Structural Transformation and the Transmission of Monetary Policy*

Tiago Bernardino†

IIES, Stockholm University

March 12, 2026

Abstract

This paper studies how the long-term sectoral shift in economic activity toward services has increased the effectiveness of monetary policy in the United States. I study the role of sectoral differences between goods and services in price rigidity and in demand composition across the income distribution—two features I document in the data. I develop a two-sector heterogeneous-agent model with sector-specific nominal rigidities and non-homothetic preferences to quantify their implications for monetary policy transmission. The shift toward services between 1970 and 2019 strengthened the effectiveness of monetary policy by 21%. As services prices are less responsive than goods prices, structural transformation raises aggregate price rigidity, flattens the Phillips curve, and amplifies the short-run transmission of monetary policy. Low-wealth households bear the largest welfare loss after a contractionary shock, with structural transformation amplifying these distributional effects. Finally, while structural transformation amplifies the impact of monetary policy shocks, it dampens the effects of supply shocks by shifting activity toward a less price-volatile sector.

Keywords: Structural Transformation; Monetary Policy; Heterogeneous-Agent Models; Non-homothetic Preferences; Price Rigidities

JEL classification: D11; D21; D31; E21; E23; E32; E52; O41

*I am grateful to my advisor Kurt Mitman for all of his guidance and mentorship. I also thank Timo Boppart, Saman Darougheh, Adrien d'Avernas (discussant), Ricardo Duque Gabriel, Daria Finocchiaro, Brian Higgins, Philipp Hochmuth, Per Krusell, Kieran Larkin, Juan Llavador, Ricardo Marto, Matthias Meier, João Quelhas, Tiago Paúl, Thore Petersen, Mariana Pires, Ricardo Reis, Amalia Repele, Verónica Salazar, Grégoire Sempé, Márcia Silva-Pereira, Leiv Storesletten, Philipp Wangner, Yimei Zou, and participants at the ENTER Jamboree, PhD-EVS, the 18th Annual Meeting of the Portuguese Economic Journal, the Graduate Workshop on Heterogeneous Agent Macroeconomics, SUDSWEC 2025, the Doctoral Workshop on Quantitative Dynamic Economics, the 14th National PhD Workshop in Finance, the 4th Banco de Portugal Winter Research Workshop, and seminars at Banco de España, Banco de Portugal, CBS, IIES, ISEG, UAB, U. Mannheim, and Sveriges Riksbank for valuable comments, discussions, and suggestions.

†Web: tiagobernardino.com; Email: tiago.bernardino@iies.su.se

1 Introduction

Modern economic growth has coincided with a process of structural transformation that has shifted economic activity toward the services sector. In the United States, the share of consumer expenditure allocated to services rose from 50% in 1970 to 68% in 2019, according to the Bureau of Economic Analysis. While this transformation has been widely studied in the context of long-run growth, its implications for short-run economic dynamics, including monetary policy transmission, have received much less attention. Structural transformation is a slow-moving process that spans decades, whereas monetary policy predominantly focuses on stabilizing the economy over months or years. Yet this separation overlooks an important interaction: changes in the composition of demand between goods and services can reshape how the economy responds to interest rate changes, and thereby alter monetary policy transmission over time.

In this paper, I study how structural transformation has changed the transmission of monetary policy in the United States. I document three facts. First, services prices adjust less frequently than goods prices—they exhibit a higher price rigidity. Second, economies with a higher services expenditure share display stronger responses to monetary policy shocks, both across countries and in the time-series. Third, low-income households are more exposed to price fluctuations because they spend relatively more on goods, whereas high-income households spend relatively more on services. To rationalize these facts, I then develop a two-sector heterogeneous-agent model with sector-specific nominal rigidities and non-homothetic preferences. The framework can explain why the Phillips curve, a statistical relationship between economic activity and inflation, has flattened in recent decades: As the economy shifts toward services, aggregate price rigidities rise, strengthening the response of consumption to monetary policy. Quantitatively, I find that the increase in the services share between 1970 and 2019 made the consumption response to monetary policy 21% larger.

I document my first empirical fact using data on the frequency and magnitude of price changes compiled by Nakamura and Steinsson (2008). I show a large and persistent gap in price adjustment across the goods and services sectors. At the monthly frequency, only 11% of service prices are reset, versus 35% for goods. When prices are reset, absolute price changes are 9% for services and 19% for goods. These differences are stable over time and are observed across a range of countries. Second, using time-series data on consumption and well-identified monetary policy shocks, I find that the short-run consumption response to monetary policy has increased

over time. This is evidence of the flattening of the Phillips curve. I also show suggestive evidence that countries with a larger consumption share of services exhibit systematically stronger responses. Third, using detailed household expenditure data for the United States, I show that the services share rises with income: households in the top income quintile spend, on average, roughly 12 percentage points more of their consumption on services than those in the bottom quintile.

Rising services share, combined with higher price stickiness, provides a simple mechanism for the observed increase in monetary policy effectiveness and the cross-country relationship. Structural transformation raises aggregate price rigidity, which leads to a flatter Phillips curve, amplifying the short-run real impact of nominal interest rate changes. Therefore, my paper provides a new explanation for the flattening of the Phillips curve in recent decades.[1]

To rationalize this mechanism and quantify its impact, I develop a two-sector dynamic general equilibrium model with household and sector heterogeneity. The production side features two final consumption sectors: services and goods. Each of these final producers aggregates a continuum of differentiated intermediate inputs, produced using labor and a sector-specific productivity term in a linear production technology. Intermediate input producers face *sector-specific price adjustment costs* consistent with the observed empirical pattern. On the household side, the model builds on the standard incomplete markets framework, incorporating endogenous labor supply choices, idiosyncratic productivity shocks, and a borrowing limit. In comparison with standard heterogeneous-agent models, I introduce non-homothetic preferences over goods and services, which creates *heterogeneous demand composition*: Households not only choose how much to work, consume, and save, but also how to allocate their consumption.

The goal is to compare the effects of a monetary policy shock in a high-services economy (the 2019 economy) with those in a low-services economy (the 1970 economy). I first calibrate the model to the 2019 economy to match the cross-sectional heterogeneity of households in income, wealth, and demand composition, as well as the sectoral heterogeneity in price rigidities and the aggregate services share. The model reproduces an empirically consistent Engel curve for services, along with realistic wealth distributions and marginal propensities to consume, which are not directly targeted. Using estimates of sectoral productivity growth, I then solve

[1] Leading explanations for the flattening of the Phillips curve include the anchoring of expectations (Coibion and Gorodnichenko, 2015, Ball and Mazumder, 2019, Hazell et al., 2022); globalization (Forbes, 2019); FED credibility (McLeay and Tenreyro, 2020); measurement issues (Stock and Watson, 2020); and particular features of the 2008 financial crisis (Gilchrist et al., 2017).

backward to construct counterfactual economies for earlier years, all the way back to 1970. The resulting differences in sectoral productivities generate the historical rise in the services share, as well as the observed increase in the relative price of services and the decline in hours worked.

Using the estimated model for both the low-services and high-services economies, I compare the effects of a conventional monetary policy shock that raises the nominal interest rate by 100 basis points. Focusing on the peak consumption response—the largest component of GDP—the same shock reduces consumption by 21% more in 2019 than in 1970. In contrast, sectoral prices respond only modestly more: 14% for goods and 11% for services. The stronger reaction of consumption relative to prices reflects the flattening of the Phillips curve, whereby real activity becomes more sensitive to monetary policy. Quantitatively, each percentage point increase in the services share amplifies the peak consumption response by about 1.2%. The contractionary effect on employment is also substantially stronger, being 32% larger in 2019 than in 1970.

To quantify the rise in aggregate price rigidities mechanism generated by structural transformation, I build a counterfactual where price rigidities are homogeneous across sectors. I find that the sectoral differences in price adjustment explain 80% of the expansion of monetary policy transmission between 1970 and 2019. Structural transformation shifts the economy to a more price-sticky sector, which flattens the Phillips curve and enlarges consumption responses to interest rate changes. Crucially, non-homothetic preferences also play a role by dampening monetary policy effects by 17%. Comparing with a homothetic preferences counterfactual, I find that the introduction of different demand income elasticities generates an additional precautionary savings motive, inducing wealth accumulation which decreases the aggregate marginal propensity to consume and, in turn, the aggregate effects of monetary policy.

While structural transformation amplifies the effects of monetary policy, it dampens the effects of supply shocks. A negative supply shock raises prices overall, but because price rigidities are stronger in services than in goods, goods prices rise more sharply, leading to a larger decline in goods consumption. As the services share increases, a smaller portion of the consumption basket is exposed to these sharp price hikes, and aggregate consumption contracts less. In contrast to demand shocks, structural transformation makes the economy less sensitive to supply shocks. This may help explain the modest impact of recent supply shocks relative to professional forecasters' expectations (Chahad et al., 2023).

The distributional effects of monetary policy are a concern for central banks (Yellen, 2016, Constâncio, 2017). I also study how structural transformation has changed these distributional

impacts of monetary policy. Since households adjust in multiple dimensions, I focus on welfare that summarizes the different adjustments in a single statistic. For each household, I compute the welfare loss associated with a monetary contractionary shock by calculating the percentage increase in assets that renders them indifferent to the steady-state. I find that monetary policy became 5% more costly in 2019 than in 1970, from a utilitarian perspective. This masks a significant difference across households. Low-wealth households are the ones that bear the highest welfare loss when interest rates rise. With a larger services share, their welfare loss is even larger, rising by 7.2% for households at the bottom of the wealth distribution. On the other hand, households at the top of the distribution benefit from the interest rate increase and, in an economy more concentrated in the services sector, their gains from the contractionary monetary policy shock are slightly higher. Overall, this means that structural transformation amplifies the inequality in welfare generated by interest rate increases.

The findings offer a new perspective on the heterogeneous effects of monetary policy within common currency areas, where differences in the services share of consumption can account for variation in policy transmission across economies.[2] Furthermore, as population aging continues to shift demand toward services—given that older households consume relatively more services—the effectiveness of monetary policy may strengthen further.

Related Literature and Contribution. This paper contributes to three strands of the literature. The first studies how long-run trends affect monetary policy transmission. This literature has focused on how changes in inflation expectations (Boivin and Giannoni, 2006); population aging (Leahy and Thapar, 2022, Mangiante, 2025); labor market structure (Pancrazi and Vukotić, 2019); or changes in the production networks (Galesi and Rachedi, 2019) affect monetary policy transmission. I complement this literature by focusing on the structural transformation trend. Specifically, I focus on the effects of a shift toward a sector with higher income elasticity of demand and higher price rigidity.

Second, this paper contributes to the literature that studies how household heterogeneity matters in the propagation of monetary policy, whose recent advances are summarized by Kaplan et al. (2018) and Auclert et al. (2025). This literature's focus has been on the role of differences in income sources (Gornemann et al., 2016); household balance sheets (Kaplan et al.,

[2]Several authors have examined the factors underlying these heterogeneous responses. Recent examples include Burriel and Galesi (2018), Almgren et al. (2022), Corsetti et al. (2022), and Pica (2023).

2018, Auclert, 2019, Slacalek et al., 2020, Luetticke, 2021); sectoral price rigidities (Clayton et al., 2018, Cravino et al., 2020); and unemployment risk (Challe, 2020, Bonciani and Oh, 2021).[3] A related literature studies the dynamics of monetary transmission with heterogeneous firms under production networks (Carvalho, 2006, Castro, 2019, Pasten et al., 2020, La'O and Tahbaz-Salehi, 2022, Rubbo, 2023, Ghassibe, 2024). This paper shows that accounting for the sectoral allocation of consumption across households dampens the aggregate effects of monetary policy.

Third, I contribute to the growing literature that studies non-homothetic preferences. A large literature has studied structural transformation in the context on non-homothetic preferences (see Herrendorf et al. (2014) for a review).[4] More recently, a series of papers have studied the role of non-homothetic preferences in business cycle fluctuations and public policy. Olivi et al. (2024) use them for optimal monetary policy. Andreolli and Surico (2025) and Orchard (2025) show how consumption baskets heterogeneity between necessities and luxuries interacts with business cycles. Boehnert et al. (2025) study monetary policy transmission through differences between tradable and non-tradable goods. De Nardi and Fella (2017) use this class of preferences to account for wealth inequality. Jaimovich et al. (2019) and Becker (2024) use them to account for differences in the product quality of consumption.[5] My contribution is distinct in modeling heterogeneity in both income elasticities and consumption basket composition across goods and services with differing price rigidities. This allows me to examine how such heterogeneity alters the transmission of monetary policy with rising services consumption.[6]

Outline. The paper is structured as follows. Section 2 documents the empirical facts that motivate the paper. Section 3 describes the model used to study monetary policy transmission. In Section 4, I estimate the model and its fit. Section 5 performs the main analysis, studying the role of structural transformation in changing the propagation of monetary policy shocks. In Section 6, I study how structural transformation changes the economic responses to negative supply shocks. Section 7 concludes.

[3] A complementary literature uses tractable heterogeneous-agent, or two-agent models to analyze the implications of household heterogeneity for monetary policy transmission (Bilbiie, 2020, 2025, Debortoli and Galí, 2025).

[4] Other recent applications that look at long-run implications of using non-homothetic preferences include Hochmuth et al. (2025) that model heterogeneous energy use across the income distribution, and Arvai and Mann (2025) that model consumption differences in digitally produced products.

[5] See Danieli (2020), Sonnervig (2025) and Bernardino et al. (2025) for three other recent applications.

[6] Non-homothetic preferences have also been used in the finance literature to explain the equity premium puzzle (Ait-Sahalia et al., 2004), to study asset prices (Pakos, 2004), or to explain differences in portfolio shares (Wachter and Yogo, 2010).

2 Empirical Analysis

In this section, I document three empirical facts about the price of services, demand composition, and monetary policy transmission in the United States. These facts motivate and guide the model and calibration that I describe in the next sections.

First, I show that the price rigidity is higher in the services sector than in the goods sector. Second, I provide evidence of the flattening of the Phillips curve and how the sectoral composition, in terms of services, can explain it. Third, I show that the Engel curve for services exhibits a positive income gradient, indicating that high-income households allocate a larger share of their total consumption to the services sector than low-income households.

2.1 Sectoral Price Rigidity

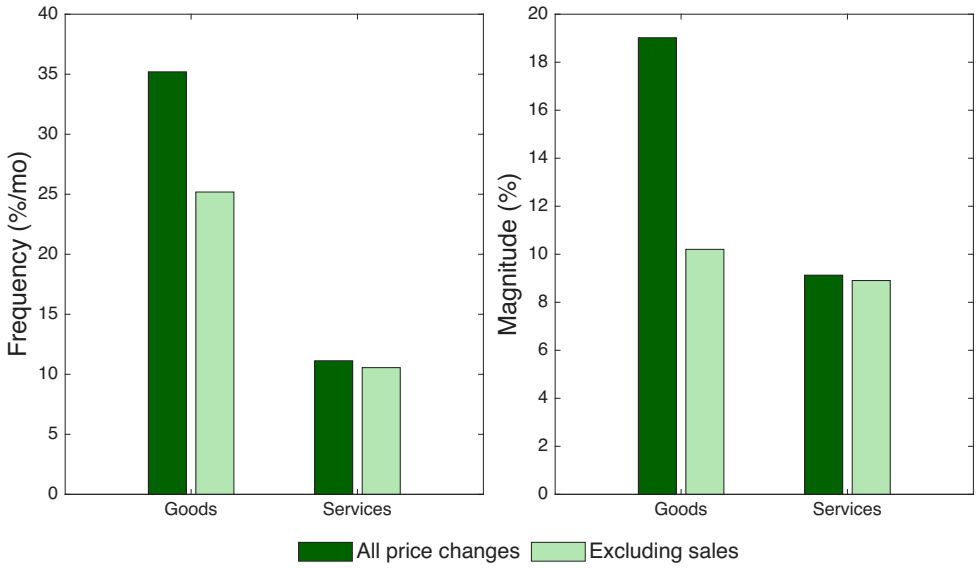
My first piece of empirical evidence concerns the differences in price rigidity between goods and services, where prices in the services sector adjust less than those in the goods sector. To document this fact, I use the data compiled by Nakamura and Steinsson (2008) on price dynamics. In this paper, the authors assembled a dataset with the frequency and magnitude of price adjustments for several categories of goods and services. It is built based on the confidential monthly product-level price data used to construct the CPI. The data cover the prices of around 70% of consumer expenditures (excluding housing prices) between 1998 and 2005 in the U.S. In particular, I use the median frequency of price changes and the median absolute size of changes in consumer prices by consumption category between 1998 and 2005. In Appendix A.1, I describe in detail the data.

Based on the reported frequency and magnitude of price adjustments for aggregated expenditure categories, I aggregate the different categories into goods or services, following the U.S. Bureau of Economic Analysis (BEA) classification. The aggregation is a weighted mean, with the average budget shares as weights.

In the left panel of Figure 1, I show the median percentage of goods and services that have a change in price within a month. In the right panel, I show the median magnitude of price changes of items classified as goods and services. I compute this number for all price changes (dark green bars) and regular price changes that exclude sales (light green bars).

The figure shows a marked difference in price adjustment frequency between goods and services: 35% of goods change prices within a month, compared with only 11% of services.

Figure 1: Frequency and Magnitude of Price Changes for Goods and Services



Notes: The left panel shows the median frequency of price changes in a month for goods and services. The right panel shows the median magnitude price change. Data are sourced from Nakamura and Steinsson (2008).

Assuming a Poisson process, this means that the implied median duration of a price spell is 2.3 months for goods and 8.5 months for services. The same difference is present when it comes to magnitudes. The median absolute price change of goods is 19%, whereas for services it is 9%. When examining only price changes that do not result from a sale, the difference between goods and services is smaller; however, it is clear that goods adjust their prices more than services.[7]

Importantly, this empirical fact is not specific to the dataset, the time considered, or the U.S. In Appendix A.3, I use an alternative dataset assembled by Bils and Klenow (2004) and find a similar pattern. Furthermore, using prices for the euro area, Gautier et al. (2024) and Dhyne et al. (2006) also demonstrate that the prices of services adjust less than those of food and industrial goods.

There are potential reasons for these observed differences. First, most categories included in the goods sector are tradable and therefore face more competition, which creates pressure for more frequent price adjustments (Cavalleri et al., 2019). A second reason is related to the production structure of goods vs. services: The production of services uses a higher share of labor than goods. Since wages are rigid, the higher labor share in the services sector could imply

[7] As Nakamura and Steinsson (2008) point out, temporary sales play an important role in generating price flexibility for retail prices. These temporary sales are more common in goods than services and are associated, for instance, with the perishability of goods.

that service prices adjust less frequently (Bobeica et al., 2019). In a review paper, Nakamura and Steinsson (2013) discuss reasons for the observed differences in price stickiness across sectors. In this paper, I do not aim to explain the differences observed, but rather evaluate their implications for monetary policy transmission. Therefore, the model I present below assumes heterogeneous sectoral price stickiness, which is exogenous. All of the factors that explain these differences are captured in a reduced form by differences in the parameter of the cost of adjusting prices that firms face.

There is weak evidence of an increase in the frequency of price adjustments over time. Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008) find minimal differences in the overall frequency of price changes when comparing adjustments between 1988–1997 and 1998–2005 for the U.S. Gautier et al. (2024) compare their results with those of Dhyne et al. (2006) and find, in a similar way, very small differences over time between 1996–2004 and 2011–2017 for the euro area. Because of these, in the next section, I will assume that the frequency of price adjustment at the sectoral level is fixed.

Structural transformation reallocates economic activity toward the services sector. Because the differences in price flexibility are persistent, the overall price rigidity increases as the services share of the economy rises. This means that the non-neutrality of monetary policy enlarges with structural transformation. Next, I look at direct evidence of the increase of monetary policy transmission and how that correlates with the services share.

2.2 The Flattening of the Phillips Curve and the Services Share

The second fact that I document is the flattening of the Phillips curve, a formal statement about the relationship between inflation and output, and how structural transformation can explain that. At least since Atkeson et al. (2001) that there is evidence about the decline in the slope of the Phillips curve. In this sub-section, I first show that the responses of real economic activity have increased over time, consistent with the flattening of the Phillips curve. Second, in a cross-country comparison, I show that countries with larger monetary non-neutrality effects are also the countries with higher services shares.

