Networks, Phillips Curves, and Monetary Policy

Elisa Rubbo*

January 22, 2020

Click here for the most recent version

Abstract

I develop an analytical framework to study monetary policy in a disaggregated economy, with multiple sectors and a general input-output network. I derive a generalized version of the Phillips curve and the welfare loss function (the backbone of the New Keynesian framework), expressing their parameters as a function of production primitives. This delivers important new insights. (i) The presence of input-output linkages flattens the Phillips curve. (ii) There is an endogenous tradeoff between stabilizing output and stabilizing consumer prices. (iii) I construct a novel inflation index that eliminates this tradeoff, which I refer to as the "divine coincidence" inflation index. (iv) Monetary policy faces a tradeoff between stabilizing the aggregate output gap and relative output across sectors. While targeting the "divine coincidence" index only stabilizes the output gap, the optimal policy can be implemented by targeting an alternative inflation index that incorporates this tradeoff. I calibrate the model to the U.S. economy, and validate my framework by showing that the divine coincidence inflation rate provides a better fit for Phillips curve regressions than conventional consumer price specifications. The calibrated model also matches empirical estimates of the slope of the consumer price Phillips curve, and predicts a 30% decline over the past 70 years, arising from increased intermediate input flows in production. The baseline strategy of targeting consumer inflation leads to a welfare loss of 1.12% of per-period GDP. Switching to the optimal policy reduces this loss to 0.28%, but cannot fully eliminate it. Targeting the output gap almost replicates the optimal policy.

*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, Gadi Barlevy, Kirill Borusyak, Ashley Craig, Francois Gourio, Greg Mankiw, Ken Rogoff, Raphael Schoenle, Anna Stansbury, 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. It has two key implications. From a positive point of view, inflation and the output gap are linearly related through the Phillips curve. From a normative point of view, welfare depends on inflation and the output gap through a quadratic loss function. The baseline model assumes only one sector of production, whereas in reality an economy has multiple and heterogeneous sectors, which trade in intermediate inputs. How do the Phillips curve and the loss function depend on the disaggregated structure of the economy? This question has crucial implications for monetary policy and inflation targeting.

Inflation targeting is grounded in the "divine coincidence": prices are stabilized if and only if the output gap is closed. Therefore monetary policy can achieve the first-best outcome by stabilizing the price level. The "divine coincidence", however, is an artifact of the one-sector model.[1] Whenever it fails, a tradeoff emerges between stabilizing inflation and closing the output gap. In addition, with multiple sectors there are many possible ways to measure aggregate inflation–that is, to map sector-level inflation rates into an average index. The one-sector model cannot tell us which one is correct. In practice consumer price inflation has been taken as the real-world counterpart of inflation in the one-sector model,[2] but this choice has no theoretical backing.[3] Characterizing the optimal response to the inflation-output tradeoff, and the correct indicators to rely on, requires a theoretical framework that relates the Phillips curve and welfare with the production structure.

This paper provides such a framework. I augment the baseline model with a realistic representation of production: there are multiple sectors, which have arbitrary neoclassical production functions and are arranged in an input-output network. Sectors face idiosyncratic productivity shocks and heterogeneous pricing frictions. In this general setting, I analytically characterize the Phillips curve, welfare and optimal policy. I solve the model as a function of three variables, the output gap, sectoral inflation and productivity, and a set of steady-state parameters. I then study its quantitative implications, by constructing time series of these variables and calibrating the parameters to actual US input-output data. The results show that taking into account the disaggregated structure of the economy is important, both from a theoretical and a quantitative

[1] The "divine coincidence" breaks in simple extensions with two sectors or sticky wages. See Blanchard and Gali (2007), Gali (2008).

[2] Statistical agencies release several different measures of consumer prices: in the US they are the consumer price index (CPI), personal consumption expenditures (PCE), their core versions (excluding sectors with very volatile prices, such as food and energy) and the GDP deflator. Central banks look at all of these measures, and various others (such as wage inflation, commodity prices, import prices, exchange rates...), but they lack a theoretical framework to aggregate them into a proxy for the output gap or into an interest rate target.

[3] For example, for a small open economy producer price inflation is the relevant statistic (see Gali and Monacelli (2005), Gali (2008)).

point of view.

The exposition is organized around two sets of results, positive and normative. The positive analysis is concerned with the response of inflation to the output gap and productivity. I first derive it at the sector level, and then aggregate sector-level responses into the Phillips curve. In this way I can characterize the Phillips curve associated with any given inflation index, corresponding to a specific weighting of sectoral inflation rates. The normative analysis focuses on the central bank's problem. I derive the welfare loss as a function of the output gap and sectoral inflation rates, and solve for optimal monetary policy. While consumer prices are not stabilized under zero output gap or when the optimal policy is implemented (they do not preserve the "divine coincidence"), I construct two distinct inflation measures that restore each of these properties. These indicators respectively provide an inflation proxy for the output gap and an optimal policy target.

The Phillips curve relates aggregate inflation (π) with the output gap ($\tilde{y}$):

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

where ρ is the discount factor, κ is the slope and u_t is a residual. For any given aggregate inflation index, I derive the parameters (κ and u) of the corresponding Phillips curve as a function of productivity shocks and model primitives. This provides a framework to interpret Phillips curve regressions. The model predicts how the estimated slope and residuals should vary across specifications and over time. Specifications where the slope changes with the production structure and where productivity fluctuations generate large cost-push shocks will have larger residuals and be more sensitive to the sample period. An example of this is the consumer-price Phillips curve. I then construct a "divine coincidence" inflation index that is always stabilized under zero output gap. The corresponding Phillips curve has no endogenous residual, and its slope does not depend on the production structure.

To build to these results, I first provide formulas relating sectoral inflation rates to the output gap and productivity. The slope of the Phillips curve is obtained by aggregating the elasticities of sector-level prices to the output gap, according to the weighting prescribed by the chosen inflation index. I characterize sectoral elasticities as a function of the input-output structure and sector-level frequencies of price adjustment. I demonstrate that these elasticities, and therefore the slope of the aggregate Phillips curve, are decreasing in the size of intermediate input flows. The output gap acts as a labor demand shifter, thereby affecting real wages. In turn, wages are passed-through into marginal costs and prices. Sectoral responses are given by the product of the elasticity of real wages with respect to the output gap (which does not depend on the network structure) times its pass-through into prices, scaled by a general equilibrium multiplier. Sectors are affected by wage changes directly (if they hire workers) and through intermediate input prices. Because of price rigidities, suppliers do not fully pass-through wage changes into their prices. As a result the indirect pass-through is smaller than the direct one, and the overall pass-through into sectoral and aggregate inflation decreases with intermediate input flows. Correspondingly, sectoral elasticities with respect to the output gap decline. The slope of the Phillips curve, which aggregates these elasticities, also declines.

From a quantitative point of view, accounting for intermediate input flows leads to very different predictions with respect to the benchmark model. In the calibration I focus on the consumer-price Phillips curve, given its prominence as a standard measure. 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. The multi-sector model also allows to evaluate the role of the input-output structure for the evolution of the Phillips curve over time. Empirical estimates suggest that the slope has declined. I fit the model to historical input-output tables for each year between 1947 and 2017.[4] The predicted slope is about 30% smaller in 2017 than in 1947.[5] Mirroring the slope of the Phillips curve, the network model features significantly more monetary non-neutrality than the baseline. The predicted response of inflation to monetary shocks also decreased over time. Intuitively, a smaller slope of the Phillips curve implies that consumer inflation is less sensitive to the output gap. In turn, the output gap is set by monetary policy through interest rates. This is why intermediate input flows reduce the inflation response to monetary (interest rate) shocks.

In the multi-sector model the "divine coincidence" does not hold at the sector level: prices cannot be stabilized in every sector, even if the output gap is closed. With input-output linkages and/or idiosyncratic shocks, sectors are differentially exposed to productivity changes. Those that are directly hit face a larger change in marginal costs. In addition, under zero output gap real wages adjust to reflect the aggregate productivity change. Wages also have an asymmetric pass-through into sectoral marginal costs. I derive the elasticities of sectoral prices with respect to productivity, decomposing it into the direct component and the indirect effect through wages. The residual of the Phillips curve, u_t, aggregates these elasticities. From equation (1), the "divine coincidence" holds in the aggregate if and only if the Phillips curve has no endogenous residual ($u_t \equiv 0$). In this

[4]The BEA provides data at different levels of aggregation for different years. For a clean comparison I convert all the data to the most aggregated level, with 47 industries.

[5]Other authors attribute the decline in the slope of the Phillips curve to a different channel (see Blanchard (2016)): 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 Steinnsson (2013) find that the frequency of price adjustment is stable over time.

case, closing the output gap is equivalent to stabilizing aggregate inflation. Consumer prices do not preserve this property in the multi-sector framework. I use BEA-KLEMS data to construct a time series of the residual in the consumer-price Phillips curve from measured sectoral TFP shocks. The series has a standard deviation of 25 basis points, suggesting that endogenous cost-push shocks are quantitatively significant.

I construct a (unique) inflation index that restores the "divine coincidence" in the aggregate. This "divine coincidence" index weights sectoral inflation rates according to sales shares, appropriately discounting more flexible sectors. Intuitively, the output gap captures distortions in aggregate demand arising from markups. In my framework sector-level markups change endogenously, due to incomplete price adjustment. Changes in markups, in turn, are reflected in inflation rates. After a cost shock, the producers who are able to adjust their prices keep a constant markup. Those who cannot adjust instead need to absorb the shock into their markup. As a result, the same shock generates a larger inflation response and a smaller markups response in sectors with more flexible prices. Therefore these sectors need to be discounted. We then want to relate sectoral markups with the aggregate output gap. I show that the correct way to do this is by weighting sectors according to sales shares. Sales shares reflect the full role of each sector in the value added chain, in contrast with consumption shares, which only account for the final stage. I also prove that the slope of the "divine coincidence" Phillips curve only depends on labor supply elasticities, and not on the input-output structure. This is only true for my indicator: for all the others, the slope depends on the production structure.

Interestingly, in the calibration 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 (2006), 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 wage inflation relative to inflation in other sectors.

I construct a time series of the "divine coincidence" indicator for the US economy from PPI data, input-output tables from the BEA and estimates of price adjustment frequencies from Pasten, Schoenle and Weber (2016). I run Phillips curve regressions over the years 1984-2017, and compare the results with standard measures of consumer inflation. The baseline OLS specification has an R-squared of about 0.2 with the "divine coincidence" index, as opposed to about 0.05 with consumer price inflation. Rolling regressions over 20 year windows for the same sample period are always significant, as opposed to about 50% of the time with consumer prices, and the estimated coefficient is stable over time.

Targeting the "divine coincidence" indicator closes the output gap, but does not replicate the efficient equilibrium that prevails under flexible prices. Indeed, in the multi-sector framework it is impossible to achieve the first-best. I derive an expression for the welfare loss as a function of the output gap and sectoral inflation rates. The output gap captures distortions in aggregate demand, while inflation reflects distortions in relative demand across firms and sectors. Within each sector, shocks induce price distortions between the firms that adjust their prices and those that do not. Incomplete price adjustment further results in relative price distortions across sectors. The size of within- and cross-sector price distortions depends on how shocks propagate through the input-output network. Their welfare effect depends on the response of quantities demanded, which is governed by the relevant elasticities of substitution in production and consumption.

Closing the output gap is not constrained optimal: while stabilizing aggregate demand, it does not minimize misallocation. 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. The optimal policy can still be implemented via inflation targeting, that is, by stabilizing an appropriate inflation indicator which incorporates the two sides of the tradeoff. The optimal target weights the (inflation proxy for the) output gap against sectoral inflation rates according to the relative marginal benefit and marginal cost of distorting aggregate demand to alleviate the misallocation associated with inflation.

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. This loss decreases to 0.28% under the optimal policy, but it cannot be fully eliminated due to imperfect stabilization. Closing the output gap yields a small additional loss with respect to the optimal policy, therefore the output gap is a good target from a quantitative point of view. Intuitively, monetary policy is a blunt instrument to correct misallocation, because it can only change relative prices proportionately to their elasticity with respect to the output gap. The cost of distorting aggregate demand is larger than the gain, and in practice the optimal output gap is close to zero.

Related literature From a conceptual standpoint, my work is related with the literature on markup distortions and welfare. I consider an economy with sticky prices, where producers who cannot adjust their price absorb cost shocks into their markup. By generating markup distortions, incomplete price adjustment results in a welfare loss. My framework is close to Baqaee and Farhi (2017), who study the effect of markup changes on aggregate productivity. The key difference is that in their setup markup shocks are exogenous, whereas I solve for the endogenous markup changes induced by productivity shocks and price rigidities.

Two main strands of the literature have extended the one-sector New Keynesian framework to incorporate realistic elements of the production structure. Earlier papers focus on simple extensions of the baseline model, maintaining its analytical approach. More recent works consider richer input-output structures, shifting the focus away from an analytical characterization of the Phillips curve and welfare, and towards quantitative implications. I combine the benefits of the two approaches: I provide an analytical characterization of the Phillips curve, welfare and optimal policy in a fully general production structure, and explore its quantitative implications.

Early extensions of the baseline model account for wage rigidities (Blanchard and Gali (2007), Gali (2008)), two sectors (Aoki (2001), Woodford (2003)), or a horizontal economy (Benigno (2004), Gali and Monacelli (2008)). These papers demonstrate the failure of the "divine coincidence". These models are easy to solve analytically, but they are too stylized to capture many features of a realistic production network.

These are explored in more recent works, which focus on quantitative models with richer production structures. On the positive side, a large literature shows that input-output linkages and heterogeneous pricing frictions increase monetary non-neutrality (Basu (1995), Carvalho (2006), Carvalho and Nechio (2011), Nakamura and Steinsson (2013) and Pasten, Schoenle and Weber (2016)). Pasten, Schoenle, Weber (2017) simulate the response of real GDP to sector level productivity shocks and find that they produce sizable fluctuations. On the normative side, Castro-Cienfuegos (2019) calibrates a multi-sector model to the US input-output structure, to study the cost of trend inflation. He shows that this cost is higher with input-output linkages, and isolates the components of the welfare loss arising from within- and cross-sector price dispersion. He also finds a flattening of the Phillips curve in the calibration that accounts for input-output linkages. Several of my results are common with this literature. My analytical approach allows to demonstrate the flattening of the Phillips curve and the presence of an inflation-output tradeoff in full generality, beyond a specific parameterization. I can study how the parameters of the Phillips curve and welfare function depend on the network structure, and build on these results to derive optimal indicators and characterize the optimal monetary policy response to sector-level fluctuations.

The normative section of my analysis is also related to parallel and independent work by La'O and Tahbaz-Salehi (2019), who study optimal monetary policy with production networks. 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, sectoral weights have different determinants in the optimal targeting rule (the information structure versus substitution elasticities).

Previous works show that consumer inflation is not the correct indicator for the Phillips curve and monetary policy. From a theoretical perspective, open economy models imply the correct indicator is the PPI. Gali and Monacelli (2005) derive this result for a small open economy. Benigno (2004) finds that in a currency area countries should be weighted according to their size, discounting those where prices are more flexible. This closely resembles the "divine coincidence" indicator, in the special case of a two-country, two-sector model, with no intermediate inputs. In a closed economy, Mankiw and Reis (2003) derive the optimal inflation target in a multi-sector environment, under the assumption that the central bank seeks to maximize output stability. They do not consider input-output linkages, but introduce several elements of sectoral heterogeneity, such as size, volatility and price flexibility. They also find that a high weight should be placed on wage inflation, which is indeed what the "divine coincidence" indicator does. Also for a closed economy, Eusepi, Hobjin and Tambalotti (2011) calibrate an inflation index that captures the welfare cost of nominal distortions a in a multi-sector horizontal economy. I move beyond these contributions in three dimensions: first, I consider a more general input-output structure; second, my optimal inflation target is based on a "microfounded" objective function, derived by maximizing the welfare of the representative consumer; third, I provide an analytical characterization of the optimal sectoral weights as a function of the production structure.

The limitations of consumer prices as an indicator are also well documented empirically, and many studies seek to construct indicators that perform better. Stock and Watson (1999) and Bernanke and Boivin (2003) propose better predictors of future inflation trends, and Stock and Watson (2015) show that using sectoral inflation rates improves the prediction of trend inflation. Orphanides and Van Norden (2002) show that even ex-post measures of the output gap do not predict consumer inflation out of sample. A large empirical literature documents the instability of the Phillips curve: while Gali and Gertler (1999) use real marginal costs as a proxy for the output gap and estimate significant coefficients (and with the correct sign), follow-up studies found that their results do not hold for different time periods or data vintages. Mavroeidis, Plagborg-Muller and Stock (2014) summarize this literature, and document that estimates of the Phillips curve coefficients are very sensitive to the way inflation expectations are accounted for, to the specific measure of output gap used and to the sample period. I contribute to this literature by showing that regressions with the "divine coincidence" specification yield more stable and significant estimates over time and across specifications.

The rest of the paper is organized as follows: Section 2 introduces the setup; Section 3 contains some key definitions; Section 4 presents the positive analysis; Section 5 derives the normative results; Section 6 calibrates the model to the US economy; Section 7 discusses results from Phillips curve regressions. Section 8 concludes. Two extensions of the basic model, that introduce exogenous cost-push shocks and dynamics, are presented in Appendix C. An online Appendix collects more detailed proofs and empirical results.

2 Setup

This section introduces the key elements of the network model and lays out the assumptions about preferences, timing and policy instruments. Following the approach in Baqaee and Farhi (2017), I define competitive equilibrium for a given change in sectoral markups. In my setup markups change endogenously because some firms cannot adjust their price, and therefore absorb changes in marginal costs into their markup. Those that can adjust instead keep their markups constant. I assume that fiscal policy eliminates markup distortions in steady-state, and log-linearize the model around the efficient equilibrium. I introduce the variables and steady-state parameters that characterize the log-linearized system.

2.1 Timing and policy instruments

In the main text I consider a one-period model. The economy starts in steady state. Then sectors are hit by unanticipated productivity shocks, and monetary policy is implemented; some firms can adjust their price, while others cannot. The world ends after production and consumption take place. The dynamic version of the model is presented in Appendix D. In the static setup money supply is the only policy instrument (to be replaced with interest rates in the dynamic version), and works through a cash-in-advance constraint.

2.2 Preferences

Consumers derive utility from consumption and leisure, and have homothetic preferences over a bundle of all goods produced in the economy. Their utility function is

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

where C is a constant returns to scale aggregator of all goods produced in the economy and L is labor supply.

Consumers maximize utility subject to the cash in advance constraint

$$PC \leq M$$

and the budget constraint

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

where P is the price of the consumption bundle, M is total money supply, w is the wage, Π are firm profits (rebated lump-sum to households) and T is a lump-sum transfer (that the government uses to finance input subsidies to firms).

2.3 Production

There are N sectors in the economy (indexed by $i \in \{1, ..., N\}$), and within each sector there is a continuum of firms, producing differentiated varieties. Customers buy a CES bundle of these varieties.

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

$$Y_i = A_i F_i(L_i, \{x_{ij}\})$$

where L_i is labor, x_{ij} is the quantity of good j used as input, and A_i is a Hicks-neutral, sector-specific productivity shock.[6] I assume that labor is freely mobile across sectors.

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

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

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

In the initial steady-state all firms set their price optimally, solving

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

where D_i and P_i are the sector-level demand and price index, and ϵ_i is the elasticity of substitution between varieties within sector i. The government provides input subsidies, in order to eliminate markup distortions (i.e. the markup over pre-subsidy marginal costs is 1). Input subsidies cannot change in response to productivity shocks, and are constrained to be the same for all firms within the same sector. Under the optimal subsidy

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

[6] 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.

firms charge price

$$p_i^* = mc_i \tag{2}$$

in steady-state.

With price rigidities, firms within the same sector charge different prices outside of steady-state. Firms who can adjust their price do not change their markup. Firms that cannot adjust must accept a change in their markup, given by

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

2.4 Equilibrium

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

The log-linearized model is 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 steady-state parameters, which capture the input-output structure and sector-level pricing frictions. They are defined below.

3.1 Steady-state parameters

3.1.1 Price rigidity parameters

I model price rigidities by assuming that only a fraction δ_i of the firms in each sector i can adjust their price after observing the shocks. I denote by Δ the diagonal matrix whose i-th element is δ_i.

The Calvo assumption together with firms' optimal pricing equation (2) yield the following mappings between inflation, marginal costs and markups:

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

Remark 1. In this setup wage rigidities can be modeled by adding a labor sector, which sells labor services to all the other sectors and has sticky prices.

3.1.2 Input-output definitions

The input-output structure is characterized by steady-state consumption, labor and input-output shares. I introduce two additional objects, the Leontief inverse and the vector of sales shares, that are derived from the input-output matrix and consumption shares.

Consumption shares The $N \times 1$ vector β denotes 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 element i, j given by the expenditure share on input j in i's 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 $(I - \Omega)^{-1}$.

It holds that, while ω_{ij} is the fraction of sector i revenues spent on goods from sector j "directly", 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 i's purchases of j's product through intermediate inputs from its suppliers, its suppliers' suppliers, etc. Formally, the Leontief inverse is given by the geometric sum:

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

where the k-th power captures the expenditure of i on j through paths of length k.

In my setup labor is the only factor of production. Therefore labor and intermediate input shares must sum to one in every sector:

$$\alpha_i + \sum_j \omega_{ij} = 1$$

Or equivalently, denoting by $\mathbf{1}$ the column vector with all entries equal to 1:

$$\alpha + \Omega \mathbf{1} = \mathbf{1}$$

This yields the following relation between labor shares and the Leontief inverse:

$$(I - \Omega)^{-1} \alpha = \mathbf{1} \tag{4}$$

The right hand side in equation (4) captures "total" (direct and indirect) sector-level labor shares. Intuitively, producers hire labor directly, or buy it indirectly through intermediate inputs. Given that labor is the only factor of production, equation (4) tells us that the total labor share must be equal to one in every sector.

"Adjusted" Leontief inverse I refer to the matrix $(I - \Omega\Delta)^{-1}$ as the "adjusted" Leontief inverse.

We established that the Leontief inverse captures the direct and indirect expenditure of sector i on goods from sector j, as a share of i's revenue. With flexible prices this is also the elasticity of i's marginal cost to j's marginal cost, because firms charge constant markups. With price rigidities this is no longer true, as marginal cost changes do not get fully passed-through into input prices. In this case the "direct" elasticity of i's marginal cost with respect to j's is $\omega_{ij}\delta_j$, which discounts ω_{ij} by the fraction δ_j of producers in j that adjust their price. Correspondingly, the total (direct and indirect) elasticity of i's marginal cost with respect to j's is given by the (i, j) element of the "adjusted" Leontief inverse.

Sales shares The vector λ of sectoral sales shares in total GDP has components

$$\lambda_i = \frac{P_i Y_i}{PC}$$

Section 4.3 argues that sales shares, as opposed to consumption shares, capture the full role of each sector in the value added chain. Therefore the inflation index in our "divine coincidence" specification of the Phillips is based on sales shares. It holds that sales shares are related with consumption shares and the Leontief inverse as follows:

$$\lambda^T = \beta^T (I - \Omega)^{-1}$$

Elasticities of substitution While the log-linearized model only depends on the input-output concepts defined above, elasticities of substitution matter for the welfare loss derived in Section 5. 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_{ij}^k, and their elasticity of substitution in consumption by σ_{ij}^C; the elasticity of substitution between good i and labor in the production of good k is denoted by θ_{iL}^k.

3.2 Variables

3.2.1 Aggregate output gap

The output gap captures distortions in aggregate demand, which I express in terms of final output. I define it in the standard way.

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

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

Lemma 1 derives the change in efficient output induced by a given productivity shock, and we can solve for the change in realized output as a function of productivity shocks and labor supply using Lemma 2. A consequence of this Lemma is that the output gap can be equivalently expressed in terms of labor demand, as in equation (6).

Lemma 1. *The change in efficient output after a productivity shock $d \log A$ is given by*

$$y^{nat} = \frac{1+\varphi}{\gamma+\varphi} \lambda^T d \log A \tag{5}$$

Proof. See Appendix A1. □

The intuition for this result is simple. From Hulten's theorem, under flexible prices the first-order change in aggregate productivity is an average of sector-level productivity shocks, with weights given by sales shares λ:

$$d \log A_{AGG} = \lambda^T d \log A$$

From the consumption-leisure trade-off, in the flex-prices equilibrium labor supply changes by

$$d \log L^{nat} = \frac{1-\gamma}{\gamma+\varphi} \lambda^T d \log A$$

Equation (5) follows immediately from these two relations, together with the definition of aggregate productivity:

$$y^{nat} = d \log L^{nat} + d \log A_{AGG}$$

Lemma 2. *Around the undistorted steady-state, to a first order the change in aggregate productivity in the economy with price rigidities is the same as in the flex-price equilibrium.*[7]

Proof. See appendix A1. □

Lemma 2 allows to write the output gap as a function of productivity shocks and labor demand. We have:

$$\widetilde{y} = d \log L - d \log L^{nat} = d \log L - \frac{1-\gamma}{\gamma+\varphi} \lambda^T d \log A \tag{6}$$

3.2.2 Sectoral inflation rates

The $N \times 1$ vector of inflation rates is denoted by

$$\pi = \begin{pmatrix} \pi_1 \\ ... \\ \pi_N \end{pmatrix}$$

Remark 2. While the output gap captures distortions in aggregate demand, Section 5.1 shows that the welfare cost of "relative sectoral output gaps" can be written as a function of sectoral inflation rates (see Proposition 4).

4 The Phillips curve

The Phillips curve is a linear relation between aggregate inflation π^{AGG} and the output gap $\tilde{y}$. In a one-period model it is given by

$$\pi^{AGG} = \kappa \tilde{y} + u \tag{7}$$

[7] There is a second order productivity loss due to incomplete price adjustment. See Section 4.2

where κ is the slope and u is the residual. A key result in the one-sector model is that the Phillips curve has no endogenous residual ($u \equiv 0$). This is another way to state that there is no endogenous tradeoff between stabilizing output and inflation (the "divine coincidence"). With multiple sectors instead the slope and residual of the Phillips curve depend on the choice of an aggregate inflation index π^{AGG}. Different indicators are characterized by their weighting of sectoral inflation rates. The corresponding Phillips curves can be steeper or flatter, and the residual has different variance.

