



Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jmonecoThe inflationary effects of sectoral reallocation[☆]

Francesco Ferrante, Sebastian Graves*, Matteo Iacoviello

Federal Reserve Board, United States

ARTICLE INFO

Article history:

Received 11 November 2022

Revised 13 March 2023

Accepted 13 March 2023

Available online xxx

Keywords:

Sectoral reallocation

Inflation

Input-output models

Moment-matching exercise

ABSTRACT

The COVID-19 pandemic has led to an unprecedented shift of consumption from services to goods. We study this demand reallocation in a multi-sector model featuring sticky prices, input-output linkages, and labor reallocation costs. Reallocation costs hamper the increase in the supply of goods, causing inflationary pressures. These pressures are amplified by the fact that goods prices are more flexible than services prices. We estimate the model allowing for demand reallocation, sectoral productivity, and aggregate labor supply shocks. The demand reallocation shock explains a large portion of the rise in U.S. inflation in the aftermath of the pandemic.

© 2023 Published by Elsevier B.V.

1. Introduction

The COVID-19 pandemic has led to a large, abrupt, and unprecedented increase in the demand for goods relative to services in the United States, interrupting a secular decline in the share of spending on goods. A popular narrative is that this sudden reallocation of demand has strained supply chains, leading to bottlenecks and labor shortages in a number of key sectors, thus contributing to a buildup of inflationary forces. Figure 1 illustrates the recent behavior of consumption, inflation, and employment in the U.S. economy. The share of consumption expenditures on goods rose from 31 percent in the last quarter of 2019 to more than 35 percent by the middle of 2021, and has remained high thereafter.¹ Personal Consumption Expenditures inflation reached almost six percent by the end of 2021, primarily driven by a surge in goods inflation, while services inflation has been more muted. Finally, employment collapsed and rebounded, remaining significantly below the pre-pandemic trend by the end of the sample, driven by a decline in labor market participation. Figure 2 shows that these aggregate movements mask even larger movements in more disaggregated data, illustrating how the COVID-19 pandemic has been accompanied by an unprecedented increase in the dispersion of output, prices, and employment across industries.

In this paper, we develop a multi-sector New Keynesian model of the U.S. economy to quantify the aggregate and cross-sectional implications of this reallocation of demand. The model features input-output linkages between sectors, hetero-

[☆] We are grateful to Ambrogio Cesa-Bianchi, Michael Weber, Mishel Ghassebi, and Omar Rachedi for insightful discussions of our paper. We also thank Andrea Prestipino and Johannes Pfeifer for numerous conversations as well as seminar participants at UBC, the Midwest Macroeconomics Conference, the Federal Reserve Board, the Bank of Finland, the Paris School of Economics, the Federal Reserve Bank of Cleveland, the European Commission and the DNB for helpful comments. We are also grateful to Michael Weber and Raphael Schoenle for sharing data on the frequency of price adjustment across industries. All errors and omissions are our own. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System.

* Corresponding author.

E-mail addresses: francesco.ferrante@frb.gov (F. Ferrante), sebastian.h.graves@frb.gov (S. Graves), matteo.iacoviello@frb.gov (M. Iacoviello).

¹ Throughout this paper, we use data available through the first half of 2022.

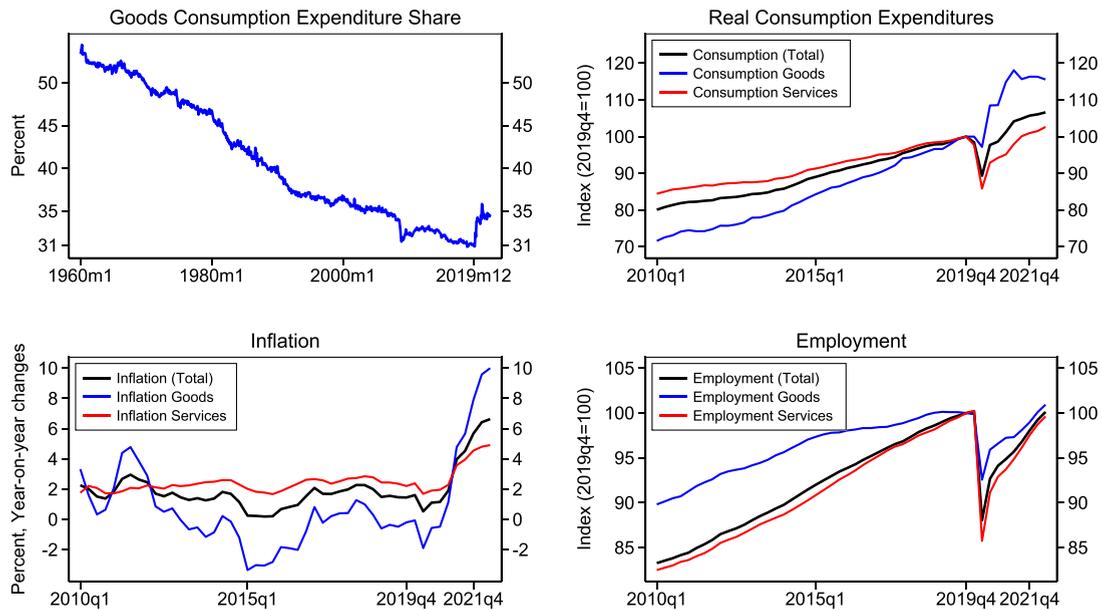


Fig. 1. Consumption, Inflation and Employment in the Goods and Services Sectors. The COVID-19 pandemic led to an unprecedented increase in the demand for goods relative to services in the United States (top panels). Personal Consumption Expenditures inflation has risen, more for goods than for services (bottom left panel). Employment has initially declined before recovering, more in the goods than in the service sector (bottom right panel). In the top left panel, the monthly goods share is expressed as the share in total PCE of nominal goods consumption.

generosity in sectoral price rigidity, and costs of reallocating inputs across sectors.² In particular, we assume that firms face convex hiring costs when increasing their labor input; as our model does not include capital, these hiring costs capture a variety of frictions affecting a firm's ability to expand its productive capacity. Based on the aggregate and cross-sectional developments outlined in Figs. 1 and 2, we allow for three shocks: a preference shock that alters the relative demand for goods and services; sectoral productivity shocks; and an aggregate labor supply shock. Using aggregate and cross-sectional data, we then estimate the parameters governing hiring costs and production function elasticities as well as the size of the aggregate labor supply shock. The estimated model allows us to quantify the role that each shock has played in driving aggregate and cross-sectional developments in the aftermath of the COVID-19 pandemic.

We study the implications of each of the three shocks individually and then examine how well the model fits the data when all the shocks occur at once. We find that the demand reallocation shock is able to explain a large portion—3.5 percentage points—of the increase in U.S. inflation post-pandemic.³ In the model, inflation occurs in response to a reallocation shock for two main reasons. First, because of the hiring costs, firms in goods-producing sectors can increase their labor input only gradually. While these firms could adjust production by using more intermediate inputs, these are only imperfect substitutes for labor, causing a slow adjustment in quantities and a large rise in prices. Furthermore, since goods produced by one sector are also used as intermediate inputs by others, the inflationary pressures propagate across sectors through the production network. In contrast, service-producing sectors reduce production swiftly, with only modest declines in prices. Second, the inflationary effects of the shift in demand are amplified by the heterogeneity in price rigidity that exists across sectors. A key feature of the data is that industries that produce goods have more flexible prices than those that produce services. We find that allowing for heterogeneity in price rigidity across sectors increases the inflationary effects of the preference shock by around 25 percent.

At the industry level, we show that our demand reallocation shock is able to explain a good proportion of the cross-sectional evolution of prices and quantities since the onset of the pandemic. Not only does the shock explain why goods prices have risen more than services prices, but it also accounts for the observed heterogeneity within goods-producing and within services-producing industries, despite the fact that it affects final demand for goods and services uniformly. Both input-output linkages and sectoral heterogeneity in price stickiness contribute to this result. In the model as in the data, sectors producing goods which are directly consumed by households or selling inputs which are heavily used in the production of these goods experience a larger increase in inflation. Furthermore, sectors with more flexible prices exhibit larger price changes, all else equal.

² We model the industry structure after the U.S. input-output tables provided by the BEA as in Baqaee and Farhi (2022). We calibrate the heterogeneity in price rigidity as in Pasten et al. (2020). We estimate the cost of reallocating inputs using the strategy discussed in Section 3.

³ As shown in Fig. 1, inflation rose by 4.2 percentage points between 2019:Q4 and 2021:Q4.

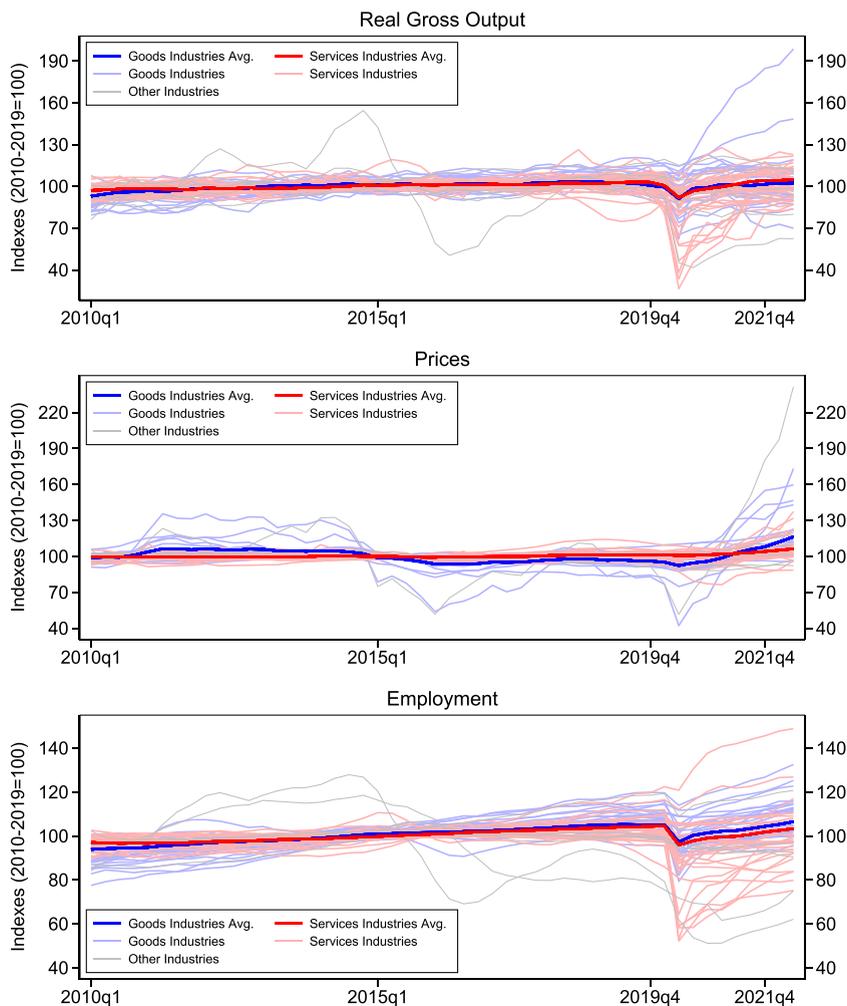


Fig. 2. Output, Prices and Employment across 66 Private Industries. Each line denotes the evolution since 2010 of the 66 private industries for which BEA publishes quarterly data on gross output, prices, and intermediate inputs. Individual industries and averages (weighted by industry gross output) are indexed to 100 in the 2010–2019 period. Employment data at the 3-digit NAICS code level are aggregated at the BEA industry level using the concordance described in <https://www.uspto.gov/sites/default/files/documents/occe-ip-economy-supplement.pdf>. Variables at the industry level are detrended by calculating for each industry a log-linear time trend from 2005:Q1 through 2019:Q4. Gray lines denote sectors (“other” sectors) for which no output is directly consumed.