For the first part, I use monthly data on total expenditure from the BEA and monetary policy shocks by Romer and Romer (2023), the longest time-series of these shocks. The shocks are identified using a narrative approach based on readings and analysis of the historical minutes

and transcripts of the Federal Open Market Committee (FOMC) meetings. The authors identify seven monetary policy shocks since 1960. To compare across periods, I scale the shocks by the associated change in the interest rate. I estimate impulse responses using local projections (Jordà, 2005):

$$\Delta \log C_{t+h|t-1} = \alpha_h + \beta_h \epsilon_t^M + \gamma_h X_t + \varepsilon_{t+h}, \text{ for } h = \{0, 1, ..., 12\}, \tag{1}$$

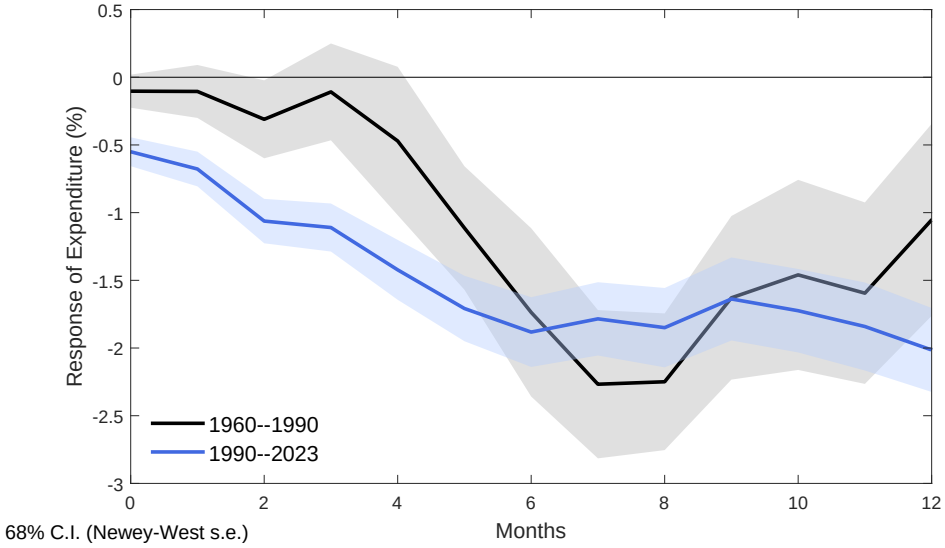
where $\log \Delta C_{t+h|t-1}$ is the log consumption change between $t-1$ and $t+h$, ϵ^M is the monetary policy shock scaled by the interest rate, and X_t is a matrix of controls that include 12 months lags of the monetary policy shock and the dependent variable. The set $\{\beta_h\}$ gives the estimates of the impulse response function for each horizon h. Figure 2a plots the estimated impulse response functions to a contractionary monetary policy shock, splitting the sample between 1960–1990 and 1990–2023.

A contractionary monetary policy shock reduces total consumption. A higher interest rate creates a negative income effect through higher interest payments and a negative substitution effect that postpones current consumption and increases savings. The peak of the response is achieved 6 months after the shock. There are, however, differences in the two impulse response functions estimated. The response between 1960 and 1990, for the first 6 months, is smaller in magnitude than the response between 1990 and 2023. This indicates that the same monetary policy shock has a different impact over time. In the more recent period, monetary non-neutrality is stronger than in the earlier period. I interpret this as direct evidence of the flattening of the Phillips curve.

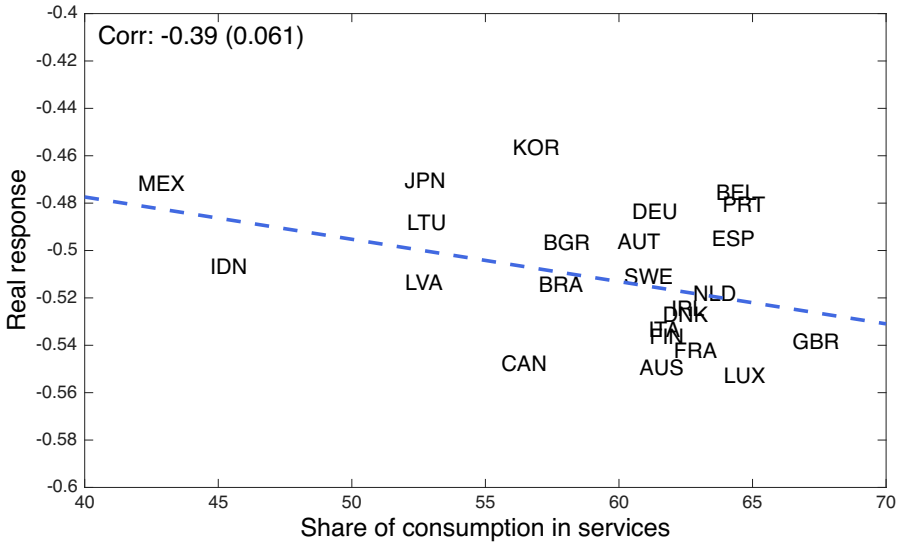
In the second part of this analysis, I look into the cross-country relationship between monetary non-neutrality and the services share, in the same spirit of Brinca et al. (2016) or Almgren et al. (2022). I use the output responses to monetary policy estimated by Galesi and Rachedi (2019), who estimate a structural VAR model using a panel of quarterly data on output, inflation, and interest rates for 25 countries. Monetary policy shocks are identified through sign restrictions on the impact of interest rates on output and prices. Specifically, the authors assume that a monetary policy shock raises the nominal interest rate while reducing both inflation and output. The share of consumption in services corresponds to the 2000-2020 average, constructed from national accounts data reported by each country's statistical authority—the counterpart to the U.S. BEA in Table 2.3.5.

Figure 2: Monetary Policy Transmission and the Services Share

(a) Expenditure IRF



(b) Output Response and the Services Share



Notes: Panel (a) plots the impulse response function of total expenditure to a monetary policy shock using Equation (1). The bands around the impulse responses are 68% confidence intervals using Newey and West (1987) standard errors. Panel (b) shows the relationship between the contemporaneous response of output to a 100 basis point increase in the interest rate and the share of consumption in services across countries for 25 countries. The correlation between the output response and the share of consumption in services is -0.39, with a p-value of 0.06.

Figure 2b presents a scatter plot between the contemporaneous output response to monetary policy and the share of consumption in services. The figure suggests that monetary policy has

bigger effects on output in countries with a higher share of consumption in services. In fact, the linear correlation coefficient between the output response and the share of consumption in services is -0.39, and the slope estimated in the linear regression between these two variables is also negative. Mexico has the smallest services share (42%) and with a response of -0.47%, and the United Kingdom has the highest services share (66.5%) and a response of -0.54%.

Overall, this empirical analysis shows that the Phillips curve has flattened over time in the United States, and that countries with a higher share of consumption in the services sector respond more to monetary policy. The increase in the services share, in combination with the difference in price rigidity, provides a simple and novel explanation for the observed flattening of the Phillips curve.[8] As the economic activity gets more concentrated in the services sector, the overall level of price rigidities also increases, which decreases the slope of the Phillips curve, i.e., enlarges monetary policy non-neutrality. The cross-country correlation provides additional suggestive evidence of this mechanism.

The central objective of this paper is to rationalize how structural transformation changes the transmission of monetary policy. In this next section, I develop a model that establishes this relationship through differences in sectoral price rigidity.

2.3 Heterogeneous Demand Composition

The third motivating fact concerns differences in the demand composition between low- and high-income households. Low-income households allocate a smaller share of their expenditures to services than high-income households. To document this, I use expenditure data from the Consumer Expenditure Survey (CEX) curated by the U.S. Bureau of Labor Statistics (BLS). The CEX contains household-level consumption data that track spending in all product categories, such as food, housing, utilities, transportation, health, and education. This is the data that is used to estimate expenditure weights in computing the official Consumer Price Index.

The CEX has two datasets: Quarterly interviews and diary surveys. For this paper, I use the former between 2000 and 2019.[9] This dataset consists of a rotating panel in which, in each quarter, between 5,000 and 8,000 households are interviewed; each household stays in the sample, at most, for 5 consecutive quarters. Because most income questions are only asked

[8] See Stock and Watson (2020) for a review of other reasons that justify that.

[9] Even though the CEX starts in 1980, the questionnaire had several changes in the 1980s and 1990s, which could make it harder to compare across time. I also exclude the years associated with the COVID-19 pandemic, since consumption patterns change substantially due to other circumstances beyond economic factors.

in the final interview, I only use households that participate in all interviews. Furthermore, I exclude from the sample households with negative expenditure in a given category, zero food expenditures, or negative income and when the household head is below 25 years old. This leaves me with an average of around 6,000 households per quarter. In Appendix A.1, I describe the details of all steps in the data cleaning.

As the goal is to document how consumption patterns change over the income distribution, I start by dividing households into five income groups. Household income is the sum across household members of their pre-tax income (including earnings, business income, social benefits, pensions, alimony, gifts, and gambling winnings), subtracted from personal taxes. I then regress income on household size, average age of the two main earners, and number of earners in the household to control for differences in household demographics. Based on the residuals of this regression for each quarter, I divide households into five income bins, corresponding to the 5-20, 20-40, 40-60, 60-80, and 80-95 percentiles.[10]

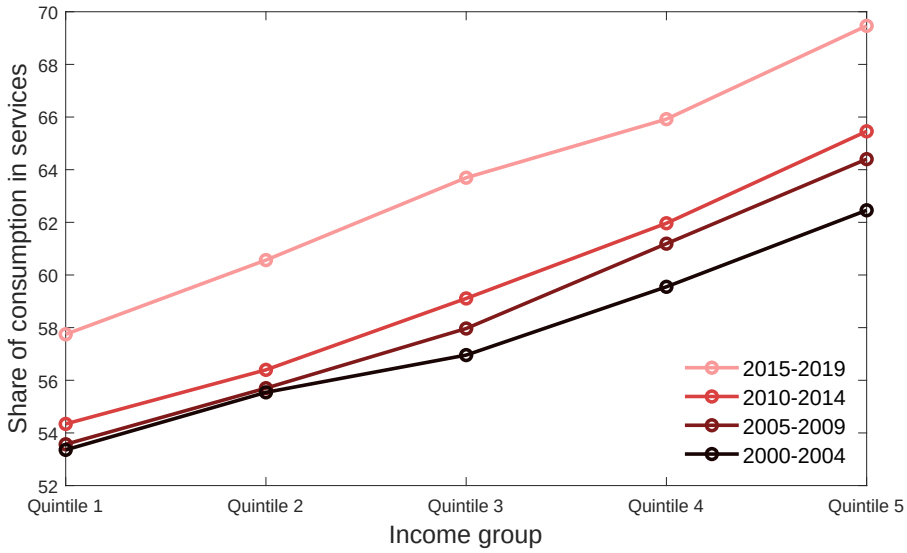
Second, I classify non-durable expenditure items by economic activity. The CEX reports expenditures on several items. I divide these items into services or goods, following the U.S. Bureau of Economic Analysis (BEA) classification. Examples of goods expenditures include food and alcohol consumed at home, whereas service expenditures include food away, public transportation, or education. In Appendix A.2, I provide the complete classification.

In the third and last step, I compute the average expenditure share on services and goods for each income group. Figure 3 plots the average service share by income group for different points in time.

Two patterns emerge in this figure. First, the share of total expenditure households allocate to services increases over the income distribution. Households in the lowest income group have a smaller share of their expenditures on services than those in the highest income group. Second, there is an upward trajectory of the aggregate services share in consumption due to structural transformation. The slope of the Engel curve over time is relatively constant, with the high-income group allocating 11-12 percentage points more of their expenditure to services than the low-income group. This empirical fact can be seen as a generalization of Engel's law, which states that as income rises, household expenditure on food increases in absolute terms but its relative proportion declines (Engel, 1857, 1895). Boppart (2014) also documents this fact.

[10]Following Aguiar and Bils (2015), I exclude the bottom and top 5% to eliminate outliers and mitigate the impact of top-coding. In Appendix A.4, I show that this exclusion does not affect the results.

Figure 3: Average Service Share by Income Quintiles



Notes: The figure shows the average service share in total expenditures for each income quintile grouped in 5 years. Table A.3 in the Appendix enumerates all the categories classified as services. The data source is the U.S. Consumer Expenditure Survey between 2000 and 2019.

This fact is robust when I exclude the old-age population from the data. This group consumes a higher share of services than the working-age population (on average, 4.3 percentage points more), and could affect the pattern above. Moreover, including the bottom and top 5% of the income distribution in the first and last income group, respectively, does not significantly affect the pattern described above. In Appendix A.4, I describe these two robustness exercises in detail.

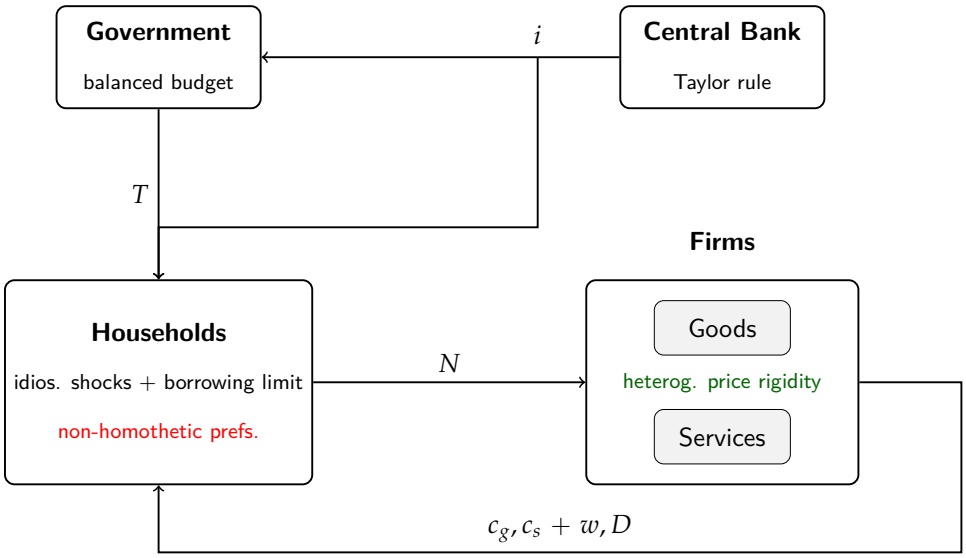
This empirical fact means that differences in demand composition across the income distribution imply that aggregate shocks—such as changes in nominal interest rates—can affect households asymmetrically. Accordingly, the model developed in the next section incorporates heterogeneous demand composition, microfounded through non-homothetic preferences. This preference structure is also essential to endogenously generate the rise in the services share.

3 Model

In this section, I describe the model used to study how changes in the composition of consumer demand, in terms of goods and services, affect the transmission of monetary policy. The model comprises two sectors: goods and services. It combines features of heterogeneous-agent

incomplete-markets models (Bewley, 1986, Imrohoroğlu, 1989, Huggett, 1993, Aiyagari, 1994) with non-homothetic preferences, as in Comin et al. (2021), and the New-Keynesian framework (Woodford, 2004, Galí, 2008). There is also a government that issues bonds on which it pays interest. A monetary authority determines the interest rate in accordance with a Taylor rule. Time is discrete, and one period in the model corresponds to one quarter.

The chart below illustrates the model schematically. Compared with the literature that studies monetary policy transmission, this framework is novel in combining heterogeneous price rigidities across sectors with non-homothetic consumer preferences.



3.1 Households

There is a continuum of infinitely lived ex ante identical households with measure 1. They derive utility from consumption and disutility from supplying labor according to:

$$U = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, h_t), \tag{2}$$

where $\beta \in (0,1)$ is the discount factor. The per-period utility function assumes CRRA preferences over consumption and additively separable preferences for leisure, using the MaCurdy (1981) specification:

$$u(c,h) = \frac{c^{1-\gamma}-1}{1-\gamma} - \chi \frac{h^{1+\eta}}{1+\eta}, \tag{3}$$

where γ is the inverse intertemporal elasticity of substitution, η is the inverse Frisch elasticity, and χ is a parameter that regulates the disutility of working. c is a consumption index that

aggregates the sectoral consumption of goods, c_g, and services, c_s. It is implicitly defined through a non-homothetic constant elasticity of substitution (CES) function, as in Comin et al. (2021):

$$1 = c^{\frac{1}{\sigma}} c_g^{\frac{\sigma-1}{\sigma}} + (\Omega c^{\epsilon})^{\frac{1}{\sigma}} c_s^{\frac{\sigma-1}{\sigma}}, \tag{4}$$

where σ is the demand price elasticity, Ω is the relative taste-shifter of services, and ϵ is a parameter that regulates the income elasticity of demand, i.e., the degree of non-homotheticity. I restrict the demand price elasticity to be positive, such that $\sigma \in (0, 1)$, which implies that goods and services are gross complements. Under this restriction, ϵ must be larger than 1 for the services Engel's curve to have a positive slope (Figure 3). Moreover, Ω is positive. Note that if $\epsilon = 1$, I recover the homothetic CES aggregator.[11]

The household's labor productivity, ω_t, follows a log-AR(1) process with persistency ρ_ω and zero-mean shocks with variance σ^2_ω. I discretize this process in a 5-state Markov chain. These shocks are not insurable, and households face a no-borrowing constraint.

Households rent their labor services, $\omega_t h_t$, for a real wage w. Also, households, at a price p_b, can buy and sell bonds, b, which are the savings vehicle. The returns on savings are given by i_t, the nominal interest rate that is set by the central bank.

Thus, at time t, households face the following budget constraint:

$$\begin{aligned} &p_{g,t}c_{g,t} + p_{s,t}c_{s,t} + p_{b,t}b_{t+1} = w_t\omega_t h_t + (p_{b,t} + i_t)b_t + T_t + D_t, \\ &\text{with } \log\omega_t = \rho_\omega \log\omega_{t-1} + \varepsilon^\omega_t, \;\; \varepsilon^\omega \sim \mathcal{N}(0, \sigma^2_\omega), \end{aligned} \tag{5}$$

where $p_{g,t}$ and $p_{s,t}$ are the sectoral prices of goods and services, respectively, T_t is a lump-sum transfer (tax if negative), and D_t are the dividends distributed to households from the firms' profits. I distributed dividends according to the skill-level as in McKay and Reis (2016), resembling the bonus payments that higher productivity workers usually receive.

The household problem can be decomposed into two sub-problems: The allocation of consumption across sectors within each period (the intratemporal problem) and the allocation of consumption and savings over time (the intertemporal problem).[12] This means that the household problem is solved sequentially: first, the consumption-savings decision, then the

[11] This formulation sets the goods sector as the base sector of this model. One could also write preferences for sector-specific ϵ and Ω. However, a uniform scaling of (ϵ_g, ϵ_s) or (Ω_g, Ω_s) would imply the same observable choice behavior. Hence, I opted for a parsimonious approach to simplify the exposition.

[12] This is an application of a two-stage budgeting approach that applies to explicitly additive preferences, as shown by Gorman (1971).

intratemporal allocation of consumption across sectors.

For the intratemporal problem, households take prices as given and minimize the total expenditure, E, subject to Equation (4), the non-homothetic CES aggregator:[13]

$$\begin{aligned} &\min_{c_g, c_s} E = p_g c_g + p_s c_s \\ &\text{s.to} \;\; c^{\frac{1}{\sigma}} c_g^{\frac{\sigma-1}{\sigma}} + (\Omega c^{\epsilon})^{\frac{1}{\sigma}} c_s^{\frac{\sigma-1}{\sigma}} = 1. \end{aligned} \tag{6}$$

The solution of the problem yields the sectoral Hicksian demands given by

$$c_g = \left(\frac{p_g}{E}\right)^{-\sigma} c \text{ and } c_s = \Omega \left(\frac{p_s}{E}\right)^{-\sigma} c^{\epsilon}, \tag{7}$$

which indicates the quantity of goods and services households consume, given sectoral prices and total expenditure. The expenditure function is given by

$$E(c_g, c_s) \equiv \sum_{m \in \{g,s\}} p_m c_m = \left[c p_g^{1-\sigma} + \Omega c^{\epsilon} p_s^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \tag{8}$$

For the intertemporal problem, households choose how much to consume, save, and work, given wages and prices, maximizing Equation (2) subject to the household budget constraint, Equation (5). The household problem can be written recursively as follows:

$$\begin{aligned} V(\omega, b; \Xi) = \max_{\{c, b', h\}} \;\; & u(c,h) + \beta \mathbb{E}\left[V(\omega', b'; \Xi')\right] \\ \text{s.to} \;\; & E + p_b b' = w\omega h + (p_b + i)b + T + D \\ & E = \left[(p_g c)^{1-\sigma} + \Omega (p_s c^{\epsilon})^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \\ & \Xi' = \Psi(\Xi) \\ & c \geqslant 0, \; b' \geqslant 0, \; h \in (0,1), \end{aligned} \tag{9}$$

where $\Xi(\omega, b) \in \mathcal{M}$ is the distribution on the space $X = W \times B$, of labor productivity $\omega \in W$, and agents' bond holdings $b \in B$ across the population, which together with the policy variables determine the equilibrium prices. Ψ is an equilibrium object that specifies the law of motion of the distribution Ξ, and $\mathcal{M}$ is the set of probability measures on X.

[13] In Appendix B.1, I derive the problem and show the relevance of this preference relation to match the heterogeneous demand composition fact reported in Section 2.

