

NETWORKS, PHILLIPS CURVES, AND MONETARY POLICY

ELISA RUBBO

Booth School of Business, University of Chicago

This paper revisits the New Keynesian framework, theoretically and quantitatively, in an economy with multiple sectors and input-output linkages. Analytical expressions for the Phillips curve and welfare, derived as a function of primitives, show that the slope of all sectoral and aggregate Phillips curves is decreasing in intermediate input shares, while productivity fluctuations endogenously generate an inflation-output tradeoff—except when inflation is measured according to the novel divine coincidence index. Consistent with the theory, the divine coincidence index provides a better fit in Phillips curve regressions than consumer prices. Monetary policy can no longer achieve the first-best, resulting in a welfare loss of 2.9% of per-period GDP under the constrained-optimal policy, which increases to 3.8% when targeting consumer inflation. The constrained-optimal policy must tolerate relative price distortions across firms and sectors in order to stabilize the output gap, and it can be implemented via a Taylor rule that targets the divine coincidence index.

KEYWORDS: Input-output networks, Phillips curve, monetary policy, price index, inflation targeting, welfare cost of business cycles.

1. INTRODUCTION

THE NEW KEYNESIAN FRAMEWORK PROVIDES THE THEORETICAL UNDERPINNING of modern monetary policy. Yet, the workhorse version of this framework rests on a simplified representation of production, with only one industry and no intermediate inputs. In reality, production happens across multiple sectors, which buy intermediate inputs from each other and sell to final consumers. To introduce these realistic features in the New Keynesian framework, this paper combines it with a multi-sector model of production, and systematically revisits its three pillars—the IS curve, the Phillips curve, and the welfare function—both theoretically and quantitatively.

While the IS curve remains unaffected, the Phillips curve and the welfare function are drastically changed, with crucial implications for monetary policy. Even when fluctuations are driven only by productivity, the central bank can no longer achieve the efficient equilibrium. Moreover, consumer price inflation is neither a good proxy for the output gap, nor a good policy target. Instead, monetary policy should aim at stabilizing a novel inflation measure—the divine coincidence index—which weights sectors according to sales shares and price adjustment frequencies, and is a sufficient statistic for the output gap.

Elisa Rubbo: elisarubbo@uchicago.edu

I am grateful to my advisors—Emmanuel Farhi, Elhanan Helpman, Gita Gopinath, and Gabriel Chodorow-Reich—for invaluable guidance and continuous support. I am indebted to Andrew Lilley and Giselle Montamat for many insightful conversations all throughout the development of this paper. I thank for their comments three anonymous referees, Mark Aguiar, Pol Antràs, Adrien Auclert, Gadi Barlevy, David Baqaee, Kirill Borusyak, Ashley Craig, Francois Gourio, Erik Hurst, Rohan Kekre, Greg Mankiw, Mikkel Plagborg-Møller, Richard Rogerson, Ken Rogoff, Raphael Schoenle, Jim Stock, and Ludwig Straub, and seminar participants at the Chicago Fed, the Cleveland Fed, the European Central Bank, Harvard, Princeton, Chicago, MIT, Columbia GSB, LSE, UCL, CREI, Yale, Boston University, UC San Diego, the San Francisco Fed, the Minneapolis Fed, the St Louis Fed, Rochester University, Brown University, Bocconi University, Mannheim University, NUS, Bank of England, Bank of Israel, NYUAD, and the Virtual Macro Seminar Series. I thank Erin Eidschun for excellent research assistance.

© 2023 The Author. *Econometrica* published by John Wiley & Sons Ltd on behalf of The Econometric Society. Elisa Rubbo is the corresponding author on this paper. This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

The multi-sector model, introduced in Section 2, allows for an arbitrary number of industries, with general constant-returns-to-scale production functions, which take as inputs labor and any combination of intermediate goods. As Section 3 illustrates, the new framework is tractable enough that its log-linearized version can be solved analytically, while at the same time it can be calibrated to fully match disaggregated input-output data.

Section 4 then turns to studying the Phillips curve. In multi-sector economies, one can construct inflation rates and Phillips curves for every sector. Since prices across sectors are differentially affected by changes in the output gap, the Phillips curve slopes are sector-specific. Moreover, in sharp contrast with the one-sector benchmark, in the multi-sector economy productivity fluctuations generate a tradeoff between closing the output gap and eliminating inflation, both at the sector level and in the aggregate. This tradeoff manifests as sector-specific, productivity-driven residuals in the Phillips curves, and it affects all inflation measures, except for the divine coincidence index.

The slopes of all sectoral Phillips curves get flatter as producers use more intermediate inputs, because price rigidities get compounded at each step along the production chain. Hence, by properly accounting for input-output linkages, the multi-sector model has the potential to reconcile relatively high micro level frequencies of price adjustment with the flat Phillips curves measured in empirical studies. Section 4 formalizes this argument,¹ deriving an expression for the slope of sectoral and aggregate Phillips curves as a function of micro level intermediate input shares and price adjustment probabilities.

A key result in the one-sector economy, often referred to as the divine coincidence, states that productivity fluctuations do not generate an inflation-output tradeoff, and therefore they do not affect the residual of the Phillips curve. In the multi-sector economy, instead, productivity shocks affect the residual term of all sectoral and aggregate Phillips curves, except the divine coincidence one. Hence, the divine coincidence index is the unique inflation-based sufficient statistic for the output gap, and the corresponding empirical Phillips curve is the only well-specified one.

In the one-sector economy, the divine coincidence holds because of an equilibrating mechanism, whereby, whenever productivity falls (increases), wages always also fall (increase), mirroring the marginal product of labor. The wage response stabilizes marginal costs, thereby offsetting any inflationary pressure. In economies with sticky wages or multiple sectors, however, the wage adjustment may be dampened, and industries may be differentially exposed to wage and productivity changes. Therefore, prices rise in sectors which suffer more from decreasing productivity than they benefit from decreasing wages, and vice versa. Section 4 analytically characterizes the effect of sectoral productivity shocks on sectoral and aggregate inflation. In particular, the multi-sector model predicts positive consumer price inflation whenever productivity falls in downstream industries—which means closer to final consumers—or in industries with more flexible prices.

Having characterized the positive properties of the multi-sector model, Section 5 then turns to welfare and optimal monetary policy. Like in the one-sector model, in multi-sector economies too there is a welfare loss whenever employment or relative prices deviate from the efficient level. The welfare loss function expresses the cost of employment and price distortions in terms of the network primitives. In multi-sector economies, relative price distortions occur both within and across sectors, and cause producers and final consumers to inefficiently shift demand towards the firms and sectors whose relative price

¹The argument was first made by Basu (1995) in a simpler model, and later by Christiano (2016). Quantitatively, Carvalho (2006) and Nakamura and Steinsson (2013) showed that input-output linkages increase monetary non-neutrality significantly.

is too low given their productivity. The demand response is governed by sectoral elasticities of substitution. For within-sector distortions, the welfare cost is the same function of sectoral inflation, price adjustment probabilities, and substitution elasticities as in the one-sector model. In the multi-sector model, the welfare function further weights each sector according to the relevant sale share. Price distortions across sectors are new to the multi-sector model, and their welfare cost depends on the interaction between distortions originating in different sectors. Two distortions offset each other if they move downstream prices in the same direction, because on net they produce smaller deviations from the efficient relative prices, and vice versa.

Beyond deriving the welfare loss function, Section 5 also draws the implications for monetary policy. Given the multiple production sectors, the central bank has considerably limited power, because it has only one instrument, the nominal interest rate. Therefore, it can only move prices along one direction—parallel to the slope of sectoral Phillips curves—which is unlikely to coincide, when there are multiple sectors, with the efficient relative price change. Moreover, given that Phillips curves are flatter with input-output linkages, the central bank must consequently move employment far away from the efficient level in order to affect prices. For all these reasons, in the multi-sector economy monetary policy should focus on closing the output gap. Nonetheless, this does not mean the end of inflation targeting, because, as Section 5 demonstrates, the central bank can close the output gap at all periods by replacing consumer price inflation with the divine coincidence index in the Taylor rule. However, since it cannot correct relative prices, the central bank must tolerate a welfare loss.

Having elaborated the theoretical aspects in Sections 4 and 5, Sections 6 and 7 explore the empirical properties of the multi-sector framework. Section 6 calibrates the multi-sector model to the U.S. input-output data. By properly accounting for intermediate input linkages and wage rigidity, in the multi-sector model the slope of the consumer price Phillips curve is reduced by one order of magnitude, reconciling micro level measures of price adjustment frequencies with macro level estimates. To evaluate how changes in the input-output structure over time affected the Phillips curve, we calibrate the multi-sector model to historical input-output data for 1947–2017. Although the slope remains one order of magnitude smaller than in the one-sector benchmark throughout the whole period, during this period it has further flattened by about 30% due to growing intermediate input flows relative to labor compensation.

Section 6 also estimates the welfare loss from business cycles, when the central bank follows alternative policy rules. In the multi-sector economy, the welfare loss is several orders of magnitude larger than in the one-sector benchmark. It is a well-known result that, in frictionless economies, the welfare loss from business cycles is tiny—about 0.05% of per-period GDP according to the famous Lucas (1977) estimate. Moreover, in the one-sector New Keynesian model, welfare is the same as in the frictionless economy, because monetary policy can replicate the efficient equilibrium. By contrast, in the multi-sector economy the central bank cannot eliminate relative price distortions—even when it adopts the optimal policy—leading to an average loss of 2.9% of per-period GDP relative to the efficient flex-price equilibrium. If instead the central bank targets consumer prices, the welfare loss increases to 3.8% of per-period GDP.

Finally, Section 7 compares Phillips curve estimates for the divine coincidence index and several common measures of consumer price inflation (the CPI, the PCE, and their core versions), based on time series data for sector-level inflation. As predicted by the theoretical model, the divine coincidence specification outperforms all consumer price specifications over the sample period 1984–2018, in several dimensions. First, the divine

coincidence specification yields a two to four times larger R-squared than the consumer price specifications. Second, estimated in 20-year rolling regressions, the Phillips curve slope always has the correct sign, and it is significant and stable over time. By contrast, when using the consumer price specifications, the slope is often insignificant and has the wrong sign. Both results are consistent with the model prediction that the divine coincidence index is more robustly correlated with the output gap, because it does not suffer from productivity-driven cost-push shocks.

Readers can refer to the Supplement for proofs of the theoretical results and more detail about the Phillips curve regressions. The additional Additional Online Materials (Rubbo (2023)) included in the replication files reports further calibration exercises and robustness checks for the empirical results, referenced throughout the text.

1.1. *Related Literature*

The framework in this paper builds on Baqaee and Farhi (2019, 2020b),² who studied the implications of exogenous markups in input-output models. Compared to Baqaee and Farhi (2020b), this paper further characterizes the endogenous evolution of sectoral markups determined by price rigidities, and the consequent inflation response.

In previous work on monetary policy in multi-sector economies, researchers focused on either small-scale analytical models (Gali (2015), Woodford (2003), Gali and Monacelli (2008), Kara (2009, 2010)), or on large-scale quantitative models (Carvalho (2006), Carvalho and Nechio (2011), Cagliarini, Robinson, and Tran (2011), Nakamura and Steinsson (2013), Pasten, Schoenle, and Weber (2017), Pasten, Schoenle, and Weber (2019), Castro Cienfuegos (2019), Hoynck (2020), Ghassibe (2021)). The current paper combines the two approaches by revisiting and extending previous analytical results in a quantitatively realistic production network. This allows us to formalize in a general network the insight that input-output linkages flatten the Phillips curve (Basu (1995), Christiano (2016)), expressing sectoral and aggregate slopes in terms of micro level input-output data and price adjustment frequencies. Moreover, our approach allows us to provide analytical expressions for the cost-push shocks which affect sectoral Phillips curves as a result of changes in productivity. Building on this novel result, we can construct the divine coincidence inflation index and show that it is a better sufficient statistic for the output gap.

Discussing the optimal monetary policy, Section 5 below generalizes previous results from New Keynesian models with two sectors or wage rigidities (Erceg, Henderson, and Levin (2000), Blanchard and Gali (2007)), and it thus confirms that consumer prices are a suboptimal policy target in multi-sector economies (Aoki (2001), Benigno (2004), Huang and Liu (2005)). Our general approach allows us to establish that the output gap is always a nearly optimal policy target, regardless of the network structure. Moreover, our approach demonstrates that one can always stabilize the output gap by targeting the divine coincidence inflation index. Compared to previous quantitative studies of the optimal inflation target (Mankiw and Reis (2003), Eusepi, Hobijn, and Tambalotti (2011)), this paper characterizes analytically the weighting scheme for the divine coincidence index, showing that it does not depend directly on micro level input-output shares, but only on sectoral sales shares and price adjustment frequencies.

In parallel and independent work, La'O and Tahbaz-Salehi (2019) also studied optimal monetary policy in multi-sector economies with an input-output network. This paper differs from La'O and Tahbaz-Salehi (2019) in two main dimensions. First, while La'O and Tahbaz-Salehi (2019) focused only on optimal policy, this paper discusses the Phillips

²Jones (2013) and Bigio and La'O (2020) also adopted similar models.

curve as well.³ Second, La’O and Tahbaz-Salehi (2019) microfounded price rigidities as arising from incomplete information, and characterized information structures under which monetary policy can replicate the efficient equilibrium. These results complement the analysis in the current paper, which adopts the Calvo model. Our modeling of price rigidities is more restrictive, but it allows us to work with a dynamic multi-period model.

Other studies have also considered New Keynesian economies with heterogeneous firms, focusing on different elements than those emphasized in this paper. Baqaee and Farhi (2020a) introduced downward nominal wage rigidity, and studied its interaction with sector-specific supply and demand shocks. Baqaee, Farhi, and Sangani (2021) instead showed that, when allowing for non-CES demand, monetary shocks have first-order effects on allocative efficiency.

A large empirical literature has documented the shortcomings of consumer price inflation in Phillips curve regressions and forecasting (Orphanides and van Norden (2002), Mavroeidis, Plagborg-Moller, and Stock (2014)), and has sought to construct indicators with better statistical properties (Stock and Watson (1999), Bernanke and Boivin (2003), Stock and Watson (2016)). This paper shows that replacing consumer prices with the divine coincidence index improves the fit of Phillips curve regressions, yielding stable and significant estimates over time and across specifications.

2. ENVIRONMENT

2.1. Consumers

There is an infinitely-lived representative agent, who enjoys consumption (C) and dislikes labor (L). The agent’s per-period preferences are described by the utility function

$$U_t = \frac{C_t^{1-\gamma}}{1-\gamma} - \frac{L_t^{1+\varphi}}{1+\varphi}. \tag{1}$$

There are N industries in the economy, and the consumer has homothetic preferences over their products. The consumption aggregator is $C \equiv \mathcal{C}(Q_1, \dots, Q_N)$, where Q_1, \dots, Q_N denote the quantities consumed of the various goods.

The agent maximizes the present discounted value of per-period utility flows, with discount factor ρ , subject to the budget constraint

$$P_t^C C_t + B_{t+1} \leq W_t L_t + \Pi_t - T_t + (1 + i_t) B_t, \tag{2}$$

where P_t^C is the price index implied by the consumption aggregator C , W_t is the nominal wage, Π_t are firm profits (rebated lump-sum to the household), T_t is a lump-sum tax, B_t is the quantity of risk-free bonds paying off in period t owned by the household, and i_t is the nominal interest rate.