We then examine the two supply shocks. The sectoral productivity shocks are motivated by the increase in the dispersion of sector-level variables shown in Fig. 2. Additionally, some sectors, such as the metals or oil industries, have experienced both significant declines in production and increases in prices, which cannot be explained by demand reallocation alone. To account for this, we measure the evolution of total factor productivity at the industry level between 2019:Q4 and 2021:Q4, and feed the estimated shocks into our multi-sector model. We find that sectoral productivity shocks dramatically improve the model’s cross-sectional fit, but dampen aggregate inflation, as aggregate productivity rose above trend over this period. The final shock we consider is a reduction in aggregate labor supply, motivated by the prolonged decline in employment shown in Fig. 1. We estimate the magnitude of this shock and find that it explains approximately two-thirds of the post-pandemic decline in employment. However, its effect on inflation is relatively limited: on its own, it would only increase inflation by around 1.5 percentage points, which is less than half the impact of the demand reallocation shock.

When we consider the effect of all three shocks simultaneously, the estimated model can explain the majority of the rise in U.S. inflation between the end of 2019 and the end of 2021, largely driven by the demand reallocation shock.⁴ The model also explains a large proportion of the cross-sectional dynamics of prices and quantities: both the demand reallocation shock and the sectoral productivity shocks are important for this finding. The labor supply shock is important for explaining

⁴ Due to the non-linearities inherent in the model, the total effect of the three shocks is not equal to the sum of the individual effects.

the persistent decline in aggregate employment, but plays a smaller role in explaining aggregate inflation and no role in accounting for the model's cross-sectional fit.

We extend our model by conducting a variety of experiments pertaining to the properties of the demand reallocation shock. We find that an unexpected reversal of the reallocation shock would be inflationary, driven by rising services prices, as services sectors struggle to increase capacity. We also consider a scenario in which households and firms are repeatedly surprised about how persistent the reallocation shock is. In such a scenario, inflationary pressures are more muted, as services-producing sectors reduce output by less, and prices by more, than in our baseline assumption in which the high persistence of the shock is known immediately. We then apply our model to two episodes not directly targeted by our estimation exercise. We show that demand reallocation during the Great Recession—away from goods and towards services—may have raised inflation by around 1.5 percentage points. Finally, we show that the model can rationalize the persistence of inflation during 2022 when we allow for productivity developments that occurred in the first half of 2022, which were negative in many sectors, particularly those producing goods.

In [Section 2](#) we describe the model, which we calibrate and estimate in [Section 3](#). [Section 4](#) studies the cross-sectional and aggregate effects of the demand reallocation shock and the two supply shocks: sectoral productivity shocks and an aggregate labor supply shock. In [Section 5](#) we study various extensions of the model, while [Section 6](#) discusses sensitivity analysis.

1.1. Related literature

The model in our paper builds on the rapidly growing literature studying the role of production networks in propagating the effects of monetary policy, such as [La'O and Tahbaz-Salehi \(2022\)](#), [Pasten et al. \(2020\)](#), [Ozdagli and Weber \(2017\)](#) and [Ghassibe \(2021\)](#). In particular, [Pasten et al. \(2020\)](#) show that sectoral heterogeneity in price stickiness significantly amplifies the real effects of monetary policy.⁵ We show how heterogeneity in price rigidity amplifies the inflationary effects of a reallocation of demand from services to goods due to the fact that services-producing sectors have stickier prices than goods-producing sectors on average.⁶ In addition, we use the COVID-19 period to estimate production function elasticities in a multi-sector model featuring input-output linkages, and find values broadly similar to those in [Atalay \(2017\)](#) despite markedly different estimation strategies.

Our model also relates to the literature documenting and estimating asymmetric labor adjustment costs at the firm level. [Ilut et al. \(2018\)](#) provide empirical evidence on the response of firms and industries to idiosyncratic shocks and find that the response of employment to positive shocks is only around 50–70 percent as large as that to negative shocks of the same size. The estimated hiring costs in our model provide asymmetric employment responses that are within this range.

In using a model of production networks to understand developments since the COVID-19 pandemic, our paper also builds on [Baqae and Farhi \(2022\)](#). While their quantitative application studies the initial lockdown phase of the pandemic, our focus is on post-lockdown dynamics, particularly on the surge in inflation that occurred in 2021. Another key difference is that they study a two-period model with no factor adjustment across sectors. In comparison, we estimate the factor adjustment costs in an infinite-horizon economy. Using this framework, we are able to study how expectations about the persistence of shocks affect labor reallocation and inflation.

Recent papers have considered the implications of a demand reallocation shock such as the one that is central to our analysis. [Guerrieri et al. \(2021\)](#) and [Fornaro and Romei \(2022\)](#) study the optimal response of monetary policy to a demand reallocation shock in sticky-wage models with two periods and two sectors. Our focus is on quantifying the contribution of the demand reallocation shock to inflation, and on contrasting the reallocation shock with other competing shocks. In related, contemporaneous work, [Anzoategui et al. \(2022\)](#) show how the effects of a demand reallocation shock depend on potentially binding capacity constraints, both domestic and foreign, and [di Giovanni et al. \(2022\)](#) use a two-period model to quantify the contributions of different shocks to the run-up in inflation in the post-lockdown period. In their two-period model with no labor adjustment across sectors, demand reallocation shocks only cause inflation in the presence of downward nominal wage rigidity. In contrast, we study an infinite-horizon model without wage rigidity where demand reallocation shocks are inflationary due to costs of reallocating labor across sectors, which we estimate using aggregate and cross-sectional data. Like [di Giovanni et al. \(2022\)](#), we also find that sectoral supply shocks explain little of the increase in U.S. inflation. However, while they attribute the rise in inflation to an aggregate demand shock, we find that the reallocation of demand from services to goods is the key driver of inflation dynamics.⁷

⁵ [Pasten et al. \(2021\)](#), [Smets et al. \(2019\)](#) and [Ruge-Mucia and Wolman \(2022\)](#) also study the effects of sectoral shocks in multi-sector New Keynesian models in the presence of heterogeneity in price stickiness.

⁶ [Galesi and Rachedi \(2019\)](#) show that the long-run shift from goods to services has important implications for the transmission of monetary policy.

⁷ Our model abstracts from any aggregate demand effect associated with fiscal or transfer policies. However, it is possible that the fiscal stimulus measures enacted during the COVID-19 pandemic may have affected the demand for goods relative to services. For example, the peak month—March 2021—for the goods share of PCE expenditures during the pandemic period coincides with the timing of the largest Economic Impact Payments. [de Soyres et al. \(2022\)](#) provide empirical evidence for this channel.

2. Model

This section describes a multi-sector New Keynesian model featuring sticky prices and input-output linkages. Time is discrete and infinite. The economy consists of K sectors. The model contains two frictions: costs to adjusting prices and costs to reallocating labor across sectors. In order to incorporate these frictions, we assume that in each sector $i = \{1, \dots, K\}$ there are three types of firms: a representative competitive producer, monopolistically competitive firms, and labor agencies. In each sector, the representative competitive producer aggregates the output of a continuum of monopolistically competitive firms. These firms use labor and intermediate inputs to produce their differentiated products, and set prices subject to quadratic adjustment costs. Sector-specific labor is supplied to these firms by agencies that hire labor from a representative household and face convex hiring costs.

Below we describe the problem faced by each type of firm before turning to the problem of the representative household. We then set out the central bank's monetary policy rule and the model's market clearing conditions.

2.1. Representative competitive producer

In each sector i , a representative competitive producer aggregates the output of a continuum of monopolistically competitive firms (indexed by s):

$$Y_t^i = \left[\int_0^1 Y_t^i(s)^{\frac{\epsilon-1}{\epsilon}} ds \right]^{\frac{\epsilon}{\epsilon-1}}, \quad (1)$$

where ϵ is the elasticity of substitution across varieties within a sector. The solution to the competitive producer's problem implies the following demand curve for differentiated products in each sector:

$$Y_t^i(s) = \left(\frac{P_t^i(s)}{P_t^i} \right)^{-\epsilon} Y_t^i. \quad (2)$$

2.2. Monopolistically competitive firms

In each sector, a continuum of firms supply differentiated products to the representative competitive producer subject to price adjustment costs. These differentiated products are produced according to the following production function:

$$Y_t^i(s) = A_t^i \left(\alpha_i^{\frac{1}{\epsilon_Y}} (M_t^i(s))^{\frac{\epsilon_Y-1}{\epsilon_Y}} + (1 - \alpha_i)^{\frac{1}{\epsilon_Y}} (L_t^i(s))^{\frac{\epsilon_Y-1}{\epsilon_Y}} \right)^{\frac{\epsilon_Y}{\epsilon_Y-1}}, \quad (3)$$

where ϵ_Y denotes the elasticity of substitution between labor and intermediate inputs. In order to study sectoral productivity shocks, we allow productivity in each sector, A_t^i , to vary over time. $L_t^i(s)$ denotes labor hired by firm s in sector i at time t . Intermediate inputs, $M_t^i(s)$, are a CES bundle of the outputs of the K sectors of the economy:

$$M_t^i(s) = \left(\sum_{j=1}^K \Gamma_{i,j}^{\frac{1}{\epsilon_M}} (M_{j,t}^i(s))^{\frac{\epsilon_M-1}{\epsilon_M}} \right)^{\frac{\epsilon_M}{\epsilon_M-1}}, \quad (4)$$

where ϵ_M is the elasticity of substitution among the different inputs in each sectors intermediate inputs bundle. The economy's input-output matrix is encoded in the parameters $\Gamma_{i,j}$ (where $\sum_{j=1}^K \Gamma_{i,j} = 1$), which determine the importance of the output of sector j as an input of production in sector i . The problem of a monopolistically competitive firm can be split into two stages: a cost minimization problem and a price-setting problem.