3.2 Production

3.2.1 Sectoral Producer

Each sector $m \in \{g, s\}$ has a final representative producer that aggregates a continuum of intermediate inputs indexed by $j \in [0, 1]$ and with prices $p_m(j)$:

$$Q_m = \left(\int_0^1 q_m(j)^{\frac{\theta_m - 1}{\theta_m}} dj \right)^{\frac{\theta_m}{\theta_m - 1}}, \tag{10}$$

where θ_m is the elasticity of substitution across intermediate inputs for sector m. Given a level of sectoral demand Q_m, the final producer cost minimization problem yields the demand for intermediate input j, which is given by[14]

$$q_m(j) = q(p_m(j); P_m, Q_m) = \left(\frac{p_m(j)}{P_m} \right)^{-\theta_m} Q_m, \tag{11}$$

where P_m is the equilibrium price of the final good m and can be expressed as

$$P_m^{1-\theta_m} = \int_0^1 p_m(j)^{1-\theta_m} dj. \tag{12}$$

3.2.2 Intermediate Input Producers

Each input j is produced by an intermediate producer that operates under monopolistic competition using labor input $n_m(j)$ hired in a competitive labor market at the wage rate of w. The production technology is linear in labor:

$$q_m(j) = Z_m n_m(j), \tag{13}$$

where Z_m is the sector m productivity. This functional form of the production function implicitly assumes that capital is fixed. Since I am interested in comparing the responses to a monetary policy shock in the short run, and capital is a slow-moving variable, this assumption does not entail a significant cost.

Intermediate input producers face price adjustment costs as in Rotemberg (1982). These adjustment costs are measured in nominal terms and given by a quadratic function of the

[14] In Appendix B.2, I show the full derivation of the intermediate input demand.

change in prices:

$$\Phi_m\left(p_{m,t}(j), p_{m,t-1}(j)\right) = \frac{\theta_m}{2\kappa_m}\left[\log\left(\frac{p_{m,t}(j)}{p_{m,t-1}(j)}\right)\right]^2 Q_{m,t}p_{m,t}, \tag{14}$$

where κ_m is a sector-specific parameter that regulates the cost of price adjustment that is set to match the sector-specific frequency of price adjustment.

Firm's j optimizes nominal profits by solving the following pricing problem

$$\begin{aligned} V_{m,t}(p_{m,t-1}(j)) = \max_{\{p_{m,t}(j)\}} & \; p_{m,t}(j)q_{m,t}(j) - w_t n_{m,t}(j) - \Phi_m(p_{m,t}(j), p_{m,t-1}(j)) \\ & + \frac{1}{1+i_t} V_{m,t+1}(p_{m,t}(j)) \\ \text{s. to} \quad & \text{Equations (11) and (13).} \end{aligned} \tag{15}$$

Let $1 + \pi_{m,t} \equiv \frac{P_{m,t}}{P_{m,t-1}}$, the sector-specific inflation rate. In a symmetric equilibrium, the optimal price-setting rule of sector m is

$$\log(1+\pi_{m,t}) = \frac{\kappa_m}{\theta_m}\left(1 - \theta_m + \theta_m \frac{w_t}{Z_m P_{m,t}}\right) + \frac{1}{1+i_t}(1+\pi_{m,t+1})\log(1+\pi_{m,t+1})\frac{Q_{m,t+1}}{Q_{m,t}}, \tag{16}$$

which is the New-Keynesian Phillips curve (NKPC) of sector m.[15] Note that in the steady state, where there are no price adjustments, it is simply equal to the markup multiplied by the marginal cost—in this case, the wage rate:

$$P_m = \frac{\theta_m}{\theta_m - 1}\frac{w}{Z_m}. \tag{17}$$

The adjustment costs are not resource costs, but act as if they affected the price setter's utility. This assumption is to avoid counterfactual price-adjustment booms that would arise from large resource costs associated with large price movements, as in Hagedorn et al. (2019). The profits of intermediate producers of sector i are distributed to households and given by

$$d_{m,t} = \int_0^1 d_{m,t}(j)dj = P_{m,t}Q_{m,t} - w_t N_{m,t}, \tag{18}$$

with $N_m = \int_0^1 n_m(j)\, dj$.

[15] In Appendix B.3, I show in detail the derivation of the New-Keynesian Phillips curve.

3.3 Government

The government issues nominal bonds denoted by B that households can buy at the price of p_b. The total value of public debt in the economy is fixed. The government levies a lump-sum tax on households to finance interest payments on bonds. The government budget is balanced at each period and reads as

$$p_{b,t}B = \int (p_{b,t} + i_t) b_t d\Xi + T_t. \tag{19}$$

3.4 Monetary Authority

The monetary authority sets the nominal interest rate following a Taylor (1993) rule:

$$i_t = i_{SS} + \phi_\pi \pi_{t-1} + \varepsilon_t^M, \tag{20}$$

where i_{SS} is the steady-state nominal interest rate, π_t is the aggregate inflation rate in period t, which is the weighted average of the sectoral inflation rates, using sectoral consumption as weights. ε_t^M is the monetary policy shock that follows an AR(1) process with persistency ρ_M.

3.5 Market Clearing and Competitive Equilibrium

Market clearing occurs when firms' demand for labor equals the total supply of labor from households, households' demand for bonds matches the quantity of outstanding bonds, and households consume all of the goods produced by firms in each sector:

$$\int \hat{h}_t(\omega, b)\, d\Xi_t = N_{g,t} + N_{s,t} \tag{21}$$

$$\int \hat{b}_t(\omega, b)\, d\Xi_t = B \tag{22}$$

$$\int \hat{c}_{m,t}(\omega, b) d\Xi_t = Q_{m,t},\ m \in \{g, s\}, \tag{23}$$

where $\hat{h}_t(\omega, b), \hat{b}_t(\omega, b)$ and $\hat{c}_{m,t}(\omega, b)$ are the household policy functions of labor, savings, and sectoral consumption.

Definition. A competitive equilibrium is a sequence of lump-sum transfers T_t; interest rates i_t; value functions V_t with policy functions $\hat{c}_{g,t}, \hat{c}_{s,t}, \hat{h}_t$ and $\hat{b}_t$; prices $p_{b,t}$, $p_{g,t}$, $p_{s,t}$, and w_t; profits $\pi_{g,t}$ and $\pi_{s,t}$; and a law of motion Ψ, such that:

1. V_t satisfies the Bellman equation (9), with the solution given by the policy functions $\hat{c}_{g,t}, \hat{c}_{s,t}, \hat{h}_t$ and $\hat{b}_t$ given sequences of lump-sum taxes, prices, interest rate, and dividends.

2. Firms maximize profits according to Equation (15), which are distributed in the form of dividends to households.

3. The government runs a balanced budget, as in Equation (19).

4. For all Ξ_t, market-clearing conditions (21)–(23) are satisfied.

5. The aggregate law of motion of the distribution, Ψ, is generated by the savings policy function.

4 Estimation and Model Fit

This section describes the estimation and fit of the model as a representation of the U.S. economy. The model is calibrated to capture the cross-section heterogeneity of wealth, income, and, novel in this class of models, demand composition. I then vary the sectoral productivities, which generate different services shares, and compare the effects of monetary policy across economies with different sectoral compositions.

4.1 Estimation Strategy

To calibrate the model described above, I fix 2019 as the reference year. For this year, the model is required to match the observed cross-sectional heterogeneity across households with respect to wealth, income, and consumption. These moments matter for two reasons. First, the income and wealth distributions shape the marginal propensity to consume—a central object in heterogeneous-agent models that governs the magnitude and propagation of aggregate shocks, such as monetary policy (Kaplan and Violante, 2022). Second, to assess welfare, it is important to match the distribution of service expenditures across households (the Engel curve for services, as documented in Section 2), since different sectors respond differently to a monetary shock, which may affect households heterogeneously.

In a second step, I generate different services consumption shares by changing the sectoral productivities.[16] Structural transformation is believed to be driven by two separate forces. First,

[16]The shift toward services also occurs in value-added and employment (Kuznets, 1973, Herrendorf et al., 2014).

due to productivity growth differentials, relative prices adjust, creating a substitution effect that in turn affects expenditure shares. This is the well-known cost-disease channel described by Baumol (1967). Second, an income effect operates through the non-homotheticity of preferences. A higher income increases the consumption of luxuries (goods with positive income elasticities of demand); in this case, services (Falkinger, 1990, Matsuyama, 1992).[17]

Finally, I compare the dynamics around the steady-state. In each steady-state, I generate an unexpected monetary policy shock and study the consumption and price responses to that shock. Specifically, I focus on a low and a high services share economy that match the 1970 and the 2019 services share, respectively. This approach is akin to that of Da-Rocha and Restuccia (2006), Moro (2012), and Galesi and Rachedi (2019).

Estimation of the 2019 parameters proceeds in three steps. First, I derive the relative demand for services, which allows me to estimate the price and income elasticities of demand, ε and σ. Second, I fix a set of parameters identified from external evidence, such as the income process and the price adjustment cost. Third, I use SMM to estimate the disutility of work, χ, and the taste-shifter parameter, Ω, to ensure that the model matches the aggregate hours worked and the share of services expenditure.

4.2 Model Estimation

4.2.1 Demand Estimation

To estimate the income and price demand elasticities, ϵ and σ respectively, I follow the estimation strategy proposed by Comin et al. (2021) and applied in other papers such as Buiatti et al. (2024). In summary, the non-homothetic CES preference relation—Equation (4)—allows me to derive the demand system described in Equation (7) that can be mapped to the data, allowing for the estimation of ϵ and σ.

Empirical Strategy. Take the Hicksian demands derived in Section 3–Equation (7). I can combine both and write the relative demand for services as

$$\log\left(\frac{\nu_{s,t}}{\nu_{g,t}}\right) = (1-\sigma)\log\left(\frac{p_{s,t}}{p_{g,t}}\right) + (1-\sigma)(\epsilon-1)\log\left(\frac{E_t}{p_{g,t}}\right) + (\epsilon-1)\log\nu_{g,t} + \log(\Omega),$$

[17] The structural transformation literature has emphasized the quantitative role of these two channels. For example, Boppart (2014) finds that income and substitution effects have roughly equal importance, whereas Comin et al. (2021) argue that income effects dominate. Acemoglu and Guerrieri (2008) propose differences in factor proportions and capital deepening as an alternative channel of structural transformation.

where $\nu_{i,t}$ is the share of expenditure of sector i in period t and all other letters have the same meaning as before.

This equation shows that the relative demand for services depends on the relative price of services, total expenditure, the share of goods in consumption, and a constant taste shifter. Moreover, this expression is written only in terms of observables and model parameters. By exploiting the cross-section and time variation of the expenditure shares, sectoral prices, and total expenditure, I can estimate the value of σ and ϵ. The empirical counterpart of the above equation is

$$\log\left(\frac{\nu^n_{s,t}}{\nu^n_{g,t}}\right) = (1-\sigma)\log\left(\frac{p^n_{s,t}}{p^n_{g,t}}\right) + (1-\sigma)(\epsilon-1)\log\left(\frac{E^n_t}{p^n_{g,t}}\right) + (\epsilon-1)\log\nu^n_{g,t} + \zeta^n + \xi^n_t, \quad (24)$$

where the superscript n is the household identifier, ζ^n is a household fixed effect that accounts for the taste parameter, and ξ^n_t denotes the error term of the equation.

Furthermore, two additional assumptions pertain to the structure of Equation (24). First, the term ζ^n is assumed to be a linear function of household characteristics (such as age, size, and number of earners) and regional characteristics. Including these controls accounts for taste differences in the estimates. Second, the error term may include common sector-time fixed effects across households. Including sector-time fixed effects thus captures potential aggregate consumption shocks.

I estimate σ and ϵ on this equation using the generalized method of moments (GMM). There is also a constraint in the value of the coefficients of the equation: The product of the coefficient on relative prices, $(1-\sigma)$, and expenditure share on goods, $(\epsilon-1)$, has to be equal to the coefficient on expenditure, $(1-\sigma)(\epsilon-1)$.

To account for potential measurement error due to self-reported consumption data, I instrument household expenditures with the annual household income after taxes and the income quintile of the household. This follows Aguiar and Bils (2015) and aims to capture the permanent household income that is correlated with household expenditures and is not affected by transitory measurement error in total expenditures.

Furthermore, the above equation suffers from potential endogeneity issues arising from regressing prices on quantities. I instrument household relative prices with a "Hausman" relative-price instrument (Hausman et al., 1994, Hausman, 1996) that is constructed in two steps. First, for each expenditure category included in a sector, I compute the average price

across regions, excluding the own region. Second, the aggregation is done with the average region's expenditure shares in each category as weights. These instruments aim to capture the common trend in U.S. prices (relevance condition) while alleviating endogeneity concerns due to regional shocks (exogeneity condition).

Data. I use the U.S. CEX quarterly interview dataset between 2000 and 2010. The data construction is as described in Section 2, which follows Aguiar and Bils (2015) and Comin et al. (2021) closely. I classify the non-durable expenditure between goods and services, following the U.S. BEA categorization.[18]

The consumption data are combined with regional quarterly price series from the BLS's urban CPI. For each household and sector, I construct a price index using household-specific expenditure weights, as in Comin et al. (2021). This way of building the price indexes, to a certain extent, takes into account that different households may face different effective prices.

Results. Table 1 shows the results. Column (1) reports results of the simplest estimation without any additional controls, column (2) controls for region fixed effects, and column (3) controls for region and year times quarter fixed effects. The estimated ϵ is larger than 1, which means that the Engel curve of the services sector is positively sloped. In other words, services are a luxury good, and the share of expenditure that households allocate to this sector increases with income. This is what I showed in Figure 3 in Section 2. Moreover, the estimated σ is smaller than 1, implying that goods and services are gross complements in household preferences.[19] In Appendix C.3, I re-estimate the parameters in two subsamples, finding similar values which indicate that these parameters do not change over time.

4.2.2 Pre-estimated Parameters

Other Preference Parameters. I set the discount factor, β, to 0.99, as is standard in models in which a period corresponds to a quarter. The risk-aversion parameter, γ, is set to 1.2 on the bulk part of values used in the HANK literature (e.g. Gaillard and Wangner, 2022). The inverse Frisch elasticity, η, is set to 1, following micro-estimates (Chetty et al., 2011).

[18] In Appendix A.1, I describe the dataset construction in detail, and in Appendix A.2 I provide a complete classification of goods and services.

[19] Comin et al. (2021) also find similar values even though they consider three sectors (agriculture, manufacturing, and services) instead of two.

Table 1: Demand Estimation

	(1)	(2)	(3)
σ	0.209	0.176	0.234
	(0.044)	(0.039)	(0.051)
ϵ	1.619	1.667	1.731
	(0.061)	(0.058)	(0.080)
Region FE	N	Y	Y
Year × Quarter FE	N	N	Y

Notes: The table presents results of the GMM estimation of Equation (24) with 60,932 observations. The number in parentheses corresponds to the robust standard error clustered at the household level. All regressions include demographic household controls.

Income Process. The idiosyncratic productivity process is given by a log-AR(1) process that depends on two parameters: The persistence of an idiosyncratic shock, ρ_ω, and the variance of the idiosyncratic shock, σ^2_ω. I use the estimates of Krueger et al. (2016), based on the Panel Study of Income Dynamics (PSID), for both parameters, 0.99 and 0.1, respectively. These numbers are similar to other estimates used in the heterogeneous-agent literature, such as Storesletten et al. (2004) or Karahan and Ozkan (2013). The log-AR(1) process is discretized into a 5-state Markov chain with the Tauchen (1986) method using the approach by Flodén (2008).

Sectoral Markups. The elasticity of substitution in the model has a direct mapping to markups, as Equation (17) illustrates. Marto (2024) estimates an average markup for the services sector of 1.27 and an average markup of 1.21 for the non-services sector. This implies an elasticity of substitution of 4.70 in the services sector and 5.76 for the goods sector.

Sectoral Price Adjustment Costs. The adjustment cost parameter on prices is set using the first-order equivalence of Rotemberg- and Calvo-type adjustment frictions.[20] For each sector, I set the adjustment cost parameter such that the NKPC slope in the model, θ_i/κ_i, matches the slope of the corresponding equation implied by a Calvo (1983) price adjustment, where the frequency of price adjustments corresponds to the one in Section 2. As such, κ_s is set to 89.2 and κ_g to 8.5.

[20]Up to a first-order approximation, a Calvo (1983) and a Rotemberg (1982) setting yield the same results, as shown by Roberts (1995). Only in exceptional cases does the equivalence not hold—as, for example, when there is trend inflation (Ascari and Rossi, 2012).

Sectoral Productivity. The sectoral productivities of goods and services are the only two parameters that differ between the 1970 and the 2019 steady state. For the 2019 steady state, I normalize Z_g to one and choose Z_s to match the relative price between goods and services (note that this ratio only depends on model parameters). For the 1970 steady state, I back out the productivity levels by applying average sector-specific productivity growth rates observed in the U.S. postwar period—1.1% for services and 2.1% for goods (Duarte and Restuccia, 2010).

Government. The total number of bonds in the economy, B, corresponds to the outstanding public debt. These bonds have a very short maturity, one period, and do not carry any risk. The total amount is set to 1, which corresponds to approximately 90% of the model's annual GDP.

Monetary Authority. Inflation in the steady state is 0, and the nominal interest rate in the steady state is set to 0.75% to match an annual interest rate of 3%. Following the New-Keynesian literature (e.g., Auclert et al., 2024), the central bank follows a Taylor rule as in Equation (20), with corresponding weights on inflation $\phi_\pi = 1.5$ and zero weight on the output gap. Finally, the innovations of the monetary policy follow an AR(1) process with a persistency of 0.5 (e.g., Galí, 2008).

Summary. Table 2 lists the values of all pre-estimated parameters, as well as the source.

4.2.3 Simulated Method of Moments

The last step uses SMM to estimate the remaining parameters: The disutility from working, χ, and the taste-shifter, Ω. I choose the values of these two parameters to minimize the loss function below:

$$\mathcal{L}(\Omega, \chi) = ||M_m - M_d||, \tag{25}$$

which corresponds to the squared Euclidean norm between model moments, M_m, and data moments, M_d. I have two parameters to estimate, and I target two moments in the data, which yield an exactly identified system: The average fraction of hours worked between 1980 and 2023 and the consumption share in 2019. The first data moment is sourced from the Organization for Economic Co-operation and Development (OECD) and the second from the BEA.

The model moments are computed in the steady state. In Appendix C, I describe the

Table 2: Exogenous Parameters

Parameter	Description	Value	Source
I. Household			
β	Discount factor	0.99	Standard (quarterly model)
γ	CRRA	1.20	Standard
η	Inverse Frisch elasticity	1.00	Chetty et al. (2011)
ρ_z	Persistence of idiosync. productivity	0.99	Krueger et al. (2016)
σ_z	Std. dev. of idiosync. productivity	0.10	Krueger et al. (2016)
II. Firm			
θ_g	Elasticity of substitution (goods)	5.8	Marto (2024)
θ_s	Elasticity of substitution (services)	4.7	Marto (2024)
κ_g	Price adjustment cost (goods)	8.5	Data
κ_s	Price adjustment cost (services)	89.2	Data
Z_g^{2019}	Goods productivity	1.0	standardized
Z_s^{2019}	Services productivity	0.6	match 2019 relative price
III. Government and Monetary Authority			
B	Total assets	1	90% GDP
i_{SS}	Steady-state interest rate	0.75%	3% annual rate
ϕ_π	Taylor rule weight on inflation	1.5	Standard NK literature
ρ_M	Persistence of MP innovations	0.5	Standard NK literature

Notes: The table shows the values of all parameters that are calibrated exogenously, i.e., taking the value of empirical studies that estimate them. The last column contains the respective source.

algorithm I follow to solve for the steady state and compute the necessary model moments. Table 3 displays the estimated parameter values, the model moments, and their data target.

Table 3: SMM: Calibration Fit

Moment	Model Mom.	Data Mom.	Data Source	Parameter	Param. Value
Average hours worked	0.217	0.212	OECD	χ	30.0
Average service share 2019	0.673	0.678	BEA	Ω	7.0

Notes: The table shows the fit of the parameters calibrated using SMM. The values of the parameters are chosen to minimize the loss function $\mathcal{L}$, from Equation (25).

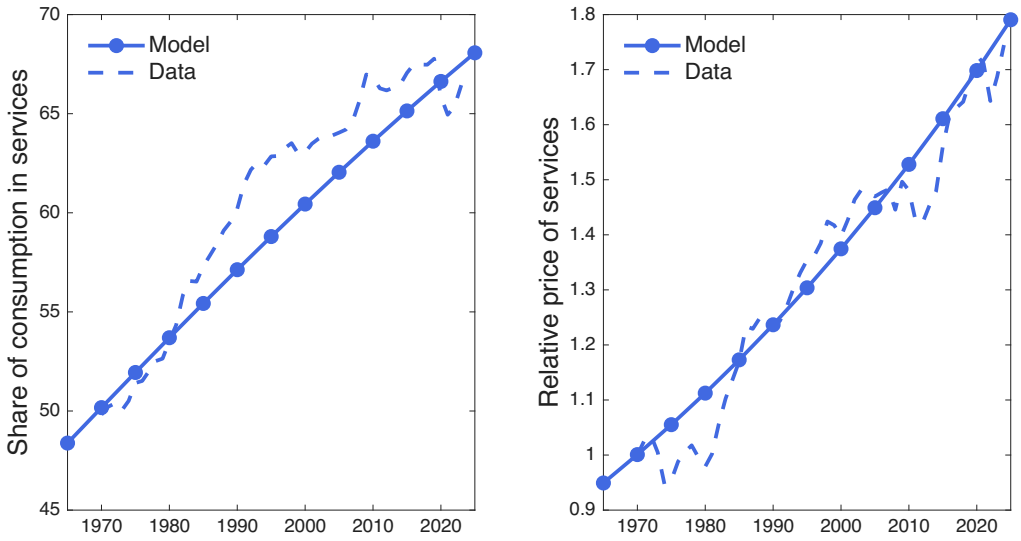
4.3 Model Fit

The model described in Section 3 aims to study the implications of the change in demand composition for monetary policy transmission. It is thus relevant that it matches both the long-run trend of the shift in the services share in the economy and the household micro-behavior that is central in the class of heterogeneous-agent models (Kaplan and Violante, 2022).

