Networks, Phillips Curves, and Monetary Policy

Elisa Rubbo*

June 30, 2020

Click here for the most recent version

Abstract

I develop an analytical framework for monetary policy in a multi-sector economy with a general input-output network. I derive the Phillips curve and welfare as a function of the underlying production primitives. Building on these results, I characterize (i) the correct definition of aggregate inflation and (ii) how the optimal policy trades off inflation in different sectors, based on the production structure. I construct two novel inflation indicators. The first yields a well-specified Phillips curve. Consistent with the theory, this index provides a better fit in Phillips curve regressions than conventional specifications with consumer prices. The second is an optimal policy target, which captures the tradeoff between stabilizing aggregate output and relative output across sectors. Calibrating the model to the U.S. economy I find that targeting consumer inflation generates a welfare loss of 0.8% of per-period GDP relative to the optimal policy, while targeting the output gap is close to optimal.

*Harvard University, Department of Economics (email: elisarubbo@g.harvard.edu). I am grateful to my advisors–Emmanuel Farhi, Elhanan Helpman, Gita Gopinath and Gabriel Chodorow-Reich– for invaluable guidance and continuous support. I thank for their comments Pol Antras, Adrien Auclert, Gadi Barlevy, Kirill Borusyak, Ashley Craig, Francois Gourio, Greg Mankiw, Ken Rogoff, Raphael Schoenle, Jim Stock and Ludwig Straub, and seminar participants at the Chicago Fed, the Cleveland Fed, and the European Central Bank. I am indebted to Andrew Lilley and Giselle Montamat for many insightful conversations all throughout the development of this paper.

1 Introduction

The New Keynesian framework informs the central banks' approach to monetary policy, and constitutes the theoretical foundation underpinning inflation targeting. The baseline New Keynesian model assumes only one sector of production, whereas in reality an economy has multiple and heterogeneous sectors, which trade in intermediate inputs. There are crucial issues that the model is silent about. What is the correct definition of aggregate inflation, based on the production structure? How should central banks trade-off inflation in different sectors, depending on their position in the input-output network?

I extend the New Keynesian framework to account for multiple sectors, arranged in an input-output network. Sectors have arbitrary neoclassical production functions, and face idiosyncratic productivity shocks and heterogeneous pricing frictions. I solve the model analytically, providing an exact counterpart of traditional results in the multi-sector framework.

I derive the two key objects which constitute the "backbone" of the optimal policy problem: the Phillips curve and the welfare loss function. Building on this result, I construct two novel indicators. The first inherits the positive properties of inflation in the one-sector model, and therefore can be viewed as its natural extension to a multi-sector economy. Specifically, this index yields a well-specified Phillips curve and it is stabilized together with aggregate output (a property which is referred to as the "divine coincidence"). The second indicator instead serves as an optimal policy target. These two indicators are distinct, and they are both different from consumer price inflation.[1]

I then explore the quantitative and empirical implications of the model. My representation of production is fully general, and can match any input-output structure. The evolution of the economy is characterized by three variables (the output gap, sectoral inflation and productivity) and a set of steady-state parameters which depend on the production structure and sectoral pricing frictions. I construct time series of these variables and calibrate the parameters for the US economy. The analysis shows

[1]Traditionally, researchers and policy makers take consumer price inflation as the relevant real-world counterpart of inflation in the one-sector model. This choice, however, has no theoretical backing. Previous works argue that consumer prices are not necessarily the relevant indicator for monetary policy. For example, (Galí and Monacelli, 2008) and (Galí, 2015) show that in a small open economy the relevant statistic for the Phillips curve and monetary policy is producer price inflation.

that taking into account the disaggregated structure of the economy is important, not just from a theoretical but also from a quantitative point of view.

The positive analysis is presented in Section 4. I provide a general expression for the Phillips curves associated with any given inflation index. The Phillips curve describes the joint evolution of aggregate inflation ($\bar{\pi}$) and the output gap ($\tilde{y}$):

$$\bar{\pi}_t = \rho \mathbb{E} \pi_{t+1} + \kappa \tilde{y}_t + u_t \tag{1}$$

where ρ is the discount factor, κ is the slope and u_t is a residual. In a multi-sector economy one can construct different measures of aggregate inflation, depending on the weighting of sectoral inflation rates. The slope and residual of the Phillips curve depend on the inflation index $\bar{\pi}_t$ on the left-hand-side, and on the production structure. I derive κ and u_t for a generic choice of $\bar{\pi}_t$. I show that in general these Phillips curves are misspecified, because the residual u_t has an endogenous component which depends on sectoral productivity shocks. Notably, this is true for the traditional Phillips curve specification with consumer prices on the left-hand-side as well. I also show that the presence of intermediate input flows flattens the consumer price Phillips curve, and approximating the multi-sector economy with a one-sector model always leads to overestimating its slope. I then construct a novel inflation measure (the "divine coincidence" index) which instead yields a well-specified Phillips curve, with no endogenous residual and a slope that is independent of the production structure.

To build to these results, I first derive sectoral inflation rates as a function of the output gap and productivity. I then show that the slope of the Phillips curve aggregates sector-level elasticities with respect to the output gap, while the residual aggregates sector-level elasticities with respect to productivity.

The slope of the Phillips curve captures the price response to changes in aggregate demand. When demand is above the efficient level labor supply must also increase, and this requires higher real wages. While the network structure does not affect the relation between aggregate demand and real wages, it is crucial for the pass-through of wages into prices. I demonstrate that intermediate input flows reduce this pass-through, thereby flattening the Phillips curve. Wage changes filter through the network until they reach final prices. If intermediate input prices are sticky, only part of the shock is transmitted to the next producer along the chain. Price rigidities thus get "compounded", reducing the slope of the Phillips curve.

The residual u_t is a time-varying wedge between aggregate output and aggregate prices. In the one-sector benchmark the "divine coincidence" tells us that this wedge cannot result from productivity fluctuations: output is stabilized whenever prices are stabilized. Intuitively, a negative productivity shock increases marginal costs and prices, but this direct effect is perfectly offset by a fall in equilibrium wages (reflecting a lower marginal product of labor). With multiple sectors these two forces no longer offset each other, because sectoral marginal costs are asymmetrically exposed to productivity and wage changes. As a consequence neither sector-level nor aggregate inflation are stabilized under zero output gap.

I then derive the (unique) inflation index that restores the "divine coincidence" in the aggregate. This index weights sectoral inflation rates according to sales shares, appropriately discounting more flexible sectors. Total sales shares, and not final consumption shares, capture the full role of each sector in the production network. Sectors with more flexible prices need to be discounted, because here the same shock generates a larger inflation response.[2]

Section 6 illustrates the quantitative relevance of these results, with a focus on the consumer-price Phillips curve. The network model predicts a slope of around 0.1, consistent with empirical estimates (usually between 0.1 and 0.3). By contrast, the one-sector model implies a slope of about 1. Based on historical input-output tables, the multi-sector model also predicts that the slope has declined by about 30% between 1947 and 2017, a result consistent with empirical estimates.[3] I use sectoral TFP shocks measured in the BEA-KLEMS dataset to construct a time series for the endogenous residual. The series has a standard deviation of 25 basis points, suggesting that endogenous cost-push shocks explain a significant fraction of the variation in consumer inflation.

[2]Interestingly, in the calibrated model the "divine coincidence" index assigns the highest weight (of 18%) to wage inflation. This is because labor has the highest sales share, and wages are quite rigid. Previous contributions (Mankiw and Reis (2003), Blanchard and Gali (2007), Blanchard (2016)) also suggest using wage inflation as an indicator. I provide a formal argument, and characterize the correct weight for wages relative to other sectors.

[3]See for example Blanchard (2016). Blanchard (2016) and other authors attribute the decline in the slope of the Phillips curve to a different channel: with better monetary policy inflation is more stable, therefore firms adjust prices less often. This dampens the response of inflation and reduces the slope of the Phillips curve. I mute this channel by assuming constant frequencies of price adjustment. For many sectors it is impossible to track their evolution over time, due to lack of data. For sectors where data are available, Nakamura and Steinsson (2013) find that the frequency of price adjustment is stable over time.

Section 7 provides an empirical validation of the framework, showing that the "divine coincidence" index is a better indicator of the output gap than consumer prices. I construct a time series of the "divine coincidence" index for the US economy over the years 1984-2017, and compare Phillips curve regressions with this index to standard specifications with consumer prices. In the baseline specification the R-squared is about 0.05 with consumer prices and about 0.2 with the "divine coincidence" index. Rolling regressions over 20 year windows have a stable coefficient and are always significant with the "divine coincidence" index, versus about 50% of the time with consumer prices.

The normative analysis is presented in Section 5. I derive welfare as a function of the output gap and sectoral inflation rates, solve for optimal monetary policy, and construct the inflation target which implements this policy. Targeting the "divine coincidence" index closes the output gap, but this does not implement the optimal policy. While the output gap captures distortions in aggregate demand, with multiple sectors there are also distortions in relative demand across firms and sectors. Relative demand distortions cannot be fully eliminated, thus monetary policy cannot replicate the efficient equilibrium that emerges under flexible prices. These distortions however can be alleviated, at the cost of deviating from the optimal aggregate demand. Closing the output gap therefore is not constrained optimal. In this sense the "divine coincidence" does not hold from a normative point of view, unlike in the baseline model.

Monetary policy has only one instrument (interest rates or money supply), therefore it needs to trade off aggregate demand against allocative efficiency. We argued before that the "divine coincidence" inflation index moves one-to-one with the aggregate output gap. I show that the welfare cost of distortions in relative demand across firms and sectors can also be inferred from sectoral inflation rates. The optimal policy therefore can still be implemented via inflation targeting. The size of relative price distortions depends on how shocks propagate through the input-output network, and their welfare cost depends on the response of quantities demanded, which is governed by the relevant elasticities of substitution in production and consumption. Sectoral weights in the optimal policy target are determined by these two elements.

Targeting consumer inflation, as prescribed by the baseline model, leads to a welfare loss of 1.12% of per-period GDP with respect to a world without pricing frictions. Switching to the optimal policy brings this loss down to 0.28%, but does not fully eliminate it. Closing the output gap instead is almost optimal. Intuitively, the output gap is a good target because monetary policy, being one-dimensional, is a blunt instrument to correct misallocation. Therefore the cost of distorting aggregate demand is larger than the gain in allocative efficiency, and in practice the optimal output gap is close to zero.

