbb{E}_t[\hat{\pi}_{t+1}]\} + \kappa\{\mathcal{B}(L)\hat{mc}_t\} + \{\mathcal{B}(L)\varepsilon^p_t\}$. Thus, the slope remains to be κ. Furthermore, because monetary policy shocks $\{\varepsilon^r_t\}$ are orthogonal to the basis of the residual (i.e., price markup shocks), $\{\varepsilon^p_t\}$, the instruments' exclusion restriction is satisfied.

Table 4: Slopes of the Phillips curves

	(1)	(2)	(3)	(4)	(5)	(6)
	Early sample (1966-99)			Late sample (2000-19)		
	Cobb–Douglas	Nested CES	Translog	Cobb–Douglas	Nested CES	Translog
Panel A. Calvo price stickiness parameter						
ξ_p	0.66	0.72	0.65	0.84	0.89	0.75
	(0.53, 0.79)	(0.59, 0.85)	(0.53, 0.78)	(0.80, 0.89)	(0.89, 0.97)	(0.72, 0.77)
Panel B. Slopes of the price Phillips curves						
marginal cost	0.023	0.013	0.026	0.006	0.003	0.017
	(0.007, 0.057)	(0.004, 0.037)	(0.009, 0.059)	(0.003, 0.010)	(0.000, 0.003)	(0.014, 0.023)
output gap	0.016	0.010	0.028	0.006	0.001	0.005
	(0.004, 0.027)	(0.003, 0.021)	(0.009, 0.049)	(0.003, 0.009)	(0.000, 0.003)	(0.004, 0.007)
labor gap	0.027	0.020	0.039	0.006	0.002	0.005
	(0.007, 0.044)	(0.005, 0.037)	(0.013, 0.066)	(0.004, 0.010)	(0.000, 0.003)	(0.003, 0.006)
labor share	0.047	0.044	0.118	0.009	0.005	0.009
	(-0.140, 0.197)	(-0.084, 0.143)	(-0.481, 0.406)	(0.005, 0.015)	(0.000, 0.006)	(0.007, 0.012)
labor share,	0.023	0.016	0.047	0.006	0.003	0.006
υ adjusted	(0.007, 0.057)	(0.005, 0.041)	(0.010, 0.197)	(0.003, 0.010)	(0.000, 0.003)	(0.004, 0.009)

Notes: Columns (1)-(3) and (4)-(6) illustrate the results from the Cobb–Douglas, nested CES, and translog models estimated using the early and late samples, respectively. Panel A shows the Calvo price stickiness parameter in the posterior mode and the equal-tailed 95% credible intervals. Panel B focuses on the slopes of the price Phillips curves with different forcing variables, computed using Equation (3.9).

Panel A in Table 4 shows the Calvo price stickiness parameter, ξ_p. For the early sample, the three models feature ξ_p values of approximately two-thirds. As shown in Columns (4) and (5), the estimated Cobb–Douglas and nested CES models rely on large ξ_p values for the late sample, corresponding to excessively sticky prices. For example, the average duration of prices implied by ξ_p in Column (5) is longer than nine quarters. In contrast, the translog model uses reasonable values of ξ_p to match the data. In the posterior mode, ξ_p is 0.75 in Column (6), which is consistent with the empirical evidence on the degree of price stickiness in Bils and Klenow (2004) and Nakamura and Steinsson (2008).

ξ_p under the translog model modestly increased in the posterior mode from 0.65 to 0.75 between the two sample periods. As a result, κ in Column (6), 0.017, is slightly smaller than that in Column (3), 0.026. This decline is statistically insignificant given that 0.017 is included in the 95% credible interval in Column (3), (0.009, 0.059). Furthermore, 0.017 still corresponds to a significantly positive slope of the marginal cost Phillips curve. We obtain this result despite highly stable inflation during the late sample, exemplified by the missing disinflation and reinflation episodes (Coibion and Gorodnichenko, 2015; Ball and Mazumder, 2020). Finally, 0.017 is largely consistent with the empirical results in Gagliardone et al. (2025) when estimation uncertainty is taken into consideration, despite the different methods,

countries, and covered industries in the data.[18]

The increase in ξ_p and the resulting decrease in κ may reflect other potential changes affecting the slopes of the Phillips curves in a reduced-form manner.[19] For example, building on Basu (1995), Rubbo (2023, section 6) shows that changes in the input-output structure have decreased the slopes of the Phillips curves by a modest amount (about 30%) via macroeconomic complementarities. This magnitude is largely comparable to our results for the slopes of the marginal cost Phillips curve, κ, but is insufficient to rationalize the flattening of the conventional Phillips curves.

In short, through the lens of the translog model, which matches the data better than the other two models do, the New Keynesian price Phillips curve is not flat and is alive well. Furthermore, despite this result for the marginal cost Phillips curve, the translog model can generate substantially smaller κ_x values for the conventional activity measures in Column (6) than κ_x in Column (3), i.e., the flattening of the conventional Phillips curves. Thus, we conclude that the translog model with structural changes, captured by the differences in translog parameters between the two sample periods, can rationalize the two seemingly contradictory empirical results of a steep marginal cost Phillips curve (Gagliardone et al., 2025) and flattened conventional Phillips curves (Del Negro et al., 2020; Stock and Watson, 2020; Hazell et al., 2022; Inoue et al., 2024; Smith et al., 2025).

The results differ for the other production functions. The Cobb–Douglas and nested CES models predict quite flat marginal cost Phillips curves for the late sample, given the large values of ξ_p. Those quite flat marginal cost Phillips curves translate into similarly flat conventional Phillips curves, as shown in Columns (4) and (5). Thus, in this case, the flattening of the conventional Phillips curves mirrors the similar flattening of the marginal cost Phillips curve, contradicting the empirical evidence in Gagliardone et al. (2025). We obtain similar results when we simulate the translog model and estimate the Cobb–Douglas

[18] The two-standard-error bands reported in Gagliardone et al. (2025, table 2, panel c) include 0.017 for three out of four regression specifications (their unrestricted model and the two models controlling for unobserved confounding factors using industry-by-time fixed effects at the granular industry level).

[19] These changes may include (i) sectoral transformation toward services (Galesi and Rachedi, 2019; Cotton and Garga, 2022) and the fact that service prices are stickier than manufacturing prices (Bils and Klenow, 2004; Nakamura and Steinsson, 2008; Imbs et al., 2011; Carvalho et al., 2021), (ii) trade and globalization (Sbordone, 2007), (iii) market concentration (Wang and Werning, 2022; Fujiwara and Matsuyama, 2023; Baqaee et al., 2024), (iv) the production network and an increase in the intermediate input share (Basu, 1995; Rubbo, 2023), (v) occupational composition in the labor market and an increase in the share of non-routine jobs (Siena and Zago, 2021), and (vi) the growing competition among online retailors (Cavallo, 2018).

and nested CES models using the simulated data (see Appendix C.4).

3.4 Marginal Costs and Cyclical Returns to Scale

Next, we discuss the mechanism involved. Specifically, we relate the non-flattening of the marginal cost Phillips curve and the flattening of the conventional Phillips curves to more procyclical returns to scale in recent decades.

Like Equation (2.4) of the illustrative model, the marginal costs consist of factor prices, TFP, and returns to scale owing to the translog structure:

$$\hat{mc}_t = \underbrace{\alpha_{ks}\hat{r}_{st} + \alpha_{ke}\hat{r}_{et} + \alpha_s\hat{w}_{st} + \alpha_u\hat{w}_{ut}}_{\text{factor prices}} - \underbrace{\varepsilon_t^a}_{\text{tfp}} - \underbrace{(\beta_k.\hat{k}_{et} + \beta_s.\hat{l}_{st} + \beta_u.\hat{l}_{ut})}_{\widehat{rts}_t^{tl}}. \tag{3.10}$$

Thus, cyclical variation in returns to scale weakens the connection between marginal costs and their conventional components, such as factor prices and TFP. Consider an expansionary monetary policy shock increasing the output gap. Because the wage Phillips curves are steep in the model (see the estimates of ζ_s and ζ_u in Tables C.1-C.2 in Appendix C.3) and in the data (Galí and Gambetti, 2020; Heise et al., 2022; Bernanke and Blanchard, 2023), factor prices respond procyclically to the shock, paralleling the response of the output gap. However, the response of the marginal costs could be mitigated given the procyclical response of rts_t^{tl}. In this case, inflation fluctuates less than the output gap (i.e., the output gap Phillips curve is quite flat), not because the pass-through of the marginal costs on inflation is weak (i.e., the marginal cost Phillips curve is steep) but because the marginal costs are relatively stable owing to the offsetting of procyclical factor prices and returns to scale conditional on monetary policy shocks. This prediction of the translog model is consistent with the missing pass-through of labor costs to price inflation in the recent US data (Peneva and Rudd, 2017).

To illustrate this mechanism, we plot the impulse responses of the relevant variables to one-standard-deviation expansionary monetary policy shocks in Figure 2. We first consider models without cyclical variation in $\widehat{rts}_t^{tl}$. Panel (b) shows that the marginal costs increase similarly under the nested CES model across the sample period. However, because of substantially stickier prices in the late sample, the inflation responses are much smaller for the late sample than for the early sample in Panel (a). Similarly, for the Cobb–Douglas model, the pass-through of the marginal costs on inflation has been weak in recent decades (not shown), reflecting the estimated flattening of the marginal cost Phillips curve.