4.3.1 Model Fit across Time

Figure 4 evaluates long-run model fit by comparing the services share and the relative price in the model and data. Model values correspond to different steady states that only differ in terms of relative sectoral productivities. The values for 2019 are targeted in the calibration. However, the services share and relative price for other years are untargeted. They are generated through the income effect from the combination of the non-homothetic preferences and economic growth, and through the substitution effect generated by differences in the productivity growth rates of the two sectors that change the relative price.

Figure 4: Long-run Model Fit: Services Share and Relative Price



Notes: The figure shows the long-run fit of the model. The left panel shows the share of services, and the right panel shows the relative price. Full line is the model estimated values, and the dashed line corresponds to what is observed in the data.

The model replicates the data closely, with untargeted long-run moments aligning well with their empirical counterparts. In 1970, the model-implied services share exceeds the observed

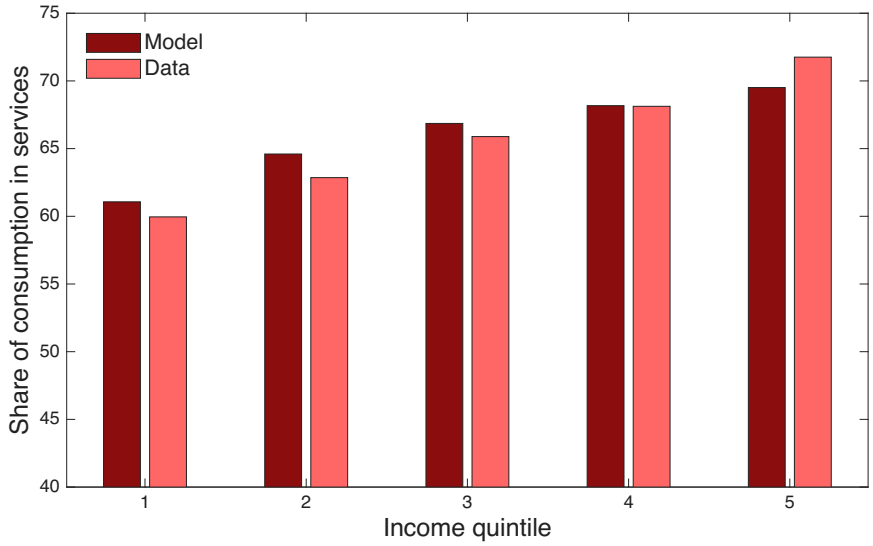
value by only 1.5 percentage points, and the relative price of services is marginally higher than in the data. Since the comparison is across steady states, the framework cannot capture the full transitional dynamics of structural transformation. The model also produces an average decline in hours worked of 0.1 percent per year, in line with the empirical trend from 1980 to 2020.

4.3.2 Model Fit in the Cross-Section

I now examine model fit in terms of household micro-behavior. I start by looking at the Engel curve of the services sector, which summarizes the differences in sectoral consumption allocations across households. I then look at the wealth and marginal propensity to consume distributions. These objects summarize the household heterogeneity in the model, which is key to understanding the aggregate response to monetary policy.

The Engel Curve of Services. Figure 5 shows the average consumption share allocated to the services sector for each income quintile, in both the model (dark red bars) and the data (light red bars). Income corresponds to the right-hand side of the household budget constraint in Equation (5). Data shares come from the CEX dataset, after applying the data cleaning described in Section 2.

Figure 5: The Engel Curve of Services: Data vs. Model



Notes: The figure shows the average expenditure share in the services sector for each quintile of income. Dark red bars correspond to the share obtained in the model, where income corresponds to the sum of earnings, dividends, and capital income—i.e., the right-hand side of Equation (5). Light red bars correspond to the estimated shares using the CEX data for the U.S. as described in Section 2.

The model replicates the positive slope of the Engel curve documented in Empirical Fact 1 in Section 2, with lower-income households allocating a smaller share of expenditure to services relative to higher-income households. Furthermore, quantitatively, the model-implied slope is close to its empirical counterpart, though slightly attenuated. The difference between the model and the data services share is less than 1 percentage point. Importantly, this relationship is not directly targeted in the calibration: Estimation of the income and price elasticities of demand, ϵ and σ, does not mechanically determine the slope of the Engel curve, except for the sign. This heterogeneity in the composition of consumption baskets means that an aggregate shock that has different sectoral impacts could affect households differently.

Wealth Distribution. Table 4 shows different statistics regarding the wealth distribution in the model and the data. The model does a relatively good job of matching the model and data wealth moments, especially considering that these are untargeted moments. The ratio of mean wealth to mean earnings in the model is in the same order of magnitude as in the data. The share of wealth owned by the bottom 50% of the wealth distribution is also in the same range observed in the data, as well as the share owned by the top 10% of the wealth distribution.

Table 4: Model Wealth Distribution

Wealth Statistic	Data	Model
Mean wealth	4.1	4.4
Median wealth	1.5	1.8
Wealth, bottom 50%	2.5%	3.1%
Wealth, top 10%	49.9%	48.6%
HtM share	17.3%	23.4%

Notes: The table shows some moments of the wealth distribution in the data and the model. The mean wealth corresponds to the ratio of the mean wealth to the mean earnings. HtM stands for Hand-to-Mouth and corresponds to the share of households with zero savings. All data moments come from the 2019 wave of the Survey of Consumer Finances, computed by Kaplan and Violante (2022) and Kaplan et al. (2014).

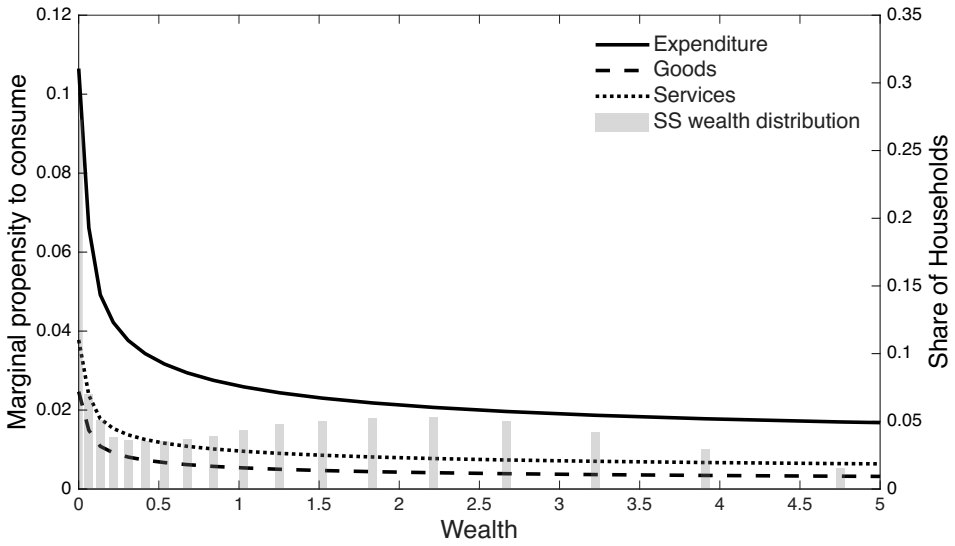
An important dimension of the wealth distribution is the share of hand-to-mouth (HtM) households—agents with zero savings. Achieving a close match to the wealth distribution, and in particular to the share of HtM households, is essential for capturing the economy's response to aggregate shocks such as monetary policy shocks. HtM agents exhibit strong consumption and labor responses because they do not have savings that allow them to smooth shocks over time. The model slightly overestimates this share relative to the data, as reported in the table.

Marginal Propensity to Consume. I now look at the marginal propensity to consume (MPC). This object is important in determining the transmission of aggregate shocks, such as monetary policy (Kaplan et al., 2018, Kaplan and Violante, 2022). The MPC is the fraction of a small, unanticipated one-time additional income that a household spends within a given period. For a household with productivity ω and asset holdings b, the quarterly MPC out of an unexpected windfall of money, M, is computed as

$$MPC^q = \frac{c_m(\omega, b + M) - c_m(\omega, b)}{M}. \tag{26}$$

Figure 6 plots the MPC for total expenditure (full line), goods (dashed line), and services (dotted line). The aggregate annual MPC of total expenditure is 28%, which is within the range of empirical estimates between 20% and 60% out of transitory income changes between \$500 and \$1,000.[21] For an overview of recent estimates of MPCs, refer to Havranek and Sokolova (2020). Boehm et al. (2025) have a recent empirical application using a randomized control trial in France, and Orchard et al. (2025) argue that the general equilibrium MPC that is consistent with the observed macro effects of the 2008 rebates is only 0.1.

Figure 6: Quarterly Marginal Propensity to Consume



Notes: The figure shows the marginal propensity to consume (MPC) out of a 5% of the average labor income unexpected transfer. MPCs are computed according to Equation (26) and averaged across productivity. The solid line shows the MPC for total expenditure, the dashed line the MPC for goods consumption, and the dotted line the MPC for services consumption. Gray bars in the background correspond to the mass of agents with a given wealth.

[21] I compute the annual MPC via the formula $MPC_a = 1 - (1 - MPC^q)^4$. The quarterly MPC is 8%.

The figure also shows that the MPC out of a transitory windfall is higher for services than for goods. In fact, the aggregate MPC of services is 5%, whereas the aggregate MPC of goods is 2%. This difference is rooted in non-homothetic preferences. Consumption in the goods sector is a necessity, whereas consumption in the services sector is a luxury. When households receive an unexpected windfall of money, they spend more on services than on goods, as their necessities for goods are already satisfied. The figure also shows a decaying MPC along the wealth distribution. Poorer households have a higher MPC than richer households. For households with no assets, the MPC of expenditure is 10%. As we move along the wealth distribution, the MPC converges to 0.5%, the MPC under no uncertainty.[22]

5 Structural Transformation and the Transmission of Monetary Policy

The primary objective of this paper is to quantify the impact of the shift in economic activity toward services on the transmission of monetary policy. To do that, I consider two economies that are similar to each other except for their sectoral productivities, which endogenously generate different services shares, as described in Section 4. The first economy has a services share of 50%, which corresponds to the share observed in 1970, and the second economy has a services share of 68%, corresponding to the service share in 2019.[23]

In each economy, I consider an experiment in which, at time $t = 0$, there is an unexpected monetary shock that corresponds to a quarterly innovation to the Taylor rule (20) of $\varepsilon_0^M = 0.01$—i.e., 100 basis points—which means-revert at rate ρ_M. I solve for the transition paths of prices and real aggregates where the shock is treated as an "M.I.T. shock": The change in the interest rate is a surprise, but once it is realized, its path is known. Figure D.8 in the Appendix plots the path of the monetary policy shock. This then generates impulse response functions to the monetary policy shock, which are the main object of interest in this section.[24] In Appendix C, I describe the algorithm used in detail.

[22] This corresponds to $i/(i + p_b)$.

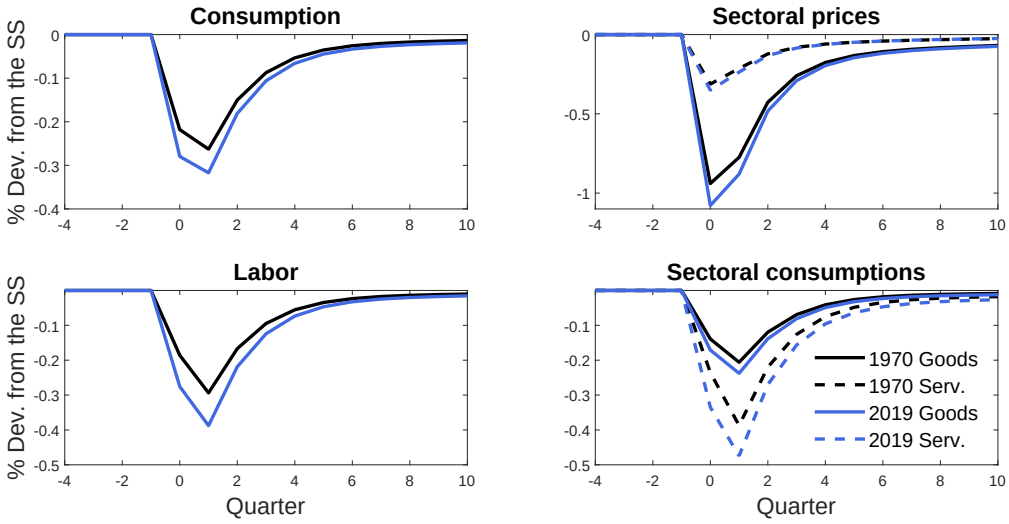
[23] The policy functions and wealth distribution of these two economies can be found in Appendix D.1.

[24] Studying the economy's deterministic response to shocks that are ex ante unexpected is useful for understanding its response to recurring aggregate shocks; see Boppart et al. (2018).

5.1 Aggregate Response to a Monetary Shock

Figure 7 plots the aggregate impulse response functions to a one percentage point contractionary monetary policy for both economies. The top-left panel shows the consumption index response, the top-right panel the path of sectoral prices, the bottom-left panel the path of aggregate labor, and the bottom-right panel the path of sectoral consumption. The black line corresponds to the 1970 calibration and the blue line to the 2019 calibration.

Figure 7: Aggregate Responses to a Monetary Policy Contraction: 1970 vs. 2019



Notes: The figure shows the aggregate response after a monetary policy contraction of 100 basis points for two model calibrations. In black, I show the responses when the services share is equal to the one observed in 1970, and in blue when the services share is equal to the share observed in 2019. All responses are shown in percentage deviations from steady-state values.

A contractionary monetary policy shock induces a negative demand contraction, which generates recessionary effects on both consumption and output. These effects operate through both direct and indirect channels. The direct channel arises because the nominal interest rate enters the household budget constraint (5), causing consumption to respond immediately to changes in the interest rate. This response reflects both an intertemporal substitution effect—where higher interest rates incentivize saving due to increased returns— and an income effect, as assets are in positive net supply and thus yield higher returns.

The indirect channel operates through the general equilibrium effects of monetary policy on wages and taxes. As the nominal interest rate increases, non-hand-to-mouth households reduce their consumption. In response, firms lower their labor demand, putting downward pressure on wages, which leads households to further reduce consumption. In addition, a rise

in i_t increases government interest payments and requires higher taxes to maintain a balanced budget. The resulting decline in disposable income further depresses household consumption.

Focusing on the 1970 economy, represented by the black lines in Figure 7, aggregate consumption falls by 0.2% and aggregate labor by 0.3%. These aggregate responses qualitatively align with empirical evidence from U.S. aggregate time series VARs by Christiano et al. (2005) or, more recently, using Norwegian administrative microdata by Holm et al. (2021).

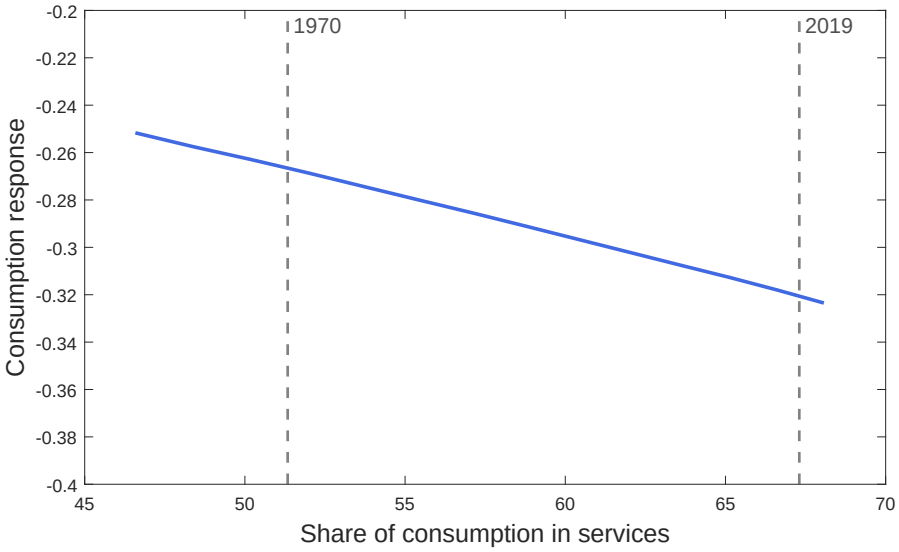
There are, however, differences across sectors. In response to the 100 basis point increase in nominal interest rate, the price of services falls by less than the price of goods. In fact, on impact, the price of services falls by 0.31% and the price of goods falls by almost 1% compared with their steady-state values. This is consistent with Empirical Fact 1 in Section 2, which shows nominal price rigidities are stronger in the services sector than in the goods sector. Because of that, and through the New-Keynesian channel, when prices do not adjust sufficiently for markets to clear, quantities have to change. As such, consumption in the services sector falls more than consumption in the goods sector, as prices in the services sector adjust less than in the goods sector. While the fall in services consumption is 0.4%, goods consumption only falls by 0.2% compared with their steady-state value. These differences across sectors are grounded in the different slopes of the sectoral New-Keynesian Phillips curves, due to the differences in the cost of price adjustment and markups/elasticity of substitution, and in differences in the demand income elasticity generated by the non-homothetic preferences. I decompose the differences in the responses, in Appendix D.2.

When I compare the 1970 responses with the 2019 responses, I find that the aggregate responses are larger when economic activity is more concentrated in the services sector. Aggregate consumption responds 21% more on impact in the 2019 economy than in the 1970 economy, and employment responds 32% more. Sectoral consumption and prices also react more strongly with a higher services share. In the goods sector, prices respond 14% more and consumption 15% more on impact in the 2019 economy relative to the 1970 economy. In the services sector, prices fall 11% more and consumption 23% more on impact when the services share in the economy is larger.

Figure 8 generalizes the above comparison, plotting the aggregate response of consumption on impact to the same monetary policy shock of 100 basis points for different services shares. As the composition of the economic activity changes and the share of services rises, the recessionary effects of monetary policy also get larger. In fact, a one percentage point higher services share

raises the consumption response to a monetary policy contractionary shock by 1.2%. Between 1970 and 2019, the services share rose by almost 18%, leading to the 21% higher response reported above.

Figure 8: Aggregate Consumption Responses to a Monetary Policy Contraction



Notes: The figure shows the aggregate response after a monetary policy contraction of 100 basis points for different shares of consumption in services. Responses are shown in percentage deviations from steady-state values. The two dashed vertical lines indicate the services share observed in 1970 and 2019, respectively.

5.2 Mechanisms

5.2.1 Income and Substitution Effects of Structural Transformation

The theory of structural transformation implicit in my model shifts economic activity toward services through two separate forces. First, sectoral productivities grow at different rates, which affects the relative price of services and changes expenditure shares (Baumol, 1967). Second, growth generates an income effect through non-homothetic preferences, which shift consumption toward luxuries—i.e., services (Falkinger, 1990, Matsuyama, 1992).

Table 5 compares the importance of these two channels for the transmission of monetary policy. Column (1) shows the baseline response differential between 1970 and 2019, column (2) shows the consumption response if there were income effects only between 1970 and 2019, and column (3) shows the consumption response if substitution effects were only operating between 1970 and 2019. The income effect counterfactual is obtained by increasing the overall

income size of the economy while keeping the relative price of 1970. On the other hand, the substitution effect counterfactual is obtained by changing the relative price but keeping overall income constant. I get this by changing sectoral productivities.

Table 5: Counterfactuals of Monetary Policy Channels: Income and Substitution Effects of Structural Transformation

	(1) 1970 (bsl)	(1) 2019 (bsl)	(2) Income effect	(3) Substitution effect
Consump. response (vs. 1970)	–	20.64	11.52	14.02
Relative price	1.00	1.68	1.00	1.68
SS consumption	0.03	0.05	0.05	0.03
Service share	51.34	67.30	58.09	61.22

Notes: The table shows how much the consumption response to a monetary policy shock changes in the baseline and under an income effect and a substitution effect counterfactuals. The income effect counterfactual is obtained by changing the sectoral productivities such that aggregate income is the same as in 2019, but the relative price is that of 1970. The substitution effect counterfactual is obtained by changing the sectoral productivities such that aggregate income is the same as in 1970 and the relative price is equal to the one in 2019.

Income effects lead to a 12% increase in the consumption response to monetary policy, whereas substitution effects are responsible for a 10% increase in the consumption response. This is related to how much the services share increases due to each of these effects. The income effect is responsible for a 6.75 percentage point increase in this share, whereas the substitution effect is responsible for almost a 10 percentage point increase in the services share.