Related literature My framework is closely related with Baqaee and Farhi (2019, 2020b), who study markup distortions, aggregate output and welfare in production networks. In the presence of price rigidities markups change endogenously due to productivity and monetary shocks. I solve for the equilibrium response of markups to these shocks, and relate it with the Phillips curve and welfare.

A large literature extends the one-sector New Keynesian model to incorporate realistic elements of the production structure. One branch of this literature studies the Phillips curve and optimal policy in the presence of wage rigidities or in a multi-sector horizontal economy (Aoki (2001), Woodford (2003), Benigno (2004), Blanchard and Gali (2007), Galí (2015), Galí and Monacelli (2008)). Another branch combines nominal rigidities with richer production network models (Basu (1995), Carvalho (2006), Carvalho and Nechio (2011), Nakamura and Steinsson (2013), Pasten et al. (2017, 2019), Castro Cienfuegos (2019), Hoynck (2020), Baqaee and Farhi (2020a)). I contribute to this literature by providing an analytical characterization of the Phillips curve, welfare and optimal monetary policy with a general input-output structure.

My paper is also related with previous works deriving optimal indicators, based on theoretical (Benigno (2004), Galí and Monacelli (2008)) or quantitative arguments (Mankiw and Reis (2003), Eusepi et al. (2011)). I extend previous theoretical results to a more general setup, and my analytical approach allows to relate sectoral weights with the underlying production primitives.

In parallel and independent work, La'O and Tahbaz-Salehi (2019) perform a similar "normative" analysis. In their setup price rigidities are microfounded as arising from incomplete information, while production functions are restricted to be Cobb-Douglass. Because of these modeling differences, the sectoral weights in their optimal targeting rule are determined by the information structure rather than by substitution elasticities.

A large empirical literature documents the limitations of consumer price inflation for Phillips curve regressions and forecasting (Orphanides and van Norden (2002), Mavroeidis et al. (2014)). Many studies seek to construct indicators with better statistical properties (Stock and Watson (1999), Bernanke and Boivin (2003), Stock and Watson (2016)). I show 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 Setup

This section lays out the key elements of the network model and the assumptions about preferences, timing and policy instruments. Section 2.5 introduces the equilibrium concept, which is designed to account for the endogenous evolution of markups under price rigidities.

2.1 Timing and policy instruments

In the main text I consider a one-period model. The dynamic version is presented in Online Appendix 2.

The timing is as follows: before the world begins, firms set prices based on their expectations of productivity and money supply; then sectoral productivities are realized, and the central bank sets money supply; some firms have the possibility to adjust their price after observing the realized productivity and money supply, while others do not; the world ends after production and consumption take place. Inflation is defined as the change in prices with respect to the pre-set ones.

In the static setup money supply is the only policy instrument (to be replaced with interest rates in the dynamic version). I impose that nominal consumption expenditure cannot exceed the aggregate money supply M, so that with incomplete price adjustment an increase in M raises aggregate demand and output.

2.2 Preferences

Consumers derive utility from consumption and leisure, with utility function

$$U = \frac{C^{1-\gamma}}{1-\gamma} - \frac{L^{1+\varphi}}{1+\varphi}$$

L is labor supply. There are N goods produced in the economy, and agents have homothetic preferences over all of these goods. $C(c_1, ..., c_N)$ is their utility from consumption, defined over bundles $(c_1, ..., c_N)$.

Consumers maximize utility subject to the budget constraint

$$PC \leq wL + \Pi - T$$

where P is the price index of the consumption bundle, w is the nominal wage, Π are firm profits (rebated to households) and T is a lump-sum transfer from the government.

In addition, nominal consumption expenditure PC cannot exceed the aggregate money supply M.

2.3 Production

There are N sectors in the economy (indexed by $i \in \{1, ..., N\}$). Within each sector there is a continuum of firms, producing differentiated varieties.

All firms f in sector i have the same constant returns to scale production function

$$Y_{if} = A_i F_i(L_{if}, \{x_{ijf}\})$$

where L_{if} is the amount of labor hired by firm f in sector i, x_{ijf} is the quantity of good j that it uses as input, and A_i is a Hicks-neutral, sector-specific productivity shock.[4] Labor is freely mobile across sectors.

Customers buy a CES bundle of sectoral varieties, with elasticity of substitution ϵ_i.

Cost minimization and markups All producers in sector i solve the cost-minimization problem

$$C_i = min_{\{x_{ij}\}, L_i} \ wL_i + \sum_j p_j x_{ij} \ \ s.t. \ A_i F_i(L_i, \{x_{ij}\}) = \bar{y}$$

Under constant returns to scale marginal costs are the same for all firms, and they use inputs in the same proportions.

[4] Note that this is without loss of generality: factor-biased productivity shocks can be modeled by introducing an additional sector which simply purchases and sells the factor, and letting a Hicks-neutral shock hit this sector.

Before the world begins, all firms set their price optimally based on their expected marginal cost. They solve

$$max_{p_i} \ \mathbb{E} D_i \left(p_i - \left(1-\tau_i\right) mc_i\right) \left(\frac{p_i}{P_i}\right)^{-\epsilon_i}$$

where D_i and P_i are the sector-level demand and price index, and τ_i is an input subsidy provided by the government. The subsidies τ_i are set in order to eliminate the distortions that arise under the CES demand structure, where firms have constant desired markup given by $\mu_i^* = \frac{\epsilon_i}{\epsilon_i - 1} > 1$. This is inefficient, since there are no fixed costs. The optimal subsidies are set so that the resulting markup over pre-subsidy marginal costs is 1, and firms price at marginal costs:

$$1 - \tau_i = \frac{\epsilon_i - 1}{\epsilon_i} \Rightarrow p_i^* = \mathbb{E} mc_i \tag{2}$$

Input subsidies cannot change in response to shocks, and are constrained to be the same for all firms within the same sector.

After productivity and money supply are realized, firms in the same sector end up charging different prices. Those who can adjust their price keep a constant markup equal to the desired one. All other firms need to keep constant prices, and must accept a change in markup given by

$$d \log \mu_{if}^{NA} = -d \log mc_i$$

2.4 Government

The government provides input subsidies to firms, financing them through lump-sum taxes on consumers. It runs a balanced budget, so that the lump-sum transfer equals total input subsidies.

2.5 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. For given sectoral markups I impose market clearing, and further require that the evolution of markups is consistent with the realization of productivity and monetary shocks.

For given output gap, sectoral probabilities of price adjustment δ_i and sectoral productivity shifters, general equilibrium is given by a vector of firm-level markups, a vector of prices p_i, a nominal wage w, labor supply L, a vector of sectoral outputs y_i, a matrix of intermediate input quantities x_{ij}, and a vector of final demands c_i, such that: a fraction δ_i of firms in each sector i adjust their price; markups are optimally chosen by adjusting firms, while they are such that prices stay constant for the non-adjusting firms; consumers maximize utility subject to the budget and cash-in-advance constraint; producers in each sector i minimize costs and charge the relevant markup; and markets for all goods and labor clear.

3 Definitions

I approximate the model by solving for first and second order log-deviations from an equilibrium where productivity and money supply are equal to their expected value. The Phillips curve and welfare are fully characterized by three variables (the output gap, the vector of sectoral inflation rates and the vector of sectoral productivity shifters), and a set of equilibrium parameters, which capture the input-output structure and sector-level pricing frictions. These variables and parameters are defined below.

3.1 Variables

3.1.1 Aggregate output gap

Definition 1. The aggregate output gap $\tilde{y}$ is the log-difference between realized output y and efficient output y_{nat}:

$$\tilde{y} = y - y_{nat}$$

Section A1 in the Supplemental Material derives natural output as a function of productivity.

3.1.2 Sectoral inflation rates

The $N \times 1$ vector of inflation rates is denoted by $\pi = \begin{pmatrix} \pi_1 & ... & \pi_N \end{pmatrix}^T$.

Remark 1. While the output gap captures distortions in aggregate demand, Proposition 4 in Section 5.1 shows that the welfare cost of relative demand distortions across sectors (which is related with sectoral output gaps) can be written as a function of sectoral inflation rates.

3.2 Steady-state parameters

3.2.1 Price rigidity parameters

To model price rigidities, I assume that only a fraction δ_i of the firms in each sector i can adjust their price after observing money supply and productivity. I collect these price adjustment parameters into a diagonal matrix Δ.

Remark 2. This Calvo-style assumption, together with the firms' optimal pricing equation (2), yields a mapping between inflation, marginal costs and markups. The fraction δ_i of firms in each sector i who can adjust prices fully passes-through changes in sectoral marginal costs $d\log mc_i$ into their price.[5] The remaining fraction $1-\delta_i$ is constrained to keep its price fixed, therefore it fully absorbs cost changes into its markup. At the sector level, this implies a markup response equal to $d\log\mu_i = -(1-\delta_i)\, d\log mc_i$, and a change in price given by $\pi_i = \delta_i d\log mc_i$. Therefore, the following relation holds:

$$\pi = \Delta d\log mc = -\Delta\left(I-\Delta\right)^{-1} d\log\mu \tag{3}$$

where $d\log mc$ is the vector of sectoral marginal cost changes, and $d\log\mu$ is the vector of sectoral markups.

Remark 3. Wage rigidities can be easily incorporated into this setup, by adding a labor sector which collects labor services and sells them to all the other sectors. While there still is a flexible underlying wage (paid by the labor sector to workers), the market wage, defined as the price charged by the labor sector, is sticky.

3.2.2 Input-output definitions

The input-output structure is characterized by steady-state consumption, labor and input-output shares. We also introduce two useful derived objects, the Leontief inverse

[5] Remember that desired markups are constant under the CES assumption (see Section 2.3).

and the vector of sales shares, constructed from the input-output matrix and the vector of consumption shares.

Consumption shares The $N \times 1$ vector β denotes sectoral expenditure shares in total consumption, and has components $\beta_i = \frac{p_i c_i}{PC}$.

Labor shares Sector-level labor shares in total sales are encoded in the $N \times 1$ vector α, with components $\alpha_i = \frac{wL_i}{p_i y_i}$.

Input-output matrix The input-output matrix Ω is an $N \times N$ matrix, with elements $\omega_{i,j}$ given by the expenditure share on input j in i's total sales: $\omega_{ij} = \frac{p_j x_{ij}}{p_i y_i}$.

Leontief inverse The Leontief inverse of the input-output matrix Ω is the matrix $(I - \Omega)^{-1}$.

While ω_{ij} is the fraction of sector i revenues directly spent on goods from sector j, the Leontief inverse captures the total (direct and indirect) expenditure of sector i on goods from sector j (again as a share of i's revenues). The indirect component comes from the fact that j's product can be embedded in i's intermediate inputs, if i's suppliers, or i's suppliers' suppliers, etc., use good j in production.