2.2.1. Cost minimization

Given the CES aggregator in Eq. (4), the cost minimization problem implies the following price index for intermediate inputs:

$$P_t^{M,i} = \left(\sum_{j=1}^K \Gamma_{i,j} (P_t^j)^{1-\epsilon_M} \right)^{\frac{1}{1-\epsilon_M}}. \quad (5)$$

Given this price index for intermediate inputs, $P_t^{M,i}$, and a price of labor in sector i , $P_t^{L,i}$, the marginal cost of production in sector i is:

$$MC_t^i = \frac{1}{A_t^i} \left(\alpha_i (P_t^{M,i})^{1-\epsilon_Y} + (1 - \alpha_i) (P_t^{L,i})^{1-\epsilon_Y} \right)^{\frac{1}{1-\epsilon_Y}}. \quad (6)$$

2.2.2. Price setting

Given the marginal cost just derived, firms set prices subject to non-pecuniary, quadratic adjustment costs. The recursive form of their problem is:

$$V_t^i(P_{t-1}^i(s)) = \max_{\frac{P_t^i(s)}{P_t^i}} \left(\frac{P_t^i(s)}{P_t^i} \right)^{-\epsilon} \frac{Y_t^i}{P_t^i} (P_t^i(s) - MC_t^i) - \frac{\kappa_i}{2} \left(\frac{P_t^i(s)}{P_{t-1}^i} - 1 \right)^2 P_t^i \frac{Y_t^i}{P_t^i} + E_t [\mathcal{M}_{t+1} V_{t+1}^i(P_t^i(s))], \quad (7)$$

where κ_i is the sector-specific price adjustment cost, and \mathcal{M}_{t+1} is the stochastic discount factor of the representative household. The solution to the price setting problem is the following sector-level New Keynesian Phillips curve:

$$1 - \epsilon + \epsilon \frac{MC_t^i}{P_t^i} - \kappa_i (\Pi_t^i - 1) \Pi_t^i + \kappa_i E_t \left(\mathcal{M}_{t+1} \frac{(\Pi_{t+1}^i)^2}{\Pi_{t+1}^i} (\Pi_{t+1}^i - 1) \frac{Y_{t+1}^i}{Y_t^i} \right) = 0, \quad (8)$$

where $\Pi_t^i = \frac{P_t^i}{P_{t-1}^i}$ denote the gross inflation rate at the sector level and $\Pi_t = \frac{P_t}{P_{t-1}}$ denotes the aggregate inflation rate.

2.2.3. Labor agencies

In each sector, labor is supplied to the monopolistically competitive firms by a representative labor agency that hires labor from the representative household. We assume that these agencies face convex hiring costs denoted in units of labor, the size of which is key to our results and which we estimate in Section 3.⁸ In contrast, agencies are able to freely decrease employment in each sector. The recursive form of the labor agency's problem is

$$V_t^i(L_{t-1}^i) = \max_{L_t^i} \frac{P_t^{L,i}}{P_t} L_t^i - \frac{W_t}{P_t} L_t^i \left(1 + \mathbb{1}_{L_t^i > L_{t-1}^i} \frac{c}{2} \left(\frac{L_t^i}{L_{t-1}^i} - 1 \right)^2 \right) + E_t [\mathcal{M}_{t+1} V_{t+1}^i(L_t^i)], \quad (9)$$

where c is the hiring cost and $\mathbb{1}_{L_t^i > L_{t-1}^i}$ is a function indicating positive hiring. The solution to this problem is the following dynamic equation for sectoral labor demand:

$$\frac{P_t^{L,i}}{P_t} = \frac{W_t}{P_t} + \mathbb{1}_{L_t^i > L_{t-1}^i} \frac{W_t}{P_t} \left(\frac{c}{2} \left(\frac{L_t^i}{L_{t-1}^i} - 1 \right)^2 + c \left(\frac{L_t^i}{L_{t-1}^i} - 1 \right) \frac{L_t^i}{L_{t-1}^i} \right) - E_t \left(\mathbb{1}_{L_{t+1}^i > L_t^i} \mathcal{M}_{t+1} c \frac{W_{t+1}}{P_{t+1}} \left(\frac{L_{t+1}^i}{L_t^i} - 1 \right) \left(\frac{L_{t+1}^i}{L_t^i} \right)^2 \right). \quad (10)$$

This equation shows how current or future expected hiring costs introduce a wedge between the aggregate wage and the price of labor in each sector. Such a wedge generates flow dividends that are distributed to the household.⁹

2.3. Households

A representative household consumes a bundle of goods and of services:

$$C_t = \left(\frac{C_t^g}{\omega_t} \right)^{\omega_t} \left(\frac{C_t^s}{1 - \omega_t} \right)^{1 - \omega_t}. \quad (11)$$

We allow the preference parameter for goods, ω_t , to vary over time. The solution to the household's cost minimization problem implies:

$$P_t^g C_t^g = \omega_t P_t C_t, \quad (12)$$

$$P_t = (P_t^g)^{\omega_t} (P_t^s)^{1 - \omega_t}. \quad (13)$$

Equation (12) implies that ω_t equals the expenditure share on goods. Figure 1 shows that ω_t rose from 0.31 before the pandemic to above 0.35 in early 2021. Thus this is the size of the shift in ω_t that we will study in Section 4.1.

⁸ Our formulation echoes the literature studying convex hiring costs in models of the labor market, such as Merz and Yashiv (2007) and Gertler and Trigari (2009).

⁹ A common way of introducing frictions to labor mobility assumes that the disutility of labor supply depends both on the aggregate quantity of labor supplied and its composition across sectors, as in Horvath (2000) and Bouakez et al. (2020). Such a formulation does not lend itself to studying questions such as how the reallocation of labor depends on the expected persistence of shocks.

Goods consumption and services consumption are both bundles of the consumption of output from each of the K sectors:

$$C_t^g = \prod_{i=1}^K \left(\frac{C_{i,t}}{\gamma_i^g} \right)^{\gamma_i^g}, \quad (14)$$

$$C_t^s = \prod_{i=1}^K \left(\frac{C_{i,t}}{\gamma_i^s} \right)^{\gamma_i^s}. \quad (15)$$

where $\sum_{i=1}^K \gamma_i^g = 1$ and $\sum_{i=1}^K \gamma_i^s = 1$. Again, the solution to the cost-minimization problem implies:

$$P_t^g = \prod_{i=1}^K (P_t^i)^{\gamma_i^g}, \quad (16)$$

$$P_t^s = \prod_{i=1}^K (P_t^i)^{\gamma_i^s}. \quad (17)$$

Turning to the household's dynamic problem, the household has preferences over total consumption, C_t , and hours worked, N_t :

$$U_t = \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\gamma}}{1-\gamma} - \chi_t \frac{N_t^{1+\psi}}{1+\psi} \right). \quad (18)$$

To incorporate a labor supply shock, we allow the disutility of labor supply, χ_t , to vary over time around a steady-state value $\bar{\chi}$. The representative household maximizes utility subject to the nominal budget constraint:

$$P_t C_t + B_{t+1} = W_t N_t + (1 + i_{t-1}) B_t + div_t, \quad (19)$$

where div_t denotes profits from monopolistically competitive firms and labor agencies and B_t are nominal bondholdings (paying interest rate i_{t-1}). The solution of the household's problem gives the following first-order conditions:

$$1 = E_t \left(\mathcal{M}_{t+1} \frac{1 + i_t}{\Pi_{t+1}} \right), \quad (20)$$

$$C_t^{-\gamma} \frac{W_t}{P_t} = \chi_t N_t^\psi, \quad (21)$$

where the stochastic discount factor is $\mathcal{M}_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}$.

2.4. Monetary policy and market clearing

Monetary policy follows a Taylor rule which responds only to aggregate inflation:

$$\log(1 + i_t) = \log \frac{1}{\beta} + \phi \log \Pi_t. \quad (22)$$

The model's market clearing conditions are as follows. First, the markets for sectoral output clear when:

$$Y_t^i = C_{i,t} + \sum_{j=1}^K M_{i,t}^j \quad \forall i. \quad (23)$$

Second, the aggregate labor market clearing condition is:

$$\sum_{i=1}^K L_t^i \left(1 + \mathbb{1}_{L_t^i > L_{t-1}^i} \frac{c}{2} \left(\frac{L_t^i}{L_{t-1}^i} - 1 \right)^2 \right) = N_t. \quad (24)$$

Finally, the bond market clears when:

$$B_{t+1} = 0. \quad (25)$$

3. Taking the model to the data

In order to bring the model to the data, we posit that the U.S. economy has been hit by three distinct shocks during the COVID-19 pandemic. First, a demand reallocation shock—an increase in ω_t . Second, an aggregate labor supply shock—an increase in χ_t . And finally, sectoral productivity shocks—changes in A_t^i across industries. We will show the inclusion of these three shocks allows the model to account for movements in both aggregate and cross-sectional variables in the 2019:Q4–2021:Q4 period. It should be noted that by focusing on the overall changes from the end of 2019 through the end of 2021 we are abstracting from the sharp movements in macroeconomic variables that took place in 2020:Q2, in the most acute phase of the pandemic and the associated lockdown measures.

We assume that these shocks occur simultaneously, and that, following the shocks, the driving terms revert back to their steady-state values following AR(1) processes:

$$\omega_{t+1} = (1 - \rho_\omega)\bar{\omega} + \rho_\omega\omega_t, \quad (26)$$

$$\chi_{t+1} = (1 - \rho_\chi)\bar{\chi} + \rho_\chi\chi_t. \quad (27)$$

$$A_{t+1}^i = (1 - \rho_A) + \rho_A A_t^i. \quad (28)$$

We proceed by externally calibrating a number of the model's parameters, along with the size of the demand reallocation shock and the sectoral productivity shocks. We then estimate: (i) the production function elasticities, (ii) the hiring cost parameter, and (iii) the magnitude of the aggregate labor supply shock. Given the non-linearities inherent in the model—in particular the large sectoral movements induced by idiosyncratic productivity shocks and the asymmetries caused by the labor hiring cost—we estimate these parameters and show impulse response functions for versions of the model that we solve using nonlinear methods.¹⁰

3.1. Calibrated parameters and shocks

We study a 66 sector version of the model. The model's input-output matrix, $\Gamma_{i,j}$, and the shares of intermediates in production, α_i , are calibrated using the BEA's input-output tables. We use the BEA's bridge between PCE categories and NAICS industries to calibrate the sectoral consumption shares γ_i^g and γ_i^s . We label sectors as services-producing if more of their output is directly consumed as services than as goods. This classification leaves us with 32 services-producing sectors, 28 goods-producing sectors, and 6 sectors that produce neither goods nor services, as none of their output is directly consumed.¹¹

We calibrate price adjustment costs at the sectoral level using data from [Pasten et al. \(2020\)](#).¹² We convert the frequency of price adjustment at the sector level from their paper to the value of the Rotemberg cost parameter, κ_i , that implies the same slope of the New Keynesian Phillips curve. A key feature of the price adjustment data is that the prices of industries that produce goods are more flexible than those of industries that produce services.