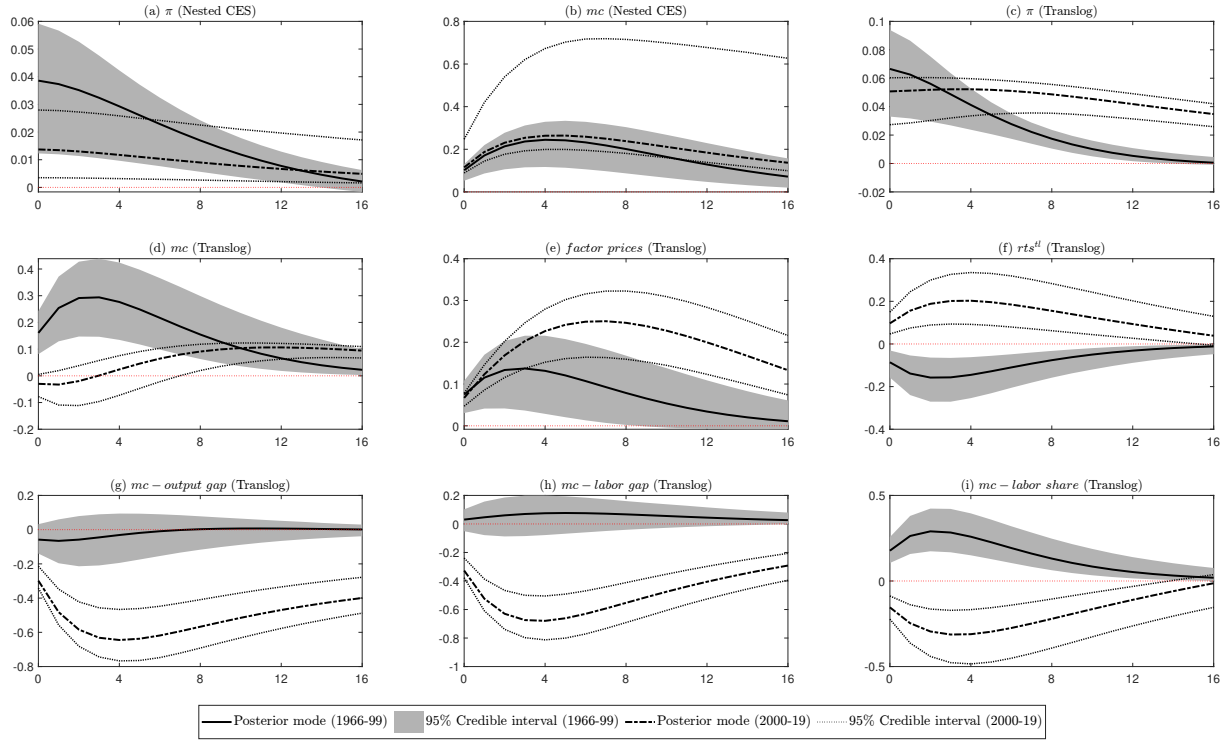


Figure 2: Impulse responses to expansionary monetary policy shocks

Notes: This figure plots the impulse responses to one-standard-deviation expansionary monetary policy shocks. We show the model-predicted impulse responses in the posterior mode with the equal-tailed 95% credible intervals for each sample period. Panels (a) and (b) are based on the nested CES model. The other panels regard the translog model.

The results differ in the translog model with a cyclical rts_t^{tl}. When estimated using the late sample, monetary policy shocks generate persistent inflation responses with a smaller impact response than those in the early sample (Panel (c)). These distinct inflation dynamics reflect the dissimilar responses of marginal costs between the two sample periods (Panel (d)). In contrast to the hump-shaped responses of marginal costs before 2000, the short-term responses within a year are muted after 2000. These close-to-zero responses lead to small changes in the present discounted values of marginal costs in the short term and, thus, small variations in inflation responses.[20]

These distinct dynamics of marginal costs mostly arise from more procyclical rts_t^{tl} after 2000 than before (see Hyun et al., 2024; **?**, for empirical evidence of time-varying returns to scale). Panels (e) and (f) decompose the response of marginal costs into that of factor prices

[20]Equation (3.8) implies that $\hat{\pi}_t = \kappa \sum_{\tau=0}^{\infty} (\tilde{\beta}\gamma)^\tau \mathbb{E}_t[\hat{mc}_{t+\tau}] + \sum_{\tau=0}^{\infty} (\tilde{\beta}\gamma)^\tau \mathbb{E}_t[\varepsilon^p_{t+\tau}]$. Thus, the responses of $\hat{\pi}_{t+h}$ to a monetary policy shock at time t for $h \geq 0$ depend on the discounted values of the expected responses of marginal costs, $\kappa \sum_{\tau=0}^{\infty} (\tilde{\beta}\gamma)^\tau \mathbb{E}_t[\hat{mc}_{t+h+\tau}]$.

and rts_t^{tl} considering Equation (3.10). Note that procyclical rts_t^{tl} after 2000 partially cancel out procyclical factor prices, yielding smaller responses in marginal costs.

Finally, Panels (g)-(i) show the responses of the gap between marginal cost and other measures of economic activity, constituting an omitted variable in conventional Phillips curve regressions, as illustrated in Proposition 1. More countercyclical responses in recent periods generate more negative omitted variable biases in the identified slope, κ_x. The flattening of the conventional Phillips curves follows from this result.[21]

3.5 Determinants of Inflation

The above results imply that the inflation dynamics and their determinants could vary across the models and the sample periods. From the results of the variance decomposition analysis, we find that the translog model in recent periods relies more on risk premium shocks, probably capturing the Great Recession, and less on price markup shocks than the other models do. Furthermore, inflation expectation is a more important driver of inflation in the post-2000 period than before.

Table 5 shows the variance of inflation at the business cycle frequencies (Panel A) and its decomposition into the contribution of the nine structural shocks (Panel B) in the posterior mode of each model. Given more stable inflation data in recent periods, the estimated models using the post-2000 data predict smaller variances than those using the early sample. Furthermore, all three models similarly predict that price and wage markup shocks explain the majority of the cyclical variations in inflation in Columns (1)-(3). However, for the late sample, different models attribute inflation volatility to different shocks. Given a quite flat Phillips curve, the nested CES model relies heavily on price markup shocks. The Cobb–Douglas model adds investment-specific shocks to the set of meaningful drivers of inflation, yielding a larger cyclical variance of inflation than the nested CES model does. In contrast to the results in Columns (1)-(5), the translog model captures aggregate fluctuations due to risk premium shocks and their contribution to inflation. Finally, wage markup shocks have minimal contributions to overall inflation volatility in recent periods in all three models.

[21] The conventional measures of economic activity mostly reflect factor prices in the marginal costs. Consider the posterior mode of the translog model in recent periods. Conditional on monetary policy shocks, the cyclical correlations between factor prices ($\alpha_{ks}\hat{r}_{st} + \alpha_{ke}\hat{r}_{et} + \alpha_s\hat{w}_{st} + \alpha_u\hat{w}_{ut}$) and the output gap, labor gap, labor shares, and fixed cost adjusted labor shares are 0.95, 0.94, 0.98, and 0.97, respectively. In contrast, $\text{corr}_{c,mp}(\hat{mc}_t, \text{factor prices})$ is only 0.31, reflecting the cyclical returns to scale in the marginal costs (see Equation (3.10)).

Table 5: Variance decomposition of inflation

	(1)	(2)	(3)	(4)	(5)	(6)
	Early sample (1966-99)			Late sample (2000-19)		
	C–D	NCES	Translog	C–D	NCES	Translog
Panel A. Variance of inflation at the cyclical frequencies						
$\text{var}_c(\hat{\pi}_t)$	0.086	0.076	0.084	0.050	0.018	0.057
Panel B. Contribution of each structural shock, %						
Productivity (Hicks neutral)	3	2	3	3	1	1
Risk premium	1	0	1	1	6	40
Government spending	1	0	1	2	0	0
Investment-specific productivity	3	1	7	41	0	12
Monetary policy	5	4	9	9	2	8
Price markup	55	69	50	38	91	34
Wage markup, skilled	1	4	1	2	0	1
Wage markup, unskilled	32	20	27	2	1	2
Skilled worker population share	0	0	0	0	0	1

Notes: This table shows the variance of inflation at the business cycle frequencies corresponding to the periodicity of 6-32 quarters (Panel A) and its decomposition into the contribution of the nine structural shocks (Panel B) in the posterior mode.