The representative worker is hired by labor unions, who then sell his labor services to other firms. Labor is undifferentiated and fully mobile across industries and labor unions; therefore, all unions pay the same flexible wage to the worker. The labor unions are treated like any other production sector (see Section 2.2), and in particular, there is a

³The modeling of production is also slightly different between the two papers, as this paper allows for arbitrary constant-returns-to-scale production functions, while La’O and Tahbaz-Salehi (2019) imposed Cobb–Douglas. This matters for the welfare cost of relative price distortions across sectors, which is proportional to the elasticities of substitution between inputs in production and between consumption goods.

unit mass of unions, facing nominal rigidities. As a result, there are many wages: the flexible wage W_t paid by the unions to the worker, and the prices charged by the various unions to the rest of the economy. The worker owns all the labor unions, and hence he earns the profit rebates, so that in practice, labor income coincides with the average sticky price charged by the unions. This model of wage rigidity is slightly different from previous ones (Erceg, Henderson, and Levin (2000), Blanchard and Gali (2007)), but it generates the same implications, as Example 2 below illustrates.

The optimal consumption-leisure tradeoff satisfies the first-order condition

$$\frac{W_t}{P_t^C} = C_t^\gamma L_t^\varphi. \quad (3)$$

Intertemporal optimization yields the Euler equation

$$U_{ct} = \rho(1 + i_{t+1})\mathbb{E} \left[U_{ct+1} \frac{P_t^C}{P_{t+1}^C} \right]. \quad (4)$$

2.2. Producers

There are N industries in the economy, indexed by $i, j, k \in \{1, \dots, N\}$, and within each industry, there is a unit mass of firms. Differentiated varieties from all the firms f within an industry i are combined into an industry level output with a CES aggregator (omitting time subscripts for legibility)

$$Y_i = \left(\int_0^1 Y_{if}^{\frac{\epsilon_i-1}{\epsilon_i}} df \right)^{\frac{\epsilon_i}{\epsilon_i-1}}.$$

The implied sectoral price index is

$$P_i = \left(\int_0^1 P_{if}^{1-\epsilon_i} df \right)^{\frac{1}{1-\epsilon_i}}.$$

All firms within an industry have the same constant-returns-to-scale production function, which takes as inputs labor (L_{if}) and intermediate goods X_{ijf} from all industries $j = 1, \dots, N$ (again omitting time subscripts)

$$Y_{if} = A_i F_i(L_{if}, \{X_{ijf}\}_{j=1}^N),$$

where A_i is a Hicks-neutral, industry-specific TFP shifter.⁴ Without loss of generality, we define the functions F_i so that $A_i = 1$ for all sectors i in the initial equilibrium.

Aggregate real output is defined as the consumption aggregate of the representative agent

$$\mathcal{Y}_t = \mathcal{C}(Q_{1t}, \dots, Q_{Nt}).$$

The government pays proportional input subsidies τ_i to the firms, which are constant over time, and set at the industry level so that the profit-maximizing price is equal to

⁴Input-biased productivity shocks can be modeled by introducing an artificial sector which purchases some input j and resells it to sector i , while facing Hicks-neutral productivity shocks.

pre-subsidy marginal costs:

$$1 - \tau_i = \frac{\epsilon_i - 1}{\epsilon_i}.$$

This assumption is standard in the New Keynesian literature (Gali (2015), Woodford (2003)), and it eliminates the incentive for monetary policy to correct static inefficiencies coming from monopoly power.

All firms choose the cost-minimizing combination of inputs, which under constant returns to scale, is the same for all firms within an industry. Industry level marginal costs are given by

$$MC_{it} = \min_{\{X_{ijt}, L_{it}\}} W_t L_{it} + \sum_j P_{jt} X_{ijt} \quad \text{s.t.} \quad A_{it} F_i(L_{it}, \{X_{ijt}\}) = 1. \tag{5}$$

Price rigidities are modeled à la Calvo. At each period, only a fraction δ_i of firms in sector i are allowed to update their price. The optimal reset price P_{it}^* maximizes the present discounted value of profits over all future periods $t + s$, in the event that the firm has not been able to change its price until $t + s$:

$$P_{it}^* = \operatorname{argmax}_{P_{it}} \mathbb{E} \left[\sum_{s=0}^{\infty} \text{SDF}_{t+s} (1 - \delta_i)^s Y_{it+s}(P_{it}) (P_{it} - (1 - \tau_i) MC_{it+s}) \right]. \tag{6}$$

In equation (6), $\text{SDF}_{t+s} = \rho^s \frac{U_{ct+s}}{U_{ct}} \frac{P_t^c}{P_{t+s}^c}$ is the households' stochastic discount factor, and producers of each differentiated variety face the demand function $Y_{ift+s}(P_{ft}) = Y_{it+s} \left(\frac{P_{ft}}{P_{it+s}}\right)^{-\epsilon}$. The solution is given by

$$P_{it}^* = \frac{\mathbb{E} \sum_s [\text{SDF}_{t+s} (1 - \delta_i)^s Y_{ift+s}(P_{it}^*) MC_{it+s}]}{\mathbb{E} \sum_s [\text{SDF}_{t+s} (1 - \delta_i)^s Y_{ift+s}(P_{it}^*)]}. \tag{7}$$

The firms who cannot reset their price keep it unchanged ($P_{ft} = P_{ft-1}$).

2.3. Policy Instruments

The central bank sets the nominal interest rate i_{t+1} at each period t . As explained in Section 2.2, the government also provides input subsidies to firms. Subsidies are financed through lump-sum taxes T_t levied on households (see equation (2)), whose amount is set to balance the government's budget.

2.4. Equilibrium

The equilibrium concept adapts the definition in Baqaee and Farhi (2020b) to account for the endogenous determination of markups given pricing frictions and shocks. Given sectoral markups, all markets must clear; and the evolution of markups must be consistent with Calvo pricing and the realization of sectoral productivities and monetary policy.

DEFINITION 1: At each period t , for given sectoral probabilities of price adjustment δ_i , sectoral productivity shifters A_{it} , and interest rate i_{t+1} , the *general equilibrium* is given by a

vector of firm-level markups \mathcal{M}_{ijt} , a vector of sectoral prices P_{it} , a nominal wage W_t , labor supply L_t , a vector of sectoral outputs Y_{it} , a matrix of intermediate input quantities X_{ijt} , and a vector of final demands C_{it} such that: (i) a fraction δ_i of firms in each sector i charges the profit-maximizing price given by (6); (ii) the markup charged by adjusting firms is given by the ratio of the profit-maximizing price and marginal costs, while the markups of non-adjusting firms are such that their price remains constant; (iii) consumers maximize utility subject to their budget constraint; (iv) producers in each sector i minimize costs and charge the relevant markup; and (v) markets for all goods and labor clear.

REMARK 1: This equilibrium concept nests the standard one with flexible prices, which is obtained as a special case when $\delta_i = 1$ for every sector i .

3. LOG-LINEARIZED MODEL

The log-linearized model is characterized by three equations: the IS equation, the Phillips curve, and the policy rule. Section 3.1 introduces the variables and parameters that characterize the evolution of the log-linearized economy. Section 3.2 then derives some important preliminary results, while Section 3.3 presents the full dynamic system.

3.1. Notation

Table I defines the model variables. Throughout the analysis, we compare three equilibria: the sticky-price economy, the flex-price economy, and a steady state with no changes in productivity and monetary policy, around which we log-linearize. Lowercase letters denote log-differences between the sticky-price equilibrium and the steady state; the subscript nat denotes log differences between the flex-price equilibrium and the steady state; and tildes denote gaps, which means log-differences between the sticky-price and the flex-price equilibrium. Vectors are in bold.

REMARK 2: We solve the model in terms of the aggregate output gap, without keeping track of sector-level output gaps. Section 4 shows that, with constant returns to scale and perfectly mobile labor, sectoral Phillips curves only depend on the aggregate gap. Relative gaps do matter for misallocation and welfare; however, Section 6.3 shows that these can be more conveniently expressed in terms of relative price wedges.

Up to the second order, utility and production functions are fully characterized by equilibrium labor, input, and consumption shares, and by the relevant Allen elasticities of substitution. Table II introduces these input-output parameters, and defines the Leontief inverse and Domar weights. The table also introduces notation for the price adjustment

TABLE I
MODEL VARIABLES.

Aggregate output gap	$\tilde{y}_t \equiv y_t - y_{t,\text{nat}}$
Sectoral inflation rates	$\boldsymbol{\pi}_t = (\pi_{1t} \dots \pi_{Nt})^T, \pi_{it} \equiv p_{it} - p_{it-1}$
Sectoral markups	$\boldsymbol{\mu}_t = (\mu_{1t} \dots \mu_{Nt})^T, \mu_{it} \equiv p_{it} - mc_{it}$
Sectoral productivity changes	$\log \mathbf{A}_t = (\log A_{1t} \dots \log A_{Nt})^T$

TABLE II
INPUT-OUTPUT DEFINITIONS.

Consumption shares	$\beta \in \mathbb{R}^N, \beta_i = \frac{P_i C_i}{P^C}$
Labor shares	$\alpha \in \mathbb{R}^N, \alpha_i = \frac{W L_i}{M C_i Y_i}$
Input-output matrix	$\Omega \in \mathbb{R}^N \times \mathbb{R}^N, \omega_{ij} = \frac{P_j x_{ij}}{M C_i Y_i}$
Leontief inverse	$(I - \Omega)^{-1}$
Domar weights	$\lambda^T = \beta^T (I - \Omega)^{-1}$
Substitution elasticities	$\{\sigma_{ij}^C\}$ consumption $\{\theta_{jk}^i, \theta_{jL}^i\}$ between inputs in sector i 's production
Price adjustment parameters	ϵ_i between varieties $\Delta = \text{diag}(\hat{\delta}_1, \dots, \hat{\delta}_N)$

parameters. Note that in the matrix Δ the primitive Calvo parameters δ_i are replaced with the increasing and convex function

$$\hat{\delta}_i(\delta_i, \rho) \equiv \frac{\delta_i(1 - \rho(1 - \delta_i))}{1 - \rho\delta_i(1 - \delta_i)},$$

which appears in the derivation of the Phillips curve.

The elements Ω_{ij} of the input-output matrix encode the direct elasticity of sector i 's cost to a change in sector j 's price. The Leontief inverse, instead, captures the total elasticity, directly and indirectly through intermediate inputs. To see this, one can write an expansion of the Leontief inverse where the n th term is the exposure of i to j through paths of length n :

$$(I - \Omega)_{ij}^{-1} = I_{ij} + \Omega_{ij} + \Omega_{ij}^2 + \dots$$

In a similar way, while the direct elasticity of the consumer price index P^C to changes in the price of the various goods is given by the consumption shares β , the total elasticity is given by the product of the consumption shares times the Leontief inverse, which is equal to the Domar weights λ , and it also coincides with the sectors' total sales as a fraction of GDP.⁵ The Leontief inverse also maps direct labor shares in sectoral value added, α ,⁶ into total shares, $[(I - \Omega)^{-1}\alpha]$, which account for the labor content of a sector's suppliers, its suppliers' suppliers, etc.

REMARK 3: With constant returns to scale and labor being the only factor of production, labor must account for the entirety of value added. Hence, it holds that $[(I - \Omega)^{-1}\alpha]_i = 1$ for every sector i .⁷ Furthermore, the total share of labor in aggregate value added is also $\lambda^T \alpha = 1$.

⁵Sectoral sales shares $s_i \equiv \frac{P_i Y_i}{\sum_j P_j C_j}$ satisfy $s_i = \beta_i + \sum_j s_j \Omega_{ji}$. Rearranging this equation implies $s^T = \lambda^T = \beta^T (I - \Omega)^{-1}$.

⁶In the flex-price equilibrium, producers charge non-unit markups on post-subsidy marginal costs, because of monopolistic competition and CES demand. However, subsidies are set so that prices equal pre-subsidy marginal costs. This effectively eliminates markups, and guarantees that α indeed corresponds to the labor share in value added.

⁷To see this, note that labor and intermediate input shares must sum to 1 in all sectors: $\Omega \mathbf{1} + \alpha = \mathbf{1}$. Rearranging this equation implies $(I - \Omega)^{-1} \alpha = \mathbf{1}$.

REMARK 4: Our modeling of wage rigidity via the labor unions imposes a particular structure on α and Ω . Only the union hires workers directly; therefore, it is the only industry with non-zero labor share ($\alpha_i = 0 \forall i \neq \text{union}$). Moreover, since labor is its only input, the union has unit labor share ($\alpha_{\text{union}} = 1, \Omega_{\text{union},j} = 0 \forall j$). For all other sectors, instead, the direct labor share (as measured in the national accounts) is encoded in $\Omega_{i,\text{union}}$. When wages are flexible ($\delta_{\text{union}} = 1$), this definition of the input shares α and Ω is equivalent to eliminating the union sector and setting $\alpha_i = \Omega_{i,\text{union}} \forall i$.⁸

3.2. Basic Results

This section states three basic results: Lemma 1 extends Hulten’s theorem to an economy with elastic labor supply; Lemma 2 establishes that the output gap and the employment gap always coincide; and Lemma 3 shows that the output gap is proportional to the sales-weighted sum of sectoral markups.⁹

LEMMA 1: *In the log-linearized model, the natural employment and output only depend on aggregate (sales-weighted) productivity, and are given by*

$$\begin{aligned} l_{t,\text{nat}} &= \frac{1 - \gamma}{\gamma + \varphi} \boldsymbol{\lambda}^T \log \mathbf{A}_t, \\ y_{t,\text{nat}} &= \frac{1 + \varphi}{\gamma + \varphi} \boldsymbol{\lambda}^T \log \mathbf{A}_t. \end{aligned} \tag{8}$$

PROOF: Supplement Section 2.1.

Q.E.D.

Lemma 1 tells us that natural employment and output in the network model are the same function of aggregate productivity as in the one-sector model, where aggregate productivity is the sales-weighted average of sectoral productivities. Lemma 2 then relates the change in productivity in the flex-price and sticky-price equilibria, up to a first-order approximation.

LEMMA 2: *Around an undistorted steady state, the first-order change in aggregate labor productivity in the economy with nominal rigidities is the same as in the efficient economy with flexible prices:*

$$y_t - l_t = y_{t,\text{nat}} - l_{t,\text{nat}} = \boldsymbol{\lambda}^T \log \mathbf{A}_t. \tag{9}$$

Lemma 2 follows immediately from the envelope theorem. In the flex-price equilibrium, labor is optimally allocated within and across sectors, and small deviations from the optimal allocation have no first-order effect on aggregate productivity.

REMARK 5: An important consequence of Lemma 2 is that the output and employment gap always coincide. Rearranging equation (9) immediately yields

$$\tilde{y}_t \equiv y_t - y_{t,\text{nat}} = l_t - l_{t,\text{nat}} \equiv \tilde{l}_t.$$

⁸Note, however, that eliminating the union sector would slightly complicate the notation for the welfare loss from cross-sector misallocation in Section 6.3.

⁹Baqee and Farhi (2020b) and Bigio and La’O (2020) derived similar results in terms of the aggregate labor share (in place of the employment gap).

Finally, Lemma 3 relates the output gap with the sales-weighted sum of sectoral markups.

LEMMA 3: *The output gap is proportional to the sales-weighted aggregate markup:*

$$(\gamma + \varphi)\tilde{y}_t = - \sum_i \lambda_i \mu_{it}. \tag{10}$$

PROOF: Supplement Section 2.1.

Q.E.D.

The intuition for Lemma 3 is as follows: when markups are high (low), the share of revenues that is paid out to workers through wages is small (large). This reduces (increases) labor supplied, resulting in a negative (positive) output gap. The contribution of each sector to the aggregate output gap is proportional to its sale share.

3.3. Equilibrium Conditions

The evolution of the economy is characterized by three equations: the IS curve, the Phillips curve, and the policy rule. This section derives the IS curve, states the policy rule, and lays the foundations for the Phillips curve, which is derived in Section 4.