"Adjusted" Leontief inverse The "adjusted" Leontief inverse is the matrix $(I - \Omega\Delta)^{-1}$. The (i, j) element of this matrix is the elasticity of i's marginal cost with respect to j's marginal cost. With price rigidities it is different from the Leontief inverse, because marginal cost changes are not fully passed-through into prices. In this case the "direct" elasticity of i's marginal cost with respect to j's is $\omega_{ij}\delta_j$, which discounts the input share ω_{ij} by the fraction δ_j of producers in j that adjust their price. The total (direct plus indirect) elasticity is then given by $(I - \Omega\Delta)^{-1}_{ij}$.

Sales shares The vector λ of sectoral sales shares in total GDP has components $\lambda_i = \frac{P_i Y_i}{PC}$. It is a well known result that $\lambda^T = \beta^T (I - \Omega)^{-1}$.

Elasticities of substitution The log-linearized model only depends on the input and consumption shares introduced above. Elasticities of substitution in production and consumption instead matter for the second-order welfare loss derived in Section 5.1. I denote the elasticity of substitution between varieties from sector i by ϵ_i, the elasticity of substitution between goods i and j in the production of good k by θ^k_{ij}, and their elasticity of substitution in consumption by σ^C_{ij}; the elasticity of substitution between good i and labor in the production of good k is denoted by θ^k_{iL}.

4 The Phillips curve

The Phillips curve describes the joint evolution of aggregate inflation π^{AGG} and the output gap $\tilde{y}$. The standard New-Keynesian Phillips curve is given by (see for example Galí (2015)):

$$\pi_t^{AGG} = \rho \mathbb{E} \pi_{t+1}^{AGG} + \kappa \tilde{y}_t + u_t \tag{4}$$

where $\mathbb{E}\pi_{t+1}^{AGG}$ is expected future inflation, κ is the slope, ρ is the discount factor and u_t is a residual. In the main text I focus on a one-period model, where the Phillips curve has no forward-looking term:[6]

$$\pi_t^{AGG} = \kappa \tilde{y}_t + u_t \tag{5}$$

The slope κ captures the percentage change in prices when output raises by 1% above the efficient level. The residual u_t captures a time-varying wedge between output and prices. With multiple sectors there are several possible ways to define aggregate inflation, depending on the weighting of sectoral inflation rates. Accordingly, the slope and residual of the Phillips curve are different for different measure of aggregate inflation π^{AGG} on the left-hand-side of (5). For a given inflation index, the slope and residual depend on the production network, and the residual also depends on the realization of sectoral productivity shocks.

Section 4.1 derives the slope and residual of the Phillips curves corresponding to any given aggregate inflation index, as a function of the production structure. Two implications are worth noting. First, for any given weighting of sectoral inflation rates in π_{AGG} the slope of the Phillips curve is decreasing in the size of intermediate input flows. Second, for a generic choice of π_{AGG} productivity shocks result in an endogenous residual $u_t \neq 0$ (i.e. they generate endogenous "cost-push" shocks).

Thus, the multi-sector model provides a "microfoundation" for the flattening of the Phillips curve and for the inflationary effects of productivity shocks. Both are

[6]The dynamic version of the model is derived in Online Appendix 2.

consistent with empirical evidence, but at odds with the one-sector model. Notably, it is widely recognized that shocks to certain sectors, such as the oil sector, raise consumer inflation even if output is stabilized. A key result in the one-sector model however states there can be no endogenous tradeoff between stabilizing output and prices (this result is sometimes referred to as the "divine coincidence", see Blanchard and Gali (2007)). Therefore the inflationary effect of oil shocks can only be represented in a stylized way as an exogenous shock to producers' desired markups (a "cost-push" shock), which makes the model ill-suited for policy analysis.

The results in Section 4.1 imply that the Phillips curves corresponding to a generic inflation index π_{AGG} are mis-specified. Proposition 4.2 in Section 4.2 derives our main "positive" result: it constructs the unique inflation index which yields a well-specified Phillips curve with constant slope and no endogenous residual. This index preserves the "positive divine coincidence", and more in general it inherits all the positive properties of inflation in the one-sector model.

The main text outlines the methods and fundamental results. All the proofs are reported in Sections A2 and A3 of the Supplemental Material, together with additional results.

4.1 Phillips curves for a generic inflation index

4.1.1 Notation and aggregation

To derive the Phillips curves corresponding to any inflation index π_{AGG} I first solve for the equilibrium response of sector-level inflation to productivity and monetary shocks. I then combine sector-level inflation rates into the aggregate Phillips curve based on the weighting prescribed by π_{AGG}.

Formally, I express the vector π of sector-level inflation rates as a function of productivity shocks ($d \log A$) and the output gap ($\widetilde{y}$):

$$\underbrace{\pi}_{N\times 1} = \underbrace{\mathcal{B}}_{N\times 1} \widetilde{y} + \underbrace{\mathcal{V}}_{N\times N} \underbrace{d \log A}_{N\times 1} \tag{6}$$

Here I denote by $\mathcal{B}$ the $N \times 1$ vector whose components $\mathcal{B}_i$ are the elasticities of sector i's price with respect to the output gap, and by $\mathcal{V}$ the $N \times N$ matrix whose elements $\mathcal{V}_{ij}$ are the elasticities of sector i's price with respect to a productivity shock to sector j. Changes in productivity, marginal costs and prices are defined with respect to the flex-price equilibrium where productivity and money supply are equal to their expected value. The elasticities $\mathcal{B}$ and $\mathcal{V}$ are derived in Propositions 1 and 2.

To derive equation (6) I start from the pricing equation (3), which states that sector-level inflation is given by the change in sectoral marginal costs, times the fraction of adjusting firms. This is because firms would like to fully pass-through changes in marginal costs into their prices, but only a fraction Δ of them has the opportunity to do so. Marginal costs in turn depend on wages and productivity, either directly, or indirectly through input prices. While productivity shocks are exogenous, we can solve for the equilibrium wage as a function of the output gap and productivity.[7]

For a given inflation index π_{AGG}, the corresponding Phillips curve is obtained by aggregating both sides of Equation (6). Our inflation index is characterized by the vector of weights ϕ that it assigns to sectoral inflation rates:

$$\pi^{AGG} \equiv \phi^T \pi = \sum_i \phi_i \pi_i$$

Weighting both sides of Equation (6) according to ϕ we obtain the Phillips curve:

$$\pi^{AGG} = \underbrace{\phi^T \mathcal{B}}_{\text{slope}} \tilde{y} + \underbrace{\phi^T \mathcal{V} d \log A}_{\text{residual}} \tag{7}$$

The slope is the aggregate elasticity with respect to the output gap, while the residual is the aggregate elasticity with respect to productivity. Consumer inflation π^C is a special case, obtained by weighting sectoral inflation rates according to consumption shares ($\phi = \beta$).

4.1.2 Slope of the Phillips curve

Proposition 1 derives the elasticities of prices with respect to the output gap sector-by-sector, and aggregates them into the slope of the consumer-price Phillips curve.

[7]This is done the proof of Propositions 1 and 2, reported in Section A2 of the Supplemental Material.

Proposition 1. *The elasticity of sectoral prices with respect to the output gap is*

$$\mathcal{B} = \frac{\overbrace{\Delta\left(\left(I-\Omega\Delta\right)^{-1}\alpha\right)}^{\text{nominal wage pass-through}}}{\underbrace{1-\bar{\delta}_w}_{\text{GE multiplier}}}\underbrace{(\gamma+\varphi)}_{\text{real wage}} \tag{8}$$

where

$$\bar{\delta}_w \equiv \beta^T \Delta \left(I-\Omega\Delta\right)^{-1}\alpha \tag{9}$$

is the pass-through of nominal wages into consumer prices.

The slope κ^C of the consumer-price Phillips curve is given by

$$\kappa^C = \frac{\bar{\delta}_w}{1-\bar{\delta}_w}(\gamma+\varphi) \tag{10}$$

The vector $\mathcal{B}$ and the slope κ^C are the elasticities of sectoral and consumer prices with respect to the output gap. Intuitively, if output is above potential then labor demand must increase. This puts upwards pressure on wages and prices (so that $\mathcal{B} > \mathbf{0}$ and $\kappa^C > 0$). The term $(\gamma + \varphi)$ on the right hand side of (8) and (10) is the effect on real wages, which is governed by the parameters of the labor supply curve and does not depend on the production structure. The remaining component in (8) and (10) is the pass-through of real wages into prices, which instead depends on the input-output network.

Intuitively, wage shocks percolate through the production network all the way to final prices. The vector of labor shares α gives the direct effect of the shock, while its propagation is captured by the adjusted Leontief inverse $(I - \Omega\Delta)^{-1}$ introduced in Section 3.2.2. When producer prices are sticky suppliers do not fully pass-through the wage change to their customers, so that part of the shock is absorbed at every node along the path. As a result the effect on final prices is dampened ($(I - \Omega\Delta)^{-1}\alpha < \mathbf{1}$), which is reflected in a flatter Phillips curve. Corollary 1 in Section A2 of the Supplemental Material formally proves this result.

Without intermediate inputs ($\Omega = \mathbb{O}$, $\alpha = \mathbf{1}$) we have $\bar{\delta}_w = \mathbb{E}_\beta(\delta)$. Corollary 1 shows that in the presence of input-output linkages $\bar{\delta}_w < \mathbb{E}_\beta(\delta)$. The slope of the consumer-price Phillips curve then is

$$\kappa^C = (\gamma + \varphi) \frac{\bar{\delta}_w}{1 - \bar{\delta}_w} < (\gamma + \varphi) \frac{\mathbb{E}_\beta (\delta)}{1 - \mathbb{E}_\beta (\delta)} \tag{11}$$

The right hand side of Equation (11) is the slope predicted by standard calibrations, which directly map the one-sector model into the data without accounting for input-output linkages. The difference between the left and right hand sides of Equation (11) is quantitatively large. Section 6.3 evaluates it for the US economy, finding that the left hand side is one order of magnitude smaller (~ 0.1 against ~ 1).

4.1.3 Endogenous cost-push shocks

Proposition 2 derives the elasticities of sectoral prices with respect to productivity shocks, and aggregates them into the endogenous residual of the consumer-price Phillips curve.