Table 6: Variance decomposition of inflation using the Phillips curve

	(1)	(2)	(3)	(4)	(5)	(6)
	Early sample (1966-99)			Late sample (2000-19)		
	$\mathbb{E}_t[\hat{\pi}_{t+1}]$	$\hat{mc}_t$	ε_t^p	$\mathbb{E}_t[\hat{\pi}_{t+1}]$	$\hat{mc}_t$	ε_t^p
Panel A. Decomposition of $var_c(\hat{\pi}_t)$ *using the Phliips curve:* $\hat{\pi}_t = \tilde{\beta}\gamma\mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa\hat{mc}_t + \varepsilon_t^p$						
$\text{cov}_c(\hat{\pi}_t, \tilde{\beta}\gamma\mathbb{E}_t[\hat{\pi}_{t+1}]$ or $\kappa\hat{mc}_t$ or $\varepsilon_t^p)$	0.062	0.004	0.018	0.048	0.002	0.008
slope ($\tilde{\beta}\gamma$ or κ or 1)	0.997	0.026	1	0.999	0.017	1
$\text{corr}_c(\hat{\pi}_t, \mathbb{E}_t[\hat{\pi}_{t+1}]$ or $\hat{mc}_t$ or $\varepsilon_t^p)$	0.978	0.409	0.684	0.969	0.301	0.530
$\text{std}_c(\hat{\pi}_t)$	0.290	0.290	0.290	0.239	0.239	0.239
$\text{std}_c(\mathbb{E}_t[\hat{\pi}_{t+1}]$ or $\hat{mc}_t$ or $\varepsilon_t^p)$	0.219	1.428	0.089	0.206	1.391	0.061
Panel B. Contribution of each structural shock, %						
Productivity (Hicks neutral)	3	17	-	1	4	-
Risk premium	1	7	-	49	-2	-
Government spending	1	6	-	0	0	-
Investment-specific productivity	9	22	-	14	30	-
Monetary policy	11	20	-	9	-3	-
Price markup	40	-18	100	23	37	100
Wage markup, skilled	1	0	-	1	17	-
Wage markup, unskilled	34	44	-	2	15	-
Skilled worker population share	0	0	-	1	2	-

Notes: Panel A decomposes $\text{var}_c(\hat{\pi}_t)$ through the lens of the Phillips curve. Equation (3.8) implies that $\text{var}_c(\hat{\pi}_t) = \text{cov}_c(\hat{\pi}_t, \tilde{\beta}\gamma\mathbb{E}_t[\hat{\pi}_{t+1}]) + \text{cov}_c(\hat{\pi}_t, \kappa\hat{mc}_t) + \text{cov}_c(\hat{\pi}_t, \varepsilon_t^p)$. These three terms are shown in the first row based on the translog models in the posterior mode for the pre- and post-2000 samples. The remaining part of Panel A decomposes each covariance into the slope (e.g., κ), the correlation between $\hat{\pi}_t$ and, e.g., $\hat{mc}_t$, and standard deviations of these two variables. Panel B presents conditional covariances on each structural shock relative to unconditional covariances. For example, according to column (5), 4% of $\text{cov}_c(\hat{\pi}_t, \kappa\hat{mc}_t)$ is because of the covariation driven by TFP shocks.

Table 6 decomposes $\text{var}_c(\hat{\pi}_t)$ through the lens of the Phillips curve. Equation (3.8) implies that $\text{var}_c(\hat{\pi}_t) = \text{cov}_c(\hat{\pi}_t, \tilde{\beta}\gamma\mathbb{E}_t[\hat{\pi}_{t+1}]) + \text{cov}_c(\hat{\pi}_t, \kappa\hat{mc}_t) + \text{cov}_c(\hat{\pi}_t, \varepsilon_t^p)$. These three terms equal

0.062, 0.004, and 0.018 (the first three numbers in the first row in Panel A) in the posterior mode of the translog model for the pre-2000 sample. Similarly, for the translog model using the post-2000 data, $\text{var}_c(\hat{\pi}_t)$ is decomposed into 0.048 (inflation expectation), 0.002 (marginal cost), and 0.008 (price markup shock). These numbers imply that the contribution of inflation expectation to inflation volatility, $\frac{\text{cov}_c(\hat{\pi}_t, \tilde{\beta}\gamma \mathbb{E}_t[\hat{\pi}_{t+1}])}{\text{var}_c(\hat{\pi}_t)}$, increased from 74% to 84% between the two sample periods. This increased importance of inflation expectation or, relatedly, a decrease in $\text{cov}_c(\hat{\pi}_t, \kappa \hat{mc}_t)$ can be decomposed into changes in the slope (κ), the correlation between $\hat{\pi}_t$ and $\hat{mc}_t$, and the standard deviations of these two variables, as shown in Panel A. Comparing Columns (2) and (5) suggests that the decrease in $\text{cov}_c(\hat{\pi}_t, \kappa \hat{mc}_t)$ is mostly due to the moderate flattening of the marginal cost Phillips curve (κ) and a decrease in the correlation between inflation and the marginal cost.

The decrease in this correlation results from procyclical returns to scale that dampen the short-term responses of marginal costs conditional on demand shocks. Consider the impulse responses of inflation and marginal costs to monetary policy shocks (Figure 2(c)-(d)). The impact response of marginal costs is weak. In contrast, inflation responds significantly because it reflects the present discounted value of the current and future marginal costs. Given similar mechanisms for other demand shocks, the contemporaneous correlation between the marginal costs and inflation decreases after 2000 compared with before 2000.

Panel B illustrates conditional covariances on each structural shock relative to unconditional covariances. For example, according to Column (5), 4% of $\text{cov}_c(\hat{\pi}_t, \kappa \hat{mc}_t)$ is because of the comovement driven by TFP shocks. Comparing Columns (2) and (5), we find that the contribution of demand shocks (risk premium, government spending, and monetary policy) to $\text{cov}_c(\hat{\pi}_t, \kappa \hat{mc}_t)$ indeed decreased between the two periods. The contemporaneous covariance conditional on monetary policy shocks is negative in Column (5), reflecting a negative short-run response of marginal costs (or procyclical price markups) in Figure 2(d) for the recent period. A similar result holds for risk premium shocks (not shown) as a result of procyclical returns to scale conditional on demand shocks.

Another significant difference between Columns (2) and (5) is the correlation between inflation and the marginal costs conditional on price markup shocks. Before 2000, a positive price markup shock increased inflation and decreased the marginal cost, generating a negative correlation between them. In contrast, after 2000, price makeup shocks led to positively correlated variations between inflation and the marginal cost. This change is again caused by more procyclical returns to scale. When a price markup shock induces high inflation and

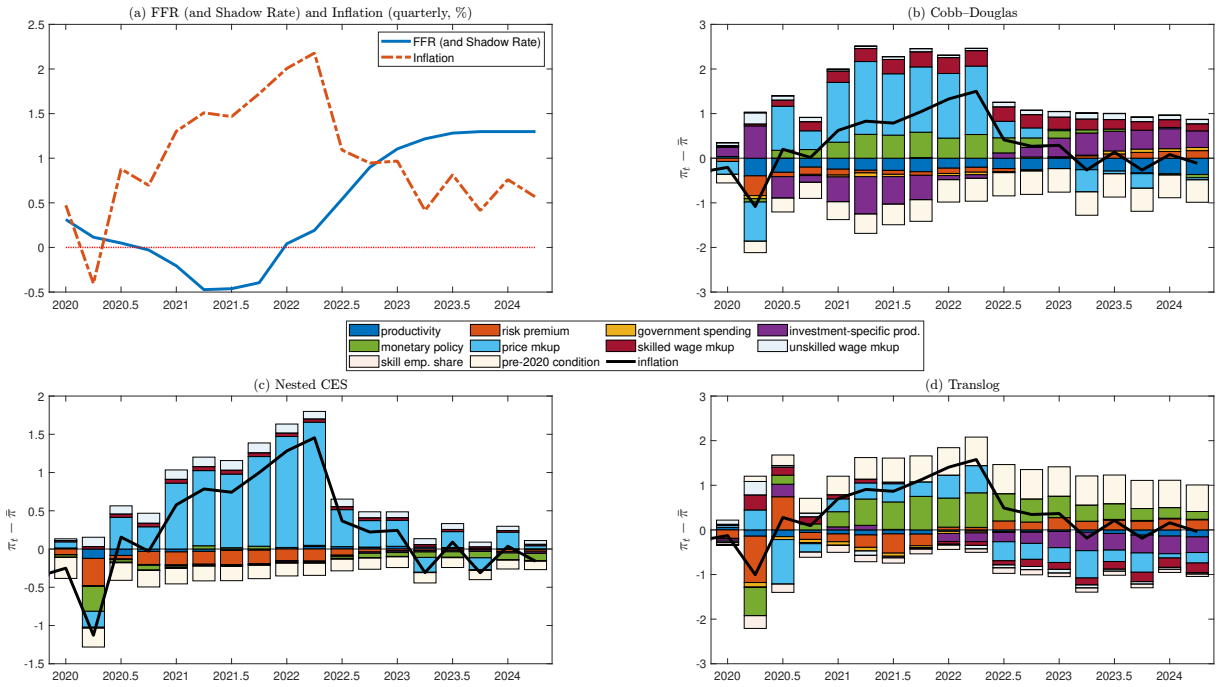


Figure 3: Historical decomposition of the pandemic-era inflation

Notes: Panel (a) presents the federal funds rate (augmented with the shadow rate of Wu and Xia (2016) when the zero lower bound was binding) and quarterly GDP deflator inflation. Panels (b)-(d) show the historical decomposition of the realized quarterly inflation into the contribution of each structural shock and the pre-2020 conditions for the Cobb–Douglas, nested CES, and translog models, respectively.

low output, decreased factor inputs endogenously decrease returns to scale, pushing marginal costs upward.