IS Curve

Log-linearizing the Euler equation (4) and applying the definition of output gap yields the IS curve:

$$\tilde{y}_t = \mathbb{E}[\tilde{y}_{t+1}] - \frac{1}{\gamma}(i_{t+1} - \mathbb{E}[\pi_{t+1}^C] - r_{t+1, \text{nat}}). \tag{11}$$

The natural interest rate $r_{t+1, \text{nat}}$ is given by $r_{t+1, \text{nat}} = -\log \rho + \gamma \lambda^T \mathbb{E}(\log \mathbf{A}_{t+1} - \log \mathbf{A}_t)$, which only depends on the discount factor ρ and on changes in aggregate productivity.

Policy Rule

The central bank sets the nominal interest rate i_{t+1} following a Taylor rule, which targets an aggregate inflation index $\bar{\pi}_t$:¹⁰

$$i_{t+1} = r_{t+1}^{\text{nat}} + \zeta \bar{\pi}_t. \tag{12}$$

The parameter ζ governs the responsiveness to the target.

Marginal Costs and Prices

The key building blocks towards deriving sectoral Phillips curves are the firms' cost minimization problem (5), the optimal reset price condition (7), and the agent's consumption-leisure tradeoff (3). Different from the one-sector benchmark, the multi-sector economy has $N - 1$ state variables, given by lagged sectoral relative prices. Hence, their law of motion is part of the Phillips curve block, and it is given by

$$\mathbf{p}_t = \mathbf{p}_{t-1} + \boldsymbol{\pi}_t. \tag{13}$$

¹⁰As long as there are no exogenous shocks to the firms' desired markups, we can omit the output gap from the policy target without loss of generality, because Proposition 1 shows that it can be written as a function of the divine coincidence inflation index.

Log-linearizing the firms' optimal reset price condition (7) yields the law of motion for sectoral inflation rates:¹¹

$$\pi_t = \Delta(\mathbf{m}c_t - \mathbf{p}_{t-1}) + \rho(I - \Delta)\mathbb{E}\pi_{t+1}. \quad (14)$$

Equation (14) is the same as in the one-sector model (Gali (2015), Woodford (2003)). The evolution of marginal costs, however, crucially depends on the production network, and it follows immediately from the firms' cost minimization problem (5):

$$\mathbf{m}c_t = \alpha w_t + \Omega \mathbf{p}_t - \log \mathbf{A}_t. \quad (15)$$

Sector i 's marginal cost responds to nominal wage inflation π_t^w in proportion to its labor share α_i , and to input prices, in proportion to the input shares $\{\Omega_{ij}\}_{j=1}^N$. Marginal costs also decrease with productivity $\log \mathbf{A}_{it}$.

Plugging (15) into (14) yields the following law of motion for sectoral inflation rates:

$$(I - \Delta\Omega)\pi_t = \Delta(\alpha w_t - (I - \Omega)\mathbf{p}_{t-1} - \log \mathbf{A}_t) + \rho(I - \Delta)\mathbb{E}\pi_{t+1}. \quad (16)$$

Note that, in an efficient economy where prices equal marginal costs ($\mathbf{p}_t = \mathbf{m}c_t$) and nominal wages remain constant ($w_t = 0$), equation (15) implies

$$(I - \Omega)\mathbf{p}_t + \log \mathbf{A}_t = \mathbf{0}. \quad (17)$$

That is, sectoral prices move proportionately to their network-adjusted productivities. Hence, in equation (16), current prices adjust to offset the wedge $(I - \Omega)\mathbf{p}_{t-1} + \log \mathbf{A}_t$ between the prices inherited from the previous period and the efficient relative prices, given current sectoral productivities.

Log-linearizing the optimal consumption-leisure tradeoff (3), and using Lemmas 1 and 2, allows us to connect real wages ($w_t - \beta^T \mathbf{p}_t$) with the output gap and aggregate productivity:¹²

$$w_t - \beta^T \mathbf{p}_t = (\gamma + \varphi)\tilde{y}_t + \lambda^T \log \mathbf{A}_t. \quad (18)$$

Equation (18) tells us that, when aggregate output is at the efficient level ($\tilde{y} = 0$), real wages change proportionately to the aggregate productivity $\lambda^T \log \mathbf{A}_t$, which—up to the first order—is also equal to the marginal product of labor.

Equations (16), (13), and (18), together with the IS curve (11) and the policy rule (12), pin down the dynamics of the economy. Section 4.2 builds on equations (16) and (18), to construct a multi-sector counterpart of the structural Phillips curve from the one-sector model.

4. THE PHILLIPS CURVE

The Phillips curve is one of the three key equations that govern the dynamic evolution of the economy, together with the IS curve (11) and the policy rule (12). In the one-sector model, it is given by

$$\pi_t = \rho\mathbb{E}\pi_{t+1} + \kappa\tilde{y}_t + u_t. \quad (19)$$

¹¹The derivation is the same as in the one-sector model (Gali (2015)). See Section 2.1 of the Supplement for the details.

¹²See Supplement Section 2.1 for the full derivation.

In equation (19), current inflation π_t is increasing in expected future inflation $\mathbb{E}\pi_{t+1}$, in the current output gap \tilde{y}_t , and in a time-varying residual u_t , which is often referred to as a cost-push shock. This section revisits the slope κ , the residual u_t , and the relation between current and future inflation in the multi-sector economy.

Whenever the residual is $u_t \neq 0$, inflation is not zero even if the central bank closes the output gap in all periods ($\tilde{y}_s \equiv 0 \forall s \geq t$), thereby creating a tradeoff between stabilizing output and stabilizing inflation. A well-known result, commonly referred to as the divine coincidence, states that, in the benchmark one-sector model, productivity shocks do not affect the residual u_t (Gali (2015), Woodford (2003)).

In the multi-sector economy, one can construct inflation rates and Phillips curves for every sector, and then aggregate them in different ways. Section 4.1 shows that, in sharp contrast with the one-sector benchmark, changes in productivity affect the residual u_t , thereby generating a tradeoff between closing the output gap and eliminating inflation. This happens both at the sector level and in the aggregate, except when inflation is measured according to a novel index, called the divine coincidence index. It is not surprising that inflation cannot be stabilized sector-by-sector, because it is efficient for relative prices to change in response to asymmetric productivity shocks. However, it is surprising that there is a tradeoff between closing the output gap and stabilizing almost all measures of aggregate inflation, including consumer prices.

Section 4.2 then derives the slopes and residuals of all sectoral Phillips curves, while Section 4.3 shows how they can be combined using different weights to obtain aggregate Phillips curves. These weights and the input-output structure determine the slope and residual of the aggregate curves. Since prices across sectors are differentially affected by changes in the output gap, the slope of the Phillips curve depends on the way inflation is measured. Nonetheless, Proposition 2 and Corollary 2 show that all Phillips curves become flatter in the presence of multiple sectors and intermediate inputs. The two propositions also show that different inflation measures face a different inflation-output tradeoff, whose sign depends on their differential exposure to changes in productivity versus changes in wages. Moreover, in the multi-sector economy, current inflation depends on expected inflation in all sectors, and not just on the expected future value of the reference index.

Section 4.4 illustrates these results in simple input-output networks.

4.1. The Divine Coincidence Index

Definition 2 introduces the divine coincidence index.

DEFINITION 2: Assume that no sector has fully rigid prices ($\delta_i \neq 0 \forall i$). The *divine coincidence index* weights sectors according to their sales shares, and discounts those with more flexible prices. It is defined as

$$\pi_t^{DC} \equiv \sum_i \frac{\lambda_i \frac{1 - \hat{\delta}_i}{\hat{\delta}_i}}{\sum_j \lambda_j \frac{1 - \hat{\delta}_j}{\hat{\delta}_j}} \pi_{it}.$$

Proposition 1 illustrates the properties of this index.

PROPOSITION 1: *The divine coincidence Phillips curve is given by*

$$\pi_t^{DC} = \rho \mathbb{E} \pi_{t+1}^{DC} + \frac{\gamma + \varphi}{\sum_j \lambda_j \frac{1 - \hat{\delta}_j}{\hat{\delta}_j}} \tilde{y}_t. \quad (20)$$

Notably, current divine coincidence inflation only depends on future divine coincidence inflation, proportionately to the discount factor ρ , and it is not affected by productivity fluctuations. If at least one sector has sticky prices ($\Delta \neq 1$), the divine coincidence index is the unique inflation index with this property.

PROOF: The proposition builds on two results: Lemma 3, and the relation between markups and inflation rates

$$\mu_{it} = -\frac{1 - \hat{\delta}_j}{\hat{\delta}_j} [\pi_{it} - \rho \mathbb{E} \pi_{it+1}], \quad (21)$$

which follows immediately from equation (14). Substituting equation (21) into equation (10) from Lemma 3—which shows that the output gap is proportional to the sales-weighted sum of sectoral markups—yields the divine coincidence Phillips curve. Uniqueness is proved in Section 2.2 of the Supplement. *Q.E.D.*

Lemma 3 is the key insight behind the divine coincidence result: the output gap is proportional to the sales-weighted aggregate markup, so that high (low) markups correspond to low (high) real wages and a negative (positive) output gap. Furthermore, the Calvo assumption allows us to connect markups with inflation, as equation (21) shows. Intuitively, when the firms that are allowed to reset prices are randomly selected, equation (14) tells us that a fraction $\hat{\delta}_i$ of the change in industry i 's marginal cost is transmitted into current inflation (net of inflation expectations). The remaining fraction $1 - \hat{\delta}_i$ is instead absorbed into the industry's average markup. Hence, for given observed inflation, the implied markup change is decreasing in the probability of price adjustment δ_i .

REMARK 6: Note that the slope of the divine coincidence Phillips curve depends on the input-output network only through the sectoral sales shares λ . Like in the one-sector model, the slope is increasing in the elasticity of labor supply ($\gamma + \varphi$), and decreasing in the amount of price stickiness, as measured by the sectoral parameters $\frac{1 - \hat{\delta}_i}{\hat{\delta}_i}$. Moreover, the divine coincidence Phillips curve is flatter when there are more intermediate inputs, because the sales shares λ_j in the denominator of (20) are larger. In the special case where all sectors have the same frequency of price adjustment δ , the denominator becomes $\frac{1 - \hat{\delta}}{\hat{\delta}} \sum_i \lambda_i$, which is proportional to the aggregate sales-to-GDP ratio $\sum_i \lambda_i$.¹³

¹³In the model and calibration, the ratio of sales to GDP is the same as sales over payments to primary factors (i.e., labor). This is because the profits generated by markups are canceled out by the taxes needed to pay for input subsidies. In the data, however, sales over payments to labor are not the same thing as sales over GDP, and they have not changed in the same way over time (sales over payments to labor have increased, sales over GDP are roughly stable).

4.2. Sector-Level Phillips Curves

Proposition 2 derives sector-level Phillips curves as a function of disaggregated input shares and price adjustment probabilities.

PROPOSITION 2: *Sector-level Phillips curves are given by*

$$\boldsymbol{\pi}_t = \rho(I - \mathcal{V})\mathbb{E}\boldsymbol{\pi}_{t+1} + \mathcal{B}(\gamma + \varphi)\tilde{y}_t - \mathcal{V}\boldsymbol{\chi}_t, \tag{22}$$

where

$$\begin{aligned} \mathcal{B} &\equiv \frac{\Delta(I - \Omega\Delta)^{-1}\boldsymbol{\alpha}}{1 - \boldsymbol{\beta}^T\Delta(I - \Omega\Delta)^{-1}\boldsymbol{\alpha}}, \\ \mathcal{V} &\equiv [\Delta(I - \Omega\Delta)^{-1} - \mathcal{B}(\lambda^T - \boldsymbol{\beta}^T\Delta(I - \Omega\Delta)^{-1})](I - \Omega), \\ \boldsymbol{\chi}_t &\equiv \mathbf{p}_{t-1} + (I - \Omega)^{-1}\log \mathbf{A}_t. \end{aligned}$$

The matrix \mathcal{V} is such that $\sum_j \mathcal{V}_{ij} = 0$, $(I - \mathcal{V})_{ij} \in [0, 1]$, and, as long as no sector has fully flexible prices ($\delta_i < 1 \forall i$), $I - \mathcal{V}$ is invertible.

PROOF: Equation (22) is obtained by combining the law of motion for prices (16) with the relation between output gap and real wages in (18). To make real wages appear in (16), subtract $\Delta\boldsymbol{\alpha}\boldsymbol{\beta}^T\boldsymbol{\pi}_t$ from both sides of the equation to obtain

$$\begin{aligned} (I - \Delta\Omega - \Delta\boldsymbol{\alpha}\boldsymbol{\beta}^T)\boldsymbol{\pi}_t \\ = \Delta(\boldsymbol{\alpha}(w_t - \boldsymbol{\beta}^T\mathbf{p}_t) - (I - \Omega - \boldsymbol{\alpha}\boldsymbol{\beta}^T)\mathbf{p}_{t-1} - \log \mathbf{A}_t) + \rho(I - \Delta)\mathbb{E}\boldsymbol{\pi}_{t+1}. \end{aligned}$$

Plugging in equation (18) allows us to express real wages as a function of the output gap and aggregate productivity, and inverting the coefficient on inflation on the left-hand side yields the result.¹⁴ The properties of the matrix \mathcal{V} are derived in the Supplement, Section 2.2. Q.E.D.

Equation (22) generalizes the Phillips curve in (19), to express the slope, residual, and expectation terms explicitly as a function of the primitives of the economy (price adjustment probabilities, input shares, productivity, and labor supply elasticities).

All of these terms depend on the rigidity-adjusted Leontief inverse $(I - \Omega\Delta)^{-1}$. Similarly to the usual Leontief inverse, this matrix governs the propagation of cost shocks through the network. The (i, j) th element corresponds to the elasticity of sector i 's marginal cost with respect to sector j 's marginal cost, directly, and indirectly through the input-output network. The importance of each sector depends on its input shares Ω and on its price flexibility Δ . Sticky-price sectors only partially transmit cost shocks along the production chain, and therefore they are discounted. Remark 7 formally states that the transmission of shocks is dampened when prices are stickier.

¹⁴Note that

$$I - \Delta\Omega - \Delta\boldsymbol{\alpha}\boldsymbol{\beta}^T = (I - \Delta\Omega)(I - \Delta(I - \Omega\Delta)^{-1}\boldsymbol{\alpha}\boldsymbol{\beta}^T)$$

and

$$(I - \Delta(I - \Omega\Delta)^{-1}\boldsymbol{\alpha}\boldsymbol{\beta}^T)^{-1} = I + \frac{\Delta(I - \Omega\Delta)^{-1}\boldsymbol{\alpha}\boldsymbol{\beta}^T}{1 - \boldsymbol{\beta}^T\Delta(I - \Omega\Delta)^{-1}\boldsymbol{\alpha}}.$$

REMARK 7: Element-by-element, the matrix $(I - \Omega\Delta)^{-1}$ is increasing in the price adjustment probabilities $\{\delta_i\}_{i=1}^N$.¹⁵

Let us now analyze the slope, residual, and expectation terms in detail. The slope of sectoral Phillips curves corresponds to the elasticity of sector-level prices with respect to the output gap, and it is given by the vector $\mathcal{B}(\gamma + \varphi) \in \mathbb{R}^{N \times 1}$. This vector has three components: the elasticity of labor supply $\gamma + \varphi$, the pass-through of nominal wages into prices $\Delta(I - \Omega\Delta)^{-1}\alpha$, and an equilibrium multiplier $\frac{1}{1 - \beta^T \Delta(I - \Omega\Delta)^{-1}\alpha}$. Let us discuss each in order. The elasticity $\gamma + \varphi$ is the slope of the labor supply curve, which tells us by how much real wages must increase to sustain a 1% increase in employment. Intuitively, as output increases, labor demand also increases, and workers must be compensated with higher real wages.