Proposition 2. *The elasticity of sectoral prices with respect to productivity shocks is given by*

$$\mathcal{V} = \Delta (I - \Omega\Delta)^{-1} \left[\frac{\alpha\lambda^T - \beta^T \Delta (I - \Omega\Delta)^{-1}}{1 - \bar{\delta}_w} - I \right] \tag{12}$$

The residual in the consumer-price Phillips curve is given by

$$u^C = \frac{\bar{\delta}_w - \bar{\delta}_A}{1 - \bar{\delta}_w} \lambda^T d \log A \tag{13}$$

where

$$\bar{\delta}_A (d \log A) \equiv \frac{\beta^T \Delta (I - \Omega\Delta)^{-1} d \log A}{\lambda^T d \log A} \tag{14}$$

is the pass-through of the productivity shock into consumer prices, scaled by the aggregate shock.

The elasticity $\mathcal{V}$ captures a direct and an indirect effect of productivity shocks on marginal costs. If aggregate productivity falls, marginal costs increase (direct effect). However equilibrium wages fall (indirect effect), thereby reducing marginal costs. In the flex-price economy real wages are equal to aggregate productivity $\lambda^T d \log A$, which is also the marginal product of labor. As long as output is at the efficient level ($\tilde{y} = 0$), in the sticky-price economy real wages are the same as in the flex-price equilibrium.[8] In the one-sector model the direct and indirect effect exactly offset each other when $\tilde{y} = 0$. This is the key intuition behind the "divine coincidence". With multiple sectors instead marginal costs are asymmetrically exposed to wages and productivity.

Formally, the direct effect of productivity on sectoral prices is given by the second term in (12):

$$\text{direct component} = -\Delta \left(I - \Omega\Delta\right)^{-1} d\log A \tag{15}$$

The adjusted Leontief inverse captures the shock propagation into marginal costs, and the price response is obtained by multiplying the change in marginal costs times the adjustment probability Δ, according to the pricing Equation (3).

The indirect effect through wages is given by the first term in Equation (12):

$$\text{wage component} = \underbrace{\Delta \left(I - \Omega\Delta\right)^{-1} \alpha}_{\text{pass-through}} \ \underbrace{\frac{1 - \bar{\delta}_A}{1 - \bar{\delta}_w}}_{\text{GE multiplier}} \ \underbrace{\lambda^T d\log A}_{\text{real wages}}$$

The term $\lambda^T d\log A$ is the change in real wages, which equals the aggregate productivity shock. The general equilibrium multiplier $\frac{1-\bar{\delta}_A}{1-\bar{\delta}_w}$ maps real wages into nominal wages. The term $\Delta \left(I - \Omega\Delta\right)^{-1} \alpha$ is the pass-through of nominal wages into sectoral prices, which we derived in Section 4.1.2. Note that, while the direct component depends on the full distribution of sectoral productivity shocks, the wage component only depends on the aggregate shock.

In general, at the sector level the wage and productivity pass-through are different. Online Appendix 1 provides some illustrative examples. As a result inflation is not stabilized sector-by-sector, even if the output gap is closed. Proposition 2 shows that consumer inflation is not stabilized either. Its response depends on the relative pass-through of wages and productivity into consumer prices, given by the difference $\bar{\delta}_w - \bar{\delta}_A$.[9] From Equation (13) we see that following a negative shock ($\lambda^T d\log A < 0$) consumer inflation is positive if and only if the productivity pass-through is larger than the wage pass-through ($\bar{\delta}_A > \bar{\delta}_w$). This is the case whenever downstream or flexible sectors are hit by a "worse" shock than the average, as the examples in Online

[8] This result is derived in the proof of Proposition 2 (see Supplemental Material A2).

[9] The productivity pass-through $\bar{\delta}_A$ is defined in Equation (14), mirroring the wage pass-through $\bar{\delta}_w$ introduced in Section 4.1.2. Note that $\bar{\delta}_A$ is scaled by the aggregate shock, and it depends on the full distribution of sectoral productivity shocks (while $\bar{\delta}_w$ is a constant).

Appendix 1 illustrate.

A natural question at this point is whether there are shocks after which prices are stabilized sector-by-sector under zero output gap. Corollary 2 in Section A2 of the Supplemental Material shows that the only shock with this property is an aggregate labor augmenting shock, which in this setup is equivalent to a TFP shock proportional to sectoral labor shares α. This result implies that perfect stabilization is impossible not only in the presence of asymmetric sector-level shocks, but also after an aggregate TFP shock.[10] Indeed, aggregate TFP shocks generate large cost-push shocks in the calibrated model: a 1% negative shock increases consumer inflation by 0.26% under zero output gap.

4.2 The "divine coincidence" inflation index

Section 4.1.3 shows that, for a generic inflation index, the corresponding Phillips curve is mis-specified: the slope changes with the input-output structure, and productivity fluctuations generate endogenous "cost-push" shocks. This is true also of the consumer price Phillips curve. Proposition 3 constructs a novel inflation statistic, the "divine coincidence" index (DC) which eliminates both of these issues.

Proposition 3. *Assume that no sector has fully rigid prices ($\delta_i \neq 0\ \forall i$). Then the inflation statistic*

$$DC \equiv \lambda^T (I - \Delta) \Delta^{-1} \pi$$

satisfies

$$DC = (\gamma + \varphi) \widetilde{y} \tag{16}$$

Unless prices are fully flexible in all sectors ($\Delta = I$), DC is the only aggregate inflation statistic which yields a Phillips curve with no endogenous cost-push term.3

The "divine coincidence" index weights sectors according to sales shares, and discounts more flexible sectors. Sales shares, and not consumption shares, fully capture the value added by each sector to final consumption. Prices respond more to any given cost shock in flexible sectors, thereby the need to discount them.[11]

[10] Except in a horizontal economy without input-output linkages, where aggregate TFP shocks and labor augmenting shocks coincide.

[11] The weights in DC are all positive. We know from Corollary 2 (see Supplemental Material A2) that in general π_i cannot be zero in every sector, therefore we can have $\lambda^T (I - \Delta) \Delta^{-1} \pi = 0$ only if

The proof of Proposition 3 shows that the output gap is inversely proportional to a sales-weighted sum of sectoral markups:

$$(\gamma + \varphi)\,\tilde{y} = -\lambda^T d\log\mu$$

Intuitively, aggregate demand is lower when markups are high, resulting in a negative output gap. Proposition 3 shows that the correct way to aggregate sectoral markups is according to sales shares.[12] As explained in Remark 2 above, we can infer changes in sector-level markups from inflation rates:

$$d\log\mu = -\,(I-\Delta)\,\Delta^{-1}\pi \tag{17}$$

This is because inflation rates reflect the optimal price change implemented by adjusting firms, which is the opposite of the markup change faced by non-adjusting firms. The proportion of adjusting relative to non-adjusting firms in each sector is increasing in the price adjustment probability Δ, therefore for given inflation rates the corresponding markup change is smaller in flexible sectors.

While the relation between the output gap and markups in equation (17) does not rely on the specific pricing assumptions (ex. Calvo), the mapping between markups and inflation rates in equation (17) does depend on the Calvo assumption and on the CES demand structure within sectors.[13] Nonetheless, the Calvo-CES benchmark highlights important forces that are at play also in richer setups. The empirical results in Sections 7.2 and 7.3 show that the "divine coincidence" index based on this model provides a good fit in Phillips curve regressions, much better than consumer price inflation.

π_i is positive in some sectors and negative in others. This implies that under zero output gap there are always sectors where inflation is positive and sectors where it is negative.

[12]This argument is closely related to Proposition 3 in Baqaee and Farhi (2020b).

[13]Crucially, in the Calvo-CES framework the wedge between changes in prices and markups is exogenous and constant (it is given by $(I-\Delta)\,\Delta^{-1}$). This is no longer true under different pricing models, either because the share of adjusting firms is endogenous (as for example in menu cost models), or because the desired pass-through from marginal costs into prices is endogenous (this happens with fixed menu costs, variable adjustment costs or non-CES demand). In general there is no closed form solution for this endogenous wedge.

5 Welfare function and optimal policy

Pricing frictions result in three types of distortions. First, aggregate output is not at the efficient level (i.e. the output gap is $\tilde{y} \neq 0$). Second, adjusting and non-adjusting firms within each sector charge different prices, even though they face the same marginal cost. Customers inefficiently substitute towards the cheaper varieties, resulting in distortions in their relative output. Third, sectoral prices do not fully adjust to reflect their relative productivities, so that relative output across sectors is also distorted.[14] The welfare loss function derived in Proposition 4 gives a formal representation of these three channels.

In the one sector benchmark there are no cross-sector distortions. Moreover the "divine coincidence" implies that stabilizing aggregate output also eliminates within-sector distortions, thereby replicating the efficient allocation. This result no longer holds in the multi-sector model. Even though the "divine coincidence" inflation index is stabilized together with aggregate output, inflation is not stabilized sector-by-sector, and relative prices within and across sectors are distorted.[15] Monetary policy has one instrument (money supply or interest rates) to address all three types of distortions, therefore it cannot replicate the first-best. In this sense the "divine coincidence" fails from a normative point of view.

Specifically, targeting the "divine coincidence" index replicates the efficient aggregate output, but it ignores relative price distortions. Section 5.2 characterizes the optimal monetary policy response to this tradeoff. Section 5.3 shows that the optimal policy can still be implemented by stabilizing an appropriate inflation index, which trades off the "divine coincidence" index against an inflation statistic that captures the effect of monetary policy on relative price distortions. The examples in Online Appendix 1 illustrate the optimal monetary policy in three simple networks.

Remark 4. I derive optimal policy in terms of the aggregate output gap, even though the actual policy instrument is money supply. The two are equivalent, because the consumer-price Phillips curve and the cash-in-advance constraint yield a one-to-one

[14]The second and third channel are conceptually the same. If we considered a fully disaggregated model, where sectors are identified with individual firms, they could be unified into the cross-sector component. For expositional purposes however it is useful to keep them distinct, to facilitate the comparison with the one-sector benchmark.

[15]Corollary 2 in the Supplemental Material A2 shows that perfect stabilization can be achieved only after an aggregate labor augmenting shock.

mapping between output gap and money supply:

$$d\log M = \pi_C + y = \left(1+\kappa^C\right)\widetilde{y} + u^C + \underbrace{\frac{1+\varphi}{\gamma+\varphi}\lambda^T d\log A}_{y_{nat}}$$

All the proofs are reported in Sections B1 and B2 of the Supplemental Material.