In this section I address two questions. First, for a given aggregate inflation index I derive the parameters of the corresponding Phillips curve (slope and residual). Second, I look for an index that yields a "divine coincidence" Phillips curve, with no endogenous residual. I show that this index is unique, and the slope of the corresponding Phillips curve does not change with the production structure. By contrast, for all other aggregate inflation indicators the slope depends on input-output shares and pricing frictions. These two properties imply that the Phillips curve associated with the "divine coincidence" index is the only well-specified one.

I proceed as follows. In Section 4.1 I derive the elasticity of prices sector-by-sector with respect to the output gap and sectoral productivity shocks, as a function of network primitives. For a given weighting of sectoral inflation rates, I then aggregate these elasticities into the slope and residual of the corresponding Phillips curve. I especially focus on consumer price inflation, given its prominence as the most common measure. I show that the slope of the consumer-price Phillips curve decreases with intermediate input flows, and productivity shocks generate an endogenous residual. In this sense the Phillips curve as usually measured is misspecified. The examples in Section 4.2 illustrate why this is the case. Section 4.3 then derives the "divine coincidence" inflation index, and shows that it is unique.

4.1 General expression for given inflation index

4.1.1 Notation and aggregation

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.

Formally, we can express inflation as a function of the output gap and productivity as follows:

$$\underbrace{\pi}_{N\times 1} = \underbrace{\mathcal{B}}_{N\times 1} \widetilde{y} + \underbrace{\mathcal{V}}_{N\times N} \underbrace{d\log A}_{N\times 1} \tag{8}$$

An aggregate inflation index π^{AGG} is characterized by the vector of weights ϕ that it assigns to sectoral inflation rates:

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

For a given indicator, the corresponding Phillips curve is obtained by aggregating both sides of Equation (8):

$$\pi^{AGG} = \underbrace{\phi^T \mathcal{B}}_{\text{slope}} \tilde{y} + \underbrace{\phi^T \mathcal{V} d \log A}_{\text{residual}} \tag{9}$$

Correspondingly, the slope is the aggregate elasticity with respect to the output gap, while the residual is the aggregate elasticity with respect to productivity (as in Equation (9)).

Consumer inflation π^C is a special case, obtained by weighting sectoral inflation rates according to consumption shares ($\phi = \beta$):

$$\pi^C = \beta^T \pi$$

The corresponding Phillips curve is

$$\pi^C = \beta^T \mathcal{B} \tilde{y} + \beta^T \mathcal{V} d \log A \tag{10}$$

so that the slope and residual are given by

$$\kappa^C = \beta^T \mathcal{B}$$
$$u^C = \beta^T \mathcal{V} d \log A$$

Sections 4.1.2 and 4.1.3 below characterize the elasticities $\mathcal{B}$ and $\mathcal{V}$, and derive the slope and residual of the consumer-price Phillips curve as a corollary.[8] Since we are log-linearizing the model around the initial steady-state, all changes in prices, wages and productivity are relative to this steady-state.

4.1.2 Slope of the Phillips curve

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

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

$$\mathcal{B} = \frac{\Delta \left((I - \Omega \Delta)^{-1} \alpha \right)}{1 - \bar{\delta}_w} (\gamma + \varphi) \tag{11}$$

[8]Propositions 1 and 2 can be seen as an application of Proposition 10 in Baqaee and Farhi (2017), recast in terms of sectoral probabilities of price adjustment and in the special case of an efficient initial equilibrium.

where

$$\bar{\delta}_w \equiv \beta^T \Delta \left(I - \Omega\Delta\right)^{-1} \alpha \tag{12}$$

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

Proof. See Appendix A2 □

Corollary 1. *The slope of the consumer-price Phillips curve is given by*

$$\kappa^C = \frac{\bar{\delta}_w}{1 - \bar{\delta}_w} \left(\gamma + \varphi\right) \tag{13}$$

Proof. The result follows immediately from Proposition 1 and Equation 9. □

The vector $\mathcal{B}$ and the slope κ^C are the elasticities of sectoral and consumer prices with respect to the output gap. They have two components: the first depends on the parameters of the labor supply curve (γ and φ), and it captures the elasticity of real wages with respect to the output gap; the second depends on the network structure, and corresponds to the pass-through of real wages into prices.

Equation (6) shows that the output gap acts as a labor demand shifter. Intuitively, if output is above potential then labor demand must increase. The effect on real wages is governed by the parameters of the labor supply curve. I show that the elasticity of real wages to the output gap is the sum of the inverse Frish elasticity φ and the wealth effect in labor supply γ (see the proof of Proposition 1), which gives us the first component on the right hand side of (11) and (13). Wages in turn affect marginal costs and prices. Correspondingly, the wage pass-through into prices gives us the second component in (11) and (13). In turn, the real wage pass-through can be decomposed into the nominal wage pass-through, given by $\Delta\left(\left(I - \Omega\Delta\right)^{-1} \alpha\right)$ for sectoral inflation rates and by $\bar{\delta}_w$ for consumer prices, and a general equilibrium multiplier $(1 - \bar{\delta}_w)$. The general equilibrium multiplier maps changes in real wages into changes in nominal wages. The two differ precisely because of the response of consumer prices, given by $\bar{\delta}_w$.

Thus, the nominal wage pass-through is the key object. In the one-sector benchmark we have $\Omega = \mathbb{O}$ and $\alpha = 1$, so that marginal costs have unit elasticity with respect to wages. Therefore the pass-through is simply given by the price rigidity parameter δ:

$$\bar{\delta}_w = \Delta\left(\left(I - \Omega\Delta\right)^{-1} \alpha\right) = \delta$$

With input-output linkages instead this pass-through gets dampened, as stated in Corollary 2.

Corollary 2. *As long as some sector uses an intermediate input with sticky prices, the pass-through of wages into marginal costs is less than one:*

$$\exists i, j \text{ such that } \omega_{ij}\delta_j < \omega_{ij} \Rightarrow (I - \Omega\Delta)^{-1}\alpha < \mathbf{1} \tag{14}$$

Correspondingly, the aggregate pass-through $\bar{\delta}_w$ is lower than the average price rigidity $\mathbb{E}_\beta(\delta)$:

$$\exists i, j \text{ such that } \omega_{ij}\delta_j < \omega_{ij} \Rightarrow \bar{\delta}_w < \mathbb{E}_\beta(\delta) \tag{15}$$

A reduction in labor shares compensated by a uniform increase in input shares also reduces $\bar{\delta}_w$:

$$d\alpha_i < 0,\ d\omega_{ij} = d\omega_{ik}\ \forall j, k,\ \exists j \text{ such that } \omega_{ij}\delta_j < \omega_{ij} \Rightarrow d\bar{\delta}_w < 0 \tag{16}$$

Proof. See Appendix A2. □

The intuition is as follows. Marginal costs and prices are affected by wages directly, or indirectly through input prices. The direct exposure depends on own labor shares, while the indirect one is given by the suppliers' labor share, the suppliers' suppliers labor share, etc. The indirect pass-through is dampened, because of incomplete price adjustment: some suppliers cannot change their price in response to wages. Formally, suppliers need to be discounted by the fraction of firms that update their prices in each sector. The overall pass-through into marginal costs is given by the adjusted Leontief inverse, multiplied by the vector α of steady-state labor shares:

$$\frac{d\log mc}{d\log w} = (I - \Omega\Delta)^{-1}\alpha \tag{17}$$

We can then use the pricing equation (3) to translate changes in marginal costs into inflation:

$$\frac{d\log p}{d\log w} = \Delta\frac{d\log mc}{d\log w} = \Delta\,(I - \Omega\Delta)^{-1}\alpha \tag{18}$$

This gives the pass-through of nominal wages into sectoral inflation rates in equation (11). Corollary 2 states that with input-output linkages at least one component of the vector $(I - \Omega\Delta)^{-1}\alpha$ is strictly smaller than one, reflecting the incomplete wage pass-through from price rigidities. Therefore the corresponding price pass-throughs are smaller than the adjustment probabilities. Also note that different sectors have different pass-through. The pass-through is higher in sectors with a large direct labor share and flexible prices, and whose suppliers have a large direct labor share and flexible prices, etc. To obtain the overall pass-through into consumer prices $\bar{\delta}_w$ we simply aggregate the sectoral responses in (18) according to consumption shares:

$$\bar{\delta}_w = \beta^T \frac{d\log p}{d\log w} = \beta^T \Delta \left(I - \Omega\Delta\right)^{-1} \alpha \tag{19}$$

The inequality $(I - \Omega\Delta)^{-1} \alpha < 1$ from Corollary 2 implies that the pass-through $\bar{\delta}_w$ is smaller than the average price rigidity:

$$\bar{\delta}_w < \mathbb{E}_\beta\left(\delta\right)$$

As a result, intermediate input flows reduce the slope of the consumer-price Phillips curve:

$$\kappa^C = (\gamma + \varphi) \frac{\bar{\delta}_w}{1 - \bar{\delta}_w} < (\gamma + \varphi) \frac{\mathbb{E}_\beta\left(\delta\right)}{1 - \mathbb{E}_\beta\left(\delta\right)} \tag{20}$$

The right hand side of Equation (20) is the slope predicted by standard calibrations, which directly map the one-sector model into the data. The Calvo parameter in the model is calibrated by simply averaging sectoral price rigidities according to consumption shares. Quantitatively, the difference between the left and right hand sides of Equation (20) is important. Section 6.3 calibrates the model to the US economy, and finds 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. Corollary 3 aggregates them into the endogenous residual of the consumer-price Phillips curve. In contrast with the one-sector model, in the multi-sector model closing the output gap does not stabilize either sector-level prices or consumer prices. In other words, the "divine coincidence" fails. Corollary 4 shows that the only shock for which the "divine coincidence" still holds is an aggregate labor augmenting shock.

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

$$\mathcal{V} = \Delta \left(I - \Omega\Delta\right)^{-1} \left[\frac{1 - \bar{\delta}_A}{1 - \bar{\delta}_w} \alpha \lambda^T - I\right] \tag{21}$$

where

$$\bar{\delta}_A \left(d\log A\right) \equiv \frac{\beta^T \Delta \left(I - \Omega\Delta\right)^{-1} d\log A}{\lambda^T d\log A} \tag{22}$$

is the pass-through of the productivity shock into consumer prices, relative to the aggregate shock.

Proof. See Appendix A2 □

Corollary 3. *The residual in the consumer-price Phillips curve is given by*

$$u^C = \frac{\bar{\delta}_w - \bar{\delta}_A}{1 - \bar{\delta}_w} \lambda^T d \log A \tag{23}$$

Proof. See Appendix A2 □

Proposition 2 goes back to the core of the "divine coincidence". The key intuition is that, under zero output gap, wages and productivity have competing effects on marginal costs and inflation. Real wages must adjust to reflect the change in the marginal product of labor: from Equation (6), the output gap is closed if and only if labor supply is at its efficient level. This can happen only if the change in real wages is equal to the change in the marginal product of labor, which in this setup is given by the change in aggregate productivity.

If aggregate productivity falls ($\lambda^T d \log A < 0$), then wages also fall, and the two have opposite effects on marginal costs. In the one sector model marginal costs are symmetrically exposed to wages and productivity (they have unit elasticity with respect to both), so that these two forces exactly counteract each other. This is the key mechanism behind the "divine coincidence". With multiple sectors, instead, marginal costs are not equally exposed to wages and productivity. Therefore they are not necessarily stabilized under zero output gap, and they can move in opposite directions in different sectors.

The change in real wages only depends on aggregate productivity $\lambda^T d \log A$. The general equilibrium multiplier $\frac{1-\bar{\delta}_A}{1-\bar{\delta}_w}$ maps it into the change in nominal wages. We derived the pass-through of nominal wages into sectoral prices in Section 4.1.2. The overall inflation response is given by the first term in Equation (21):

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

The direct effect of productivity instead depends on the specific distribution of sectoral shocks, not just on the aggregate. Productivity affects marginal costs directly, and indirectly through input prices. The propagation is asymmetric across sectors: a productivity shock to sector i has a larger impact on i's customers, its customers customers, etc., and the propagation along the chain is larger if prices are more flexible. As for wages, the overall pass-through of productivity into marginal costs is given by the adjusted Leontief inverse:

$$\frac{\partial \log mc}{\partial \log A} = -\left(I - \Omega\Delta\right)^{-1} \tag{24}$$

The pricing Equation (18) shows that changes in marginal cost are mapped into changes in prices by multiplying by the adjustment probabilities Δ. Together, this result and Equation (24) yield the second term of (21). Whenever the two components of (21) are different it is impossible to stabilize inflation at the sector level, even if the output gap is closed.

Corollary 3 shows that consumer inflation is not stabilized either. Overall, its response depends on the relative pass-through of wages and productivity into consumer prices, given by the difference $\bar{\delta}_w - \bar{\delta}_A$. We introduced the wage pass-through $\bar{\delta}_w$ in Section 4.1.2. The productivity pass-through $\bar{\delta}_A$ is defined in Equation (22). Note that, different from $\bar{\delta}_w$, $\bar{\delta}_A$ is not a constant, but it depends on the specific realization of sectoral productivity shocks. The pass-through $\bar{\delta}_A$ simply aggregates sectoral elasticities according to consumption shares, and scales the overall response by the aggregate shock.

From Equation (23) we see that, for a negative shock, consumer inflation is positive if and only if $\bar{\delta}_A > \bar{\delta}_w$. That is, if and only if the pass-through of the productivity shock into consumer prices is larger than the pass-through of nominal wages. The examples in Section 4.2 illustrate that this is the case whenever downstream or flexible sectors are hit by a "worse" shock than the average.

A natural question at this point is whether there are shocks after which prices are stabilized in every sector under zero output gap. Corollary 4 below states 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.

Corollary 4. *It holds that $\mathcal{V}\alpha = \mathbf{0}$, and α is the only vector with this property.*

Proof. See Appendix A2 □

A consequence of Corollary 4 is that perfect stabilization is impossible not only in the presence of asymmetric sector-level shocks, but also after an aggregate TFP shock - except in a horizontal economy,[9] where aggregate TFP shocks and labor augmenting shocks are the same. Quantitatively, aggregate TFP shocks generate a significant inflation-output tradeoff. In the calibrated model a 1% negative shock increases consumer inflation by 0.26% under zero output gap.

4.2 Examples

This section illustrates with three examples the main mechanisms through which the "divine coincidence" fails. The vertical chain isolates the effect of input-output linkages, while heterogeneous adjustment frequencies are the only force at play in the horizontal economy. The oil economy combines the two.

[9]That is, an economy without intermediate inputs.

In the vertical chain consumer price inflation is dominated by the downstream sector, whose relative consumption share is larger than the relative sales share. In the horizontal economy instead consumer inflation overrepresents more flexible sectors. The two channels interact in the oil economy. This last example rationalizes the common wisdom that oil shocks create an endogenous "cost-push" term in the Phillips curve, highlighting the crucial role of wage rigidities and heterogeneous adjustment probabilities in generating this outcome.

Example 1. Vertical chain

Consider an economy made of two sectors, which we label U (for "upstream") and D (for "downstream"). Both sectors use labor, and D also uses U as an intermediate input. Only D sells to final consumers. As the downstream good is the only consumption good, the residual in the consumer-price Phillips curve is given by inflation in this sector under zero output gap.

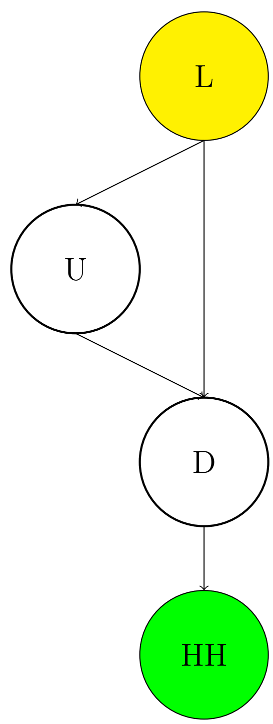

Figure 1: Vertical chain

Let's start by considering a negative productivity shock that hits D. The elasticities $\mathcal{V}_{UD}$ and $\mathcal{V}_{DD}$ of U and D's prices with respect to a productivity shock in D can be derived from Proposition 2, and are given by:

$$\mathcal{V}_{UD} = \overbrace{\underbrace{\delta_U}_{\text{pass-through}} \underbrace{\frac{1-\delta_D}{1-\bar{\delta}_w}}_{\text{multiplier}}}^{\text{nominal wages}} > 0$$

$$\mathcal{V}_{DD} = \left[\overbrace{\underbrace{\bar{\delta}_w}_{\text{pass-through}} \frac{1-\delta_D}{\underbrace{1-\bar{\delta}_w}_{\text{multiplier}}}}^{\text{multiplier}} - \overbrace{\delta_D}^{\text{productivity}} \right] < 0$$

where

$$\bar{\delta}_w = \delta_D \left(\overbrace{\alpha_D}^{\text{direct pass-through}} + \overbrace{(1-\alpha_D)\,\delta_U}^{\text{through } U} \right)$$

The drop in productivity leads to higher marginal costs for D. However the fall in real wages exactly compensates the change in D's productivity. In the one-sector case the two effects cancel out, and marginal costs remain unchanged. This is not true in our example: as long as there is some price stickiness in U, the change in wages is not fully passed-through into D's marginal cost:

$$\alpha_D + (1-\alpha_D)\,\delta_U < 1$$

Thus producers in D want to increase their price. In this example consumer prices coincide with the price of the downstream sector, so that consumer inflation is positive. Note that inflation has the opposite behavior in U, as nominal wages drop and productivity is unchanged.

The asymmetric exposure of sectoral marginal costs is not merely a result of the asymmetric nature of the shock (it hits only one sector). Indeed, the "divine coincidence" breaks also under an aggregate Hicks-neutral shock. The intuition is as follows: the upstream sector U is more exposed to the change in wages, while the downstream sector is more exposed to the productivity shock. On one side U benefits from the full wage decline, which it doesn't completely pass-through. On the other side, D suffers from both his own negative shock and the component of U's shock that gets passed-through. Therefore under zero output gap marginal costs fall in U and increase in D.

Formally, the pass-through of productivity into D's price is

$$\bar{\delta}_A = \delta_D \left(\overbrace{1}^{\text{direct}} + \overbrace{(1-\alpha_D)\,\delta_U}^{\text{through } U} \right)$$

Therefore we have

$$\bar{\delta}_A > \bar{\delta}_w = \delta_D \left(\alpha_D + (1-\alpha_D)\,\delta_U\right)$$

and consumer inflation is positive.

In this example, the "divine coincidence" does not hold for consumer prices due to the presence of input-output linkages. Consumer inflation is not stabilized because it fails to account for the decrease in upstream prices. This is a general result: Proposition 3 below shows that to restore the "divine coincidence" we need to weight sectors by sales shares, which account for their full contribution to total value added.

Example 2. Horizontal economy

This example introduces another channel through which the "divine coincidence" breaks with consumer inflation. Here there are no input-output linkages, but sectors face idiosyncratic shocks and heterogeneous pricing frictions. In this setup, consumer inflation overrepresents more flexible sectors.

Consider the horizontal economy in Figure (2): there are N sectors, $\{1, ..., N\}$, with consumption shares $\beta_1, ..., \beta_N$ and adjustment probabilities $\delta_1, ..., \delta_N$.

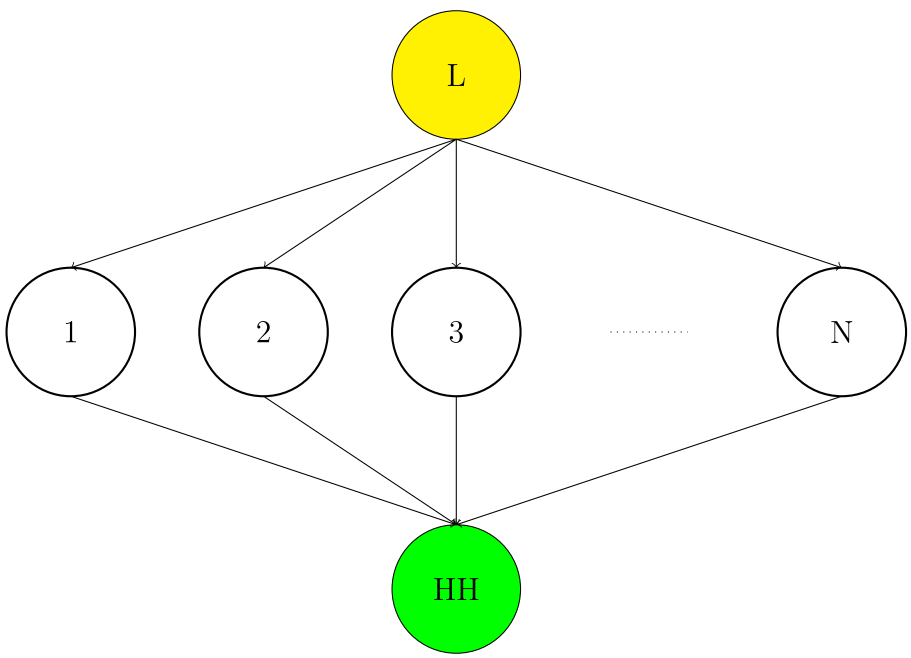

Figure 2: Horizontal economy

We argued that under zero output gap wages adjust to reflect the "average" change in productivity $\mathbb{E}_\beta (dlogA)$. From Proposition 2, inflation in each sector i satisfies

$$\pi_i = \delta_i \left(\underbrace{\frac{1-\bar{\delta}_A}{1-\bar{\delta}_w} \mathbb{E}_\beta (dlogA)}_{\text{wage}} \quad - \quad \underbrace{dlogA_i}_{\text{productivity}} \right) \tag{25}$$

where

$$\bar{\delta}_w = \mathbb{E}_\beta(\delta)$$

$$\bar{\delta}_A = \frac{\mathbb{E}_\beta(\delta dlogA)}{\mathbb{E}_\beta(dlogA)}$$

We see from (25) that inflation increases in sectors that received a worse shock than the "average" $\frac{1-\bar{\delta}_A}{1-\bar{\delta}_w}\mathbb{E}_\beta(dlogA)$, and vice versa.

Consumer inflation is obtained by aggregating sectoral inflation rates according to consumption shares:

$$\pi^C = -\frac{Cov_\beta(\delta, dlogA)}{1-\mathbb{E}_\beta(\delta)} \tag{26}$$

Therefore consumer inflation is negative if productivity increased more than the average in flexible sectors, and vice versa. Thus the divine coincidence does not hold, because consumer inflation overrepresents more flexible sectors. This is a general result. Proposition 3 below derives the correct way to discount sectors based on price flexibility.

Example 3. Oil shocks and consumer inflation

This example presents a stylized "oil economy", to explore the channels through which a negative oil shock can increase consumer prices under zero output gap. Section 6 quantifies these channels in the actual US economy, finding that a 10% negative shock raises consumer prices by 0.22% under zero output gap.

The oil sector is among those with the highest sales share ($\simeq .09$, 99th percentile), although it has a relatively small consumption share ($\simeq .0009$, 69th percentile). This is consistent with the fact that oil is an important intermediate input. To capture this feature I study a simplified economy, given by a combination of a vertical chain and a horizontal economy (see Figure 3).

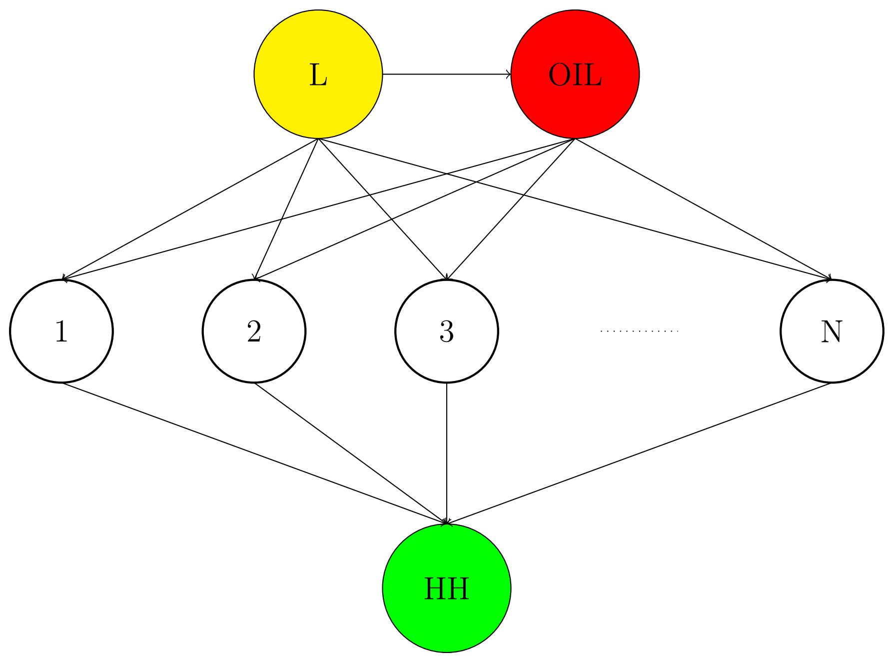

Figure 3: Oil economy

Labor is the first stage, with sticky wages.[10] Then comes oil, and finally the last stage is broken down into multiple sectors. These sectors have heterogeneous consumption shares (β_i), oil input shares ($\omega_{i,oil}$) and adjustment frequencies (δ_i). In the data oil prices are very flexible, therefore oil shocks are (almost) fully passed-through to the final goods sector. On the other hand wages are rigid. As a result oil shocks act rather like downstream shocks, in spite of the role of oil as an intermediate input. We know from Example 1 that negative downstream shocks generate positive consumer inflation. Furthermore, oil input shares and adjustment frequencies are positively correlated in the data. Example 2 shows that this further increases the pressure on consumer prices.