In turn, higher wages raise marginal costs and prices. The pass-through of nominal wages into prices, $\Delta(I - \Omega\Delta)^{-1}\alpha$, combines the direct effect of wages on marginal cost, given by the labor shares α , with the indirect effect through the input-output network, obtained by multiplying the direct effect times the rigidity-adjusted Leontief inverse. The price adjustment probabilities Δ then give us the response of prices to changes in costs. Since prices increase together with wages, real wages increase by less than nominal wages. Specifically, in order to raise real wages by 1%, nominal wages must increase by $\frac{1}{1 - \beta^T \Delta(I - \Omega\Delta)^{-1}\alpha} > 1$, to compensate for the response of consumer prices.

Following Remark 7, all sectoral Phillips curves are flatter in economies with stickier prices or larger intermediate input shares. Corollary 1 formalizes this result, and Example 2 below provides a concrete illustration.

COROLLARY 1: *Component-by-component, \mathcal{B} is weakly increasing in each price adjustment probability $\{\delta_i\}_{i=1}^N$. Moreover, a fall in a sector's labor share, compensated by a uniform increase in all its input shares, weakly reduces all the components of \mathcal{B} . In particular, all sectoral Phillips curves are flatter than in an economy with no intermediate inputs ($\Omega = \mathbb{O}$, $\alpha = \mathbf{1}$):*

$$\frac{\Delta(I - \Omega\Delta)^{-1}\alpha}{1 - \beta^T \Delta(I - \Omega\Delta)^{-1}\alpha} \leq \frac{\Delta\mathbf{1}}{1 - \mathbb{E}_\beta(\delta)},$$

where $\mathbb{E}_\beta(\hat{\delta})$ denotes a weighted average of the parameters $\hat{\delta}$ using the consumption shares β as weights.

Let us now go back to equation (22), and let us focus on the residual $-\mathcal{V}\chi_t$. Recall from equation (17) that the efficient relative prices equal network-adjusted productivities $(I - \Omega)^{-1} \log \mathbf{A}_t$. Hence, the vector χ_t is the wedge between lagged prices \mathbf{p}_{t-1} and efficient relative prices. The elements \mathcal{V}_{ij} of the matrix $\mathcal{V} \in \mathbb{R}^{N \times N}$ describe the effect of a wedge in sector j on the price of sector i .

REMARK 8: The wedges χ_t affect sectoral inflation only to the extent that they are not proportional across sectors ($\chi_t \neq \mathbf{1}$). This follows immediately from the result that

¹⁵To see this, note that the adjusted Leontief inverse can be written as

$$(I - \Omega\Delta)^{-1} = I + \Omega\Delta + (\Omega\Delta)^2 + \dots$$

All the elements of this sum are (weakly) increasing in Δ ; therefore, we have $\frac{\partial (I - \Omega\Delta)^{-1}}{\partial \delta_l} \geq 0 \forall i, j, l$.

$\sum_j \mathcal{V}_{ij} = 0$. When initial prices are as in the steady state ($\mathbf{p}_{t-1} = \mathbf{0}$), we have $\boldsymbol{\chi}_t \propto \mathbf{1} \iff \log \mathbf{A}_t \propto \boldsymbol{\alpha}$. That is, the only change in productivity which does not generate cost-push shocks on sectoral inflation is the labor-augmenting one.

For simplicity, let us assume that $\mathbf{p}_{t-1} = \mathbf{0}$, and focus on the effect of productivity changes. The intuition for the general case is the same. Productivity has two offsetting effects on inflation. First, through a direct effect described by the term $-\Delta(I - \Omega\Delta)^{-1}$, marginal costs and prices are decreasing in sectoral productivities. Second, in a counterbalancing mechanism described by the term $\mathcal{B}(\boldsymbol{\lambda}^T - \boldsymbol{\beta}^T \Delta(I - \Omega\Delta)^{-1})$, nominal wages are also increasing in productivity, with an offsetting effect on marginal costs. In fact, as stated in Lemma 1, to maintain output at the efficient level ($\tilde{y} = 0$), real wages must increase by the same amount as aggregate productivity $\boldsymbol{\lambda}^T \log \mathbf{A}_t$. This can be achieved via an increase in nominal wages, or via a decline in consumer prices. From the direct effect, we know that consumer prices change by $-\boldsymbol{\beta}^T \Delta(I - \Omega\Delta)^{-1} \log \mathbf{A}_t$. Hence, nominal wages need to increase by $(\boldsymbol{\lambda}^T - \boldsymbol{\beta}^T \Delta(I - \Omega\Delta)^{-1}) \log \mathbf{A}_t$. Like for the slope, the effect of wages on sectoral prices is given by the vector \mathcal{B} . Overall, prices fall in the sectors that benefit more from higher productivity than they suffer from higher wages, and vice versa. Formally, an increase in the productivity of sector j reduces sector i 's price ($\mathcal{V}_{ij} > 0$) if and only if i 's marginal cost is more exposed to a 1% increase in j 's price than to an increase in the real wage proportional to j 's sale share.

Finally, like in the one-sector model, in equation (22) current inflation depends on future expected inflation, because producers preemptively raise their current price when they expect a future price increase. In the multi-sector model, however, a price increase in one sector propagates to all the others, as captured by the matrix \mathcal{V} , so that current inflation in one sector depends on expected inflation in all other sectors.

4.3. Aggregation

Building on Proposition 2, Corollary 2 expresses the slope and residual of any aggregate Phillips curve as a function of the model primitives. Formally, an aggregate price index $\bar{\pi}_t^\phi$ is an average of sectoral inflation rates, according to some weights $\boldsymbol{\phi}^T = (\phi_1, \dots, \phi_N)$. The slope and residual of the Phillips curve are also an average of sectoral ones, with the same weights $\boldsymbol{\phi}$, and can be expressed as a function of the aggregation operator $\boldsymbol{\delta}^\phi$.

DEFINITION 3: The aggregation operator $\boldsymbol{\delta}^\phi$ is defined as

$$\boldsymbol{\delta}^\phi \equiv \boldsymbol{\phi}^T \Delta(I - \Omega\Delta)^{-1}.$$

The operator $\boldsymbol{\delta}^\phi \mathbf{x}$ describes the propagation of a shock to sectoral marginal costs parallel to the vector \mathbf{x} into the price index defined by the weights $\boldsymbol{\phi}$. As usual, the rigidity-adjusted Leontief inverse describes the propagation of the shock through the network, while the price adjustment probabilities Δ govern the price response. Finally, sectoral prices are aggregated according to the given weights $\boldsymbol{\phi}$.

COROLLARY 2: Given a weighting $\boldsymbol{\phi}$ of sectoral inflation rates, the corresponding aggregate Phillips curve is given by

$$\bar{\pi}_t^\phi = \rho \boldsymbol{\phi}^T (I - \mathcal{V}) \mathbb{E} \boldsymbol{\pi}_{t+1} + \kappa^\phi (\gamma + \varphi) \tilde{y}_t + u_t^\phi,$$

where the slope and residual are

$$\begin{aligned} \kappa^\phi &= \frac{\delta^\phi \alpha}{1 - \delta^\beta \alpha}, \\ u_t^\phi &= \left(\delta^\phi \alpha \frac{\lambda^T - \delta^\beta}{1 - \delta^\beta \alpha} - \delta^\phi \right) (I - \Omega) \chi_t. \end{aligned} \tag{23}$$

Corollary 2 expresses the parameters of the structural Phillips curve in equation (19) as a function of micro level primitives. Following Corollary 1, the slope κ^ϕ is flatter in economies with larger intermediate input flows, and in particular, it is always weakly smaller than in an economy with the same consumption shares but no intermediate inputs:

$$\kappa^\phi \leq \frac{\mathbb{E}_\phi \Delta}{1 - \mathbb{E}_\beta(\Delta)}.$$

The consumer price Phillips curve is an important special case. Its slope only depends on the pass-through of wages into consumer prices $\delta^\beta \alpha$, and it is given by

$$\kappa^C = \frac{\delta^\beta \alpha}{1 - \delta^\beta \alpha} (\gamma + \varphi). \tag{24}$$

Corollary 1 implies that $\delta^\beta \alpha \leq \mathbb{E}_\beta(\Delta)$; therefore, the consumer price Phillips curve is flatter than in an economy with no intermediate inputs.

Unlike in the one-sector model, the residual u_t^ϕ is no longer independent of productivity, and changes in productivity generate a tradeoff between closing the output gap and stabilizing a generic inflation index π_t^ϕ . Following the same reasoning as in Section 4.2, the behavior of $\bar{\pi}_t^\phi$ depends on the counterbalancing effects of productivity and wages. When $\mathbf{p}_{t-1} = \mathbf{0}$, inflation falls after a positive shock ($\lambda^T \log \mathbf{A}_t > 0$) if and only if the pass-through $\delta^\phi \log \mathbf{A}_t$ of productivity into the price index $\bar{\pi}_t^\phi$ is larger than the pass-through of the counterbalancing wage change, $\delta^\phi \alpha \frac{\lambda^T - \delta^\beta}{1 - \delta^\beta \alpha} \log \mathbf{A}_t$. This is the case when flex-price sectors, or sectors with higher weight ϕ_i , benefit more from the productivity change than they lose from the increase in wages.

This is also true of consumer price inflation. The residual in the consumer price Phillips curve is given by

$$u_t^C = \frac{\delta^\beta \alpha \lambda^T - \delta^\beta}{1 - \delta^\beta \alpha} (I - \Omega) \chi_t. \tag{25}$$

For $\mathbf{p}_{t-1} = \mathbf{0}$, consumer inflation falls after a positive productivity shock $\lambda^T \log \mathbf{A}_t > 0$ if and only if the wage pass-through $\delta^\beta \alpha \lambda^T \log \mathbf{A}_t$ is smaller than the productivity pass-through $\delta^\beta \log \mathbf{A}_t$. This happens when downstream sectors—which have larger consumption shares—or sectors with flexible prices receive a better shock than the average.

Also note that current consumer inflation does not depend only on future consumer inflation, but it also depends on the consumption-weighted price wedges

$$\beta^T \mathcal{V} \mathbb{E} \pi_{t+1} = \frac{\delta^\beta \alpha \lambda^T - \delta^\beta}{1 - \delta^\beta \alpha} (I - \Omega) \mathbb{E} \pi_{t+1}$$

created by the producers' preemptive price adjustment.

4.4. Examples

This section illustrates how the slope and residual of the Phillips curves depend on the input-output network in four simple economies. To facilitate the intuition, only for this section and Section 5.4 we assume $\mathbf{p}_{t-1} = \mathbf{0}$, and, unless otherwise noted, we let α denote the sectoral labor shares, rather than encoding them into $\Omega_{i,\text{union}}$.

EXAMPLE 1—Roundabout Economy: Based on a simple yet powerful model of input-output linkages, the roundabout economy, Basu (1995), Carvalho (2006), Nakamura and Steinsson (2010), Cagliarini, Robinson, and Tran (2011), Christiano (2016) pointed out that intermediate inputs generate real rigidities and increase monetary non-neutrality. This example revisits the roundabout model as a special case, and evaluates whether it is a good approximation of a generic input-output framework. As the example illustrates, one would need to know the full input-output structure in order to evaluate the accuracy of calibrations based on the roundabout economy. Nonetheless, it is fair to say that the roundabout model usually provides a good approximation of the slope of the Phillips curve, while it is ill-suited to predict the propagation of sectoral productivity shocks, and to study welfare (see Remark 9).

Consider a Cobb–Douglas roundabout economy, where all sectors $i = 1, \dots, N$ have the same price adjustment probability δ , and they take as inputs labor, with the same share $\bar{\alpha}$,¹⁶ and the same bundle of intermediate goods. The sectors j have heterogeneous shares γ_j in the input bundle, which can be calibrated to replicate observed sectoral sales shares:

$$Y_i = L_i \bar{\alpha} X_i^{1-\bar{\alpha}}, \quad X = \prod_j X_j^{\gamma_j}, \quad \sum_j \gamma_j = 1.$$

In the roundabout economy, the slope of the consumer-price Phillips curve is given by

$$\kappa_{\text{round}} = \frac{\delta_{\text{round}}^\beta \bar{\alpha}}{1 - \delta_{\text{round}}^\beta \bar{\alpha}} (\gamma + \varphi) = \bar{\alpha} \frac{\bar{\delta}}{1 - \bar{\delta}} (\gamma + \varphi). \tag{26}$$

In equation (26), the labor share and the price rigidity parameter $\bar{\delta}$ are calibrated to their (consumption-weighted) averages: $\bar{\delta} = \mathbb{E}_\beta(\Delta)$, $\bar{\alpha} = \mathbb{E}_\beta(\alpha)$. This calibration provides a good approximation of the actual slope, as long as price adjustment probabilities do not differ systematically between intermediate and final producers, or between sectors that use sticky-price versus flex-price intermediate inputs.

Specifically, the roundabout model overestimates κ^{CPI} whenever $\delta_{\text{round}}^\beta \bar{\alpha} > \delta^\beta \alpha$. This is the case if (i) intermediate good producers have stickier prices than final good producers, (ii) producers with more flexible prices use stickier intermediate inputs, or (iii) sectors that use fewer intermediate inputs have stickier prices. In fact, it holds that

$$\delta_{\text{round}}^\beta \bar{\alpha} = \bar{\delta} \bar{\alpha} \mathbb{E}_\beta((I - \bar{\delta} \Omega_{\text{round}})^{-1} \mathbf{1}),$$

whereas

$$\delta_{\text{true}}^\beta \alpha = \mathbb{E}_\beta(\Delta(I - \Omega \Delta)^{-1} \alpha).$$

Calibrating $\bar{\delta} = \mathbb{E}_\beta(\Delta)$, we have $\kappa_{\text{round}} > \kappa_{\text{true}}$ when $\text{Cov}_\beta(\Delta, (I - \Omega \Delta)^{-1} \mathbf{1}) < 0$, that is, when producers with more flexible prices use stickier intermediate inputs. We also have

¹⁶Note that, in this economy, $\bar{\alpha}$ also coincides with the inverse of the sales-to-GDP ratio: $\bar{\alpha} = \frac{1}{\lambda^T \mathbf{1}}$.

$\kappa_{\text{round}} > \kappa_{\text{true}}$ when the roundabout model overestimates the pass-through of wages into marginal costs ($\bar{\alpha} \mathbb{E}_\beta((I - \bar{\delta} \Omega_{\text{round}})^{-1} \mathbf{1}) > \mathbb{E}_\beta((I - \Omega \Delta)^{-1} \boldsymbol{\alpha})$), which is the case if intermediate inputs have stickier prices than final goods, or if sectors which use fewer intermediate inputs (i.e., with larger α) have stickier prices.

Implementing the calibration in (26) for the U.S. economy, we obtain $\kappa_{\text{round}} = 0.14$, which is of the same order of magnitude as $\kappa_{\text{true}} = 0.09$ in the full calibration (see Section 6).

EXAMPLE 2—Vertical Chain (or Sticky Wages): This example focuses on one of the simplest input-output networks, the two-stage vertical chain, to illustrate our modeling of wage rigidity; the compounding of nominal rigidities along the production chain; and why an increase in the productivity of final (intermediate) good producers improves (worsens) the inflation-output tradeoff in the consumer price Phillips curve, while this does not happen with the divine coincidence index.

The economy has two sectors: sector I produces intermediate goods, and sector F produces final goods. Producers in I only use labor, while producers in F use labor and intermediate inputs from I with shares α_F and $\Omega_{FI} = 1 - \alpha_F$. Consumers only purchase final goods ($\beta_F = 1, \beta_I = 0$).