The top portion of [Table 1](#) details the other externally calibrated parameters. The Frisch inverse labor supply elasticity parameter ψ is set at 1, and the inverse of the intertemporal elasticity of substitution parameter γ is set at 2. We assume a discount factor β of 0.995 and a response coefficient of interest rates to inflation $\phi = 1.5$, consistent with the Taylor principle. The steady-state goods expenditure share $\bar{\omega}$ is set at 0.31 in line with its value in 2019, and the elasticity of substitution ϵ across varieties is 10.

Given the assumption on household preferences, the expenditure share on goods in the model is simply equal to ω_t . We calibrate the size of the demand reallocation shock ($\Delta\omega = 0.045$) to match the peak increase in the goods expenditure share between 2019:Q4 and 2021:Q4. We calibrate the size of the sectoral productivity shocks to match changes in sectoral TFP over the same period, the measurement of which we describe in the supplementary materials. We set $\rho_\omega = 0.975$, to mimic the slow decline in the goods expenditure share following its spike in 2020. We set the persistence of productivity and labor supply shocks to 0.95.

3.2. Estimated parameters and shocks

We estimate the hiring cost c , the elasticity of substitution between intermediate inputs ϵ_M , and the elasticity of substitution between labor and intermediate inputs ϵ_Y . We also estimate the size of the labor supply shock $\Delta\chi$. We group these parameters in the vector θ and estimate them by minimizing the distance between various cross-sectional and aggregate moments from data, and their model counterparts.

¹⁰ We solve the model using the perfect foresight solver in Dynare (version 4.5.6). Such approach has the advantage of capturing the full nonlinear dynamics of the model, albeit at the expense of abstracting from uncertainty. See [Adjemian et al. \(2022\)](#).

¹¹ Few sectors produce both goods and services: only 12 of the 66 sectors have both $\gamma_i^g > 0$ and $\gamma_i^s > 0$.

¹² The use of the PPI data to construct their estimates of the frequency of price adjustment at the sector level is discussed in more detail in [Gorodnichenko and Weber \(2016\)](#).

Table 1
Parameter Values.

| Calibrated Parameters | Symbol | Value/Range | Target/Source |
|--|----------------|---------------------|----------------------------------|
| Inverse Elasticity of Substitution | γ | 2 | Standard |
| Labor Supply Disutility | $\bar{\chi}$ | 1 | Normalization |
| Inverse Labor Supply Elasticity | ψ | 1 | Standard |
| Taylor Rule Coefficient on Inflation | ϕ | 1.5 | Standard |
| Discount Factor | β | 0.995 | Standard |
| Elasticity Across Varieties | ϵ | 10 | Standard |
| Goods Expenditure Share | $\bar{\omega}$ | 0.31 | BEA |
| Intermediate Input Share (Range) | α_i | 0.11 to 0.83 | BEA |
| Price Adjustment Cost (Range) | κ_i | 0.05 to 99.9 | Pasten et al. (2020) |
| Reallocation Shock Persistence | ρ_ω | 0.975 | Goods Expenditure Share |
| Labor Supply Shock Persistence | ρ_χ | 0.95 | Standard |
| Sectoral TFP Shock Persistence | ρ_A | 0.95 | Standard |
| Size of Reallocation Shock | $\Delta\omega$ | 0.045 | Δ Goods Expenditure Share |
| Sectoral TFP Shocks (Range) | ΔA_i^t | -0.29 to 0.25 | Measured Sectoral TFP |
| Estimated Parameters | Symbol | Value (s.e.) | Target/Source |
| Hiring Cost | c | 19.1 (12.6) | Estimated |
| Elasticity Across Intermediates | ϵ_M | 0.13 (0.24) | Estimated |
| Elasticity Between Intermediates & Labor | ϵ_Y | 0.82 (0.08) | Estimated |
| Labor Supply Shock Size | $\Delta\chi$ | 0.09 (0.04) | Estimated |

The top panel shows parameters that we calibrate externally. The bottom panel shows parameters that we estimate as described in Section 3.2. For the intermediate input share, price adjustment cost, and sectoral TFP shocks we report the range across industries. Industries with lowest and highest values of α_i are “Housing” and “Funds, Trusts, and Other Financial Vehicles,” respectively. Industries with lowest and highest values of κ_i are “Oil and Gas Extraction” and “Legal Services,” respectively.

Our cross-sectional moments are based on industry output, inflation, and employment developments. For each of the 66 sectors, we calculate the percent change in gross output between 2019:Q4 and 2021:Q4 relative to a sector-specific trend.¹³ We repeat the same procedure for price indexes and employment and stack these cross-sectional changes in three vectors: $\mathbf{y}_d, \mathbf{p}_d, \mathbf{l}_d$.

We also target two aggregate moments, both shown in Fig. 1. Between 2019:Q4 and 2021:Q4, goods inflation rose by 6 percentage points, whereas services inflation rose by 1 percentage point. We target the differential rise in the two inflation rates and set $\Delta\pi_d^G - \Delta\pi_d^S = 5\%$. Second, we target the change in total employment. Employment declined 4 percent relative to trend between 2019:Q4 and 2021:Q4, so that $\Delta L_d = -4\%$. The estimated parameters solve the following problem:

$$\theta = \arg \min_{\theta} [\psi(\theta)]' \mathbf{W} [\psi(\theta)], \quad (29)$$

where:

$$\psi(\theta) = \begin{bmatrix} \sigma(\mathbf{y}_d^g) - \sigma(\mathbf{y}_m^g(\theta)) \\ \sigma(\mathbf{p}_d^g) - \sigma(\mathbf{p}_m^g(\theta)) \\ \sigma(\mathbf{l}_d^g) - \sigma(\mathbf{l}_m^g(\theta)) \\ \sigma(\mathbf{y}_d^s) - \sigma(\mathbf{y}_m^s(\theta)) \\ \sigma(\mathbf{p}_d^s) - \sigma(\mathbf{p}_m^s(\theta)) \\ \sigma(\mathbf{l}_d^s) - \sigma(\mathbf{l}_m^s(\theta)) \\ \rho(\mathbf{y}_d, \mathbf{y}_m(\theta)) \\ \rho(\mathbf{p}_d, \mathbf{p}_m(\theta)) \\ \rho(\mathbf{l}_d, \mathbf{l}_m(\theta)) \\ \Delta L_d - \Delta L_m(\theta) \\ (\Delta\pi_d^G - \Delta\pi_d^S) - (\Delta\pi_m^G(\theta) - \Delta\pi_m^S(\theta)) \end{bmatrix}. \quad (30)$$

In the equation above, $\sigma(\mathbf{y}_d^g)$, for instance, denotes the cross-sectional standard deviation of the percent change in output for goods-producing sectors between 2019:Q4 and 2021:Q4, and $\sigma(\mathbf{y}_m^g(\theta))$ denotes the model counterpart. By the same token, $\rho(\mathbf{y}_d, \mathbf{y}_m(\theta))$ denotes the correlation between industry changes in output and the corresponding model objects, which we calculate one year after the shocks occur. We construct measures of dispersion separately for goods-producing and services-producing sectors as there is significant heterogeneity in the data: goods prices are much more dispersed than services prices, whereas the opposite is true for labor. This is informative for our estimation procedure. Finally, \mathbf{W} is a weighting matrix: we use the identity matrix, implying that all moments have equal weight.¹⁴

¹³ We calculate the trend over the 2005–2019 period, as 2005 is the first year for which BEA produces quarterly GDP-by-industry data.

¹⁴ We calculate each of the standard deviations weighting by sectoral gross output.

Before turning to the parameter estimates, we discuss the relationship between these moments and the parameters that we estimate. There is clearly a direct link between the size of the labor supply shock and the decline in aggregate employment. The size of the hiring cost is closely related to difference in goods and services price inflation. As we will show in the next section, with no hiring costs there would be no change in relative prices in response to a demand reallocation shock. On the other hand, if hiring is costly, goods production will increase more slowly, and the relative price of goods will rise. Finally, the production function elasticities are important in determining how each of the shocks that hit the model propagate through the production network. The parameters ϵ_Y and ϵ_M also affect how stringent hiring costs are, since a high elasticity of substitution would imply that firms can avoid labor costs by using intermediate inputs. Hence, c , ϵ_M and ϵ_Y jointly affect the sectoral dynamics of output, prices and labor, and the cross-sectional moments from the data help us discipline these parameters.

The estimated parameters are reported in the bottom portion of [Table 1](#). The production function elasticities are in line with the values estimated using very different approaches (e.g. [Atalay, 2017](#)). As will be discussed in [Section 4.3](#), we find an important role for the aggregate labor supply shock in accounting for the aggregate decline in employment. The hiring costs that we estimate are relatively modest: for example, these imply that the labor agency would need to pay hiring costs of around 0.2% of its payroll in order to increase employment by 1% in a given quarter. In practice these costs are small in aggregate: when we subject the model to all shocks, the total hiring costs paid are equal to 0.15% of output in the period when the shocks occur, 0.08% of output in the next quarter, and quickly converge to zero thereafter. We discuss the robustness of the estimation strategy in [Section 6](#).

4. Results

With the estimated parameters in hand, we now consider the role of each shock individually, before simulating the model with all three shocks turned on.

4.1. The COVID-19 demand reallocation shock

First, we turn off the aggregate labor supply shock ($\Delta\chi = 0$) and the sectoral TFP shocks ($\Delta A_i^j = 0 \forall i$), and we consider our main experiment, which looks at the effect of an increase in demand for goods relative to services. In order to highlight important features of the model, we contrast the effect of this shock in the baseline model with that which would occur: (i) if there were no labor adjustment costs, and (ii) if price stickiness were homogeneous across sectors.

[Figure 3](#) undertakes the first comparison and plots the response of key variables to the demand reallocation shock. The reallocation of demand leads to a large increase in goods consumption and a corresponding decline in services consumption. The dotted lines show that, absent hiring costs, these changes would offset each other leaving aggregate prices, consumption and employment unchanged. Once we introduce hiring costs, the increase in employment in goods-producing industries is much slower, constraining goods supply and resulting in a smaller increase in goods consumption compared with the frictionless model. As a consequence of the costs of increasing production, goods prices jump, resulting in year-over-year goods inflation peaking around 6 percent after one year.