5.2.2 The Role of Heterogeneous Price Rigidities and Non-homothetic Preferences

In Section 2, I showed that goods and services differ from each other in two respects. First, prices in the services sector adjust less frequently than in the goods sector. In the model, this is reflected in a higher price adjustment cost for the services sector, and in turn a flatter Phillips curve for this sector. Since the economy is more concentrated in the services sector, the aggregate Phillips curve is also flatter, which translates into stronger effects of monetary policy. Second, I have also shown that the demand composition between goods and services changes along the income distribution. I model this through non-homothetic preferences. This second difference means that the sectoral marginal propensity to consume differs across sectors and influences the aggregate demand response to monetary policy.

Table 6 compares aggregate responses of the baseline calibration and two counterfactual exercises designed to isolate the role of heterogeneous price adjustment costs and non-

homothetic preferences. Column (1) reports the baseline, replicating the analysis from the previous subsection. The rise in the services share between 1970 and 2019 implies that the same monetary policy shock generates a 21% larger response of aggregate consumption, but only a 14% larger response in the price of goods, and 11% for the price of services.

Table 6: Counterfactuals of Monetary Policy Channels: Homotheticity and Homogeneous Sectoral Costs of Price Adjustment

	(1) Baseline		(2) Homog. κ_m		(3) Homothetic	
	1970	2019	1970	2019	1970	2019
Service share	51.3	67.3	51.3	67.3	51.0	67.2
MPC	8.1	7.6	8.1	7.6	8.6	8.4
Consump. response (% change vs. 1970)	20.6		3.5		24.1	
Price of goods response (% change vs. 1970)	13.7		5.9		6.3	
Price of serv. response (% change vs. 1970)	10.7		5.9		3.5	

Notes: The table compares the response of consumption and sectoral prices to a contractionary monetary policy shock in an economy with the services share equal to the one observed in 2019 and an economy with the services share equal to the one observed in 1970. Column (1) shows the response for the baseline scenario, column (2) shows the response when the cost of price adjustment is homogeneous across sectors, and column (3) when preferences are homothetic.

Column (2) shows the responses for the counterfactual, in which the costs of price adjustments are homogeneous across sectors. I set $\kappa_g = \kappa_s = 56$, equal to the value used by Hagedorn et al. (2019). As this only affects the agents' decisions after a monetary policy shock, the 1970 and 2019 steady states are equal to the baseline, and hence the same share of hand-to-mouth households and marginal propensity to consume. However, having homogeneous costs of price adjustment has consequences for the transmission of monetary policy. In fact, consumption responses are now only 3.5% larger in 2019 than in 1970. This demonstrates the importance of heterogeneous price rigidities in understanding how sectoral composition and structural transformation impact the transmission of monetary policy.

Column (3) shows responses for the counterfactual when preferences are homothetic.[25] To build this counterfactual, I make preferences homothetic by setting $\epsilon = 1$ and recalibrate the taste-shifter, Ω, to match the baseline model services share in 1970 (51.3%) and 2019 (67.3%). This way, the production side of the economy is unchanged compared with the baseline, and all differences come from changing to a homothetic preference relation.[26] The aggregate difference

[25] In Appendix D.1, I plot the steady-state policy functions of this counterfactual exercise.

[26] Note that Ω does not affect households' decisions but rather the optimal preferred mix between services and

in consumption responses between 1970 and 2019 increases from 21% to 24%, which indicates that differences in the demand composition dampen the effects on aggregate demand that monetary policy has.

The intuition is that non-homothetic preferences generate different income elasticities of demand. Goods are necessities whose share in household consumption decreases with income, while services are luxuries whose share increases with income. As a result, being at the borrowing constraint has an additional downside: households may be unable to consume enough necessities. To prevent this, they raise their precautionary savings, leading to greater wealth accumulation. Because households are now wealthier, their marginal propensity to consume is lower in the non-homothetic world than in the homothetic one. The table above illustrates this mechanism. Consequently, the overall responsiveness of aggregate demand to monetary policy shocks is lower.

5.3 Distributional Responses to a Monetary Shock and Welfare Implications

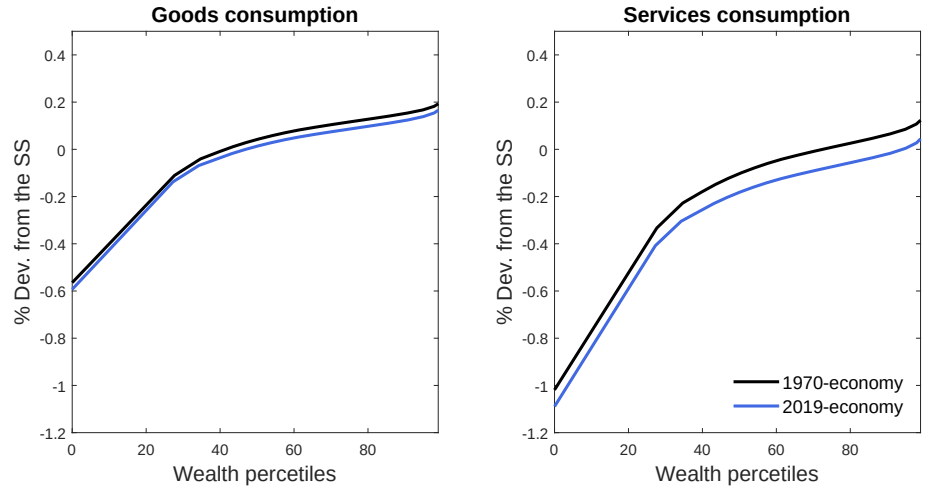
I now look into the distributional effects of a monetary policy shock and show that the share of services in the economy matters for welfare consequences. Figure 9 shows the responses of goods and services consumption over the wealth distribution. Similarly to the above, black lines correspond to responses when the services share is equal to 51% (the 1970 economy) and the blue lines to responses when the services share is equal to 68% (the 2019 economy).

Focusing on the left panel with goods consumption heterogeneous responses, there is a clear divide between low- and high-wealth households. Households in the bottom part of the wealth distribution respond negatively to the contractionary shocks, whereas households in the top part of the distribution increase their consumption of goods. In particular, the bottom 20% of the wealth distribution, which includes HtM households, decreased their consumption of goods between 0.4% and 0.6% compared with what they were consuming before the shock. The top 20% of the wealth distribution, on the other hand, increase their consumption of goods between 0% and 0.4% compared with the steady-state value.

The right panel shows a similar pattern, although the differences along the distribution of wealth are larger. For this sector, the bottom 20% of the wealth distribution decreases their consumption between 0.6% and 1%, whereas the top 20% of the wealth distribution increases

goods. This strategy is similar to that of Hochmuth et al. (2025) to disentangle the effects of non-homothetic preferences on the distributional consequences of becoming climate neutral.

Figure 9: Sectoral Consumption Responses to a Monetary Policy Contraction by Wealth Percentile



Notes: The figure shows consumption of goods and the consumption of services responses after a monetary policy contraction of 100 basis points by wealth percentile. Black lines depict the response for the 1970 economy and blue lines for the 2019 economy. All values are expressed as percentage deviations from steady-state values.

their consumption by an amount similar to the goods consumption: 0–0.2%.

As shown above, a higher share of services in the economy increases its responsiveness to monetary policy shocks. This pattern also holds across the wealth distribution. For the consumption of goods, the decrease between 1970 and 2019 is more modest than for the consumption of services. The difference in goods consumption responses between 1970 and 2019 is relatively constant along the wealth distribution. For services consumption, the difference enlarges along the wealth distribution, with high-wealth households decreasing their services consumption more than low-wealth households. Figure D.9 in the Appendix complements this analysis and plots the impulse response functions of consumption and labor supply. For the labor response, in particular, it shows a decrease between the 1970 and the 2019 economy. However, since low-wealth households have positive responses and high-wealth households have negative responses to the monetary policy tightening shock, this means that those at the top are better off in this regard than those at the bottom of the distribution.

The Role of Non-homothetic Preferences for Distributional Consequences. Figure D.10 in the Appendix contrasts sectoral consumption responses across the wealth distribution in the baseline economy and in a homothetic counterfactual. The counterfactual sets $\epsilon = 1$ and recalibrates Ω to match the aggregate services share, as in the previous subsection. With

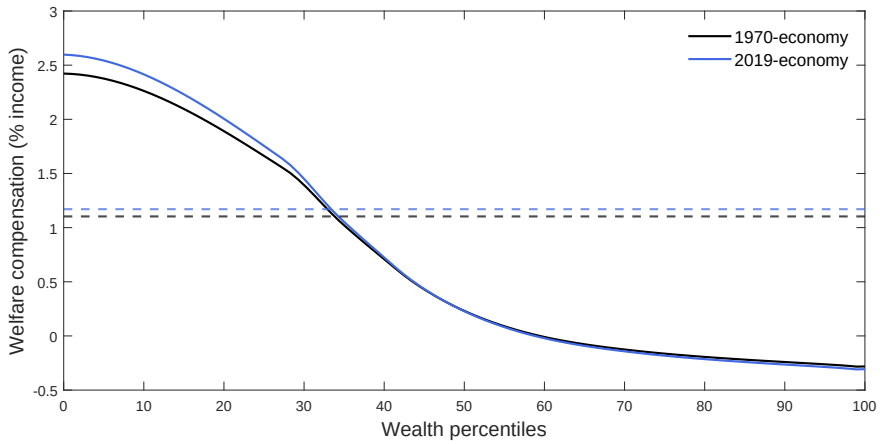
homothetic preferences, the differences along the wealth distribution in goods consumption responses are larger relative to the baseline. Under non-homothetic preferences, households at the bottom of the distribution reduce goods consumption less, and those at the top also reduce it less. For services, non-homotheticities render the consumption responses uniformly stronger—especially at the bottom of the distribution, since services are considered luxuries.

Welfare. These differences in responses to a monetary shock, together with differences in the consumption baskets in terms of goods and services, motivate a welfare analysis. For each household, I compute the welfare loss associated with a monetary contractionary shock by calculating the asset transfer, m, in the period when the shock happens that renders them indifferent to the steady-state:

$$V^0(\omega, b) = V^1(\omega, b + m).$$

Figure 10 shows the value o m in percentage of income, for the median productivity ω state, along the wealth distribution. A positive value means that a positive transfer is needed to restore welfare, which implies that monetary policy is welfare-harmful.

Figure 10: Welfare Compensation by Wealth Position



Notes: The figure plots the welfare compensation to a 100 basis points monetary policy tightening over wealth. Positive values mean that a positive transfer is required to restore the steady-state welfare—i.e., that welfare decreased with the increase in the nominal interest rate. Values are in percentage of the household income. Dashed horizontal lines correspond to the aggregate welfare loss. 1970 economy values are represented in black, and 2019 economy values in blue.

Independent of the services share level in the economy, monetary policy tightening leads

to a loss of welfare, such that a positive compensation of around 1% of the income is needed. Households at the bottom part of the wealth distribution are the ones more negatively affected by the shock since they are the ones with higher losses of consumption, as shown above. On the other hand, households at the top of the wealth distribution benefit from the shock.[27]

This means that a monetary shock increases welfare inequality. In general, an increase in the interest rate affects households in two ways: A direct effect through an increase in the nominal return of assets and an indirect effect as the government has to increase taxes to finance the higher interest payments. These two channels explain why low-wealth households are worse off from the negative shock—the increase in taxes is higher than the increase in the return—and high-wealth households are better off—the increase in the nominal return compensates the higher tax payments.

Compared with the 2019 economy, the shift in economic activity toward the services sector also changed the welfare implications of monetary policy: The aggregate welfare cost increases by 5% of the income. In terms of the welfare distributional effects of the monetary shock, households with lower wealth levels now have a higher cost. Those with no assets need to receive compensation 7% larger in the 2019 economy compared with the 1970 economy. On the other hand, households with higher wealth levels are slightly better off with the monetary shock in the 2019 economy than they were in the 1970 economy. This implies that as the service share increases, monetary policy becomes more powerful—which also increases the inequality generated.

6 Structural Transformation and Supply Shocks

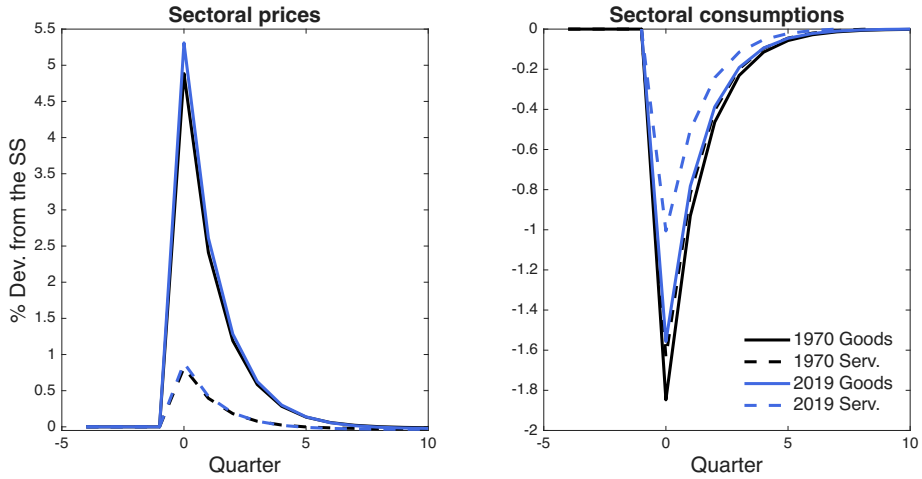
In this section, I extend the analysis by comparing the effects of supply shocks in economies with different services shares in consumption. I consider the same two economies of Section 5: The 1970 economy and the 2019 economy that have a services share of 51% and 67%, respectively. Except for differences in the shares of services that are generated by differences in relative sectoral productivity, these economies are identical. These two years are particularly relevant in this context, because each experienced significant supply shocks—the oil crises in the 1970s and COVID-related disruptions in 2020 and successive years with the Ukraine invasion.[28]

[27] Coibion et al. (2017) find that contractionary monetary policy increases income and consumption inequality.
[28] COVID is commonly thought of as a supply shock that caused demand shortages (Guerrieri et al., 2022).

In each economy, I consider an experiment in which, at time $t = 0$, there is an unexpected TFP shock. This shock is common to both sectors, such that $\Delta Z_m / Z_m = -5\%$. The shock follows an AR(1) process with persistence 0.5. I solve for the transition paths of prices and real aggregates, treating the shock as an "M.I.T. shock".

Aggregate Effects. Figure 11 plots the aggregate impulse response functions to this shock for both the low- and high-services economy. The top left panel shows the consumption index response, the top right panel the path of sectoral prices, the bottom left panel the path of aggregate labor, and the bottom right panel the path of sectoral consumption. The black line corresponds to the 1970 calibration and the blue line to the 2019 calibration.

Figure 11: Aggregate Responses to a Negative Supply Shock: 1970 vs. 2019



Notes: The figure shows the aggregate response after a negative TFP shock of 5% for two model calibrations. In black, I show responses when the services share is equal to the one observed in 1970, and in blue when the services share is equal to the share observed in 2019. All responses are presented as percentage deviations from steady-state values.

A negative supply shock leads to an increase in production costs. In response, firms increase prices. However, due to differences in the frequency of price adjustment between the two sectors, the price of goods adjusts more than the price of services. This means that consumption falls, with goods consumption decreasing more than services consumption. This generates an increase in the services share, given that the relative price of services decreased. These responses align with the estimated impulse response functions to oil shocks by Känzig (2021) for the U.S. and Broer et al. (2025) for Germany.

In fact, focusing on the 1970 economy's responses, I find that the negative TFP shock leads

to a contemporaneous increase in the price of goods by 4.9% and the price of services by 0.8%. As a result, goods consumption falls by 1.8% and services consumption by 1.6%. The smaller decline in goods consumption, relative to the large price increase, reflects the non-homothetic nature of preferences: Goods are a necessity, so households cannot reduce their consumption proportionally to the price change. Overall, this implies that the consumption index—a proxy for welfare—declines by 1.3% immediately following the negative TFP shock.

Comparing the 1970 economy responses with 2019 economy responses, I find that the overall recessionary effects of the negative TFP shock are milder. Aggregate consumption responses are 37% smaller and sectoral consumption responses are 16% smaller for goods and 38% for services. Since goods are the sector in which prices rise more steeply, a larger services share implies that a smaller portion of the consumption basket is exposed to this adjustment. Consequently, aggregate consumption contracts less, and output falls by a smaller amount, as the services share in the economy increases.

Mechanisms. Table 7 compares the aggregate responses of the baseline calibration and the same two counterfactual exercises that isolate the effect of heterogeneous price rigidities and non-homothetic preferences. Column (1) reports the baseline, corresponding to the responses reported above. Consumption responses are 37% smaller in the 2019 economy than in the 1970 economy. On the other hand, price responses are slightly larger.

Table 7: Counterfactuals: Homotheticity and Homogeneous Sectoral Costs of Price Adjustment

	(1) Baseline		(2) Homog. κ_m		(3) Homothetic	
	1970	2019	1970	2019	1970	2019
Service share	51.3	67.3	51.3	67.3	51.0	67.2
MPC	8.1	7.6	8.1	7.6	8.6	8.4
Consump. response (% change vs. 1970)	-37.3		-9.2		-19.9	
Price of goods response (% change vs. 1970)	8.3		0.9		4.9	
Price of serv. response (% change vs. 1970)	6.9		0.9		4.1	

Notes: The table compares the response of consumption and sectoral prices to a negative supply shock in an economy with the services share equal to the one observed in 2019 and an economy with the services share equal to the one observed in 1970. Column (1) shows it for the baseline scenario, column (2) when the cost of price adjustment is homogeneous across sectors, and column (3) for the case in which preferences are homothetic.

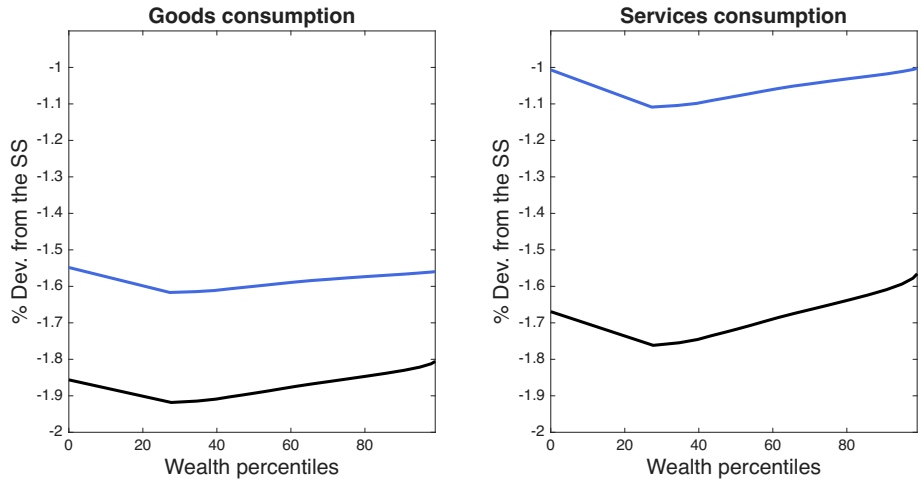
Column (2) shows responses for the counterfactual when the costs of price adjustments are homogeneous across sectors. As in Section 5, I set $\kappa_g = \kappa_s = 56$. Having homogeneous costs of

price adjustment renders the responses in the 2019 economy more similar to responses in the 1970 economy. Heterogeneous price rigidities contribute to 75% of the effects—a value very similar to what I find in the case of the transmission of monetary policy.

Column (3) shows responses for the counterfactual when preferences are homothetic. To build this counterfactual, I make preferences homothetic by setting $\epsilon = 1$ and recalibrate the taste-shifter, Ω, to match the services share in 1970 (51.3%) and 2019 (67.3%). This way, the production side of the economy is unchanged compared with the baseline, and all differences come from changing to a homothetic preference relation. The aggregate consumption responses difference between 1970 and 2019 decrease from -37% to -20%. This means that non-homothetic preferences also contribute to smoothing the recessionary supply shocks' effects.

Heterogeneous Effects. Figure 12 shows the responses of goods and services consumption along the wealth distribution. The black line corresponds to 1970 economy responses when the services share is low, and the blue line to the 2019 economy when the services share is high.

Figure 12: Sectoral Consumption Responses to a Negative Supply Shock by Wealth Percentile



Notes: The figure shows consumption of goods and the consumption of services responses after a negative aggregate TFP shock of 5% by wealth percentile. Black lines depict the response for the 1970 economy and blue lines for the 2019 economy. All values are expressed as percentage deviations from steady-state values.

The figure shows that, contrary to the monetary policy shock, all households decrease their consumption after the negative supply shock. Differences across the wealth distribution are also much smaller. However, the difference in consumption responses between 1970 and 2019 is larger for low-wealth than for high-wealth households.

7 Conclusion

Over the past 50 years, the share of services in total expenditure has risen from about 50% to nearly 70%. At the same time, the Phillips curve has flattened, indicating stronger monetary effectiveness. This paper argues that structural transformation can explain the stronger monetary policy transmission: The demand shift toward services increases aggregate price rigidity because services prices adjust less frequently than goods prices. The resulting higher rigidity flattens the Phillips curve and amplifies the short-run real effects of monetary policy.

I rationalize this in a dynamic two-sector heterogeneous-agent model. The model features non-homothetic preferences and heterogeneous nominal price rigidities, and it is estimated for the United States. By changing the sectoral productivity, the model endogenously generates economies with different shares of services. In each of the economies, I study the transmission of monetary policy by doing a policy experiment that increases the nominal interest rate.