Sticky wages. To model an economy with sticky prices and sticky wages, one can simply relabel the intermediate good producer in this example as a labor union. The union’s price coincides with the sticky wage. In this case, the direct labor share of the downstream sector is $\alpha_F = 0$, while $\Omega_{F,\text{union}} = 1$. The implications for the Phillips curve and welfare, illustrated below and in Example 5, are the same as in conventional sticky-wage models (Erceg, Henderson, and Levin (2000), Blanchard and Gali (2007)). The modeling device of adding labor unions as the most upstream sector, which sells labor services to all the others, can be adopted in any production network, as illustrated in Example 4.

Intermediate inputs flatten the Phillips curve. Recall from equation (24) that the slope of the consumer price Phillips curve is increasing in the pass-through of nominal wages into prices, $\delta^\beta \boldsymbol{\alpha}$, which in our example is given by

$$\delta^\beta \boldsymbol{\alpha} = \delta_F (\alpha_F + (1 - \alpha_F) \delta_I) < \delta_F = \mathbb{E}_\beta(\delta).$$

In the vertical economy, $\delta^\beta \boldsymbol{\alpha}$ is always smaller than the consumption-weighted average price stickiness δ_F , and it is decreasing in the intermediate input share $1 - \alpha_F$. Price rigidity in the intermediate good sector reduces the pass-through of wages into the final good producers’ marginal cost, which is given by the sum of the direct pass-through α_F , and the indirect pass-through via intermediate good prices, $(1 - \alpha_F) \delta_I$. The sum is less than 1 whenever intermediate good prices are sticky ($\delta_I < 1$). The effect on final good prices is then further discounted by the fraction δ_F of final producers who can reset their price. Hence, the Phillips curve is flatter than in an economy with no intermediate inputs. In the extreme case where final producers do not hire workers directly ($\alpha_F = 0$), we have $\delta^\beta \boldsymbol{\alpha} = \delta_F \delta_I \leq \min\{\delta_I, \delta_F\}$. In this case, calibrating $\delta^\beta \boldsymbol{\alpha}$ to any weighted average of micro level price adjustment frequencies would lead to an overestimate of the slope of the Phillips curve.

Cost-push shocks. In the vertical economy, changes in productivity generate a tradeoff between closing the output gap and stabilizing consumer prices. As an illustrative example, let us study the effect of an increase in the productivity of the final sector ($\log A_F > 0$). Note that consumer prices coincide with the price of the final good. Let us then consider how the shock affects marginal costs in the intermediate and final good sectors. We know

that wages increase, to reflect the higher marginal product of labor, while productivity increases only in the final good sector. Hence, the intermediate good sector does not benefit from the productivity change, but must pay higher wages; therefore, it faces positive inflation $\pi_I > 0$. If intermediate good prices are not fully flexible, however, the increase in wages is not fully passed through to final good producers, who instead enjoy the full benefit of the productivity shock. Thus, the pass-through of wages into final good prices is smaller than the pass-through of productivity:

$$\delta^\beta \alpha = \delta_F(\alpha_F + (1 - \alpha_F)\delta_I) < \delta_F = \delta_A.$$

Corollary 2 then implies $u^C < 0$, that is, consumer prices fall under zero output gap. Instead, the divine coincidence index is stabilized, because it averages out inflation in the intermediate and final good sectors, in such a way that the two exactly offset each other.

EXAMPLE 3—Horizontal Economy: This example focuses on the cost-push shocks induced by changes in productivity, and shows why a fall (an increase) in the relative productivity of flex-price sectors worsens (improves) the inflation-output tradeoff in the consumer-price Phillips curve. By contrast, the divine coincidence index does not suffer from an inflation-output tradeoff, because this index discounts flex-price sectors. Consider a horizontal economy with N sectors that hire labor and sell to final good consumers. There are no intermediate inputs ($\alpha = \mathbf{1}$, $\Omega = \mathbb{O}$), and final consumption shares are given by β_1, \dots, β_N . Price adjustment probabilities $\delta_1, \dots, \delta_N$ are different across sectors, which face idiosyncratic productivity changes $\log A_1, \dots, \log A_N$. In this economy, the residual in the consumer price Phillips curve is inversely proportional to the covariance of productivity shocks and price adjustment probabilities:

$$u^C \propto -\text{Cov}_\beta(\hat{\delta}, \log A).$$

As usual, this result comes from the offsetting effect of wages and productivity on marginal costs and prices. To close the output gap, real wages must change by the same amount as aggregate productivity, $\mathbb{E}_\beta(d \log A)$. The sectors that received a better shock compared to the average see a decrease in marginal costs and hence lower their price, and vice versa. The consumption-weighted average marginal cost does not change; however, prices respond more in flexible sectors. Therefore, controlling for the output gap, consumer inflation increases if the productivity of flexible sectors falls relative to the average, and vice versa.

EXAMPLE 4—Oil Economy: This example formalizes the conventional wisdom that oil price shocks generate a tradeoff between stabilizing output and stabilizing consumer price inflation. The mechanism combines insights from the vertical chain and the horizontal economy. Our oil economy has a vertical component, to account for the fact that wages are sticky, while oil prices are very flexible. In the model economy, a sticky-wage labor union takes the place of the intermediate good sector. The oil sector is in between the labor union and the final good sectors, and it uses labor to produce an intermediate input sold to final good firms. Final producers also hire labor directly. We assume $\delta_{\text{oil}} = 1$; therefore, the oil sector fully transmits wage and productivity changes to final good producers. In addition, like in the horizontal economy, there are many final good sectors, with different oil input shares $\{\omega_{i,\text{oil}}\}$ and price adjustment probabilities δ_i . The oil shares govern the heterogeneous exposure of final good producers to oil shocks. In the U.S. data, oil input shares are positively correlated with price adjustment probabilities.

Correspondingly, residual u_t^C in the consumer price Phillips curve has common elements with both the vertical and horizontal economies. It is given by

$$u^C = - \frac{\overbrace{\text{Cov}_\beta(\delta_i, \omega_{i,\text{oil}})}^{\text{horizontal}} + \overbrace{(1 - \delta_L)\mathbb{E}_\beta(\delta_i)\mathbb{E}_\beta(\omega_{i,\text{oil}})}^{\text{vertical}}}{1 - \delta_L\mathbb{E}_\beta(\delta_i)} \log A_{\text{oil}}.$$

Since $\delta_{\text{oil}} = 1$, changes in oil productivity are fully passed-through to final good producers, and therefore can be thought of as a negative downstream shocks. Like in the vertical chain, a negative shock $\log A_{\text{oil}} < 0$ generates a cost-push shock $u^C > 0$. Intuitively, when wages are sticky ($1 - \delta_L$ is large), there is no adjustment mechanism whereby wages fall together with productivity; therefore, marginal costs and prices increase. Moreover, sectors with more flexible prices face a larger cost increase, because of the positive correlation between price flexibility and oil input shares. Like in the horizontal economy, this generates an even larger cost-push shock.

5. WELFARE AND OPTIMAL MONETARY POLICY

This section derives the welfare loss function in the multi-sector economy (Section 5.1) and discusses the optimal monetary policy (Sections 5.2 and 5.3). The examples in Section 5.4 illustrate the optimal monetary policy in the simple economies introduced in Section 4.4.

5.1. Welfare Function

Proposition 3 derives a second-order approximation of the welfare loss with respect to the efficient equilibrium that would emerge in an economy with flexible prices. Welfare is negatively affected by the output gap, by price dispersion within sectors, and by relative price distortions across sectors. The substitution operators, introduced in Definition 4, describe the welfare cost of price distortions across sectors, which comes from inefficient expenditure switching by consumers and producers towards sectors whose relative price is too low compared to their productivity.

PROPOSITION 3: *A second-order approximation of the welfare loss with respect to the flex-price equilibrium is given by*

$$\mathbb{W} = \frac{1}{2} \sum_t \rho^t \left[(\gamma + \varphi) \tilde{y}_t^2 + \sum_i \lambda_i \frac{1 - \hat{\delta}_i}{\hat{\delta}_i} \epsilon_i \pi_{it}^2 + \Phi_C(\tilde{\chi}_t^T, \tilde{\chi}_t) + \sum_s \lambda_s \Phi_s(\tilde{\chi}_t^T, \tilde{\chi}_t) \right]. \quad (27)$$

The vector $\tilde{\chi} \in \mathbb{R}^{N,1}$ has components

$$\tilde{\chi}_{it} \equiv p_{it} - \sum_j (I - \Omega)_{ij}^{-1} \log A_{jt},$$

denoting wedges between current sectoral prices and network-adjusted productivities.

PROOF: Supplement Section 2.3.

Q.E.D.

DEFINITION 4: The *substitution operator*¹⁷ for sector $s \in \{1, \dots, N\}$ is a symmetric, bi-linear¹⁸ operator $\Phi_s : \mathbb{R}^{N \times L} \times \mathbb{R}^{N \times M} \mapsto \mathbb{R}^{L \times M}$ (where L and M are generic integers), such that

$$[\Phi_s(X^T, Y)]_{l,m} \equiv \frac{1}{2} \sum_k \sum_h \omega_{sk} \omega_{sh} \theta_{kh}^s (X_{k,l} - X_{h,l})(Y_{k,m} - Y_{h,m}). \tag{28}$$

Similarly, the substitution operator for final consumers has elements

$$[\Phi_C(X^T, Y)]_{l,m} \equiv \frac{1}{2} \sum_k \sum_h \beta_k \beta_h \sigma_{kh}^C (X_{k,l} - X_{h,l})(Y_{k,m} - Y_{h,m}). \tag{29}$$

The welfare function in equation (27) depends on three terms: the output gap, price dispersion within sectors, and price distortions across sectors. The first two terms are common to the one-sector model, while the third is novel. The squared output gap measures deviations of employment from the efficient level (see Remark 5). The squared sectoral inflation rates instead capture within-sector price dispersion. All the firms within a sector face the same marginal cost; therefore, it is inefficient for them to charge different prices. The welfare loss is proportional to the within-sector substitution elasticity ϵ_i , which governs the extent to which customers misallocate demand towards the varieties with lower price. While the expression for sectoral losses is the same as in the one-sector economy, equation (27) also tells us that they should be aggregated according to sales shares.

The wedges $\tilde{\chi}_i$ between sectoral prices and productivities are novel to the multi-sector economy. They arise because a share of the firms cannot adjust their price in response to changes in productivity and input prices, thereby distorting sectoral relative prices. The welfare cost is captured by the third term in equation (27), where sectoral price wedges are taken as inputs into the substitution operators. The operators $\Phi_s(\tilde{\chi}_i^T, \tilde{\chi}_i)$ and $\Phi_C(\tilde{\chi}_i^T, \tilde{\chi}_i)$ describe the resulting demand misallocation by producers in each sector s and by final consumers, which causes sectoral total factor productivities (TFP) to decline. In equations (28) and (29), changes in the relative demand for inputs and consumption goods are proportional to the relevant elasticities of substitution of each producer ($\{\theta_{kh}^s\}_{k,h=1}^N$) and of final consumers ($\{\sigma_{kh}^C\}_{k,h=1}^N$), while the effect of distortions in the demand for each input on TFP is proportional to the inputs' cost shares.

The substitution operators are similar to the variance of price wedges across inputs (consumption goods), with probability weights given by input (consumption) shares $\{\omega_{sk}\}_{k=1}^N$ ($\{\beta_k\}_{k=1}^N$). Differently from a regular variance, the operators place higher weight on input pairs with larger substitution elasticity. In the special case where substitution elasticities are uniform across inputs ($\theta_{kh}^s \equiv \theta^s$, $\sigma_{kh}^C \equiv \sigma^C$), the substitution operators are exactly equal to the variances $\theta^s \text{Var}_{\Omega(s)}(\tilde{\chi})$ and $\sigma^C \text{Var}_\beta(\tilde{\chi})$.

¹⁷This definition was first introduced by Baqaee and Farhi (2018). Baqaee and Farhi (2018) used the substitution operators to describe the first-order effect of markups on allocative efficiency around a distorted steady state, while in Proposition 3 the same operators characterize the second-order effect around an efficient steady state.

¹⁸Bi-linear means that, for every $A \in \mathbb{R}^{L_1 \times L_2}$, $X \in \mathbb{R}^{N \times L_1}$,

$$A^T \Phi_s(X^T, Y) = \Phi_s(A^T X^T, Y),$$

and for every $B \in \mathbb{R}^{M_1 \times M_2}$, $Y \in \mathbb{R}^{N \times M_1}$,

$$\Phi_s(X^T, Y)B = \Phi_s(X^T, YB).$$

The same holds for Φ_C .

In equation (27), sector-level TFP losses are then weighted by sales shares and added up to obtain an aggregate Harberger triangle, which—as shown in the proof of Proposition 3—captures the difference between the fall (increase) in the aggregate labor share and the increase (fall) in firm profits coming from higher (lower) markups.

REMARK 9: The Cobb–Douglas roundabout economy, which is commonly used to approximate the input-output structure, would provide a poor approximation of the welfare function (27). If we believe that inputs are complementary, restricting production technologies to be Cobb–Douglas would overestimate the welfare loss from cross-sector misallocation. Assuming a roundabout structure further exacerbates the issue, if the sectors with larger sale share are more essential inputs, that is, more complementary with the others (the energy sector is a clear example).

5.2. Optimal Monetary Policy

The optimal monetary policy minimizes the welfare loss function (27) subject to the IS curve (11), the sectoral Phillips curves (22), and the evolution of log prices (13). Monetary policy can freely set the nominal interest rate; therefore, the IS curve can be eliminated from the set of constraints. Moreover, it is convenient to rewrite the welfare loss function using the following notation:

$$\begin{aligned} \mathcal{L}^{\text{within}} &\equiv \text{diag}(\boldsymbol{\lambda}) \text{diag}(\boldsymbol{\epsilon})(I - \Delta)\Delta^{-1}, \\ \mathcal{L}^{\text{across}} &\equiv \Phi_C(I_N, I_N) + \sum_s \lambda_s \Phi_s(I_N, I_N), \\ \mathcal{L} &\equiv \mathcal{L}^{\text{within}} + \mathcal{L}^{\text{across}}. \end{aligned}$$

Proposition 4 states and solves the optimal policy problem with commitment.

PROPOSITION 4: *The optimal monetary policy with commitment solves*

$$\min_{\{\tilde{y}_{t+r}, \pi_{t+r}\}_{r=0}^{\infty}} \frac{1}{2} \sum_t \rho^t [(\gamma + \varphi)\tilde{y}_t^2 + \boldsymbol{\pi}_t^T \mathcal{L} \boldsymbol{\pi}_t + \boldsymbol{\chi}^T \mathcal{L}^{\text{across}} \boldsymbol{\chi}]$$

subject to the vector of sectoral Phillips curves (22) and the law of motion for prices

$$\mathbf{p}_t = \mathbf{p}_{t-1} + \boldsymbol{\pi}_t.$$

The first-order condition is

$$\tilde{y}_t + \mathcal{B}^T \mathcal{L} \boldsymbol{\pi}_t = -\mathcal{B}^T [\boldsymbol{\zeta}_t^p - (I - \mathcal{V})\boldsymbol{\zeta}_{t-1}^{\text{PC}}], \tag{30}$$

where the Lagrange multipliers $\boldsymbol{\zeta}_t^{\text{PC}}$ and $\boldsymbol{\zeta}_t^p$ follow the laws of motion

$$\begin{cases} \boldsymbol{\zeta}_t^{\text{PC}} = (I - \mathcal{V})^T \boldsymbol{\zeta}_{t-1}^{\text{PC}} + \boldsymbol{\zeta}_t^p - \mathcal{L} \boldsymbol{\pi}_t, \\ \rho \boldsymbol{\zeta}_{t+1}^p = \boldsymbol{\zeta}_t^p + \rho \mathcal{V} \boldsymbol{\zeta}_{t+1}^{\text{PC}} - \rho \mathcal{L}^{\text{across}} \boldsymbol{\chi}_{t+1}. \end{cases}$$

To build intuition, let us start by examining the static case with $\rho = 0$, also assuming that the economy is in steady state at time $t - 1$, so that $\boldsymbol{\zeta}_{t-1}^{\text{PC}} = \mathbf{0}$. The first-order condition (30) becomes

$$\tilde{y}_t + \mathcal{B}^T \mathcal{L} \boldsymbol{\pi}_t = 0. \tag{31}$$

In equation (31), the optimal monetary policy trades off achieving a zero output gap against improving allocative efficiency within and across sectors. The effect of the output gap on sectoral prices is given by the slope \mathcal{B} of sector-level Phillips curves. The matrix \mathcal{L} instead captures the effect of relative price changes on allocative efficiency, for given initial inflation π .