In contrast, employment in services-producing sectors falls immediately, as such firms face no costs in reducing their workforce.¹⁵ The asymmetry caused by hiring costs is key in understanding the inflationary effects of this shock: in services-producing sectors, the decline in demand translates largely into a fall in quantities rather than prices. In contrast, in goods-producing sectors the increase in demand pushes up prices due to the costs firms face in increasing their capacity. While services inflation initially declines, it then also rises, peaking around 3 percent after 5 quarters. Taken together, the dynamics of sectoral inflation result in aggregate inflation peaking at 3.5 percent after one year, which represents a sizeable portion of the increase in aggregate inflation shown in [Fig. 1](#). The demand reallocation shock can also explain a roughly 1.5 percent decline in both aggregate consumption and employment in the baseline model.

In [Fig. 4](#) we repeat the experiment but assuming that all sectors have the same price stickiness (equal to the sector-weighted average stickiness in our baseline calibration). As goods prices tend to be more flexible than services prices, this assumption raises price stickiness in goods-producing sectors and lowers it in services-producing sectors, on average. Higher price stickiness in the goods sectors results in a lower path for goods inflation, causing a peak aggregate inflation 0.8 percentage points lower than in our baseline. Hence, heterogeneous price stickiness is an important element to explain the inflationary effects of the demand reallocation shock.

Despite the simplicity of the demand reallocation shock, the model contains rich predictions on the dynamics of sectoral prices and quantities. [Figure 5](#) shows that this relative demand shock is able to explain a good fraction of the dispersion in industry-level inflation rates and output growth. The positive correlation between inflation in the model and the data holds not only across all sectors but also within the sets of goods-producing or services-producing sectors. Both the input-output structure in the model and heterogeneity in price rigidity across sectors are important for this result, as we show in the more detail in the supplementary material. For example, despite the negative shock to final demand for services, prices and

¹⁵ Despite the absence of costs to cutting employment, labor in service sectors declines less than in the frictionless model as firms internalize the prospect of future hiring costs.

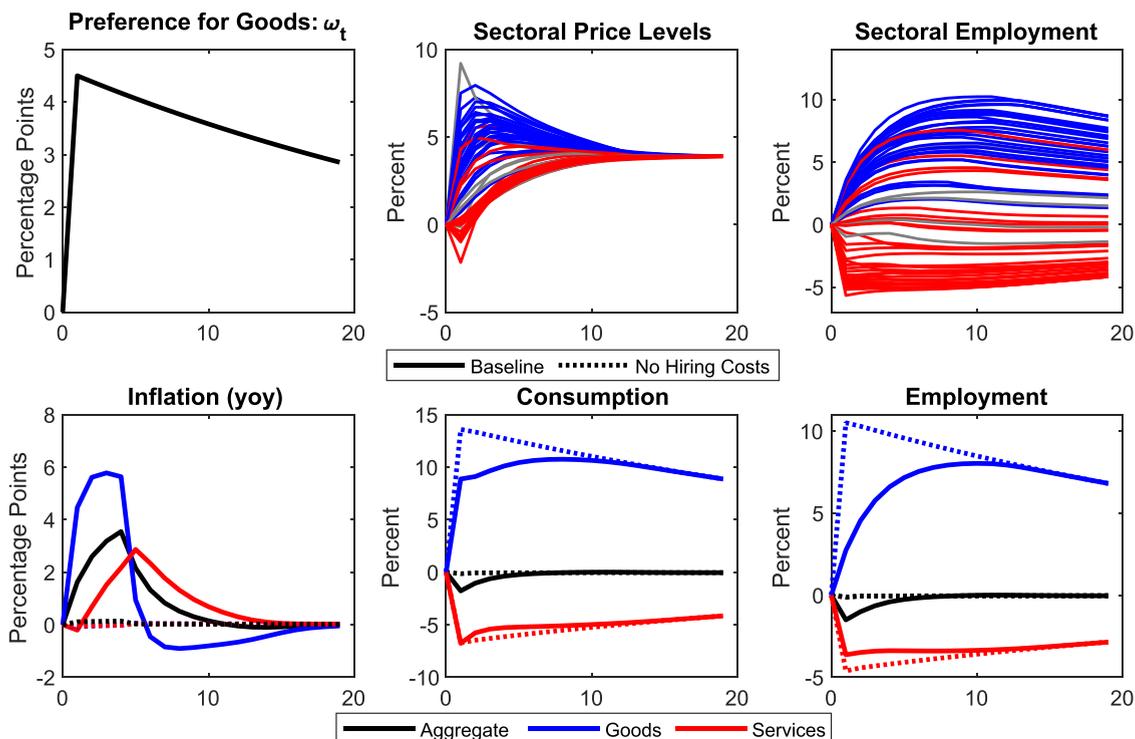


Fig. 3. Aggregate Effects of the Demand Reallocation Shock. This figure plots the impulse response of key variables to the demand reallocation shock that increases the value of the preference parameter for goods (ω_t) in period 1. Each period is one quarter. Solid lines denote the baseline model. Dotted lines denote the response of aggregates if there were no hiring costs. For clarity, we only plot sectoral variables in the baseline model. Gray lines denote sectors (“other” sectors) for which no output is directly consumed.

quantities rise in a number of services sectors, such as the warehousing sector, which are heavily used as intermediates for goods production.

4.2. Sectoral productivity shocks

There are a number of sectors for which price and quantity dynamics are harder to reconcile solely with the dynamics following an aggregate reallocation shock. One striking example is the “Motor Vehicle Parts and Dealer” sector, which has experienced a 40% decline in quantities and a 50% rise in prices between 2019:Q4 and 2021:Q4. Such evidence is suggestive of the importance of pandemic-related supply distributions in some sectors.¹⁶

To understand the importance of such disruptions, we now consider in isolation the role of sectoral productivity shocks. By linking industry data on employment from the BLS with data on output and intermediate inputs from the BEA, we measure the evolution of total factor productivity at the industry level between 2019:Q4 and 2021:Q4 and feed the estimated sectoral component of the productivity series into the model. Details of our measurement of sectoral TFP are provided in the supplementary material, where we show that sectoral TFP shocks can explain a significant fraction of the cross-sectional evolution of both prices and quantities. However, their effect on aggregate inflation is actually slightly negative. This occurs as sectoral TFP growth was above trend, on average, between 2019:Q4 and 2021:Q4.¹⁷

4.3. Labor supply shock

While the demand reallocation and sectoral productivity shocks explain a significant fraction of both sectoral and aggregate price and quantity dynamics, together they explain less than half of the decline in employment experienced in the United States. This is the motivation for introducing a negative shock to labor supply in our estimation exercise. As in a standard New Keynesian model, such a shock lowers employment and consumption, while putting upward pressure on wages

¹⁶ Our closed-economy model abstracts from disruptions to global supply chains, although such disruptions may indirectly show up as negative domestic sectoral productivity shocks.

¹⁷ Our estimates of industry productivity dynamics are close to those of Fernald and Li (2022). We plot the estimated productivity shocks by sector in the supplementary material.

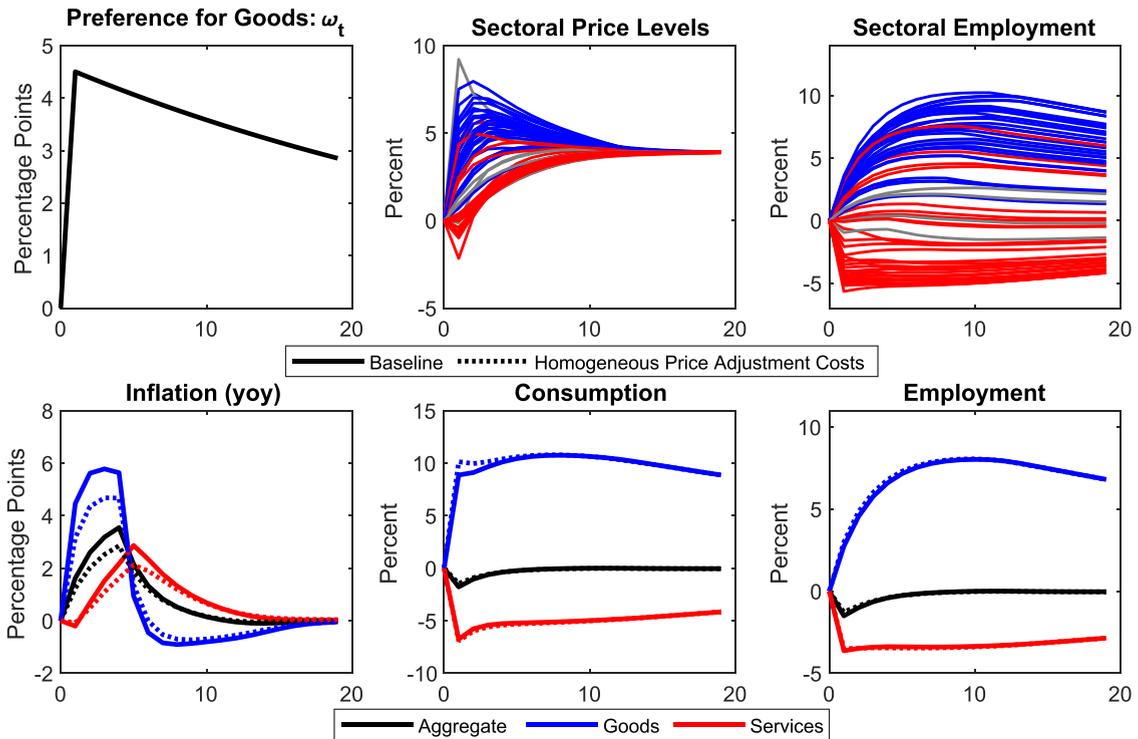


Fig. 4. Demand Reallocation Shock: Heterogeneous vs Homogeneous Price Stickiness. This figure plots the impulse response of key variables to the demand reallocation shock that increases the value of the preference parameter for goods (ω_t) in period 1. Each period is one quarter. Solid lines denote the baseline model. Dotted lines denote the response of variables if price adjustment costs were homogeneous across industries. For clarity, we only plot sectoral variables in the baseline model. Gray lines denote sectors (“other” sectors) for which no output is directly consumed.

