Business Cycles with Cyclical Returns to Scale[*]

Jay Hyun[†] — HEC Montréal

Ryan Kim[‡] — Johns Hopkins University

Byoungchan Lee[§] — HKUST

June 23, 2023

Abstract

We study business cycles with cyclical returns to scale. Contrary to tightly parameterized production functions (e.g., Cobb-Douglas and constant elasticity of substitution), we empirically identify strong input complementarity that leads to procyclical returns to scale. We, therefore, propose a flexible translog production function that allows complementarity-induced procyclical returns to scale. We integrate this function into a standard medium-scale dynamic stochastic general equilibrium (DSGE) model. Our estimated model with input complementarity (i) features procyclical returns to scale and acyclical price markups, (ii) better matches the cyclicality of factor shares, and (iii) significantly decreases the contribution of markup shocks to output fluctuations relative to those of the standard model.

***JEL* Codes**: C11, E23, E31, E32

Keywords: Business cycles, Translog production function, DSGE model, Returns to scale, Markup, Markup shocks

---

[*]We are grateful to David Andolfatto, David Berger, Mark Bils, Chris Conlon, Andres Drenik, Heitor S. Pellegrina (discussant), Katherine Ho, Yang Jiao, Loukas Karabarbounis, Amit Khandelwal, Jennifer La'O, Jae Won Lee, Andreas Mueller, Monica Morlacco (discussant), Seunghoon Na, Emi Nakamura, Hyunseung Oh, Zhesheng Qiu (discussant), Ricardo Reis, Bernard Salanie, Stephanie Schmitt-Grohé, Jón Steinsson, Martin Uríbe, Carlos Vegh, Jonathan Vogel, David Weinstein, Michael Woodford, Mark Wynne (discussant), Pierre Yared, Tack Yun, and the seminar participants at the Columbia colloquia, Asian Econometric Society Meeting, EconCon, EGSC, Midwest Macro Meeting, Tsinghua Workshop in Macroeconomics, WEAI, YES, KERIC, and HKUST-Jinan Macro Workshop for their helpful comments and suggestions. This paper subsumes previous papers circulated under the titles "Price-Cost Markup Cyclicality: New Evidence and Implications" and "Business Cycles with Input Complementarity." All errors are our own.

[†]jungsik.hyun@hec.ca, Jungsik Hyun. Department of Applied Economics, HEC Montréal, 3000, chemin de la Côte-Sainte-Catherine Montréal, QC H3T 2A7, Canada

[‡]rkim59@jhu.edu, Sung Ryong Kim. Johns Hopkins University, School of Advanced International Studies, 1717 Massachusetts Ave NW, Office 749, Washington DC 20036

[§]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

Standard business cycle models make strong *a priori* structural assumptions on the shape of the production function. The most widely used production function in macroeconomics is the Cobb-Douglas production function. Despite its convenient tractable features, it imposes an excessively restrictive structure on how firms substitute their inputs (elasticity of substitution), the productivity of each input (marginal product of input), and the productivity of all inputs together (returns to scale). This production function was often justified by the Kaldor (1957) growth facts, but the recent decline in the labor share (e.g., Karabarbounis and Neiman 2014) calls this justification into question. Many researchers acknowledge this limitation and have started to adopt a more general constant elasticity of substitution (CES) production function, but even this production function has important restrictions: one single constant parameter governs the elasticity of substitution among inputs, and returns to scale is typically assumed to be constant and fixed over time.

We empirically assess the plausibility of these restrictions by imposing and estimating a flexible translog production function (Christensen et al. 1973, 1975). Compared to CES, the translog is another generalization of the Cobb-Douglas production function that allows more flexibility in input substitution, the marginal product of input, and returns to scale. Similar to the nonparametric production function estimation technique developed in the industrial organization literature (Gandhi et al. 2020), we utilize the first-order condition of firms to estimate the marginal product of input and to assess the variability in returns to scale. We employ standard panel data techniques with detailed industry-level panel data for the estimation.

In our estimation, we find strong complementarity between labor and energy that leads to time-varying *procyclical* returns to scale. The idea of time-varying returns to scale is striking yet simple. It reflects the idea that when firms employ more factors during boom periods, there are synergies among these factors that lead to larger aggregate marginal product of inputs and returns to scale than in recession periods. The procyclical movement in returns to scale also induces a procyclical wedge between the marginal product of input and the real input price, which is tightly connected to the price markup cyclicality in standard macroeconomic models.

Motivated by our empirical evidence, we estimate a medium-scale dynamic stochastic general equilibrium (DSGE) model as in Smets and Wouters (2007), incorporating a flexible translog production function. Given the empirical importance of complementarity between labor and energy in generating procyclical returns to scale, we include energy input and allow a translog substitution parameter between labor and energy. To understand the implications of using the translog production function, we estimate two other models that are nested in our benchmark model: (i) a standard two-factor Cobb-Douglas production function with labor and capital, i.e., the Smets and Wouters (2007) model, and (ii) a three-factor Cobb-Douglas production function that additionally includes energy input. Comparing the marginal data densities across the three models, we confirm that our model generally outperforms the other two in terms of data fit.

The estimated model with the translog production function generates procyclical returns to scale, consistent with our empirical findings, and acyclical price markups. We confirm significant complementarity between labor and energy, inducing procyclical returns to scale. In turn, the procyclical returns to scale in our otherwise standard DSGE model leads to more procyclical price markups than the conventional models; the large returns to scale during expansions decreases marginal costs and allows price markups to rise. On the contrary, models with the Cobb-Douglas production functions—regardless of including or excluding energy input—generate countercyclical returns to scale and price markups, which are inconsistent with our empirical findings.[1]

Furthermore, we document that the model with a translog structure better matches the empirical cyclicality of input shares than other models with Cobb-Douglas production functions. The data show that labor shares are countercyclical, and energy shares are procyclical. However, the Cobb-Douglas production function with fixed costs cannot match the different cyclicalities of input shares because all input shares are perfectly positively correlated. We break this tight link between input shares and generate data-consistent labor and energy share cyclicality with the flexible translog production function. Moreover, our model generates procyclical capital and profit shares, as in Smets and Wouters (2007).

Finally, in our model with procyclical returns to scale, the contributions of price and wage markup shocks to output fluctuations are substantially smaller than those in models with the Cobb-Douglas production function. As in Smets and Wouters (2007), we conduct the forecast error variance decomposition exercise for each of the three different models with the corresponding production functions. Although integrating the energy input into the conventional Cobb-Douglas production function has a negligible effect on the decomposition results, adding the translog structure with procyclical returns to scale reduces nearly one-third of the contribution of the markup shocks to output. The comparisons of the Bayesian estimation results and the impulse responses of output to price and wage markup shocks across the different models reveal the importance of procyclical returns to scale in suppressing the markup shocks. Having procyclical returns to scale changes the Calvo parameters and amplifies the responses of real variables, which in turn reduces the residual variations that have previously been attributed to price and wage markup shocks. Thus, the estimated markup shock processes are less persistent and feature a smaller impetus, making the markup shocks less important drivers of US business cycles than those in previous studies. The variance decomposition of price markups reveals that these depressed markup shocks, in addition to the changes in the responsiveness of price markups to other shocks, render the price markups more procyclical in our benchmark model than in Cobb-Douglas counterparts.

[1]The conventional models feature countercyclical returns to scale because of the fixed cost of production. In our model with translog production, however, the countercyclical effects of the fixed costs on returns to scale are dominated by the procyclical effects of the input complementarity. See Section 3.3 for details.

To the best of our knowledge, this paper is the first to investigate the role of procyclical returns to scale in a business cycle framework. Our paper is closely related to previous studies that go beyond constant returns to scale in business cycle analysis (Benhabib and Farmer 1994, 1996; Schmitt-Grohé 2000) and a growing literature that generalizes an aggregate Cobb-Douglas production function (e.g., Antras (2004); Chrinko (2008); Karabarbounis and Neiman (2014); Cantore et al. (2015)).[2] Most previous studies reject a Cobb-Douglas production function and find complementarity among inputs beyond what Cobb-Douglas technology implies, similar to our analysis.[3] To integrate the general production function into the DSGE framework, we normalize the translog production function such that it preserves dimensionless parameters, as pioneered by De Jong (1967) and first incorporated into the DSGE framework by Cantore and Levine (2012) and Cantore et al. (2014). Relatedly, Gechert et al. (2022) highlights the importance of normalizing the production function. Regarding time-varying parameters, Koh and Santaeulàlia-Llopis (2017) propose a CES production function that features a time-varying elasticity of substitution. We complement previous studies by proposing a translog production function with time-varying returns to scale.

The procyclical returns to scale speak to the considerable literature that studies the countercyclicality of price markup, which is a first-order building block in many subfields of macroeconomics.[4] Despite its importance, existing empirical evidence on price markup cyclicality is mixed.[5] This paper finds that integrating procyclical returns to scale leads to a novel procyclical margin in price markups in the standard medium-scale DSGE model. Our emphasis on procyclical returns to scale differs from previous papers that emphasize other sources of relatively more procyclical price markup, such as wage rigidity (Nekarda and Ramey 2020), time-varying demand elasticity (Stroebel and Vavra 2019), and endogenous assortment (Anderson et al. 2020). As a complementary mechanism to ours, Drautzburg et al. (2021) considers bargaining shocks, which resemble wage markup shocks, and

[2]There are important studies that microfound the aggregate production function with heterogeneous industry or firm models (e.g., Atalay 2017; Raval 2019; Oberfield and Raval 2021; Smirnyagin 2022). In particular, Baqaee and Farhi (2021) show that aggregate returns to scale can vary over time due to the change in allocative efficiency. Our paper instead extends the production function itself with the translog structure and infers the aggregate implications of using this more flexible production function.

[3]One notable exception is Karabarbounis and Neiman (2014), who find that labor and capital are substitutes; however, when estimating their parameters, they study long-term trends rather than business cycle movements. See, e.g., Hassler et al. (2019) for a discussion of different substitution patterns across short- and long-run horizons among energy and other inputs.

[4]In the context of models with nominal rigidity, countercyclical markups conditional on demand changes are necessary to explain both procyclical wages and countercyclical unemployment (Rotemberg and Woodford 1991, Rotemberg 2013). In the study of monetary policy, many New Keynesian models suggest that central banks should target a constant average markup for price stability (Goodfriend and King 1997). In the scholarship on price dynamics, countercyclical markups conditional on financial distortion explain missing disinflation during the Great Recession (Gilchrist et al. 2017). Ravn et al. (2006) find that the introduction of deep habit formation substantially affects the cyclicality of price markups. Finally, Bils et al. (2018) find that unconditional countercyclical markups explain at least half of the cyclicality in the labor wedge.

[5]Some studies find that price markups are countercyclical, and other studies find that price markups are procyclical or acyclical. See, e.g., Bils (1987), Rotemberg and Woodford (1991), Rotemberg and Woodford (1999), Gali et al. (2007), Bils et al. (2013), and Bils et al. (2018) for countercyclicality and Hall (2013), Stroebel and Vavra (2019), and Anderson et al. (2020) for procyclicality.

explains the factor (capital) share cyclicality using a Bayesian VAR and a structural model. Having procyclical returns to scale also decreases the contributions of markup shocks to output fluctuations compared to Smets and Wouters (2007). These results imply that using a translog production function may alleviate the concerns raised in previous studies about the excessive importance of markup shocks in New Keynesian models (see, e.g., Chari et al., 2009; Justiniano et al., 2010).

The remainder of this paper is structured as follows. Section 2 presents the industry-level data and the estimation results with a translog production function. Section 3 presents and estimates a medium-scale DSGE model with the translog production function and discusses the business cycle implications. Section 4 concludes the paper.

2 Empirical Analyses

This section estimates production function coefficients and assesses the variability in returns to scale under the translog production function. We present the data used in this analysis, the empirical framework, and the estimation results.

2.1 Data

The main dataset used in this paper is the annual six-digit North American Industry Classification System (NAICS) industry-level data from the NBER-CES Manufacturing Industries Database. This database records detailed information on 473 manufacturing industries from 1958 to 2009. The information is compiled from the Annual Survey of Manufacturers and the Census of Manufacturers. The variables in this database include gross output (value of shipment), value-added, and 4-factor inputs (labor, capital, material, and energy) for each industry over time. These data also include industry-specific deflators for output, material, energy, investment, and wage bills for total employees. Appendix A reports the summary statistics of the data (see also Bartelsman et al. 2000).

The most significant advantage of the NBER-CES data over aggregate data is that they allow us to exploit both time-series and cross-sectional variations and corresponding panel data techniques to estimate production function parameters. Substantial variation in the data is especially important for our analysis, which seeks to relax strong functional form assumptions. The advantage comes at a cost, as our estimates come only from manufacturing sectors. As a supplementary analysis, we also use the Integrated Industry-Level Production Account (KLEMS) database, which covers the entire US private economy but for a smaller number of aggregate sectors and a shorter period. Appendix A presents the summary statistics and more information about this database. To further support the representativeness and robustness of the estimated production function parameters, we re-estimate and confirm our results with the DSGE model using time-series data for the entire US economy, as shown in Section 3.2.

2.2 Empirical Framework and Estimation Results

Estimating a translog production function is challenging because it has an excessive number of parameters.[6] For example, for the four inputs available in the NBER-CES and KLEMS data, we must estimate fourteen parameters with a translog production function, many more than the four parameters in a Cobb-Douglas production function and the five parameters in a CES production function. Even with detailed industry-level data for many years, it is difficult to estimate all fourteen parameters in the translog production function because of multicollinearity.

To overcome the challenge of estimating many parameters, we exploit a firm's first-order condition in the spirit of the non-parametric identification method developed in the IO literature (Gandhi et al. 2020).[7] Firms' optimality conditions generate a relationship between the marginal product of each input and its price. Using this relationship and a panel data estimation technique, we recover the part of production associated with the marginal product of a specific input, which has significantly fewer parameters.

In estimating the first-order condition, we address potential estimation concerns by choosing the following three specifications. First, we use energy input to tightly link the marginal product of input (energy) with the real input (energy) price in the first-order condition. Second, we apply log-linearization and demean the variables to make the equations linear and address potential endogeneity concerns. Third, we use lagged input prices as instrumental variables to address non-classical measurement errors and rule out other endogeneity issues. We conduct various robustness checks to address other potential concerns and report the results in Appendix B.