Note that monetary policy has only one instrument. Therefore, it can influence relative prices only along one dimension, proportionately to the slope \mathcal{B} of sectoral Phillips curves. In practice, the slope \mathcal{B} is small, and likely orthogonal to the efficient price change. Therefore, the effect of monetary policy on misallocation ($\mathcal{B}^T \mathcal{L} \pi$) is small in general,¹⁹ and the central bank should focus on closing the output gap instead.

Using Proposition 1, we can replace the output gap in (30) with the divine coincidence index. For $\rho = 0$, this allows us to express the optimal policy as an inflation targeting rule:

$$\pi_t^* \equiv \left[\frac{\lambda^T (I - \Delta) \Delta^{-1}}{\gamma + \varphi} + \mathcal{B}^T \mathcal{L} \right] \pi_t = 0. \tag{32}$$

In the dynamic model, the optimal policy has additional forward and backward looking components, on top of the static tradeoff between output gap and allocative efficiency. Forward looking considerations are captured by the Lagrange multiplier ζ^p , and reflect the fact that the current policy affects future misallocation. The multiplier ζ_t^p captures the cost of future price distortions, which respond to the current output gap proportionately to the slope of sectoral Phillips curves \mathcal{B} . The backward looking term $\mathcal{B}^T (I - \mathcal{V})^T \zeta_{t-1}^{PC}$ instead reflects the fact that monetary policy can improve misallocation by influencing inflation expectations. This creates gains from commitment, and makes current policy dependent on past shocks. In equation (30), the Lagrange multiplier ζ_{t-1}^{PC} captures the cost of inflation in the previous period, while $I - \mathcal{V}$ is the response of past inflation to current inflation, and \mathcal{B} is the elasticity of current inflation to the output gap.

The presence of these new terms does not change the conclusions from the static setting. The price change that monetary policy can achieve by distorting the output gap is always proportional to the slope \mathcal{B} ; therefore, it is small and likely orthogonal to the desired one. Thus, stabilizing the output gap remains nearly optimal in the dynamic framework as well.

5.3. Output Gap Targeting

Section 5.2 shows that there is a tradeoff between closing the output gap and reducing relative price wedges across sectors. However, deviating from zero output gap has a small effect on sectoral wedges; therefore, the optimal policy is very similar to closing the output gap at all periods. Lemma 4 proves that this can be achieved in every period by following a Taylor rule (12), with the divine coincidence index as the inflation target.

LEMMA 4: *Assume that sectoral productivity shocks follow AR1 processes:*

$$\log A_{it} = \eta_A \log A_{it-1} + \xi_{it}$$

¹⁹This does not mean that the welfare loss from misallocation is small. To the contrary, Section 6.3 shows that it is large quantitatively.

with $\mathbb{E}\xi_{it} = 0$ and $\eta_A < 1$. Then there exists a unique path of inflation rates and markups such that $\tilde{y}_t = 0 \forall t$. This equilibrium can be implemented via the interest rate rule

$$i_{t+1} = \underbrace{r_{t+1, \text{nat}} + \mathbb{E}(\pi_{t+1}^{CPI} | \tilde{y} \equiv 0)}_{\text{nominal rate under zero output gap}} + \zeta \pi_t^{DC} \tag{33}$$

with $\zeta > 0$.

PROOF: Supplement Section 2.3.

Q.E.D.

5.4. Examples

The examples in this section illustrate three general principles in the simple economies from Section 4.4. First, the optimal monetary policy has a bias towards stabilizing wages, or upstream prices (vertical chain example). Second, the central bank faces a tradeoff between correcting misallocation within and across sectors (horizontal economy example). And third, the central bank should accommodate inflationary supply shocks, but the optimal deviation from the efficient output is small (oil economy example). While the vertical chain and oil economy examples highlight reasons why the optimal policy should deviate from perfect output stabilization, it is important to reiterate that the optimal output gap is very close to zero in most practical applications, as explained in Section 5.2.

To streamline the exposition, we eliminate dynamic considerations by assuming $\rho = 0$. Equation (31) and the definition of the matrix \mathcal{L} allow us to isolate a within-sector and an across-sector component of the optimal output gap:

$$\tilde{y}_t^* = \tilde{y}_{\text{within},t}^* + \tilde{y}_{\text{across},t}^*; \quad \tilde{y}_{\text{within},t}^* \equiv -\mathcal{B}^T \mathcal{L}^{\text{within}} \boldsymbol{\pi}_t; \quad \tilde{y}_{\text{across},t}^* \equiv -\mathcal{B}^T \mathcal{L}^{\text{across}} \boldsymbol{\pi}_t.$$

EXAMPLE 5—Vertical chain: Let us return to the vertical chain in Example 2, and again consider a negative productivity shock to the final good sector. In this case, the optimal monetary policy keeps output above the efficient level, and tolerates positive consumer price inflation, in order to stabilize the price of intermediate goods. Note that, if we interpret the upstream sector as a labor union, the vertical chain model represents an economy with sticky prices and sticky wages. This example then tells us that optimal monetary policy prioritizes wage over price stabilization.

Formally, the within-sector component of the optimal output gap is given by

$$\tilde{y}_{\text{within}}^* = -\epsilon \frac{(1 - \alpha_F)(1 - \delta_I)(1 - \delta_F)}{(1 - \delta^\beta \alpha)^2} [\delta_I - \delta^\beta \alpha] \log A_F. \tag{34}$$

This expression is positive for $d \log A_F < 0$. By bidding up wages, a positive output gap reduces the desired price adjustment of intermediate good firms, thereby reducing price dispersion among intermediate goods. This comes at the cost of increasing price dispersion in the final good sector, where lower productivity is no longer compensated by lower wages. Nonetheless, this policy is constrained optimal, because intermediate good prices are more responsive than final good prices to monetary policy, as they are more directly exposed to wage changes. In the equation, raising the output gap reduces price dispersion by δ_I in the intermediate good sector, and increases it by $\delta^\beta \alpha < \delta_I$ in the final good sector.

The across-sector component is also positive for $d \log A_F < 0$. Due to price rigidity in the intermediate good sector, wages fall by more than intermediate good prices. There-

fore, final good producers inefficiently demand too much labor relative to intermediate inputs. Pushing output above the natural level mitigates the fall in wages, thereby closing the inefficient gap between wages and intermediate good prices. Hence, we have

$$\tilde{y}_{\text{across}}^* = -\theta_L^F \frac{(1 - \alpha_F)\alpha_F}{(1 - \delta^\beta \alpha)^2} (1 - \delta_I)^2 (1 - \delta_F) \log A_F > 0.$$

EXAMPLE 6—Horizontal economy: Consider the horizontal economy in Example 3. This example highlights the tradeoff between implementing the correct relative prices within and across sectors. The economy has a nested CES structure, with an outer nest given by the final consumption aggregator, and an inner nest given by the aggregators of sectoral varieties. Eliminating price dispersion in the inner nest would require zero inflation for all firms, because some of them are constrained to keep a fixed price. By contrast, in the outer nest, relative prices should adjust to reflect changes in relative productivities across sectors.

Thus, the within- and across-sector components of the optimal output gap have opposite sign, and in this example they are exactly proportional. Their relative magnitude is given by the ratio of the substitution elasticities in the inner and outer nest (ϵ and σ), which govern the cost of price distortions within each. The optimal output gap trades off the two components, and is given by

$$\tilde{y}^* = (\sigma - \epsilon)(\gamma + \varphi)\mathbb{E}_{\beta(1-\delta)} \boldsymbol{\pi}, \tag{35}$$

where $\mathbb{E}_{\beta(1-\delta)}$ is the expectation computed with probability weights $\{\frac{\beta_i(1-\delta_i)}{\sum_j \beta_j(1-\delta_j)}\}_{i=1,\dots,N}$. One can then solve for the response of inflation in (35) as a function of sectoral productivity shocks. Following the same reasoning as in Example 3, this depends on the covariance between productivity shocks and price adjustment frequencies:

$$\mathbb{E}_{\beta(1-\delta)} \boldsymbol{\pi} = \text{Cov}_{\beta(1-\delta)}(\delta, \log A). \tag{36}$$

Note that the sectors with highest weight in the covariance are those with an intermediate degree of price stickiness (so that the product $\delta_i(1 - \delta_i)$ is highest). The intuition is particularly clear in the case of within-sector misallocation. Sectors where no firm can adjust ($\delta_i \approx 0$), or all firms can adjust ($\delta_i \approx 1$), face little price dispersion, because a large fraction of the firms end up charging the same price. Hence, the sectors which face the most severe price dispersion are those where about half of the firms can adjust, and the other half cannot (i.e., $\delta_i \approx \frac{1}{2}$).

EXAMPLE 7—Oil economy: Consider the oil economy in Example 4. Consistent with the conventional wisdom, after a negative oil shock the optimal monetary policy should tolerate some inflation, and keep output slightly above the flex-price level. To see this, recall that oil shocks have commonalities with downstream shocks in a vertical chain, and heterogeneous shocks in a horizontal economy, as explained in Example 4. The optimal policy is mainly determined by the vertical component. Example 5 then tells us that the central bank should implement a positive output gap, in an effort to reduce misallocation in the labor market. Formally, the component of the optimal output gap that comes from misallocation within the labor sector is

$$\tilde{y}_{\text{vert,within}}^* = -\epsilon(\gamma + \varphi) \frac{\delta_L \mathbb{E}_\beta(1 - \delta)}{1 - \delta_L \mathbb{E}_\beta(\delta)} \left(1 - \frac{\delta_L(1 - \mathbb{E}_\beta(\delta))}{1 - \delta_L \mathbb{E}_\beta(\delta)} \right) \mathbb{E}_{\beta(1-\delta)}(\omega_{i,\text{oil}}) \log A_{\text{oil}}.$$

This expression is positive for $\log A_{\text{oil}} < 0$, and it becomes 0 when wages are fully rigid or fully flexible ($\delta_L = 0$ or $\delta_L = 1$, respectively). In fact, when $\delta_L = 0$ or $\delta_L = 1$, there is no misallocation across workers, and when $\delta_L = 0$, monetary policy cannot influence wages or prices. Also note that there is no across-sector misallocation coming from the vertical component of the economy ($\bar{y}_{\text{vert.across}}^* = 0$), because of two reasons. First, labor is the only input in the oil sector, so that oil producers cannot substitute across inputs. Second, the oil price is fully flexible, and thus oil prices are never distorted relative to wages. Therefore, producers of the final goods always use the efficient combination of labor and oil.

Section 2.3 of the Supplement derives the horizontal component of the optimal output gap, which is similar to equation (36). Its policy implications are ambiguous *ex ante*, because they depend on the covariance between oil input shares and price adjustment frequencies, and on the relative magnitude of within- and across-sector elasticities of substitution. When calibrating sectoral oil shares and price adjustment frequencies to the U.S. data, the covariance is small quantitatively; therefore, the optimal policy is driven by the vertical component.

6. QUANTITATIVE RESULTS

6.1. Data

To calibrate the multi-sector model, we need to assign values to the preference and production parameters in Table II. In addition, to compute the welfare estimates in Section 6.3 and the time series of cost-push shocks in Section 6.2.2, we also need an estimate of sectoral productivity changes over time and of their covariance matrix.

The preference parameters are set to standard values in the literature: the wealth effect in labor supply is $\gamma = 1$, the inverse Frish elasticity is $\varphi = 2$, and the discount factor is $\rho = 0.9975$.

Disaggregated input, labor, and final consumption shares are calibrated to match the BEA input-output accounts. Following the standard BEA methodology explained in Horowitz and Planting (2006, Chapter 12), we net out the non-labor component of value added from total costs in the construction of labor and input shares (α and Ω). The baseline calibration is based on data for the year 2012, at the 405-sector level. The calibration of the Phillips curve slope over time (Section 6.2.1) instead is based on historical data for the years 1947–2017, with the level of aggregation varying slightly across years (46 to 71 industries).

We use estimates of sector-level probabilities of price adjustment computed by Pasten, Schoenle, and Weber (2017) based on firm-level price series underlying the BLS PPI data. There are no estimates for personal services, repair and maintenance, and government. In the baseline calibration, we set the adjustment probability of these sectors equal to the mean.²⁰ We also calibrate the quarterly probability of wage adjustment to 0.25, in line with Barattieri, Basu, and Gottschalk (2014) and Beraja, Hurst, and Ospina (2019).

The elasticities of substitution in production and consumption are calibrated to consensus values in the literature. The substitution elasticity between consumption goods is $\sigma = 0.9$,²¹ the elasticity of substitution between labor and intermediate inputs is

²⁰The missing sectors account for 10% of total sales, of which 8% comes from the government. None of the results changes significantly in an alternative calibration which excludes these sectors.

²¹Atalay (2017), Herrendorf, Rogerson, and Valentinyi (2013), and Oberfeld and Raval (2014) estimated it to be slightly less than 1.

$\theta_L = 0.5$,²² the elasticity of substitution across intermediate inputs is $\theta = 0.1$,²³ and the elasticity of substitution between varieties within each sector is $\epsilon = 8$.²⁴

Finally, productivity shocks are measured as the growth rate of the Multifactor Productivity (MFP) index at the sector level,²⁵ reported in the BEA Integrated Industry-Level Production Account data for the period 1988–2016.²⁶ The covariance matrix of sectoral productivity shocks is also computed based on these estimates.

6.2. Phillips Curve

6.2.1. Slope

As a key result, in the multi-sector model, the slope of the consumer price Phillips curve is decreasing in intermediate input shares. To illustrate the quantitative importance of this mechanism, Table III reports the slope in the baseline calibration and under alternative input-output structures.

The first column presents the benchmark New Keynesian model with flexible wages and no input-output linkages, while the second column corresponds to the baseline calibration of the multi-sector model. In the multi-sector economy, the slope is one order of magnitude smaller than in the one-sector model. The third and fourth columns tease out the contribution of sticky wages and intermediate inputs, showing that both are important. The third column reports the slope in a model where input-output linkages are as in the baseline and wages are flexible, while the fourth assumes sticky wages but no input-output linkages.

The final column shows that eliminating heterogeneity in price adjustment frequencies does not change the slope. This result is not true in general, but it depends on the correlation (or lack thereof) between labor, input, and consumption shares.²⁷ Previous quantitative studies (Carvalho (2006), Nakamura and Steinsson (2010)) instead found that heterogeneous price adjustment frequencies increase monetary non-neutrality. The difference with respect to these previous results comes from how the average price rigidity is computed. Specifically, we first computed the transformation $\hat{\delta}$ derived in Section 3.1 sector-by-sector, and then averaged the sectoral $\hat{\delta}$'s. Previous studies did the opposite. By Jensen's inequality, this mechanically leads to a lower value of the average price rigidity,

TABLE III
PHILLIPS CURVE SLOPE IN THE MAIN AND ALTERNATIVE CALIBRATIONS.