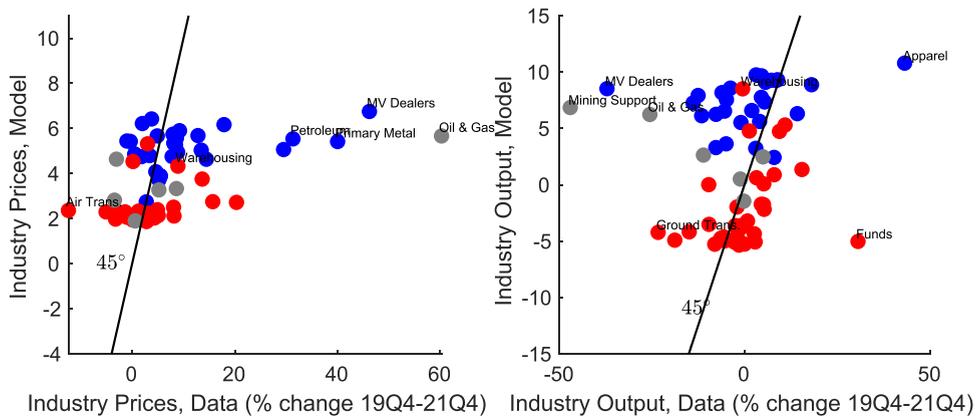


Fig. 5. Model and Data: Sectoral Responses to Demand Reallocation Shock. This figure compares the cross-sectional implication of the model with the data in response to a demand reallocation shock that increases preferences for goods. Each dot is one industry. On the x-axis we plot inflation rates (percent change in the industry chain-type price index) and real gross output growth for the 66 private industries for which BEA publishes GDP-by-industry data. On the y-axis we plot the model counterparts one year after the reallocation shock. Services-producing industries are shown in red and goods-producing industries are shown in blue. Gray dots denote sectors (“other” sectors) for which no output is directly consumed.

and prices. In the supplementary material we show that this shock leads to a rise in inflation of 1.5 percentage points, less than half of that seen in response to the demand reallocation shock.

4.4. All 3 COVID-19 shocks

Having considered the three types of shock in isolation, we now show their effects when they occur simultaneously (as assumed in our estimation procedure). Figure 6 plots the impulse response functions in this case. Overall our model suggests that these shocks are responsible for an increase in inflation of slightly less than 3.5 percentage points, close to that which

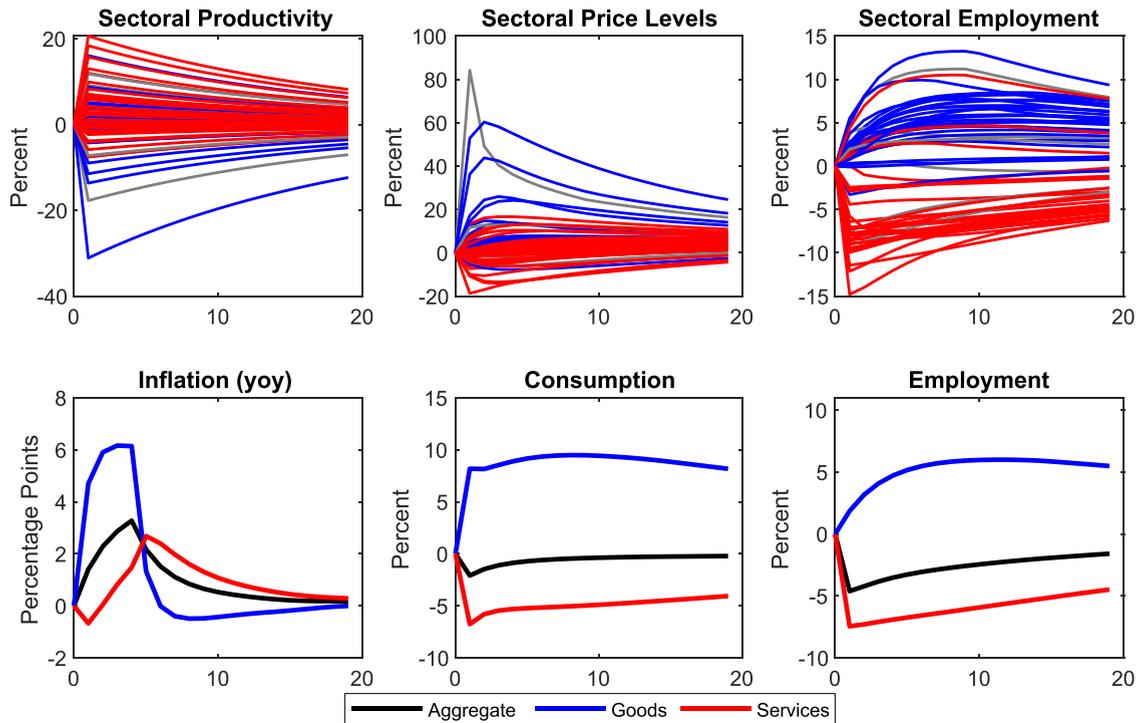


Fig. 6. Aggregate Effects of All Shocks. This figure plots the impulse response of key variables to three combined shocks: (1) a demand reallocation shock that increases preferences for goods, (2) estimated sectoral TFP shocks, and (3) a negative labor supply shock. Each period is one quarter. Gray lines denote sectors (“other” sectors) for which no output is directly consumed.

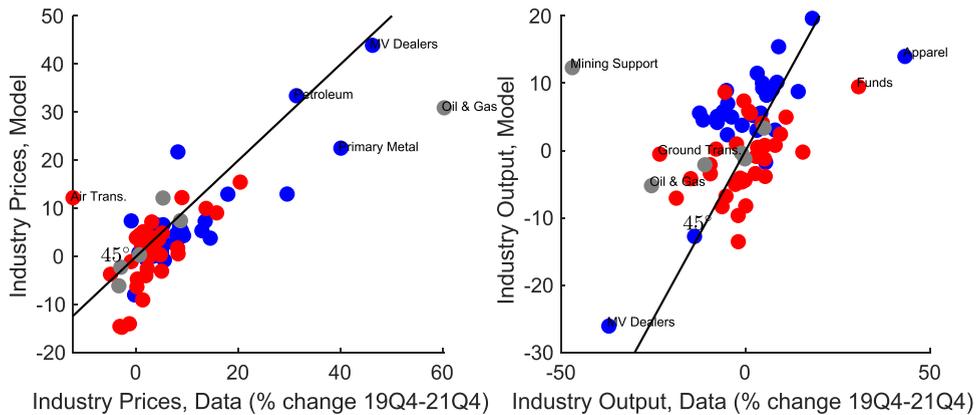


Fig. 7. Model and Data: Sectoral Responses to All Shocks. This figure compares the cross-sectional implication of the model with the data in response to three combined shocks: (1) a demand reallocation shock that increases preferences for goods, (2) estimated sectoral TFP shocks, and (3) a negative labor supply shock. Each dot is one industry. On the x-axis we plot inflation rates (percent change in the industry chain-type price index) and real gross output growth for the 66 private industries for which BEA publishes GDP-by-industry data. On the y-axis we plot the model counterparts one year after the shocks. Services-producing industries are shown in red and goods-producing industries are shown in blue. Gray dots denote sectors (“other” sectors) for which no output is directly consumed.

was observed in the data. Thus, the deflationary effects of the sectoral productivity shocks appear to offset the inflationary effects of the labor supply shock. However, the model exhibits significant non-linearities: summing the inflationary effects of the individual shocks would lead to an increase in inflation around 30 percent larger than seen in Fig. 6. This occurs as the negative labor supply shock reduces the expansion in hiring that occurs in goods-producing sectors in response to the demand reallocation shock, and consequently the run-up in hiring costs that such firms face. In the supplementary material we provide an alternative decomposition based on considering the effect of removing shocks one at a time. Our finding that the demand reallocation shock is the key driver of inflation is robust to this approach.

Turning to the cross-section, Fig. 7 shows that the combination of the three shocks provides an excellent description of cross-sectional developments in prices and quantities. For example, the correlation between sectoral inflation rates in the

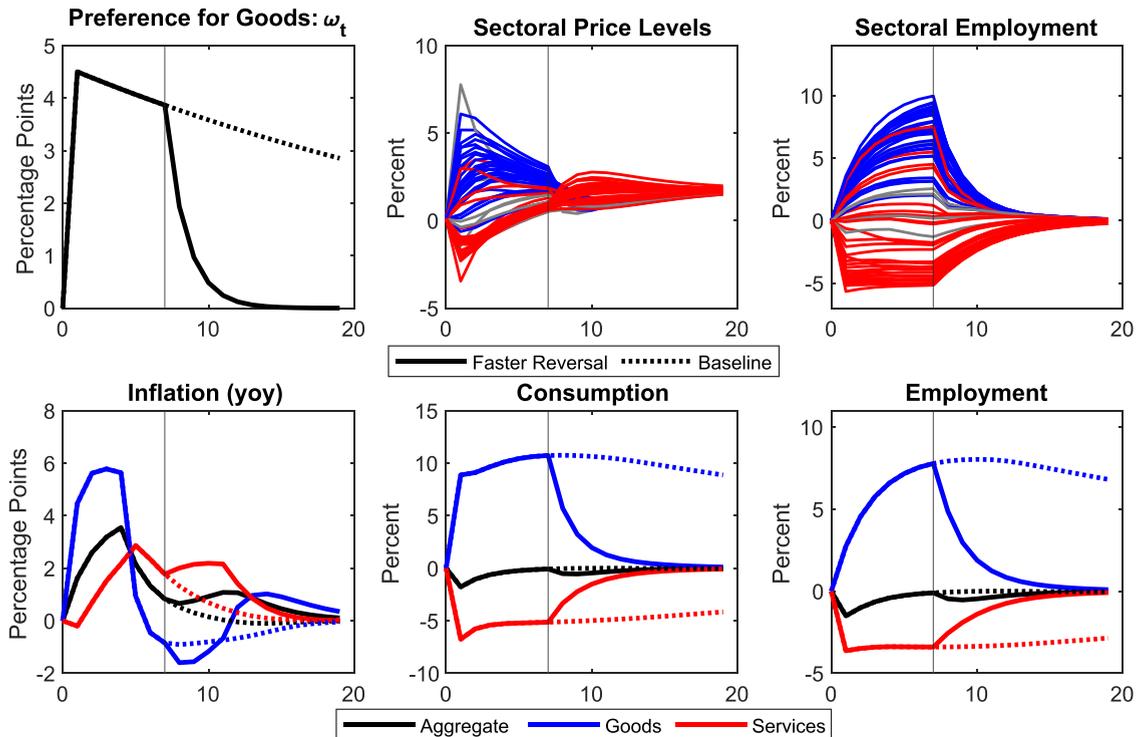


Fig. 8. Aggregate Effects of Reversal of Demand Reallocation Shock. This figure plots the impulse response of key variables to the demand reallocation shock that increases the value of the preference parameter for goods (ω_t) in period 1. The solid lines show outcomes if the persistence unexpectedly declines from 0.95 to 0.5 after two years (denoted by the vertical line). The dotted lines shows the baseline persistence. Each period is one quarter. Gray lines denote sectors (“other” sectors) for which no output is directly consumed.

model and the data is 0.80. Even if one is only interested in aggregate developments, we consider this to be strong evidence in favor of the channels in this paper.