Estimation Framework. For simplicity, consider the following translog production function with only two inputs, labor and capital:

$$\ln(Y) = \underbrace{\varepsilon^a + \beta_l \ln(L) + \beta_k \ln(K)}_{\text{Cobb-Douglas}} + \underbrace{\beta_{lk} \ln(L) \ln(K) + \frac{\beta_{ll}}{2} \ln(L) \ln(L) + \frac{\beta_{kk}}{2} \ln(K) \ln(K)}_{\text{second-order terms}}, \tag{2.1}$$

where Y is output, L is labor, K is capital, and ε^a is the log of total factor productivity. The first part of the production function is a conventional Cobb-Douglas function, which is a first-order approximation of a general production function. A translog production function extends this approximation to the second order. Assuming $\beta_{lk} = 0, \beta_{ll} = 0$, and $\beta_{kk} = 0$ recovers a Cobb-Douglas production function.

[6] For example, Syverson (2011) write: "many researchers also use the translog form... is more flexible, though more demanding of the data." It is also difficult to calibrate parameters given that no previous work integrates a translog production function into the business cycle model.

[7] The other way to proceed is to impose more structure in the estimation, as we did in Section 3.2.

By generalizing Equation (2.1) to four different inputs, we have:

$$\ln(Y) = \underbrace{\varepsilon^a + \sum_i \beta_i \ln(V^i)}_{\text{Cobb-Douglas}} + \underbrace{\sum_i \sum_k \frac{\beta_{ik}}{2} \ln(V^i) \ln(V^k)}_{\text{second-order terms}} \text{ with } \beta_{ik} = \beta_{ki}, \tag{2.2}$$

where V denotes one of four different inputs indexed by i and k, namely, energy (e), labor, capital, and material. Here, we allow for a flexible substitution structure among the four inputs, in contrast to the conventional macroeconomic models that use either two inputs (labor and capital) or three inputs (labor, capital, and material) with implicitly imposed restrictions on the substitution pattern among these inputs.

The simplest method to estimate the parameters in Equation (2.2) is to regress log output on log inputs and treat the residual as the unobserved productivity. This approach has two key problems for our purpose. First, as we have already emphasized, estimating all fourteen parameters in this specification is extremely challenging, even with the rich variation available in the panel data. Second, as widely documented in the productivity estimation literature (e.g., Hall 1988; Evans 1992; Fernald 2014), flexible inputs are likely to be correlated with productivity, generating inconsistent estimates of the parameters. For example, productive industries are likely to use more inputs than other industries in the industry-level data.

To avoid these two concerns, we exploit a firm's first-order condition. Consider a firm's first-order condition with respect to an input V^i:

$$\underbrace{\frac{P^i}{P}}_{\text{real input price}} = \tau^i \underbrace{\left[\beta_i + \sum_k \beta_{ik} \ln(V^k)\right] \frac{Y}{V^i}}_{\text{marginal product of input}}, \tag{2.3}$$

where P^i is the nominal price of the input V^i, and τ^i is the wedge or gap between the real input price P^i/P and the marginal product of input V^i. The input-specific wedge τ^i allows Equation (2.3) to be consistent with a large class of models that deviates from the frictionless economy. Without friction, the marginal product of input equals the real input price, and τ^i could be treated as a classical measurement error or the ex post productivity shock as in Gandhi et al. (2020). Once researchers allow frictions, such as input adjustment costs (Hall 2004), imperfect competition in output (Rotemberg and Woodford 1999) and input (Berger et al. 2019) markets, and financial frictions (Jermann and Quadrini 2012; Arellano et al. 2019; Bigio and La'O 2020), a wedge arises between the marginal product of input and the real input price. This wedge is part of the labor wedge, which is the difference between the marginal product of labor and the marginal rate of substitution (see, e.g., Chari et al. 2007; Karabarbounis 2014; Bils et al. 2018). Note that assuming $\beta_{ik} = 0$ for all $i = 1, ..., 4$ recovers the conventional first-order condition under the Cobb-Douglas production function: $\frac{P^i}{P} = \tau^i \beta_i \frac{Y}{V^i}$.

To avoid the extra structural assumptions required to measure real output and the output price index (see, e.g., Hottman et al. 2016), we rewrite Equation (2.3) as follows:

$$s^i = \tau^i \left[\beta_i + \sum_k \beta_{ik} \ln(V^k) \right], \tag{2.4}$$

where $s^i = \frac{P^i V^i}{PY}$ is the input expenditure share out of total sales for input V^i. The right-hand side is the wedge τ^i multiplied by an output elasticity with respect to the input V^i, which is a unit-free measure of the marginal product of input. The special cases of Equation (2.4) are used to calibrate parameters or inform price markup cyclicality in previous macroeconomic models. In the frictionless economy with Cobb-Douglas technology, an exponent of each input in the production function equals the corresponding input share: $s^i = \beta_i$. This restriction is often used to recover the Cobb-Douglas production function parameters from observed income shares. In models of imperfect competition with the Cobb-Douglas production function, the wedge is interpreted as the inverse of the price-cost markup, and the inverse of the input share identifies the markup up to a constant: $s^i = \beta_i \frac{1}{\text{markup}}$. This restriction allows researchers to inform on markup behavior with the input share.

To input the data into Equation (2.4), we log-linearize the equation around the steady state and allow the input share, wedge, and all inputs to vary across industries and over time, which are indexed by industry j and time t, respectively:

$$\hat{s}^i_{jt} = \sum_k \delta_{ik} \hat{V}^k_{jt} + \hat{\tau}^i_{jt}, \tag{2.5}$$

where $\delta_{ik} \equiv \beta_{ik} \left(\frac{\bar{\tau}^i}{\bar{s}^i} \right)$, $\hat{x}$ denotes the log-deviation from the steady state value, and $\bar{x}$ denotes the steady state value for any variable x. The log-linearization facilitates the estimation by making the equation linear in parameters and is consistent with the DSGE analysis in Section 3.1.

We estimate Equation (2.5) using a standard panel data technique. We double-demean the variables across industries and over time to eliminate industry-specific and time-specific components, including the aggregate trend. The empirical counterpart of Equation (2.5) is:

$$\hat{\hat{s}}^i_{jt} = \sum_k \delta_{ik} \hat{\hat{V}}^k_{jt} + \hat{\hat{\tau}}^i_{jt}, \tag{2.6}$$

where $\hat{\hat{x}}_{jt} = \ln x_{jt} - \frac{1}{J} \sum_{j=1}^{J} \ln x_{jt} - \left[\frac{1}{T} \sum_{t=1}^{T} \left(\ln x_{jt} - \frac{1}{J} \sum_{j=1}^{J} \ln x_{jt} \right) \right]$ for any variable x. Technically, the double-demeaning is identical to allowing industry and time fixed effects. Note that although we demean the variables and remove the aggregate components from all variables, our parameters of interest, δ_{ik}, match the aggregate parameters.

Since the input share s^i_{jt} and all four inputs V^k_{jt} are observed in the data, Equation (2.6) can be estimated by regressing the input share on all four inputs and treating the wedge $\hat{\tilde{\tau}}^i_{jt}$ as a residual. However, there are two major problems in estimating Equation (2.5) directly for every input i. First, the wedge term τ^i_{jt} may contain industry-time-varying components, such as an adjustment cost, that are correlated with inputs and generate a confounding relationship. For example, consider any positive aggregate shock that raises inputs in production. If an industry faces higher input adjustment costs relative to other industries, this industry is likely to utilize fewer inputs relative to other industries.[8] Second, since the input share s^i contains the input V^i, there is a positive mechanical correlation between the input share $s^i = V^i \times \frac{P^i}{PY}$ and the input V^i when V^i has a measurement error. See, for example, Berman et al. (2015) for the formal derivation of such a mechanical correlation.

To address the first estimation concern, we choose energy input as a choice variable and focus on estimating energy efficiency (output elasticity with respect to energy). As a result of using energy input, there are fewer components in the wedge that can be correlated with V^k, particularly regarding the adjustment cost. The energy input and the intermediates in general are known to have smaller adjustment costs than other inputs and are typically assumed away (e.g., Basu 1995; Bils et al. 2018). In addition, other potential concerns related to monopsony power or heterogeneous input quality are mitigated when we focus on energy shares.[9]

In addition to using energy input, we utilize lagged double-demeaned input prices, where demeaning over time only uses past input price information, as instrumental variables to avoid mechanical correlation and relax concerns related to the remaining wedge term. As previously discussed, the input share and the input usage generate a positive mechanical correlation when the variables are measured with error. Unless a researcher has unusually detailed micro-level data, the input variables in any data have measurement error problems. For example, it is difficult to allow for bulk discounts or quality differences in material inputs or to control for the education, experience, and specific skills of labor input. Capital input is known to have a large measurement error even at the firm level (Collard-Wexler and De Loecker 2020), and the perpetual inventory method in the NBER-CES data requires an assumption on initial capital stock. Instrumenting inputs with lagged input prices, which does not involve input usage, solves the mechanical correlation problem that arises from these measurement errors.

[8] Note that this problem resembles the issue in estimating Equation (2.2), which arises from the correlation between productivity and inputs.

[9] Regarding monopsony power, there are fewer concerns on how firms exercise market power in markets for energy inputs, and such friction would not appear as a wedge term. Previous studies have documented such frictions in labor input (e.g., Berger et al. 2019). In contrast, energy production in the U.S. has faced heavy regulation and other restrictions by Congress, such as tax preferences, spending subsidies, and environmental regulations. As a result, it is highly unlikely that firms exert market power in their energy inputs. Regarding heterogeneous input quality, energy input is likely to have homogeneous quality across industries relative to other inputs. Since input prices partially reflect the quality of inputs, the higher input prices and input shares might reflect a higher quality of inputs, which will appear as a wedge in Equation (2.5). See, for example, De Loecker et al. (2016) for the structural treatment of input quality differences in the IO literature.

Regarding the remaining energy wedge $\hat{\hat{\tau}}^e_{jt}$, even after eliminating time- and industry-specific terms, there might exist industry-time-varying wedge components correlated with inputs. By using the lagged input prices as instruments, we assume that the idiosyncratic $\hat{\hat{\tau}}^e_{jt}$ is not correlated with the idiosyncratic component of the predetermined input prices. In doing so, we demean the input prices across industries in a standard panel data method but demean input prices across time using only the past price information to avoid using forward price information; Appendix B.1 shows that using the standard double-demeaning method for the instrumental variables generates similar results.[10] If the idiosyncratic $\hat{\hat{\tau}}^e_{jt}$ is serially uncorrelated, the lagged double-demeaned input prices satisfy the exogeneity assumption since they do not affect the current energy wedge. Furthermore, the instrumental variables satisfy the relevance condition if they are autocorrelated and are correlated with input usage; Appendix A shows that these instruments are highly correlated with $\hat{\hat{V}}^i_{jt}$.[11]

In Appendices B.3 and B.5, we additionally conduct robustness exercises by applying specific interpretations in $\hat{\hat{\tau}}^e_{jt}$ based on previous macroeconomic models and find that controlling for such elements in $\hat{\hat{\tau}}^e_{jt}$ has limited effects on the estimation results. For example, consider a multi-industry business cycle model with heterogeneous price rigidity across industries (e.g., Nakamura and Steinsson 2010). In this case, the energy wedge is a price markup and can vary over time and across industries. Moreover, if industries facing greater price rigidity alter their future markups (inverse wedge) and affect their input prices by changing the input usage relative to other industries facing weaker price rigidity, the exogeneity assumption of instruments could be violated. To address such a concern, we control for industry-time-varying measures of market power following previous studies, such as price-cost markups (De Loecker et al. 2020) and the Lerner index (Gutierrez and Philippon 2017). We also include the measures of price rigidity and inventory-to-sales ratio, which are known to be closely related to price markup, as well as a measure of financial frictions, adjustment costs, and fixed costs in production. Our main empirical results do not change with these alternative specifications, likely because double-demeaning at the detailed industry-time level already eliminates most of the variation in the energy wedge originating for reasons emphasized in previous models.

Equation (2.6) clearly illustrates which variation in the data identifies the substitution pattern among energy and other inputs. Suppose that the coefficient of labor in Equation (2.6) is positive; the energy share increases with an increase in labor input, holding other inputs constant. Under the translog technology, such an increase in the share of energy is interpreted as a result of an increase in energy efficiency that arises from an increase in labor input. In this case, the energy and labor inputs are complements, and the coefficient δ_{ek} captures the strength of the complementarity. Note that a large magnitude of δ_{ek} translates to a considerably smaller magnitude of $\beta_{ek} = \delta_{ek}\frac{\bar{s}^e}{\bar{\tau}^e}$ due to

[10] We are grateful to an anonymous referee for this excellent suggestion.

[11] Note that in using the KLEMS database, the material input is not highly correlated with the lagged input prices once we double-demean the variables, potentially due to the smaller number of observations available in the KLEMS data. Accordingly, we use the NBER-CES as our main dataset and the KLEMS data for supplementary analyses.

the small energy share $\bar{s}^e$ in the data.

Since our empirical strategy is heavily motivated by the productivity estimation literature on industrial organizations, comparing our estimation technique with those in the literature is worthwhile. One key difference in our approach is that we rely on the representative production function, as in a typical DSGE model. Since our exercise is intended to discuss the aggregate parameters, we do not believe that this assumption is particularly worrisome. Additionally, in Section 3.1, we show that a model counterpart of the industry-level energy share (2.5) can be derived by aggregating the corresponding firm-level first-order condition with an assumption on the firm-level translog production function. Similarly, Appendix C.5 illustrates how the aggregate translog production function arises from the explicit aggregation of the firm-level translog production functions. Within the structure of our model, across different levels of aggregation, the interpretations of the production parameters are the same and the estimation assumptions are similar.

Given the industry-level production function, our use of the first-order condition follows Gandhi et al. (2020), who estimate the first-order condition to allow a flexible substitution pattern in a production function. One advantage of using industry-level data is the availability of entity-level input price measures, which are rarely available in more micro-level data. We use input price deflators as instrumental variables to alleviate the mechanical correlation problems and endogeneity concerns. In addition, we use them to measure the quantity of inputs at the entity level and therefore avoid input price bias in using the product- or firm-level data (see, e.g., De Loecker and Goldberg 2014). Double-demeaning, instrumenting, and controlling for variables to address the potential endogeneity of the energy wedge in estimating the first-order condition is similar to the methods that address the endogeneity of total factor productivity in estimating the production function. Our use of instrumental variables is similar to the method of Doraszelski and Jaumandreu (2013), who use the first-order conditions and lagged input prices to address the issue of simultaneity. The double-demeaning with an instrumental variable approach resembles the dynamic panel data method (e.g., Arellano and Bond 1991; Blundell and Bond 1998), and controlling for the potential energy wedge in our robustness exercise is similar to the control function approach (e.g., Olley and Pakes 1996; Levinsohn and Petrin 2003; Ackerberg et al. 2015).