	Textbook	Full model	Flex wage	No IO	$\delta = \text{mean}$
Slope	1.16	0.09	0.32	0.22	0.08

²²This is consistent with Atalay (2017), who estimated this parameter to be between 0.4 and 0.8.

²³See Atalay (2017).

²⁴This is consistent with estimates of the variety-level elasticity of substitution from the industrial organization and international trade literatures.

²⁵The MFP is constructed taking into account labor, capital, and intermediate inputs from manufacturing and services. Therefore, this index captures changes in gross output TFP, which is the correct empirical counterpart of the sector-level TFP shocks in the model.

²⁶<https://www.bea.gov/data/special-topics/integrated-industry-level-production-account-klems>.

²⁷For example, ignoring heterogeneity would overestimate the slope if sectors with large labor share have stickier prices.

because $\hat{\delta}$ is a convex transformation.²⁸ Adopting the same procedure as previous studies would imply a much smaller calibrated slope of 0.03.

Section 2.3 of the additional Additional Online Materials discusses the implications for monetary non-neutrality.

Evolution Over Time. We also consider how changes in the input-output structure of the economy affected the slope of the Phillips curve. The two panels of Figure 1 report the slope of the consumer-price and divine coincidence Phillips curves as predicted by the model, based on historical input-output data for the period 1947–2017. Importantly, even at the beginning of the period, the slopes are much smaller than implied by the benchmark one-sector model. Therefore, changes in the input-output structure cannot rationalize a regime shift from steep to flat Phillips curves (see, e.g., Stock and Watson (2007), Blanchard (2016), Del Negro, Lenza, Primiceri, and Tambalotti (2020)). Nonetheless, the model still predicts that the CPI and DC Phillips curves have flattened by about 30% over the past 70 years.

The calibration relies on some important assumptions. First, we assume that sectoral price rigidities remained constant over time, due to lack of historical data.²⁹ Second, as




FIGURE 1.—Slope of the Phillips curve over time.

²⁸Nakamura and Steinsson (2010) provided intuition for why $\hat{\delta}$ is a convex transformation. Consider two economies, both with the same average probability of price adjustment across sectors. In the first economy, all sectors have the same adjustment probability, while in the second, some sectors are more flexible and some are stickier. As long as the discount factor is large enough, producers reset their prices to be an “average” of the optimal prices over the period before their next opportunity to adjust. If all sectors have the same adjustment probability, the producers who can adjust know that many others will also have changed their price by the time they get to adjust again. Therefore, they preemptively adjust more. If instead some sectors adjust very infrequently, producers in the flexible sectors know that they will likely have another opportunity to reset prices before the stickier sectors also get to change theirs. Therefore, it is optimal to wait. The expectations channel gets muted as the discount factor goes to zero.

²⁹For sectors where data are available, Nakamura and Steinsson (2013) found that the frequency of price adjustment remained stable since the 1990s. Nonetheless, it is possible that changes in the frequency of price adjustment may be an important reason behind the flattening of the Phillips curve over longer horizons. For

explained in Section 6.1, we net out the non-labor component of value added from total costs in the computation of labor and input shares. The implications for the slope are not sensitive to this assumption if the non-labor component of value added corresponds to profits, and markups are determined independently of the output gap. If instead non-labor value added also includes rental payments to fixed factors (such as capital and land), our calibration is an upper bound to how much the Phillips curve has flattened.³⁰ The blue lines depict the calibrated slopes over time. The red and green lines isolate the contributions of changes in intermediate input versus consumption shares. The decomposition is based on equation (24), which expresses the slope of the Phillips curve as a function of the pass-through of nominal wages into consumer prices, $\delta^\beta \alpha$. This pass-through in turn can be decomposed into a term coming from consumption shares, and a term coming from the input-output structure:

$$\delta^\beta \alpha = \underbrace{\beta^T}_{\text{consumption}} \underbrace{\Delta(I - \Omega\Delta)^{-1} \alpha}_{\text{input-output}}$$

The red line assumes that input shares remained constant at their 1947 values, and lets consumption shares evolve as observed in the data (setting $\delta^\beta_{\text{cons}} \alpha \equiv \beta_t^T \Delta(I - \Omega_{1947} \Delta)^{-1} \alpha_{1947}$). Vice versa, the green line fixes consumption shares at their 1947 values and lets input shares evolve as observed in the data (setting $\delta^\beta_{\text{input}} \alpha \equiv \beta_{1947}^T \Delta(I - \Omega_t \Delta)^{-1} \alpha_t$). Changes in intermediate input shares explain most of the overall decline. Changes in the allocation of consumption expenditures across final goods have an effect on the slope of the Phillips curve after 1980, as consumers started buying fewer manufacturing goods and more services, which have a stickier price.

Wage Phillips Curve. The multi-sector model also allows us to compare the wage and the consumer-price Phillips curve. The wage Phillips curve is much steeper, and the calibrated slope has remained almost unchanged over time (0.78 in 1947 against 0.77 in 2017). This result is consistent with empirical studies (see, e.g., Hooper, Mishkin, and Sufi (2019)), and can be easily rationalized in light of the theory. From Proposition 2, the slope of the wage Phillips curve is $\kappa^{\text{wage}} = \frac{\delta_{\text{wage}}}{1 - \delta^\beta \alpha}$. Note that κ^{wage} depends on input-output linkages only through general equilibrium effects, captured by the denominator $1 - \delta^\beta \alpha$. Therefore, it is not surprising that it remained unaffected by changes in the input-output structure. Moreover, the wage adjustment probability δ_{wage} is much larger than the overall pass-through of wages into consumer prices, $\delta^\beta \alpha$, which is further reduced by sticky intermediate and final good prices.

6.2.2. Cost-Push Shocks

This section evaluates the quantitative importance of the inflation-output tradeoff generated by changes in productivity, in two ways. First, Figure 2 plots a time series of the

example, firms may adjust prices less often today than in the 1970s, because with better monetary policy, inflation has become more stable (see Blanchard (2016)).

³⁰In the baseline calibration, the labor share is computed as the ratio of payments to labor over the sum of payments to labor plus intermediate input purchases. Section 2.1 of the additional Additional Online Materials reports an alternative calibration which replaces payments to labor with total value added. In this calibration, the slope declines only slightly (from 0.145 to 0.115 for the CPI Phillips curve, and from 0.24 to 0.2 for the DC Phillips curve), consistent with the fact that sales over total value added have increased by less than sales over labor compensation.




FIGURE 2.—Productivity component of the residual.

model-implied residual in the consumer price Phillips curve, based on observed changes in sectoral TFPI. Second, Table IV calibrates the response of consumer inflation to oil shocks, which are commonly believed to cause an inflation-output tradeoff.

Time Series. Figure 2 plots the productivity component of the residual u^C defined in equation (25).³¹

The time series has a mean of -0.16 and a standard deviation of 0.25 , which are both large relative to the calibrated slope of the consumer-price Phillips curve, suggesting that productivity-induced cost-push shocks are a significant driver of variation in consumer price inflation. This is consistent with the results in Supplement Section 1.2, whereby adding the productivity-driven residual to Phillips curve regressions increases the R-squared in the CPI and PCE specifications. Interestingly, the R-squared does not increase for core inflation or the divine coincidence index, suggesting that the model-based residual u^C is not correlated generically with economic activity, but only with the inflation indices which it is designed to track. This validates our modeling of how sectoral productivity changes propagate into different measures of aggregate inflation.

Inflation Response to Oil Shocks. Table IV computes the inflation response to oil shocks.

TABLE IV
CONSUMER INFLATION AFTER A 10% NEGATIVE SHOCK TO THE OIL SECTOR (FULL MODEL).

	$\delta = \text{actual}$	$\delta \equiv \delta_{\text{mean}}, \delta_{\text{oil}} = 1$	$\delta \equiv \delta_{\text{mean}}$
Sticky wages	0.25	0.07	-0.03
Flexible wages	0.19	-0.03	-0.17

³¹Note that the time series in the plot does not account for lagged deviations of prices from steady state.

In the full calibration, consumer prices increase by 0.25% in response to a 10% fall in oil productivity. The simplified oil economy in Example 4 helps us to identify the drivers of this result: wages are sticky, oil prices are flexible, and there is a positive correlation between price flexibility and oil input shares across sectors. The second row and the second and third columns shut down each of these channels, leading to a much smaller inflation response.

Based on Proposition 4, we also compute the optimal monetary policy. The central bank should raise the output gap by 0.08% on impact, which is a small response compared to the 0.69% decline in natural output. Hence, while the central bank should partially accommodate the shock, it should still let output fall substantially.

6.3. Welfare

The textbook New Keynesian model implies that changes in productivity generate no welfare loss compared to an economy with flexible prices. Moreover, the Lucas (1977) estimate tells us that, in the flex-price economy, the welfare loss induced by productivity fluctuations is very small, equal to 0.05% of per-period GDP. We now revisit these conclusions in the multi-sector economy, showing that, when combined with price rigidities and multiple sectors, productivity fluctuations generate sizable welfare losses, even under the optimal monetary policy.

As explained in Section 5, monetary policy has very limited power to correct relative prices within and across sectors; therefore, it must tolerate a welfare loss. Figure 3 compares welfare under different input-output structures, policy rules, and correlations between sectoral productivity shocks.

The height of the bars corresponds to the average welfare loss in units of per-period GDP, relative to the efficient equilibrium. The left panel calibrates input-output linkages according to the BEA data. In the first set of bars, the covariance of sectoral productivity

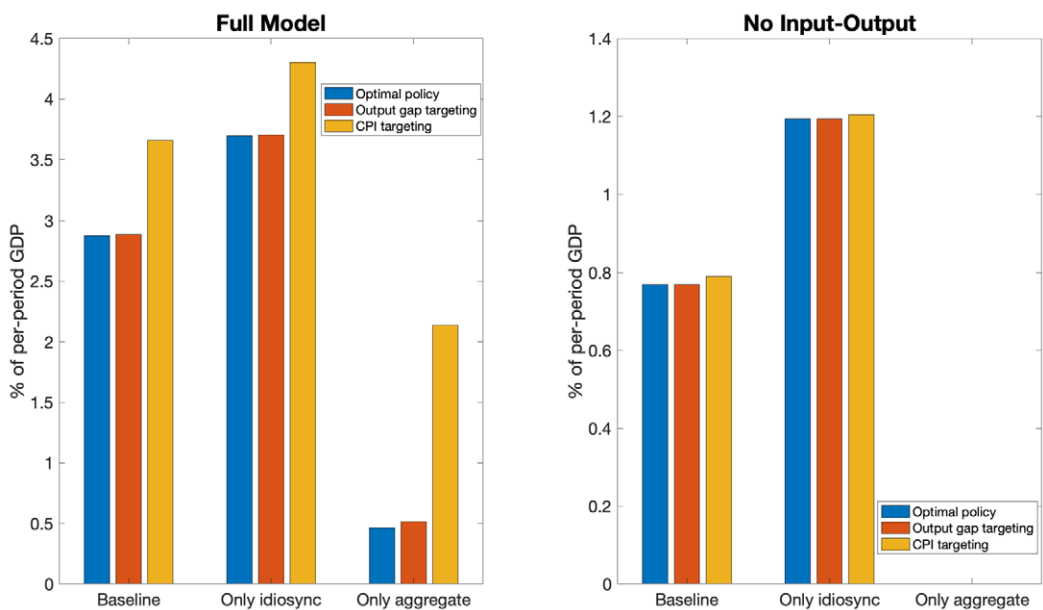


FIGURE 3.—Welfare loss from business cycles.

shocks is as measured in the KLEMS data, while in the second and third set, productivity shocks are either independent or fully correlated across sectors (keeping the variance of aggregate productivity constant). As the figure shows, most of the welfare loss comes from the idiosyncratic component.

Bars of different colors correspond to different policy rules (blue for optimal policy, red for output gap targeting, and yellow for consumer price targeting). Under the optimal policy, the expected welfare loss is 2.9% of per-period GDP—orders of magnitude larger than the Lucas estimate. Targeting zero output gap yields a negligible additional loss, consistent with the discussion in Section 5.2. Targeting consumer prices instead increases the loss to 3.7% of per-period GDP. Under the optimal policy, the welfare loss comes almost entirely from misallocation, not from inefficiently high or low employment, because this policy is close to implementing zero output gap. Section 2.2 of the additional Additional Online Materials further decomposes the loss from misallocation into its within- and across-sector components.

The right panel of Figure 3 instead considers an economy without input-output linkages. Here, the welfare loss is much smaller (0.8% of per-period GDP), and targeting consumer prices is close to optimal.

7. PHILLIPS CURVE REGRESSIONS

This section compares traditional empirical specifications of the Phillips curve, based on consumer prices, against the novel divine coincidence specification.

7.1. Data

The sample period is 1984–2018. The time series of consumer price inflation (CPI, PCE, and their core versions), wage inflation, the output gap, the employment gap, and the unemployment rate are provided by the FRED database. The time series for the divine coincidence index is constructed by aggregating sector-level PPI series, provided by the BLS. The weight assigned to each sector is a function of the relevant Domar weight and price adjustment probability, as in Definition 2. Domar weights are computed based on BEA input-output data, while sectoral price adjustment probabilities are from Pasten, Schoenle, and Weber (2017).

The Supplement, Section 1.1 provides a comparison between the weighting schemes of the divine coincidence index versus the PCE. A key difference is that consumer prices do not include wage inflation, which in the divine coincidence index instead has the highest weight of 18%. The divine coincidence index also assigns a high weight to large intermediate good sectors with sticky prices, such as professional services, financial intermediation, and durable goods. By contrast, the sectors with highest weight in the PCE are healthcare, real estate, and nondurable goods. These sectors have large consumption shares, but relatively small total sale shares. Section 1.1 of the Supplement also plots a time series of the divine coincidence index against various measures of consumer and producer price inflation.

7.2. Full-Sample Regressions

Table V shows regression coefficients and R-squareds for the Phillips curve regression

$$\bar{\pi}_t = \rho \mathbb{E} \bar{\pi}_{t+1} + \kappa \tilde{y}_t + v_t \quad (37)$$

TABLE V
REGRESSIONS OF YEARLY INFLATION ON THE CBO'S UNEMPLOYMENT GAP, 1984–2018.

	–		Mich		SPF CPI		GMM
	Gap	R ²	Gap	R ²	Gap	R ²	Gap
DC	–0.17 (0.03)	0.25	–0.19 (0.03)	0.31	–0.17 (0.03)	0.31	–0.16 (0.05)
CPI	–0.27 (0.07)	0.10	–0.29 (0.07)	0.12	–0.26 (0.07)	0.12	–0.12 (0.04)
coreCPI	–0.18 (0.06)	0.06	–0.25 (0.06)	0.27	–0.16 (0.06)	0.19	–0.04 (0.04)
PCE	–0.17 (0.06)	0.06	–0.2 (0.06)	0.08	–0.16 (0.06)	0.08	–0.06 (0.02)
corePCE	–0.12 (0.05)	0.03	–0.16 (0.05)	0.17	–0.1 (0.05)	0.12	–0.02 (0.01)
PPI	–0.32 (0.25)	0.01	–0.14 (0.24)	0.11	–0.37 (0.25)	0.04	–0.2 (0.13)

with different inflation measures $\bar{\pi}_t$ on the left-hand side. The independent variable \tilde{y}_t is the CBO employment gap. The results are similar when \tilde{y}_t is the CBO output gap or the unemployment rate (see Section 1.2 of the Supplement).³²

Each row in the table corresponds to a different inflation index (CPI, PCE, core CPI, core PCE, DC, and PPI), while the various columns adopt different methods to account for inflation expectations. The first column omits the expectation term in (37). The second and third columns use survey-based proxies for expected inflation, from the Michigan Survey of Consumers and the Survey of Professional Forecasters. The fourth column is based on a moment condition, which imposes that the difference between actual and expected inflation must be orthogonal to all available information at the time when expectations are formed.³³ Regardless of how one controls for inflation expectations, the R-squared is much higher in the divine coincidence specification. This is consistent with the model implication that productivity changes do not generate cost-push shocks in the divine coincidence Phillips curve.