Formally, if $\delta_{oil} = 1$ and workers adjust wages with probability δ_L, under zero output gap consumer inflation is given by

$$\pi^C = -\frac{\overbrace{Cov_\beta(\delta, \omega_{oil})}^{\text{horizontal}} + \overbrace{(1-\delta_L)\,\mathbb{E}_\beta(\delta)}^{\text{vertical}}\,\mathbb{E}_\beta(\omega_{oil})}{1-\delta_L\mathbb{E}_\beta(\delta)} d\log A_{oil} \tag{27}$$

It is immediate to see from (27) that, for $d\log A_{oil} < 0$, consumer inflation π_C is decreasing in wage flexibility δ_L and increasing in the covariance between oil shares and adjustment frequencies.

Table 2 in Section 6.4.2 reports the calibrated response of inflation to an oil shock in the US

[10] See Remark 1 in Section 3.1.1 for a discussion of how I model wage rigidities.

network, under different assumptions about price and wage rigidity. Even if the full network is more complex than the simple economy in this example, the simple model captures well the mechanisms at play. Both wage rigidities and the correlation between oil shares and adjustment frequencies lead negative oil shocks to increase consumer inflation. If instead we assumed wages to be flexible ($\delta_L = 1$) and price adjustment frequencies to be uniform ($\delta_{oil} = \delta_i \equiv \delta_{mean}\ \forall i$), while still accounting for the actual input-output shares, an oil shock would act like an upstream shock, and consumer prices would decrease under zero output gap. The intuition is the same as in Example 1.

4.3 "Divine coincidence" specification

Section 4.1.3 shows that input-output linkages reduce the slope of the consumer-price Phillips curve, and productivity fluctuations result in endogenous cost-push shocks. Examples 1 and 2 suggest that this happens because consumption shares do not capture the contribution of upstream sectors to value added, and consumer prices fail to account for the fact that flexible sectors respond more to a given cost shock. Building on this intuition, Proposition 3 derives an inflation measure that yields a Phillips curve with constant slope and no endogenous residual.

Proposition 3 constructs our "divine coincidence" indicator. It proceeds in two steps. First, Lemma 3 aggregates sector-level markups into the output gap. Second, the pricing equation (2) shows that sector-level markups can be inferred from inflation rates, appropriately corrected for adjustment probabilities. Intuitively, the output gap is proportional to a notion of aggregate markup because of "factor suppression": when markups are higher wages correspond to a smaller fraction of workers' marginal product, which reduces labor supply and output. It turns out that the correct way to aggregate sector-level markups into the output gap is to weigh them according to sales shares.[11]

Changes in markups can then be inferred from inflation rates, because both are determined by the same cost shock. Producers who can adjust their price respond to the shock one-for-one, and this is reflected in inflation rates. Those who cannot adjust instead absorb the shock into their markup. For a given cost shock inflation is higher in flexible sectors, and the change in average markup is smaller. Therefore these sectors need to be discounted.

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

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

[11] This argument is closely related to Proposition 3 in Baqaee and Farhi (2017).

satisfies

$$SW = (\gamma + \varphi)\,\widetilde{y} \tag{28}$$

Moreover, unless prices are fully flexible in all sectors ($\Delta = I$), SW is the only aggregate inflation statistic yields a Phillips curve with no endogenous residual.

Proof. Lemma 3 states that the output gap is proportional to a notion of "aggregate" markup, that weights sector level markups according to sales shares.

Lemma 3. *The following relation holds between the output gap and sector-level markups:*

$$(\gamma + \varphi)\,\tilde{y} = -\lambda^T d\log\mu \tag{29}$$

Proof. See Appendix A3. □

Markups can be inferred from inflation rates and price adjustment probabilities from the pricing equation (2):

$$-\,d\log\mu = (I - \Delta)\,\Delta^{-1}\pi \tag{30}$$

Together, Equations (29) and (30) yield the sales-weighted Phillips curve:

$$\lambda^T (I - \Delta)\,\Delta^{-1}\pi = -\lambda^T dlog\mu = (\gamma + \varphi)\,\tilde{y}$$

Lemma 4 implies that $SW = \lambda^T (I - \Delta)\,\Delta^{-1}\pi$ is the only aggregate inflation statistic such that the corresponding Phillips curve has no endogenous residual.

Lemma 4. *If $\Delta \neq I$ then $\lambda^T (I - \Delta)\,\Delta^{-1}$ is the only vector ν that satisfies*

$$\nu^T \mathcal{V} = 0$$

Proof. See appendix A3. □

□

Remark 3. The weights in SW are all positive. Therefore we can have $\lambda^T (I - \Delta)\,\Delta^{-1}\pi = 0$ only if π_i is positive in some sectors and negative in others, so that under zero output gap there are always sectors where inflation is positive and sectors where it is negative.[12]

[12] We know from Lemma 4 that in general π_i cannot be zero in every sector.

5 Welfare function and optimal policy

Targeting the "divine coincidence" index derived in Proposition 3 closes the output gap. With multiple sectors, however, inflation cannot be stabilized everywhere, even if the output gap is closed (see Corollary 4). Section 5.1 derives the welfare loss as a function of the output gap and sectoral inflation rates, and demonstrates that monetary policy faces a trade-off between closing the output gap and minimizing the distortions associated with sectoral inflation. Targeting the "divine coincidence" index stabilizes the output gap, but ignores the negative consequences of inflation for allocative efficiency. Section 5.2 characterizes the optimal response of monetary policy to this tradeoff, and Section 5.3 shows that it can be implemented by stabilizing an appropriate inflation index, different from both consumer prices and the "divine coincidence" index. Specifically, the optimal target trades off the "divine coincidence" indicator against an inflation statistic that captures the effect of monetary policy on allocative efficiency.

Remark 4. I derive optimal policy in terms of the aggregate output gap, although the actual policy instrument is money supply. This is without loss of generality, because there is a one-to-one mapping between the two. This mapping can be derived from the consumer-price Phillips curve and the cash-in-advance constraint:

$$dlogM = \pi_C + y = \pi_C + \underbrace{\widetilde{y} + \frac{1+\varphi}{\gamma+\varphi}\lambda^T dlogA}_{y_{nat}} =$$

$$= \left(1+\kappa^C\right)\widetilde{y} + u^C + \frac{1+\varphi}{\gamma+\varphi}\lambda^T dlogA$$

5.1 Welfare function

Proposition 4 derives a second-order approximation of the welfare loss with respect to the flex-price outcome. The approximated loss is a quadratic function of inflation and the output gap. The output gap corresponds to a distortion in aggregate demand (see Equation (6)), while inflation is associated with relative price distortions.[13]

In the baseline one-sector model the welfare loss is also a quadratic function of inflation and

[13]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 (30)).

the output gap, given by

$$\mathbb{W} = \frac{1}{2}\left[\ (\gamma+\varphi)\,\tilde{y}^2 + \epsilon\tfrac{1-\delta}{\delta}\pi^2\ \right] \tag{31}$$

In this model inflation reflects relative price distortions between adjusting and non-adjusting firms: the two face the same marginal cost, but end up charging different prices. Customers buy too much of the varieties whose relative price is lower than in the efficient equilibrium. For a given price distortion, quantities respond more if the elasticity of substitution ϵ is higher. Therefore the welfare cost is increasing in ϵ (see Equation (31)).

In the network model instead the welfare cost of inflation comes from both within-sector and cross-sector misallocation. The cross-sector channel is new, while within-sector effects need to be appropriately aggregated. Proposition 4 characterizes the within- and cross-sector components as a function of primitives.

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

$$\mathbb{W} = \frac{1}{2}\left[\ (\gamma+\varphi)\,\tilde{y}^2 + \pi^T\mathcal{D}\pi\ \right] \tag{32}$$

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

$$d^1_{ii} = \lambda_i\epsilon_i\frac{1-\delta_i}{\delta_i} \tag{33}$$

$\mathcal{D}_2$ is positive semidefinite. It can be written as a function of the substitution operators in production and consumption defined below.[14]

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

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

$$+\alpha_t\sum_k \omega_{tk}\theta^t_{kL}X_kY_k$$

and

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

[14] 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.

The elements of $\mathcal{D}_2$ are given by

$$d_{ij}^2 = \frac{1-\delta_i}{\delta_i}\frac{1-\delta_j}{\delta_j}\left(\Phi_C\left((I-\Omega)^{-1}_{(i)}, (I-\Omega)^{-1}_{(j)}\right) + \sum_t \lambda_t \Phi_t\left((I-\Omega)^{-1}_{(i)}, (I-\Omega)^{-1}_{(j)}\right)\right) \tag{34}$$

Proof. See appendix B1 □

Although the within- and cross-sector components of the welfare loss are not conceptually different,[15] it is convenient to consider them separately.

From equation (33) 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 should then be aggregated by sales shares, discounting flexible sectors. The intuition is the same as in Proposition 3. Overall, the within-sector component of the total welfare loss is given by

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

Equation (35) re-writes Equation (34) to highlight that the total welfare loss from cross-sector misallocation is equivalent to a combination of negative sector-level productivity shocks. 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 aggregate loss is given by

$$\pi^T \mathcal{D}_2 \pi = \sum_t \underbrace{\lambda_t}_{\text{aggregation}} \underbrace{\sum_{i,j} \Phi_t(i,j)}_{\text{productivity loss in sector } t} \tag{35}$$

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

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

The aggregate effect is obtained by weighting sectors according to sales shares λ_t, as in Hulten's formula. To derive the productivity loss for each sector t we proceed in two steps. First we isolate the distortionary component of sectoral inflation rates and track its propagation across the network,

[15] They could be unified into the cross-sector component if we considered a fully disaggregated model, where sectors are identified with individual firms.

which results in relative price distortions across t's inputs. Intuitively, inflation is associated with a distortion only to the extent that it mirrors the change in the markup of non-adjusting firms. We want to extract this inefficient component, and study its propagation into the relative prices of buyer sectors. Second, we translate relative price distortions into t's productivity loss from inefficient substitution. This is given by the substitution operators.

I define relative prices with respect to nominal wages. Lemma 5 provides the mapping between inflation rates and relative price distortions.

Lemma 5. *Relative price distortions with respect to the flex-price outcome are given by*

$$d\log p - d\log w = (I-\Omega)^{-1}(I-\Delta)\Delta^{-1}\pi \tag{36}$$

Proof. See Appendix B1 □

From Equation (36), relative price distortions can be decomposed into a direct and a propagation effect:

$$d\log p - d\log w = \underbrace{(I-\Omega)^{-1}}_{\text{propagation}}\underbrace{(I-\Delta)\Delta^{-1}\pi}_{\text{markup (direct)}}$$

Here is the intuition. A distortion in the relative price of a sector k can come either directly from a change in k's markup, or indirectly from a change in the markup of some of its inputs. The mapping between markups and inflation rates is given by equation (30):

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

This corresponds to the direct effect. The Leontief inverse $(I-\Omega)^{-1}$ captures the propagation effect: the price distortion induced in sector k by a change in i 's markup is given by $(I-\Omega)^{-1}_{ik}\,d\log\mu_i$, or equivalently by $(I-\Omega)^{-1}_{ik}\frac{1-\delta_i}{\delta_i}\pi_i$, using the mapping between markups and inflation rates. From Lemma 5 we can also derive the relative price distortion between each sector pair (k,h) associated with inflation in sector i. This is given by

$$\left((I-\Omega)^{-1}_{ki}-(I-\Omega)^{-1}_{hi}\right)\frac{1-\delta_i}{\delta_i}\pi_i$$

Thus, Lemma 5 characterizes the relative price distortions generated by sectoral inflation. The productivity effect on each sector t is captured by the corresponding substitution operator Φ_t (see Definition 2), and depends on the interaction between distortions associated with inflation different sectors. Intuitively, distortions from 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. This is captured by $\Phi_t(i,j)$, which measures the productivity loss of sector t induced by a 1% increase in i's inflation, given that j's also increased by 1%. $\Phi_t(i,j)$ weights each pair (k,h) by the relevant input shares ω_{ti} and ω_{tj}, and the substitution elasticity θ^t_{kh}:

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

When elasticities of substitution are uniform ($\theta^t_{kh} \equiv \theta^t$), the substitution operator is simply given by 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 $\{\omega_{tk}\}_{k=1..N}$:

$$\Phi_t\left(i,j\right) = \theta^t Cov_{\Omega_t}\left(\underbrace{(I-\Omega)^{-1}_{(i)}\frac{1-\delta_i}{\delta_i}\pi_i}_{\text{distortion from } i}, \underbrace{(I-\Omega)^{-1}_{(j)}\frac{1-\delta_j}{\delta_j}\pi_j}_{\text{distortion from } j}\right)$$

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 obtained from Hulten's formula (see Equation (35)).

5.2 Optimal policy

Optimal monetary policy minimizes the welfare loss, subject to the response of inflation to the output gap and productivity. In the one-sector model this constraint is given by the Phillips curve, and the central bank solves

$$\begin{aligned} min_{\pi,\tilde{y}} \;\; \mathbb{W} = \tfrac{1}{2}\left[\; (\gamma+\varphi)\,\tilde{y}^2 + \epsilon\tfrac{1-\delta}{\delta}\pi^2 \;\right] \\ s.t.\; \pi = \kappa\tilde{y} \end{aligned} \tag{37}$$

The "divine coincidence" holds, and optimal policy achieves the first best by setting $\pi = \tilde{y} = 0$. With multiple sectors the optimal policy problem extends the baseline in two dimensions. First, the inflation term is replaced by the more complex misallocation loss derived in Proposition 4.

Second, the constraint is not given by the aggregate Phillips curve, but it takes into account the full response of inflation rates sector-by-sector. Thus the problem becomes:

$$\begin{aligned} min_{\tilde{y},\pi} \ \tfrac{1}{2} \Big[\ (\gamma+\varphi)\, \tilde{y}^2 + \pi^T \mathcal{D} \pi \ \Big] \\ s.t. \ \pi = \mathcal{B}\tilde{y} + \mathcal{V} dlogA \end{aligned} \tag{38}$$

In the multi-sector model there is a trade-off between within- and cross-sector misallocation. Within sectors, all firms face the same marginal cost, because they have the same production function. Therefore they should all charge the same price. Given that some firms cannot change their price, within-sector distortions are eliminated only if nobody adjusts. However keeping sector-level prices fixed is inefficient if relative productivities have changed. In the efficient equilibrium, relative prices are given by

$$\pi = (I-\Omega)^{-1}\, d\log A$$

As a result, even with a full set of sector-level subsidies it would be impossible to replicate the efficient outcome. When the output gap is closed inflation is not stabilized, and sector-level relative prices are different from the efficient equilibrium. They are given by

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

Monetary policy is constrained to use a one-dimensional instrument (interest rates or money supply), that jointly affects the output gap and misallocation. This is reflected in the constraint, which shows that monetary policy can only implement relative price changes that are proportional to the vector $\mathcal{B}$ of sectoral elasticities with respect to the output gap. As a result, the central bank needs to tradeoff distortions in aggregate demand (associated with the output gap) against allocative efficiency considerations. Proposition 5 characterizes the solution to the optimal policy problem.

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

$$\widetilde{y}^* = -\frac{\mathcal{B}^T \mathcal{D} \mathcal{V} dlogA}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}} \tag{39}$$

Proof. The result follows immediately from solving the minimization problem (38). □

The optimal output gap trades off the marginal cost and benefit of deviating from the efficient aggregate demand level. The denominator in equation (39) reflects the marginal cost, and it is always positive. The cost comes directly from distortions in aggregate demand (whose welfare effect is proportional to the labor supply elasticity $(\gamma+\varphi)$), and indirectly from relative price distortions caused by the output gap (given by $\mathcal{B}^T\mathcal{D}\mathcal{B}\tilde{y}^2$). The numerator in (39) is the marginal benefit. Here $\tilde{\pi}^T\mathcal{D}\pi$ is the marginal cost of inducing inflation $\tilde{\pi}$ for given current inflation π. From the constraint we see that the output gap affects prices proportionately to their elasticity $\mathcal{B}$. Therefore the marginal benefit of increasing the output gap on misallocation given that current inflation is π is $-\mathcal{B}^T\mathcal{D}\pi$. Under zero output gap, the productivity shock induces inflation

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

so that the overall marginal benefit is given by $-\mathcal{B}^T\mathcal{D}\mathcal{V}d\log A$.

5.3 Inflation targeting

We argued that in the multi-sector model closing the output gap is not optimal, because it ignores the negative consequence of sector-level inflation for allocative efficiency. Targeting the "divine coincidence" index stabilizes the output gap, therefore it is not constrained efficient. Nonetheless, Proposition 6 demonstrates that the optimal policy can still be implemented by stabilizing an appropriate inflation index, which incorporates the tradeoff between aggregate demand and allocative efficiency.

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 $\widetilde{y} > \widetilde{y}^*$.

This vector is given by

$$\phi^T = \lambda^T (I-\Delta)\Delta^{-1} + \mathcal{B}^T\mathcal{D} \tag{40}$$

Proof. See Appendix B2 □

From the "divine coincidence" Phillips curve we can re-write the target in terms of inflation and the output gap:

$$\phi^T\pi = (\gamma+\varphi)\tilde{y} + \mathcal{B}^T\mathcal{D}\pi \tag{41}$$

Equation (41) shows that the optimal target weights the output gap against sectoral inflation rates according to the relative marginal benefit and marginal cost of distorting aggregate demand to counteract the misallocation associated with inflation. From our derivation of the welfare function in Proposition 4 we have that the marginal benefit is $\mathcal{B}^T\mathcal{D}\pi$, while the marginal cost is $\gamma+\varphi$.

This result extends with minimal modifications to the dynamic setup (see Appendix C). Here the optimal policy can be implemented via a Taylor rule which targets the inflation statistic in Proposition (6), with an additional correction for inflation expectations.

5.4 Examples

The examples below illustrate optimal monetary policy in the three simple networks introduced in Section 4.2 (vertical chain, horizontal economy and oil economy).

Example 4. Optimal policy in the vertical chain

Consider a negative downstream shock in a two-stage vertical chain, as in Example 1. Can monetary policy do better than implementing a zero output gap? Which sector shall it seek to stabilize?

Recall that marginal costs increase downstream and fall upstream under zero output gap. A positive output gap then increases inflation downstream, while it stabilizes the upstream sector. I argue that this is what the optimal policy should do, because distortions are more costly in the upstream sector, and this sector is also easier to stabilize.

In the vertical chain, the only source of cross-sector misallocation is substitution by D between labor and the intermediate input produced by U. There is no misallocation across consumption goods, because there is only one of them, nor across U's inputs, because it only uses labor. Misallocation between U and labor happens because U's price does not fully adjust to reflect the change in labor costs. Monetary policy can offset this effect by stabilizing wages, and thereby reducing U's desired price adjustment. After a negative productivity shock this entails implementing a positive output gap. This is reflected in the cross-sector component of the optimal output gap:

$$-\frac{(\gamma+\varphi)\,\theta_L^D}{(1-\delta_D\,(1-(1-\alpha_D)\,\delta_U))^2}\left[\underbrace{(1-\alpha_D)\,\alpha_D}_{\text{input shares}}\quad\underbrace{(1-\delta_U)^2\,(1-\delta_D)}_{\propto\frac{1-\delta_U}{\delta_U}\pi_U}\right]dlogA_D>0$$

The within-sector component of the optimal output gap instead is

$$-\epsilon\frac{(\gamma+\varphi)\,(1-\alpha_D)\,(1-\delta_U)\,(1-\delta_D)}{(1-\delta_D\,(1-(1-\alpha_D)\,\delta_U))^2}\left[\underbrace{\delta_U}_{\text{benefit for }U}\quad-\underbrace{\delta_D\,(1-(1-\alpha_D)\,\delta_U)}_{\text{cost for }D}\right]dlogA_D$$

Here $\frac{\delta_D}{\delta_U}$ is the relative cost of within-sector price dispersion in D and U, and $1-(1-\alpha_D)\,\delta_U$ is the relative marginal effect of monetary policy on inflation in the two sectors. We argued that stabilizing sector U entails implementing a positive output gap. This policy reduces the overall within-sector misallocation if and only if the benefit for U (δ_U) is greater than the loss for D ($\delta_D\,(1-(1-\alpha_D)\,\delta_U)$). This is always the case if adjustment frequencies are the same in the two sectors.

Example 5. Optimal policy in the horizontal economy

Consider the horizontal economy in Example 2. The tradeoff between within- and cross-sector misallocation is particularly clear in this setup. On one hand, stabilizing within-sector misallocation would require all firms to charge the same price, which must be the same as in the initial equilibrium, because some firms cannot adjust. On the other, hand under zero output gap sectors that faced a worse shock than the average see an increase in their marginal cost, and vice versa. This should be reflected in relative prices. Therefore any improvement in within-sector misallocation comes at the cost of worse cross-sector misallocation (and vice versa).

Indeed, the relative welfare effect of the output gap on these two channels is given by the ratio of within- and cross-sector elasticities of substitution. The two have opposite sign, reflecting the fact that correcting within- vs cross-sector misallocation requires opposite price adjustments. The marginal gain of increasing the output gap for within-sector misallocation is given by

$$\epsilon\mathbb{E}_{\beta(1-\delta)}\pi$$

and the corresponding cross-sector component is

$$-\sigma\mathbb{E}_{\beta(1-\delta)}\pi$$

Here I denoted by $\mathbb{E}_{\beta(1-\delta)}$ the expectation computed with probability weights

$$\left\{\frac{\beta_i\,(1-\delta_i)}{\sum_j \beta_j\,(1-\delta_j)}\right\}_{i=1,\ldots,N}$$

Stabilizing within-sector misallocation is optimal if and only if the corresponding elasticity is larger than the cross-sector one. Sectors with large consumption share are given higher weight. Flexible sectors are discounted, because only non-adjusting firms face markup distortions.

We can derive the marginal gain of increasing the output gap as a function of productivity by solving for the sector-level inflation induced by the shock under zero output gap. The cross-sector component is

$$-\sigma(\gamma+\varphi)\,Cov_{\beta(1-\delta)}(\delta, d\log A) \tag{42}$$

The within-sector component can be obtained from (42), multiplying times $-\frac{\epsilon}{\sigma}$. Like the cost-push shock in Example 2, the optimal output gap depends on the correlation between productivity shocks and adjustment frequencies. As discussed in Example 2, this correlation reflects the competing effect of wages and productivity on marginal costs, prices and markups. For the optimal policy, however, the relevant "probability weights" are not just given by plain consumption shares, but they also account for the fraction of non-adjusting firms. Intuitively, fully flexible sectors price at their marginal cost, therefore they do not face any distortion and do not need to be stabilized.

Example 6. Optimal policy in the oil economy

Let's go back to the simple model of an oil economy in Example 3, given by a combination of a vertical chain and a horizontal economy.

Here all the cross-sector misallocation comes from distortions across final goods, and not across the three "stages" (labor, oil and final goods), because oil uses only labor as an input, and we assumed that oil prices are fully flexible. Therefore the cross-sector component of the optimal output gap is the same as the horizontal economy, given by:

$$-\sigma(\gamma+\varphi)\,\mathbb{E}_{\beta(1-\delta)}\pi$$

Writing inflation as a function of the productivity shock, we obtain:

$$\sigma(\gamma+\varphi)\frac{\delta_L(1-\mathbb{E}_\beta(\delta))}{1-\delta_L\mathbb{E}_\beta(\delta)}\left[\mathbb{E}_{\beta(1-\delta)}(\delta\omega_{oil})-\frac{\delta_L(1-\mathbb{E}_\beta(\delta))}{1-\delta_L\mathbb{E}_\beta(\delta)}\mathbb{E}_{\beta(1-\delta)}(\delta)\,\mathbb{E}_{\beta(1-\delta)}(\omega_{oil})\right]dlogA_{oil} \tag{43}$$

With respect to the horizontal economy (see equation (42)), we introduced sticky wages. When wages are fully rigid ($\delta_L = 0$) monetary policy has no effect on marginal costs and therefore on markup distortions, so that the optimal output gap is zero. In addition, in Examples 2 and 5 we argued that the covariance between productivity shocks and adjustment frequencies captures the countervailing effects of wages and productivity on marginal costs and markups. In equation (43) sectoral productivity shocks are replaced by oil input shares, that reflect the pass-through of the oil shock into sectoral marginal costs. When $\delta_L < 1$ the wage channel is muted, and the productivity effect becomes more important. As a result we have:

$$\mathbb{E}_{\beta(1-\delta)}(\delta\omega_{oil})-\frac{\delta_L(1-\mathbb{E}_\beta(\delta))}{1-\delta_L\mathbb{E}_\beta(\delta)}\mathbb{E}_{\beta(1-\delta)}(\delta)\,\mathbb{E}_{\beta(1-\delta)}(\omega_{oil}) > Cov_{\beta(1-\delta)}(\delta,\omega_{oil})$$

Within-sector misallocation comes from both final goods and labor. Consistent with the intuition from the horizontal economy, within- and cross-sector misallocation have a common component, with opposite sign. From Example (5) we know that, if $\epsilon > \sigma$, optimal policy stabilizes within-sector misallocation. In addition, there is within-sector misallocation in the labor stage as well. From the vertical chain example we know that this can be stabilized with a positive output gap: given that aggregate productivity is lower, wages should decrease. A positive output gap increases labor demand, with an offsetting effect on wages. As a result, the distortion between adjusting and non-adjusting workers becomes smaller, thereby reducing misallocation in the labor sector. The overall effect is in equation (44):

$$\epsilon\left(\gamma+\varphi\right)\frac{\delta_L\mathbb{E}_\beta\left(1-\delta\right)}{1-\delta_L\mathbb{E}_\beta\left(\delta\right)}\left[\underbrace{\left(1-\frac{\delta_L\left(1-\mathbb{E}_\beta\left(\delta\right)\right)}{1-\delta_L\mathbb{E}_\beta\left(\delta\right)}\right)\mathbb{E}_{\beta(1-\delta)}\left(\omega_{oil}\right)}_{\text{vertical chain}}\right.$$
$$\left.-\left[\underbrace{\mathbb{E}_{\beta(1-\delta)}\left(\delta\omega_{oil}\right)-\frac{\delta_L\left(1-\mathbb{E}_\beta\left(\delta\right)\right)}{1-\delta_L\mathbb{E}_\beta\left(\delta\right)}\mathbb{E}_{\beta(1-\delta)}\left(\delta\right)\mathbb{E}_{\beta(1-\delta)}\left(\omega_{oil}\right)}_{\text{horizontal economy}}\right]\right]dlogA_{oil} \tag{44}$$

Quantitatively, the calibration in Section 6.4.2 shows that the optimal output gap is positive.