Estimation Results. Table 1 presents the estimated parameters in Equation (2.6). Columns (1)-(3) present the results using the NBER-CES database, and Columns (4)-(6) present the results using the KLEMS database. In using the NBER-CES database, we do not impose weights on each observation for our baseline analyses, but weighting the observation with the industry-specific output leads to similar estimation results. We use both one- and two-year-lagged input prices as instrumental variables to improve the relevance condition. In using the KLEMS database, we explicitly weighted the observation by an average industry output, given that the motivation for

Table 1: Estimation of Equation (2.6)

	Dependent Variable: Energy Share					
Data	NBER-CES			KLEMS		
	(1)	(2)	(3)	(4)	(5)	(6)
Labor	1.923***	1.845***	1.784***	2.112*	2.343**	1.828**
	(0.556)	(0.526)	(0.502)	(1.104)	(1.001)	(0.923)
Energy	-0.612**	-0.559**	-0.557**	1.338	1.293*	1.663***
	(0.268)	(0.244)	(0.239)	(0.843)	(0.722)	(0.542)
Material	0.180	0.176		1.074	0.596	
	(0.233)	(0.175)		(1.581)	(1.260)	
Capital	0.035		0.184	0.444		0.319
	(0.339)		(0.231)	(0.632)		(0.425)
CES test	7.35	8.67	8.05	5.41	5.27	9.09
(p-value)	.01	0	0	.02	.02	0
J-test	3.94	3.28	3.94			
(p-value)	.41	.66	.56			
Observations	22759	22759	22759	1062	1062	1062

Note. Columns (1)-(3) present the IV regression results using the NBER-CES database, and Columns (4)-(6) present the IV regression results using the KLEMS database. All four inputs (labor, energy, material, and capital) and the energy input share are logged and double-demeaned across industries and time. The lagged double-demeaned input prices, which are demeaned using only past price information, for all four inputs are used as instrumental variables; both $t-1$ and $t-2$ lagged input prices are used for the NBER-CES database, and only $t-1$ lagged input prices are used for the KLEMS database. In using the KLEMS data, the observations are weighted by industry-specific output to inform the aggregate representative parameters. For the implementation, we use the GMM specification with the weighting matrix that accounts for the arbitrary correlation among observations within industries. The standard errors in parentheses are clustered at the industry level. The CES test statistic and the corresponding p-value regard the null hypothesis of the (nested) CES functional form (see Appendix B.6), and the J-test and the corresponding p-value refer to the Hansen's J-statistics and p-value for overidentifying restrictions, respectively. *, **, and *** indicate significance at the 0.1, 0.05, and 0.01 levels, respectively.

using the KLEMS data is to inform the aggregate economy beyond the manufacturing sectors. We only use one-year-lagged input prices since two-year-lagged input prices are not highly correlated with inputs with double-demeaning (with past information). However, adding two-year-lagged input prices as instrumental variables does not alter the main results. Appendix B.2 shows that the main results remain robust to using two- and three-year lagged double-demeaned input prices and allowing the correlation of $\hat{\tilde{\tau}}^e_{jt}$ with idiosyncratic lagged input prices up to one year in both databases.

Column (1) is based on the NBER-CES database with all four inputs. The coefficient in front of labor is the most economically and statistically significant estimate, reflecting the complementarity between labor and energy. The estimated parameter shows that a one-percent increase in labor leads to an increase in the energy share of approximately 1.9 percent. The strong complementarity between labor and energy is robust to excluding material or capital input, as shown in Columns (2)-(3), and to using the KLEMS database, as shown in Columns (4)-(6). The coefficient in front of energy is negative in Columns (1)-(3) but positive in Columns (4)-(6), potentially because the energy input becomes less efficient as the manufacturing sectors use more energy but becomes more efficient as the economy as a whole utilizes more energy. The coefficients of capital and material are not statistically significant regardless of using different specifications or data, consistent with the unit elasticity of substitution between energy and material or capital featured in Cobb-Douglas production functions.

Our results suggest the need for a more flexible production function with respect to energy input for macroeconomic models. The coefficient of labor is clearly economically and statistically different from zero, formally rejecting the Cobb-Douglas production function that requires the energy input share to be invariant with respect to any factor input. Furthermore, the empirical results do not support a nested CES production function (see also Appendix B.6). For the robustness checks, Appendix B revisits the main empirical results by considering various other specifications, such as adjusting for the fixed costs in production and controlling for more variables. The key complementarity result between labor and energy, $\delta_{el} > 0$, remains robust across different specifications and is largely consistent with the Bayesian estimation results in Section 3.2.

2.3 Returns to Scale Cyclicality

This section formally defines the returns to scale of the translog production function and assesses its cyclicality. Conceptually, returns to scale measures by what percentage output increases when all inputs increase by one percent. Specifically, by deriving the local elasticity of scale (Hanoch 1975; Epifani and Gancia 2006) for the translog production function $F(\{V^i\}; \varepsilon^a)$ in Equation (2.2), the industry-time-specific returns to scale are expressed as follows:

$$rts_{jt} = \sum_i \left[\beta_i + \sum_k \beta_{ik} \ln(V^k_{jt}) \right], \tag{2.7}$$

where $rts_{jt} \equiv \frac{\partial \log\left[F(\{\lambda V^i_{jt}\}; \varepsilon^a_t)\right]}{\partial \log(\lambda)}|_{\lambda=1}$ denotes the returns to scale and λ reflects the proportional changes in all inputs. Under the conventional Cobb-Douglas production function, the returns to scale do not depend on inputs: $rts_{jt} = \sum_i \beta_i$. Under the translog production function, however, the returns to scale change with input usage. The degree of change depends on the parameters $\{\beta_{ik}\}$, governing the substitution pattern among inputs. If inputs are complements (substitutes), an increase in one input raises (lowers) both the efficiency of other inputs and the returns to scale. The returns to scale in Equation (2.7) nest the constant returns to scale ($rts = 1$) embedded in the CES and Cobb-Douglas production functions.

To examine the returns to scale with the estimated parameters, we follow previous studies on returns to scale (Hall 1990; Basu and Fernald 1997) and additionally assume that the wedge does not differ across inputs: $\tau^i_{jt} = \tau_{jt}$. This assumption still allows the common components of the wedge across different inputs, such as the price markup and fixed cost of production, and it is consistent with the DSGE model presented in Section 3. Despite its consistency with previous work and the DSGE model, one potential concern about this assumption is the presence of an input-specific adjustment cost that is likely to differ across different inputs. In Appendix B.4, we explicitly integrate the input-specific adjustment cost into the analysis and show that the returns to scale cyclicality results are robust to this concern.

With the common wedge assumption, rewriting the returns to scale by combining Equation (2.7) with the sum of Equation (2.4) across all input shares and log-linearizing the resulting equation, we have:

$$\widehat{rts}_{jt} = -\hat{\tau}_{jt} + \hat{s}^{\text{all}}_{jt}, \tag{2.8}$$

where $s^{\text{all}}_{jt} \equiv \sum_i s^i_{jt}$ denotes the sum of all four input expenditure shares of total sales for industry j at time t. $\hat{s}^{\text{all}}_{jt}$ is observed in the data, and $\hat{\tau}_{jt}$ can be recovered as a residual of Equation (2.6) with the estimated parameters reported in Column (1) of Table 1. Similar to what we have done for Equation (2.5), we use double-demeaned variables to recover the log-linearized variables in the data. Due to the demeaning, our estimation strategy only identifies the wedge up to a constant and does not shed light on the *level* of returns to scale. However, we can still analyze the association between returns to scale and value-added and infer the *cyclicality* of the returns to scale.

Figure 1 presents a strong positive relationship between returns to scale and value-added despite using the two different databases. We interpret this result as empirical evidence that supports the notion of *procyclical* returns to scale; when industries experience larger (smaller) value-added, they preserve larger (smaller) returns to scale. The procyclical notion of returns to scale mainly arises from the strong input complementarity between energy and labor, increasing the marginal product of energy in expansions. As a result, production becomes more efficient in expansions with larger returns to scale and an inverse wedge. Since a large part of the inverse wedge is the price markup in canonical DSGE models, our results have important implications for price markup cyclicality. We explore this link more carefully in a standard medium-scale DSGE model in Section 3.

3 Macroeconomic Implications

Motivated by our empirical results, we explore the macroeconomic implications of the complementarity-induced procyclical returns to scale. We integrate the flexible translog production function into a standard medium-scale DSGE model (Smets and Wouters 2007). We re-estimate the model using a Bayesian method with aggregate time-series data and confirm the procyclical returns to scale that arise from the complementarity between labor and energy. By comparing our benchmark model with

Figure 1: Returns to scale

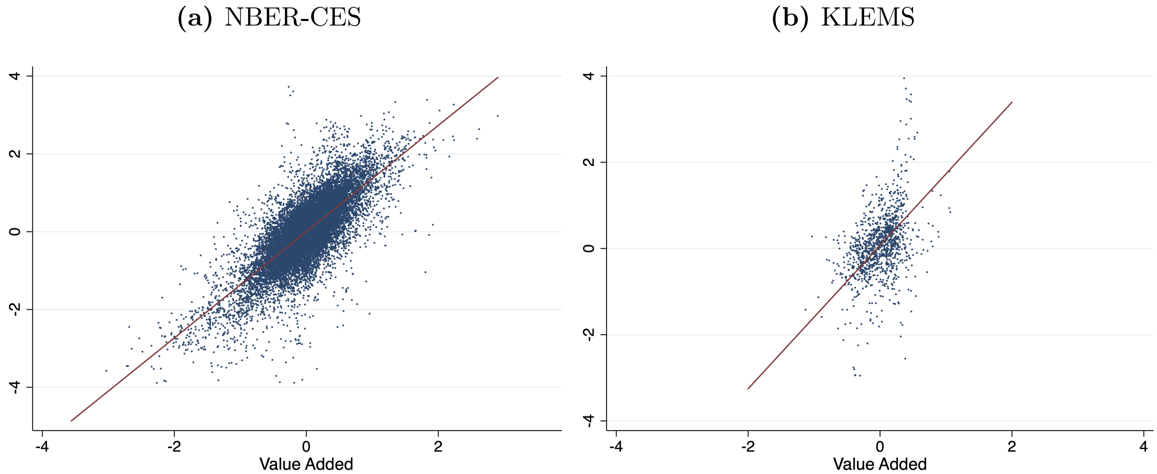

Note. Figure 1 shows the results based on both the NBER-CES (1a) and the KLEMS data (1b). The outliers are excluded for the visibility of the results. The y-axis is the returns to scale, and the x-axis is the value-added. The returns to scale is recovered based on Equation (2.8). All the variables are in double-demeaned logged values, capturing how much percentage the returns to scale increase on average when the value-added increases by one percent. The slopes of the linear lines in Figures 1a and 1b are 1.36 and 1.39, respectively.

models assuming Cobb-Douglas production functions, we find that the model with the procyclical returns to scale (i) generates acyclical price markups instead of countercyclical price markups as in the standard models, (ii) matches the different cyclicality of input shares, and (iii) decreases approximately one-third of the contribution of price and wage markup shocks to output fluctuation.

3.1 Model

This section describes how we extend and nest the Smets and Wouters (2007) model. Our discussion focuses on the key differences from the standard model and the relationship with the empirical framework in Section 2. In particular, we characterize the translog production function with energy input, the corresponding changes in the first-order conditions, and the modeling of the energy market. The other structure of the model follows Smets and Wouters (2007) closely. There exist households, labor unions, final good producers, intermediate goods producers, the government and central bank, and global energy consumers and suppliers in the model. The model features sticky prices, sticky wages, costly capacity utilization, investment adjustment costs, and consumption habits. To focus on the cyclical properties, the model equations below are written in detrended variables using the growth rate on the balanced growth path. The details of the model are relegated to Appendix C.

Production Function. Consider the following detrended translog production function for intermediate good i at time t:

$$y_t(i) = \underbrace{\exp(\varepsilon_t^a)[k_t^s(i)]^{\beta_k}[l_t(i)]^{\beta_l}[e_t(i)]^{\beta_e}}_{\text{Cobb-Douglas}} \underbrace{\left(\frac{l_t(i)}{\bar{l}}\right)^{\beta_{el}\log(e_t/\bar{e})} \left(\frac{e_t(i)}{\bar{e}}\right)^{\beta_{el}\log(l_t/\bar{l})}}_{\text{second-order terms}} - \underbrace{\upsilon}_{\text{fixed costs}} \tag{3.1}$$

$$\text{with } \beta_k + \beta_l + \beta_e = 1,$$

where $k_t^s(i)$ represents capital services used in production, $l_t(i)$ is labor, and $e_t(i)$ is energy. l_t and e_t are aggregate labor and energy, respectively, which individual firms take as given when maximizing their profits, and $\bar{e}$ and $\bar{l}$ are steady-state values of e_t and l_t, respectively. Aggregate productivity $\exp(\varepsilon_t^a)$ follows an exogenous process, and υ is the fixed cost in production. The first part of the production function has a conventional Cobb-Douglas form with constant returns to scale, and the second part captures the second-order terms. All variables in Equation (3.1) are expressed with lowercase letters, denoting the detrended variables around the balanced growth path.[12]

We extend the standard two-factor Cobb-Douglas function with labor and capital to integrate the core of the empirical findings. Given the economically and statistically significant estimate in Table 1, indicating the complementarity between labor and energy, we introduce an energy input $e_t(i)$ and allow the translog substitution structure between labor and energy via the substitution parameter β_{el} in the second-order terms. It is straightforward that our production function nests the three-factor Cobb-Douglas with $\beta_{el} = 0$ and the two-factor Cobb-Douglas used in Smets and Wouters (2007) with $\beta_{el} = 0$ and $\beta_e = 0$. We utilize these two special cases of our general translog production function to emphasize how our new production function changes the traditional business cycle results.