I find that in the last 50 years, monetary policy became 21% more powerful when analyzing the consumption responses, with the sectoral differences in price rigidities explaining 80% of the amplification result. Furthermore, structural transformation has also amplified the unequal effects of monetary policy. A contractionary monetary policy shock hurts low-wealth households, while benefiting high-wealth households. The shift toward services widens this gap. Finally, in an extension, I look at the consumption responses to negative supply shocks. I show that a larger service share makes consumption more resilient to adverse supply shocks, as prices in services are more stable.

These findings carry implications for policymakers, especially central bankers. In the United States and in the euro area, a single monetary authority sets the interest rate for economies that differ widely in their sectoral composition. The results help explain why monetary policy does not affect all countries in the same way and shed light on the heterogeneous responses observed across Europe. Finally, as populations age and older households consume relatively more services, demand will continue to shift toward services, and monetary policy transmission is likely to strengthen further.

References

Acemoglu, D. and V. Guerrieri (2008). Capital Deepening and Nonbalanced Economic Growth. *Journal of Political Economy 116*(3), 467–498. *(Cited on page 21.)*

Aguiar, M. and M. Bils (2015). Has Consumption Inequality Mirrored Income Inequality? *American Economic Review 105*(9), 2725–2756. *(Cited on pages 12, 22, 23, and 56.)*

Ait-Sahalia, Y., J. A. Parker, and M. Yogo (2004). Luxury Goods and the Equity Premium. *The Journal of Finance 59*(6), 2959–3004. *(Cited on page 5.)*

Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics 109*(3), 659–684. *(Cited on page 14.)*

Almgren, M., J.-E. Gallegos, J. Kramer, and R. Lima (2022). Monetary Policy and Liquidity Constraints: Evidence from the Euro Area. *American Economic Journal: Macroeconomics 14*(4), 309–340. *(Cited on pages 4 and 9.)*

Andreolli, M. and P. Surico (2025). Less Is More: Consumer Spending and the Size of Economic Stimulus Payments. *American Economic Journal: Macroeconomics*. *(Cited on page 5.)*

Arvai, K. and K. Mann (2025). Consumption Inequality in the Digital Age. *Banque de France Working Paper* (890). *(Cited on page 5.)*

Ascari, G. and L. Rossi (2012). Trend Inflation and Firms Price-Setting: Rotemberg versus Calvo. *The Economic Journal 122*(563), 1115–1141. *(Cited on page 24.)*

Atkeson, A., L. E. Ohanian, et al. (2001). Are Phillips Curves Useful for Forecasting Inflation? *Federal Reserve Bank of Minneapolis Quarterly Review 25*(1), 2–11. *(Cited on page 8.)*

Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review 109*(6), 2333–2367. *(Cited on page 5.)*

Auclert, A., M. Rognlie, and L. Straub (2024). The Intertemporal Keynesian Cross. *Journal of Political Economy 132*(12), 4068–4121. *(Cited on page 25.)*

Auclert, A., M. Rognlie, and L. Straub (2025). Fiscal and Monetary Policy with Heterogeneous Agents. *Annual Review of Economics 17*, 539–562. *(Cited on page 4.)*

Ball, L. and S. Mazumder (2019). A Phillips Curve with Anchored Expectations and Short-Term Unemployment. *Journal of Money, Credit and Banking 51*(1), 111–137. *(Cited on page 2.)*

Baumol, W. J. (1967). Macroeconomics of Unbalanced Growth: the Anatomy of Urban Crisis. *The American Economic Review 57*(3), 415–426. *(Cited on pages 21 and 34.)*

Becker, J. (2024). Do Poor Households Pay Higher Markups in Recessions? *Working paper*. *(Cited on page 5.)*

Bernardino, T., P. M. S. Brinca, S. Darougheh, and M. Silva Pereira (2025). The Heterogeneous Effects of Supply Shocks in Necessity Goods. *Available at SSRN 5549678*. *(Cited on page 5.)*

Bewley, T. (1986). Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers. *Contributions to Mathematical Economics in Honor of Gérard Debreu 79*. *(Cited on page 14.)*

Bilbiie, F. O. (2020). The New Keynesian Cross. *Journal of Monetary Economics 114*, 90–108. *(Cited on page 5.)*

Bilbiie, F. O. (2025). Monetary Policy and Heterogeneity: An Analytical Framework. *Review of Economic Studies 92*(4), 2398–2436. *(Cited on page 5.)*

Bils, M. and P. J. Klenow (2004). Some Evidence on the Importance of Sticky Prices. *Journal of Political Economy 112*(5), 947–985. *(Cited on pages 7, 54, and 56.)*

Bobeica, E., M. Ciccarelli, and I. Vansteenkiste (2019). The Link between Labor Cost and Price Inflation in the Euro Area. *ECB Working Paper* (2235). *(Cited on page 8.)*

Boehm, J., E. Fize, and X. Jaravel (2025). Five Facts about MPCs: Evidence from a Randomized Experiment. *American Economic Review 115*(1), 1–42. *(Cited on page 30.)*

Boehnert, L., S. de Ferra, K. Mitman, and F. Romei (2025). High Non-Tradable Consumption, Inequality, and Weak Monetary Policy. *Working Paper*. *(Cited on page 5.)*

Boivin, J. and M. P. Giannoni (2006). Has Monetary Policy Become More Effective? *The Review of Economics and Statistics 88*(3), 445–462. *(Cited on page 4.)*

Bonciani, D. and J. Oh (2021). Unemployment Risk, Liquidity Traps, and Monetary Policy. *Staff Working Paper No. 920*. *(Cited on page 5.)*

Boppart, T. (2014). Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences. *Econometrica 82*(6), 2167–2196. *(Cited on pages 12 and 21.)*

Boppart, T., P. Krusell, and K. Mitman (2018). Exploiting MIT Shocks in Heterogeneous-Agent Economies: the Impulse Response as a Numerical Derivative. *Journal of Economic Dynamics and Control 89*, 68–92. *(Cited on pages 31 and 64.)*

Brinca, P., H. A. Holter, P. Krusell, and L. Malafry (2016). Fiscal Multipliers in the 21st Century. *Journal of Monetary Economics 77*, 53–69. *(Cited on page 9.)*

Broer, T., J. Kramer, and K. Mitman (2025). The Distributional Effects of Oil Shocks. *IMF Economic Review*. *(Cited on page 41.)*

Buiatti, C., J. Duarte, and F. Saenz (2024). Europe Falling Behind: Structural Transformation and Labor Productivity Growth Differences Between Europe and the U.S. *Cambridge-INET Working Paper Series* (2017/04). *(Cited on page 21.)*

Burriel, P. and A. Galesi (2018). Uncovering the Heterogeneous Effects of ECB Unconventional Monetary Policies across Euro Area Countries. *European Economic Review 101*, 210–229. *(Cited on page 4.)*

Calvo, G. A. (1983). Staggered Prices in a Utility-maximizing Framework. *Journal of Monetary Economics 12*(3), 383–398. *(Cited on page 24.)*

Carvalho, C. (2006). Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks. *The B.E. Journal of Macroeconomics 6*(3), 1–58. *(Cited on page 5.)*

Castro, N. (2019). The Importance of Production Networks and Sectoral Heterogeneity for Monetary Policy. *Available at SSRN 3499902*. *(Cited on page 5.)*

Cavalleri, M. C., A. Eliet, P. McAdam, F. Petroulakis, A. Soares, and I. Vansteenkiste (2019). Concentration, Market Power and Dynamism in the Euro Area. *ECB Working Paper*. *(Cited on page 7.)*

Chahad, M., A.-C. Hofmann-Drahonsky, A. Page, and M. Tirpák (2023). An Updated Assessment of Short-Term Inflation Projections by Eurosystem and ECB Staff. *Economic Bulletin Boxes 1*. *(Cited on page 3.)*

Challe, E. (2020). Uninsured Unemployment Risk and Optimal Monetary Policy in a Zero-Liquidity Economy. *American Economic Journal: Macroeconomics 12*(2), 241–283. *(Cited on page 5.)*

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins. *American Economic Review 101*(3), 471–475. *(Cited on page 23.)*

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy 113*(1), 1–45. *(Cited on page 33.)*

Clayton, C., X. Jaravel, and A. Schaab (2018). Heterogeneous Price Rigidities and Monetary Policy. *Available at SSRN 3186438*. *(Cited on page 5.)*

Coibion, O. and Y. Gorodnichenko (2015). Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation. *American Economic Journal: Macroeconomics 7*(1), 197–232. *(Cited on page 2.)*

Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2017). Innocent Bystanders? Monetary Policy and Inequality. *Journal of Monetary Economics 88*, 70–89. *(Cited on page 40.)*

Comin, D., D. Lashkari, and M. Mestieri (2021). Structural Change with Long-run Income and Price Effects. *Econometrica 89*(1), 311–374. *(Cited on pages 14, 15, 21, and 23.)*

Constâncio, V. (2017). Inequality and Macroeconomic Policies. Intervention by Vítor Constâncio, Vice-President of the ECB, at the Annual Congress of the European Economic Association Lisbon. *(Cited on page 3.)*

Corsetti, G., J. B. Duarte, and S. Mann (2022). One Money, Many Markets. *Journal of the European Economic Association 20*(1), 513–548. *(Cited on page 4.)*

Cravino, J., T. Lan, and A. A. Levchenko (2020). Price Stickiness along the Income Distribution and the Effects of Monetary Policy. *Journal of Monetary Economics 110*, 19–32. *(Cited on page 5.)*

Da-Rocha, J. M. and D. Restuccia (2006). The Role of Agriculture in Aggregate Business Cycles. *Review of Economic Dynamics 9*(3), 455–482. *(Cited on page 21.)*

Danieli, A. (2020). Expenditure Elasticity and Inequality Over the Business Cycle. *Available at SSRN 3593499*. *(Cited on page 5.)*

De Nardi, M. and G. Fella (2017). Saving and Wealth Inequality. *Review of Economic Dynamics 26*, 280–300. *(Cited on page 5.)*

Debortoli, D. and J. Galí (2025). Heterogeneity and Aggregate Fluctuations: Insights from TANK Models. *NBER Macroeconomics Annual 39*(1), 307–357. *(Cited on page 5.)*

Dhyne, E., L. J. Alvarez, H. L. Bihan, G. Veronese, D. Dias, J. Hoffmann, N. Jonker, P. Lünnemann, F. Rumler, and J. Vilmunen (2006). Price Changes in the Euro Area and the United States: Some Facts from Individual Consumer Price Data. *Journal of Economic Perspectives 20*(2), 171–192. *(Cited on pages 7 and 8.)*

Duarte, M. and D. Restuccia (2010). The Role of the Structural Transformation in Aggregate Productivity. *The Quarterly Journal of Economics 125*(1), 129–173. *(Cited on page 25.)*

Engel, E. (1857). Die Produktions- und Consumtionsverhältnisse des Königreichs Sachsen. *Zeitschrift des Statistischen Büreaus des Königlich Sächischen Ministeriums des Innern 8-9*(2), 1–54. *(Cited on page 12.)*

Engel, E. (1895). *Die Lebenskosten belgischer Arbeiter-Familien früher und jetzt*. Heinrich. *(Cited on page 12.)*

Falkinger, J. (1990). On Growth along a Hierarchy of Wants. *Metroeconomica 41*(3), 209–223. *(Cited on pages 21 and 34.)*

Flodén, M. (2008). A Note on the Accuracy of Markov-Chain Approximations to Highly Persistent AR(1) Processes. *Economics Letters 99*(3), 516–520. *(Cited on page 24.)*

Forbes, K. J. (2019). Inflation Dynamics: Dead, Dormant, or Determined Abroad? *Brookings Papers on Economic Activity*, 257–319. *(Cited on page 2.)*

Gaillard, A. and P. Wangner (2022). Inequality, Business Cycles, and Growth: A Unified Approach to Stabilization Policies. *Available at SSRN 4292481*. *(Cited on page 23.)*

Galesi, A. and O. Rachedi (2019). Services Deepening and the Transmission of Monetary Policy. *Journal of the European Economic Association 17*(4), 1261–1293. *(Cited on pages 4, 9, and 21.)*

Galí, J. (2008). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and its Applications*. Princeton University Press. *(Cited on pages 14 and 25.)*

Gautier, E., C. Conflitti, R. P. Faber, B. Fabo, L. Fadejeva, V. Jouvanceau, J.-O. Menz, T. Messner, P. Petroulas, P. Roldan-Blanco, et al. (2024). New Facts on Consumer Price Rigidity in the Euro Area. *American Economic Journal: Macroeconomics 16*(4), 386–431. *(Cited on pages 7 and 8.)*

Geary, R. C. (1950). A Note on "A Constant-Utility Index of the Cost of Living". *The Review of Economic Studies 18*(1), 65–66. *(Cited on page 59.)*

Ghassibe, M. (2024). Endogenous Production Networks and Non-Linear Monetary Transmission. *BSE, Barcelona School of Economics*. *(Cited on page 5.)*

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek (2017). Inflation Dynamics during the Financial Crisis. *American Economic Review 107*(3), 785–823. *(Cited on page 2.)*

Gorman, W. M. (1971). Two-Stage Budgeting. *Unpublished Paper, the London School of Economics and Political Science*. *(Cited on page 15.)*

Gornemann, N., K. Kuester, and M. Nakajima (2016). Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy. *CEPR Discussion Paper No. DP11233*. *(Cited on page 4.)*

Guerrieri, V., G. Lorenzoni, L. Straub, and I. Werning (2022). Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages? *American Economic Review 112*(5), 1437–1474. *(Cited on page 40.)*

Hagedorn, M., J. Luo, I. Manovskii, and K. Mitman (2019). Forward Guidance. *Journal of Monetary Economics 102*, 1–23. *(Cited on page 36.)*

Hagedorn, M., I. Manovskii, and K. Mitman (2019). The Fiscal Multiplier. *NBER Working Paper* (25571). *(Cited on page 18.)*

Hausman, J., G. Leonard, and J. D. Zona (1994). Competitive Analysis with Differenciated Products. *Annales d'Economie et de Statistique*, 159–180. *(Cited on page 22.)*

Hausman, J. A. (1996). *Valuation of New Goods under Perfect and Imperfect Competition*. National Bureau of Economic Research Cambridge, Mass., USA. *(Cited on page 22.)*

Havranek, T. and A. Sokolova (2020). Do Consumers Really Follow a Rule of Thumb? Three Thousand Estimates from 144 Studies Say "Probably Not. *Review of Economic Dynamics 35*, 97–122. *(Cited on page 30.)*

Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson (2022). The Slope of the Phillips Curve: Evidence from US States. *The Quarterly Journal of Economics 137*(3), 1299–1344. *(Cited on page 2.)*

Herrendorf, B., R. Rogerson, and A. Valentinyi (2014). Growth and Structural Transformation. *Handbook of Economic Growth 2*, 855–941. *(Cited on pages 5 and 20.)*

Hobijn, B. and D. Lagakos (2005). Inflation Inequality in the United States. *Review of Income and Wealth 51*(4), 581–606. *(Cited on page 54.)*

Hochmuth, P., P. Krusell, and K. Mitman (2025). Distributional Consequences of Becoming Climate-Neutral. *IZA Discussion Paper 17861*. *(Cited on pages 5 and 37.)*

Holm, M. B., P. Paul, and A. Tischbirek (2021). The Transmission of Monetary Policy under the Microscope. *Journal of Political Economy 129*(10), 2861–2904. *(Cited on page 33.)*

Huggett, M. (1993). The Risk-free Rate in Heterogeneous-agent Incomplete-insurance Economies. *Journal of Economic Dynamics and Control 17*(5-6), 953–969. *(Cited on page 14.)*

Imrohoroğlu, A. (1989). Cost of Business Cycles with Indivisibilities and Liquidity Constraints. *Journal of Political Economy 97*(6), 1364–1383. *(Cited on page 14.)*

Jaimovich, N., S. Rebelo, and A. Wong (2019). Trading Down and the Business Cycle. *Journal of Monetary Economics 102*, 96–121. *(Cited on page 5.)*

Jordà, Ò. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review 95*(1), 161–182. *(Cited on page 9.)*

Känzig, D. R. (2021). The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements. *American Economic Review 111*(4), 1092–1125. *(Cited on page 41.)*

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary Policy According to HANK. *American Economic Review 108*(3), 697–743. *(Cited on pages 4 and 30.)*

Kaplan, G., G. Violante, and J. Weidner (2014). The Wealthy Hand-to-Mouth. *Brookings Papers on Economic Activity 1*, 77–153. *(Cited on page 29.)*

Kaplan, G. and G. L. Violante (2022). The Marginal Propensity to Consume in Heterogeneous Agent Models. *Annual Review of Economics 14*(1), 747–775. *(Cited on pages 20, 27, 29, and 30.)*

Karahan, F. and S. Ozkan (2013). On the Persistence of Income Shocks over the Life Cycle: Evidence, Theory, and Implications. *Review of Economic Dynamics 16*(3), 452–476. *(Cited on page 24.)*

Klenow, P. J. and O. Kryvtsov (2008). State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation? *The Quarterly Journal of Economics 123*(3), 863–904. *(Cited on page 8.)*

Krueger, D., K. Mitman, and F. Perri (2016). Macroeconomics and Household Heterogeneity. In *Handbook of Macroeconomics*, Volume 2, pp. 843–921. Elsevier. *(Cited on page 24.)*

Kuznets, S. (1973). Modern Economic Growth: Findings and Reflections. *The American Economic Review 63*(3), 247–258. *(Cited on page 20.)*

La'O, J. and A. Tahbaz-Salehi (2022). Optimal Monetary Policy in Production Networks. *Econometrica 90*(3), 1295–1336. *(Cited on page 5.)*

Leahy, J. V. and A. Thapar (2022). Age Structure and the Impact of Monetary Policy. *American Economic Journal: Macroeconomics 14*(4), 136–173. *(Cited on page 4.)*

Luetticke, R. (2021). Transmission of Monetary Policy with Heterogeneity in Household Portfolios. *American Economic Journal: Macroeconomics 13*(2), 1–25. *(Cited on page 5.)*

MaCurdy, T. E. (1981). An Empirical Model of Labor Supply in a Life-Cycle Setting. *Journal of Political Economy 89*(6), 1059–1085. *(Cited on page 14.)*

Mangiante, G. (2025). Demographic Trends and the Transmission of Monetary Policy. *Available at SSRN 4149657*. *(Cited on page 4.)*

Marto, R. (2024). Structural Change and the Rise in Markups. *Federal Reserve Bank of St. Louis, Research Division*. *(Cited on page 24.)*

Matsuyama, K. (1992). Agricultural Productivity, Comparative Advantage, and Economic Growth. *Journal of Economic Theory 58*(2), 317–334. *(Cited on pages 21 and 34.)*

McKay, A. and R. Reis (2016). The Role of Automatic Stabilizers in the U.S. Business Cycle. *Econometrica 84*(1), 141–194. *(Cited on page 15.)*

McLeay, M. and S. Tenreyro (2020). Optimal Inflation and the Identification of the Phillips Curve. *NBER macroeconomics annual 34*(1), 199–255. *(Cited on page 2.)*

Moro, A. (2012). The Structural Transformation between Manufacturing and Services and the Decline in the U.S. GDP Volatility. *Review of Economic Dynamics 15*(3), 402–415. *(Cited on page 21.)*

Nakamura, E. and J. Steinsson (2008). Five Facts about Prices: A Reevaluation of Menu Cost Models. *The Quarterly Journal of Economics 123*(4), 1415–1464. *(Cited on pages 1, 6, 7, 8, 53, and 54.)*

Nakamura, E. and J. Steinsson (2013). Price Rigidity: Microeconomic Evidence and Macroeconomic Implications. *Annual Review of Economics 5*(1), 133–163. *(Cited on page 8.)*

Newey, W. K. and K. D. West (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica 55*(3), 703. *(Cited on page 10.)*

Nie, J. and A. Gautam (2020). Spending Patterns and Cost of Living for Younger versus Older Households. *The Federal Reserve Bank of Kansas City Economic Review*, 5–21. *(Cited on page 56.)*

Olivi, A., V. Sterk, and D. Xhani (2024). Optimal Monetary Policy during a Cost-of-Living Crisis. *mimeo, UCL. 1*. *(Cited on page 5.)*

Orchard, J. (2025). Non-homothetic Demand Shifts and Inflation Inequality. *Working Paper*. *(Cited on page 5.)*

Orchard, J. D., V. A. Ramey, and J. F. Wieland (2025). Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates. *The Quarterly Journal of Economics*, 2001–2052. *(Cited on page 30.)*

Pakos, M. (2004). Asset Pricing with Durable Goods and Non-Homothetic Preferences. *MPRA Paper No. 26167*. *(Cited on page 5.)*

Pancrazi, R. and M. Vukotić (2019). Inflation Sensitivity to Monetary Policy: What Has Changed Since the Early 1980s? *Oxford Bulletin of Economics and Statistics 81*(2), 412–436. *(Cited on page 4.)*