5.1 Welfare function

Proposition 4 derives a second-order approximation of the welfare loss relative to the efficient equilibrium with flexible prices. The loss function is quadratic in the output gap (which captures distortions in aggregate output) and inflation (which is associated with distortions in relative output within and across sectors).[16]

Definition 2. The substitution operators Φ_t (for sector t) and and Φ_C (for final consumption) are symmetric operators from $\mathbb{R}^N \times \mathbb{R}^{\mathbb{N}}$ to $\mathbb{R}$, defined as[17]

$$\Phi_t\left(X,Y\right) = \frac{1}{2}\sum_k\sum_h \omega_{tk}\omega_{th}\theta^t_{kh}\left(X_k - X_h\right)\left(Y_k - Y_h\right) + \alpha_t\sum_k \omega_{tk}\theta^t_{kL}X_kY_k$$

and

$$\Phi_C\left(X,Y\right) = \frac{1}{2}\sum_k\sum_h \beta_k\beta_h\sigma^C_{kh}\left(X_k - X_h\right)\left(Y_k - Y_h\right)$$

Proposition 4. *The second-order welfare loss with respect to the flex-price efficient outcome is*

$$\mathbb{W} = \frac{1}{2}\left[\ \left(\gamma+\varphi\right)\tilde{y}^2 + \pi^T\mathcal{D}\pi\ \right] \tag{18}$$

The matrix $\mathcal{D}$ can be decomposed as $\mathcal{D} = \mathcal{D}_1 + \mathcal{D}_2$, where $\mathcal{D}_1$ captures the productivity loss from within-sector misallocation and $\mathcal{D}_2$ captures the productivity loss from

[16] Interestingly, the loss function does not depend on sectoral productivity shocks directly. Intuitively, misallocation is determined by markup distortions. I derive the welfare function around an efficient steady-state, therefore there is no interaction between the productivity shock and initial misallocation (the envelope theorem holds). The welfare loss is entirely driven by the change in markups induced by the shock, which we can infer from sectoral inflation rates (see equation (17)).

[17] Φ_C and Φ_t are the same as in Baqaee and Farhi (2018). They apply these operators to sector-level price changes and labor shares around a distorted steady-state, to derive the first-order change in allocative efficiency. I work around an efficient steady-state where markup shocks have no first-order effect on allocative efficiency, while the substitution operators applied to sector level price changes characterize the second-order loss.

cross-sector misallocation. $\mathcal{D}_1$ is diagonal with elements

$$d^1_{ii} = \lambda_i \epsilon_i \frac{1-\delta_i}{\delta_i} \tag{19}$$

$\mathcal{D}_2$ is positive semidefinite, with elements given by

$$d^2_{ij} = \frac{1-\delta_i}{\delta_i}\frac{1-\delta_j}{\delta_j}\left(\Phi_C\left((I-\Omega)^{-1}_{(i)},(I-\Omega)^{-1}_{(j)}\right) + \sum_t \lambda_t \Phi_t\left((I-\Omega)^{-1}_{(i)},(I-\Omega)^{-1}_{(j)}\right)\right) \tag{20}$$

In the baseline one-sector model the welfare loss is given by

$$\mathbb{W} = \frac{1}{2}\left[\ (\gamma+\varphi)\,\tilde{y}^2 + \tfrac{1-\delta}{\delta}\epsilon\pi^2\ \right] \tag{21}$$

Here inflation only captures within-sector distortions. For a given price distortion, quantities respond more if the elasticity of substitution ϵ is higher. Therefore the welfare cost in Equation (21) is increasing in ϵ. In the network model instead the welfare loss associated with inflation comes from both cross-sector and within-sector price distortions, and the latter need to be appropriately aggregated.

From equation (19) we see that the price dispersion loss within each sector is $\epsilon_i \pi_i^2$, the same as in the one-sector model. Sector-level losses are then aggregated by sales shares, discounting flexible sectors. The intuition is the same as for the "divine coincidence" index in Proposition 3. Overall, the within-sector component of the total welfare loss is given by

$$\pi^T \mathcal{D}_1 \pi = \sum_i \lambda_i \frac{1-\delta_i}{\delta_i}\epsilon_i \pi_i^2$$

The welfare loss from cross-sector misallocation in Equation (20) can be expressed as a weighted sum of sector-level productivity losses:

$$\pi^T \mathcal{D}_2 \pi = \sum_t \underbrace{\lambda_t}_{\text{aggregation}} \underbrace{\sum_{i,j} \Phi_t(i,j)}_{\text{productivity loss in sector } t} \tag{22}$$

Here we treated final consumption as an additional sector with $\lambda_C = 1$, and with some abuse of notation we defined

$$\Phi_t(i,j) \equiv \Phi_t\left((I-\Omega)^{-1}_{(i)}\frac{1-\delta_i}{\delta_i}\pi_i, (I-\Omega)^{-1}_{(j)}\frac{1-\delta_j}{\delta_j}\pi_j\right)$$

Intuitively, relative price distortions induce producers in each sector t to substitute towards the inputs whose relative price is lower than in the efficient equilibrium. The welfare consequence of this misallocation is equivalent to a negative TFP shock for sector t. The total loss is obtained by aggregating sector-level contributions according to sales shares, as in Hulten's formula.

Lemma 11 in section B1 shows that the relative price distortion between inputs k and h induced by inflation in sector i is given by

$$d\log p_k - d\log p_h = \underbrace{\left((I-\Omega)^{-1}_{ki} - (I-\Omega)^{-1}_{hi}\right)}_{\text{propagation}}\underbrace{\frac{1-\delta_i}{\delta_i}\pi_i}_{\text{markup}} \tag{23}$$

Equation (23) highlights an "impulse" component, given by the distortion in i's markup associated with inflation (see Remark 2), and a propagation component, captured by the Leontief inverse $(I-\Omega)^{-1}$.

Producers buy too much of the inputs whose relative price is lower than in the efficient equilibrium. The productivity loss for each sector t is captured by the corresponding substitution operator Φ_t (see Definition 2), and it depends on the interaction between inflation in different sectors. More precisely, $\Phi_t(i,j)$ measures the productivity loss of sector t induced by a 1% increase in i's inflation, given that j's also increased by 1%. Intuitively, the distortions associated with π_i and π_j reinforce each other if they produce similar relative price changes across input pairs (k,h), especially those with higher input shares or higher elasticity of substitution. Correspondingly, $\Phi_t(i,j)$ weights each pair (k,h) by the relevant input shares ω_{ti} and ω_{tj}, and the substitution elasticity θ^t_{kh}:[18]

$$\Phi_t(i,j) = \underbrace{\omega_{tk}\omega_{th}}_{\text{input shares}}\underbrace{\theta^t_{kh}}_{\text{substitution}}\underbrace{\left((I-\Omega)^{-1}_{ki} - (I-\Omega)^{-1}_{hi}\right)\frac{1-\delta_i}{\delta_i}\pi_i}_{\text{distortion from } i}\underbrace{\left((I-\Omega)^{-1}_{kj} - (I-\Omega)^{-1}_{hj}\right)\frac{1-\delta_j}{\delta_j}\pi_j}_{\text{distortion from } j}$$

[18] In the interesting special case where elasticities of substitution are uniform ($\theta^t_{kh} \equiv \theta^t$), the substitution operator coincides with the covariance between the price distortions induced by i and j across sector pairs (k,h), with probability weights given by t's input shares $\{\omega_{tk}\}_{k=1..N}$.

The total productivity loss in sector t is obtained by summing the contributions of all pairs (i, j):

$$\text{Loss in } t = \sum_{i,j} \Phi_t(i, j)$$

and the aggregate productivity loss is given by Hulten's formula, as in Equation (22).

5.2 Optimal policy

Optimal monetary policy minimizes the welfare loss derived in Proposition 4, subject to the response of inflation to the output gap and productivity shocks.

In the one-sector model the central bank solves

$$\begin{aligned} min_{\pi,\tilde{y}} \quad \mathbb{W} = \tfrac{1}{2}\left[\ (\gamma+\varphi)\,\tilde{y}^2 + \epsilon\tfrac{1-\delta}{\delta}\pi^2\ \right] \\ s.t.\ \pi = \kappa\tilde{y} \end{aligned} \tag{24}$$

Here the constraint is given by the aggregate Phillips curve. The "divine coincidence" implies that there is no tradeoff between stabilizing output and stabilizing prices, therefore the optimal policy achieves the first best by setting $\pi = \tilde{y} = 0$.

With multiple sectors the optimal policy problem extends this baseline in two dimensions. First, the inflation term is replaced by the more complex misallocation loss derived in Proposition 4, which captures both within- and cross-sector distortions. Second, the constraint is not just the aggregate Phillips curve, but it is given by the full vector of sectoral Phillips curves. Thus the problem becomes:

$$\begin{aligned} min_{\tilde{y},\pi}\ \tfrac{1}{2}\left[\ (\gamma+\varphi)\,\tilde{y}^2 + \pi^T\mathcal{D}\pi\ \right] \\ s.t.\ \pi = \mathcal{B}\tilde{y} + \mathcal{V}d\log A \end{aligned} \tag{25}$$

Proposition 5. *The value of the output gap that minimizes the welfare loss is*

$$\widetilde{y}^* = -\frac{\mathcal{B}^T\mathcal{D}\mathcal{V}d\log A}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}} \tag{26}$$

Proposition 5 follows immediately from the first order conditions of the minimization problem (25). The optimal policy trades off the marginal cost and benefit of deviating from the efficient aggregate output. The denominator in equation (26) reflects the marginal cost, and it is always positive. It comes from distortions in aggregate demand (whose welfare effect is proportional to the labor supply elasticities $(\gamma+\varphi)$), and from the relative price distortions that they trigger (captured by the term $\mathcal{B}^T\mathcal{D}\mathcal{B}$).

The numerator in (26) is the marginal gain. For given current inflation π, the marginal benefit of inducing inflation $\tilde{\pi}$ is $-\tilde{\pi}^T\mathcal{D}\pi$. For $\tilde{y}=0$ and a given productivity shock $d\log A$, inflation is given by

$$\pi = \mathcal{V}d\log A$$

Increasing the output gap raises inflation by $\tilde{\pi}=\mathcal{B}$. Therefore the overall marginal gain is given by

$$-\tilde{\pi}^T\mathcal{D}\pi = -\mathcal{B}^T\mathcal{D}\mathcal{V}d\log A$$

The constraint tells us that monetary policy has limited effect on misallocation, because it can only implement relative price changes which are proportional to the vector $\mathcal{B}$ of sectoral elasticities with respect to the output gap.