We parsimoniously adapt the translog structure to business cycle models. When the model (3.1) and the empirical (2.2) production functions are compared, the two major changes are apparent. First, we propose the three-factor translog production function with labor, capital, and energy; we do not additionally include material input with the input-output structure. Although adding more inputs is a potentially exciting margin to explore, this is unnecessary for the procyclical returns to scale results and could complicate the already complex analyses of the medium-scale DSGE model. Additionally, Columns (3) and (6) in Table 1 show that there still exists strong complementarity between labor and energy even when material input is excluded. Because the labor-energy complementarity is the most significant, robust, and necessary empirical estimate for the time-varying returns to scale, we focus on β_{el} and the corresponding translog structure in the model and abstract away from the other translog parameters.[13] This minimal adjustment from the conventional framework highlights the role

[12] Specifically, $y_t(i) = \frac{Y_t(i)}{\gamma^t}$, $e_t(i) = \frac{E_t(i)}{\gamma^t}$, $l_t(i) = L_t(i)$, and $k_t^s(i) = \frac{K_t^s(i)}{\gamma^t}$, where γ denotes the steady-state gross growth rate.

[13] Specifically, we do not allow the translog structure between energy and the other inputs (and energy itself).

of the translog structure while avoiding overcomplicating theoretical investigations.

Second, we normalize the production function. The inputs in the second-order terms appear as deviations from the steady-state value. This formulation carefully follows previous studies that integrate the generalized production function into business cycle models to make the production function parameters dimensionless or unit-free (see, e.g., Cantore and Levine 2012; Koh and Santaeulàlia-Llopis 2017). By expressing the production function (3.1) as $\log \frac{y_t(i)+\upsilon}{\bar{y}+\upsilon}$, it is straightforward that the production function parameters are dimensionless with our normalization. Appendix C.6 explicitly shows that the production parameters depend on units of variables without making such a normalization. The normalization also makes the second-order term disappear in steady state and makes our production function directly comparable to the standard Cobb-Douglas specifications without changing the long-run balanced growth path.

When compared to the empirical specification (2.2), there are other minor changes in the normalized translog production function (3.1) that closely follows the specifications in Smets and Wouters (2007). We allow individual firm i's translog production function in the economy and introduce the second-order terms related to the returns to scale, l_t and e_t, as an aggregate externality for individual firms so that they do not choose their own returns to scale.[14] In addition, as in previous work, we incorporate fixed costs in production υ so that firms earn zero profit in steady state. Given that our main empirical analyses in Section 2 do not allow these costs, Appendix B.5 revisits the empirical estimation and finds that the input complementarity and procyclical returns to scale are robust to the existence of the fixed costs. Finally, we include a labor-augmenting deterministic growth rate in the economy so that we can compare the model's outcome to the time-series data with a trend. Appendix C.6 presents the other general properties of the normalized translog production function (3.1).

First-order Conditions. Given real wage w_t, real price of capital service r_t^k, and real price of energy p_t^e, firm i solves its cost-minimization problem subject to the translog production function (3.1). The first-order conditions with respect to energy, labor, and capital are given by:

$$p_t^e = mc_t(i)\frac{y_t(i)+\upsilon}{y_t(i)}\left(\beta_e + \beta_{el}\hat{l}_t\right)\frac{y_t(i)}{e_t(i)}, \tag{3.2}$$

$$w_t = mc_t(i)\frac{y_t(i)+\upsilon}{y_t(i)}\left(\beta_l + \beta_{el}\hat{e}_t\right)\frac{y_t(i)}{l_t(i)}, \tag{3.3}$$

Note that the substitution parameter between energy and capital (δ_{ek}) or between energy and material (δ_{em}) is not statistically significant in Table 1. Also, the sign of the energy square term parameter (δ_{ee}) is not robust to using different datasets: δ_{ee} is estimated to be negative based on the NBER-CES data but positive based on the KLEMS data.

[14]This specification is similar to the externality assumption in the increasing returns to scale literature (e.g., Baxter and King 1991) or the redistributive shock introduced in Rios-Rull and Santaeulalia-Llopis (2010).

$$r_t^k = mc_t(i)\frac{y_t(i)+\upsilon}{y_t(i)}\beta_k\frac{y_t(i)}{k_t^s(i)}, \tag{3.4}$$

where $mc_t(i)$ is the real marginal cost of production or the inverse price markup that arises from monopolistic competition and $\hat{x}$ is the log deviation from the steady state value $\bar{x}$ for any variable x. As in our empirical analyses (Equation (2.3)), Equation (3.2) shows that the real energy price equals the marginal product of energy multiplied by the wedge that includes the real marginal costs and the term related to the fixed cost. Although we evaluate Equation (3.2) at the aggregate level in this model, evaluating the same equation at the industry level recovers the industry-time-varying first-order condition used in Section 2.[15]

Compared to the model with a Cobb-Douglas specification with energy input, the only extensions with the translog are $\beta_{el}\hat{l}_t$ and $\beta_{el}\hat{e}_t$ in the first-order conditions (3.2) and (3.3). Clearly, β_{el} has a first-order effect on factor demand. Assuming complementarity between labor and energy ($\beta_{el} > 0$), an increase in labor or energy during expansions raises the marginal product of the other input more than in the standard case. In contrast, there is a smaller increase in the marginal product of input if labor and energy are substitutes ($\beta_{el} < 0$). Note that changing β_{el} has no direct first-order effect on production (3.1). Our production function is identical to the Cobb-Douglas specification with log-linearization, as explicitly shown in Appendix C.6. On contrary, the model's first-order conditions are different from their Cobb-Douglas counterparts. This feature of the model is desirable given that our empirical framework in Section 2 relies solely on the first-order conditions to identify the complementarity between labor and energy inputs. Within the structure of the otherwise standard model, we reassess our empirical results by re-estimating the key parameter β_{el} with the aggregate data.

Energy Market. Given that we introduce an energy input into the production function, we need to specify the energy market. We introduce a global energy market into our model where the real energy price p_t^e is determined subject to energy demand and supply shocks. We impose a parsimonious structure on the energy market to focus on the general production function and minimize deviations

[15] Specifically, consider a continuum of firms, which is indexed by i and operates in industry j. $y_t(i,j)$ is given by $f(k_t^s(i,j), l_t(i,j), e_t(i,j); l_t(j), e_t(j), \varepsilon_t^a)$, where $l_t(j)$ and $e_t(j)$ are aggregated labor and energy at the industry level, respectively, and $f(k_t^s(i), l_t(i), e_t(i); l_t, e_t, \varepsilon_t^a)$ denotes the production function $y_t(i)$ in Equation (3.1). Each firm (i,j) takes $l_t(j)$ and $e_t(j)$ as given. Under this condition, the firm-level cost minimization problem implies that $s_t^e(i,j) = \tau_t(i,j)\left(\beta_e + \beta_{el}\hat{l}_t(j)\right)$, where $s_t^e(i,j)$ and $\tau_t(i,j)$ are the energy share and the wedge at the firm-industry-time-level, respectively. When firms in the same industry j are exposed to the same industry-level realization of Calvo shocks, these firms become symmetric, and we obtain $s_t^e(j) = \tau_t(j)\left(\beta_e + \beta_{el}\hat{l}_t(j)\right)$, which features industry-level labor in the marginal product of energy. Furthermore, cross-sectional aggregation implies that, up to log-linearization, $s_t^e = \tau_t\left(\beta_e + \beta_{el}\hat{l}_t\right)$, which is the same aggregate equation that we can derive from Equation (3.2). Although we allow the input dispersion to link the theory with the empirical analyses tightly, alternatively, we can assume that production takes place at a representative firm and that intermediaries repackage goods and are subject to Calvo-type price-setting friction. This specification removes the dispersion of factor inputs across i and yields an identical aggregate dynamic to that of the current model up to the first order.

from the benchmark Smets and Wouters (2007) model.

The global energy demand, excluding US industrial usage e_t, is denoted by e_t^d. The energy market clearing condition is given by:

$$e_t + e_t^d = e_t^s, \tag{3.5}$$

where e_t^s is the global energy supply. The global energy supply is determined by the energy price and exogenous disturbances:

$$\frac{e_t^s}{\bar{e}^s} = \left(\frac{p_t^e}{\bar{p}^e}\right)^{\kappa_s} \exp(\varepsilon_t^{es}), \tag{3.6}$$

where κ_s is the price elasticity of e_t^s, and the exogenous supply disturbances ε_t^{es} follow an AR(1) process: $\varepsilon_t^{es} = \rho_{es}\varepsilon_{t-1}^{es} + \eta_t^s$, where $\eta_t^s \sim (0, \sigma_{es}^2)$. We assume that ϕ_e fraction of e_t^s is produced domestically. Thus, the US net energy import is given by $e_t - \phi_e e_t^s$, which implies that the gross domestic product (GDP) is $y_t^{GDP} = y_t - (e_t - \phi_e e_t^s)$.

Global energy demand, excluding US industrial usage, depends on the energy price p_t^e, exogenous disturbances to demand ε_t^{ed}, US GDP y_{t-1}^{GDP}, and real interest rates $\mathbb{E}_{t-1}[R_{t-1}/\Pi_t]$, where R_t and Π_t represent gross nominal interest rates and inflation, respectively:

$$\frac{e_t^d}{\bar{e}^d} = \left(\frac{y_{t-1}^{GDP}}{\bar{y}^{GDP}}\right)^{\rho_{ey}} \left(\frac{\mathbb{E}_{t-1}[R_{t-1}/\Pi_t]}{\bar{R}/\bar{\Pi}}\right)^{\rho_{err}} \left(\frac{p_t^e}{\bar{p}^e}\right)^{-\kappa_d} \exp(\varepsilon_t^{ed}), \tag{3.7}$$

where κ_d is the price elasticity of global energy demand. The exogenous demand disturbances ε_t^{ed} follow an AR(1) process: $\varepsilon_t^{ed} = \rho_{ed}\varepsilon_{t-1}^{ed} + \eta_t^d$, where $\eta_t^d \sim (0, \sigma_{ed}^2)$. Because global economic activity positively affects energy demand (Kilian, 2009; Balke and Brown, 2018), we include lagged US GDP and real interest rates on the right-hand side as proxies for global economic activity. Furthermore, real interest rates capture the states of financial markets that might affect energy prices, as discussed in Kilian (2014) and Basak and Pavlova (2016).

The remaining parts of the model are identical to those in the Smets and Wouters (2007) model. See Appendix C for the full structure of the model.

3.2 Bayesian Estimation

This section presents the Bayesian estimation, consisting of the likelihood function, the prior, and the posterior. For the objective comparisons across models, we apply the same Bayesian techniques, prior, and data as in Smets and Wouters (2007) to both our benchmark model with the three-factor translog function and the model with the three-factor Cobb-Douglas function ($\beta_{el} = 0$). We also bring in the energy data and make relevant prior assumptions to estimate the new parameters introduced regarding the energy market. The estimation results not only confirm the input complementarity ($\beta_{el} > 0$) we find in Section 2 but also reveal more rigid prices and wages and less persistent markup shocks. These changes are important in understanding the sources of the business cycle, which we analyze in Section 3.3.3.

We use a Kalman filter to compute the likelihood of a state-space system. For the state equation, we employ the algorithm suggested in Sims (2002). Our observation equation consists of nine variables, including the growth rate of the global energy supply $\Delta\log(E_t^s) \equiv \log(E_t^s) - \log(E_{t-1}^s)$, the logarithm of the real energy price $\log p_t^e$, and the seven major macroeconomic variables used in Smets and Wouters (2007, Equation (15)):

$$\begin{pmatrix} dlGDP_t \\ dlCONSUMPTION_t \\ dlINVESTMENT_t \\ dlWAGE_t \\ lHOURS_t \\ dlPRICE_t \\ FEDFUNDS_t \\ lENERGYPRICE_t \\ dlENERGY_t \end{pmatrix} = \begin{pmatrix} \bar{\gamma} \\ \bar{\gamma} \\ \bar{\gamma} \\ \bar{\gamma} \\ \bar{l} \\ \bar{\pi} \\ \bar{r} \\ \bar{p}^e \\ \bar{\gamma} \end{pmatrix} + \begin{pmatrix} \hat{y}_t^{GDP} - \hat{y}_{t-1}^{GDP} \\ \hat{c}_t - \hat{c}_{t-1} \\ \hat{i}_t - \hat{i}_{t-1} \\ \hat{w}_t - \hat{w}_{t-1} \\ \hat{l}_t \\ \hat{\pi}_t \\ \hat{r}_t \\ \hat{p}_t^e \\ \hat{e}_t^s - \hat{e}_{t-1}^s \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \nu_t \end{pmatrix}, \tag{3.8}$$

where $\hat{c}_t$, $\hat{i}_t$, $\hat{\pi}_t$, and $\hat{r}_t$ denote the log deviations of (detrended) consumption, investment, gross price inflation, and gross nominal risk-free return, respectively.

We obtain the global energy quantity data $\{E_t^s\}$ from BP Energy (2020). The real energy price is computed by the ratio of the producer price index for total energy to the GDP deflator. Because it is an index, we demean the logarithm of the energy price in Equation (3.8). For the other macroeconomic variables, we extend the dataset constructed by Smets and Wouters (2007) to later periods. Our sample spans from 1966:q1 to 2019:q4. Given the binding zero lower bound, we replace the federal funds rate with the shadow rate of Wu and Xia (2016) from 2009:q1 to 2015:q4. Because E_t^s data are available only at an annual frequency, we interpolate the annual series to construct a quarterly measure and introduce a measurement error ν_t in the observation equation. The standard deviation of ν_t is denoted by σ_ν. In total, we have nine observables, nine structural shocks (global energy supply and demand shocks and the seven macroeconomic shocks in Smets and Wouters 2007), and one measurement error. See Appendix D for further data details.