Pasten, E., R. Schoenle, and M. Weber (2020). The Propagation of Monetary Policy Shocks in a Heterogeneous Production Economy. *Journal of Monetary Economics 116*, 1–22. *(Cited on page 5.)*

Pica, S. (2023). Housing Markets and the Heterogeneous Effects of Monetary Policy across the Euro Area. *Available at SSRN 4060424*. *(Cited on page 4.)*

Roberts, J. M. (1995). New Keynesian Economics and the Phillips Curve. *Journal of Money, Credit and Banking 27*(4), 975–984. *(Cited on page 24.)*

Romer, C. D. and D. H. Romer (2023). Presidential Address: Does Monetary Policy Matter? The Narrative Approach after 35 Years. *American Economic Review 113*(6), 1395–1423. *(Cited on page 8.)*

Rotemberg, J. J. (1982). Sticky Prices in the United States. *Journal of Political Economy 90*(6), 1187–1211. *(Cited on pages 17 and 24.)*

Rubbo, E. (2023). Networks, Phillips Curves, and Monetary Policy. *Econometrica 91*(4), 1417–1455. *(Cited on page 5.)*

Slacalek, J., O. Tristani, and G. L. Violante (2020). Household Balance Sheet Channels of Monetary Policy: A Back of the Envelope Calculation for the Euro Area. *Journal of Economic Dynamics and Control 115*, 103879. *(Cited on page 5.)*

Sonnervig, M. (2025). Unequal Business Cycles. *Working Paper*. *(Cited on page 5.)*

Stock, J. H. and M. W. Watson (2020). Slack and Cyclically Sensitive Inflation. *Journal of Money, Credit and Banking 52*(S2), 393–428. *(Cited on pages 2 and 11.)*

Stone, R. (1954). Linear Expenditure Systems and Demand Analysis: an Application to the Pattern of British Demand. *The Economic Journal 64*(255), 511–527. *(Cited on page 59.)*

Storesletten, K., C. I. Telmer, and A. Yaron (2004). Cyclical Dynamics in Idiosyncratic Labor Market Risk. *Journal of Political Economy 112*(3), 695–717. *(Cited on page 24.)*

Tauchen, G. (1986). Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions. *Economics Letters 20*(2), 177–181. *(Cited on page 24.)*

Taylor, J. B. (1993). Discretion versus Policy Rules in Practice. In *Carnegie-Rochester Conference Series on Public Policy*, Volume 39, pp. 195–214. *(Cited on page 19.)*

Wachter, J. A. and M. Yogo (2010). Why do Household Portfolio Shares Rise in Wealth? *The Review of Financial Studies 23*(11), 3929–3965. *(Cited on page 5.)*

Woodford, M. (2004). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press. *(Cited on page 14.)*

Yellen, J. (2016). Macroeconomic Research After the Crisis. At "The Elusive 'Great' Recovery: Causes and Implications for Future Business Cycle Dynamics" 60th annual economic conference sponsored by the Federal Reserve Bank of Boston, Boston, Massachusetts. *(Cited on page 3.)*

Young, E. R. (2010). Solving the Incomplete Markets Model with Aggregate Uncertainty using the Krusell–Smith Algorithm and Non-Stochastic Simulations. *Journal of Economic Dynamics and Control 34*(1), 36–41. *(Cited on page 63.)*

Structural Transformation and the Transmission of Monetary Policy

Appendix

Tiago Bernardino
IIES, Stockholm University

A Data and Empirical Motivation

A.1 Data Sources and Cleaning

To compute the frequency of price adjustment, I use the summary statistics compiled by Nakamura and Steinsson (2008). They are built based on the confidential monthly product-level price data used to compute the CPI in the United States. Table A.1 shows the median frequency and magnitude of price adjustment for different categories, using only regular price changes and considering only regular price changes. I then aggregate by goods and services, as described in the main text.

Table A.1: Frequency and Magnitude of Price Changes

Category	Weight	Freq. (Reg)	Freq. (All)	Magnitude (Reg)	Magnitude (All)
Processed food	8.2	10.5	25.9	13.2	26.5
Unprocessed food	5.9	25.0	37.3	14.2	27.1
Household furnishing	5.0	6.0	19.4	8.7	20.8
Apparel	6.5	3.6	31.0	11.5	30.2
Transportation goods	8.3	31.3	31.3	6.1	6.1
Recreation goods	3.6	6.0	11.9	10.1	18.9
Other goods	5.4	15.0	15.5	7.3	10.0
Utilities	5.3	38.1	38.1	6.3	6.3
Vehicle fuel	5.1	87.6	87.6	6.4	6.4
Travel	5.5	41.7	42.8	21.6	21.9
Services (excl. travel)	38.5	6.1	6.6	7.1	7.3
All sectors	100.0	8.7	19.4	8.5	10.7

Notes: This table shows the relative weight in the consumption basket, the median frequency of regular price changes within a month, and the median magnitude of price changes within a month. It shows only regular price changes (Reg.) and for all price changes (All). Source: Table 2 of Nakamura and Steinsson (2008).

To document the differences in the demand composition across households, I use the CEX survey. This is a well-known survey used to compute the CPI expenditure weights each year. In my analysis, I consider the years between 2000 and 2019, which include 622,266 observations. I then exclude households that report negative expenditures in a given expenditure category (mostly health-related expenditures), households that report zero food consumption (both at home and away), households younger than 25 years old, and finally, households that report a negative income. This leaves me with 558,848 households, which corresponds to an average

number of 6,009 households per quarter. Table A.2 summarizes these steps.

Table A.2: Data-Cleaning Steps of the CEX sample

Step	Description	Observations	HH/quarter
0	Initial dataset	622,266	6,691
1	Exclude HH with negative expenditures	618,533	6,651
2	Exclude HH with zero food expenditures	615,289	6,616
3	Exclude HH below 25 years old	577,224	6,207
4	Exclude HH with negative income	558,848	6,009

Notes: This table shows the cleaning steps performed in the CEX data. Each observation corresponds to a household in a given quarter. The last column shows the average number of households per quarter.

A.2 Classification into Goods and Services

From the several expenditure items reported in the CEX data, I use the aggregation suggested in Hobijn and Lagakos (2005) with only small differences that allow me a better categorization into goods and services. I then follow the Bureau of Economic Analysis to classify these categories into goods and services. Table A.3 shows this classification.

A.3 Robustness to Differences in the Sectoral Price Rigidity

In this subsection of the Appendix, I show that the differences in the sectoral price rigidity are also present when using a different dataset.

Using the Bils and Klenow (2004) dataset. I use an alternative dataset, assembled by Bils and Klenow (2004), to compare the frequency of price changes between goods and services. In this paper, they consider the underlying microdata used to compute the CPI between 1995 and 1997. Similar to Nakamura and Steinsson (2008), this covers about 70% of consumer expenditures. I follow the same aggregation procedure described in the main text to compute the median price frequency and the median price duration in the goods and services sector.

In Figure A.1, I show the median frequency of price changes within a month for goods and services. The same pattern as the one described in the main text appears, although for this dataset, services adjust slightly more frequently than in the data assembled by Nakamura and Steinsson (2008). Within a month, almost 30% of the price quotes of goods change, whereas

Table A.3: Classification into Goods and Services

Category	CEX code	Classification
Household operations	HOUSOPPQ	Services
Reading	READPQ	Services
Entretainment	ENTERTPQ	Services
Food away	FDAWAYPQ	Services
Education	EDUCAPQ	Services
Other vehicles	EOTHVEHP	Services
	OTHVEHPQ	
Public transportation	TRNOTHPQ	Services
	TRNTRPPQ	
Personal care	PERSCAPQ	Services
Medical services	HLTHINPQ	Services
	MEDSRVPQ	
Other lodging	OTHLODPQ	Services
Utilities: natural gas	NTLGASPQ	Services
Utilities: electricity	ELCTRCPQ	Services
Utilities: fuel	ALLFULPQ	Services
Utilities: telephone	TELEPHPQ	Services
Utilities: water	WATRPSPQ	Services
Medical drugs and supplies	PREDRGPQ	Goods
	MEDSUPPQ	
Tobacco	TOBACCPQ	Goods
Gasoline	GASMOPQ	Goods
Food at home	FDHOMEPQ	Goods
Apparel	APPARPQ	Goods
Alcoholic beverages	ALCBEVPQ	Goods

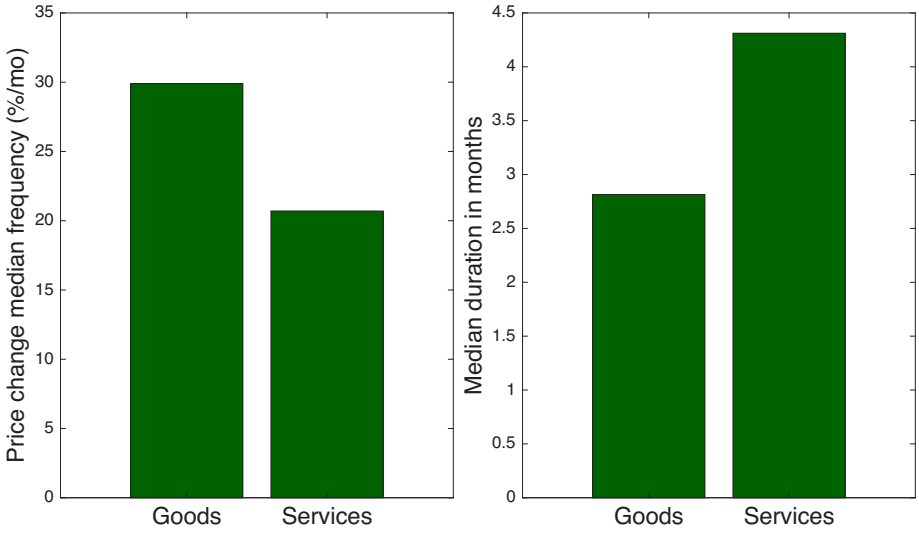
Notes: The table shows, for the different categories considered, the corresponding variable code in the Consumer Expenditure Survey and the corresponding classification by the Bureau of Economic Analysis into goods and services.

only 20.7% for the services. This implies that the median duration of a good price spell is 2.8 months for a good and 4.3 months for a service.

A.4 Robustness to Heterogeneous Demand Composition

In this subsection of the Appendix, I show that the heterogeneous demand composition pattern documented in Section 2 is robust to different sample constructions.

Figure A.1: Robustness to Fact 1: using Bils and Klenow (2004) Dataset



Notes: The left panel shows the median frequency of price changes in a month for goods and services. The right panel shows the implied median price duration of a price spell in months. Duration is equal to $-1/\log(1-f)$, where f is the frequency of price changes. Data are sourced from Bils and Klenow (2004).

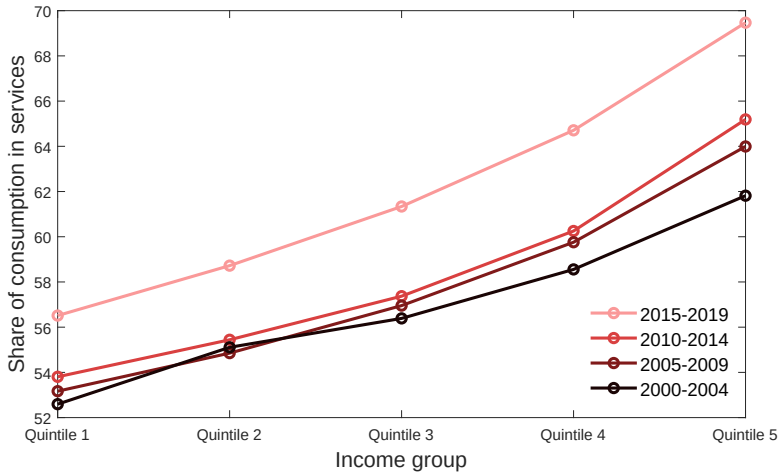
Excluding households older than 65 years old. Consumption patterns change along the life cycle. Nie and Gautam (2020) show that the implicit average inflation rate between 1984 and 2018 faced by older households is 0.4 percentage points larger than it is for younger households. All these differences come from the fact that they allocate their expenditures differently, and expenditure categories have different price changes.

I show that excluding households older than 65 years old does not qualitatively affect the heterogeneity of the demand composition documented across the income distribution. Figure A.2 shows that old household consumption patterns do not drive the result, as the difference between the first and the fifth income groups is 10 percentage points, similar to the baseline figure in the main text. There is, however, a downward shift of the average share, which indicates that older households consume a higher share of services (namely, healthcare services and nursing homes).

Including households in the bottom and top 5% of the income distribution. The CEX dataset is constructed to have a representative sample of the American population that self-reports their expenditures. It might have, therefore, data points that are outliers. Furthermore, the BLS uses top-coding rules on some variables, such as income.

To account for potential problems arising from outliers and censored data, I follow Aguiar

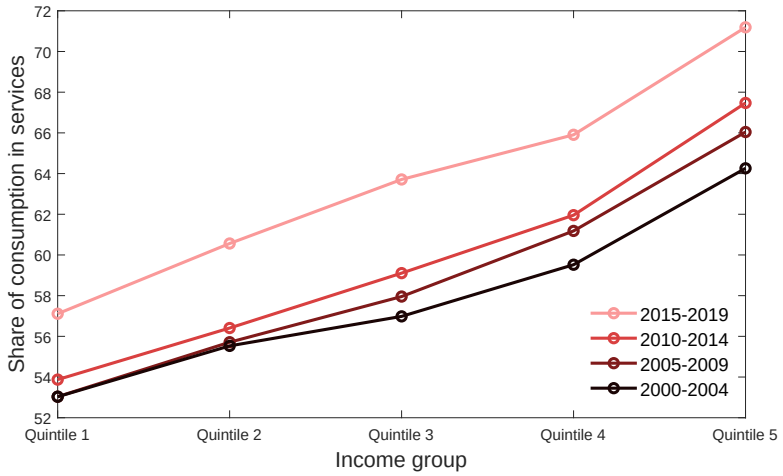
Figure A.2: Robustness to Heterogeneous Demand Composition: Excluding Old Households



Notes: The figure shows the average service share for each income group, excluding households above 65 years old. Table A.3 in the Appendix enumerates all the categories classified as services. The data source is the Consumer Expenditure Survey for the U.S. between 2000 and 2019.

and Bils (2015) and exclude each quarter's top and bottom 5% of income earners in the main text. In Figure A.3, I reproduce the differences in demand composition when including these extreme income earners. The pattern is the same, with high-income households allocating a larger share of their expenditures to services than low-income households. The major difference is in the highest income quintile, which has an even higher share of consumption in services. For the years 2015–2019, which reaches 71%, whereas when I exclude the extreme values from the analysis, the value is 69%.

Figure A.3: Robustness to Heterogeneous Demand Composition: Including the Bottom and Top 5% of the Income Distribution



Notes: The figure shows the average service share for each income group, including households at the bottom and top 5% of the income distribution. Table A.3 in the Appendix enumerates all the categories classified as services. The data source is the Consumer Expenditure Survey for the U.S. between 2000 and 2019.

B Additional Model Derivations

B.1 Household Intratemporal Problem

Each period, households choose how to allocate their total expenditure between goods and services. For that, they solve the expenditure minimization problem given by

$$\min_{c_g, c_s} E = p_g c_g + p_s c_s$$
$$\text{s.to } c^{\frac{1}{\sigma}} c_g^{\frac{\sigma-1}{\sigma}} + (\Omega c^{\epsilon})^{\frac{1}{\sigma}} c_s^{\frac{\sigma-1}{\sigma}} = 1.$$

The first-order conditions are

$$p_g = \lambda c^{\frac{1}{\sigma}} \frac{\sigma-1}{\sigma} c_g^{\frac{-1}{\sigma}},$$

and

$$p_s = \lambda (\Omega c^{\epsilon})^{\frac{1}{\sigma}} \frac{\sigma-1}{\sigma} c_s^{\frac{-1}{\sigma}},$$

where λ is the Lagrange multiplier associated with the restriction of the problem.

Solving for c_g and c_s yields the Hicksian demands:

$$c_g = \left(\frac{p_g}{E}\right)^{-\sigma} c \text{ and } c_s = \Omega \left(\frac{p_s}{E}\right)^{-\sigma} c^{\epsilon},$$

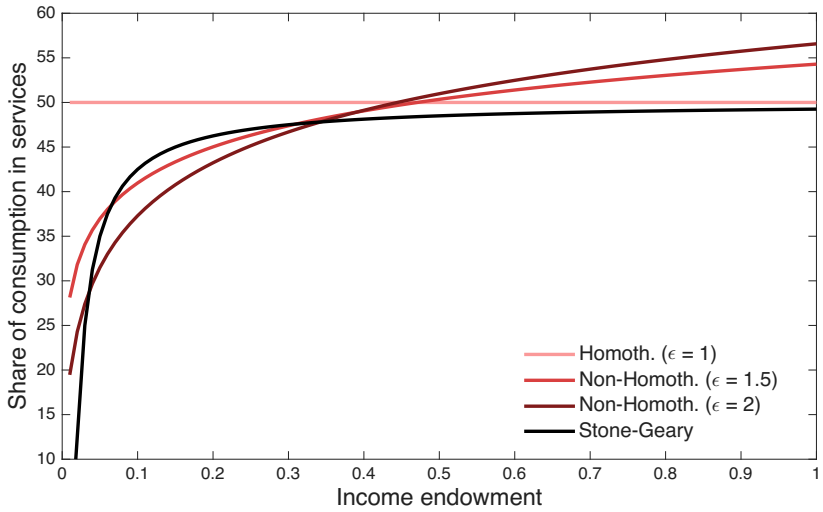
corresponding to Equation (7) in the main text. Furthermore, the expenditure function is

$$E(c_g, c_s) \equiv \sum_{m \in \{g,s\}} p_m c_m = \left[c p_g^{1-\sigma} + \Omega c^{\epsilon} p_s^{1-\sigma} \right]^{\frac{1}{1-\sigma}}.$$

Note that when $\epsilon = 1$, the consumption aggregator is the standard homothetic CES. For the relevant case when goods and services are gross complements with $\sigma \in (0, 1)$, when $\epsilon > 1$, services are a "luxury consumption good" and when $\epsilon < 1$, services are instead a "necessity consumption good". As Section 2 points out, the slope of the Engel curve of services is positive, meaning that high-income households consume relatively more services than low-income households, i.e., services are a "luxury consumption good".

Figure B.4 illustrates the relevance of modeling preferences with a non-homothetic relation for the household micro-behavior consumption heterogeneity. The figure shows the service share in consumption for different income endowments when the CES aggregator is homothetic ($\epsilon = 1$), when the CES aggregator is non-homothetic ($\epsilon = 1.5$ and $\epsilon = 2$), and when there is a subsistence point (Stone-Geary), the most used form of non-homothetic preferences in the literature (Geary, 1950, Stone, 1954).[29]

Figure B.4: Illustration of the Non-Homothetic CES Aggregator



Notes: The figure plots the service share in consumption for different levels of income. In red, it shows for the homothetic CES case, orange and pink for the non-homothetic CES case with different values of ϵ, and in green for the Stone-Geary case. Both CES cases use $\sigma = 0.2$ and $\Omega = 1$. The Stone-Geary case uses an additive log specification with a subsistence point on goods equal to 0.015.

[29] I set $\sigma = 0.2$ and $\Omega = 1$ for the homothetic and non-homothetic CES cases. For the Stone-Geary preferences, I use an additive log specification with a subsistence point on goods equal to 0.015.

For the homothetic case, independently of the income, households always spend half of their budget on services. However, for the non-homothetic cases, the service share increases with income. For the nhCES, a higher ϵ implies a higher slope of the Engel curve. For the Stone-Geary case, the slope occurs mainly in the first income levels, and asymptotically converges to the homothetic case of a service share equal to 50%.

This figure justifies the modeling choice of the non-homothetic CES preference relation. With the homothetic relation, I am unable to reproduce the household heterogeneity documented in the data, nor the shift in economic activity toward services. But even among the different non-homothetic preference relations possible, the non-homothetic CES gives a better fit of the service share heterogeneity.

B.2 Final Producer Problem

The final producer of sector m aggregates a continuum of intermediate inputs indexed by $j \in [0,1]$, according to Equation (10). Given a level of aggregate demand Q_m, the final producer solves:

$$\min_{\{q_m(j)\}} \int_0^1 p_m(j) q_m(j) dj$$
$$\text{s.to } Q_m = \left(\int_0^1 q_m(j)^{\frac{\theta_m - 1}{\theta_m}} dj \right)^{\frac{\theta_m}{\theta_m - 1}}.$$

The first-order condition for input j is

$$p_m(j) = \lambda \left(\frac{Q_m}{q_m(j)} \right)^{\frac{1}{\theta_m}},$$

where λ is the Lagrange multiplier. For two different intermediate inputs, I have that

$$\left(\frac{p_m(j)}{p_m(\ell)} \right)^{\theta_m} = \frac{q_m(\ell)}{q_m(j)} \Leftrightarrow \int_0^1 p_m(\ell) q_m(\ell) d\ell = \int_0^1 \left(\frac{p_m(j)}{p_m(\ell)} \right)^{\theta_m} q_m(j) p_m(\ell) dl.$$

Note that the LHS of the above equation corresponds to the expenditure done in sector m. Hence,

$$P_m Q_m = p_m(j)^{\theta_m} q_m(j) \int_0^1 p_m(\ell)^{1-\theta_m} d\ell.$$

The integral in the above equation is the definition of the price index of sector m as defined

in Equation (12). Therefore, I obtain the demand for intermediate input j as in Equation (11):

$$q_m(j) = \left(\frac{p_m(j)}{P_m}\right)^{-\theta_m} Q_m.$$

B.3 Intermediate Input Producer Problem

In this subsection, I solve the intermediate input producer dynamic pricing problem. I omit the subscript for the sector, m, for convenience. The intermediate firm j operates under monopolistic competition and sets the input j price, $p(j)$. These firms face a demand given by Equation (11). Moreover, they have a cost for adjusting prices given by Equation (14).