5. Model extensions

In this section we undertake a number of extensions. First, we consider the implications of the demand reallocation shock under different assumptions about its persistence and how persistent it was expected to be. Next, we consider some out-of-sample experiments: we study the demand reallocation that occurred around the time of the Great Recession and we finish by estimating the effect of sectoral TFP shocks that occurred during the first half of 2022.

5.1. A reversal of the COVID-19 demand reallocation shock

What would happen to inflation if demand shifts away from goods back to services faster than anticipated? To consider such hypothesis, we perform the following exercise. Initially, the economy is hit by the baseline reallocation shock from services to goods studied in the previous section. After eight quarters, the economy is hit by an unexpected reversal in demand from goods back to services. We model such a reversal by assuming that the persistence of the baseline shock unexpectedly drops from 0.975 to 0.5 in period 8.

Figure 8 compares outcomes in this reversal experiment with those that occur in the baseline experiment when the demand reallocation shock is highly persistent. We find that such a reversal would raise inflation by around a percentage point relative to the no-reversal baseline. In particular, the reversal leads to renewed inflationary pressures, primarily driven by services-producing sectors which struggle to increase capacity in response to their unexpectedly fast increase in demand.

5.2. Unexpected persistence of the COVID-19 demand reallocation shock

Our baseline experiment assumes that the agents are immediately aware of the persistence of the demand reallocation shock. An alternative hypothesis is that the persistence of the shift in demand from services to goods turned out to be higher than initially anticipated. To investigate this, we now consider a demand reallocation shock that is “unexpectedly” persistent. In particular, we assume that agents initially believe that the shock has a quarterly persistence of 0.5, even though

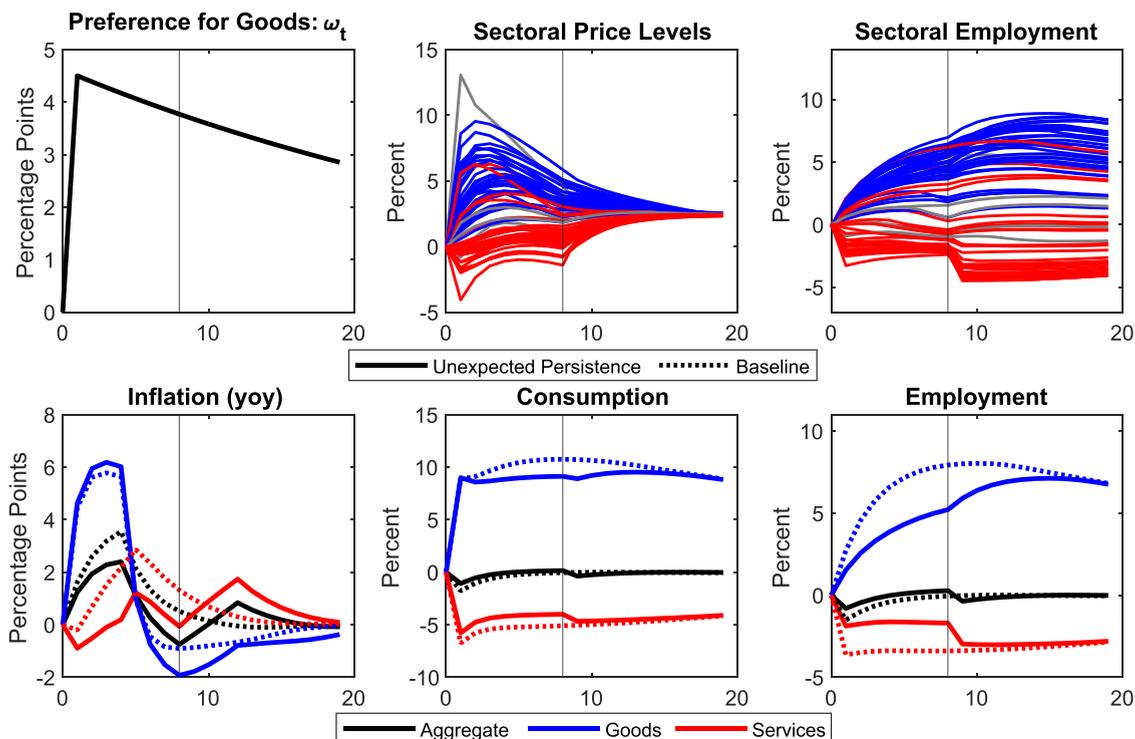


Fig. 9. Aggregate Effects of Unexpected Persistence of Demand Reallocation Shock. This figure plots the impulse response of key variables to the demand reallocation shock that increases the value of the preference parameter for goods (ω_t) in period 1. The solid lines show outcomes if agents expect the shock to have a lower persistence of 0.5 for the first eight quarters and thus are repeatedly surprised about its persistence. After eight quarters (denoted by the vertical line) agents learn the true persistence. The dotted lines show outcomes if the persistence is known immediately. Each period is one quarter. Gray lines denote sectors (“other” sectors) for which no output is directly consumed.

the relative demand for goods, ω_t , follows the same ex post path as in our baseline experiment. Consequently, for the first two years, agents are repeatedly surprised by the persistence of ω_t . After two years, we assume that agents learn the true persistence of the shock.

Figure 9 plots the response of key variables in our model to such a sequence of shocks. This shows that in such a scenario less labor is shed in services-producing sectors, while fewer employees are hired by goods-producing sectors. An implication of this reduction in reallocation is that price dispersion is higher than in the baseline. In particular, prices in services-producing sectors fall much more than in the baseline, as their decline in demand feeds less into quantities than it does in the baseline.

The bottom-left panel of Fig. 9 shows that the lower services price inflation in this scenario is largely responsible for lower total inflation. Aggregate inflation peaks at around 2.5 percent, as opposed to around 3.5 percent under our baseline assumption on expectations. On the other hand, when agents finally realize the persistence of the shock, there is a second bout of inflation, as services-producing sectors lay off workers and raise prices.

5.3. Demand reallocation during the great recession

Our reallocation shock is inflationary primarily due to asymmetric labor adjustment costs, regardless of whether it shifts demand from services to goods or vice versa. We prove this with an application to the Great Recession, the other recent episode with a large shift in the composition of consumption expenditures. Between 2008:Q2 and 2009:Q1, the goods expenditure share fell from 34 to 31.8 percent. We model such a shift as a shock to the relative demand for goods ω_t that is half the size and the opposite sign of our baseline reallocation shock.

Figure 10 shows the effects of this shift in demand, both in our baseline calibration and in a version with homogeneous price adjustment costs. The inflationary effect of the reallocation shock during the Great Recession is proportionally smaller than in our baseline experiment, with inflation peaking at 1.4 percent. The dampened effect is explained by the heterogeneity in price adjustment costs. As goods prices are more flexible on average than those of services, heterogeneity in price stickiness amplifies the effects on inflation of a shift in demand towards goods, but dampens the effects on inflation of a shift in demand towards services. Despite this dampening, our model suggests that demand real-

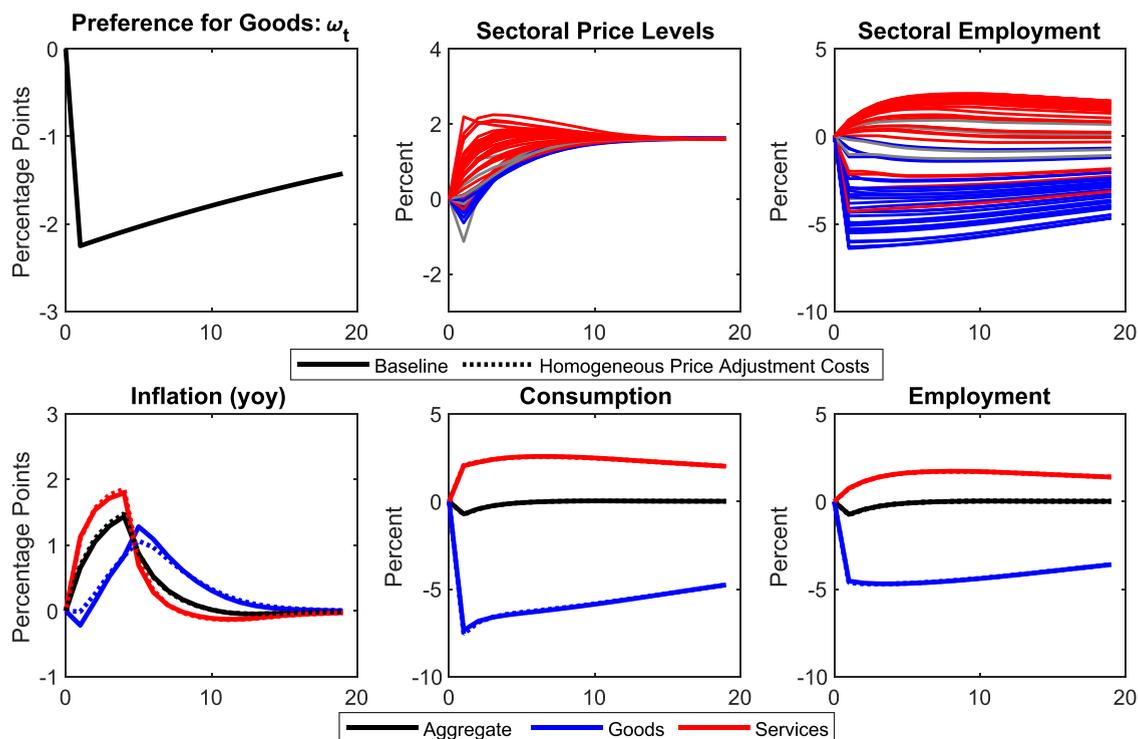


Fig. 10. Demand Reallocation Shock During the Great Recession. This figure plots the impulse response of key variables to a demand reallocation shock that decreases the value of the preference parameter for goods (ω_t) in period 1. Each period is one quarter. Solid lines denote the baseline model. Dotted lines denote the response of variables if price adjustment costs were homogeneous across industries. For clarity, we only plot sectoral variables in the model with heterogeneous price adjustment costs. Gray lines denote sectors (“other” sectors) for which no output is directly consumed.

location during the Great Recession could partly explain the “missing deflation” that has been the focus of a large literature.¹⁸

5.4. Additional productivity shocks during 2022

In Section 4 we considered shocks that occurred between 2019:Q4 and 2021:Q4. Absent further shocks, our model would have predicted that inflation should have declined significantly in 2022, particularly in goods-producing sectors. This is at odds with the data, as inflation remained persistently high during 2022. A number of possible explanations have been proposed for this persistence, such as renewed supply shortages caused by the war in Ukraine and continued lockdowns in China.