6 Quantitative analysis

6.1 Data

I calibrate the input-output matrix Ω and the labor and consumption shares α and β based on input-output tables published by the BEA. 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 6.3.3 also relies on less disaggregated historical input-output data (46 industries), always from the BEA input-output accounts, to study the slope of the Phillips curve and monetary non-neutrality over time.

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 use industry-level frequencies of price adjustment constructed by Pasten, Schoenle and Weber (2017). For sectors with missing data I set the adjustment probability equal to the mean. I calibrate the quarterly probability of wage adjustment to 0.25, in line with Barattieri, Basu and Gottschalk (2014) and Beraja, Hurst and Ospina (2016).

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 $\sigma = 0.9$,[16] the elasticity of substitution between labor and intermediate inputs to $\theta_L = 0.5$,[17] the elasticity of substitution across intermediate inputs to $\theta = 0.001$,[18] and the elasticity of substitution between varieties within each sector to $\epsilon = 8$.[19]

I obtain estimates of annual industry-level TFP changes for the period 1988-2016 from the BEA Integrated Industry-Level Production Account data.

6.2 Welfare loss from business cycles

As discussed in Section 5, in a multi-sector economy monetary policy cannot replicate the flex-price efficient outcome, because it faces a tradeoff between stabilizing aggregate demand and reducing misallocation. As a result, business cycle fluctuations generate a welfare loss with respect to an efficient economy without pricing frictions. In this section I calibrate the expected welfare loss under different policy rules. I assume that productivity shocks are normally distributed, with zero mean and covariance Σ.

I calibrate Σ from sector-level TFP shocks as measured in the BEA KLEMS data.[20] I refer to the Multifactor Productivity (MFP) measure, and calibrate productivity shocks as the growth rate of this index at the sector level. 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.

I compare the results with two counterfactual scenarios, one where industry-level shocks are iid, and the other where they are perfectly correlated. I keep the variance of aggregate productivity constant across the three specifications. Analytical expressions for the welfare loss under various

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

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

[18]See Atalay (2017).

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

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

policy rules are provided in the online Appendix E.

The results are plotted in Figures 4 and 5. The bars correspond to the percentage of per-period GDP that consumers would be willing to forego in exchange of going from a world with sticky prices to one without pricing frictions, for the given monetary policy rule. Bars of different colors represent different rules (blue for optimal policy, red for output gap targeting, and yellow for consumer inflation targeting).

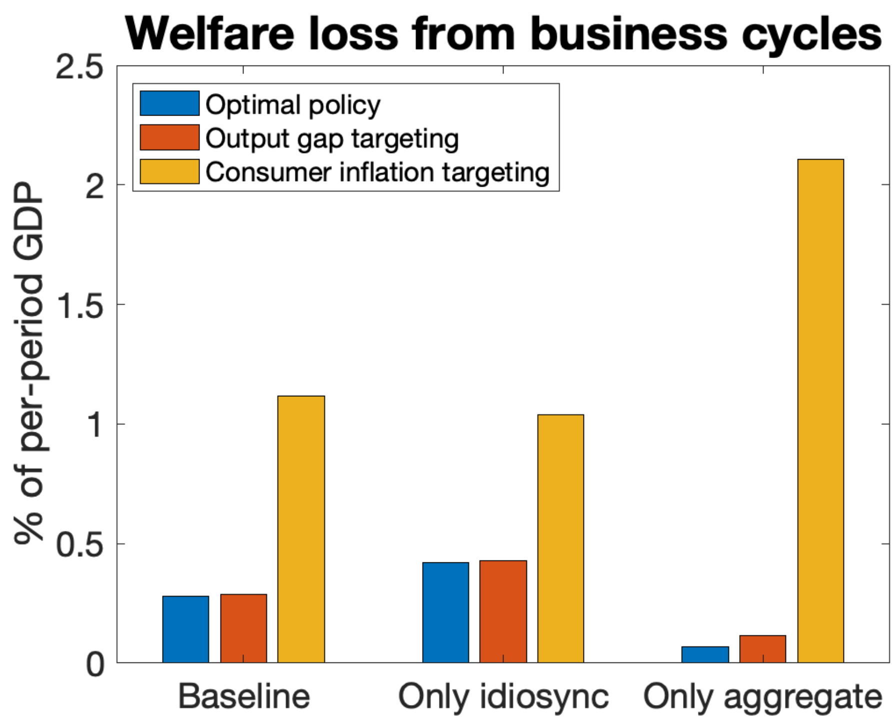

Figure 4: Actual input-output network; different calibrations keep the variance of aggregate output constant

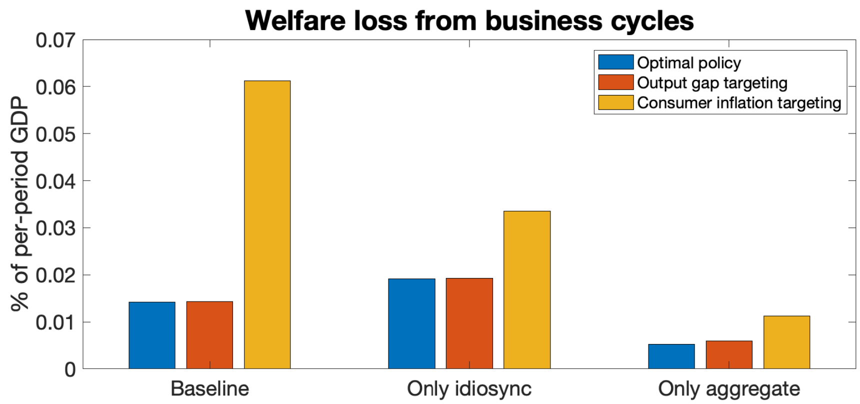

Figure 5: Model with no input-output linkages; different calibrations keep the variance of aggregate output constant

6.2.1 Optimal policy

The blue bars in Figure 4 represent the welfare loss under the optimal policy. In the full calibration business cycles generate a welfare loss of 0.28% of per-period GDP, with respect to an economy without pricing frictions. The idiosyncratic component is the main driver. This result suggests that pricing frictions significantly amplify the negative welfare consequences of business cycle fluctuations. In an efficient economy without pricing frictions we have the well-known Lucas' estimate of the welfare cost of business cycles, which is about 0.05% of per-period GDP.[21] The Lucas' estimate is one order of magnitude smaller than the additional loss induced by price rigidities. This is in stark contrast with the one-sector model, where monetary policy can stabilize inflation and the output gap at the same time, thereby replicating the first best. In this model therefore pricing frictions generate no additional welfare loss with respect to the efficient economy.

Input-output linkages are key for this result. Figure 5 shows that the loss is much smaller in an economy with the same wage rigidity and heterogeneous price adjustment frequencies, but without input-output linkages.[22]

6.2.2 Targeting the output gap

The red bars in Figures 4 and 5 show that on average targeting zero output gap yields a small loss with respect to the optimal policy. This result is consistent with our discussion in Section 5.1. Monetary policy does need to tradeoff aggregate demand and misallocation. However it can only affect relative prices proportionately to their elasticity with respect to the output gap. This makes

[21] Here the welfare cost comes from the uncertainty generated by fluctuations in consumption.

[22] Here consumption shares are calibrated to replicate relative sales shares.

monetary policy a blunt tool to correct relative price distortions. Therefore it should place higher weight on aggregate demand, and stabilize the output gap.

We reach a similar conclusion when comparing the behavior over time of the "divine coincidence" SW, which is our inflation proxy for the output gap, and the optimal policy target (see Figure 6).

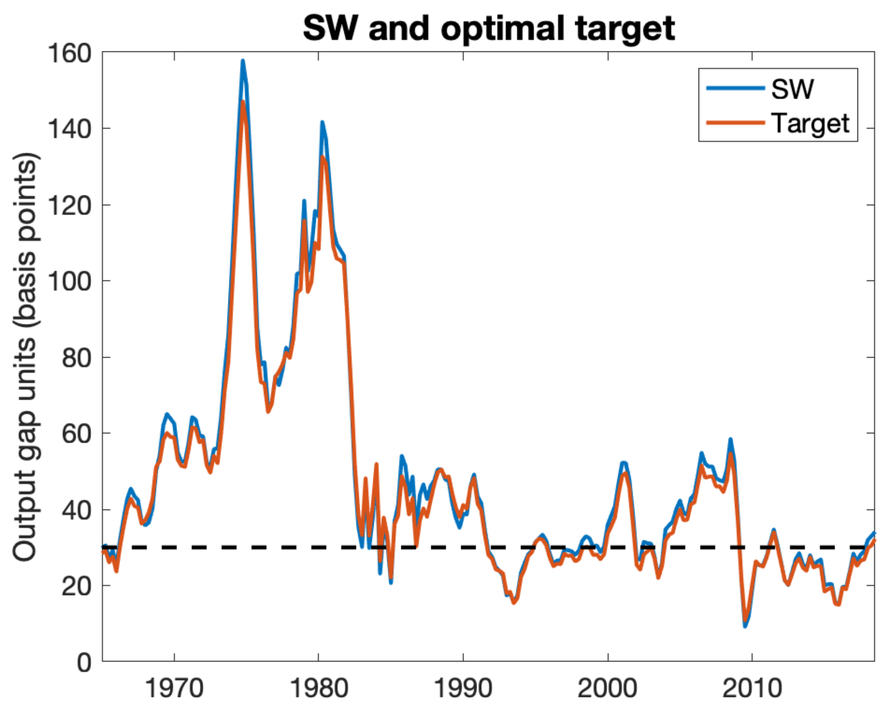

Figure 6: Time series of aggregate inflation (SW) and the optimal policy target

The two series move closely together, which means that the optimal target almost coincides with the output gap. The target however is often a few basis points lower than SW, suggesting that the optimal policy should be slightly more expansionary than output gap targeting.

6.2.3 Targeting consumer inflation

The welfare loss under consumer inflation targeting is represented by the yellow bars in Figures 4 and 5. While stabilizing consumer inflation is optimal in the baseline model, in the multi-sector model it yields a much larger welfare loss than the optimal policy. This result crucially depends on the input-output structure. With the full input-output network (Figure 4), targeting consumer prices yields poor results even with only aggregate shocks. Figure 5 instead shows that the loss is much smaller in the calibration without input-output linkages, regardless of the distribution of the shocks.

6.2.4 Within- versus cross-sector misallocation

Section 5.1 shows that the welfare loss from misallocation has two components, coming from relative price distortions within and across sectors. Figure 7 compares the relative magnitude of these two components. The within-sector loss dominates in the main calibration, which assumes higher substitutability between varieties from the same sector than across goods from different sectors. This result is reversed in an alternative calibration, which assumes the same elasticity of substitution within and across sectors.

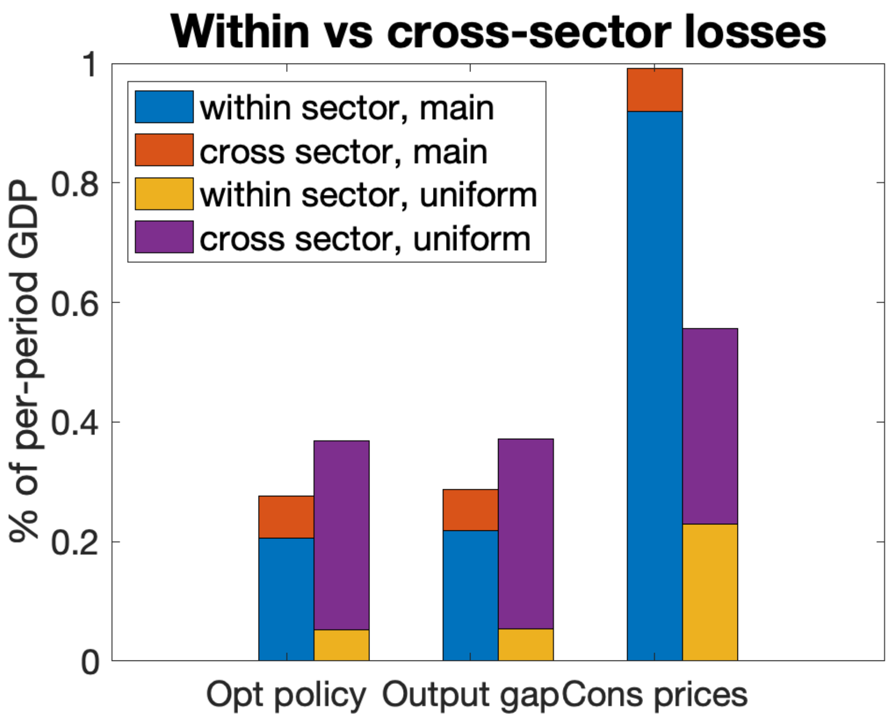

Figure 7: Main calibration: $\epsilon = 8$, $\sigma = 0.9$, $\theta_L = 0.5$, $\theta = 0.001$; uniform elasticities: $\epsilon = \sigma = \theta_L = \theta = 2$

6.3 Slope of the Phillips curve and monetary non-neutrality

This section evaluates the quantitative importance of input-output linkages and heterogeneous pricing frictions for the slope of the Phillips curve and monetary non-neutrality. I measure monetary non-neutrality from the impulse response of inflation to a given real rate shock.[23] I compare

[23]Monetary non-neutrality is closely related with the slope of the Phillips curve. See section 6.3.2 below for a discussion.

the slope and impulse responses in the full calibration and in calibrations that ignore input-output linkages and/or heterogeneous adjustment frequencies. I use historical input-output data to study their evolution over time.

6.3.1 Slope of the Phillips curve

Section 4.1.2 demonstrates that input-output linkages flatten the Phillips curve. The calibration suggests that this effect is quantitatively large. Table 1 shows the slope of the Phillips curve in the full calibration and in alternative calibrations, that assume away input-output linkages, wage rigidities and heterogeneous pricing frictions. The slope in the full calibration is more than one order of magnitude smaller than in the baseline, which ignores both input-output linkages and wage rigidities. Input-output linkages alone account for a large share of the effect: from the second column we read that the slope is 60% smaller in the full calibration than in a model without input-output linkages, but with the same price and wage rigidities. The third column instead shows that the slope is about the same in the full model and in a calibration that sets all adjustment frequencies equal to the mean. This is not a general result, but it is driven by the specific joint distribution of labor shares and adjustment frequencies that we observe in the data. We find that adjustment frequencies are not correlated with labor shares, so that setting them equal to the mean does not affect the wage pass-through $\bar{\delta}_w$ (see Section 4.1.2), which in turn determines the slope of the Phillips curve. Heterogeneity in price stickiness instead matters in the dynamic version of the model, where it reduces the inflation response to monetary shocks (see Section 6.3.2).[24]

	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: Slope in the full calibration and alternative calibrations

6.3.2 Monetary non-neutrality

This section studies the effect of input-output linkages and heterogeneous pricing frictions for monetary non-neutrality, using the dynamic version of the model presented in Appendix C. Before

[24]This is because in the dynamic setting the inflation response does not only depend on the slope of the Phillips curve, which captures an instantaneous response, but it also depends on the expected speed of adjustment over time.

presenting the calibration, I introduce some key theoretical results. I first derive the law of motion of sectoral inflation rates, and then aggregate it into the law of motion of consumer prices. I discuss how these differ from the one-sector benchmark, and what are the channels that dampen the response of consumer prices in the multi-sector model.

The law of motion of sectoral inflation is given by

$$\pi_t = \mathcal{B}\tilde{y}_t + \mathcal{V}\log\mu_{t-1} + \rho\mathcal{M}\mathbb{E}\pi_{t+1} \tag{45}$$

Different from the one-sector model, with multiple sectors past markups are a state variable. Therefore they appear in Equation (45). The matrix $\mathcal{M}$ (derived in Appendix C) captures the propagation of markups shocks through the network over time. Iterating forward the expression in (45) we obtain

$$\pi_t = \sum_{s\geq 0} (\rho\mathcal{M})^s \left[\mathcal{B}\mathbb{E}_t\tilde{y}_{t+s} + \mathcal{V}\log\mu_{t+s-1}\right] \tag{46}$$

Aggregating both sides of (46) according to consumption shares we find

$$\pi_t^C = \sum_{s\geq 0} \beta^T (\rho\mathcal{M})^s \left[\mathcal{B}\mathbb{E}_t\tilde{y}_{t+s} + \mathcal{V}\log\mu_{t+s-1}\right] \tag{47}$$

In contrast with the baseline model, the impulse response of consumer inflation to real rate shocks is not fully characterized by the slope of the Phillips curve, for two reasons. First, the response of current inflation π_t^C to future discounted output gaps $\rho^s\tilde{y}_{t+s}$ is not constant in s. Second, while the output gap captures an appropriately defined aggregate markup (see Lemma 3), inflation also responds to relative markups $\{\mathcal{V}\log\mu_{t+s-1}\}_{s=0}^{\infty}$.

The first term in Equation (47) captures the response of current inflation π_t^C to discounted future output gaps $\rho^s\tilde{y}_{t+s}$: it is the dynamic analogue of the slope of the Phillips curve. In the one-sector benchmark the elasticity of current prices with respect to discounted future output gaps is precisely the slope of the Phillips curve. With multiple sectors the response to future output gaps is mediated by the powers of the matrix $\mathcal{M}$. Intuitively, output gaps have a differential effect across sectors. Marginal costs respond more in upstream sectors, and it takes time for the shock to propagate downstream. The response of current inflation π_t^C to future output gaps $\rho^s\tilde{y}_{t+s}$ takes this into account, therefore it is not constant with respect to the time horizon s.[25]

The second term in Equation (47) captures anticipated changes in relative markups, which arise because of heterogeneous shock pass-through and adjustment frequencies. Indeed, it is through

[25] In the calibration it is bell-shaped. This is intuitive: a change in the current output gap today mostly affects upstream sectors, which have little impact on consumer prices. As time passes the shock propagates downstream, generating a larger response of consumer prices.

this second channel that heterogeneous adjustment frequencies increase monetary non-neutrality in the dynamic setting, whereas we found in Section 6.3.1 that they have no effect in the static model. To gain intuition consider two scenarios, both with the same average probability of price adjustment across sectors. In the first scenario 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 can adjust very infrequently, producers in the flexible sectors know that they will likely have another opportunity to reset their price before the stickier sectors also get to change theirs. Therefore it is optimal for them to wait. This is captured by the expected markup term in equation (47): slower adjustment implies larger anticipated markup distortions. The expectations channel gets muted as the discount factor goes to zero. This is why heterogeneous adjustment frequencies play no role in the static setting.

Figure 8 plots the impulse response of consumer inflation to a 1% real rate shock. It shows that both heterogeneous adjustment frequencies and input-output linkages dampen the response of consumer prices.[26]

[26] my results are consistent previous works, such as Carvalho (2006), Nakamura and Steinnson (2010).

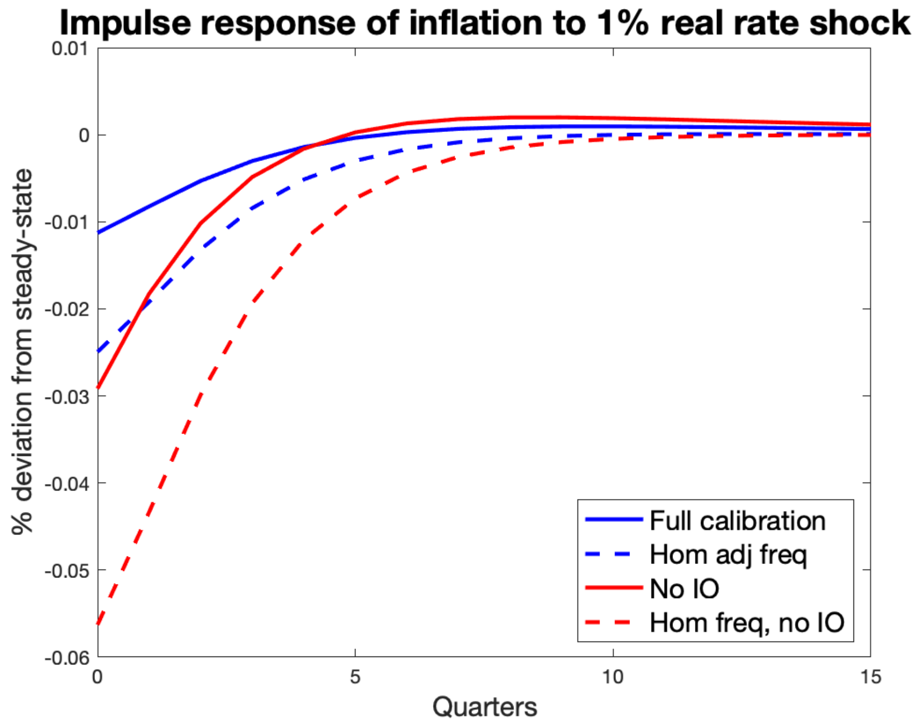

Figure 8: Response to a 1% interest rate shock under a Taylor rule with $\varphi_\pi = 1.24$ and $\varphi_y = .33/12$. Persistence $= 0.5$

6.3.3 Phillips curve and monetary non-neutrality over time

The model allows to predict how the Phillips curve evolves with the production structure. I calibrate the slope using historical input-output data for each year between 1947 and 2017. Due to lack of data, I assume the frequencies of price adjustment to be constant over time. The results are plotted in Figure 9.

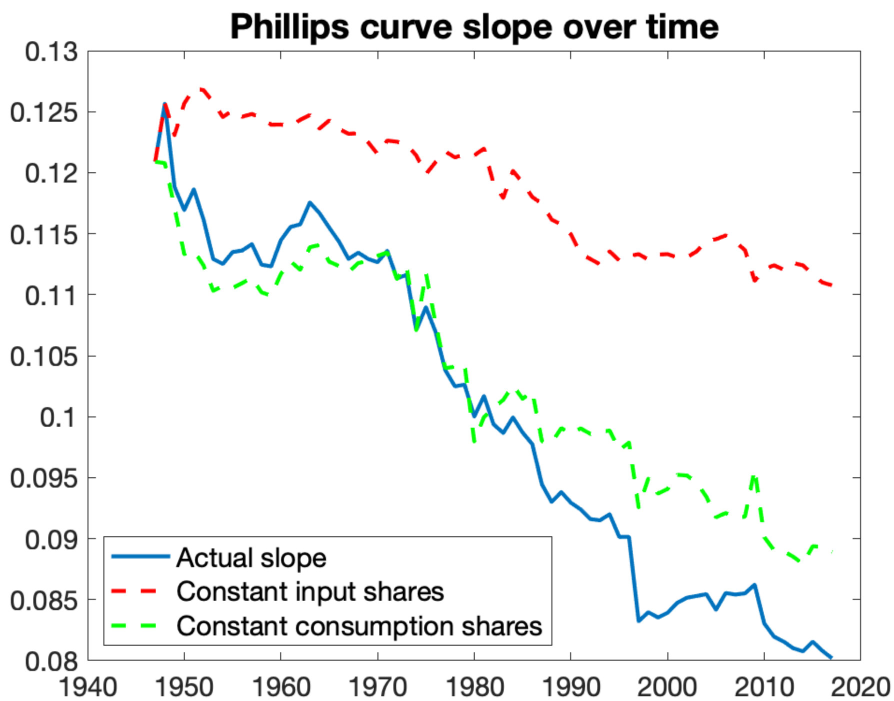

Figure 9: Slope of the Phillips curve over time

The solid line depicts the calibrated slope. We see that it decreased by about 30% over this time period. The total change can be decomposed into two effects. The first comes from an increase in intermediate input flows, while the second comes from a shift in consumption shares away from manufacturing and towards services. Section 6.3.1 shows that the slope is determined by the pass-through of nominal wages into consumer prices, $\bar{\delta}_w$. I will therefore study how these two channels affect $\bar{\delta}_w$. Using the results from Section 6.3.1, we can immediately decompose $\bar{\delta}_w$ into a term related with consumption shares, and a term related with the input-output structure:

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

The dashed green line in Figure 9 plots the slope for constant consumption shares (set at their 1947 value), while letting the input-output matrix change over time. The dashed red line instead keeps input flows constant at their 1947 value and lets consumption shares change. The figure shows that increasing intermediate input flows account for most of the decline until 1980, whereas shifts in consumption shares also have a role after 1980.

Formally, we can split the overall change in $\bar{\delta}_w$ into a change in sector-level pass-through for constant consumption shares, and a change in consumption shares for constant pass-through:

$$\bar{\delta}_w^{2017} - \bar{\delta}_w^{1947} = \frac{\beta_{1947}^T + \beta_{2017}^T}{2} \left(PT_{2017} - PT_{1947}\right) +$$

$$+ \left(\beta_{2017}^T - \beta_{1947}^T\right) \frac{PT_{1947} + PT_{2017}}{2}$$

where I used the notation

$$PT \equiv \Delta \left(I - \Omega\Delta\right)^{-1} \alpha$$

I find that 79% of the overall decline in $\bar{\delta}_w$ can be attributed to increasing intermediate input shares, while the remaining effect comes from shifts in the composition of the consumption basket. Figure 6.3.3 further decomposes the changes in consumption and input-output shares across sectors.