3.6 Pandemic-Era Inflation through the Lens of the Model

Given the different inflation dynamics under the translog model compared with those under the Cobb–Douglas and nested CES models, this section investigates recent inflation data through the lens of all three models. Specifically, we augment the sample with the eight quarterly variables used for the Bayesian estimation in Section 3.2 from 2020:q1 to 2024:q2.[22] For each model, given the parameter values in the posterior mode in Section 3.2, we use the Kalman smoother to estimate the realized structural shocks during this extended sample period from 2020:q1-2024:q2.

Figure 3(b)-(d) shows the historical decomposition of the realized quarterly inflation into the contributions of each structural shock and the pre-2020 economic conditions for

[22]We use the skilled employment share until 2023:q4. Its values in 2024:q1-q2 are treated as missing.

the Cobb–Douglas, nested CES, and translog models, respectively. Panel (a) presents the federal funds rate (augmented with the shadow rate of Wu and Xia (2016) when the zero lower bound was binding) and quarterly GDP deflator inflation. Clearly, US monetary policy was loose between 2021:q1 and 2022:q2 when inflation was increasing.

Panel (d) shows that the translog model emphasizes the contribution of risk premium shocks at the onset of the pandemic in 2020:q2, reflecting elevated demand for safe assets (Fisher, 2015), and loose monetary policies contributing to high inflation during 2021:q1-2022:q2 in addition to price markup shocks, probably reflecting global supply chain disruptions (Di Giovanni et al., 2022; Bai et al., 2024a), the global energy crisis (Bernanke and Blanchard, 2023), and a (unmodeled) nonlinear Phillips curve (see, e.g., Harding et al., 2023; Benigno and Eggertsson, 2023). Because this model features less sticky prices than the other two models do, demand shifters such as monetary policy and risk premium shocks appear to be meaningful sources of recent inflation fluctuations. Note that this result holds despite quite flat conventional Phillips curves in recent periods.

In contrast, the Cobb–Douglas (Panel (b)) and nested CES (Panel (c)) models rely heavily on price markup shocks to match the realized inflation dynamics after 2020. Because the Phillips curves are quite flat in these models, demand disturbances contribute little to inflation, as shown in Table 5.[23] Instead, these models shift the Phillips curve using supply shocks to match the inflation fluctuations.

In short, the translog model provides a more comprehensive interpretation of pandemic-era inflation than the Cobb–Douglas and nested CES models do. The historical decomposition in Panel (d) reflects several aggregate developments that occurred during the period instead of relying solely on supply-side disturbances. Furthermore, this result emphasizes that monetary policies have been highly relevant to inflation in recent years, which is consistent with the results based on different methods (see Comin et al., 2023; Gagliardone and Gertler, 2023; Bocola et al., 2024; Giannone and Primiceri, 2024). In contrast, the Cobb–Douglas and nested CES models do not predict this policy implication.

Our historical decomposition analysis using a structural model is complementary to other explanations for recent inflation. Given the Ricardian nature of the model, fiscal shocks had limited contributions to inflation, as shown in Figure 3. For models emphasizing the role of massive fiscal stimulus packages and increases in unfunded government debt, see Chen et al. (2022); Bianchi et al. (2023); Barro and Bianchi (2023); Di Giovanni et al. (2023).

[23] See L'Huillier et al. (2022); Beaudry et al. (2024b) for other policy implications of quite flat Phillips curves.

Crump et al. (2024) focuses on labor market conditions and the measurement of the natural rate of unemployment, which is absent in our model. Note that our analysis assumes full-information, linear rational expectations models. Beaudry et al. (2024a) deviates from this assumption and introduces limited information and bounded rational beliefs that affect inflation through inflation expectations. Finally, several authors have questioned linearity and proposed convex price and wage Phillips curves (see Daly and Hobijn, 2014; Ball et al., 2022; Boehm and Pandalai-Nayar, 2022; Harding et al., 2022, 2023; Benigno and Eggertsson, 2023; Schmitt-Grohé and Uribe, 2023; Blanco et al., 2024a,b). In contrast, Beaudry et al. (2025) shows that the evidence in support of a convex Phillips curve is not robust. Additionally, Kocherlakota (2024) argues that the substitution effect among intermediate goods yields a concave Phillips curve. Interestingly, a nonlinear version of our model can generate convex Phillips curves even in an environment similar to that of Kocherlakota (2024) if input complementarity is sufficiently strong (see Appendix A and Gagliardone et al. 2025, for a related discussion on the basis of macroeconomic complementarities).

Robustness checks. Our quantitative model in this section features six translog parameters. Although the data and estimation methods yield reasonably precise estimates of these parameters, we further check the robustness of the results to two simpler models with fewer additional degrees of freedom. Specifically, we consider (i) a four-factor model where a single parameter disciplines the cyclicality of returns to scale and their co-movement with the output gap, and (ii) a two-factor model à la Smets and Wouters (2007) with three translog parameters. The cyclicality of the returns to scale and its implications for the slopes of the Phillips curves largely remain unchanged. However, the MDD significantly prefers the translog model over the simpler output gap-based model (see Appendices C.6 and C.7).

4 Empirical Analyses

This section estimates production function parameters using industry-level panel data and IV methods. We show that the estimates are consistent with the results based on the structural estimation of the translog DSGE models in Section 3.2, corroborating stronger input complementarity at the aggregate level in recent decades than in earlier decades.

4.1 Data

We utilize data from the EU KLEMS Growth and Productivity Accounts, which provide detailed, harmonized, industry-level information across a wide range of European countries. We focus on EU KLEMS data for five main reasons. First, the macroeconomic dynamics motivating our analysis are not unique to the US. Smith et al. (2025, tables 1 and 5) documented that the conventional Phillips curves in the Euro area flattened around 2000, similarly to the US case (see also Musso et al., 2009; Oinonen et al., 2013; Oinonen and Vilmi, 2021). Moreover, the recent disconnect between wage growth and price inflation is a shared phenomenon in both regions (Peneva and Rudd, 2017; Eser et al., 2020). Second, disaggregated labor data by industry and educational attainment are essential for our analysis of input complementarity. The EU KLEMS dataset provides this information, which is collected consistently over a long sample period. Third, comparable datasets for the US economy, e.g., the US data in EU KLEMS and NBER-CES (Bartelsman and Gray, 1996), do not include the disaggregated labor input measures. This feature makes the European data particularly well-suited for the empirical analysis in this section. Fourth, the EU KLEMS dataset covers a broad set of countries and, in particular, industries. Such wide coverage and rich variation can be leveraged to represent production technology at the aggregate level, enabling a straightforward comparison with the quantitative analyses in Section 3. Finally, the EU KLEMS data are publicly available, widely used, and well-documented, ensuring transparency and replicability (O'Mahony and Timmer, 2009; Bontadini et al., 2023). In short, the EU KLEMS dataset provides a valuable source for empirical validation of our hypothesis, complementary to the structural analysis in Section 3.

We construct a multicountry, industry-level panel dataset by integrating the EU KLEMS national, labor, and capital accounts. From the national accounts, we extract data on the volumes and prices of output, intermediate inputs, total labor hours, and compensation across industries, countries, and years. The labor accounts decompose labor inputs by education, allowing us to calculate hours worked by college graduates (skilled labor) and others (unskilled labor), as well as wages for each group. We classify college graduates as skilled workers, which is consistent with the DSGE estimation in Section 3.2.1 (see also Krusell et al., 2000; Ohanian et al., 2023). Finally, we obtain volumes and prices for the real capital stock from the capital accounts.

We rely on three releases of the EU KLEMS dataset: March 2008, October 2012, and

February 2023. We combine these three releases to achieve the most comprehensive coverage of labor data, particularly labor usage disaggregated by educational attainment. We use the March 2008 release for the pre-2000 sample (1970-1999). By combining the 2012 and 2023 releases, we create a continuous dataset for the post-2000 sample (2002-2019). Our sample period for recent decades begins in 2002 because labor data disaggregated by educational attainment are limited for 2000 and 2001.[24] Appendix D.1 reports the summary statistics of the data (see also Bontadini et al., 2023).

4.2 Empirical Framework

Estimating a translog production function is challenging because of the large number of parameters involved. To address this issue, we use the FOC for the firm's cost-minimization problem following Gandhi et al. (2020), **?**, and Hyun et al. (2024). This condition establishes a relationship between the marginal product of an input (e.g., skilled labor) and its price (e.g., the skilled wage rate). Using this condition, we identify the shape of the production function associated with the input considered. We utilize the time series and cross-sectional variations in (lagged) input prices as IVs to further enhance the identification, following Doraszelski and Jaumandreu (2013, 2018) and Hyun et al. (2024).