Table 2 presents the prior and posterior distributions of the parameters regarding the translog structure and the energy market. We assume the same prior as Smets and Wouters (2007) for the other parameters. Among the new parameters, $\frac{\bar{e}}{\bar{e}^s}$, the domestic share of global energy usage in

Table 2: The prior and the posterior of the new parameters

Parameter		Prior			HKL		HKL-CD	
Coeff.	Description	Mean	Std.	Family	Posterior Mode	Credible Set (5%, 95%)	Posterior Mode	Credible Set (5%, 95%)
β_e	SS energy shares	0.05	0.02	Gamma	0.012	(0.008, 0.019)	0.011	(0.007, 0.018)
β_{el}	Input complementarity	0.5	$1/\sqrt{12}$	Uniform	0.030	(0.008, 0.052)	-	-
ρ_{ey}	Elasticity of e^d w.r.t. y^{GDP}	1	0.8	Gamma	0.17	(0.07, 0.35)	0.24	(0.12, 0.37)
ρ_{err}	Elasticity of e^d w.r.t. R/Π	1	0.8	Gamma	0.09	(0.03, 0.49)	0.10	(0.04, 0.46)
κ_d	Price elasticity of e^d	0.1	0.08	Gamma	0.009	(0.003, 0.201)	0.009	(0.005, 0.086)
κ_s	Price elasticity of e^s	0.1	0.08	Gamma	0.10	(0.04, 0.12)	0.10	(0.05, 0.11)
σ_{ed}	Std. of e^d shocks	1	2	Inv. Gamma	0.74	(0.69, 1.74)	0.75	(0.71, 1.14)
σ_{es}	Std. of e^s shocks	1	2	Inv. Gamma	0.72	(0.52, 0.82)	0.72	(0.55, 0.80)
ρ_{ed}	Persistence of e^d shocks	0.5	0.25	Beta	0.9997	(0.9968, 0.9998)	0.9997	(0.9997, 0.9998)
ρ_{es}	Persistence of e^s shocks	0.5	0.25	Beta	0.9996	(0.9973, 0.9999)	0.9996	(0.9996, 0.9996)
$\bar{p}^e$	SS real energy prices	0	2	Normal	-0.12	(-3.35, 3.15)	-0.01	(-3.21, 3.32)
σ_ν	Std. of measurement errors	0.1	0.1	Inv. Gamma	0.05	(0.03, 0.19)	0.05	(0.04, 0.19)
ϕ_e	Domestic share of e^s	0.0265	0.0094	Normal	0.03	(0.01, 0.04)	0.03	(0.01, 0.04)

Notes: HKL denotes our benchmark model with the translog production function. HKL-CD refers to the Cobb-Douglas specification with energy but without complementarity ($\beta_{el} = 0$). SS stands for steady state. e^d, e^s, y^{GDP}, and R/Π are global energy demand, supply, US GDP, and US gross real interest rates, respectively. σ_ν is the standard deviation of ν_t in Equation (3.8).

steady state cannot be separately identified from the other parameters that characterize the dynamics of the energy market. Thus, we set $\frac{\bar{e}}{\bar{e}^s}$ at the average value of $\frac{E_t}{E_t^s}$ in the data during our sample period, which is 0.07. Based on previous studies on the energy shares of value-added, we assume that β_e has a Gamma distribution with a 5% mean and 2% standard deviation.[16] To facilitate the comparison between the aggregate estimate of the input complementarity parameter β_{el} and the micro estimate of δ_{el} in Section 2 (Table 1), we use an uninformative prior of β_{el}. Specifically, β_{el} is assumed to have a standard uniform distribution between 0 and 1.

For the energy market parameters, we use the following priors. We demean $\log p_t^e$ in our observation equation and thus assume that $\bar{p}^e$ has a normal distribution with mean zero and standard deviation two. This choice is similar to the prior of $\bar{l}$ in Smets and Wouters (2007). For the standard deviation of the measurement error, we choose the mean and standard deviation of 0.1. The global share of the US energy production ϕ_e is given by $\frac{E_t - E_t^{ni}}{E_t^s}$, where E, E^{ni}, and E^s represent US energy demand, US net energy imports, and global energy supply, respectively. The sample average of this ratio during the sample period and its standard errors are used as the prior mean and standard deviation of ϕ_e, respectively. We set loose priors of the energy shock parameters $\sigma_{ed}, \sigma_{es}, \rho_{ed}$, and ρ_{es}, similar to the shock parameters in Smets and Wouters (2007). For the elasticities in the energy supply and demand equations (3.6)-(3.7), we rely on a preliminary regression analysis. Specifically,

[16] The energy shares of value added are assumed to be 4% in Rotemberg and Woodford (1996), 10% in Backus and Crucini (2000), 4.3% in Finn (2000), and 5.17% in Dhawan and Jeske (2008).

we estimate $\log E_t^s = c + dt + \kappa^s \log p_t^e + \varepsilon_t^{es}$ and a similar equation for $\log E_t^d$ by using the lagged values of Romer and Romer (2004) monetary policy shocks as instrumental variables. Based on the point estimates and standard errors, we set the prior mean and standard deviations of $\rho_{ey}, \rho_{err}, \kappa_d$, and κ_s. See Appendix D.2 for these regression results.

The posterior is computed using a random walk Metropolis-Hastings algorithm with a chain length of 500,000. We consider three different models: our benchmark model (HKL), the model without input complementarity in production ($\beta_{el} = 0$; HKL-CD), and the Smets and Wouters (2007) model ($\beta_{el} = 0, \beta_e = 0$; S&W). The acceptance rates of the chains are 27%, 31%, and 28%, respectively.

As shown in Table 2, the following parameters are newly introduced and estimated relative to Smets and Wouters (2007). The energy share in steady state β_e is estimated to be 1.2%, and the posterior mode of β_{el} is 0.03. To relate this value to the empirical estimate of δ_{el} in Table 1, we note that the model structure implies the following:

$$\delta_{el} = \beta_{el} \frac{1 - \beta_e}{\beta_e} - (\Phi - 1)\beta_l, \tag{3.9}$$

where Φ is the gross price markup in steady state.[17] At the posterior mode of β_{el}, Φ, β_l, and β_e, this equation yields $\delta_{el} = 2.06$, which is within the range of one standard error of $\hat{\delta}_{el}$ in Table 1. The price elasticities of global energy demand and supply, κ_d and κ_s, are small at the posterior mode, consistent with the results in Hamilton (2009), Kilian (2009), and Kilian and Murphy (2012). Both energy supply and demand shocks are estimated to be highly persistent (ρ_{ed} and ρ_{es}) and larger (σ_{ed} and σ_{es}) than the other structural shocks to match substantial fluctuations in energy prices during the sample period. Finally, the measurement errors (σ_ν) induced by the interpolation are estimated to be relatively small.

Among the other parameters common to the Smets and Wouters (2007) model, Table 3 reports those that show the most noticeable changes across the three models: the parameters associated with the Calvo price and wage stickiness and the markup shocks. The estimation results show that our benchmark model features more rigid prices and wages and less persistent markup shocks than the two other models with Cobb-Douglas production functions.[18] More rigid prices are intuitive, given the input complementarity ($\beta_{el} > 0$) reported in Table 2. As explicitly shown in Section 3.3.1, the input complementarity adds large countercyclical fluctuations to the marginal costs. As a result, the Bayesian estimator of our benchmark model prefers the stickier prices and the flatter price Philips curve to match the correlation of inflation and GDP observed in the data. Given the stickier prices,

[17] Equations (3.2), (C.70) and (C.80) are used in the derivation.

[18] Although the Calvo parameters in our benchmark model are different from those in the two other models, they are broadly consistent with the empirical estimates in Nakamura and Steinsson (2008) and Barattieri et al. (2014) and the model-based estimates in Justiniano et al. (2010, 2011).

Table 3: The Calvo and the markup shock parameters in the three models

Parameter		Prior			Posterior Mode		
Coeff.	Description	Mean	Std.	Family	HKL	HKL-CD	S&W
ξ_p	Calvo sticky price	0.5	0.1	Beta	0.82	0.78	0.77
ξ_w	Calvo sticky wage	0.5	0.1	Beta	0.83	0.79	0.79
ρ_p	Price markup shocks:	0.5	0.2	Beta	0.89	0.95	0.94
μ_p	$\hat{\lambda}_t^p = \rho_p \hat{\lambda}_{t-1}^p + \eta_t^p - \mu_p \eta_{t-1}^p,$	0.5	0.2	Beta	0.79	0.87	0.85
σ_p	where $\eta_t^p \sim N(0, \sigma_p^2)$.	0.1	2	Inv. Gamma	0.13	0.13	0.13
ρ_w	Wage markup shocks:	0.5	0.2	Beta	0.97	0.98	0.98
μ_w	$\hat{\lambda}_t^w = \rho_w \hat{\lambda}_{t-1}^w + \eta_t^w - \mu_w \eta_{t-1}^w$	0.5	0.2	Beta	0.95	0.96	0.96
σ_w	where $\eta_t^w \sim N(0, \sigma_w^2)$.	0.1	2	Inv. Gamma	0.36	0.37	0.37

Notes: HKL denotes our benchmark model with the translog production function. HKL-CD refers to the Cobb-Douglas specification with energy but without complementarity ($\beta_{el} = 0$). S&W is the Smets and Wouters (2007) model.

the stickier wages in our benchmark model are also intuitive. As the Bayesian estimator matches the real wage fluctuation in the data (Equation (3.8)), it changes the wage Philips curve parallel to the flatter price Philips curve and makes nominal wages more rigid.[19] Then, the flatter Philips curves amplify the real variables' responses to structural shocks, yielding smaller residual variations to be explained by markup shocks. Consistent with this intuition, the contribution of price and wage markup shocks to output fluctuations is substantially smaller in the benchmark model than in the two other models, as shown in Section 3.3.3. For the rest of the parameters, the estimation results are similar across the three models (see Appendix D.3).

Our benchmark model (HKL) compares favorably with the two other models regarding the fit to the data. Table 4 shows the marginal data densities for the nine observables in Equation (3.8) and for the seven variables in Smets and Wouters (2007). When the original set of seven variables in the Smets and Wouters (2007) model is used, HKL and HKL-CD feature similarly larger marginal data densities than S&W. When the information in the energy prices and quantities is also considered, HKL outperforms HKL-CD in terms of the data fit. This result implies that the input complementarity in our translog framework can provide useful insights into US macroeconomic fluctuations.

3.3 Business Cycle Implications

This section illustrates the business cycle implications of the normalized translog production function. Despite its parsimonious structure, our benchmark model features procyclical returns to scale

[19]See Equations (C.88) and (C.92) in Appendix C for the price and wage Phillips curves.

Table 4: Marginal data densities

Log marginal data density	(1) HKL	(2) HKL-CD ($\beta_{el} = 0$)	(3) S&W
log(MDD) for the nine observables in Equation (3.8)	-320.6	-326.3	-
log(MDD) for the seven observables in Smets and Wouters (2007)	169.1	170.2	159.4

Notes: HKL denotes our benchmark model with the translog production function. HKL-CD refers to the Cobb-Douglas specification with energy but without complementarity ($\beta_{el} = 0$). S&W is the Smets and Wouters (2007) model. We use the algorithm proposed by Sims et al. (2008) to compute the marginal data densities.

and acyclical price markups rather than the countercyclical returns to scale and price markups in the models with Cobb-Douglas production functions. Moreover, the benchmark model generates countercyclical labor shares and procyclical energy shares consistent with the data and maintains the countercyclical capital shares and procyclical profit shares, similar to previous studies. Finally, the variance decomposition exercises show that our framework leads to less important markup shocks for output fluctuations than the Cobb-Douglas models. These depressed markup shocks, in addition to the presence of the input complementarity parameter ($\beta_{el} > 0$), make the price markups more procyclical in our benchmark model than the Cobb-Douglas models.

3.3.1 Returns to Scale and Price Markups

Armed with the estimated model using aggregate data, we revisit the returns to scale procyclicality discussed in Section 2 and investigate price markup cyclicality. Given the normalized translog production function (3.1), denoted by $y_t(i) = f(k_t^s(i), l_t(i), e_t(i); l_t, e_t, \varepsilon_t^a)$, the returns to scale that firm i faces are given by $rts_t(i) \equiv \frac{\partial \log[f(\lambda k_t^s(i), \lambda l_t(i), \lambda e_t(i); l_t, e_t, \varepsilon_t^a)]}{\partial \log(\lambda)}|_{\lambda=1} = [1 + \beta_{el}(\hat{l}_t + \hat{e}_t)] \frac{y_t(i)+\upsilon}{y_t(i)}$. Note that firms take the aggregate labor l_t and energy e_t as given when they change their scales. Since the cross-sectional dispersion of $\log(y_t(i))$ is of the second order, the average returns to scale across firms, up to the first order, are given by:

$$\begin{aligned} \widehat{rts}_t = & \quad \widehat{rts}_t^{trans} + \widehat{rts}_t^{fix} \\ & = \beta_{el}\left(\hat{l}_t + \hat{e}_t\right) - \frac{\Phi - 1}{\Phi}\hat{y}_t, \end{aligned} \tag{3.10}$$

where $rts_t^{trans} \equiv 1 + \beta_{el}\left(\hat{l}_t + \hat{e}_t\right)$ denotes the part of the returns to scale that arises from the translog structure and $rts_t^{fix} \equiv \frac{y_t + \upsilon}{y_t}$ denotes the part of the returns to scale that emerges because of the fixed costs in production. The second line of Equation (3.10) is derived by relating the fixed cost υ to the price markup Φ under the zero-profit condition in steady state, as in Smets and Wouters

(2007). Note that the rts_t is conceptually the same as our empirical measure of returns to scale rts_{jt} in Section 2.3 since both of them are based on the cross-sectional average of the firm-level returns to scale. As a result of the adjustment in the production function (3.1), compared with Equation (2.7), Equation (3.10) has only one production parameter β_{el} and additionally features the term (rts_t^{fix}) relevant to the fixed costs in production.

The returns to scale vary over the business cycle with the translog structure and fixed costs. In the simplest case with the Cobb-Douglas specification ($\beta_{el} = 0$) and without fixed costs ($\upsilon = 0$), the returns to scale are time-invariant: $rts_t = \beta_k + \beta_l + \beta_e = 1$. Allowing fixed costs in production recovers the returns to scale in Smets and Wouters (2007), $rts_t^{fix} = \frac{y_t+\upsilon}{y_t}$, which are countercyclical. During recessions, firms face relatively higher fixed costs and thus larger returns to scale than during expansions. Additionally, allowing the translog structure makes the returns to scale more procyclical if inputs are complements ($\beta_{el} > 0$) and more countercyclical if inputs are substitutes ($\beta_{el} < 0$). Depending on the degree of input complementarity, the returns to scale may not be negatively correlated with output, even in the presence of fixed costs.