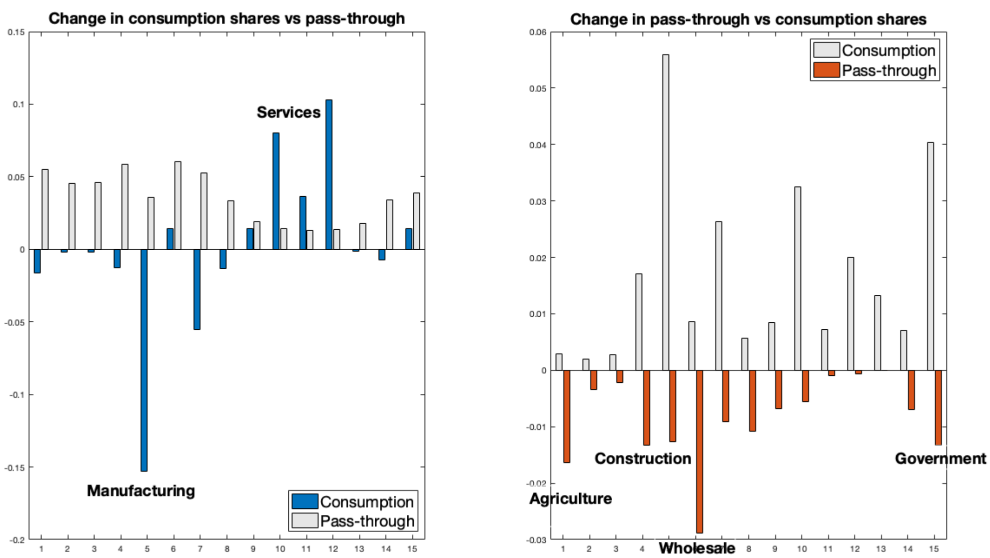

Figure 10: Upper panel: change in consumption shares and average wage pass-through.
Lower panel: change in pass-through and average consumption shares.

The grey bars in the two plots respectively represent average pass-through and average consumption shares. The bars in color represent changes in consumption shares and wage pass-through into sectoral prices. From the upper plot we see that consumption shifted away from manufacturing (where wages have high pass-through) towards services (where they have lower pass-through). The lower plot shows that increasing intermediate input shares reduced the pass-through in all sectors, and more so in sectors with high consumption shares (such as construction, manufacturing and government). Both channels reduced the slope, although quantitatively the change in input shares accounts for most of the effect.

I find similar results in the dynamic setting. Figure 11 plots the impact response of inflation to a 1% shock in real rates between 1947 and 2017. Inflation has become less responsive to monetary policy over time, and again the increase in intermediate input shares explains most of the effect.

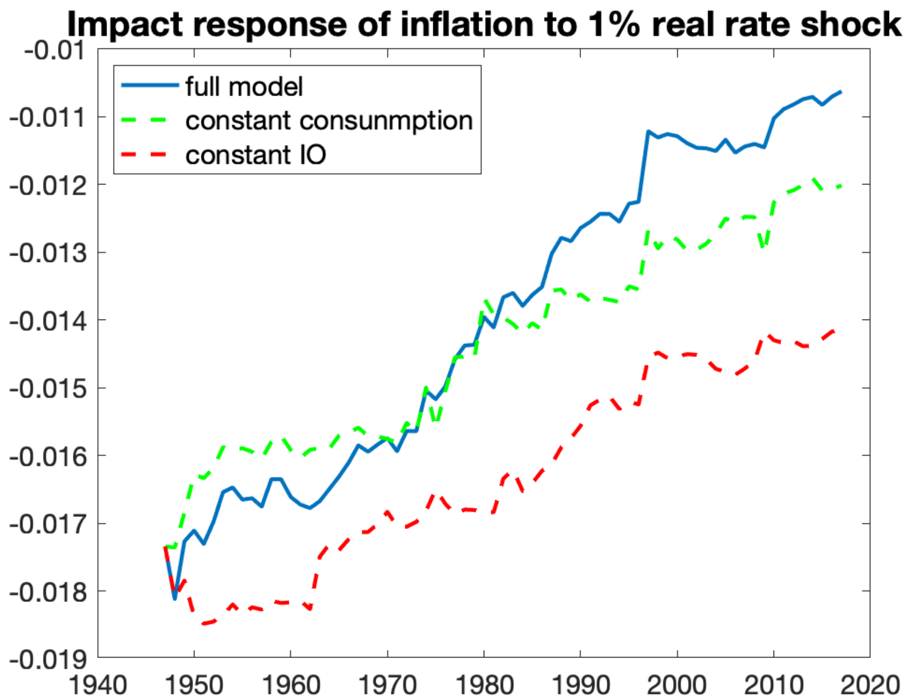

Figure 11: Impact response of consumer inflation and output to a 1% real rate shock

6.3.4 Wage Phillips curve vs consumer price Phillips curve

The calibrated model predicts a flat consumer price Phillips curve, and a substantial decline of its slope over time. By contrast, the wage Phillips curve is much steeper. The calibrated slope was 0.78 in 1947, and it has remained almost constant over time (it declined to 0.77 in 2017). These results from the quantitative model are consistent with empirical evidence (see for example Hooper, Mishkin, Sufi (2019)), which points at a flattening of the consumer price Phillips curve, while the wage Phillips curve is stronger and has not changed over time.

6.4 Endogenous cost-push shocks

6.4.1 Time series

I construct a time series for the endogenous cost-push term in the consumer-price Phillips curve using sector-level measures of yearly TFP shocks from the BEA KLEMS data. Shocks are then aggregated into the Phillips curve residual based on Corollary 3 in Section 4.1.3. Figure 12 plots the results. The residual has a mean of -0.16 and a standard deviation of 0.25, both of which are large relative to the calibrated slope (that is 0.09). I also find that the residual tracks oil prices quite closely, as shown in the figure.

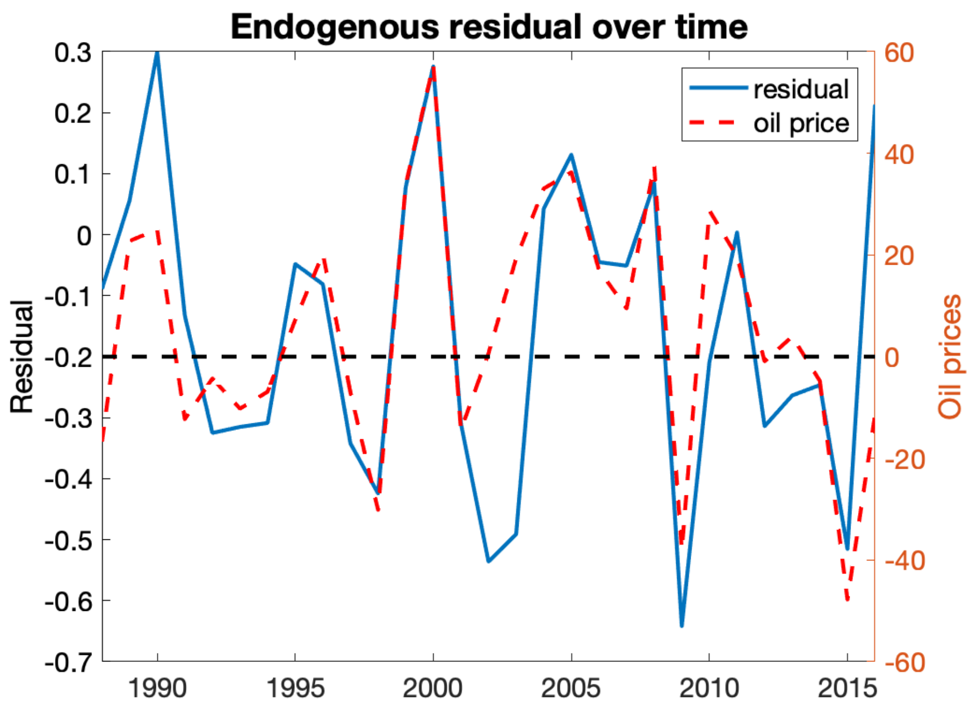

Figure 12: Time series of the endogenous residual and oil prices

6.4.2 Oil shocks

Consistent with the common wisdom in the policy practice, Figure 12 above and Examples 3 and 6 in Sections 4.2 and 5.2 point at the importance of oil shocks in explaining inflation dynamics.

Table 2 reports the calibrated response of inflation to an oil shock in the US network, under different assumptions about price and wage rigidity.

	$\delta =$ 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 productivity shock to the oil sector, full model

Even if the network structure is more complex than in Example 3, we see that the simple model from the example captures well the mechanisms at play. Both wage rigidities and the correlation between oil shares and adjustment frequencies contribute to the effect of oil shocks on consumer inflation. Even though oil is an "upstream" input, if wages are rigid and oil prices are flexible an oil shock acts like a downstream shock. Wages do not adjust fast enough to compensate an increase in oil prices, and inflation is positive under zero output gap. This channel is reinforced by a positive correlation between oil shares and price flexibility: prices respond strongly in the sectors that are more exposed to the shock, further increasing consumer inflation. If instead we assumed wages to be flexible ($\delta_L = 1$) and price adjustment frequencies to be uniform ($\delta_{oil} = \delta_i \equiv \delta_{mean}\ \forall i$), under zero output gap consumer prices would slightly decline.

To complement the discussion in Example 6, Table 3 presents the optimal monetary policy response to a 10% negative oil shock. I express optimal policy in terms of the output gap (in percentage points). In the main calibration I find that the central bank should implement a positive output gap. The implied percentage change in output is obtained by adding the log change in natural output, $y_{nat} = -0.69$.

	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

7 Phillips curve regressions

In this section I run Phillips curve regressions with the "divine coincidence" inflation SW implied by the model and with various measures of consumer prices. I construct a time series of SW for the US economy based on PPI data from the BLS. In the main text I focus on a specification with no lags and a proxy for inflation expectations (which is consistent with the model), whereas results for other specifications (including lags and/or excluding expectations) are reported in Appendix E.

7.1 Data

I use sector-level PPI data from the BLS, and aggregate it based on BEA input-output weights and price adjustment frequencies from Pasten, Schoenle and Weber (2018). Appendix $E1$ describes more in detail how I construct sector-level price series. I measure inflation as the percentage change from the same quarter of the previous year.

Figure 13 plots SW against two measures of consumer price inflation (CPI and PCE) and against aggregate producer price inflation (PPI), from 1984 to 2018. I refer the reader to Appendix $E1$ for a more detailed comparison between the weighting of sectoral inflation rates in the PCE (or core PCE) and in SW. In the appendix I also provide a breakdown of the weights attributed to the top-15 sectors in SW. A kew difference between the two is that the PCE puts no weight on wage inflation, while instead it has an 18% weight in SW. Other important sectors in SW are professional services, financial intermediation and durable goods, whereas the PCE places high weight on health care, real estate and nondurable goods.

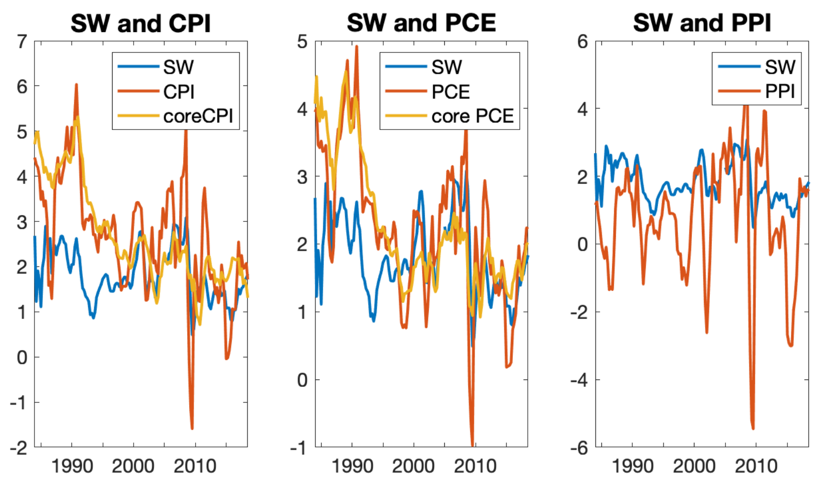

Figure 13: Comparison of SW against consumer and producer prices (1965-2018)

I construct a proxy for inflation expectations by exploiting the statistical properties of the inflation process, whose changes are well approximated by an IMA(1,1) (as discussed in 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. The forecast series are plotted in Appendix $E1$. 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.

Appendix $E1$ also reports scatterplots of the output gap against SW and consumer inflation.

7.2 Regressions over the full sample period

I run Phillips curve regressions over the period January 1984 - July 2018 for different inflation measures (SW, CPI, PCE, core CPI, core PCE). I take the CBO unemployment gap as a measure of the output gap. Appendix $E2$ reports results for different measures.

Table 5 reports results for a plain regression with just inflation and the output gap:

$$\pi_t = c + \kappa y_t + u_t \tag{48}$$

Table 6 reports results for the preferred specification with inflation expectations:

$$\pi_t = c + \rho \mathbb{E}_t \pi_{t+1} + \kappa y_t + u_t \tag{49}$$

Two results are worth noting. First, the calibrated model predicts well the estimated slope, for both consumer prices and the "divine coincidence" index SW (see Table 4). Second, the R-squared is much higher for SW. This is also consistent with the model. According to the theory, the consumer-price Phillips curve has large and volatile endogenous residuals, whereas SW yields a Phillips curve with no endogenous residuals. Correspondingly, regressions with SW should indeed have higher R-squared. Results for other specifications (with lags and/or other measures of the output gap, and for gap levels versus inflation changes) and residual plots are reported in Appendix $E2$.

	SW	consumer prices
κ	-3.00	-0.09

Table 4: Calibrated slope of the Phillips curve ($\gamma = 1$, $\varphi = 2$)

	SW	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 5: CBO unemployment gap

	SW	CPI	core CPI	PCE	core PCE
gap	-1.1054**	-0.1613**	-0.0344	-0.062	0.0047
	(0.3275)	(0.0809)	(0.052)	(0.0487)	(0.0368)
inflation expecations	0.8287**	0.4846**	0.5446**	0.6364**	0.6406**
	(0.0383)	(0.1557)	(0.0559)	(0.0621)	(0.045)
intercept	0.3484**	1.3851**	1.3193**	0.5522**	0.8388**
	(0.0789)	(0.5021)	(0.1818)	(0.196)	(0.1228)
R-squared	0.8234	0.159	0.4425	0.4635	0.6072

Table 6: CBO unemployment gap

7.3 Rolling regressions

I run rolling Phillips curve regressions (with a 20 year window) over the period January 1984 - July 2018 for different inflation measures (SW, CPI, PCE, core CPI, core PCE). Here I report results for the preferred specification (49) with inflation expectations. I take the CBO unemployment gap as a measure of the output gap. Appendix $E3$ reports results for different measures of the output gap and for other specifications.

Figure 14 compares the strength and stability of the estimated relation for different choices of the dependent variable (the "divine coincidence" index SW versus consumer prices). 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 gives a measure of the stability of the estimated coefficient, proxied by its standard deviation relative to the mean. I find that SW dominates consumer prices along all these dimensions: the R-squareds are consistently higher, the estimated coefficient is always significant and the variance of the estimate is lower. Plots of the estimated coefficients and confidence intervals are reported in Appendix $E3$.

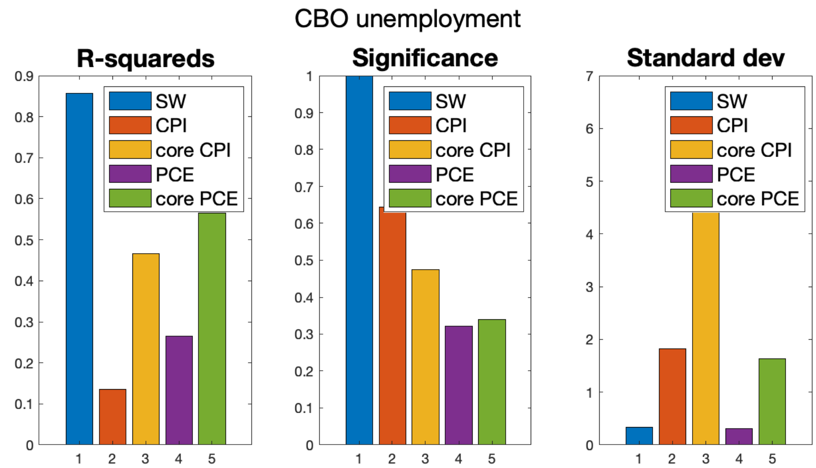

Figure 14: Summary statistics for rolling Phillips curve regressions (20y window, 1984-2018, unemployment gap as dependent variable)

8 Conclusion

This paper revisits the "positive" and "normative" implications of the New Keynesian model under a realistic production structure, with multiple sectors, input-output linkages and heterogeneous shocks and pricing frictions. I show that the consumer-price Phillips curve is misspecified: the slope changes with the input-output structure and productivity shocks create an endogenous residual. I propose a correct specification based on the model, which aggregates sectoral inflation rates according to sales shares and appropriately discounts flexible sectors. I compare Phillips curve regressions with consumer prices and with my indicator, finding that the estimated coefficients are consistent with the calibrated model and the R-squared is significantly higher with my preferred indicator.

Monetary policy cannot replicate the first best, but the optimal policy can still be implemented by targeting an appropriately defined inflation indicator. I evaluate the performance of the two standard targets in the Taylor rule, output gap and consumer inflation, finding that targeting the output gap is close to optimal, while stabilizing consumer prices generates an expected loss of 0.8% of per-period GDP with respect to the optimal policy.

References

Acemoglu, D., U. Akcigit, and W. Kerr (2016). Networks and the macroeconomy: An empirical exploration. NBER Macroannual 30(1), 273–335.

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The network origins of aggregate fluctuations. Econometrica 80(5), 1977–2016.

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2017). Microeconomic origins of macroeconomic tail risks. The American Economic Review 107(1), 54–108.

Alvarez, F., H. Le Bihan, and F. Lippi (2016). The real effects of monetary shocks in sticky price models: A sufficient statistic approach. The American Economic Review 106 (10), 2817–2851.

Aruoba, S.B., F.X. Diebold and C. Scotti (2012). Real-Time Measurement of Business Conditions. Journal of Business and Economic Statistics, 417-427.

Atalay, E. (2017). How important are sectoral shocks? American Economic Journal: Macroeconomics (Forthcoming).

Baqaee, D. R. (2016). Cascading failures in production networks.

Baqaee, D. R. and E. Farhi (2017). Productivity and Misallocation in General equilibrium.

Baqaee, D. R. and E. Farhi (2017). The macroeconomic impact of microeconomic shocks: Beyond Hulten's Theorem.

Barrot, J.-N. and J. Sauvagnat (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. The Quarterly Journal of Economics 131(3), 1543–1592.

Bartelme, D. and Y. Gorodnichenko (2015). Linkages and economic development. Technical report, National Bureau of Economic Research.

Basu, S. (1995). Intermediate goods and business cycles: Implications for productivity and welfare. The American Economic Review, 512–531.

Benigno, P. (2004). Optimal monetary policy in a currency area. Journal of International Economics, 63 (2), 293-320.

Bernanke, B. (2008), Outstanding Issues in the Analysis of Inflation. Speech at the Federal Reserve Bank of Boston's 53rd Annual Economic Conference. https://www.federalreserve.gov/newsevents/speech

Bernanke, B and J. Boivin (2003). Monetary Policy In A Data-Rich Environment. Journal of Monetary Economics, 2003, v50(3,Apr), 525-546.

Bigio, S. and J. La'O (2016). Financial frictions in production networks. Technical report.

Bils, M. and P. J. Klenow (2004). Some evidence on the importance of sticky prices. Journal of Political Economy 112(5), 947–985.

Blanchard, O. and J. Galí (2007). Real Wage Rigidities and the New Keynesian Model. Journal of Money, Credit and Banking, Blackwell Publishing, vol. 39(s1), 35-65, 02

Blanchard, O. (2016). The US Phillips Curve: Back to the 60s?. Policy brief, Peterson Institute for International Economics, PB16-1

Boehm, C., A. Flaaen, and N. Pandalai-Nayar (2014). Complementarities in multinational production and business cycle dynamics. Technical report, Working paper, University of Michigan.

Boivin, J., M. P. Giannoni, and I. Mihov (2009). Sticky prices and monetary policy: Evidence from disaggregated U.S. data. The American Economic Review 99(1), 350–384.

Bouakez, H., E. Cardia, and F. J. Ruge-Murcia (2009). The transmission of monetary policy in a multisector economy. International Economic Review 50(4), 1243–1266.

Bouakez, H., E. Cardia, and F. Ruge-Murcia (2014). Sectoral price rigidity and aggregate dynamics. European Economic Review 65, 1–22.

Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2014). The impact of regional and sectoral productivity changes on the U.S. economy. Technical report, National Bureau of Economic Research.

Carvalho, V. M. (2014). From micro to macro via production networks. The Journal of Economic Perspectives 28(4), 23–47.

Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi (2016). Supply chain disruptions: Evidence from the great east Japan earthquake. Unpublished Manuscript.

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

Di Giovanni, J., A. A. Levchenko, and I.Méjean (2014). Firms, destinations, and aggregate fluctuations. Econometrica 82(4), 1303–1340.

Di Giovanni, J., A. A. Levchenko, and I. Méjean (2016). The micro origins of international business cycle comovement. Technical report, National Bureau of Economic Research.

Dixit, A. K. and J. E. Stiglitz (1977). Monopolistic competition and optimum product diversity. The American Economic Review, 297–308.

Eusepi, S., B. Hobjin, A. Tambalotti (2011). CONDI: A Cost-Of-Nominal-Distortions Index. Americal Economic Journal: Macroeconomics, Vol.3, No. 3, 53-91

Foerster, A., P. Sarte, and M. Watson (2011). Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production. Journal of Political Economy, 119:1, 1-38

Domar, E. D. (1961). On the measurement of technological change. The Economic Journal 71(284), 709–729.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. Econometrica 79 (3), 733–772.

Galì, J. and M. Gertler (1999). Inflation Dynamics: A Structural Econometric Analysis. Journal of Monetary Economics, 44 (2),195-222.

Galì, J. (2008). Monetary policy, inflation, and the business cycle: An introduction to the New Keynesian Framework.

Galì, J. and T. Monacelli (2008). Optimal Monetary and Fiscal Policy in a Currency Union. Journal of International Economics 76, 1. 116-132.

Goldberg, P. P. and R. Hellerstein (2011). How rigid are producer prices? FRB of New York Staff Report, 1–55.

Gopinath, G. and O. Itskhoki (2010). Frequency of price adjustment and pass-through. The Quarterly Journal of Economics 125(2), 675–727.

Gopinath, G. and O. Itskhoki (2011). In search of real rigidities. NBER Macroeconomics Annual 25(1), 261–310.

Gopinath, G., Ş. Kalemli-Özcan, L.Karabarbounis, and C.Villegas-Sanchez (2017). Capital allocation and productivity in South Europe. The Quarterly Journal of Economics, qjx024.

Grassi, B. (2017). Io in io: Competition and volatility in input-output networks. Unpublished Manuscript, Bocconi University.

Horvath, M. (1998). Cyclicality and sectoral linkages: Aggregate fluctuations from independent sectoral shocks. Review of Economic Dynamics 1(4), 781–808.

Horvath, M. (2000). Sectoral shocks and aggregate fluctuations. Journal of Monetary Economics 45(1), 69–106.

Hulten, C. R. (1978). Growth accounting with intermediate inputs. The Review of Economic Studies 45(3), 511–518.

Kelly, B., H. Lustig, and S. Van Nieuwerburgh (2013). Firm volatility in granular networks. Technical report, National Bureau of Economic Research.

Klenow, P. J. and O. Kryvtsov (2008). State-dependent or time-dependent pricing: does it matter for recent us inflation? The Quarterly Journal of Economics 123(3), 863–904.

La'O, J., and A. Tahbaz-Salehi (2019). Optimal monetary policy with production networks. Working paper.

Long, J. B. and C. Plosser (1983). Real business cycles. The Journal of Political Economy 91 (1), 39–69.

Lucas, R. E. (1977). Understanding business cycles. In Carnegie-Rochester Conference Series on Public Policy, Volume 5, pp. 7–29. Elsevier.

Mavroeidis S,. M. Plagborg-Moller M and J.H. Stock (2014). Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve. Journal of Economics Literature 52 (1) :124-188.

Mankiw, N.G and R. Reis (2003). What Measure of Inflation Should a Central Bank Target?, Journal of the European Economic Association, MIT Press, vol. 1(5), pages 1058-1086, 09

Midrigan, V. (2011). Menu costs, multiproduct firms, and aggregate fluctuations. Econometrica 79(4), 1139–1180.

Montoro, C. (2012). Oil shocks and optimal monetary policy. Macroeconomic Dynamics, vol. 16, issue 02, 240-277.

Nakamura, E. and J. Steinsson (2008). Five facts about prices: A reevaluation of menu cost models. Quarterly Journal of Economics 123(4), 1415–1464.

Nakamura, E. and J. Steinsson (2010). Monetary non-neutrality in a multisector menu cost model. Quarterly Journal of Economics 125(3), 961–1013.

Nakamura, E. and J. Steinsson (2013). Price Rigidity: Microeconomic Evidence and Macroeconomic Implications. Annual Review of Economics, 2013. 5:133–63

Orphanides A. and S. van Norden (2002). The Unreliability of Output-Gap Estimates in Real Time. The Review of Economics and Statistics Vol. 84, No. 4, 569-583.

Ozdagli, A and M. Weber (2017). Monetary Policy Through Production Networks: Evidence from the Stock Market. NBER Working Paper No. w23424

Pasten, E., R. Schoenle, and M.Weber (2016). The propagation of monetary policy shocks in a heterogeneous production economy.

Pasten, E., R. Schoenle, and M. Weber (2017). Price rigidity and the granular origin of aggregate fluctuations.

Stock, J. and M. Watson (1999), Forecasting Inflation. Journal of Monetary Economics, v44(2,Oct), 293-335.

Stock, J. and M. Watson (2016), Core Inflation and Trend Inflation. Review of Economics and Statistics, vol 98(4), pages 770-784.

Appendix A: Positive analysis

A1: Natural output and output gap

Proof of Lemma 1**:**

Given that the flex-price equilibrium is efficient, it solves the problem

$$max_{C,L,\{L_i,y_i,\{x_{ij}\}\}} \frac{C\left(\{y_i\}\right)^{1-\gamma}}{1-\gamma} - \frac{L^{1+\varphi}}{1+\varphi} \quad s.t. \quad \begin{array}{c} y_i + \sum_j x_{ij} = A_i F_i\left(\{x_{ij}\}, L_i\right) \ \forall i \\ \sum_i L_i = L \end{array} \tag{50}$$

Our objective is to compute $\frac{dlogC^*}{dlogA_i}$.