To fix ideas, consider an economic environment similar to that in Section 3.1.1 with the following production function:

$$y = \exp\left(\varepsilon^a\right) f(\log \boldsymbol{x}) - \nu,$$

where y denotes the output, ε^a represents the level of productivity, and $\boldsymbol{x}$ is a vector of factor inputs. The fixed cost of production, ν, ensures the zero-profit condition on the BGP. The FOC for the cost-minimization problem with respect to an input x^i is given by:

$$\underbrace{\frac{P^i}{P}}_{\text{real input price}} = mc \times \underbrace{\frac{\tilde{y}}{x^i}\frac{\partial \log f}{\partial \log x^i}}_{\text{marginal product}},$$

where P^i represents the nominal price of x^i, $\tilde{y} = y + \nu$, and mc denotes the real marginal

[24] Our sample covers Austria, Belgium, Germany, Denmark, Spain, Finland, France, the United Kingdom, Italy, Japan, Luxembourg, the Netherlands, Portugal, and Sweden. Labor data for 2000-2001 are only available for Finland and Italy. From 2002 (2008) onward, the data are available for Austria, Germany, Spain, France, the UK, and the Netherlands (the other countries).

costs of production. Rearranging this equation yields:

$$s^i \equiv \frac{P^i x^i}{Py} = mc \frac{\tilde{y}}{y} \frac{\partial \log f}{\partial \log x^i}, \tag{4.1}$$

where s^i is the input expenditure share of total sales for input x^i.

We log-linearize this equation and relate the input i's expenditure share to factor inputs in period t as follows:

$$\hat{s}_t^i = \sum_k \delta_{ik} \hat{x}_t^k + \tau_t, \tag{4.2}$$

$$\text{where} \qquad \delta_{ik} = -\beta_{k\cdot} - (\Phi - 1)\alpha_k + \frac{\beta_{ik}}{\alpha_i}, \tag{4.3}$$

$\alpha_i = \frac{\partial \log f}{\partial \log x^i}$, $\beta_{ik} = \frac{\partial^2 \log f}{\partial \log x^k \partial \log x^i}$, $\beta_{k\cdot} = \sum_\ell \beta_{k\ell}$, and Φ is the gross markup along the BGP. τ_t is given by real factor prices and the level of productivity: $\tau_t = \sum_\ell \alpha_\ell \widehat{(P_t^\ell / P_t)} - \varepsilon_t^a$.

The key parameters in this section, δ_{ik}s, represent how x^k affects the input share, s^i. As shown in Equation (4.3), δ_{ik} consists of three terms. The first term, $-\beta_{k\cdot}$, emerges because x^k has effects on mc via rts^{tl} and, thus, s^i (see Equation (4.1)). When $\beta_{k\cdot} > 0$, x^k positively contributes to rts^{tl}, and thus, mc decreases (see also Equation (3.10)). The second component, $-(\Phi - 1)\alpha_k$, arises from the fixed cost in production, inducing countercyclical variations in $\frac{\tilde{y}}{y}$. Finally, $\frac{\beta_{ik}}{\alpha_i}$ captures how the output elasticity (ignoring the fixed costs), $\frac{\partial \log f}{\partial \log x^i}$, varies with x^k. Note that in a textbook New Keynesian model with a Cobb–Douglas production function ($\beta_{ik} = 0$ for all i and k and $\nu = 0$) (see, e.g., Galí, 2015), the input share, s^i, simplifies to $mc \times \alpha_i$. In that case, all input shares exhibit the same cyclical variations because $\hat{s}_t^i = \hat{mc}_t = \tau_t$ for all i, and δ_{ik} equals zero for all i and k.

The empirical specification for our regression analysis builds on the theoretical prediction in Equation (4.2). Specifically, we estimate the following equation:

$$\log\left(s_{cjt}^i\right) = \sum_k \delta_{ik} \log\left(x_{cjt}^k\right) + \alpha_c + \gamma_j + \tau_t + \varepsilon_{cjt}^i, \tag{4.4}$$

where c, j, and t denote country, industry, and year, respectively. We include country, industry, and time fixed effects (α_c, γ_j, and τ_t) to control for unobserved heterogeneity at these levels. Additionally, Equation (4.4) does not involve real output, y_{cjt}^i. The dependent variable, s_{cjt}^i, can be computed using the two nominal values, $P_{cjt}^i x_{cjt}^i$ and $P_{cjt} y_{cjt}$. This

property allows us to circumvent the construction of the exact output price index (see, e.g., Hottman et al., 2016). Finally, we assume that δs do not depend on industry j. That is, we consider a representative production function to inform the aggregate production function assumed in the DSGE model in Section 3.

Equation (4.4) is based on the relationship between the real input price and the marginal product of the input. A wedge between these two quantities may arise from frictions in the economy, such as input adjustment costs (Hall, 2004), imperfect competition in output and factor markets (Rotemberg and Woodford, 1999; Berger et al., 2022), and financial frictions (Jermann and Quadrini, 2012; Arellano et al., 2019; Bigio and La'O, 2020). After controlling for county, industry, and time fixed effects, the remaining idiosyncratic component of this wedge could appear in the residual, ε^i_{cjt}. This residual term may also capture measurement errors and ex post productivity shocks (Gandhi et al., 2020).

To enhance the identification of δs, following Doraszelski and Jaumandreu (2013, 2018) and Hyun et al. (2024), we utilize lagged input prices as IVs. Our identification assumption is that after controlling for fixed effects, the idiosyncratic components of the lagged input prices are orthogonal to the idiosyncratic components of the wedge in the residual. For instance, if the residual is serially uncorrelated, the lagged input prices satisfy the exogeneity condition for the IVs. The relevance condition holds when x^i is correlated with the input price and the input price is serially correlated. In addition to this benchmark specification, we perform several robustness checks considering potential sources of persistence in the residual, such as adjustment costs and persistent country-industry-specific components in productivity. For the IV relevance, we will present the first-stage F-statistics as a measure of correlations between x^i and lagged input prices.

4.3 Estimation Results

We focus on the cases where the skilled labor share is the dependent variable in the regression equation (4.4). We show that the values of δ vary between the two sample periods, consistent with the structural estimation results in Section 3.2.

We use the skilled labor share in this section for three reasons. First, the translog model estimation results in Table 3 indicate that β_{ks} and β_{su} increased significantly between the pre- and post-2000 samples. These increases in βs suggest that δ_{ks} and δ_{su} also changed accordingly between the two sample periods. Note that this prediction is testable using the

IV estimation results of Equation (4.4) when the skilled labor share is placed on the left-hand side. Second, δ_{si} includes the $\frac{\beta_{si}}{\alpha_s}$ term in Equation (4.3). Since α_s is relatively small (see Table 3), even a moderate change in β_{si} could result in a sizable shift in the DSGE model-implied δ_{si} for $i = k, u$. In contrast, δ_{ui} is less sensitive to changes in β_{ui} because $\alpha_s < \alpha_u$ (see Table 3). Thus, detecting a structural change in the degree of input complementarity (represented by different values of βs) from the estimates of δs in the empirical model is more viable when the skilled labor share is considered than the unskilled labor share. Finally, decomposing capital into structures and equipment, as Krusell et al. (2000) and Ohanian et al. (2023) did for the US aggregate data, is challenging for the EU KLEMS data because of the lack of necessary information. Thus, we do not pursue empirical specifications using the equipment (or capital) expenditure shares for the comparability of the results.

To investigate potential changes in the production function over time, we split the sample into two periods: 1970-1999 and 2002-2019. Columns (1) and (2) in Table 7 present the IV estimation results for the pre- and post-2000 samples, respectively. The two-step efficient generalized method of moments (GMM) estimator is employed, with standard errors clustered at the country level. One- and two-year lagged input prices are utilized as IVs. In Columns (3) and (4), we calculate the values of δ in the posterior mode and the posterior standard deviation of the translog model. For this purpose, we use Equation (4.3) and the Bayesian estimation results discussed in Section 3.2.

By comparing the results in Columns (1) and (2), we can evaluate how the production function has changed over time. First, the sign of the coefficient on capital shifts from negative to positive, indicating that capital has become more complementary to skilled labor over time. Similarly, the coefficient on unskilled labor is negative in Column (1) but becomes positive and insignificant in Column (2). This result suggests that skilled and unskilled labor were more substitutable in the earlier period but became less substitutable over time. In contrast, the coefficient on skilled labor decreases between the two sample periods. Overall, all these changes are consistent with the results from the structural estimations shown in Columns (3) and (4); δ_{ks} (equipment-skilled labor) and δ_{su} (skilled labor-unskilled labor) increase, whereas δ_{ss} decreases between the two sample periods in the posterior mode of the translog model. Despite the different methods, types of variations, and countries covered in the data, the results from the IV and structural estimations are largely congruent with one another.

In summary, the IV analysis in this section suggests that the production function has

Table 7: IV and structural estimation of δ

	IV Estimation		Structural Estimation	
	pre-2000	post-2000	pre-2000	post-2000
	(1)	(2)	(3)	(4)
Capital	-0.548	0.765		
	(0.091)	(0.321)		
Structure			-0.060	-0.040
			(0.008)	(0.007)
Equipment			-0.361	0.192
			(0.132)	(0.070)
Skilled Labor	0.429	-0.093	0.746	0.207
	(0.239)	(0.418)	(0.040)	(0.129)
Unskilled Labor	-0.721	0.477	-1.259	0.475
	(0.389)	(0.640)	(0.100)	(0.234)
Intermediate Goods	0.334	0.135		
	(0.322)	(0.120)		
F-stat for Capital	25.25	13.25		
F-stat for Skilled Labor	158.89	13.15		
F-stat for Unskilled Labor	61.33	23.46		
F-stat for Intermediate Goods	531.64	71.62		
Observations	2,372	2,566		

Notes: The skilled labor share is used as the dependent variable in regression equation (4.4). Columns (1) and (2) report the IV regression results for the period 1970-1999 and 2002-2019, respectively. The inputs are instrumented using lagged input prices at $t-1$ and $t-2$. The two-step efficient GMM specification is employed, with standard errors clustered at the country level, shown in parentheses. In Columns (3) and (4), we calculate the values of δ in the posterior mode and the posterior standard deviation of the translog model. For this purpose, we use Equation (4.3) and the Bayesian estimation results discussed in Section 3.2.

evolved over time, consistent with the DSGE analysis in Section 3. This empirical evidence aligns well with our explanation of the flattening of the Phillips curve, which is based on changes in the degree of input complementarity and the resulting cyclicality of returns to scale. Our results also emphasize the need for a more flexible approach beyond the standard Cobb–Douglas production function.