5.3 Inflation targeting

In the one sector model the optimal output gap is always zero, regardless of productivity. Thanks to the "divine coincidence", the optimal policy can be implemented equivalently by targeting inflation or the output gap. This is particularly useful in the policy practice, because the output gap and productivity are difficult to measure in real time.

Proposition 6 demonstrates that the multi-sector framework preserves the convenient implementation properties of the one sector model, in that the optimal policy can still be implemented by stabilizing an appropriate inflation index.

Proposition 6. *Assume that no sector has fully rigid prices. Then there exists a unique vector of weights ϕ (up to a multiplicative constant) such that the aggregate inflation*

$$\pi_\phi = \phi^T\pi$$

is positive if and only if $\widetilde{y} > \widetilde{y}^*$. The optimal target has

$$\phi^T = \lambda^T(I-\Delta)\Delta^{-1} + \mathcal{B}^T\mathcal{D} \tag{27}$$

To build intuition, note that the first order condition from the policy problem (25) can be written as

$$(\gamma + \varphi)\,\tilde{y} + \mathcal{B}^T \mathcal{D}\pi = 0 \tag{28}$$

The policy target (27) can be immediately derived from Equation (28), just replacing the output gap with the divine coincidence inflation index (see Proposition 3).

Consistent with our discussion in Section 5.2, the optimal target weights the output gap against sectoral inflation rates according to the relative marginal benefit ($-\mathcal{B}^T \mathcal{D}\pi$) and marginal cost ($\gamma + \varphi$) of distorting aggregate output to reduce misallocation. This result extends to the dynamic setup (see Online Appendix 2), just adding a correction for inflation expectations.

6 Quantitative analysis

6.1 Data

The multi-sector economy is fully characterized by the variables and parameters introduced in Section 3.2. The parameters consist in labor, input and consumption shares (α, Ω and β), sectoral frequencies of price adjustment (Δ), and elasticities of substitution in production and consumption. To compute the expected welfare loss from business cycles (see Section 6.2) we also need the variance of sectoral productivity shocks.

I calibrate labor, input and consumption shares based on the input-output tables published by the BEA.[19] I use tables for the year 2012, because this is the most recent year for which they are available at a disaggregated level (405 industries). Section 6.3.1 relies on less disaggregated historical input-output data (46 - 71 industries), always from the BEA input-output accounts, to study the slope of the Phillips curve and monetary non-neutrality over time.

[19] The BEA does not provide a direct counterpart to the input-output matrix Ω, however this can be constructed from the available data. The BEA publishes two direct requirement tables, the Make and Use table, which contain respectively the value of each commodity produced by each industry and the value of each commodity and labor used by each industry and by final consumers. In addition the BEA publishes an Import table that reports the value of commodity imports by industry. The Make and Use matrix (corrected for imports) can be combined, under proportionality assumptions, to compute the matrix Ω of direct input requirements and the labor and consumption shares α and β.

I calibrate industry-level frequencies of price adjustment based on estimates constructed by Pasten et al. (2019). For sectors with missing data I set the adjustment probability equal to the mean. I set the quarterly probability of wage adjustment to 0.25, in line with Barattieri et al. (2014) and Beraja et al. (2019).

I choose values for the elasticities of substitution across inputs and consumption goods based on estimates from the literature. I set the substitution elasticity between consumption goods to $\sigma = 0.9$,[20] the elasticity of substitution between labor and intermediate inputs to $\theta_L = 0.5$,[21] the elasticity of substitution across intermediate inputs to $\theta = 0.001$,[22] and the elasticity of substitution between varieties within each sector to $\epsilon = 8$.[23]

I calibrate sectoral TFP shocks and their covariance matrix based on estimates of annual industry-level TFP changes for the period 1988-2016 from the BEA Integrated Industry-Level Production Account data.[24] I refer to the Multifactor Productivity (MFP) measure, and calibrate productivity shocks as the growth rate of this index at the sector level.[25]

6.2 Welfare loss from business cycles

In the one-sector model the "divine coincidence" implies that productivity fluctuations do not generate additional welfare losses with respect to an efficient economy with flexible prices. In this economy the welfare loss from business cycles is very small: the well-known Lucas' estimate is about 0.05% of per-period GDP.[26]

Section 5.2 argues that in a multi-sector economy monetary policy cannot replicate the flex-price efficient outcome, which creates the potential for larger welfare losses. In this section I calibrate the loss relative to the efficient economy under different policy rules. I assume that productivity shocks are normally distributed with zero mean and covariance matrix Σ, which I calibrate from BEA-KLEMS data. I compare

[20] Atalay (2017), Herrendorf et al. (2013), and Oberfield and Raval (2014) estimate it to be slightly less than one.

[21] This is consistent with Atalay (2017), who estimates this parameter to be between 0.4 and 0.8.

[22] See Atalay (2017).

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

[24] https://www.bea.gov/data/special-topics/integrated-industry-level-production-account-klems

[25] The 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] This welfare cost comes entirely from the uncertainty generated by fluctuations in consumption.

the results with counterfactual calibrations which assume only idiosyncratic or only aggregate shocks, keeping constant the variance of aggregate output. More detailed results are reported in Section D1 of the Supplemental Material.

Quantitatively, the departures from the one sector benchmark are significant. There is a large loss from imperfect stabilization, equal to 0.28% of per-period GDP under the optimal policy. This means that the additional loss induced by price rigidities is one order of magnitude larger than the Lucas' estimate. The idiosyncratic component of productivity is the main driver. Input-output linkages are key in determining these results. In a counterfactual calibration without input-output linkages, but with the same productivity shocks and price adjustment frequencies,[27] the welfare loss is only 0.11%.

The loss increases under suboptimal policy rules. Targeting consumer prices, which is first best in the one sector model, brings it to 1.12% of per-period GDP. Again the loss is much smaller (0.12%) in the calibration without input-output linkages, regardless of the distribution of the shocks.

Targeting zero output gap instead yields a negligible additional loss on average with respect to the optimal policy. Although monetary policy faces a tradeoff between stabilizing aggregate demand (the output gap) and relative demand across sectors (see Section 5.1), the fact that it has only one instrument makes it inefficient at correcting relative price distortions. Therefore in practice the optimal policy should focus on aggregate demand. Figure 6 in Section D1 of the Supplemental Material corroborates this argument by showing that the optimal policy target tracks the "divine coincidence" index very closely over time. The optimal target however is often a few basis points below DC, suggesting that the optimal policy should be slightly more expansionary than output gap targeting.

The Supplemental Material also provide analytical expressions for the welfare loss under different policy rules, and a decomposition of the welfare loss between within- and cross-sector misallocation.

6.3 Slope of the Phillips curve and monetary non-neutrality

Section 4.1.2 establishes that the presence of intermediate inputs reduces the slope of the Phillips curve. To evaluate the quantitative importance of this result I carry out

[27] Consumption shares are calibrated to replicate relative sales shares.

two exercises. First, I compute the slope of the Phillips curve based on the input-output tables for 2012, under different assumptions about input-output linkages, wage rigidities and pricing frictions. I then compute the slope implied by the model based on historical input-output tables from 1947 to 2017. I find that the calibrated slope has flattened by 30%, due to an increase in intermediate input flows.

The slope of the Phillips curve is also related with monetary non-neutrality,[28] which is a measure of the effectiveness of monetary policy. Section 6.3.2 shows that both input-output linkages and heterogeneous pricing frictions increase monetary non-neutrality.

6.3.1 Slope of the Phillips curve

Table 1 shows that input-output linkages and wage rigidity flatten the Phillips curve, while heterogeneous adjustment frequencies play no role. In the baseline calibration (first column) the slope is 0.09, which is in the same ballpark as empirical estimates (see Section 7.2). The second column reports the slope implied by an alternative calibration which directly maps the one-sector model to the data, ignoring input-output linkages and wage rigidity. Here the slope is more than one order of magnitude larger than in the baseline. The third column reports the slope in a calibration with sticky wages, but without input-output linkages. We find that the implied slope more than doubles with respect to the baseline calibration.

	full model	no IO, flex w	no IO	δ = mean
slope	0.09	1.16	0.22	0.08
slope relative to full calibration	1.00	0.07	0.38	1.05

Table 1: Phillips curve slope in the main and alternative calibrations

Finally, the last column shows that eliminating heterogeneity in adjustment frequencies does not affect the calibrated slope. This is not a general result, but it depends on the specific joint distribution of labor shares and adjustment frequencies that we observe in the data. Heterogeneity in price stickiness instead matters in the dynamic version of the model, where it increases monetary non-neutrality (see Section 6.3.2).

[28]See section 6.3.2 below for a discussion.

Figure 1 plots the slope of the Phillips curve implied by the model for each year between 1947 and 2017, based on historical input-output data. The blue solid line

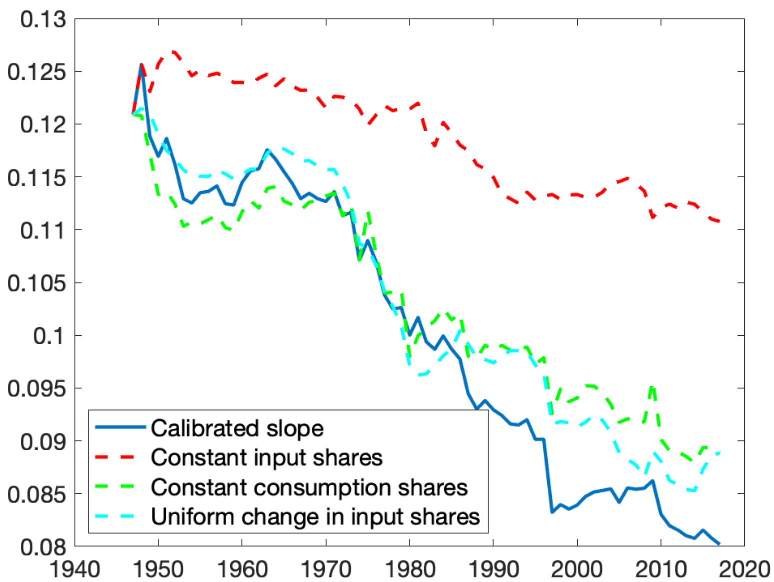

Figure 1: Slope of the Phillips curve over time

depicts the calibrated slope, which has decreased by about 30% over this time period. This result is consistent with the conventional wisdom that the Phillips curve has flattened (see for example Blanchard (2016)).[29] The flattening comes from two channels. The first is an increase in intermediate input flows (green line), while the second is a shift in consumption shares (red line), away from manufacturing and towards services.