For completeness, we define the aggregate returns to scale, considering aggregate labor and energy endogenously changing with the scale of the economy. Because y_t is obtained by aggregating $y_t(i)$, which is the same as $f(k_t^s, l_t, e_t; l_t, e_t, \varepsilon_t^a)$ up to the first-order, the aggregate returns to scale are given by:

$$\widehat{RTS}_t = 2\beta_{el}\left(\hat{l}_t + \hat{e}_t\right) - \frac{\Phi - 1}{\Phi}\hat{y}_t, \tag{3.11}$$

where $RTS_t \equiv \frac{\partial \log[f(\lambda k_t^s, \lambda l_t, \lambda e_t; \lambda l_t, \lambda e_t, \varepsilon_t^a)]}{\partial \log(\lambda)}|_{\lambda=1}$. The aggregate returns to scale are derived under the assumption that all firms change their inputs proportionally such that the aggregate inputs change by the same proportion. The difference between RTS_t and $rts_t(i)$ is as follows. For RTS_t, the exponents of the second-order term in Equation (3.1) endogenously vary with aggregate inputs when all firms adjust their inputs. In contrast, $rts_t(i)$ treats the exponents of the second-order term as constant because only firm i changes its inputs. This difference yields an additional term $\beta_{el}(\hat{l}_t + \hat{e}_t)$ in the aggregate returns to scale (3.11). Note that $RTS_t = rts_t$ without the translog structure ($\beta_{el} = 0$).

Given the structure of our model, it is straightforward to recover price markups. The changes in price markups can be most easily understood from the following expression for real marginal costs. By rearranging and combining the first-order conditions (3.2)-(3.4) and the production function (3.1), we have:

$$\hat{\Phi}_t = -\widehat{mc}_t = \beta_{el}(\hat{l}_t + \hat{e}_t) - \beta_k \hat{r}_t^k - \beta_l \hat{w}_t - \beta_e \hat{p}_t^e + \varepsilon_t^a, \tag{3.12}$$

where Φ_t is the aggregate price markup, which depends on the translog part of the returns to scale, $\widehat{rts}_t^{trans} = \beta_{el}(\hat{l}_t + \hat{e}_t)$, real input prices ($\hat{r}_t^k, \hat{w}_t$, and $\hat{p}_t^e$), and productivity shocks (ε_t^a). The price

Table 5: Cyclicality of returns to scales and price markups

Correlation with $\log(y_t^{GDP})$	(1) $\log(rts_t^{trans})$	(2) $\log(rts_t^{fix})$	(3) $\log(rts_t)$	(4) $\log(RTS_t)$	(5) $\log(\Phi_t)$
HKL	0.74	-1	0.43	0.62	-0.05
	(0.36, 0.83)	(-1, -1)	(-0.74, 0.65)	(-0.29, 0.76)	(-0.19, 0.10)
HKL-CD ($\beta_{el} = 0$)	-	-1	-1	-1	-0.26
	-	(-1, -1)	(-1, -1)	(-1, -1)	(-0.33, -0.01)
Smets and Wouters (2007)	-	-1	-1	-1	-0.27
	-	(-1, -1)	(-1, -1)	(-1, -1)	(-0.40, -0.04)

Notes: $rts_t^{trans} \equiv 1 + \beta_{el}(\hat{l}_t + \hat{e}_t)$ and $rts_t^{fix} \equiv \frac{y_t + v}{y_t}$ are the returns to scales that arise from the translog production function and fixed costs, respectively. $rts_t \equiv rts_t^{trans} \times rts_t^{fix}$ is the average returns to scale that each firm faces, $RTS_t \equiv [1 + 2\beta_{el}(\hat{l}_t + \hat{e}_t)]\frac{y_t + v}{y_t}$ is the aggregate returns to scale, and Φ_t is the gross price markup. All variables are log-linearized. For each of the three models, we report the correlations of each variable with y_t^{GDP} at the posterior mode and the 90% credible intervals. HKL denotes our benchmark model with the translog production function, HKL-CD refers to the Cobb-Douglas specification with energy ($\beta_{el} = 0$), and Smets and Wouters (2007) is the Cobb-Douglas specification without energy ($\beta_{el} = 0, \beta_e = 0$).

markup rises when firms employ more complementary inputs, face a decrease in input prices, or experience positive productivity shocks. Equation (3.12) clarifies how traditional models link price markup cyclicality to price and wage rigidity. Previous studies have focused on the relative rigidities of prices and wages, which affect the cyclicality of $\hat{w}_t$. Consider the simple case of a Cobb-Douglas specification with only labor input. The price markup then becomes $\hat{\Phi}_t = -\beta_l \hat{w}_t + \varepsilon_t^a$. In this case, conditional on any shocks other than productivity shock, the price markup cyclicality is governed by the cyclicality of real wage $\hat{w}_t$, which is tightly related to the relative rigidities of prices and wages. In our setup, we identify a new term $\beta_{el}(\hat{l}_t + \hat{e}_t)$ that additionally changes the price markup, arising from a flexible input substitution structure embedded in a general production function.

Columns (1)-(4) in Table 5 show that returns to scale are procyclical in our benchmark model (HKL), consistent with the results in Section 2, whereas they are countercyclical in the other models. Column (1) shows the cyclicality of the new returns to scale term, arising from the translog structure $rts_t^{trans} = 1 + \beta_{el}(\hat{l}_t + \hat{e}_t)$. As more inputs are used during expansions, synergies are generated from the complementarity, leading the economy to produce more. Thus, the correlation of the logarithms of rts_t^{trans} and GDP is positive (0.74). In contrast, as shown in Column (2), rts_t^{fix} features perfect or nearly perfect countercyclical returns to scale in all three models as the fixed costs become relatively larger during recessions.[20] Column (3) indicates that the total average returns to scale rts_t is procyclical in our benchmark model because the procyclical effect of the translog structure

[20] In the Smets and Wouters (2007) model, $\log(y_t^{GDP})$ and $\log\left(\frac{y_t+v}{y_t}\right)$ are perfectly negatively correlated because $y_t^{GDP} = y_t$. In models with energy input, although $y_t^{GDP} = y_t - (e_t - \phi_s e_t^s)$, we still observe a nearly perfect negative correlation between $\log(y_t^{GDP})$ and $\log\left(\frac{y_t+v}{y_t}\right)$. This is because the difference between y_t and y_t^{GDP}, the net energy import $e_t - \phi_e e_t^s$, is negligible relative to y_t.

dominates the countercyclical effect of the fixed costs. The aggregate returns to scale RTS_t feature stronger procyclicality due to the endogenous movement in aggregate labor and energy, as shown in Column (4). In the other models with Cobb-Douglas production functions, the aggregate returns to scale equal rts_t^{fix} and are countercyclical.

Column (5) shows that the price markups are acyclical in our benchmark model, whereas they are significantly countercyclical in the other models.[21] This additional procyclical variation in the price markups in HKL originates from the novel element of the returns to scale, $\widehat{rts}^{trans}$. During expansions, when firms utilize more labor and energy, the complementarity between these inputs leads to higher marginal productivity, lower real marginal costs, and larger price markups (see Equation (3.12)). The changes in the parameter estimates (Table 3) associated with the real wage rigidities in HKL also affect the price markup cyclicality. On the other hand, price markups are countercyclical in HKL-CD and S&W.[22]

3.3.2 Factor and Profit Shares

Our translog production function has a novel implications for the income distribution across factors through its effects on the first-order conditions. By rewriting Equations (3.2)-(3.4) in terms of factor shares to the first order, we have:

$$\hat{s}_t^e \equiv \hat{p}_t^e + \hat{e}_t - \hat{y}_t^{GDP} = -\hat{\Phi}_t + \widehat{rts}_t^{fix} + (\widehat{\beta_e + \beta_{el}\hat{l}_t}) + (\hat{y}_t - \hat{y}_t^{GDP}), \tag{3.13}$$

$$\hat{s}_t^l \equiv \hat{w}_t + \hat{l}_t - \hat{y}_t^{GDP} = -\hat{\Phi}_t + \widehat{rts}_t^{fix} + (\widehat{\beta_l + \beta_{el}\hat{e}_t}) + (\hat{y}_t - \hat{y}_t^{GDP}), \tag{3.14}$$

$$\hat{s}_t^k \equiv \hat{r}_t + \hat{k}_t - \hat{y}_t^{GDP} = -\hat{\Phi}_t + \widehat{rts}_t^{fix} + (\hat{y}_t - \hat{y}_t^{GDP}), \tag{3.15}$$

where s_t^e, s_t^l, and s_t^k are energy, labor, and capital shares, respectively, and $(\widehat{\beta_e + \beta_{el}\hat{l}_t}) = \frac{\beta_{el}}{\beta_e}\hat{l}_t$, $(\widehat{\beta_l + \beta_{el}\hat{e}_t}) = \frac{\beta_{el}}{\beta_l}\hat{e}_t$. The translog structure and energy input change the factor share equations from Smets and Wouters (2007) in three respects. First, the translog production function induces additional terms in the energy and labor shares, $\beta_{el}\hat{l}_t$ and $\beta_{el}\hat{e}_t$, which constitute the cyclical components of rts^{trans}. The larger complementarity ($\beta_{el} > 0$) leads to greater energy and labor demand and their

[21]The acyclical price markup in our benchmark model is broadly consistent with the empirical results in Nekarda and Ramey (2020). The correlation between GDP per capita and their markup series (band-pass filtered) varies from -0.66 to 0.37, with a mean of -0.09 across the 14 different specifications and the corresponding markup series.

[22]The length of the credible intervals under HKL reported in Table 5 is smaller when using a tighter prior of β_{el}. Specifically, based on the information in Table 1 (the empirical estimate $\hat{\delta}_{el}$ and its standard error), Equation (3.9), and Φ, β_l, and β_e being equal to their prior means, we can assume a tighter Gamma prior of β_{el} with mean 0.11 and standard deviation 0.03 than a uniform distribution. Under this specification, e.g., the 90% credible interval for $\log(RTS_t)$ becomes $(0.16, 0.78)$. Additionally, $\text{corr}(\log(y_t^{GDP}), \log(\Phi_t))$ at the posterior mode changes to 0.02 with the 90% credible interval being $(-0.08, 0.16)$.

Table 6: Cyclicality of factor and profit shares

Correlation with $\log(y_t^{GDP})$	(1) Energy Shares	(2) Labor Shares	(3) Capital Shares	(4) Profit Shares
HKL	0.72	-0.21	-0.26	0.21
	(0.41, 0.80)	(-0.32, -0.03)	(-0.39, -0.07)	(0.03, 0.32)
HKL-CD ($\beta_{el} = 0$)	-0.23	-0.23	-0.23	0.23
	(-0.33, -0.02)	(-0.33, -0.02)	(-0.33, -0.02)	(0.02, 0.33)
Smets and Wouters (2007)	-	-0.26	-0.26	0.26
	-	(-0.36, -0.06)	(-0.36, -0.06)	(0.06, 0.36)
Data	0.48	-0.26	-	-

Notes: For each of the three models, we report the correlations of the logarithm of each variable with $\log(y_t^{GDP})$ at the posterior mode and the 90% credible intervals. We compute the empirical energy and labor shares based on the data used for the Bayesian estimation in Section 3.2 and employ the Baxter and King (1999) filter with a periodicity of cycles between 6 and 32 quarters. Following Gorodnichenko and Ng (2010), we apply the same filter to the model variables and calculate the correlation coefficients by using the representation in Croux et al. (2001, Equation (8)). HKL denotes our benchmark model with the translog production function, HKL-CD refers to the Cobb-Douglas specification with energy ($\beta_{el} = 0$), and Smets and Wouters (2007) is the Cobb-Douglas specification without energy ($\beta_{el} = 0, \beta_e = 0$).

shares during expansions. Second, the energy input used in production allows us to consider the energy shares. Finally, the net energy imports make y_t^{GDP} marginally different from y_t.

Our benchmark model better matches the energy input share in the data than the model with the Cobb-Douglas production function. As shown in Column (1) of Table 6, our benchmark model features procyclical energy shares.[23] The procyclical energy shares arise from the complementarity between labor and energy ($\beta_{el} > 0$), where a positive β_{el} is motivated by the empirical analyses (Section 2) and is estimated by a Bayesian method (Section 3.2). The input complementarity's procyclical effect dominates the fixed costs' countercyclical effects (rts_t^{fix}) on energy shares. In contrast, the Cobb-Douglas production function yields identically countercyclical factor shares (HKL-CD) because it does not allow input share cyclicalities to differ across factors.

Furthermore, our benchmark model maintains countercyclical labor, capital, and profit shares, consistent with the models with Cobb-Douglas specifications. As presented in Columns (2) and (3), the labor and capital shares are similarly countercyclical across the three models because of the fixed cost in production. This cost generates a countercyclical component in the returns to scale ($\widehat{rts}_t^{fix}$), decreasing the marginal productivity of factors during expansions. For the labor share, although our benchmark model has an additional procyclical term ($\beta_l + \beta_{el}\hat{e}_t$) arising from the input complementarity, it is not large enough to overturn the countercyclical effects of the fixed cost.[24]

[23]We use monthly US industrial energy usage data from U.S. Energy Information Administration (2021) for e_t. We seasonally adjust this series using X-13 ARIMA-SEATS and aggregate it to a quarterly measure. Because this monthly measure is available from 1973, our sample for Table 6 spans from 1973:q1 to 2019:q4.

[24]Note that the magnitude of the procyclical term in the labor share ($\beta_l + \beta_{el}\hat{e}_t$) is smaller than that in the energy share ($\beta_e + \beta_{el}\hat{l}_t$), resulting in a procyclical energy share but a countercyclical labor share. This is because the energy

The labor share cyclicalities in Table 6 are also comparable to the data and the results in Rios-Rull and Santaeulalia-Llopis (2010, Table 2) and Karabarbounis (2014, Table 7). Finally, the profit share is procyclical in all three models, as shown in Column (4).

In addition to assessing the unconditional moments, we verify our benchmark model by comparing the conditional model moments with their empirical counterparts in Appendix D.4. Given that our theoretical mechanism is centered on the complementarity between labor and energy, we compare the model impulse responses of labor and energy to monetary and fiscal policy shocks with the corresponding empirical impulse responses. To do so, we use the identified structural shocks in Romer and Romer (2004), Auerbach and Gorodnichenko (2012), Ramey and Zubairy (2018), and Bauer and Swanson (2022). We find that the theoretical responses in our benchmark model are largely consistent with the empirical responses. Note that this consistency holds without directly matching the industrial energy usage e_t in the Bayesian estimation in Section 3.2.

3.3.3 Variance Decompositions

This section illustrates how the translog production function changes the relative importance of different driving forces of business cycles. For this purpose, we compute the forecast error variance decompositions (FEVDs) of output and labor using the three models and find a noticeable change in the importance of markup shocks. To understand the mechanism behind this result, we investigate the impulse responses of output to price and wage markup shocks and further decompose the unconditional price markup cyclicality into the conditional cyclicalities on each structural shock.