Robustness checks. The results in this section are robust to multiple alternative specifications and data adjustments as follows. First, we consider a value-added production function and use capital services (available only in the March 2008 release) instead of capital stock. Second, we broaden the post-2000 sample by including Finnish and Italian data from 2000 instead of 2002. Third, we compute Driscoll and Kraay (1998) standard errors. Fourth, we employ two- and three-year lagged input prices as IVs to lessen the potential influence of serial correlation in the residual of regression (4.4). Fifth, we adjust the skilled labor share by average input prices (weighted by expenditure shares) in view of $\sum_k \alpha_k \hat{p}_{kt}$ in τ_t (see Equation (4.3)). Lastly, we enhance the moment conditions to account for potentially

persistent, country-industry-specific components in the levels of productivity. In all these cases, the main results remain largely unchanged (see Appendix D.2).

5 Conclusion

This paper shows that more procyclical returns to scale in recent decades can reconcile the two seemingly divergent empirical findings: (i) the Phillips curves have flattened when conventional measures of economic activity, such as the output gap, labor gap, and labor shares, are employed (Del Negro et al., 2020; Stock and Watson, 2020; Inoue et al., 2024; Smith et al., 2025), and (ii) the Phillips curve is steep and alive when the marginal costs are used as a forcing variable (Gagliardone et al., 2025). We develop a theory based on a more flexible production function than the tightly parameterized Cobb–Douglas and CES functions. Our quantitative results emphasize how changes in the cyclicality of returns to scale help explain inflation dynamics and the Phillips curve. Furthermore, our model predicts the growing significance of inflation expectations, echoing the insights from Coibion and Gorodnichenko (2015), Hazell et al. (2022), and Meeks and Monti (2023).

This paper considers closed economy models and ignores the potential source of meaningful complementarity between imported intermediate goods and domestic factor inputs. The rise of globalization, the ongoing tariff war, and their macroeconomic implications through cyclical returns to scale under the translog framework are left for future research.

References

Alam, M. J., "Capital misallocation: Cyclicality and sources," *Journal of Economic Dynamics and Control* 112 (2020), 103831.

An, S. and F. Schorfheide, "Bayesian Analysis of DSGE Models," *Econometric reviews* 26 (2007), 113–172.

Antolin-Diaz, J. and P. Surico, "The Long-Run Effects of Government Spending," Center for Economic Policy Research Discussion Paper DP17433, 2024.

Anzoategui, D., D. Comin, M. Gertler and J. Martinez, "Endogenous technology adoption and R&D as sources of business cycle persistence," *American Economic Journal: Macroeconomics* 11 (2019), 67–110.

Arellano, C., Y. Bai and P. J. Kehoe, "Financial Frictions and Fluctuations in Volatility," *Journal of Political Economy* 127 (2019), 2049–2103.

Atkeson, A. and L. E. Ohanian, "Are Phillips curves useful for forecasting inflation?," *Federal Reserve Bank of Minneapolis Quarterly Review* 25 (2001), 2–11.

Autor, D. H., L. F. Katz and M. S. Kearney, "Trends in US Wage Inequality: Revising the Revisionists," *Review of Economics and Statistics* 90 (2008), 300–323.

Autor, D. H., F. Levy and R. J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics* 118 (2003), 1279–1333.

Bai, X., J. Fernández-Villaverde, Y. Li and F. Zanetti, "The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory," Working Paper 32098, National Bureau of Economic Research, 2024a.

Bai, Y., J.-V. Ríos-Rull and K. Storesletten, "Demand shocks as technology shocks," Working Paper 32169, National Bureau of Economic Research, 2024b.

Ball, L., D. Leigh and P. Mishra, "Understanding US inflation during the COVID-19 era," *Brookings Papers on Economic Activity* 2022 (2022), 1–80.

Ball, L. and S. Mazumder, "Inflation Dynamics and the Great Recession," *Brookings Papers on Economic Activity* 2011 (2011), 337–381.

———, "A Phillips Curve with Anchored Expectations and Short-Term Unemployment," *Journal of Money, Credit and Banking* 51 (2019), 111–137.

———, "The Nonpuzzling Behavior of Median Inflation," *Changing Inflation Dynamics, Evolving Monetary Policy* edited by Gonzalo Castex, Jordi Galí, and Diego Saravia. Santiago, Chile: Banco Central de Chile (2020), 49–70.

Baqaee, D. R., E. Farhi and K. Sangani, "The Supply-Side Effects of Monetary Policy," *Journal of Political Economy* 132 (2024), 1065–1112.

Barnichon, R. and G. Mesters, "Identifying modern macro equations with old shocks," *Quarterly Journal of Economics* 135 (2020), 2255–2298.

———, "The Phillips multiplier," *Journal of Monetary Economics* 117 (2021), 689–705.

Barro, R. J. and F. Bianchi, "Fiscal Influences on Inflation in OECD Countries, 2020-2022," Working Paper 31838, National Bureau of Economic Research, 2023.

Bartelsman, E. J. and W. Gray, "The NBER Manufacturing Productivity Database," Working Paper 205, National Bureau of Economic Research, October 1996.

Basu, S., "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare," *American Economic Review* 85 (1995), 512–31.

———, "Procyclical productivity: increasing returns or cyclical utilization?," *Quarterly Journal of Economics* 111 (1996), 719–751.

Basu, S. and J. G. Fernald, "Returns to scale in US production: Estimates and implications," *Journal of Political Economy* 105 (1997), 249–283.

———, "Aggregate productivity and aggregate technology," *European Economic Review* 46 (2002), 963–991.

Basu, S., J. G. Fernald and M. S. Kimball, "Are technology improvements contractionary?," *American Economic Review* 96 (2006), 1418–1448.

Baxter, M. and R. G. King, "Productive externalities and business cycles," *Institute for Empirical Macroeconomics, Federal Reserve Bank of Minneapolis, Discussion Paper 53* (1991).

———, "Fiscal policy in general equilibrium," *American Economic Review* 83 (1993), 315–334.

Beaudry, P., C. Hou and F. Portier, "The Dominant Role of Expectations and Broad-Based Supply Shocks in Driving Inflation," Working Paper 32322, National Bureau of Economic Research, 2024a.

———, "Monetary Policy When the Phillips Curve Is Quite Flat," *American Economic Journal: Macroeconomics* 16 (2024b), 1–28.

———, "On the Fragility of the Nonlinear Phillips Curve View of Recent Inflation," Working Paper 33522, National Bureau of Economic Research, 2025.

Beaudry, P. and F. Portier, "Understanding Noninflationary Demand-Driven Business Cycles," *NBER Macroeconomics Annual* 28 (2014), 69–130.

Benhabib, J. and R. E. A. Farmer, "Indeterminacy and increasing returns," *Journal of Economic Theory* 63 (1994), 19–41.

Benigno, P. and G. B. Eggertsson, "It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve," Working Paper 31197, National Bureau of Economic Research, 2023.

Berger, D., K. Herkenhoff and S. Mongey, "Labor market power," *American Economic Review* 112 (2022), 1147–1193.

Bergin, P. R. and R. C. Feenstra, "Staggered price setting, translog preferences, and endogenous persistence," *Journal of Monetary Economics* 45 (2000), 657–680.

Berlingieri, G., F. Boeri, D. Lashkari and J. Vogel, "Capital-skill complementarity in firms and in the aggregate economy," Working Paper 33000, National Bureau of Economic Research, 2024.

Bernanke, B. and O. Blanchard, "What caused the US pandemic-era inflation?," *Peterson Institute for International Economics Working Paper* (2023).

Bianchi, F., R. Faccini and L. Melosi, "A Fiscal Theory of Persistent Inflation," *Quarterly Journal of Economics* 138 (2023), 2127–2179.

Bianchi, F., H. Kung and G. Morales, "Growth, slowdowns, and recoveries," *Journal of Monetary Economics* 101 (2019), 47–63.

Bigio, S. and J. La'O, "Distortions in production networks," *Quarterly Journal of Economics* 135 (2020), 2187–2253.

Bilbiie, F. O., I. Fujiwara and F. Ghironi, "Optimal monetary policy with endogenous entry and product variety," *Journal of Monetary Economics* 64 (2014), 1–20.

Bilbiie, F. O., F. Ghironi and M. J. Melitz, "Endogenous Entry, Product Variety, and Business Cycles," *Journal of Political Economy* 120 (2012), 304–345.

Bils, M. and P. J. Klenow, "Some evidence on the importance of sticky prices," *Journal of Political Economy* 112 (2004), 947–985.