Figure 1 highlights that most of the flattening can be attributed to changes in the input-output structure, rather than to shifts in consumption shares. Section D2 of the Supplemental Material explains how these two channels are measured based on the model, and provides a sector-level breakdown of the consumption and input components.

[29]It is difficult to evaluate which fraction of the observed flattening is explained by changes in the input-output structure relative to other factors, given that we do not have consensus estimates of the slope of the Phillips curve at any point in time (see Mavroeidis et al. (2014)). The calibration suggests that the input-output structure played an important role. Nonetheless, the fact that the calibrated slope at the beginning of the period is low compared to conventional estimates suggests that other channels, such as the anchoring of inflation expectations, might be relevant as well.

6.3.2 Monetary non-neutrality

I use the dynamic version of the model (derived in Online Appendix 2) to study the effect of input-output linkages and heterogeneous pricing frictions on monetary non-neutrality.[30] Monetary policy is less neutral (i.e. more effective) if the same shock to the interest rate path results in a smaller inflation response and a larger output response. Figure 2 plots the impulse response of consumer inflation and output to a 10 bp shock to the nominal interest rate, with persistence 0.9, under a standard Taylor rule with $\varphi_\pi = 1.24$ and $\varphi_y = .33/12$.

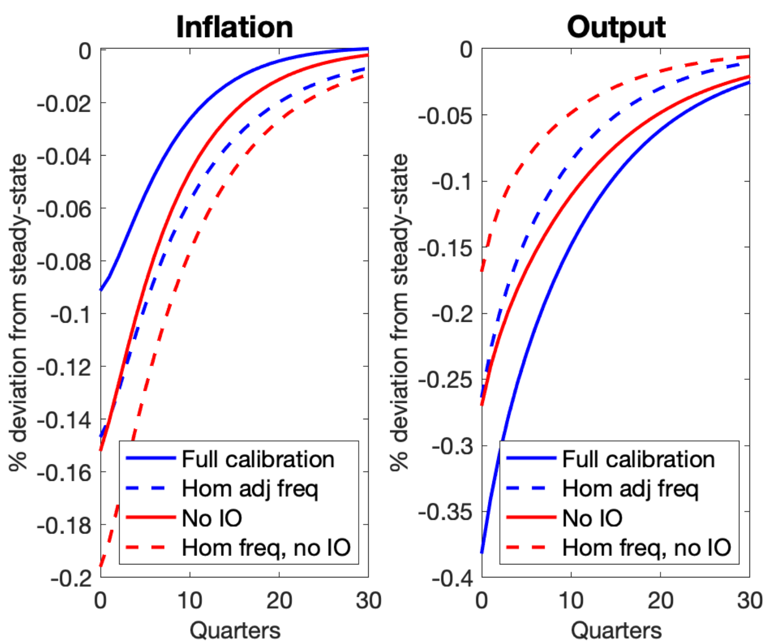

Figure 2: Impulse response to a 10 bp nominal rate shock

In the dynamic model both input-output linkages and heterogeneous adjustment frequencies increase monetary non-neutrality, even though eliminating heterogeneity in adjustment frequencies did not affect the slope of the Phillips curve in the static setup. To gain intuition 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

[30] The results in this section are consistent previous work, such as Carvalho (2006) and Nakamura and Steinsson (2010).

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. This is why heterogeneous adjustment frequencies play a different role in the dynamic versus the static setting.

Section D2 in the Supplemental Material discusses the evolution of monetary non-neutrality over time.

6.3.3 Wage Phillips curve vs consumer price Phillips curve

Empirical studies (see for example Hooper et al. (2019)) found that the wage Phillips curve is steeper than the price Phillips curve, and it has not flattened over time (or at least not as much as the price Phillips curve). This evidence is consistent with the predictions of the multi-sector model. The calibrated slope of the wage Phillips curve is 0.78 for 1947 and 0.77 for 2017, much larger than for the price Phillips curve and constant over time.

6.4 Endogenous cost-push shocks

As explained in Section 4.1.3, in the multi-sector model productivity shocks can generate an "endogenous" tradeoff between stabilizing prices and output. Section 6.4.1 below demonstrates that this phenomenon is quantitatively important in the case of oil shocks. It also shows that the optimal policy response to a negative oil shock is to implement a positive output gap, even if this raises inflation. Section 6.4.2 instead uses measured sectoral productivity shocks to construct a time series of the Phillips curve residual derived in Proposition 2, which captures the endogenous inflation-output tradeoff generated by these shocks. I find that adding this variable to otherwise standard Phillips curve regressions significantly increases the R-squared (see Section 7.2 below).

6.4.1 Oil shocks

Example 3 in Online Appendix discusses the channels through which negative oil shocks raise consumer inflation in a stylized model. Even if the actual US network is much more complex, the example captures well the mechanisms at play. Our simple model highlights three elements: the presence of wage rigidities, the presence of a positive correlation between oil shares and adjustment frequencies, and the fact that oil prices are very flexible. Table 2 compares the inflation response to oil shocks in the full calibration versus alternative calibrations that shut down each of these channels, showing that all of them are important. Overall the inflation response is sizable in the baseline calibration, equal to 0.22 for a 10% negative oil shock.

	$\delta =$ actual	$\delta \equiv \delta_{mean}$, $\delta_{oil} = 1$	$\delta \equiv \delta_{mean}$
sticky wages	0.22	0.07	-0.00
flexible wages	0.18	0.01	-0.06

Table 2: Consumer inflation after a 10% negative shock to the oil sector (full model)

To complement the discussion in Example 3, Table 3 presents the optimal monetary policy response to a 10% negative oil shock. Here policy is expressed in terms of the optimal output gap (in percentage points). The implied percentage change in output is obtained by adding the log change in natural output, $y_{nat} = -0.69$. The calibration suggests that the central bank should implement a positive output gap in response to negative oil shocks, even though this raises inflation.

	full model	$\delta = \delta_{mean}$, $\delta_{oil} = 1$	$\delta = \delta_{mean}$
sticky wages	0.11	0.16	0.18
flex wages	-0.03	0.06	0.09

Table 3: Optimal output gap (in percentage points) after a 10% negative oil shock

6.4.2 Time series

I construct a time series for the residual u^C in the consumer-price Phillips curve, derived in Proposition 2, based on BEA-KLEMS data. I proxy for productivity shocks using their sector-level measures of yearly TFP growth. Figure 3 plots the results.

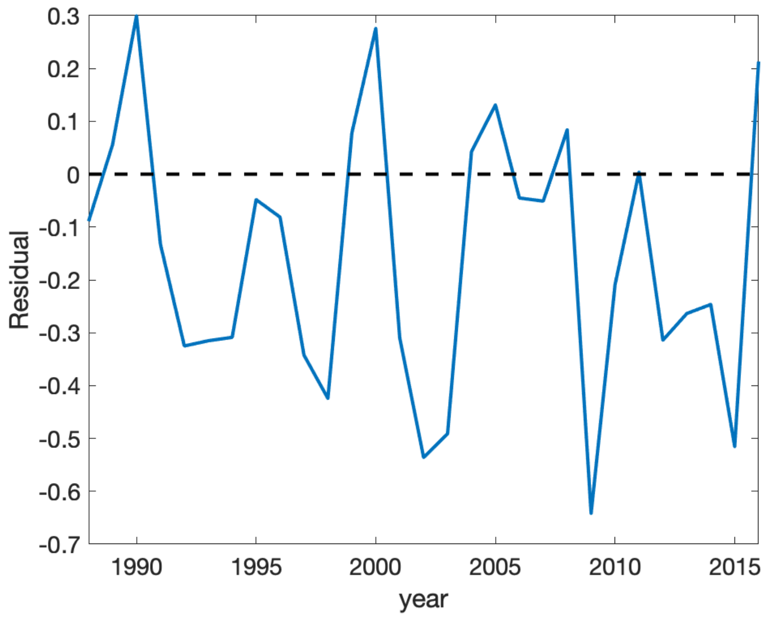

Figure 3: Time series of the endogenous residual

The estimated residual has a mean of -0.16 and a standard deviation of 0.25. Both mean and standard deviation are large relative to the calibrated slope of the consumer-price Phillips curve, which is 0.09. This suggests that endogenous "cost-push" shocks are a significant component of the variation in consumer price inflation. In Section 7.2 I add the endogenous residual u^C plotted in Figure 3 to an otherwise standard Phillips curve regression. I find that the R-squared increases significantly, getting close to the "divine coincidence" specification.

7 Phillips curve regressions

In this section I run Phillips curve regressions using different inflation measures as left-hand-side variables (various measures of consumer price inflation and the "divine coincidence" inflation index). I compare the estimated coefficients and R-squareds.

The estimation results validate my theoretical framework. First, the R-squared is 2 to 4 times higher when using the "divine coincidence" index on the left-hand-side. This is consistent with Proposition 3: the explanatory power of the output gap should be maximal for the "divine coincidence" index, because the corresponding Phillips curve is the only one without an endogenous residual. Second, the calibrated model predicts the estimated slopes correctly for both consumer prices and the "divine coincidence" index. Third, controlling for the endogenous cost-push shocks constructed in Section 6.4.2 increases the R-squared of the consumer-price Phillips curve, bringing it in the same ballpark as the "divine coincidence" specification.

Rolling regressions confirm that these results are robust to the choice of a sample period: the estimated coefficient is stable and always significant when using the "divine coincidence" index as left-hand-side variable, in contrast with traditional consumer price specifications.

7.1 Data

I construct a time series of the "divine coincidence" index DC for the US economy based on sector-level PPI data from the BLS. I measure inflation as the percentage price change from the same quarter of the previous year. I aggregate sectoral inflation rates based on sales shares implied by the BEA input-output tables, and on sector-level price adjustment frequencies constructed by Pasten et al. (2019).

A detailed comparison between the weighting of sectoral inflation rates in traditional measures of consumer prices (CPI and PCE) and in DC is reported in Section E1 of the Supplemental Material. A key difference is that consumer prices place no weight on wage inflation, which instead has a weight of 18% in DC. Other important sectors in DC are professional services, financial intermediation and durable goods. Consumer prices instead place high weight on health care, real estate and nondurable goods. The Supplemental Material also include plots of DC against consumer price inflation (CPI and PCE) and aggregate producer price inflation (PPI), and scatter-plots of the output gap against DC and consumer inflation.