The translog specification substantially decreases the contribution of price and wage markup shocks to business cycles. Table 7 presents the FEVDs of output and labor at an 8-year horizon based on all three models. Although introducing the energy input under the Cobb-Douglas framework (HKL-CD) does not meaningfully change the FEVD from that of the Smets and Wouters (2007) model, additionally introducing the translog specification (HKL) alters the results substantially. The most notable change is the contribution of markup shocks. Price and wage markup shocks explain 29% and 30% of output fluctuations in the Smets and Wouters (2007) model and the model without complementarity (HKL-CD), respectively. In contrast, the corresponding FEVD decreases to 19% in our benchmark model. The results for labor are similar. The FEVDs of labor concerning markup shocks decrease from 46% (S&W) and 47% (HKL-CD) to 35% (HKL) when we incorporate the input complementarity between labor and energy. At alternative horizons, we still observe that markup shocks are less important determinants of output and labor in our benchmark model than in the other models (see Appendix D.5). As a result, the other structural shocks, such as shocks to productivity, demand, and energy, become more important drivers of business cycles in HKL than in HKL-CD

share β_e is small, and with log-linearization, the contribution of the complementarity $\beta_e + \beta_{el}\hat{l}_t$ becomes $\frac{\beta_{el}}{\beta_e}\hat{l}_t$. This base effect generates a large procyclical variation in the energy share.

Table 7: Forecast error variance decomposition of output and labor (32 quarters)

	Output ($\log y_t^{GDP}$)			Labor ($\log l_t$)		
	HKL	HKL-CD	S&W	HKL	HKL-CD	S&W
Panel A						
Productivity (neutral)	0.39	0.34	0.36	0.01	0.01	0.01
Risk premium	0.18	0.14	0.12	0.32	0.23	0.21
Government spending	0.03	0.04	0.04	0.09	0.09	0.09
Investment-specific productivity	0.11	0.13	0.16	0.11	0.12	0.14
Monetary policy	0.06	0.05	0.04	0.11	0.08	0.08
Price markup	0.07	0.16	0.14	0.10	0.19	0.18
Wage markup	0.12	0.15	0.15	0.25	0.28	0.29
Energy demand	0.02	0.00	-	0.01	0.00	-
Energy supply	0.02	0.00	-	0.01	0.00	-
Panel B						
Productivity shocks	0.50	0.47	0.51	0.12	0.13	0.15
Demand shocks	0.27	0.22	0.20	0.52	0.40	0.38
Markup shocks	0.19	0.30	0.29	0.35	0.47	0.46
Energy shocks	0.04	0.01	-	0.01	0.00	-

Notes: Panel A decomposes the forecast error variances into the contributions of nine structural shocks in the model. Panel B summarizes the FEVDs of different types of shocks. The productivity shocks include neutral and investment-specific productivity shocks. The demand shocks include the risk premium, government spending, and monetary policy shocks. The markup shocks include price and wage markup shocks. Finally, the energy shocks include energy demand and supply shocks. HKL denotes our benchmark model with the translog production function, HKL-CD refers to the Cobb-Douglas specification with energy ($\beta_{el} = 0$), and Smets and Wouters (2007) features the Cobb-Douglas production function without energy ($\beta_{el} = 0, \beta_e = 0$). We use the posterior mode of each model for calculating the FEVDs.

and S&W.[25]

To investigate the mechanism behind the decreasing role of markup shocks, we show the impulse responses of output conditional on markup shocks for all three models. The top (bottom) panels in Figure 2 illustrate the responses to a one-standard-deviation contractionary price (wage) markup shock, and the solid, dash-dotted, and dashed lines represent the results based on our benchmark model (HKL), the model without input complementarity (HKL-CD), and the Smets and Wouters (2007) model (S&W), respectively. We do the following to understand the importance of using

[25] Although we extend the dataset constructed by Smets and Wouters (2007) to later periods, using the same sample periods as in Smets and Wouters (2007) yields a similarly depressed role of price markup shocks in HKL than the other two models. Also, in all three models, the demand shocks are estimated to be more important under the extended sample than the sample used by Smets and Wouters (2007), ending in 2004:q4. For example, the demand shocks explain 9% of output forecast error variances at the 8-year horizon (instead of 20% in Table 7) at the posterior mode in Smets and Wouters (2007). The corresponding shares for labor is 17% (instead of 38% in Table 7). It is probably because our extended sample includes the Great Recession periods when the demand shocks are known to be important (Mian et al. 2013; Mian and Sufi 2014; Benguria and Taylor 2020).

Figure 2: Impulse responses of output to price and wage markup shocks

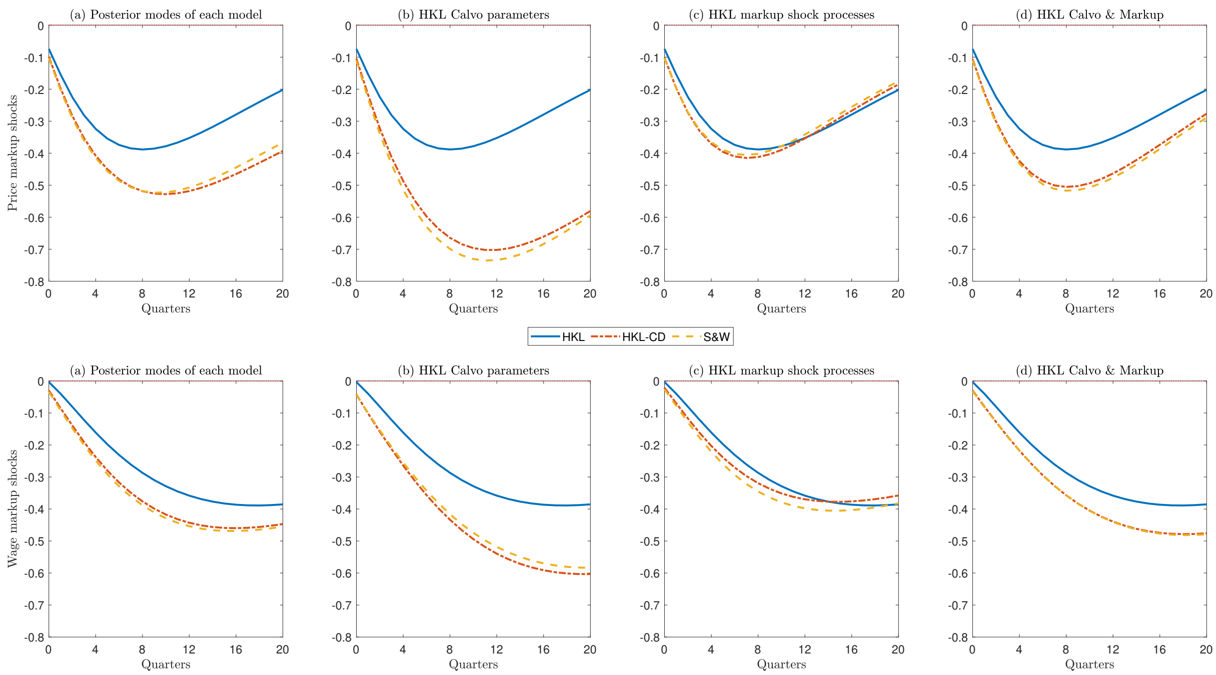

Notes. The two panels in Column (a) show the responses of output to a one-standard-deviation contractionary price and wage markup shock, respectively. The solid, dash-dotted, and dashed lines represent the results based on our benchmark model (HKL), the model without input complementarity (HKL-CD), and the Smets and Wouters (2007) model (S&W), respectively. We use the posterior mode of each model. For Column (b), we replace the Calvo sticky price and wage parameters of HKL-CD and S&W with the corresponding HKL parameters. Column (c) is based on the estimated price and wage markup shock processes of HKL in Table 3. Finally, Column (d) uses Calvo and markup shock parameters at the HKL posterior mode.

different parameters separately estimated for each model. In Column (a), we use each model's posterior mode, consistent with Table 7. In Columns (b)-(d), we fix two sets of parameters—which are notably different—across the models: (i) Calvo price and wage stickiness and (ii) persistence of markup shocks. Then we compare the impulse responses across the models to understand which parameters are responsible for the depressed role of markup shocks. For all three models, Column (b) uses the Calvo parameters in the benchmark model, Column (c) uses the markup shock parameters in the benchmark model, and Column (d) uses both the Calvo and markup shock parameters in the benchmark model.

Column (a) confirms the results in Table 7. The markup shocks have substantially smaller effects on output in our benchmark model than in the other models. The peak effects of a one-standard-deviation price markup shock on output are 0.40% (HKL), 0.53% (HKL-CD), and 0.52% (S&W) in absolute value. Similarly, for wage markup shocks, the peak effects are 0.39% (HKL), 0.46% (HKL-CD), and 0.47% (S&W). The smaller effects of both price and wage markup shocks with the translog production function are similar for investment, consumption, and labor (see Appendix

D.6).

Columns (b)-(d) show that markup shocks are less important in our benchmark model mainly because of less persistent markup shock processes. As shown in Column (b), fixing the price and wage rigidity across the three models makes the impulse response generated from HKL-CD and Smets and Wouters (2007) deviate even more from that of our benchmark model. This result emphasizes that stickier prices and wages in HKL do not decrease the importance of markup shocks as a source of business cycles. In contrast, the impulse responses of output are nearly identical across the models when we use the same markup shock parameters, as shown in Column (c). Thus, markup shock processes—estimated to be less persistent in HKL— are essential for the smaller contribution of markup shocks to output fluctuations. Column (d) fixes both Calvo and markup shock parameters, and the impulse responses are analogous to those in Column (a). Appendix D.7 separately tests whether the positive input complementarity parameter ($\beta_{el} > 0$) per se can solely depress the roles of markup shocks by fixing all the other parameters. We do not find supporting evidence for this conjecture.

As discussed in Section 3.2, the estimated parameters change because procyclical returns to scale lead to less persistent markup shocks with smaller impetus through the flatter price and wage Philips curves. Equation (3.12) shows that our benchmark model features novel procyclical elements in price markups or countercyclical components in real marginal costs because of the input complementarity. The resulting variations in the marginal costs per se do not amplify the dynamics of aggregate variables, especially given other parameters. However, it makes our Bayesian estimator select stickier prices and less responsive price inflation to the real marginal costs to match the empirical correlation of inflation and GDP. Correspondingly, nominal wages become stickier to match the real wage cyclicality in the data. Then, the stronger nominal rigidities, in turn, render real variables to respond more to structural shocks, as shown in Column (b) of Figure 2, absorbing previously unexplained variations that were attributed to price and markup shocks.

The decreasing roles of markup shocks are also critical to understand the price markup cyclicality. Table 8 revisits the price markup cyclicality in Table 5 by decomposing the unconditional covariance between GDP and price markups into the conditional covariances on each structural shock.[26] Panel A shows the unconditional cyclicality (correlation and covariance with GDP), and Panels B and C decompose this cyclicality into the contribution of nine structural shocks and four broad types of shocks, respectively. Columns (1), (2), and (4) use the posterior modes of each model, as in Table 7 and Column (a) of Figure 2. To isolate the role of the input complementarity ($\beta_{el} > 0$) from the other

[26] Because structural shocks are mutually orthogonal, $\text{cov}\left(\hat{y}_t^{GDP}, \hat{\Phi}_t\right) = \text{cov}\left(\sum_j \psi_j^{GDP}(L)\, \varepsilon_t^j, \sum_j \psi_j^{\Phi}(L)\, \varepsilon_t^j\right) = \sum_j \text{cov}\left(\psi_j^{GDP}(L)\, \varepsilon_t^j, \psi_j^{\Phi}(L)\, \varepsilon_t^j\right)$, where $\psi_j^{GDP}(L)$ and $\psi_j^{\Phi}(L)$ represent the impulse response function of $\hat{y}^{GDP}$ and $\hat{\Phi}$ to shock ε^j. Thus, the unconditional covariance $\text{cov}\left(\hat{y}_t^{GDP}, \hat{\Phi}_t\right)$ can be decomposed into the conditional covariances on each structural shock.

Table 8: Covariance decomposition of the price markup cyclicality

	(1) S&W	(2) HKL-CD	(3) HKL w/ $\beta_{el}=0$	(4) HKL	(5) (3)-(2) (%)	(6) (4)-(3) (%)
Panel A: Unconditional moments						
Correlation coefficient	-0.27	-0.26	-0.09	-0.05	-	-
Covariance	-5.44	-6.06	-1.75	-1.20	100	100
Panel B: Conditional moments I						
Productivity (neutral)	1.47	1.44	1.92	1.86	11.19	-11.95
Risk premium	-1.68	-1.98	-2.40	-2.24	-9.85	28.66
Government spending	-0.12	-0.14	-0.14	-0.12	-0.05	3.41
Investment-specific productivity	-1.17	-0.99	-0.67	-0.60	7.45	11.77
Monetary policy	-0.60	-0.62	-0.75	-0.69	-2.95	10.52
Price markup	-5.27	-5.90	-2.32	-2.28	82.91	6.19
Wage markup	1.93	2.09	2.57	2.63	11.00	11.20
Energy demand	-	0.01	0.02	0.12	0.14	19.14
Energy supply	-	0.01	0.02	0.13	0.17	21.04
Panel C: Conditional moments II						
Productivity shocks	0.30	0.45	1.26	1.26	18.63	-0.18
Demand shocks	-2.40	-2.74	-3.29	-3.06	-12.84	42.60
Markup shocks	-3.34	-3.80	0.25	0.34	93.91	17.39
Energy shocks	-	0.02	0.04	0.26	0.30	40.19

Notes: Panel A shows the unconditional correlation coefficients and covariances of GDP and price markups. Panel B decomposes the unconditional covariance into the contributions of nine structural shocks. Panel C summarizes the conditional cyclicality of different types of shocks. The productivity shocks include neutral and investment-specific productivity shocks. The demand shocks include the risk premium, government spending, and monetary policy shocks. The markup shocks include the price and wage markup shocks. Finally, the energy shocks include the energy demand and supply shocks. Column (1) regards the Smets and Wouters (2007) model, featuring the Cobb-Douglas production function without energy ($\beta_{el}=0, \beta_e=0$). HKL-CD in Column (2) refers to the Cobb-Douglas specification with energy ($\beta_{el}=0$). The results based on our benchmark model (HKL) with the translog production function are depicted in Column (4). Column (3) is based on the HKL posterior mode without the input complementarity ($\beta_{el}=0$). Column (5) compares Columns (2) and (3) to focus on the contribution of the changes in the parameter estimates due to the introduction of β_{el}. Column (6) emphasizes the role of β_{el} given the other parameters fixed by comparing Columns (3) and (4).

parameter changes in the estimates, in Column (3), we use the parameters at the HKL posterior mode except for the input complementarity parameter, which is assumed to be zero ($\beta_{el}=0$). To make this comparison explicit, Columns (5) and (6) report the percentage change from Columns (2) to (3) and from Columns (3) to (4), respectively.