The estimated slopes are of the same order of magnitude as in the calibrated model, both for consumer prices and for the divine coincidence index. Properly accounting for the input-output structure thus allows us to reconcile relatively high price adjustment frequencies at the micro level with flat empirical Phillips curves.

7.3. Rolling Regressions

Figure 4 shows summary statistics for 20-year rolling Phillips curve regressions, again following equation (37), using the CBO employment gap as independent variable and different inflation measures on the left-hand side. Section 3.3 of the additional Additional Online Materials reports the estimated coefficients and standard errors, together with results for the CBO output gap or the unemployment rate as independent variables.

³²From Remark 5, to a first order the employment gap and the output gap coincide in the model. Therefore, both specifications are coherent with the model.

³³These are standard ways to control for inflation expectations. See Mavroeidis, Plagborg-Møller, and Stock (2014) for a detailed survey.

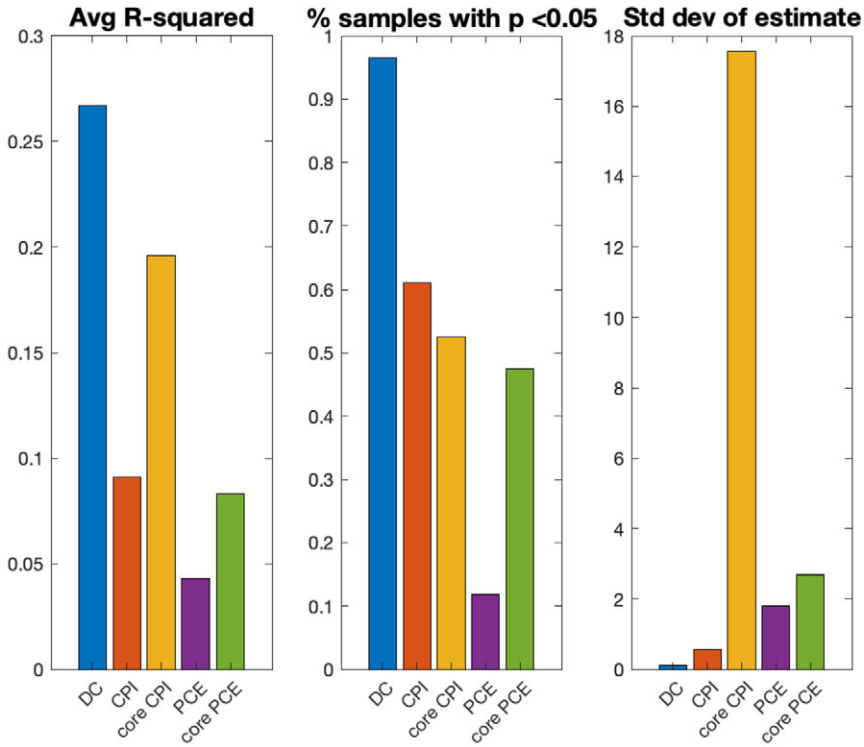


FIGURE 4.—Summary statistics for rolling Phillips curve regressions.

The panels plot three measures of the success of each specification: the R-squared (left panel), the significance of the slope coefficient (middle panel), and its stability over time, measured as the standard deviation relative to the mean (right panel).

The divine coincidence specification performs better in all three dimensions: the R-squared is higher, the slope is almost always significant, and the estimate is more stable over time. In the consumer price Phillips curve, instead, the estimated slope changes over time and is often insignificant, suggesting that this Phillips curve is misspecified, as predicted by the multi-sector model.

8. CONCLUSION

This paper shows that by properly accounting for a realistic, multi-sectoral production structure, we improve our understanding of inflation dynamics and the optimal monetary policy. To demonstrate this, it revisits two key elements of the New Keynesian framework, the Phillips curve and the welfare loss function, in a general multi-sector economy, both analytically and quantitatively.

Our analytical approach allows us to derive several novel results. First, we formalize the argument that all sectoral and aggregate Phillips curves become flatter when producers use more intermediate inputs. Quantitatively, this effect is large enough to reconcile the relatively high price adjustment frequencies measured at the micro level with the flat Phillips curves estimated from aggregate data.

Second, we characterize the inflation-output tradeoff generated by productivity shocks—which manifests as a productivity-driven residual in sectoral and aggregate

Phillips curves—and derive the unique inflation measure which does not suffer from this tradeoff, naming it the divine coincidence index. Given this property, the divine coincidence index is a better sufficient statistic for the output gap than consumer prices. Empirically, we show that, in line with the model predictions, the divine coincidence index performs better in Phillips curve regressions than traditional specifications with consumer prices, increasing the R-squared by two to four times.

These positive results have important normative implications for welfare and optimal monetary policy. Monetary policy makers have only one instrument, the nominal interest rate, which they must use to correct both the output gap and relative price distortions within and across sectors. Therefore, they cannot implement the first-best equilibrium in the multi-sector economy. We demonstrate that the constrained optimal policy aims at stabilizing the output gap rather than inflation, but nonetheless this objective can be achieved via inflation targeting, by replacing consumer price inflation with the divine coincidence index in the Taylor rule.

In the calibrated model, the welfare loss from business cycles increases by several orders of magnitude compared to the Lucas (1977) estimate, even under the constrained optimal monetary policy, due to the joint effect of multiple sectors and price rigidity.

REFERENCES

- AOKI, KOSUKE (2001): “Optimal Monetary Policy Responses to Relative-Price Changes,” *Journal of Monetary Economics*, 48, 55–80. [1420]
- ATALAY, ENGHIN (2017): “How Important Are Sectoral Shocks?” *American Economic Journal: Macroeconomics*, 9, 254–280. [1444,1445]
- BAQAEE, DAVID, AND EMMANUEL FARHI (2020a): “Supply and Demand in Disaggregated Keynesian Economies With an Application to the Covid-19 Crisis,” Working Paper 27152, Working Paper Series, National Bureau of Economic Research, series. [1421]
- BAQAEE, DAVID, EMMANUEL FARHI, AND KUNAL SANGANI (2021): “The Supply-Side Effects of Monetary Policy,” Tech. Rep. w28345, National Bureau of Economic Research. [1421]
- BAQAEE, DAVID REZZA, AND EMMANUEL FARHI (2018): “Macroeconomics With Heterogeneous Agents and Input-Output Networks,” Working Paper 24684, National Bureau of Economic Research, series: Working Paper Series. [1439]
- (2019): “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem,” *Econometrica*, 87, 1155–1203. [1420]
- (2020b): “Productivity and Misallocation in General Equilibrium,” *The Quarterly Journal of Economics*, 135, 105–163. [1420,1423,1426]
- BARATTIERI, ALESSANDRO, SUSANTO BASU, AND PETER GOTTSCHALK (2014): “Some Evidence on the Importance of Sticky Wages,” *American Economic Journal: Macroeconomics*, 6, 70–101. [1444]
- BASU, SUSANTO (1995): “Intermediate Goods and Business Cycles: Implications for Productivity and Welfare,” *The American Economic Review*, 85, 512–531. [1418,1420,1435]
- BENIGNO, PIERPAOLO (2004): “Optimal Monetary Policy in a Currency Area,” *Journal of International Economics*, 63, 293–320. [1420]
- BERAJA, MARTIN, ERIK HURST, AND JUAN OSPINA (2019): “The Aggregate, Implications of Regional Business Cycles,” *Econometrica*, 87, 1789–1833. [1444]
- BERNANKE, BEN S., AND JEAN BOIVIN (2003): “Monetary Policy in a Data-Rich Environment,” *Journal of Monetary Economics*, 50, 525–546. [1421]
- BIGIO, SAKI, AND JENNIFER LA’O (2020): “Distortions in Production Networks,” *The Quarterly Journal of Economics*, 135, 2187–2253. [1420,1426]
- BLANCHARD, OLIVIER (2016): “The Phillips Curve: Back to the ’60s?” *American Economic Review*, 106, 31–34. [1446,1447]
- BLANCHARD, OLIVIER, AND JORDI GALI (2007): “Real Wage Rigidities and the New Keynesian Model,” *Journal of Money, Credit and Banking*, 39, 35–65. [1420,1422,1436]
- CAGLIARINI, ADAM, TIM ROBINSON, AND ALLEN TRAN (2011): “Reconciling Microeconomic and Macroeconomic Estimates of Price Stickiness,” *Journal of Macroeconomics*, 33, 102–120. [1420,1435]
- CARVALHO, CARLOS (2006): “Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks,” *The B.E. Journal of Macroeconomics*, 6, 1–58. [1418,1420,1435,1445]

- CARVALHO, CARLOS, AND FERNANDA NECHIO (2011): "Aggregation and the PPP Puzzle in a Sticky-Price Model," *American Economic Review*, 101, 2391–2424. [1420]
- CASTRO CIENFUEGOS, NICOLAS (2019): "The Importance of Production Networks and Sectoral Heterogeneity for Monetary Policy," Working paper. [1420]
- CHRISTIANO, LAWRENCE J. (2016): "Comment on Networks and the Macroeconomy: An Empirical Exploration," *NBER Macroeconomics Annual*, 30, 346–373. [1418,1420,1435]
- DEL NEGRO, MARCO, MICHELE LENZA, GIORGIO E. PRIMICERI, AND ANDREA TAMBALOTTI (2020): "What's up With the Phillips Curve?" Tech. Rep. 27003, National Bureau of Economic Research, Inc, publication. NBER Working Papers. [1446]
- ERCEG, CHRISTOPHER J., DALE W. HENDERSON, AND ANDREW T. LEVIN (2000): "Optimal Monetary Policy With Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 46, 281–313. [1420,1422,1436]
- EUSEPI, STEFANO, BART HOBIJN, AND ANDREA TAMBALOTTI (2011): "CONDI: A Cost-of-Nominal-Distortions Index," *American Economic Journal: Macroeconomics*, 3, 53–91. [1420]
- GALI, JORDI (2015): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications* (Second Ed.). Princeton: Princeton University Press. [1420,1423,1428,1429]
- GALI, JORDI, AND TOMMASO MONACELLI (2008): "Optimal Monetary and Fiscal Policy in a Currency Union," *Journal of International Economics*, 76, 116–132. [1420]
- GHAASSIBE, MISHEL (2021): "Endogenous Production Networks and Non-Linear Monetary Transmission," Tech. rep. [1420]
- HERRENDORF, BERTHOLD, RICHARD ROGERSON, AND ÁKOS VALENTINYI (2013): "Two Perspectives on Preferences and Structural Transformation," *American Economic Review*, 103, 2752–2789. [1444]
- HOOPER, PETER, FREDERIC S. MISHKIN, AND AMIR SUFI (2019): "Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or Is It Just Hibernating?" Working Paper 25792, Working Paper Series, National Bureau of Economic Research. [1447]
- HOROWITZ, KAREN J., AND KAREN A. PLANTING (2006): "Concepts and Methods of the U.S. Input-Output Accounts," Tech. Rep. 0066, Bureau of Economic Analysis. publication Title: BEA Papers. [1444]
- HOYNCK, CHRISTIAN (2020): "Production Networks and the Flattening of the Phillips Curve," Working paper. [1420]
- HUANG, KEVIN X. D., AND ZHENG LIU (2005): "Inflation Targeting: What Inflation Rate to Target?" *Journal of Monetary Economics*, 52, 1435–1462. [1420]
- JONES, CHARLES I. (2013): "Misallocation, Economic Growth, and Input—Output Economics," in *Advances in Economics and Econometrics: Tenth World Congress: Volume 2: Applied Economics*. Econometric Society Monographs, Vol. 2, ed. by Daron Acemoglu, Eddie Dekel, and Manuel Arellano. Cambridge: Cambridge University Press, 419–456. [1420]
- KARA, ENGIN (2009): "Input-Output Connections Between Sectors and Optimal Monetary Policy," National Bank of Belgium Working Papers. [1420]
- (2010): "Optimal Monetary Policy in the Generalized Taylor Economy," *Journal of Economic Dynamics and Control*, 34, 2023–2037. [1420]
- LA'O, JENNIFER, AND ALIREZA TAHBAZ-SALEHI (2019): "Optimal Monetary Policy in Production Networks," *SSRN Electronic Journal*. [1420,1421]
- LUCAS, ROBERT E. (1977): "Understanding Business Cycles," *Carnegie-Rochester Conference Series on Public Policy*, 5, 7–29. [1419,1449,1453]
- MANKIW, N. GREGORY, AND RICARDO REIS (2003): "What Measure of Inflation Should a Central Bank Target?" *Journal of the European Economic Association*, 1, 1058–1086. <https://onlinelibrary.wiley.com/doi/pdf/10.1162/154247603770383398>. [1420]
- MAVROEIDIS, SOPHOCLES, MIKKEL PLAGBORG-MOLLER, AND JAMES H. STOCK (2014): "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, 52, 124–188. [1421,1451]
- NAKAMURA, EMI, AND JON STEINSSON (2010): "Monetary Non-Neutrality in a Multisector Menu Cost Model," *Quarterly Journal of Economics*, 125, 961–1013. [1435,1445,1446]
- (2013): "Price Rigidity: Microeconomic Evidence and Macroeconomic Implications," Tech. Rep. 18705, National Bureau of Economic Research, Inc, publication. Title: NBER Working Papers. [1418,1420,1446]
- OBERFIELD, EZRA, AND DEVESH RAVAL (2014): "Micro Data and Macro Technology," NBER Working Paper 20452, National Bureau of Economic Research, Inc. [1444]
- ORPHANIDES, ATHANASIOS, AND SIMON VAN NORDEN (2002): "The Unreliability of Output-Gap Estimates in Real Time," *The Review of Economics and Statistics*, 84, 569–583. [1421]
- PASTEN, ERNESTO, RAPHAEL SCHOENLE, AND MICHAEL WEBER (2017): "Price Rigidity and the Origins of Aggregate Fluctuations," Working Paper 23750, National Bureau of Economic Research, series: Working Paper Series. [1420,1444,1450]

- (2019): “The Propagation of Monetary Policy Shocks in a Heterogeneous Production Economy,” *Journal of Monetary Economics*. [1420]
- RUBBO, ELISA (2023): “Supplement to ‘Networks, Phillips Curves, and Monetary Policy,’” *Econometrica Additional Online Materials*, 91, <https://doi.org/10.3982/ECTA18654>. [1420]
- STOCK, JAMES H., AND MARK W. WATSON (1999): “Forecasting Inflation,” *Journal of Monetary Economics*, 44, 293–335. [1421]
- (2007): “Why Has U.S. Inflation Become Harder to Forecast?” *Journal of Money, Credit and Banking*, 39, 3–33. [1446]
- (2016): “Core Inflation and Trend Inflation,” *The Review of Economics and Statistics*, 98, 770–784. [1421]
- WOODFORD, MICHAEL (2003): *Interest and prices: foundations of a theory of monetary policy*. Princeton, NJ; Woodstock, Oxfordshire [England]: Princeton University Press, oCLC. arXiv:ocm52738755. [1420,1423, 1428,1429]

Co-editor Charles I. Jones handled this manuscript.

Manuscript received 2 July, 2020; final version accepted 3 February, 2023; available online 16 March, 2023.