Observe that the optimization problem in (50) can be solved in two steps: first choose $\{L_i, y_i, \{x_{ij}\}\}$ for given L, and then choose the optimal L.

Formally, solving problem (50) is equivalent to solving

$$\frac{C^*\left(L;A\right)^{1-\gamma}}{1-\gamma} = max_{\{L_i,y_i,\{x_{ij}\}\}} \frac{C\left(\{y_i\}\right)^{1-\gamma}}{1-\gamma} \quad s.t. \quad \begin{array}{c} y_i + \sum_j x_{ij} = A_i F_i\left(\{x_{ij}\}, L_i\right) \ \forall i \\ \sum_i L_i = L \end{array} \tag{51}$$

and

$$max_L \frac{C^*(L;A)^{1-\gamma}}{1-\gamma} - \frac{L^{1+\varphi}}{1+\varphi} \tag{52}$$

The solution of (52) must satisfy

$$C^*(L;A)^{\gamma} L^{\varphi} = \frac{\partial C^*}{\partial L}$$

Using the envelope theorem in problem (51) we have that

$$\frac{\partial C^*}{\partial L} = C^{*\gamma} \nu_L(A)$$

where ν_L is the Lagrange multiplier associated to the constraint $\sum_i L_i = L$. From the equation

$$L^{\varphi} = \nu_L(A)$$

we have

$$\frac{dlogL}{dlogA_i} = \frac{1}{\varphi} \frac{dlog\nu_L}{dlogA_i}$$

Finally, applying again the envelope theorem to problem (51) we have

$$\frac{dlogC^*}{dlogA_i} = C^{*\gamma}\left(\frac{\nu_L L}{\varphi C^*}\frac{dlog\nu_L}{dlogA_i} + \frac{\nu_i F_i\left(\{x_{ij}\}, L_i\right)}{C^*}\right)$$

We now need to prove that $C^{*\gamma}\frac{\nu_i F_i(\{x_{ij}\},L_i)}{C^*} = \lambda_i$ (where λ_i is the share of i's sales in GDP), and that

$$C^{*\gamma}\frac{\nu_L L}{\varphi C^*}\frac{dlog\nu_L}{dlogA_i} = \frac{1}{\varphi}\lambda_i - \frac{\gamma}{\varphi}\frac{dlogC^*}{dlogA_i} \tag{53}$$

This in turn implies that

$$\frac{dlogC^*}{dlogA_i} = \frac{1+\varphi}{\gamma+\varphi}\lambda_i$$

Let's first prove that $C^{*\gamma}\nu_i$ is equal to the price of good i relative to the CPI in the competitive equilibrium, so that $C^{*\gamma}\frac{\nu_i F_i(\{x_{ij}\},L_i)}{C^*} = \lambda_i$.

From the FOCs of (51) we have $C_i = C^{\gamma}\nu_i$ and from consumer optimization in the competitive equilibrium we have $\frac{C_j}{C_i} = \frac{p_j}{p_i}$. Thus

$$\frac{C_j}{C_i} = \frac{\nu_j}{\nu_i} = \frac{p_j}{p_i}$$

Using the fact that C is homogeneous of degree one, and normalizing $\sum_j \frac{p_j y_j}{C} = 1$, we have

$$\frac{\sum C_j y_j}{C_i} = \frac{C}{C_i} = \frac{C}{p_i} \Rightarrow p_i = C_i$$

The FOCs for (51) in turn imply that $p_i = C^{\gamma}\nu_i$.

Let's now derive equation (53). From the FOCs of (51) it holds that $C^{\gamma}\nu_L = C^{\gamma}\nu_i A_i F_{iL} = p_i A_i F_{iL} = w\ \forall i$, where the last equality follows from firm optimization in the competitive equilibrium. Moreover, from the consumers' budget constraint we have that $w = \frac{C^*}{L}$. Thus

$$C^{*\gamma}\frac{\nu_L L}{\varphi C^*}\frac{dlog\nu_L}{dlogA_i} = \frac{1}{\varphi}\left(\frac{dlogw}{dlogA_i} - \gamma\frac{dlogC^*}{dlogA_i}\right)$$

To conclude the proof we need to show that

$$\frac{dlogw}{dlogA_i} = \lambda_i$$

Using again the consumers' budget constraint we have

$$\frac{dlogw}{dlogA_i} = \frac{\partial logC^*}{\partial logA_i} + \left(\frac{\partial logC^*}{\partial logL} - 1\right)\frac{dlogL}{dlogA_i} = \lambda_i$$

Proof of Lemma 2 **:**

As explained in the proof of Lemma 1, in the undistorted steady-state productivity is maximized by optimally allocating labor both within and across sectors. Therefore around this steady-state the productivity loss induced by a change in firm-level markups is zero to a first order.

A2: Sector-level inflation

Proof of Propositions 2 **and** 1

We first solve for the change in marginal costs as a function of the change in prices, wages and productivity:

$$dlogmc_i = \widetilde{\alpha}_i dlogw + \sum_j \widetilde{\omega}_{ij} dlogp_j - dlogA_i$$

We can write the change in sectoral prices as function of the change in marginal costs using the Calvo assumption:

$$dlogmc_i = \widetilde{\alpha}_i dlogw - dlogA_i + \sum_j \widetilde{\omega}_{ij}\delta_j dlogmc_j$$

and solve for the change in marginal cost as a function of the change in wages and productivity:

$$dlogmc = \left(I - \widetilde{\Omega}\Delta\right)^{-1}(\widetilde{\alpha} dlogw - dlogA) \tag{54}$$

The change in consumer prices is

$$dlogP = \beta^T dlogp = \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}(\widetilde{\alpha} dlogw - dlogA) \tag{55}$$

From the consumption-leisure trade-off, we have

$$dlogw = dlogP + (\varphi dlogL + \gamma dlogy) =$$

$$= \left(\varphi dlogL + \gamma \tilde{y} + \gamma y^{nat} + dlogP\right) =$$

$$= \left((\gamma + \varphi)\,\widetilde{y} + \lambda^T dlogA + dlogP\right)$$

We can then use (55) to solve for the change in wages as a function of the output gap and productivity shocks. We have:

$$dlogw = \frac{(\gamma+\varphi)\,\widetilde{y} + \beta^T \left[\left(I-\widetilde{\Omega}\right)^{-1} - \Delta\left(I-\widetilde{\Omega}\Delta\right)^{-1}\right] dlogA}{1-\beta^T\Delta\left(I-\widetilde{\Omega}\Delta\right)^{-1}\widetilde{\alpha}} \tag{56}$$

Lemma 6 shows that the denominator in (56) is always well defined.

Lemma 6. $1-\beta^T\Delta\left(I-\widetilde{\Omega}\Delta\right)^{-1}\widetilde{\alpha} > 0.$

Proof. First note that, by definition of labor and input shares, it holds that $\tilde{\alpha} = (I-\Omega)\,U$, where U is a Nx1 vector with all entries equal to 1. Thus we have that

$$\beta^T\left(I-\widetilde{\Omega}\right)^{-1}\widetilde{\alpha} = \beta^T\left(I-\widetilde{\Omega}\right)^{-1}(I-\Omega)\,U =$$

$$\beta^T U = \sum_j \beta_j = 1$$

To prove Lemma 6 it is enough to show that

$$\beta^T\Delta\left(I-\widetilde{\Omega}\Delta\right)^{-1}\widetilde{\alpha} < \beta^T\left(I-\widetilde{\Omega}\right)^{-1}\widetilde{\alpha}$$

A sufficient condition for this to hold is that

$$\Delta\left(I-\widetilde{\Omega}\Delta\right)^{-1}_{ij} < (I-\Omega)^{-1}_{ij} \;\; \forall i,j$$

Note that

$$\Delta\left(I-\widetilde{\Omega}\Delta\right)^{-1}_{ij} = \delta_i\left(I-\widetilde{\Omega}\Delta\right)^{-1}_{ij} < \left(I-\widetilde{\Omega}\Delta\right)^{-1}_{ij}$$

therefore it is sufficient to prove that

$$\left(I-\widetilde{\Omega}\Delta\right)^{-1}_{ij} < (I-\Omega)^{-1}_{ij} \;\; \forall i,j$$

We can do so using the relations

$$\left(I-\widetilde{\Omega}\Delta\right)^{-1} = I + \widetilde{\Omega}\Delta + \left(\widetilde{\Omega}\Delta\right)^2 + ...$$

$$\left(I-\widetilde{\Omega}\right)^{-1} = I + \widetilde{\Omega} + \widetilde{\Omega}^2 + ...$$

This yields

$$\left(I-\widetilde{\Omega}\Delta\right)^{-1}_{ij}=\mathbb{I}(i=j)+\omega_{ij}\delta_j+\sum_k \omega_{ik}\omega_{kj}\delta_j\delta_k+...<$$

$$\mathbb{I}(i=j)+\omega_{ij}+\sum_k \omega_{ik}\omega_{kj}+...=\left(I-\widetilde{\Omega}\right)^{-1}_{ij}$$

$\square$

To find marginal costs as function of the output gap and productivity shocks, plug (56) into (54):

$$dlogmc=\frac{(\gamma+\varphi)\left(I-\widetilde{\Omega}\Delta\right)^{-1}\widetilde{\alpha}}{1-\beta^T\Delta\left(I-\widetilde{\Omega}\Delta\right)^{-1}\widetilde{\alpha}}\widetilde{y}+$$

$$\left(I-\widetilde{\Omega}\Delta\right)^{-1}\left(\frac{\widetilde{\alpha}\left[\lambda^T-\beta^T\Delta\left(I-\widetilde{\Omega}\Delta\right)^{-1}\right]}{1-\beta^T\Delta\left(I-\widetilde{\Omega}\Delta\right)^{-1}\widetilde{\alpha}}-I\right)dlogA \tag{57}$$

The expressions for the elasticities $\mathcal{B}$ and $\mathcal{V}$ in Section 3.2 follow immediately from (57).

Proof of Lemma 2**:**

In our setup labor is the only factor of production. Then labor and input shares must sum to one:

$$\alpha+\Omega\mathbf{1}=\mathbf{1}$$

therefore it holds that $(I-\Omega)^{-1}\alpha=\mathbf{1}$. The result

$$\exists i,j \text{ such that } \omega_{ij}\delta_j<\omega_{ij}\implies (I-\Omega\Delta)^{-1}\alpha<\mathbf{1}$$

follows immediately from the fact that each term in the geometric sum

$$(I-\Omega\Delta)^{-1}\alpha=\left(I+\Omega\Delta+(\Omega\Delta)^2+...\right)\alpha$$

has at least one component that is smaller than in the corresponding term of

$$(I-\Omega)^{-1}\alpha=\left(I+\Omega+\Omega^2+...\right)\alpha$$

It then follows that

$$\bar{\delta}_w=\sum_i \beta_i\delta_i\left[(I-\Omega\Delta)^{-1}\alpha\right]_i<\sum_i\beta_i\delta_i\equiv\mathbb{E}_\beta(\delta)$$

Equation 16 is obtained by differentiating (12).

Proof of Corollary 4**:**

We first show that $\mathcal{V}\alpha = 0$, that is, α belongs to $ker(\mathcal{V})$.

Recall the expression for $\mathcal{V}$:

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

Thus we have

$$\mathcal{V}\alpha = \frac{\left(I - \widetilde{\Omega}\Delta\right)^{-1} \widetilde{\alpha} \left[1 - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1} \alpha\right]}{1 - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1} \widetilde{\alpha}} - \left(I - \widetilde{\Omega}\Delta\right)^{-1} \alpha = 0$$

We then prove that α is the only element of $ker(\mathcal{V})$. Note that for every vector $x \neq 0$ such that $\mathcal{V}x = 0$ it must hold that

$$\left(I - \widetilde{\Omega}\Delta\right)^{-1} \widetilde{\alpha} \frac{\left[\lambda^T - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}\right] x}{1 - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1} \widetilde{\alpha}} = \left(I - \widetilde{\Omega}\Delta\right)^{-1} x \iff$$

$$\widetilde{\alpha}_j \frac{\left[\lambda^T - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}\right] x}{1 - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1} \widetilde{\alpha}} = x_i \ \forall i \tag{58}$$

and

$$\mathcal{C} \equiv \frac{\left[\lambda^T - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}\right] x}{1 - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1} \widetilde{\alpha}} \neq 0$$

(otherwise we would have $x = 0$). From (58) we then have that

$$\frac{\alpha_i}{\alpha_j} = \frac{x_i}{x_j} \ \forall i, j$$

so that x is proportional to the vector of labor shares α.

A3: Output gap and aggregate inflation

Proof of Lemma 3:

The output gap is a labor demand shifter. Therefore it impacts real wages, according to the parameters of the labor supply curve. From the consumers' optimal labor supply decision we have:

$$\left(logw - logw^{nat}\right) - \left(logP - logP^{nat}\right) = \gamma\tilde{y} + \varphi\left(logL - logL^{nat}\right) = (\gamma + \varphi)\,\tilde{y} \tag{59}$$

where the last equality follows from Lemma 2.

We next need to compute the left hand side of (59), which corresponds to the change in real wages induced by markup distortions.

Let's first introduce two definitions.

Definition 3. The cost-based input-output matrix $\widetilde{\Omega}$ is an $N \times N$ matrix with element i, j given by the expenditure share on input j in i's cost:

$$\widetilde{\omega}_{ij} = \frac{p_j x_{ij}}{mc_i y_i}$$

Definition 4. The sector-level steady-state labor shares in marginal costs are encoded in the $N \times 1$ vector $\tilde{\alpha}$ with components

$$\tilde{\alpha}_i = \frac{wL_i}{mc_i y_i}$$

In a steady-state with optimal subsidies it holds that $\Omega = \widetilde{\Omega}$ and $\alpha = \tilde{\alpha}$.

To solve for real wages as a function of sector-level markups we first need to consider how they impact marginal costs and prices. We have:

$$dlogmc_i = \widetilde{\alpha}_i dlogw + \sum_j \widetilde{\omega}_{ij} dlogp_j - dlogA_i$$

and

$$dlogp_i = dlogmc_i + dlog\mu_i \tag{60}$$

$$\Rightarrow dlogmc = \left(I - \widetilde{\Omega}\right)^{-1}\left(\widetilde{\alpha} dlogw - dlogA + \widetilde{\Omega} dlog\mu\right) \tag{61}$$

$$\Rightarrow dlogP = \beta^T (dlogmc + dlog\mu) = dlogw + \tilde{\lambda}^T\,(dlog\mu - dlogA)$$

It follows that

$$dlogw - dlogP = \tilde{\lambda}^T\,(dlogA - dlog\mu) \tag{62}$$

For constant productivity we then have

$$\left(logw - logw^{nat}\right) - \left(logP - logP^{nat}\right) = -\lambda^T dlog\mu \tag{63}$$

Equations (59) and (63) together give the result.

Proof of Lemma 4:

We need to prove that all the vectors $x \neq 0$ satisfying $x^T \mathcal{V} = \mathbf{0}$ are proportional to $(I - \Delta)\,\Delta^{-1}\lambda$. Proposition 3 implies that $\lambda^T (I - \Delta)\,\Delta^{-1}\mathcal{V} = \mathbf{0}$.

Consider then all vectors x such that $x^T \mathcal{V} = \mathbf{0}$. Note that

$$x^T \mathcal{V} = \mathbf{0} \Longleftrightarrow x^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1} \left[\widetilde{\alpha}\left[\lambda^T - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}\right] - \left(1 - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1} \widetilde{\alpha}\right) I\right] = \mathbf{0}$$

$$\Longleftrightarrow \tilde{x}^T \left[\widetilde{\alpha}\left[\lambda^T - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}\right] - \left(1 - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1} \widetilde{\alpha}\right) I\right] = \mathbf{0} \tag{64}$$

where $\tilde{x}^T \equiv x^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}$.

To prove the Lemma we need to show that all vectors $\tilde{x}$ satisfying (64) are proportional to $\lambda^T (I - \Delta)\left(I - \widetilde{\Omega}\Delta\right)^{-1}$.

From (64) we have the relation

$$\left(1 - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1} \widetilde{\alpha}\right) \tilde{x}_j = \tilde{x}^T \widetilde{\alpha} \left[\lambda^T - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}\right]_j \quad \forall j \tag{65}$$

from which we derive the condition

$$\frac{\tilde{x}_i}{\tilde{x}_j} = \frac{\left[\lambda^T - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}\right]_i}{\left[\lambda^T - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}\right]_j}$$

Note that we must have $\tilde{x}^T \tilde{\alpha} \neq 0$, otherwise we would get $\tilde{x}^T = \mathbf{0}$, while we want $\tilde{x} \neq 0$. Moreover, the ratio on the RHS is well defined, because $\left[\lambda^T - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}\right]_j > 0 \; \forall j$ (see Lemma 6).

Thus, $\tilde{x}^T$ must be proportional to the vector

$$\lambda^T - \beta^T \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1} = \beta^T \left[(I - \Omega)^{-1} - \Delta \left(I - \widetilde{\Omega}\Delta\right)^{-1}\right] =$$

$$= \beta^T \left[(I - \Omega)^{-1} (I - \Omega\Delta) - \Delta\right] (I - \Omega\Delta)^{-1} =$$

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

Appendix B: Optimal policy

B1: Welfare function

Proof of Lemma 5**:**

From equation (61), for each sector i we have

$$dlogp_i - dlogw = (I-\Omega)^{-1} dlog\mu$$

We can then use the pricing equation (2) to substitute for markups as a function of inflation rates.

Proof of Proposition 4**:**

In what follows, I will use the second-order approximation

$$\frac{Z-Z^*}{Z} \simeq log\left(\frac{Z}{Z^*}\right) + \frac{1}{2} log\left(\frac{Z}{Z^*}\right)^2$$

I will prove below that we can write the second-order log change in output with respect to the efficient outcome as

$$log\left(\frac{Y}{Y^*}\right) \equiv \hat{y} = \hat{l} - d$$

Using this result we can approximate the utility function around the efficient outcome as

$$\frac{U-U^*}{U_cC} \simeq \hat{y} + \frac{1}{2}\hat{y}^2 + \frac{1}{2}\frac{U_{cc}C}{U_c}\hat{y}^2 + \frac{U_lL}{U_cC}\left(\hat{l} + \frac{1}{2}\frac{U_{ll}N}{U_l}\hat{l}^2\right) =$$

$$= \hat{y} + \frac{1-\gamma}{2}\hat{y}^2 - \left(\hat{l} + \frac{1+\varphi}{2}\hat{l}^2\right) =$$

$$= \hat{y} + \frac{1-\gamma}{2}\hat{y}^2 - \left(\hat{y} + d + \frac{1+\varphi}{2}\hat{y}^2\right) =$$

$$= -\frac{\gamma+\varphi}{2}\tilde{y}^2 - d$$

where the last equality follows from the fact that, to the second order, $\hat{y}^2 = \tilde{y}^2$.

I will now derive the approximation

$$\hat{y} = \hat{l} - d$$

and the explicit expression for the second order component d.

For given labor supply and Hicks-neutral productivity shifters there is a productivity loss if firm's markups are not 1, which is the case whenever there are price rigidities and the economy is hit by a productivity shock. To a first-order this loss is zero around the efficient steady-state, so that $\hat{l}$ is the only component. Away from the efficient outcome a change in firm-level markups produces two effects. First, sector-level productivities are lower than in the efficient steady state (i.e. more labor is required to produce one unit of sectoral output), due to relative price distortions across firms within the sector. I will denote by a_i the logarithm of the "effective" productivity of sector i relative to the actual TFP A_i.

Second, there is a change in "sector-level markups", defined as

$$\mu_i = \frac{p_i}{mc_i}$$

where p_i is the sectoral price index (and note that the marginal cost is the same for all producers in sector i). This change in markups generates relative price distortions across sectors.

Both of these channels are associated with an aggregate productivity loss. I will refer to the first channel as the "within-sector" component of this loss, and to the second channel as the "cross-sector" component. I will now derive a first-order approximation of both, and then compute the second order approximation around the efficient steady-state.

Note that aggregate productivity can be expressed as a function of real wages and labor shares. Denoting the aggregate labor share by $\Lambda = \frac{wL}{GDP}$, by definition we can write aggregate output as

$$Y = \frac{1}{\Lambda}\frac{w}{P}L$$

Here $\frac{w}{P}$ is the real wage, and $\frac{1}{\Lambda}\frac{w}{P}$ acts as a measure of aggregate labor productivity. In log deviations from steady-state we have:

$$\hat{Y} = \hat{w} - \hat{P} - \hat{\Lambda} + \hat{l} \tag{66}$$

To compute the change in aggregate productivity we need to know how real wages and Λ respond in equilibrium to sectoral productivities and markups.

Real wages are derived in the proof of Lemma 3 (see equation (63)). Combining (63) with (66) we obtain the first-order approximation

$$dlogY - dlogL = \tilde{\lambda}^T (a - dlog\mu) - dlog\Lambda \tag{67}$$

We then need to compute $dlog\Lambda$ as function of the change in sectoral markups and productivities. It holds that

$$\Lambda + \lambda^T \left(1 - \frac{1}{\mu}\right) \equiv 1$$

Therefore we have

$$dlog\Lambda = -\frac{1}{\Lambda}\left(\sum_i d\lambda_i \left(1 - \frac{1}{\mu_i}\right) + \sum_i \lambda_i \frac{dlog\mu_i}{\mu_i}\right)$$

Using (67) we find that, around the efficient steady state,

$$dlogY - dlogL = \underbrace{\tilde{\lambda}^T a}_{\text{within sector}} + \underbrace{\left(\frac{\lambda^T}{\Lambda} - \tilde{\lambda}^T\right) dlog\mu}_{\text{cross-sector}} \tag{68}$$

As $\frac{\lambda^T}{\Lambda} - \tilde{\lambda}^T = 0$ around $\mu \equiv 1$, the first-order productivity loss from cross-sector misallocation is zero. To compute the second-order loss we need to derive again the cross-sector component in equation (68) .

Note that, since the first order effect on both cross-sector misallocation and sector-level productivities is zero, the second-order terms in $(dlogA)(dlog\mu)$ are also going to be zero. Therefore we only need to derive the cross-sector component with respect to sector-level markups. We have:

$$\left(d^2logY - d^2logL\right) - \tilde{\lambda}^T a = \frac{1}{\Lambda}\left(-\left(\sum_i \lambda_i \frac{dlog\mu_i}{\mu_i}\right)^2 + 2\sum_i d\lambda_i \frac{dlog\mu_i}{\mu_i} + \sum_i \frac{\lambda_i}{\mu_i}(dlog\mu_i)^2\right) - \sum_i d\tilde{\lambda}_i dlog\mu_i =$$

$$= -\frac{1}{2}\sum_i \sum_j \widetilde{d}^2_{ij} dlog\mu_i dlog\mu_j \tag{69}$$

where

$$\widetilde{d}^2_{ij} = \sum_h \sum_k \beta_h \beta_k \sigma_{hk} \left[(I-\Omega)^{-1}_{hi} - (I-\Omega)^{-1}_{ki}\right] \left[(I-\Omega)^{-1}_{hj} - (I-\Omega)^{-1}_{kj}\right] +$$

$$+ \sum_t \lambda_t \sum_h \sum_k \omega_{th}\omega_{tk}\theta^t_{hk} \left[(I-\Omega)^{-1}_{hi} - (I-\Omega)^{-1}_{ki}\right] \left[(I-\Omega)^{-1}_{hj} - (I-\Omega)^{-1}_{kj}\right] +$$

$$+ \sum_t \lambda_t \alpha_t \sum_h \omega_{th}\omega_{tk}\theta^t_{hk} (I-\Omega)^{-1}_{hi} (I-\Omega)^{-1}_{hj} =$$

$$= \Phi_C \left((I-\Omega)^{-1}_{(i)}, (I-\Omega)^{-1}_{(j)} \right) + \sum_t \lambda_t \Phi_t \left((I-\Omega)^{-1}_{(i)}, (I-\Omega)^{-1}_{(j)} \right) \tag{70}$$

To derive the welfare loss as a function of sector-level inflation rates we need to solve for the endogenous change in sector-level markups due to price rigidities.

In this setup (where price adjustment is iid across firms and there are CES preferences "within sectors") firm-markups are the same as in the efficient equilibrium for the producers who can adjust their price. For those who cannot adjust, the change in log markup is the opposite as the change in log marginal cost. Thus the change in sector-level markups is $dlog\mu = -(I-\Delta)\, dlogmc = -(I-\Delta)\, \Delta^{-1}\pi$

Therefore we can re-write (69) as

$$d^2logY - d^2logL = \tilde{\lambda}^T a - \frac{1}{2}\pi^T \mathcal{D}_2 \pi$$

with

$$d^2_{ij} = \frac{1-\delta_i}{\delta_i}\frac{1-\delta_j}{\delta_j}\tilde{d}^2_{ij}$$

It remains to compute the "within-sector" component $\lambda^T a$.

Index by t the different varieties of product i and note that, given the CES assumption, sectoral output can be written as

$$Y_i = A_i F\left(\{x_{ij}\}, L_i\right) \frac{p_i^{-\epsilon_i}}{\int p_{it}^{-\epsilon_i} dt} \tag{71}$$

where

$$x_{ij} = \int x_{ij}(t)dt$$

$$L_i = \int L_i(t)dt$$

From (71) we see that the "effective" productivity of sector i is

$$\tilde{A}_i = A_i \frac{p_i^{-\epsilon_i}}{\int p_{it}^{-\epsilon_i} dt}$$

so that

$$a_i = \epsilon_i \left[\frac{\int p_{it}^{-\epsilon_i} dlogp_{it} dt}{\int p_{it}^{-\epsilon_i} dt} - \frac{\int p_{it}^{1-\epsilon_i} dlogp_{it} dt}{\int p_{it}^{1-\epsilon_i} dt} \right] \tag{72}$$

Given the Calvo assumption, around the efficient steady state we have that

$$\frac{\int p_{it}^{-\epsilon_i} dlogp_{it} dt}{\int p_{it}^{-\epsilon_i} dt} = \frac{\int p_{it}^{1-\epsilon_i} dlogp_{it} dt}{\int p_{it}^{1-\epsilon_i} dt} = \delta dlogmc_i$$

so that $a_i = 0$.