To understand the extent to which our model can rationalize these developments, we estimate sectoral TFP shocks from 2021:Q4 to 2022:Q2 and feed these additional shocks into our model one year after the original COVID-19 shocks. While average sectoral TFP growth was positive between 2019:Q4 and 2021:Q4, it turned negative in early 2022, driven by large declines in sectors such as “Oil and Gas” and “Computer and Electronics Products”. In Fig. 11 we show that feeding these additional TFP shocks into the model causes overall inflation to continue to rise for another year, and can help explaining why inflation in goods-producing sectors remained high throughout 2022.

6. Sensitivity analysis

Our finding that the demand reallocation shock was a major cause of the rise in inflation in the post-lockdown period is robust to using alternative model specifications and different estimation strategies. These specifications are described in the supplementary material and briefly listed here.

First, we estimate a version of the model in which we allow for firing as well as hiring costs. Firing costs are estimated to be zero, while other parameters are little affected, lending support to our baseline calibration with asymmetric labor

¹⁸ See Gilchrist et al. (2017) and Harding et al. (2022). It has to be noted that, while during the COVID-19 pandemic the change in the goods expenditure share is likely due to a shift in preferences similar to our demand reallocation shock, the identification of the drivers of the decline in the goods expenditure share during the Great Recession is more tenuous, as the accompanying credit crunch may have simultaneously disrupted both aggregate demand and the goods expenditure share.

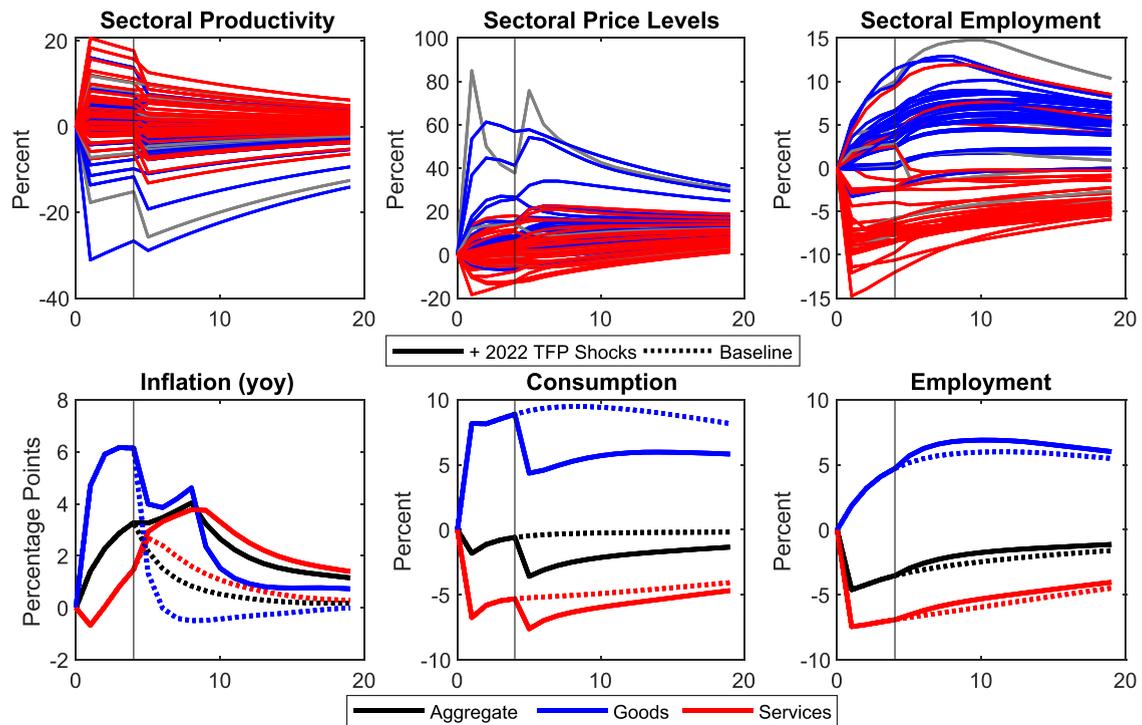


Fig. 11. Aggregate Effects of Additional TFP Shocks in 2022. This figure plots the impulse response of key variables to two sets of shocks. The dotted lines shows the response following the (1) demand reallocation shock, (2) estimated sectoral TFP shocks from 2019:Q4-2021:Q4 and (3) the negative labor supply shock (as in Fig. 6). The solid lines adds the estimated sectoral TFP shocks from 2021:Q4 to 2022Q2 after four quarters (denoted by the vertical line). Each period is one quarter. Gray lines denote sectors (“other” sectors) for which no output is directly consumed.

costs. Next, we show that if we place a much smaller weight on the cross-sectional moments in the estimation procedure we obtain much less precise estimates, supporting our approach of using cross-sectional information to identify the model parameters.

The average price stickiness in our model is roughly in line with a Calvo-style setup in which prices adjust every two quarters, which is lower than the standard price duration that is found in many estimated New Keynesian models. Hence, we re-estimate our model after scaling up the price adjustment costs to mimic an average price duration of four quarters. This alternative estimation produces results broadly in line with our baseline model, with the reallocation shock explaining a good fraction of inflation in the post-Covid period.

In another alternative we restrict production function elasticities, ϵ_M and ϵ_Y , to be equal to 1. As expected, in this case the model fit deteriorates as the model underperforms in matching cross-sectional moments. We also show that our results are robust to using a Taylor rule featuring interest rate smoothing. Finally, we show that the results change only little when we depart from Cobb-Douglas consumption preferences and use instead a more general CES specification.

7. Conclusions

In this paper, we have estimated a multi-sector model with input-output linkages in order to quantify the role that demand reallocation, sector-specific disturbances, and lower aggregate labor supply have played in driving price and quantity dynamics in the U.S. economy in the aftermath of the COVID-19 pandemic.

Our main finding is that the shift in consumption demand from services towards goods can explain a large proportion of the rise in U.S. inflation between 2019:Q4 and 2021:Q4. This demand reallocation shock is inflationary due to the costs of increasing production in goods-producing sectors and because such sectors tend to have more flexible prices than those producing services. The aggregate labor supply shock provides a smaller inflationary impulse, despite the fact that it explains the majority of the decline in employment. The sectoral productivity shocks actually lower inflation slightly, as average productivity grew strongly over this period. Our confidence in the model and its predictions is boosted by the fact that it provides an excellent description of cross-sectional developments in prices and quantities.

We have used the model to conduct a number of experiments relating to the duration and the expected persistence of the demand reallocation shock. We have also shown that the model is able to rationalize the persistence of high inflation during 2022, as many sectors, particularly those producing goods, experienced a decline in productivity in the first half of that year.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.jmoneco.2023.03.003](https://doi.org/10.1016/j.jmoneco.2023.03.003).

References

- Adjemian, S., Bastani, H., Juillard, M., Karamé, F., Mihoubi, F., Mutschler, W., Pfeifer, J., Ratto, M., Rion, N., Villemot, S., 2022. Dynare: Reference Manual Version 5. Dynare Working Papers 72. CEPREMAP.
- Anzoategui, D., Comin, D., Johnson, R., 2022. Supply Chain Constraints and Inflation. Technical Report.
- Atalay, E., 2017. How important are sectoral shocks? *Am. Econ. J. Macroecon.* 9 (4), 254–280.
- Baqee, D., Farhi, E., 2022. Supply and demand in disaggregated Keynesian economies with an application to the Covid-19 crisis. *Am. Econ. Rev.* 112 (5), 1397–1436.
- Bouakez, H., Rachedi, O., Santoro, E., 2020. The government spending multiplier in a multi-sector economy. *Am. Econ. J. Macroecon.*
- de Soyres, F., Santacreu, A.M., Young, H.L., 2022. Demand-Supply Imbalance During the Covid-19 Pandemic: the Role of Fiscal Policy. *International Finance Discussion Paper* 1353.
- di Giovanni, J., Kalemli-Özcan, S., Silva, A., Yildirim, M.A., 2022. Global Supply Chain Pressures, International Trade, and Inflation. Working Paper 30240. National Bureau of Economic Research.
- Fernald, J., Li, H., 2022. The Impact of COVID on Productivity and Potential Output. Technical Report.
- Fornaro, L., Romei, F., 2022. Monetary Policy During Unbalanced Global Recoveries. Technical Report.
- Galesi, A., Rachedi, O., 2019. Services deepening and the transmission of monetary policy. *J. Eur. Econ. Assoc.* 17 (4), 1261–1293.
- Gertler, M., Trigari, A., 2009. Unemployment fluctuations with staggered Nash wage bargaining. *J. Polit. Economy* 117 (1), 38–86.
- Ghassibe, M., 2021. Monetary policy and production networks: an empirical investigation. *J. Monet. Econ.* 119, 21–39.
- Gilchrist, S., Schoenle, R., Sim, J., Zakrajšek, E., 2017. Inflation dynamics during the financial crisis. *Am. Econ. Rev.* 107 (3), 785–823.
- Gorodnichenko, Y., Weber, M., 2016. Are sticky prices costly? evidence from the stock market. *Am. Econ. Rev.* 106 (1), 165–199.
- Guerrieri, V., Lorenzoni, G., Straub, L., Werning, I., 2021. Monetary Policy in Times of Structural Reallocation. Working Paper 2021-111. University of Chicago, Becker Friedman Institute for Economics.
- Harding, M., Lindé, J., Trabandt, M., 2022. Resolving the missing deflation puzzle. *J. Monet. Econ.* 126, 15–34.
- Horvath, M., 2000. Sectoral shocks and aggregate fluctuations. *J. Monet. Econ.* 45 (1), 69–106.
- Illut, C., Kehrig, M., Schneider, M., 2018. Slow to hire, quick to fire: employment dynamics with asymmetric responses to news. *J. Polit. Economy* 126 (5), 2011–2071.
- La'O, J., Tahbaz-Salehi, A., 2022. Optimal monetary policy in production networks. *Econometrica* 90 (3), 1295–1336.
- Merz, M., Yashiv, E., 2007. Labor and the market value of the firm. *Am. Econ. Rev.* 97 (4), 1419–1431.
- Ozdagli, A., Weber, M., 2017. Monetary Policy through Production Networks: Evidence from the Stock Market. Working Paper 23424. National Bureau of Economic Research.
- Pasten, E., Schoenle, R., Weber, M., 2020. The propagation of monetary policy shocks in a heterogeneous production economy. *J. Monet. Econ.* 116, 1–22.
- Pasten, E., Schoenle, R., Weber, M., 2021. Sectoral Heterogeneity in Nominal Price Rigidity and the Origin of Aggregate Fluctuations. Chicago Booth Research Paper 17-25. 2018–54
- Ruge-Mucia, F.J., Wolman, A.L., 2022. Relative Price Shocks and Inflation. Technical Report.
- Smets, F., Tielens, J., Van Hove, J., 2019. Pipeline Pressures and Sectoral Inflation Dynamics. Technical Report.