As shown in Panel A, unconditional price markups are more procyclical in HKL (Column (4)) than in the Cobb-Douglas models (Columns (1) and (2)), replicating the markup cyclicality results in Table 5. From the comparison across Columns (2)-(4), we show that the input complementarity affects the markup cyclicality both by altering the other parameters at the posterior mode (from Columns (2) to (3)) and by changing the propagation of structural shocks (from Columns (3) to (4)).

The further decomposition into the conditional covariances in Panels B and C emphasizes two notable changes regarding more procyclical price markups in HKL than in HKL-CD and S&W: the decreasing role of price markup shocks and the changes in the conditional cyclicality. First, as shown in Column (5) of Panel B, the most significant increase in the unconditional covariance due to the new parameter estimates originates from the changes in the conditional covariances on price markup shocks, amounting to 82.91%. This change arises mainly because the size of markup shocks shrinks significantly, rendering price markups substantially less countercyclical.[27] Second, as shown in Column (6) of Panel C, the input complementarity induces more procyclical price markups conditional on demand and energy shocks in HKL than in the other models. As the input complementarity between labor and energy raises the returns to scale and lowers the marginal cost of production during demand- and energy-based expansions, price markups increase more in HKL than in the models with Cobb-Douglas specifications.

4 Conclusion

This paper studies business cycles with a translog production function. Our empirical analyses suggest that there is complementarity between labor and energy, leading to procyclical returns to scale. Our empirical evidence is not compatible with the commonly used, tightly parameterized production functions. Thus, we introduce the normalized translog production function into a standard medium-scale DSGE model and re-estimate the input substitution parameters within the structure of our model. Our model rationalizes procyclical returns to scale, acyclical price markups, countercyclical labor shares, procyclical energy shares, and procyclical profit shares. Furthermore, we document that the contribution of price and wage markup shocks to output fluctuations in our model is substantially smaller than that in the Smets and Wouters (2007) model. The complementarity between labor and energy and the corresponding procyclical returns to scale are central to the theoretical mechanism behind the results.

Our work underscores the need to employ general forms of production functions in business cycle research. Further efforts to utilize a general functional form will extend the understanding of business cycles.

[27] Note that sticker prices make the price markup more countercyclical, given the same-sized price markup shocks. However, this effect is weaker than the depressed price markup shock channel, yielding less countercyclical price markups conditional on price markup shocks. On the other hand, for wage markup shocks, the sticker price channel is stronger than the depressed wage markup shock channel, making price markups more procyclical conditional on wage markup shocks.

References

Ackerberg, D. A., K. Caves and G. Frazer, "Structural Identification of Production Functions," *Econometrica* 83 (2015), 2411–2451.

Anderson, E., S. Rebelo and A. Wong, "Markups Across Space and Time," *Working paper* (2020).

Antras, P., "Is the U.S. aggregate production function Cobb-Douglas? New estimates of the elasticity of substitution," *Contributions to Macroeconomics* 4 (2004).

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

Arellano, M. and S. Bond, "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *The Review of Economic Studies* 58 (04 1991), 277–297.

Atalay, E., "How Important Are Sectoral Shocks," *American Economic Journal: Macroeconomics* 9 (2017), 254–280.

Auerbach, A. J. and Y. Gorodnichenko, "Measuring the output responses to fiscal policy," *American Economic Journal: Economic Policy* 4 (2012), 1–27.

Backus, D. K. and M. J. Crucini, "Oil prices and the terms of trade," *Journal of International Economics* 50 (2000), 185–213.

Balke, N. S. and S. P. Brown, "Oil supply shocks and the US economy: An estimated DSGE model," *Energy Policy* 116 (2018), 357–372.

Baqaee, D. R. and E. Farhi, "The Darwinian Returns to Scale," *Working paper* (2021).

Barattieri, A., S. Basu and P. Gottschalk, "Some evidence on the importance of sticky wages," *American Economic Journal: Macroeconomics* 6 (2014), 70–101.

Bartelsman, E. J., R. A. Becker and W. Gray, "The NBER-CES Manufacturing Industry Database," National Bureau of Economic Research Working Paper, Cambridge, MA, 2000.

Basak, S. and A. Pavlova, "A model of financialization of commodities," *Journal of Finance* 71 (2016), 1511–1556.

Basu, S., "Intermediate goods and business cycles," *American Economic Review* 85 (1995), 512–531.

Basu, S. and J. Fernald, "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy* 105 (1997), 249–283.

Bauer, M. D. and E. T. Swanson, "A reassessment of monetary policy surprises and high-frequency identification," *National Bureau of Economic Research* (2022).

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

———, "Measuring business cycles: approximate band-pass filters for economic time series," *Review of Economics and Statistics* 81 (1999), 575–593.

Benguria, F. and A. M. Taylor, "After the Panic: Are Financial Crises Demand or Supply

Shocks? Evidence from International Trade," *American Economic Review: Insights* 2 (December 2020), 509–26.

Benhabib, J. and R. E. A. Farmer, "Indeterminancy and Increasing Returns," *Journal of Economic Theory* 63 (1994), 19–41.

———, "Indeterminancy and Sector-Specific Externalities," *Journal of Monetary Economics* 37 (1996), 421–443.

Berger, D., K. Herkenhoff and S. Mongey, "Labor Market Power," Working Paper, 2019.

Berman, N., A. Berthou and J. Héricourt, "Export Dynamics and Sales at Home," *Journal of International Economics* 96 (2015), 298–310.

Bigio, S. and J. La'O, "Distortions in Production Networks," Working Paper, 2020.

Bils, M., "The Cyclical Behavior of Marginal Cost and Price," *American Economic Review* 77 (1987), 838–855.

Bils, M., P. J. Klenow and B. A. Malin, "Testing for Keynesian Labor Demand," in *NBER Macroeconomics Annual 2013* (University of Chicago Press, 2013), 311–349.

———, "Resurrecting the Role of the Product Market Wedge in Recessions," *American Economic Review* 108 (2018), 1118–1146.

Blundell, R. and S. Bond, "Initial conditions and moment restrictions in dynamic panel data models," *Journal of Econometrics* 87 (1998), 115–143.

BP Energy, "BP Statistical Review of World Energy," June, 2020.

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.

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

———, "New Keynesian Models: Not Yet Useful for Policy Analysis," *American Economic Journal: Macroeconomics* 1 (2009), 242–266.

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

Christensen, L., D. Jorgensen and L. Lau, "Transcendental logarithmic Production Frontier," *Review of Economics and Statistics* 55 (1973), 28–45.

———, "Transcendental logarithmic utility functions," *American Economic Review* 65 (1975), 367–383.

Collard-Wexler, A. and J. De Loecker, "Production Function Estimation and Capital Measurement Error," *Working paper* (2020).

Croux, C., M. Forni and L. Reichlin, "A measure of comovement for economic variables: Theory and empirics," *Review of Economics and Statistics* 83 (2001), 232–241.

De Jong, F. J., *Dimensional Analysis for Economists* (North Holland, 1967).

De Loecker, J., J. Eeckhout and G. Unger, "The Rise of Market Power and the Macroeconomic Implications," *Quarterly Journal of Economics* 135 (2020), 561–644.

De Loecker, J. and P. K. Goldberg, "Firm Performance in a Global Market," *Annual Review of Economics* 6 (2014), 201–227.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal and N. Pavcnik, "Prices, Markups and Trade Reform," *Econometrica* 84 (2016), 445–510.

Dhawan, R. and K. Jeske, "Energy price shocks and the macroeconomy: the role of consumer durables," *Journal of Money, Credit and Banking* 40 (2008), 1357–1377.

Doraszelski, U. and J. Jaumandreu, "R&D and Productivity: Estimating Endogenous Productivity," *Review of Economic Studies* 126 (2013), 1027–1084.

Drautzburg, T., J. FernÃ¡ndez-Villaverde and P. Guerrón-Quintana, "Bargaining shocks and aggregate fluctuations," *Journal of Economic Dynamics and Control* 127 (2021), 104121.

Epifani, P. and G. Gancia, "Increasing returns, imperfect competition, and factor prices," *Review of Economics and Statistics* 88 (2006), 583–598.

Evans, C. L., "Productivity shocks and real business cycles," *Journal of Monetary Economics* 29 (1992), 191–208.

Fernald, J., "A quarterly, utilization-adjusted series on total factor productivity," Federal Reserve Bank of San Francisco Working Paper 2012-19, 2014.

Finn, M. G., "Perfect competition and the effects of energy price increases on economic activity," *Journal of Money, Credit and Banking* 32 (2000), 400–416.

Gali, J., M. Gertler and J. D. Lopez-Salido, "Markups, Gaps, and The Welfare Costs of Business Fluctuations," *Review of Economics and Statistics* 89 (2007), 44–59.

Gandhi, A., S. Navarro and D. Rivers, "On the Identification of Production Functions," *Journal of Political Economy* 128 (2020), 2973–3016.

Gechert, S., T. Havranek, Z. Irsova and D. Kolcunova, "Measuring capital-labor substitution: The importance of method choices and publication bias," *Review of Economic Dynamics* 45 (2022), 55–82.

Gilchrist, S., R. Schoenle, J. W. Sim and E. Zakrajsek, "Inflation Dynamics During the Financial Crisis," *American Economic Review* 107 (2017), 785–823.

Goodfriend, M. and R. King, "The New Neoclassical Synthesis and the Role of Monetary Policy," in *NBER Macroeconomics Annual 1997, Volume 12* (MIT Press, 1997), 231–296.

Gorodnichenko, Y. and S. Ng, "Estimation of DSGE models when the data are persistent," *Journal of Monetary Economics* 57 (2010), 325–340.

Gutierrez, G. and T. Philippon, "Declining Competition and Investment in the U.S.," NBER Working Paper No. 23583, 2017.

Hall, R. E., "The relation between price and marginal cost in US industry," *Journal of Political Economy* 96 (1988), 921–947.

———, "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), 72–112.

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

———, "What the cyclical response of advertising reveals about markups and other macroeconomic wedges," Working Paper 18370, National Bureau of Economic Research, Inc. (October 2013).

Hamilton, J. D., "Understanding crude oil prices," *Energy Journal* 30 (2009), 179–206.

Hanoch, G., "The elasticity of scale and the shape of average costs," *American Economic Review* 65 (1975), 492–497.

Hassler, J., P. Krusell and C. Olovsson, "Directed technical change as a response to natural-resource scarcity," *Working paper* (2019).

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

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

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

———, "Investment shocks and the relative price of investment," *Review of Economic Dynamics* 14 (2011), 102–121.

Kaldor, N., "A Model of Economic Growth," *Economic Journal* 67 (1957), 591–624.

Karabarbounis, L., "The labor wedge: MRS vs. MPN," *Review of Economic Dynamics* 17 (2014), 206–223.

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

Kilian, L., "Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market," *American Economic Review* 99 (2009), 1053–1069.

———, "Oil Price Shocks: Causes and Consequences," *Annual Review of Resource Economics* 6 (2014), 133–154.

Kilian, L. and D. P. Murphy, "Why agnostic sign restrictions are not enough: understanding the dynamics of oil market VAR models," *Journal of the European Economic Association* 10 (2012), 1166–1188.

Koh, D. and R. Santaeulàlia-Llopis, "Countercyclical Elasticity of Substitution," Working Paper, 2017.

Levinsohn, J. A. and A. Petrin, "Estimating Production Functions Using Inputs to Control for Unobservables," *Review of Economic Studies* 70 (2003), 317–340.

Mian, A., K. Rao and A. Sufi, “Household balance sheets, consumption, and the economic slump,” *Quarterly Journal of Economics* 128 (2013), 1687–1726.

Mian, A. and A. Sufi, “What explains the 2007-2009 drop in employment?,” *Econometrica* 82 (2014), 2197–2223.

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

———, “Monetary Non-Neutrality in a Multi-Sector Menu Cost Model,” *Quarterly Journal of Economics* 125 (2010), 961–1013.

Nekarda, C. J. and V. A. Ramey, “The cyclical behavior of the price-cost markup,” *Journal of Money, Credit and Banking* 52 (2020), 319–353.

Oberfield, E. and D. Raval, “Micro data and macro technology,” *Econometrica* 89 (2021), 703–732.

Olley, S. G. and A. Pakes, “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica* 64 (1996), 1263–1297.

Ramey, V. A. and S. Zubairy, “Government spending multipliers in good times and in bad: evidence from US historical data,” *Journal of Political Economy* 126 (2018), 850–901.

Raval, D. R., “The micro elasticity of substitution and non-neutral technology,” *RAND Journal of Economics* 50 (2019), 147–167.

Ravn, M. O., S. Schmitt-Grohé and M. Uribe, “Deep Habits,” *Review of Economic Studies* 73 (2006), 195–218.

Rios-Rull, J.-V. and R. Santaeulalia-Llopis, “Redistributive shocks and productivity shocks,” *Journal of Monetary Economics* 57 (2010), 931–948.

Romer, C. D. and D. H. Romer, “A new measure of monetary shocks: Derivation and implications,” *American Economic Review* 94 (2004), 1055–1084.

Rotemberg, J. J., “Comment on Testing for Keynesian Labor Demand,” in *NBER Macroeconomics Annual 2013*volume 27 (1) (University of Chicago Press, 2013), 362–370.

Rotemberg, J. J. and M. Woodford, “Markups and the Business Cycle,” in *NBER Macroeconomics Annual 1991* (MIT Press, 1991), 63–129.

———, “Imperfect competition and the effects of energy price increases on economic activity,” *Journal of Money, Credit, and Banking* 28 (1996), 550–577.

———, “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.

Schmitt-Grohé, S., “Endogenous Business Cycles and the Dyanmics of Output, Hours, and Consumption,” *American Economic Review* 90 (2000), 1136–1159.

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," Working Paper, 2022.

Stroebel, J. and J. Vavra, "House Prices, Local Demand, and Retail Prices," *Journal of Political Economy* 127 (2019), 1391–1436.

Syverson, C., "What determines productivity?," *Journal of Economic Literature* 49 (2011), 326–365.

U.S. Energy Information Administration, *Monthly Energy Review February 2021* (Washington, DC: Office of Energy Statistics, U.S. Department of Energy, 2021).

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.