Let's now compute the second-order loss. The second derivative of $\tilde{\lambda}^T a$ with respect to sector-level productivities is $\sum_i \tilde{\lambda}_i d^2 log\left(\frac{\tilde{A}_i}{A_i}\right)$.

The expression for $d^2 log\left(\frac{\tilde{A}_i}{A_i}\right)$ can be obtained by deriving (72) with respect to $\{dlogp_{it}\}$, to find[27]

$$d^2 log\left(\frac{\tilde{A}_i}{A_i}\right) = \epsilon_i \left[\int (logp_{it} - logp_i)^2 dt - \left(\int (logp_{it} - logp_i) dt\right)^2\right] =$$

$$= \epsilon_i \frac{1-\delta_i}{\delta_i} \pi_i^2$$

We can thus express the second-order welfare loss from within-sector misallocation as

$$\frac{1}{2} \pi \mathcal{D}_1 \pi$$

where

$$d_{ij}^1 = \begin{cases} 0 & if\ i \neq j \\ \lambda_i \epsilon_i \frac{1-\delta_i}{\delta_i} & if\ i = j \end{cases}$$

Lemma 7. *The elasticity of sectoral prices with respect to initial markups $dlog\mu_{-1}$ is given by the matrix $\mathcal{V}$. The optimal output gap is*

$$\widetilde{y}^* = -\frac{\mathcal{B}^T \mathcal{D} \mathcal{V} dlog\mu_{-1}}{\gamma + \varphi + \mathcal{B}^T \mathcal{D} \mathcal{B}} \tag{73}$$

Proof. If initial sector-level markups are given by the vector $dlog\mu_{-1} \neq 0$, then sectoral inflation rates are given by

$$\pi_i = \delta_i \left(dlogmc_i - dlog\mu_{i-1}\right)$$

and sector-level markups are still given by (30). We can then proceed as in Propositions 2 and 1 to derive

$$\pi = \Delta \left(I - \Omega\Delta\right)^{-1} \left(\alpha dlogw - dlog\mu_{-1}\right)$$

[27] This is the same as in the traditional NK model (Gali (2008) Ch.4)

and

$$dlogw = \frac{\gamma + \varphi}{1 - \beta^T (I - \Omega\Delta)^{-1} \alpha} (\tilde{y} - \tilde{y}_{-1}) - \frac{\beta^T \Delta (I - \Omega\Delta)^{-1}}{1 - \beta^T (I - \Omega\Delta)^{-1} \alpha} dlog\mu_{-1}$$

From this we can immediately solve for sectoral inflation rates as a function of $\tilde{y}$,$\tilde{y}_1$ and $dlog\mu_{-1}$.

Welfare is the same function of the output gap and sectoral inflation rates as in (32), because both the variance of firm-level prices within sectors and sector-level markups are the same function of sectoral inflation rates. The optimal output gap then follows from the first order condition. □

B2: Policy target

Proof of Proposition 6:

We look for weights ϕ such that

$$\phi^T (\mathcal{B}\widetilde{y} + \mathcal{V}dlogA) > 0 \Longleftrightarrow \widetilde{y} > \widetilde{y}^* \tag{74}$$

I will first construct a vector ϕ that satisfies the condition

$$\phi^T (\mathcal{B}\widetilde{y} + \mathcal{V}dlogA) = 0 \Longleftrightarrow \widetilde{y} = \widetilde{y}^* \tag{75}$$

and then argue that this vector also satisfies (74).

Note that, as long as $\phi^T \mathcal{B} \neq 0$, we have

$$\phi^T (\mathcal{B}\widetilde{y} + \mathcal{V}dlogA) = 0 \Longleftrightarrow \widetilde{y} = -\frac{\phi^T \mathcal{V} dlogA}{\phi^T \mathcal{B}}$$

while the optimal output gap is

$$\widetilde{y}^* = -\frac{\mathcal{B}^T \mathcal{D} \mathcal{V} dlogA}{\gamma + \varphi + \mathcal{B}^T \mathcal{D} \mathcal{B}}$$

Thus (75) is satisfied for all realizations of $dlogA$ if and only if ϕ is such that

$$\frac{\phi^T \mathcal{V} dlogA}{\phi^T \mathcal{B}} = \frac{\mathcal{B}^T \mathcal{D} \mathcal{V} dlogA}{\gamma + \varphi + \mathcal{B}^T \mathcal{D} \mathcal{B}} \quad \forall dlogA$$

In turn, this is true if and only if

$$\phi^T \left[I - \frac{\mathcal{B}\mathcal{B}^T \mathcal{D}}{\gamma + \varphi + \mathcal{B}^T \mathcal{D} \mathcal{B}} \right] \mathcal{V} = 0 \tag{76}$$

that is, if and only if ϕ is a left eigenvector of the matrix $\left[I - \frac{\mathcal{B}\mathcal{B}^T \Delta \Xi \Delta}{\gamma + \varphi + \mathcal{B}^T \Delta \Xi \Delta \mathcal{B}} \right] \mathcal{V}$, relative to the eigenvalue 0.

We already proved in Lemma 3 that $\lambda^T (I-\Delta)\Delta^{-1}$ is a left eigenvector of the matrix $\mathcal{V}$ relative to the eigenvalue 0 (and it is the only such eigenvector). Therefore, as long as $\left[I - \frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right]$ is invertible, $\phi^T = \lambda^T (I-\Delta)\Delta^{-1}\left[I - \frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right]^{-1}$ is the (unique) desired eigenvector of the matrix $\left[I - \frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right]\mathcal{V}$.

The matrix $\left[I - \frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right]$ is indeed invertible: it is immediate to see that $\frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}$ has only one non-zero eigenvalue, $\frac{\mathcal{B}^T\mathcal{D}\mathcal{B}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}} < 1$, and $\mathcal{B}$ is the unique corresponding eigenvector.

Next, to satisfy condition (74) we need

$$\phi^T \left(\mathcal{B}\widetilde{y} + \mathcal{V} dlogA\right)$$

to be increasing in the output gap $\widetilde{y}$, which is true if and only if $\phi^T\mathcal{B} > 0$. To prove this we use the fact that $\mathcal{B}$ is an eigenvector of $\frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}$ relative to the eigenvalue $\frac{\mathcal{B}^T\mathcal{D}\mathcal{B}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}$. Therefore it is also an eigenvector of $\left[I - \frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right]^{-1}$, relative to the eigenvalue $\frac{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}{\gamma+\varphi} > 1$. Thus we have

$$\begin{aligned}\phi^T\mathcal{B} &= \lambda^T (I-\Delta)\Delta^{-1}\left[I - \frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right]^{-1}\mathcal{B} = \\ &= \gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B} > 0\end{aligned}$$

Finally, to obtain the formulation in (40) we observe that

$$\left[I - \frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right]^{-1} = I + \frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}} + \left(\frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right)^2 + ...$$

and

$$\left(\frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right)^n = \left(\frac{\mathcal{B}^T\mathcal{D}\mathcal{B}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right)^{n-1}\frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}$$

so that

$$\left[I - \frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right]^{-1} = I + \frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi}$$

Moreover, we have that

$$\frac{\lambda^T (I-\Delta)\Delta^{-1}\mathcal{B}}{\gamma+\varphi} = \frac{\lambda^T (I-\Delta)(I-\Omega\Delta)^{-1}\alpha}{1-\beta^T\Delta(I-\Omega\Delta)^{-1}\alpha} = 1$$

so that

$$\lambda^T (I-\Delta)\Delta^{-1}\left[I-\frac{\mathcal{B}\mathcal{B}^T\mathcal{D}}{\gamma+\varphi+\mathcal{B}^T\mathcal{D}\mathcal{B}}\right]^{-1}=\lambda^T (I-\Delta)\Delta^{-1}+\mathcal{B}^T\mathcal{D}$$

Appendix C: Cost-push shocks and Dynamics

Appendix C1: Exogenous cost-push shocks

As discussed in Section 4.1.3, productivity shocks generate an endogenous "cost-push" term in the Phillips curve. In this section I extend the analysis to allow for sector-level "exogenous" cost-push shocks, which I model as a change in producers' desired markup $d\log\mu^D$. Lemma (8) derives sectoral inflation rates and the Phillips curve.

Lemma 8. *The elasticity of sectoral prices with respect to cost-push shocks is given by*

$$\left(\frac{\mathcal{B}\lambda^T}{\gamma+\varphi}-\mathcal{V}\right) \tag{77}$$

The sales-based Phillips curve is

$$SW=(\gamma+\varphi)\,\tilde{y}+\lambda^T dlog\mu^D \tag{78}$$

while the CPI-based Phillips curve is

$$\pi^C=\kappa\tilde{y}+u+v \tag{79}$$

where

$$u=\frac{\bar{\delta}_w-\bar{\delta}_A}{1-\bar{\delta}_w}\lambda^T dlogA$$

$$v=\frac{\bar{\delta}_\mu}{1-\bar{\delta}_w}\lambda^T dlog\mu^D$$

$$\delta_\mu=\frac{\beta^T\Delta\,(I-\Omega\Delta)^{-1}\,dlog\mu^D}{\lambda^T dlog\mu^D}$$

The two Phillips curves expressed in terms of deviations from steady-state output are

$$SW=(\gamma+\varphi)\,y+\lambda^T\left(dlog\mu^D-dlogA\right) \tag{80}$$

$$\pi^C = \kappa y + \frac{\bar{\delta}_\mu \lambda^T dlog\mu^D - \bar{\delta}_A \lambda^T dlogA}{1 - \bar{\delta}_w} \tag{81}$$

Note that for $dlogA = -dlog\mu^D < 0$ we have

$$\frac{u}{\lambda^T dlogA} < \frac{v}{\lambda^T dlog\mu^D}$$

Similar to the baseline model, the central bank faces a worse trade-off after a cost-push shock than after a negative productivity shock of the same size (i.e. $dlogA = -dlog\mu^D$). This is because the change in firms' desired price is the same for the two shocks, but in the cost-push case natural output hasn't changed. In other words, inflation is the same after the two shocks for a given deviation of output from steady-state, while for a given output gap the cost-push shock generates higher inflation. Correspondingly, an additive "aggregate" cost-push term appears in the Phillips-curve.

Lemma (9) solves for the optimal policy response.

Lemma 9. *The optimal monetary policy response to a cost-push shock $dlog\mu^D$ implements output gap*

$$\tilde{y}^*_{CP} = \frac{\mathcal{B}^T \Delta \left(\mathcal{D} \left(\overbrace{\frac{\mathcal{B}\lambda^T}{\gamma + \varphi}}^{\text{inflation-output trade-off}} - \overbrace{\mathcal{V}}^{\text{propagation}} \right) - \overbrace{\mathcal{D}_2 \Delta \left(I - \Delta\right)^{-1}}^{\text{direct effect}} \right)}{(\gamma + \varphi) + \mathcal{B}^T \mathcal{D} \mathcal{B}} dlog\mu^D \tag{82}$$

Under the optimal policy the inflation target derived in Proposition 6 *takes value*

$$\pi_\phi = \left(\lambda^T - \mathcal{B}^T \mathcal{D}_2 \Delta \left(I - \Delta\right)^{-1} \right) dlog\mu^D \tag{83}$$

Comparing (82) with (39) above, we see that the optimal response to productivity and cost-push shocks has a common component (the "propagation" term in (82)). Here, monetary policy seeks to address the relative price distortions induced by the shock through input-output linkages. These are the same regardless of whether inflation responded to fluctuations in productivity or desired markups. In the case of a productivity shock this is the only effect.

Shocks in desired markups instead by definition have also a "direct" effect on relative prices. Monetary policy can partly offset this effect, because by moving the output gap it can decrease the relative markup of the sectors hit by the shock (third term in (82)). Whether this entails increasing or decreasing the output gap depends on how exposed these sectors are to monetary shocks relative to the others. Increasing the output gap is optimal when marginal costs are relatively more exposed to monetary shocks in the sectors hit by the cost-push shock than in the rest of the economy. Whenever this is the case the policy target is positive, as reflected in the second term of (83).

Finally, the first term in equation (82) comes from the fact that monetary policy faces a "worse" trade-off under the cost-push shock than under the productivity shock, because natural output has not fallen. This is also captured by the first term of the policy target (83), which has the same intuition as in the one-sector model: in the face of a cost-push shock the central bank trades off the output loss with the increase in inflation. Therefore after a cost-push shock the output gap should be lower than after an equally-sized productivity shock, while the output level and inflation should be higher. With multiple sectors however this channel is potentially counteracted by the response to the "direct" effect.

Appendix C2: Dynamics - Main results

Consumers

Consumers' intertemporal preferences are given by

$$U = \sum_{t=0}^{\infty} \rho^t \left(\frac{C_t^{1-\gamma}}{1-\gamma} - \frac{L_t^{1+\varphi}}{1+\varphi} \right)$$

where C_t is a CRS bundle of all goods produced in the economy and L_t is labor supply.

In each period consumers are subject to the budget constraint

$$P_t C_t + B_{t+1} \leq w_t L_t + \Pi_t - T_t + (1 + i_t) B_t$$

where $w_t L_t$ is labor income, Π_t are firm profits (rebated lump-sum to households), T_t is a lump-sum transfer (that the government uses to finance input subsidies to firms), B_t is the quantity of risk-free bonds paying off in period t owned by the household and i_t are nominal interest rates.

Consumer optimization yields the Euler equation

$$U_{ct} = \rho(1 + i_{t+1}) \mathbb{E} \left[U_{ct+1} \frac{P_{ct}}{P_{ct+1}} \right] \tag{84}$$

where P_{ct} are consumer prices at time t. We can log-linearize equation (84) and impose market clearing for final goods to find

$$y_t = \mathbb{E}\left[y_{t+1}\right] - \frac{1}{\gamma} \left(i_{t+1} - \mathbb{E}\left[\pi^c_{t+1}\right] - \log \rho \right)$$

In gaps, this becomes

$$\tilde{y}_t = \mathbb{E}\left[\tilde{y}_{t+1}\right] - \frac{1}{\gamma}\left(i_{t+1} - \mathbb{E}\left[\pi^c_{t+1}\right] - r^n_{t+1}\right) \tag{85}$$

where r^n_{t+1} is the natural interest rate, satisfying

$$r^n_{t+1} = \log\rho + \gamma\lambda^T\mathbb{E}\left[logA_{t+1} - logA_t\right]$$

Policy instruments

I consider a cashless economy, in which interest rates are the only policy instrument. At each period t the central bank sets the risk-free rate i_{t+1}.

Production

Within each period the production technology is as described in Section 2.3. Sectoral productivity shifters A_{it} vary across periods.

As in the one-period model, I assume that the government sets input subsidies to eliminate markup distortions. That is, in steady-state all firms optimally set their markup over pre-subsidy marginal costs at 1.

The government cannot change input subsidies in response to productivity shocks, and it cannot give different subsidies to different firms within the same sector.

All producers minimize costs given wages and input prices: producers in sector i solve

$$mc_{it} = min_{\{x_{ijt}\},L_{it}}\ w_tL_{it} + \sum_j p_{jt}x_{ijt}\ \ s.t.\ A_{it}F_i\left(L_{it}, \{x_{ijt}\}\right) = \bar{y}$$

Under constant returns to scale marginal costs are the same for all firms, and they use inputs in the same proportions. However not all firms can adjust prices. Therefore within the same sector firms charge different markups outside of steady-state.

Sector-level inflation dynamics

The firms who can update their price solve

$$p^*_{it} = max_{p_i}\ \mathbb{E}\left[\sum_t SDF_t\left(1-\delta_i\right)^t Y_{it}(p_i)\left(p_i - \left(1-\tau_i\right)mc_{it}\right)\right] \tag{86}$$

The subsidies

$$1 - \tau_i = \frac{\epsilon^*_{it} - 1}{\epsilon^*_{it}}$$

eliminate markup distortions in steady-state. With these subsidies in place the optimal reset price is

$$p_{it}^{*}=\frac{\mathbb{E}\sum_{t}\left[\frac{\epsilon_{it}}{\epsilon_{i}^{*}}SDF_{t}\left(1-\delta_{i}\right)^{t}Y_{it}(p_{i})mc_{it}\right]}{\mathbb{E}\sum_{t}\frac{\epsilon_{it}-1}{\epsilon_{i}^{*}-1}\left[SDF_{t}\left(1-\delta_{i}\right)^{t}Y_{it}(p_{i})\right]} \tag{87}$$

Log-linearizing equation (87) yields the following expression for sector-level inflation rates:

$$\pi_{it}=\frac{\delta_{i}\left(1-\rho(1-\delta_{i})\right)}{1-\delta_{i}}\left(d\log\mu_{it}^{D}-\log\mu_{it}\right)+\rho\mathbb{E}_{t}\left[\pi_{it+1}\right] \tag{88}$$

where μ_{it} is the "sector-level" markup:

$$log\mu_{it}=logp_{it}-logmc_{it}$$

and $\mu_{t}^{D}=\frac{\epsilon_{t}}{\epsilon_{t}-1}$ is the firms' desired markup.

Equilibrium

Equilibrium is defined in a similar way as in section 2.5.

For given sectoral probabilities of price adjustment δ_i, sectoral productivity shifters A_{it} and interest rates i_t for each period t, general equilibrium is given by a vector of firm-level markups μ_{fit}, a vector of 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} for each period t such that: a fraction δ_i of firms in each sector i can adjust their price in every period; markups are chosen optimally by adjusting firms (see problem (86)), while they are such that prices stay constant for the non-adjusting firms; consumers maximize intertemporal utility subject to the budget constraints; producers in each sector i minimize costs and charge the relevant markup; and markets for all goods and labor clear.

The sales-based Phillips curve

Similar to the one-period case, the Phillips curve constructed using SW as our measure of aggregate inflation inherits the properties of the Phillips curve in the one-sector model.[28]

Proposition 7. *It holds that*

$$SW_{t}\equiv\lambda^{T}\left(I-\hat{\Delta}\right)\hat{\Delta}^{-1}\pi_{t}=\rho\mathbb{E}\left(SW_{t+1}\right)+\kappa\tilde{y}_{t}+\lambda^{T}d\log\mu_{t}^{D} \tag{89}$$

[28]Note that the sales-based Phillips curve in (28) does not depend on past markups. This is a consequence of Lemma 7 in Appendix A2.

where

$$\kappa = \gamma + \varphi$$

and $\hat{\Delta}$ is a diagonal matrix with elements

$$\hat{\Delta}_{ii} = \frac{\delta_i \left(1 - \rho(1-\delta_i)\right)}{1 - \rho\delta_i(1-\delta_i)}$$

Response of inflation rates and markups to productivity and monetary shocks

Proposition (8) characterizes the evolution of sector-level inflation, inflation expectations and markups for given productivity shocks $logA_t - logA_{t-1}$, monetary policy $\tilde{y}_t$, and past markups. Note that, different from the one-sector model and the sales-based Phillips curve in Section 8, past markups are state variables.[29]

Proposition 8. *Denote by*

$$\mathcal{M} \equiv \left(\frac{\hat{\mathcal{B}}\lambda^T}{\gamma+\varphi} - \hat{\mathcal{V}}\right)\left(I - \hat{\Delta}\right)\hat{\Delta}^{-1}$$

The evolution of sectoral markups and inflation rates is given by the following system of difference equations:

$$\begin{pmatrix} \rho\mathbb{E}\pi_{t+1} \\ log\mu_t \end{pmatrix} = \begin{pmatrix} \mathcal{M}^{-1} & -\mathcal{M}^{-1}\hat{\mathcal{V}} \\ -\left(I-\hat{\Delta}\right)\hat{\Delta}^{-1}\left(I - \mathcal{M}^{-1}\right) & -\left(I-\hat{\Delta}\right)\hat{\Delta}^{-1}\mathcal{M}^{-1}\hat{\mathcal{V}} \end{pmatrix}\begin{pmatrix} \pi_t \\ log\mu_{t-1} \end{pmatrix} +$$

$$+ \begin{pmatrix} -\mathcal{M}^{-1}\left(\hat{\mathcal{B}}\tilde{y}_t + \hat{\mathcal{V}}\left(logA_t - logA_{t-1}\right)\right) - \hat{\Delta}\left(I-\hat{\Delta}\right)^{-1} d\log\mu_t^D \\ -\left(I-\hat{\Delta}\right)\hat{\Delta}^{-1}\mathcal{M}^{-1}\left(\hat{\mathcal{B}}\tilde{y}_t + \hat{\mathcal{V}}\left(logA_t - logA_{t-1}\right)\right) \end{pmatrix} \tag{90}$$

Proof. See online Appendix C3 □

The first equation in (90) extends the results from Sections 4.1.2 and 4.1.3 to the dynamic

[29]The SW-Phillips curve is independent of past markups, therefore the actual state variables are "relative" past markups $\mathcal{V}d\log\mu_{t-1}$. Given these, the system is invariant to the "aggregate" past markup $\lambda^T d\log\mu_{t-1}$.

setup. We can re-write it as

$$\pi_t = \underbrace{\hat{\mathcal{B}}\tilde{y}_t + \hat{\mathcal{V}}\left(logA_t - logA_{t-1} + log\mu_{t-1}\right)}_{\text{productivity and past markups}} + \underbrace{\left(\frac{\hat{\mathcal{B}}\lambda^T}{\gamma+\varphi} - \hat{\mathcal{V}}\right)\left(d\log\mu_t^D + \left(I - \hat{\Delta}\right)\hat{\Delta}^{-1}\rho\mathbb{E}\pi_{t+1}\right)}_{\text{cost-push shock}} \tag{91}$$

The first term contains the elasticities of sectoral inflation rates with respect to productivity and monetary shocks, which are the same as in the static setup. In addition, we now have to account for pre-existing markups, due to the fact that some producers could not adjust their price in past periods. To gain intuition, Lemma 7 in Appendix B2 shows that in the static setting the response of inflation to productivity shocks and initial markups is the same, because both induce the same desired price changes under zero output gap. Finally, expected future inflation acts as a "cost-push" shock, and indeed it has the same effect on sectoral inflation rates. To map inflation expectations into the corresponding "cost-push" shock by setting:

$$dlog\mu_{\mathbb{E}}^D = (I - \Delta)\,\Delta^{-1}\rho\mathbb{E}\pi_{t+1} \tag{92}$$

Consumer price Phillips curve

We can aggregate sectoral inflation rates into the CPI-based Phillips curve from equation (91):

$$\pi_t^C = \kappa\tilde{y}_t + \rho\mathbb{E}\pi_{t+1}^C + u_t + v_t \tag{93}$$

where

$$u_t = \frac{\bar{\delta}_w - \bar{\delta}_A}{1-\bar{\delta}_w}\lambda^T\left(dlogA_t - dlogA_{t-1}\right) + \frac{\bar{\delta}_w - \bar{\delta}_{\mu_{-1}}}{1-\bar{\delta}_w}\lambda^T dlog\mu_{t-1}$$

$$v_t = \frac{\bar{\delta}_w - \bar{\delta}_{\pi^C}}{1-\bar{\delta}_w}\rho\mathbb{E}\pi_{t+1}^C + \frac{\bar{\delta}_{\mu^D}}{1-\bar{\delta}_w}\lambda^T d\log\mu_t^D$$

$$\bar{\delta}_{\mu_{-1}} = \frac{\beta^T\Delta\left(I - \Omega\Delta\right)^{-1} dlog\mu_{t-1}}{\lambda^T dlog\mu_{t-1}}$$

$$\bar{\delta}_{\pi^C} = \frac{\beta^T\Delta\left(I - \Omega\Delta\right)^{-1}\left(I - \Omega\right)\mathbb{E}\pi_{t+1}}{\mathbb{E}\pi_{t+1}^C}$$

$$\bar{\delta}_{\mu^D} = \frac{\beta^T\Delta\left(I - \Omega\Delta\right)^{-1} d\log\mu^D}{\lambda^T d\log\mu^D}$$

Equation (93) highlights that past markups and inflation expectations also create an endogenous cost-push term in the Phillips curve, exactly in the same way as productivity shocks. This is not surprising, given that their propagation across the input-output network is the same.

Targeting the output gap

Lemma 10 shows that the central bank can implement zero output gap in all periods by targeting it directly in the nominal rate, as long as the targeting rule is "reactive enough".

Lemma 10. *There is a unique path of inflation rates such that the output gap is constantly zero. This can be implemented with the interest rate rule*

$$i_t = \underbrace{r_t^n + \beta^T \mathbb{E}\pi_{t+1}^{zg}}_{\text{nominal rate under zero output gap}} + \zeta \tilde{y}_t \tag{94}$$

with $\zeta > 0$.

Proof. See online Appendix C3. □

Optimal policy

Propositions 9 and 10 characterize the dynamic welfare loss and the central bank's problem without commitment.

Proposition 9. *Given a path* $\{y_t, \pi_t, z_{t+1}\}_{t=0}^{\infty}$ *for the output gap, sectoral inflation rates and markups, the expected second-order welfare loss is given by*

$$\sum_{t=0}^{\infty} \rho^s \mathbb{E}\left[(\gamma + \varphi)\, \tilde{y}_t^2 + \pi_t^T \mathcal{D}_1 \pi_t + z_{t+1}^T \mathcal{D}_2 z_{t+1}\right]$$

where

$$z_t \equiv -\left(I - \hat{\Delta}\right)^{-1} \hat{\Delta} log\mu_{t-1}$$

Therefore the optimal policy problem is

$$min \sum_{t=0}^{\infty} \rho^s \mathbb{E}\left[(\gamma + \varphi)\, \tilde{y}_t^2 + \pi_t^T \mathcal{D}_1 \pi_t + z_{t+1}^T \mathcal{D}_2 z_{t+1}\right]$$

subject to (90).