I focus on a regression specification with no lags and a proxy for inflation expectations, which is consistent with the dynamic model. I construct a proxy for inflation expectations based on the statistical properties of the inflation process, whose changes are well approximated by an IMA(1,1) (see Stock and Watson (2007)). I estimate the IMA(1,1) parameters and use them to construct a forecast series for each of the inflation measures that I use in the regressions. For consumer inflation it has been shown that survey measures of forecasted inflation (such as the SPF) are well approximated by this IMA(1,1) forecast. The forecast series are plotted in Section E2 of the Supplemental Material.

7.2 Regressions over the full sample period

The results presented here use the CBO unemployment gap as a right-hand-side variable. Section E3 of the Supplemental Material shows that the results are robust when using two other measures of the output gap: the CBO output gap and the unemployment rate.

Table 4 reports results for a simple specification with no lags or expectations:

$$\pi_t = c + \kappa \tilde{y}_t + u_t \tag{29}$$

A specification with inflation expectations (consistent with the dynamic model) is reported in Section E3 of the Supplemental Material.

	DC	CPI	core CPI	PCE	core PCE
gap	-3.8814	-0.2832	-0.1839	-0.1667	-0.1007
	(0.6329)	(0.0729)	(0.0642)	(0.0628)	(0.0565)
intercept	1.9842	2.9052	2.9021	2.3978	2.372
	(0.0475)	(0.1196)	(0.1052)	(0.103)	(0.0926)
R-squared	0.2154	0.0991	0.0566	0.0489	0.0227

Table 4: Regression results for the CBO unemployment gap

Two results are worth noting. First, the R-squared is much higher when using the "divine coincidence" inflation index on the left-hand-side. This is consistent with the fact that the "divine coincidence" index Phillips curve is the only one without an endogenous residual (see Proposition 3). Second, the calibrated model predicts well the estimated slope, for both consumer prices and the "divine coincidence" index. The slope implied by the calibrated model is 0.09 for the consumer-price Phillips curve and 3 for the "divine coincidence" Phillips curve.[31]

[31] The model predicts a higher slope when using the "divine coincidence" index, consistent with the fact that the weights in this index have a larger sum than for consumer prices (where they always sum to 1). The mapping between sectoral weights and the slope of the corresponding Phillips curve however is non-trivial, and relies on the propagation mechanism described in Section 4.1.2. Therefore our result can be viewed as a validation of this mechanism.

As a further validation of my theoretical framework, I run a specification that augments (29) to include the time series of the endogenous residual constructed in Section 6.4.2. Including the endogenous residual brings the R-squared for both CPI and PCE close to the "divine coincidence" specification, but it does not affect the core versions.[32] The full results are reported in Section E3 of the Supplemental Material, which also reports additional specifications including lags and inflation changes, together with residual plots.

7.3 Rolling regressions

I run rolling Phillips curve regressions with a 20 year window, over the period January 1984 - July 2018. I report results for specification (29) with inflation expectations, using the CBO unemployment gap as right-hand-side variable. Section E5 of the Supplemental Material reports results for different measures of the output gap and other specifications.

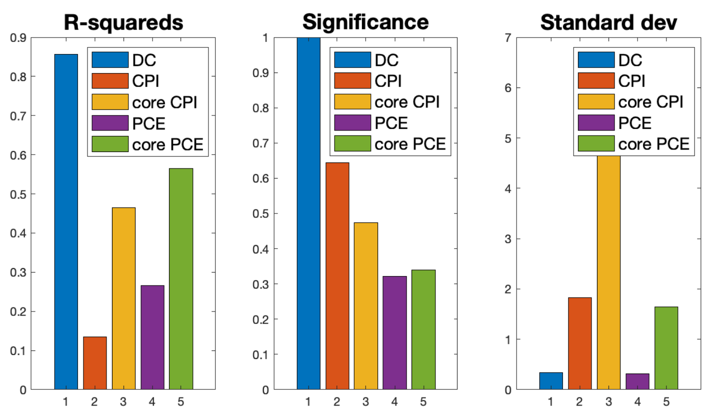

Figure 4: Summary statistics for rolling Phillips curve regressions

Figure 4 compares the strength and stability of the estimated relation for different left-hand-side variables. The left panel reports the average R-squared over the sample period, the middle panel reports the fraction of windows in which the estimated coefficient is significant, and the right panel plots the standard deviation relative to the mean of the estimated coefficient, as a measure of its stability over time. The figure shows that *DC* dominates consumer prices along all three dimensions. Plots

[32]This is consistent with the model, because core inflation excludes flexible sectors (such as food and energy) which are among the main drivers of the residual.

of the rolling coefficients and confidence intervals are reported in Section E6 of the Supplemental Material.

8 Conclusion

This paper develops a New Keynesian framework in an economy with multiple sectors, arranged in a general input-output network. I provide the exact multi-sector counterpart of traditional results. I derive analytical expressions for the Phillips curve and welfare as a function of the underlying production primitives, and construct two novel indicators (the "divine coincidence" index and the optimal policy target) which inherit the positive and normative properties of inflation in the one-sector model. I calibrate the model to the US economy, finding quantitatively important departures from the one-sector benchmark.

The consumer-price Phillips curve is flatter than in the baseline model, and productivity shocks generate an endogenous inflation-output tradeoff. These predictions are new, and consistent with empirical evidence. I further validate my framework by showing that the "divine coincidence" index implied by the model provides a better fit for Phillips curve regressions than traditional specifications with consumer prices.

I also evaluate the performance of the two standard targets in the Taylor rule, the output gap and consumer inflation, against the optimal policy. I find that targeting the output gap is close to optimal, while stabilizing consumer prices generates an expected loss of 0.8% of per-period GDP relative to the optimal policy.

References

Aoki, K. (2001): "Optimal monetary policy responses to relative-price changes," *Journal of Monetary Economics*, 48, 55–80.

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

Baqaee, D. and E. Farhi (2020a): "Supply and Demand in Disaggregated Keynesian Economies with an Application to the Covid-19 Crisis," Working Paper 27152, National Bureau of Economic Research, series: Working Paper Series.

Baqaee, D. R. and E. Farhi (2018): "Macroeconomics with Heterogeneous Agents and Input-Output Networks," Working Paper 24684, National Bureau of Economic Research, series: Working Paper Series.

——— (2019): "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 87, 1155–1203.

——— (2020b): "Productivity and Misallocation in General Equilibrium," *The Quarterly Journal of Economics*, 135, 105–163, publisher: Oxford University Press.

Barattieri, A., S. Basu, and P. Gottschalk (2014): "Some Evidence on the Importance of Sticky Wages," *American Economic Journal: Macroeconomics*, 6, 70–101.

Basu, S. (1995): "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare," *The American Economic Review*, 85, 512–531, publisher: American Economic Association.

Benigno, P. (2004): "Optimal monetary policy in a currency area," *Journal of International Economics*, 63, 293–320.

Beraja, M., E. Hurst, and J. Ospina (2019): "The Aggregate Implications of Regional Business Cycles," *Econometrica*, 87, 1789–1833.

Bernanke, B. S. and J. Boivin (2003): "Monetary policy in a data-rich environment," *Journal of Monetary Economics*, 50, 525–546.

Blanchard, O. (2016): "The Phillips Curve: Back to the '60s?" *American Economic Review*, 106, 31–34, publisher: American Economic Association.

Blanchard, O. and J. Gali (2007): "Real Wage Rigidities and the New Keynesian Model," *Journal of Money, Credit and Banking*, 39, 35–65, publisher: Blackwell Publishing.

Carvalho, C. (2006): "Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks," *The B.E. Journal of Macroeconomics*, 6, 1–58.

Carvalho, C. and F. Nechio (2011): "Aggregation and the PPP Puzzle in a Sticky-Price Model," *American Economic Review*, 101, 2391–2424.

Castro Cienfuegos, N. (2019): "The Importance of Production Networks and Sectoral Heterogeneity for Monetary Policy," *Working paper*.

Eusepi, S., B. Hobijn, and A. Tambalotti (2011): "CONDI: A Cost-of-Nominal-Distortions Index," *American Economic Journal: Macroeconomics*, 3, 53–91.

Galí, J. (2015): *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton ; Oxford: Princeton University Press, second edition ed.

Galí, J. and T. Monacelli (2008): "Optimal monetary and fiscal policy in a currency union," *Journal of International Economics*, 76, 116–132.

Herrendorf, B., R. Rogerson, and Valentinyi (2013): "Two Perspectives on Preferences and Structural Transformation," *American Economic Review*, 103, 2752–2789.

Hooper, P., F. S. Mishkin, and A. Sufi (2019): "Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or is It Just Hibernating?" Working Paper 25792, National Bureau of Economic Research, series: Working Paper Series.

Hoynck, C. (2020): "Production Networks and the Flattening of the Phillips Curve," *Working paper*.

La'O, J. and A. Tahbaz-Salehi (2019): "Optimal Monetary Policy in Production Networks," *SSRN Electronic Journal*.

Mankiw, N. G. and R. Reis (2003): "What Measure of Inflation Should a Central Bank Target?" *Journal of the European Economic Association*, 1, 1058–1086, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1162/154247603770383398.

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

Nakamura, E. and J. Steinsson (2010): "Monetary Non-Neutrality in a Multisector Menu Cost Model," *Quarterly Journal of Economics*, 125, 961–1013.

——— (2013): "Price Rigidity: Microeconomic Evidence and Macroeconomic Implications," Tech. Rep. 18705, National Bureau of Economic Research, Inc, publication Title: NBER Working Papers.

Oberfield, E. and D. Raval (2014): "Micro Data and Macro Technology," NBER Working Paper 20452, National Bureau of Economic Research, Inc.

Orphanides, A. and S. van Norden (2002): "The Unreliability of Output-Gap Estimates in Real Time," *The Review of Economics and Statistics*, 84, 569–583, publisher: The MIT Press.

Pasten, E., R. Schoenle, and M. Weber (2017): "Price Rigidity and the Origins of Aggregate Fluctuations," Working Paper 23750, National Bureau of Economic Research, series: Working Paper Series.

——— (2019): "The propagation of monetary policy shocks in a heterogeneous production economy," *Journal of Monetary Economics*.

Stock, J. H. and M. W. Watson (1999): "Forecasting inflation," *Journal of Monetary Economics*, 44, 293–335.

——— (2007): "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking*, 39, 3–33.

——— (2016): "Core Inflation and Trend Inflation," *The Review of Economics and Statistics*, 98, 770–784, publisher: MIT Press.

Woodford, M. (2003): *Interest and prices: foundations of a theory of monetary policy*, Princeton, N.J. ; Woodstock, Oxfordshire [England]: Princeton University Press, oCLC: ocm52738755.