Blanco, A., C. Boar, C. J. Jones and V. Midrigan, "The Inflation Accelerator," Working Paper 32531, National Bureau of Economic Research, 2024a.

———, "Non-linear Inflation Dynamics in Menu Cost Economies," Working Paper 32094, National Bureau of Economic Research, 2024b.

Bocola, L., A. Dovis, K. Jørgensen and R. Kirpalani, "Bond Market Views of the Fed," Working Paper 32620, National Bureau of Economic Research, 2024.

Boehm, C. E. and N. Pandalai-Nayar, "Convex supply curves," *American Economic Review* 112 (2022), 3941–3969.

Bontadini, F., C. Corrado, J. Haskel, M. Iommi and C. Jona-Lasinio, "EUKLEMS & INTANProd: industry productivity accounts with intangibles," *Sources of growth and productivity trends: methods and main measurement challenges, Luiss Lab of European Economics, Rome* (2023).

Bound, J. and A. B. Krueger, "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?," *Journal of Labor Economics* 9 (1991), 1–24.

Calvo, G. A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics* 12 (1983), 383–398.

Canova, F. and L. Sala, "Back to square one: Identification issues in DSGE models," *Journal of Monetary Economics* 56 (2009), 431–449.

Cantore, C., M. León-Ledesma, P. McAdam and A. Willman, "Shocking Stuff: Technology, Hours, and Factor Substitution," *Journal of the European Economic Association* 12 (2014), 108–128.

Cantore, C. and P. Levine, "Getting normalization right: Dealing with 'dimensional constants' in macroeconomics," *Journal of Economic Dynamics and Control* 36 (2012), 1931–1949.

Cantore, C., P. Levine, J. Pearlman and B. Yang, "CES technology and business cycle fluctuations," *Journal of Economic Dynamics and Control* 61 (2015), 133–151.

Carvalho, C., J. W. Lee and W. Y. Park, "Sectoral Price Facts in a Sticky-Price Model," *American Economic Journal: Macroeconomics* 13 (2021), 216–256.

Cavallo, A. F., "More Amazon Effects: Online Competition and Pricing Behaviors," *Changing Market Structures and Implications for Monetary Policy, Proceedings of the Federal Reserve Bank of Kansas City 2018 Jackson Hole Symposium* (2018), 291–329.

Chen, X., E. M. Leeper and C. Leith, "Strategic interactions in US monetary and fiscal policies," *Quantitative Economics* 13 (2022), 593–628.

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

Chirinko, R. S. and D. Mallick, "The Substitution Elasticity, Factor Shares, and the Low-Frequency Panel Model," *American Economic Journal: Macroeconomics* 9 (2017), 225–253.

Christensen, L. R., D. W. Jorgenson and L. J. Lau, "Transcendental Logarithmic Production Frontiers," *Review of Economics and Statistics* 55 (1973), 28–45.

———, "Transcendental Logarithmic Utility Functions," *American Economic Review* 65 (1975), 367–383.

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

Cloyne, J., J. Martinez, H. Mumtaz and P. Surico, "Taxes, Innovation and Productivity," Working Paper 30246, National Bureau of Economic Research, 2024.

Coibion, O. and Y. Gorodnichenko, "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics* 7 (2015), 197–232.

Coibion, O., Y. Gorodnichenko and R. Kamdar, "The Formation of Expectations, Inflation, and the Phillips Curve," *Journal of Economic Literature* 56 (2018), 1447–1491.

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

Comin, D. A., R. C. Johnson and C. J. Jones, "Supply Chain Constraints and Inflation," Working Paper 31179, National Bureau of Economic Research, 2023.

Cotton, C. and V. Garga, "The Role of Industrial Composition in Driving the Frequency of Price Change," FRB of Boston Working Paper 22-9, 2022.

Crump, R. K., S. Eusepi, M. Giannoni and A. Şahin, "The unemployment–inflation trade-off revisited: The Phillips curve in COVID times," *Journal of Monetary Economics* 145 (2024), 103580.

Daly, M. C. and B. Hobijn, "Downward Nominal Wage Rigidities Bend the Phillips Curve," *Journal of Money, Credit and Banking* 46 (2014), 51–93.

De Loecker, J. and F. Warzynski, "Markups and firm-level export status," *American Economic Review* 102 (2012), 2437–2471.

Del Negro, M., M. Lenza, G. E. Primiceri and A. Tambalotti, "What's Up with the Phillips Curve?," *Brookings Papers on Economic Activity* 2020 (2020), 301–373.

Di Giovanni, J., Ş. Kalemli-Özcan, A. Silva and M. A. Yildirim, "Global Supply Chain Pressures, International Trade, and Inflation," Working Paper 30240, National Bureau of Economic Research, 2022.

Di Giovanni, J., Ş. Kalemli-Özcan, A. Silva and M. A. Yildirim, "Quantifying the Inflationary Impact of Fiscal Stimulus under Supply Constraints," *AEA Papers and Proceedings* 113 (2023), 76–80.

Dixit, A. K. and J. E. Stiglitz, "Monopolistic Competition and Optimum Product Diversity," *American Economic Review* 67 (1977), 297–308.

Doraszelski, U. and J. Jaumandreu, "R&D and productivity: Estimating endogenous productivity," *Review of Economic Studies* 80 (2013), 1338–1383.

———, "Measuring the bias of technological change," *Journal of Political Economy* 126 (2018), 1027–1084.

Driscoll, J. C. and A. C. Kraay, "Consistent covariance matrix estimation with spatially dependent panel data," *Review of Economics and Statistics* 80 (1998), 549–560.

Eichenbaum, M. and J. D. Fisher, "Estimating the frequency of price re-optimization in Calvo-style models," *Journal of Monetary Economics* 54 (2007), 2032–2047.

Eisfeldt, A. L. and A. A. Rampini, "Capital reallocation and liquidity," *Journal of Monetary Economics* 53 (2006), 369–399.

Eser, F., P. Karadi, P. R. Lane, L. Moretti and C. Osbat, "The Phillips curve at the ECB," *The Manchester School* 88 (2020), 50–85.

Feenstra, R. C., "A homothetic utility function for monopolistic competition models, without constant price elasticity," *Economics Letters* 78 (2003), 79–86.

Fernald, J., "Data for: A quarterly, Utilization-Adjusted Series on Total Factor Productivity," *Working Paper 2012-19, Federal Reserve Bank of San Francisco* (2014a), accessed at http://www.frbsf.org/economic-research/files/quarterly_tfp.xlsx on August 2, 2024.

———, "A quarterly, Utilization-Adjusted Series on Total Factor Productivity," Working Paper 2012-19, Federal Reserve Bank of San Francisco, 2014b.

Fisher, J. D., "On the Structural Interpretation of the Smets–Wouters "Risk Premium" Shock," *Journal of Money, Credit and Banking* 47 (2015), 511–516.

Fitzgerald, T., C. Jones, M. Kulish and J. P. Nicolini, "Is There a Stable Relationship between Unemployment and Future Inflation?," *American Economic Journal: Macroeconomics (Forthcoming)* (2024).

Fujiwara, I. and K. Matsuyama, "Competition and the Phillips curve," Center for Economic Policy Research Discussion Paper DP17521, 2023.

Gagliardone, L. and M. Gertler, "Oil Prices, Monetary Policy and Inflation Surges," (2023).

Gagliardone, L., M. Gertler, S. Lenzu and J. Tielens, "Anatomy of the Phillips Curve: Micro Evidence and Macro Implications," *American Economic Review* forthcoming (2025).

Gagnon, E. and H. Khan, "New Phillips curve under alternative production technologies for Canada, the United States, and the Euro area," *European Economic Review* 49 (2005), 1571–1602.

Galesi, A. and O. Rachedi, "Services Deepening and the Transmission of Monetary Policy," *Journal of the European Economic Association* 17 (2019), 1261–1293.

Galí, J., *Monetary Policy, Inflation, and the Business Cycle: an Introduction to the New Keynesian Framework and its Applications* (Second ed., Princeton University Press, 2015).

Galí, J. and L. Gambetti, "Has the US Wage Phillips Curve Flattened? A Semi-Structural Exploration," *Changing Inflation Dynamics, Evolving Monetary Policy* edited by Gonzalo Castex, Jordi Galí, and Diego Saravia. Santiago, Chile: Banco Central de Chile (2020), 149–72.

Galí, J. and M. Gertler, "Inflation dynamics: A structural econometric analysis," *Journal of Monetary Economics* 44 (1999), 195–222.

Gandhi, A., S. Navarro and D. A. Rivers, "On the identification of gross output production functions," *Journal of Political Economy* 128 (2020), 2973–3016.

Giannone, D. and G. Primiceri, "The Drivers of Post-Pandemic Inflation," Working Paper 32859, National Bureau of Economic Research, 2024.

Goldin, C. and L. F. Katz, *The Race Between Education and Technology* (Harvard University Press, 2009).

Goos, M. and A. Manning, "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *Review of Economics and Statistics* 89 (2007), 118–133.

Griliches, Z., "Capital-Skill Complementarity," *Review of Economics and Statistics* 51 (1969), 465–468.