Proposition 10. *Consider an optimal policy problem without commitment, where the central bank takes future output gaps and inflation rates as given. Here the central banks chooses* $\{\tilde{y}_t, \pi_t, z_{t+1}\}$

to minimize current-period losses, solving

$$min_{\{\tilde{y}_t,\pi_t,z_{t+1}\}}\left(\gamma+\varphi\right)\tilde{y}_t^2+\pi_t^T\mathcal{D}_1\pi_t+z_{t+1}^T\mathcal{D}_2 z_{t+1}$$

*subject to (*90*). The FOCs yield*

$$\tilde{y}_t^* = -\underbrace{\frac{\mathcal{B}^T\mathcal{D}\mathcal{V}\left(dlogA_t - dlog\mu_{t-1}\right)}{\left(\gamma+\varphi\right)+\mathcal{B}^T\mathcal{D}\mathcal{B}}}_{\text{productivity and past markups}} - \underbrace{\frac{\mathcal{B}^T\left(\mathcal{D}\mathcal{M}-\mathcal{D}_2\right)}{\left(\gamma+\varphi\right)+\mathcal{B}^T\mathcal{D}\mathcal{B}}\rho\mathbb{E}\pi_{t+1}}_{\text{cost-push shock}} \tag{95}$$

Proof. See online Appendix C3 □

In the same spirit as Proposition 8, Proposition 9 decomposes the optimal output gap into the response to productivity shocks and past markups, which is the same as in the static setup (see Proposition 5), and the response to inflation expectations, which act as a "cost-push" shock (see Section 5.1). Lemma 11 and Lemma 12 below show that the optimal policy can be implemented with a targeting rule, in the same way as in the static setup (see Proposition 6).

Lemma 11. *Assume that the productivity shocks follow an AR1 process:*

$$logA_{t+1} - logA_t = \eta\left(logA_t - logA_{t-1}\right) + u_{t+1}$$

with $\mathbb{E}u_{t+1} = 0$ *and* $\eta < 1$*.*

*Then there is a unique path of inflation rates such that the optimal output gap (*95*) is implemented in each period.*

Lemma 12. *Denote by*

$$\begin{cases} \phi_t^T \equiv \frac{\lambda^T\left(I-\hat{\Delta}\right)\hat{\Delta}^{-1}+\mathcal{B}^T\mathcal{D}}{\gamma+\varphi} \\ \phi_{t+1}^T \equiv \rho\frac{\lambda^T\left(I-\hat{\Delta}\right)\hat{\Delta}^{-1}-\mathcal{B}^T\mathcal{D}_2}{\gamma+\varphi} \end{cases} \tag{96}$$

For $\zeta > \gamma$ *the interest rate rule*

$$i_t = \underbrace{\overbrace{r_t^n + \gamma\left[\mathbb{E}\widetilde{y}_{t+1}^* - \widetilde{y}_t^*\right]}^{\text{real rate under opt policy}} + \beta^T\mathbb{E}\pi_{t+1}^*}_{\text{nominal rate under optimal policy}} + \zeta\underbrace{\left(\phi_t\pi_t + \phi_{t+1}\mathbb{E}\pi_{t+1}\right)}_{\text{inflation target}} \tag{97}$$

*implements the optimal policy (*95*).*

The nominal rate can be expressed as a function of productivity shocks and current and expected inflation using the relation

$$\gamma\left[\mathbb{E}\tilde{y}^*_{t+1}-\tilde{y}^*_t\right]+\beta^T\mathbb{E}\pi^*_{t+1}=\kappa_t\pi^*_t+\kappa_{t+1}\mathbb{E}\pi^*_{t+1}+\kappa_A\left(dlogA_t-dlogA_{t-1}\right)$$

where

$$\begin{cases}\kappa_t^T\equiv-\lambda^T\left(I-\hat{\Delta}\right)\hat{\Delta}^{-1}\left(I+\left[\mathcal{M}-\frac{\mathcal{B}\mathcal{B}^T\mathcal{D}_2}{\gamma+\varphi}\right]^{-1}\hat{\mathcal{V}}\left(I-\hat{\Delta}\right)\hat{\Delta}^{-1}\right)\\ \kappa_{t+1}^T\equiv-\left(\lambda^T\left(I-\hat{\Delta}\right)\hat{\Delta}^{-1}\left(\left(1+\frac{\gamma}{(\gamma+\varphi)\rho}\right)I-\left[\mathcal{M}-\frac{\mathcal{B}\mathcal{B}^T\mathcal{D}_2}{\gamma+\varphi}\right]^{-1}\left(\frac{\gamma\left(I+\mathcal{B}\mathcal{B}^T\mathcal{D}\right)}{(\gamma+\varphi)\rho}-\hat{\mathcal{V}}\left(I-\hat{\Delta}\right)\hat{\Delta}^{-1}\right)\right)+\beta^T\right)\\ \kappa_A\equiv\eta\frac{\lambda^T\left(I-\hat{\Delta}\right)\hat{\Delta}^{-1}}{(\gamma+\varphi)}\left[\mathcal{M}-\frac{\mathcal{B}\mathcal{B}^T\mathcal{D}_2}{\gamma+\varphi}\right]^{-1}\hat{\mathcal{V}}\end{cases}$$

Proof. See online Appendix C3 □

Comparing equation (44) from the static setup with equation (96) we see that they yield a very similar inflation target. The key difference is that in the dynamic setup the central bank should not just target current inflation, but also inflation expectations. The intuition is that inflation expectations act as a negative cost-push shock, and we know that with cost-push shocks the "static" inflation target is not zero under the optimal policy. Specifically, from Lemma 9 and equation (83) we know that, after a shock $dlog\mu^D$ to desired markups, the inflation target takes value

$$\pi_T=\left(\lambda^T-\mathcal{B}^T\mathcal{D}_2\left(I-\Delta\right)^{-1}\right)dlog\mu^D$$

under the optimal policy. Correspondingly, from equation (92) we see that the time-t inflation target should be adjusted by

$$\left(\lambda^T\left(I-\Delta\right)\Delta^{-1}-\mathcal{B}^T\mathcal{D}_2\right)\rho\mathbb{E}\pi_{t+1}$$

Appendix D: Empirics

Appendix D1: Data

Sector-level PPI series

This section describes the construction of SW. The main data source for sector-level price series is PPI data from the BLS. In this dataset the sample period varies across sectors: most manufacturing series are available by the mid-1980s, while most service series are available by 2006. To make up for the incomplete series we run a Lasso regression of each incomplete series on disaggregated (338 sectors) PCE components for the period in which both are available. We then use the Lasso approximation to extend the series as far back as 1984. There are 172 incomplete series in total, and 67 missing price series with respect to the BEA classification. We make up for 40 of them by using the concordance between NAICS sectors and PCE series provided by the BEA. Summary statistics for the Lasso regressions and information about the incomplete and missing series are reported in the online Appendix.

Comparison between PCE and SW

Figure 15 compares the weights assigned to different sectors by SW and PCE, at an aggregated 21-sector level.

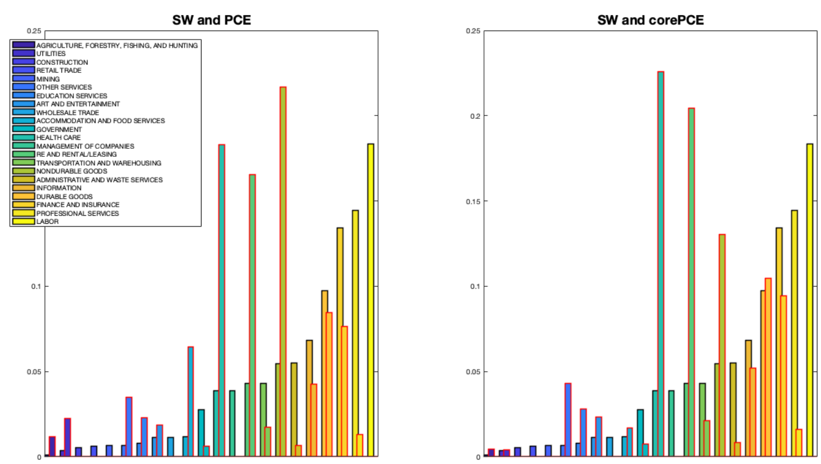

Figure 15: SW and PCE weights (The bars are ordered so that sectoral weights in SW are increasing. Those with red borders correspond to the PCE)

Detailed sectoral weights are reported in the online Appendix.

Appendix D2: Regressions over the full sample

Baseline specification (additional specifications):

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

	SW	CPI	core CPI	PCE	core PCE
gap	3.144**	0.2791**	0.1728**	0.1837**	0.1162**
	(0.5538)	(0.0618)	(0.055)	(0.0532)	(0.0482)
intercept	2.0189**	3.0193**	2.9661**	2.4878**	2.4325**
	(0.0522)	(0.1271)	(0.1131)	(0.1095)	(0.0992)
R-squared	0.1905	0.1297	0.0673	0.08	0.0407

Table 7: CBO output gap

	SW	CPI	core CPI	PCE	core PCE
gap	-3.084**	-0.1405*	0.0028	-0.036	0.0545
	(0.6645)	(0.0759)	(0.0661)	(0.0644)	(0.057)
intercept	1.9621**	2.8021**	2.7595**	2.2996**	2.2514**
	(0.0505)	(0.1259)	(0.1096)	(0.1067)	(0.0945)
R-squared	0.1359	0.0244	0	0.0023	0.0066

Table 8: unemployment rate

With oil prices:

$$\pi_t = c + \kappa \tilde{y}_t + \pi_{oil} + u_t$$

	SW	CPI	core CPI	PCE	core PCE
gap	-3.6385**	-0.2198**	-0.2038**	-0.1194**	-0.1066*
	(0.6294)	(0.0655)	(0.0643)	(0.0584)	(0.0573)
intercept	1.9532**	2.7286**	2.9576**	2.266**	2.3883**
	(0.0483)	(0.1099)	(0.1078)	(0.0979)	(0.0961)
oil prices	0.0032**	0.0185**	-0.0058*	0.0138**	-0.0017
	(0.0013)	(0.003)	(0.0029)	(0.0027)	(0.0026)
R-squared	0.2488	0.2959	0.0829	0.2049	0.0257

Table 9: CBO unemployment gap

	SW	CPI	core CPI	PCE	core PCE
gap	2.8985**	0.2137**	0.1961**	0.1351**	0.1243**
	(0.5562)	(0.0562)	(0.0553)	(0.0501)	(0.0492)
intercept	1.9843**	2.8179**	3.038**	2.3383**	2.4576**
	(0.0536)	(0.1184)	(0.1164)	(0.1055)	(0.1036)
oil prices	0.0031**	0.0179**	-0.0064**	0.0133**	-0.0022
	(0.0014)	(0.003)	(0.0029)	(0.0027)	(0.0026)
R-squared	0.2199	0.3108	0.0985	0.2221	0.0458

Table 10: CBO output gap

	SW	CPI	core CPI	PCE	core PCE
gap	-2.7813**	-0.0591	-0.0158	0.0259	0.0524
	(0.6655)	(0.0684)	(0.067)	(0.0596)	(0.0582)
intercept	1.9278**	2.601**	2.8056**	2.1465**	2.2564**
	(0.0516)	(0.116)	(0.1136)	(0.101)	(0.0987)
oil prices	0.0033**	0.0196**	-0.0045	0.0149**	-0.0005
	(0.0014)	(0.0031)	(0.0031)	(0.0027)	(0.0027)
R-squared	0.1707	0.2418	0.0155	0.1816	0.0069

Table 11: unemployment rate

With inflation expectations (additional specifications):

$$\pi_t = c + \kappa\tilde{y}_t + \rho\mathbb{E}\pi_{t+1} + \epsilon_t$$

	SW	CPI	core CPI	PCE	core PCE
gap	1.0861**	0.1881**	0.0412	0.0881**	0.0084
	(0.2714)	(0.0678)	(0.0449)	(0.0417)	(0.032)
inflation expecations	0.8297**	0.4412**	0.5398**	0.6231**	0.6365**
	(0.0368)	(0.1515)	(0.0561)	(0.0617)	(0.0455)
intercept	0.3668**	1.6124**	1.3548**	0.6459**	0.8614**
	(0.0772)	(0.4987)	(0.1892)	(0.2005)	(0.1291)
R-squared	0.8288	0.1808	0.4442	0.4744	0.6073

Table 12: CBO output gap

	SW	CPI	core CPI	PCE	core PCE
gap	-0.9404**	-0.0049	0.0781	-0.0505	0.0757**
	(0.3185)	(0.0788)	(0.0499)	(0.0477)	(0.0355)
inflation expecations	0.8468**	0.6312**	0.5668**	0.6549**	0.6432**
	(0.0372)	(0.1518)	(0.0537)	(0.0608)	(0.0434)
intercept	0.3108**	0.8344*	1.1705**	0.4941**	0.7757**
	(0.0762)	(0.4879)	(0.1711)	(0.1851)	(0.1155)
R-squared	0.8202	0.1344	0.4507	0.4616	0.6198

Table 13: unemployment rate

In changes (inflation change vs gap level):

$$\pi_t - \pi_{t-1} = c + \kappa \tilde{y}_t + u_t$$

	SW	CPI	core CPI	PCE	core PCE
gap	-0.6945**	-0.0287	-0.0212*	-0.0169	-0.0105
	(0.3258)	(0.0404)	(0.0119)	(0.0296)	(0.0123)
intercept	0.0182	0.008	-0.003	-0.0006	-0.0078
	(0.0244)	(0.0661)	(0.0195)	(0.0485)	(0.0201)
R-squared	0.0323	0.0037	0.0227	0.0024	0.0054

Table 14: CBO unemployment gap

	SW	CPI	core CPI	PCE	core PCE
gap	0.7805**	0.0472	0.02*	0.0299	0.0103
	(0.2768)	(0.0345)	(0.0102)	(0.0254)	(0.0106)
intercept	0.0365	0.0422	0.0046	0.022	-0.0035
	(0.0262)	(0.0713)	(0.0211)	(0.0524)	(0.0218)
R-squared	0.0552	0.0136	0.0273	0.0101	0.007

Table 15: CBO output gap

	SW	CPI	core CPI	PCE	core PCE
gap	-0.7464**	-0.0317	-0.0238**	-0.0199	-0.0133
	(0.3264)	(0.0405)	(0.0119)	(0.0297)	(0.0123)
intercept	0.0212	0.0113	-0.0002	0.0024	-0.0052
	(0.0247)	(0.0668)	(0.0197)	(0.049)	(0.0203)
R-squared	0.037	0.0045	0.0284	0.0033	0.0085

Table 16: unemployment rate

Appendix D3: Rolling regressions

The figures below plot estimated coefficients (with confidence intervals) and R-squareds. Alternative specifications are reported in the online Appendix.

Regression equation:

$$\pi_t = \kappa \tilde{y}_t + \rho \mathbb{E}\pi_{t+1} + \epsilon_t$$

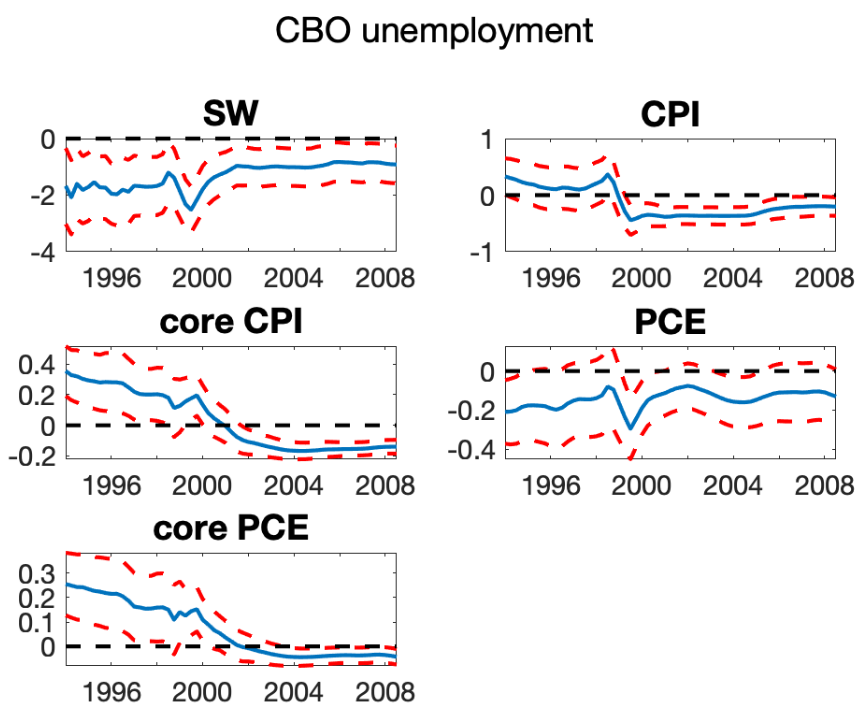

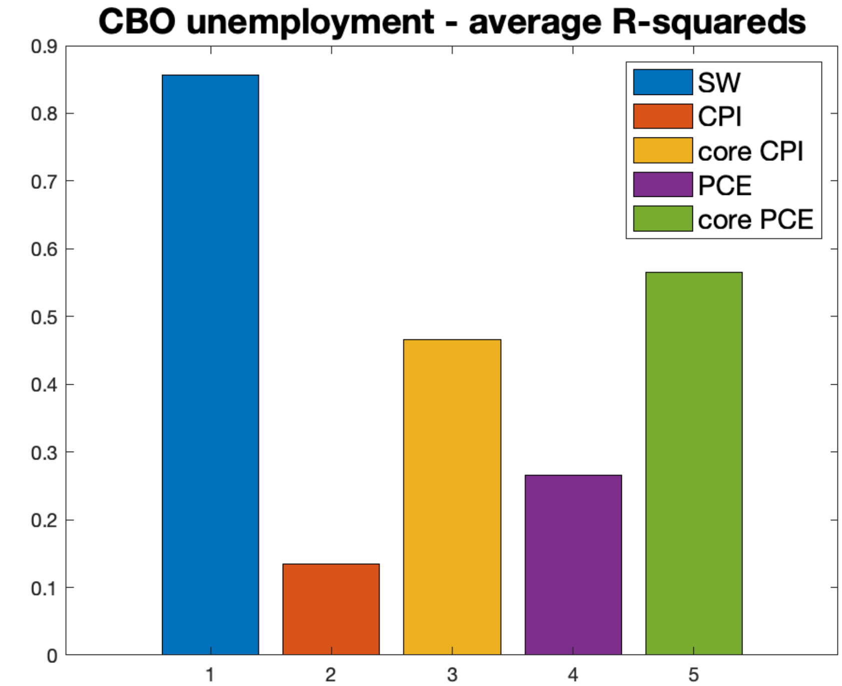
CBO unemployment - average R-squareds
SW
CPI
core CPI
PCE
core PCE
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
2
3
4
5

CBO output

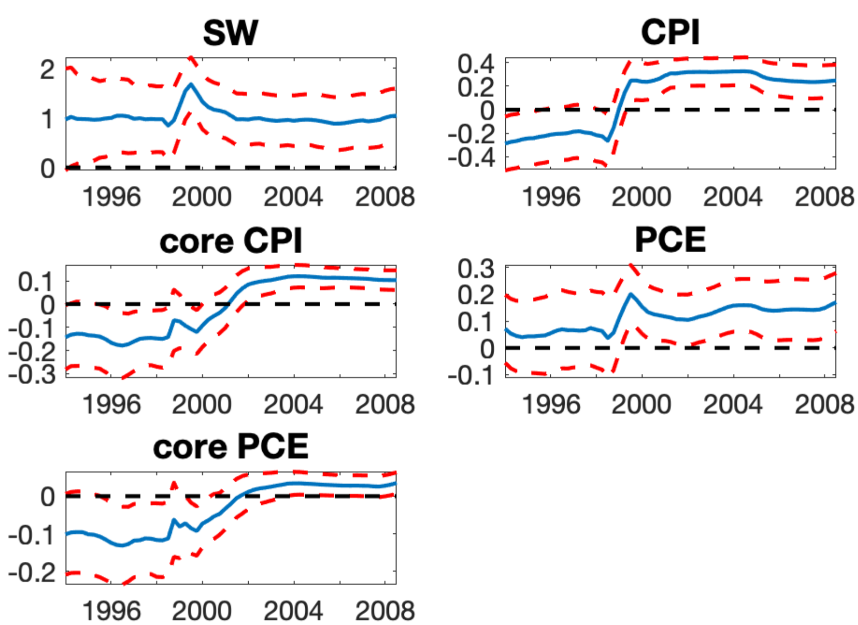
SW
CPI
core CPI
PCE
core PCE
1996
2000
2004
2008

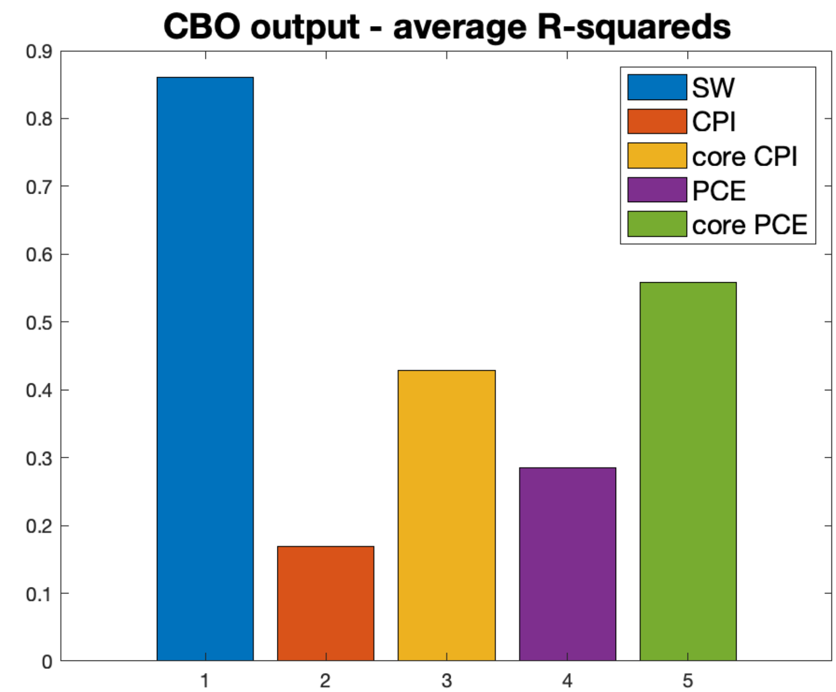
CBO output - average R-squareds
SW
CPI
core CPI
PCE
core PCE
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
2
3
4
5

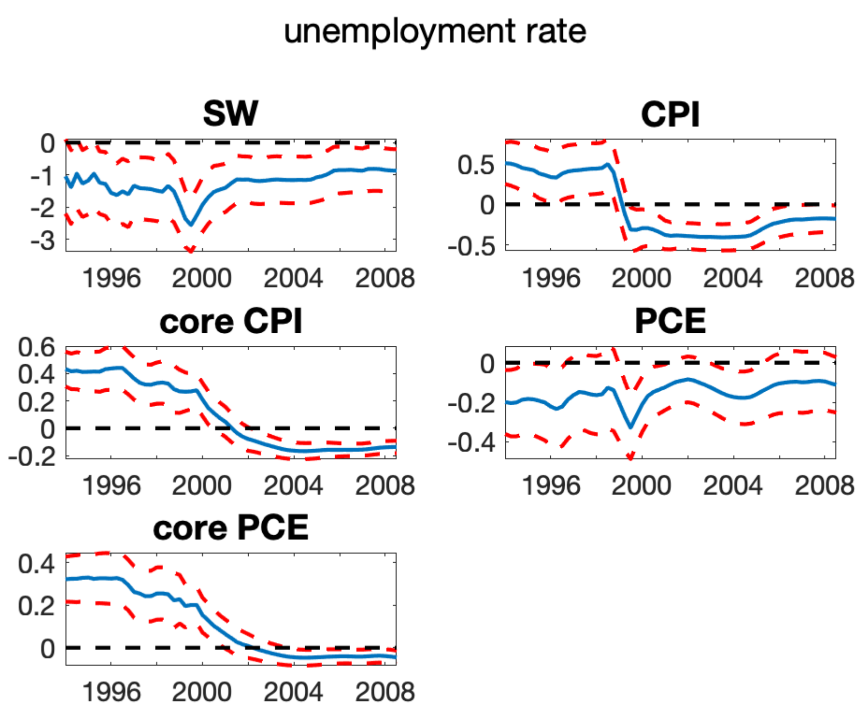
unemployment rate
SW
CPI
core CPI
PCE
core PCE
1996
2000
2004
2008

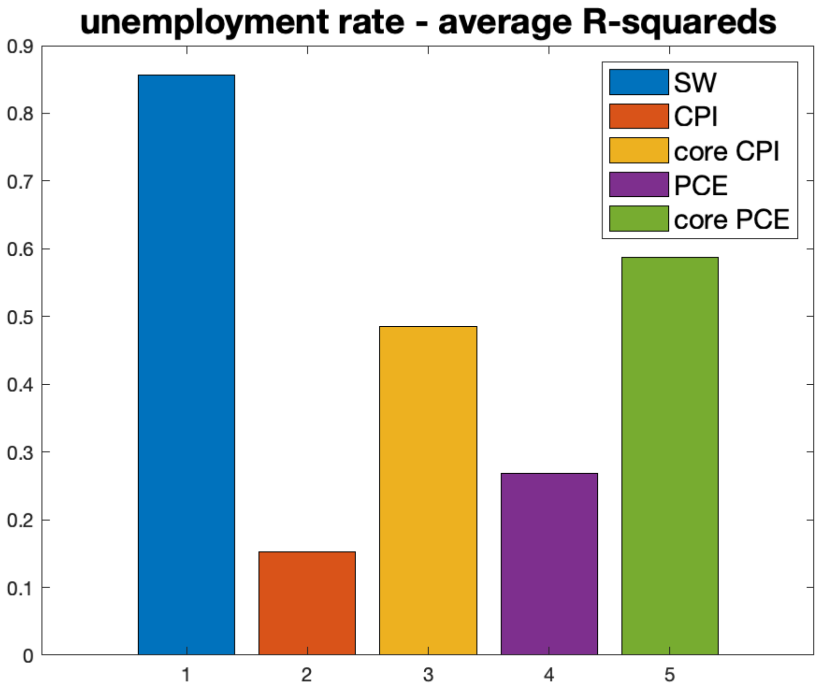
unemployment rate - average R-squareds
SW
CPI
core CPI
PCE
core PCE
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
2
3
4
5