Hall, R. E., "Invariance Properties of Solow's Productivity Residual," in P. Diamond, ed., *Growth, Productivity, Unemployment: Essays to Celebrate Bob Solow's Birthday* (Cambridge, MA: MIT Press, 1990), 71–112.

———, "Measuring Factor Adjustment Costs," *Quarterly Journal of Economics* 119 (2004), 899–927.

Harding, M., J. Lindé and M. Trabandt, "Resolving the missing deflation puzzle," *Journal of Monetary Economics* 126 (2022), 15–34.

———, "Understanding post-covid inflation dynamics," *Journal of Monetary Economics* 140 (2023), 101–118.

Havranek, T., Z. Irsova, L. Laslopova and O. Zeynalova, "Publication and Attenuation Biases in Measuring Skill Substitution," *Review of Economics and Statistics* (2024), 1–14.

Hazell, J., J. Herreno, E. Nakamura and J. Steinsson, "The Slope of the Phillips Curve: Evidence from US States," *Quarterly Journal of Economics* 137 (2022), 1299–1344.

Heise, S., F. Karahan and A. Şahin, "The Missing Inflation Puzzle: The Role of the Wage-Price Pass-Through," *Journal of Money, Credit and Banking* 54 (2022), 7–51.

Hooper, P., F. S. Mishkin and A. Sufi, "Prospects for inflation in a high pressure economy: Is the Phillips curve dead or is it just hibernating?," *Research in Economics* 74 (2020), 26–62.

Hottman, C. J., S. J. Redding and D. E. Weinstein, "Quantifying the sources of firm heterogeneity," *Quarterly Journal of Economics* 131 (2016), 1291–1364.

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

Hyun, J., R. Kim and B. Lee, "Business cycles with cyclical returns to scale," *International Economic Review* 65 (2024), 253–282.

Imbs, J., E. Jondeau and F. Pelgrin, "Sectoral Phillips curves and the aggregate Phillips curve," *Journal of Monetary Economics* 58 (2011), 328–344.

Inoue, A., B. Rossi and Y. Wang, "Has the Phillips Curve Flattened?," *Working paper* (2024).

Jermann, U. and V. Quadrini, "Macroeconomic Effects of Financial Shocks," *American Economic Review* 102 (2012), 238–271.

Jordà, Ò., S. R. Singh and A. M. Taylor, "The Long-Run Effects of Monetary Policy," Working Paper 2020-01, Federal Reserve Bank of San Francisco, 2023.

Justiniano, A., G. E. Primiceri and A. Tambalotti, "Investment shocks and business cycles," *Journal of Monetary Economics* 57 (2010), 132–145.

Kiley, M. T., "An evaluation of the inflationary pressure associated with short-and long-term unemployment," *Economics Letters* 137 (2015), 5–9.

Kimball, M. S., "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking* 27 (1995), 1241–1277.

Kocherlakota, N. R., "The Concave Phillips Curve," Working Paper 32528, National Bureau of Economic Research, 2024.

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

Krusell, P., L. E. Ohanian, J.-V. Ríos-Rull and G. L. Violante, "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica* 68 (2000), 1029–1053.

Lashkari, D., A. Bauer and J. Boussard, "Information technology and returns to scale," *American Economic Review* 114 (2024), 1769–1815.

Lenzu, S., D. A. Rivers, J. Tielens and S. Hu, "Financial Shocks, Productivity, and Prices," *Working paper* (2024).

Lewis, V. and C. Poilly, "Firm entry, markups and the monetary transmission mechanism," *Journal of Monetary Economics* 59 (2012), 670–685.

L'Huillier, J.-P., G. Phelan and W. R. Zame, "Flattening of the Phillips Curve and the Mandate of the Central Bank," Working paper, 2022.

Ma, Y. and K. Zimmermann, "Monetary Policy and Innovation," Working Paper 31698, National Bureau of Economic Research, 2023.

Mavroeidis, S., M. Plagborg-Møller and J. H. Stock, "Empirical evidence on inflation expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature* 52 (2014), 124–188.

McAdam, P. and A. Willman, "Technology, Utilization, and Inflation: What Drives the New Keynesian Phillips Curve?," *Journal of Money, Credit and Banking* 45 (2013), 1547–1579.

McLeay, M. and S. Tenreyro, "Optimal Inflation and the Identification of the Phillips Curve," *NBER Macroeconomics Annual* 34 (2020), 199–255.

Meeks, R. and F. Monti, "Heterogeneous beliefs and the Phillips curve," *Journal of Monetary Economics* 139 (2023), 41–54.

Michaels, G., A. Natraj and J. Van Reenen, "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years," *Review of Economics and Statistics* 96 (2014), 60–77.

Musso, A., L. Stracca and D. van Dijk, "Instability and Nonlinearity in the Euro-Area Phillips Curve," *International Journal of Central Banking* 5 (2009), 181–212.

Nakamura, E. and J. Steinsson, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics* 123 (2008), 1415–1464.

Ohanian, L. E., M. Orak and S. Shen, "Revisiting capital-skill complementarity, inequality, and labor share," *Review of Economic Dynamics* 51 (2023), 479–505.

Oinonen, S., M. Paloviita and L. Vilmi, "How Have Inflation Dynamics Changed over Time? Evidence from the Euro Area and USA," Discussion Papers 6/2013, Bank of Finland Research, 2013.

Oinonen, S. and L. Vilmi, "Analysing euro area inflation outlook with the Phillips curve," Technical Report 5, BoF Economics Review, 2021.

Okada, T., "Endogenous Technological Change and the New Keynesian Model," *Review of Economics and Statistics* 104 (2022), 1224–1240.

O'Mahony, M. and M. P. Timmer, "Output, input and productivity measures at the industry level: the EU KLEMS database," *Economic Journal* 119 (2009), F374–F403.

Peneva, E. V. and J. B. Rudd, "The passthrough of labor costs to price inflation," *Journal of Money, Credit and Banking* 49 (2017), 1777–1802.

Qiu, Z. and J.-V. Ríos-Rull, "Procyclical Productivity in New Keynesian Models," Working Paper 29769, National Bureau of Economic Research, 2022.

Romer, P. M., "Endogenous Technological Change," *Journal of Political Economy* 98 (1990), S71–S102.

Rotemberg, J. J. and M. Woodford, "The Cyclical Behavior of Prices and Costs," in J. B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*chapter 16 (North Holland: Elsevier Science, 1999), 1051–1135.

Rubbo, E., "Networks, Phillips Curves, and Monetary Policy," *Econometrica* 91 (2023), 1417–1455.

Sbordone, A. M., "Prices and unit labor costs: a new test of price stickiness," *Journal of Monetary Economics* 49 (2002), 265–292.

———, "Globalization and Inflation Dynamics: The Impact of Increased Competition," in *International Dimensions of Monetary Policy* (University of Chicago Press, 2007), 547–579.

Schmitt-Grohé, S. and M. Uribe, "Heterogeneous Downward Nominal Wage Rigidity: Foundations of a Nonlinear Phillips Curve," Working Paper 30774, National Bureau of Economic Research, 2023.

Schorfheide, F., "Estimation and Evaluation of DSGE Models: Progress and Challenges," Working Paper 16781, National Bureau of Economic Research, 2011.

Schorfheide, F. and D. Song, "Real-time forecasting with a mixed-frequency VAR," *Journal of Business & Economic Statistics* 33 (2015), 366–380.

———, "Real-time forecasting with a (standard) mixed-frequency VAR during a pandemic," *International Journal of Central Banking* forthcoming (2025).

Siena, D. and R. Zago, "Job Polarization and the Flattening of the Price Phillips Curve," Banque de France Working Paper 819, 2021.

Sims, C. A., "Solving Linear Rational Expectations Models," *Computational Economics* 20 (2002), 1–20.

Sims, C. A., D. F. Waggoner and T. Zha, "Methods for inference in large multiple-equation Markov-switching models," *Journal of Econometrics* 146 (2008), 255–274.

Smets, F. and R. Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review* 97 (2007), 586–606.

Smirnyagin, V., "Returns to scale, firm entry, and the business cycle," *Journal of Monetary Economics* 134 (2023), 118–134.

Smith, S. C., A. Timmermann and J. H. Wright, "Breaks in the Phillips Curve: Evidence from Panel Data," *Journal of Applied Econometrics* 40 (2025), 131–148.

Stock, J. H. and M. W. Watson, "Why has US inflation become harder to forecast?," *Journal of Money, Credit and Banking* 39 (2007), 3–33.

———, "Modeling Inflation After the Crisis," *Macroeconomic Policy: Post-Crisis and Risks Ahead, Proceedings of the Federal Reserve Bank of Kansas City 2010 Jackson Hole Symposium* (2010), 173–220.

———, "Slack and cyclically sensitive inflation," *Journal of Money, Credit and Banking* 52 (2020), 393–428.

Wang, O. and I. Werning, "Dynamic Oligopoly and Price Stickiness," *American Economic Review* 112 (2022), 2815–2849.

Woodford, M., *Interest and Prices: Foundations of a Theory of Monetary Policy* (Princeton University Press, 2003).

———, "Firm-Specific Capital and the New Keynesian Phillips Curve," *International Journal of Central Banking* 1 (2005), 1–46.

Wu, J. C. and F. D. Xia, "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit and Banking* 48 (2016), 253–291.