The dynamic problem is given in (15), but I reproduce it here for convenience:

$$\begin{aligned} V_t(p_{t-1}(j)) = \max_{\{p_t(j)\}} \; & d_t = p_t(j)q_t(j) - w_t n_t(j) - \Phi\left(p_t(j), p_{t-1}(j)\right) + \frac{1}{1+i_t} V_{t+1}(p_t(j)) \\ & \text{s.to Equations (11) and (13),} \end{aligned}$$

or equivalently,

$$\begin{aligned} V_t(p_{t-1}(j)) = \max_{\{p_t(j)\}} \; & p_t(j)\left(\frac{p_t(j)}{P_t}\right)^{-\theta} Q_t - w_t\left(\frac{p_t(j)}{P_t}\right)^{-\theta} Q_t - \Phi\left(p_t(j), p_{t-1}(j)\right) \\ & + \frac{1}{1+i_t} V_{t+1}(p_t(j)) \end{aligned}$$

The first-order condition with respect to $p_t(j)$ is

$$\begin{aligned} \frac{\partial V_t}{\partial p_t(j)} = & \left(\frac{p_t(j)}{P_t}\right)^{-\theta} Q_t - \theta p_t(j)\left(\frac{p_t(j)}{P_t}\right)^{-\theta-1}\frac{Q_t}{P_t} \\ & + \theta\frac{w_t}{Z}\left(\frac{p_t(j)}{P_t}\right)^{-\theta-1}\frac{Q_t}{P_t} - \frac{\theta}{\kappa}\log\left(\frac{p_t(j)}{p_{t-1}(j)}\right) Q_t P_t \frac{p_{t-1}(j)}{p_t(j)P_{t-1}} + \frac{1}{1+i_t}\frac{\partial V_{t+1}}{\partial p_t(j)} = 0, \end{aligned}$$

and the Envelope Condition of this problem is

$$\frac{\partial V_t}{\partial p_{t-1}(j)} = \frac{\theta}{k}\log\left(\frac{p_t(j)}{p_{t-1}(j)}\right)\frac{Q_t P_t}{p_{t-1}(j)}.$$

Combining the FOC and the Envelope Condition, I obtain the following expression:

$$\left(\frac{p_t(j)}{P_t}\right)^{-\theta} Q_t - p_t(j)\theta\left(\frac{p_t(j)}{P_t}\right)^{-\theta-1}\frac{Q_t}{P_t} + \frac{w_t}{Z}\theta\left(\frac{p_t(j)}{P_t}\right)^{-\theta-1}\frac{Q_t}{P_t} - \frac{\theta}{\kappa}\log\left(\frac{p_t(j)}{p_{t-1}(j)}\right)\frac{Q_t P_t}{p_{t-1}(j)} + \frac{1}{1+i_t}\cdot\frac{\theta}{\kappa}\log\left(\frac{p_{t+1}(j)}{P_t(j)}\right)\frac{Q_{t+1}P_{t+1}}{p_t(j)} = 0.$$

In a symmetric equilibrium, all firms choose the same prices, so $p_t(j) = P_t$. Hence, the above expression is simplified to

$$1 - \theta + \theta\frac{w_t}{Z}\frac{1}{P_t} - \frac{\theta}{\kappa}\log\left(\frac{P_t}{P_{t-1}}\right) + \frac{1}{1+i_t}\frac{\theta}{\kappa}\log\left(\frac{P_{t+1}}{P_t}\right)\frac{Q_{t+1}}{Q_t}\frac{P_{t+1}}{P_t} = 0$$

Defining $\frac{P_t}{P_{t-1}} \equiv 1 + \pi_t$, yields the optimal price setting rule, Equation (16), which is the New-Keynesian Phillips curve of sector m.

C Computational Algorithms and Calibration Details

C.1 Steady-state Equilibrium

To solve for the steady-state, I make use of my production setting with monopolistic competition and linear production. In the steady-state, prices are constant, so the optimal sectoral prices are simply given by the product of the markup and the marginal cost as in Equation (17).

Moreover, the use of non-homothetic CES preferences as in Equation (4) implies an additional step. Consumption and Expenditure are not the same as in a standard heterogeneous-agent model, i.e., the object that enters the utility function, C, and the objects that appear in the household budget constraint, E. However, I can use Equation (8) to map expenditure to consumption. In practice, for given parameters and prices, I solve for 500 values of expenditure the implied consumption index values, and then interpolate using a piecewise cubic interpolation.

The goal is to find the steady-state equilibrium. That the household's equilibrium objects — value function, policy functions of consumption, labor, sectoral consumptions and savings, and distribution — prices and dividends — w, p_g, p_s, p_b and D — and taxes, T. For that, I use the following algorithm:

1. Set parameters, normalize $w = 1$, use Equation (17) to compute p_g and p_s
2. Guess value for p_b
   (a) Guess T and D
      i. Solve the household problem with a value function iteration (VFI) algorithm, interpolating with a piecewise cubic method → get value function and policy functions (expenditure, consumption, savings, and labor supply)
      ii. Solve for the stationary distribution using the lotteries algorithm (Young, 2010)
      iii. Solve for the sectoral consumption
      iv. Aggregate and obtain aggregate consumption, expenditure, sectoral consumption, labor supply, and asset demand
      v. Solve the firms' problems → get sectoral profits, labor demand, prices, and quantities

(b) Compute the difference between total profits and guessed dividends, and check if the government budget constraint holds. Update D and T, using a convex combination update

3. Check for asset market clearing. Update p_b using a bisection update

The algorithm provides an approximate solution. I choose a tolerance criterion for each of the steps that is high enough to give precise results. The intertemporal household problem is solved with the builtin MATLAB `fminsearch` and `fminbnd` with a precision of 1e-11; the VFI algorithm is solved with a precision of 1e-10; the dividend and the lump-sum tax with a precision of 1e-8; and the price of bonds with a precision of 1e-5.

C.2 Transition Path Equilibrium

The goal is to find the sequence of equilibrium prices $\{p_t^b\}_{t=0}^T$, $\{p_t^g\}_{t=0}^T$, $\{p_t^s\}_{t=0}^T$, and $\{w_t\}_{t=0}^T$; profits $\{\pi_t^g\}_{t=0}^T$ and $\{\pi_t^s\}_{t=0}^T$; taxes $\{T_t\}_{t=0}^T$; interest rate, $\{i_t\}_{t=0}^T$; value functions, policy functions and law of motion, Ψ that solve for an exogenous shock in the nominal interest rate, ε^M.

In general solving for this is computationally challenging as the distribution of wealth evolves with the shock as well. I adapt the proposed algorithm in Boppart et al. (2018) and guess the sequence of prices along the transition.[30] Note that prices are not constant anymore and that wages cannot be normalized, which creates two additional loops. I use the following algorithm:

1. Set the nominal interest rate shock, ε^M
2. Guess the price of bonds sequence $\{p_t^b\}_{t=0}^T$
   (a) Guess the sequences of sectoral prices $\{p_t^g\}_{t=0}^T$ and $\{p_t^s\}_{t=0}^T$, and compute the sectoral inflation rates and the aggregate inflation rate
   (b) Use the Taylor rule in Equation (20) to compute the nominal interest rate
      i. Guess the wage sequence $\{w_t\}_{t=0}^T$
         A. Guess the dividend and tax sequence, $\{D_t\}_{t=0}^T$ and $\{T_t\}_{t=0}^T$
         B. Solve the household problem backwards (in $t = T+1$, the economy is in the SS)

[30] I truncate the convergence at $T = 150$, when the economy is very close to the steady-state again.

C. Solve for the distribution forwards (in $t = 0$, the economy is in the SS)

D. Solve for the sectoral consumption

E. Compute aggregate consumption, expenditure, sectoral consumptions, labor supply, and asset demand

F. Solve for the firm problem and get sectoral profits, labor demand and prices

G. Compute the difference between total profits and guessed dividends, and check if the government budget constraint holds. Update $\{D_t\}_{t=0}^{T}$ and $\{T_t\}_{t=0}^{T}$ using a convex combination update for each period

ii. Check for labor market clearing and update $\{w_t\}_{t=0}^{T}$ using a shooting method for each period

(c) Check if guessed sectoral prices are equal to firm optimal prices and update $\{p_t^g\}_{t=0}^{T}$, $\{p_t^s\}_{t=0}^{T}$ using a convex combination update for each period

3. Check for asset market clearing and update $\{p_t^b\}_{t=0}^{T}$ using a shooting algorithm for each period

The algorithm provides an approximate solution for the impulse response functions that can be used to analyse the dynamic effects of the monetary policy shock. I choose a tolerance criterion for each of the steps that is high enough to give precise results. The intertemporal household problem is solved with the builtin MATLAB `fminsearch` and `fminbnd` with a precision of 1e-11; the dividend loop and the lump-sum tax with a precision of 1e-7; the labor market clearing with a precision of 1e-6; the sectoral prices with a precision of 1e-5; and the asset market clearing with a precision of 1e-4.

C.3 Calibration Robustness

Table C.4 reports the demand estimation (Equation (24)) results when the sample is split into groups of 5 years. Even though the estimates are different, they are not different, in the statistic sense, from the one estimated with the full sample. This indicates that these parameters are quite stable over time. Note that, due to the reduction of the sample size, I cannot estimate the third specification with both region and year $\times$ quarter fixed effects.

Table C.4: Demand Estimation

	2000–2005		2006–2010	
	(1)	(2)	(1)	(2)
σ	0.436	0.082	0.224	0.215
	(0.104)	(0.112)	(0.077)	(0.062)
ϵ	2.077	1.651	1.576	1.626
	(0.365)	(0.125)	(0.094)	(0.086)
Region FE	N	Y	N	Y

Notes: The table presents results of the GMM estimation of Equation (24) with 27,440 observations for the 2000-2005 sample and 33,492 for the 2006-2011 sample. The number in parentheses corresponds to the robust standard error clustered at the household level. All regressions include demographic household controls.

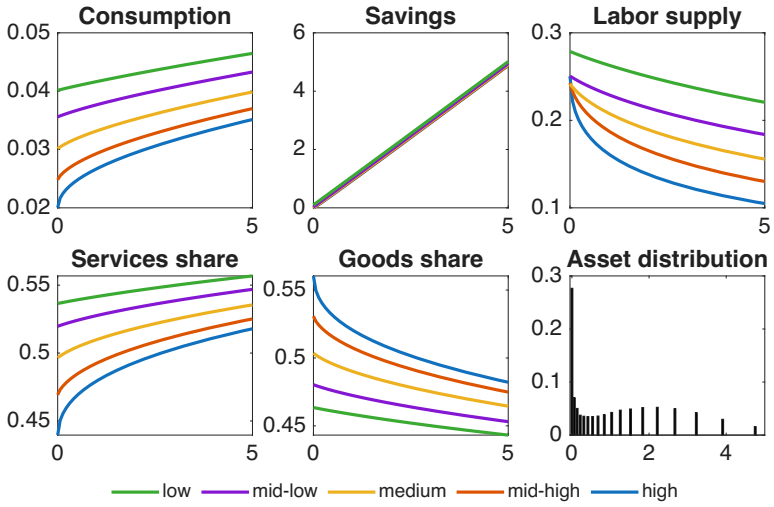
D Additional Results

D.1 Additional Figures

D.1.1 Additional Steady-State Figures

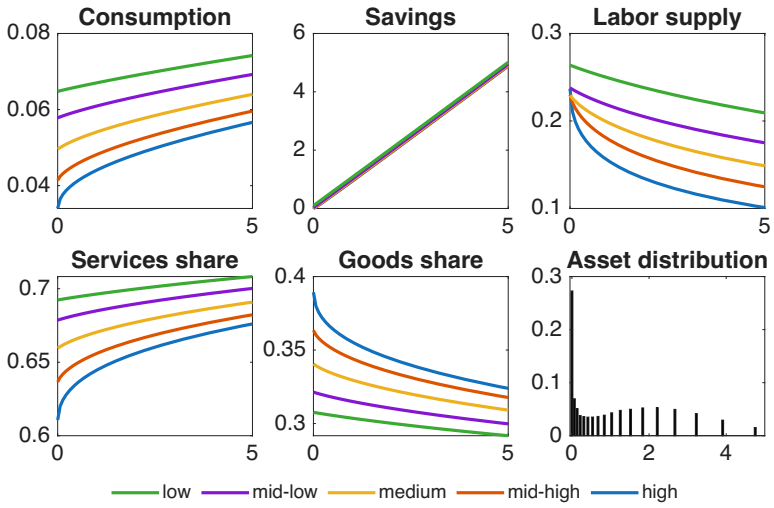
This subsection contains additional figures about the steady-state that complement the analysis in the main text. I plot the policy functions of consumption, savings, and labor supply, as well as the services and goods share and the household distribution for the different productivity states, and along the asset distribution. Figure D.5 plots them for the 1970 steady-state, figure D.6 for 2019 steady-state, and figure D.7 for the 1970 steady-state, but when I recalibrate for homothetic preferences.

Figure D.5: Steady-State Policy Functions for the 1970-Economy



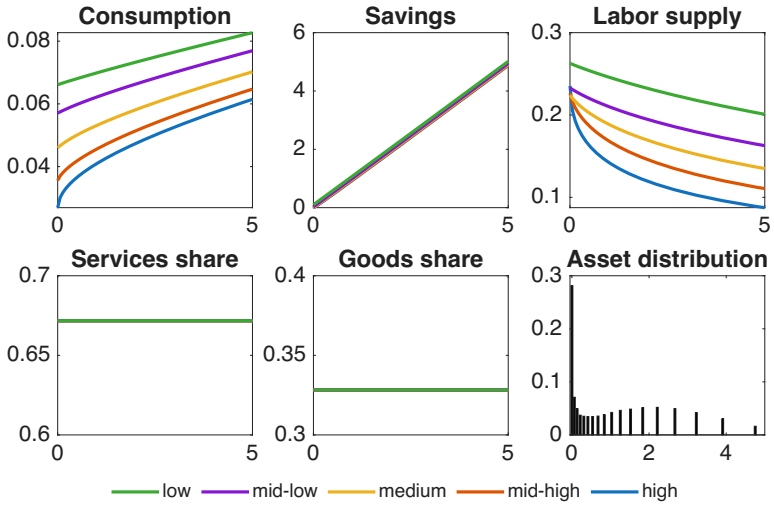
Notes: The figure shows the steady-state policy functions of consumption, savings, and labor along the state-space—asset grid (x-axis) and labor productivity (different lines). It also shows the expenditure share allocated to services and goods in the state-space and the distribution of assets.

Figure D.6: Steady-State Policy Functions for the 2019-Economy



Notes: The figure shows the steady-state policy functions of consumption, savings, and labor along the state-space—asset grid (x-axis) and labor productivity (different lines). It also shows the expenditure share allocated to services and goods in the state-space and the distribution of assets.

Figure D.7: Steady-State Policy Functions when Preferences are Homothetic (2019)



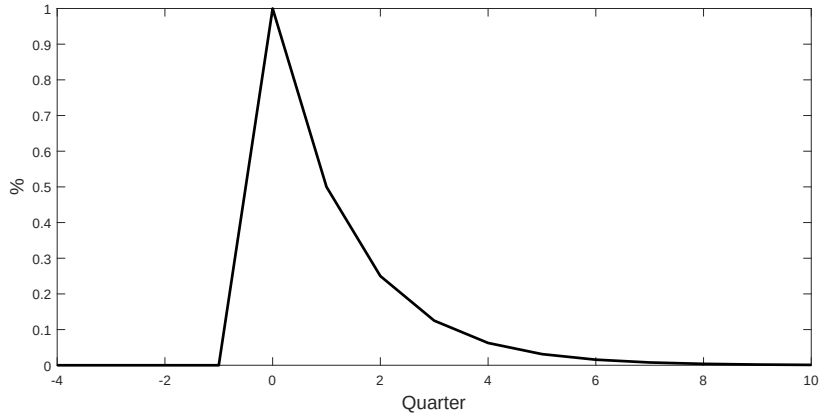
Notes: The figure shows the steady-state policy functions of consumption, savings, and labor along the state-space—asset grid (x-axis) and labor productivity (different lines). It also shows the expenditure share allocated to services and goods in the state-space and the distribution of assets.

D.1.2 Additional Monetary Policy Figures

This subsection contains additional figures regarding the responses to the monetary policy shock that complement the analysis in the main text of Section 5. Figure D.8 plots the monetary policy shock, the same one that is used for both the 1970 and the 2019 economies. Figure D.9

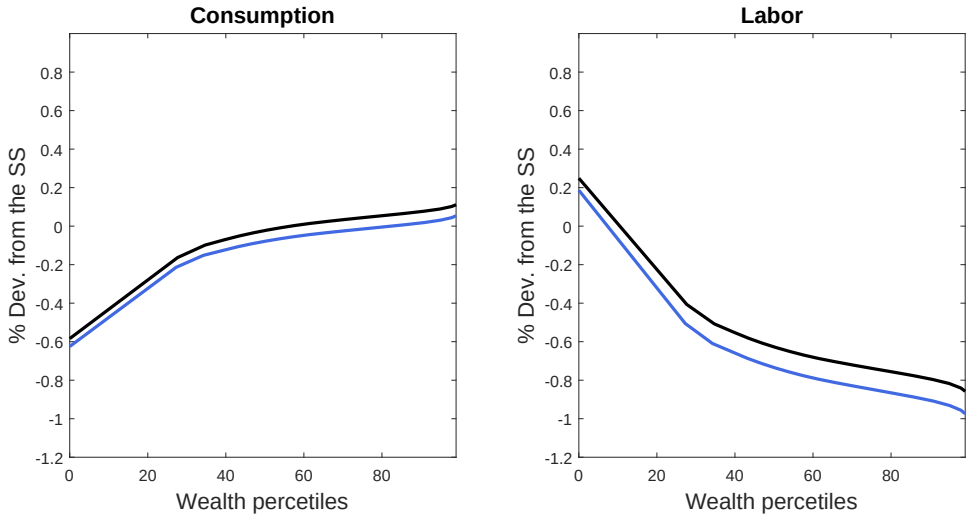
plots the response on impact of aggregate consumption and labor along the wealth distribution for the 1970 and the 2019 economies. Figure D.10 plots the heterogeneous sectoral responses on impact by wealth position for the baseline and the homothetic counterfactual.

Figure D.8: Monetary Policy Shock



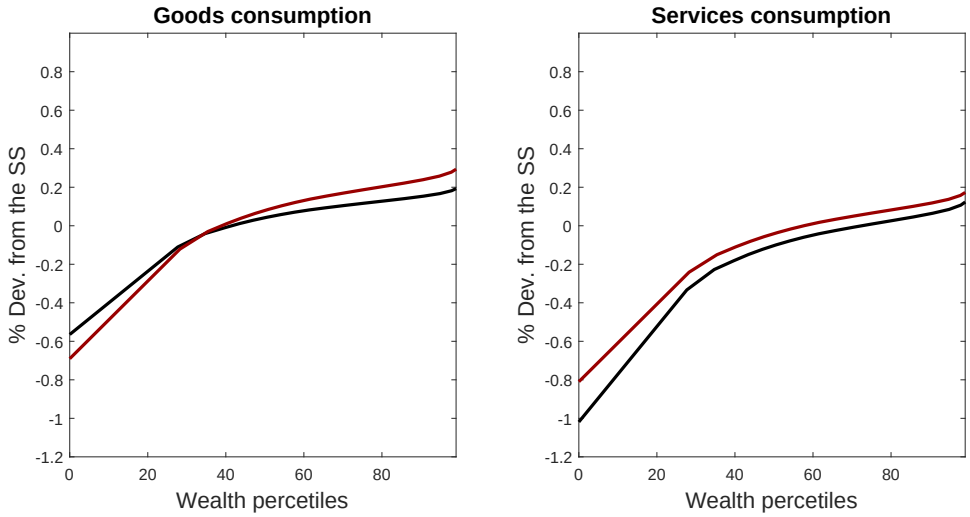
Notes: The figure shows the monetary policy shock path as described in Section 5. The values are in percentages.

Figure D.9: Consumption and Labor Responses to a Contractionary Monetary Policy Shock by Wealth Position



Notes: The figure shows the consumption index and labor responses after a monetary policy contraction of 100 basis points by wealth position. The black lines depict the response for the 1970-economy and the blue line for the 2019-economy. All values are expressed as percentage deviations from the steady-state values.

Figure D.10: Sectoral Consumption Responses by Wealth Position: Homothetic Counterfactual

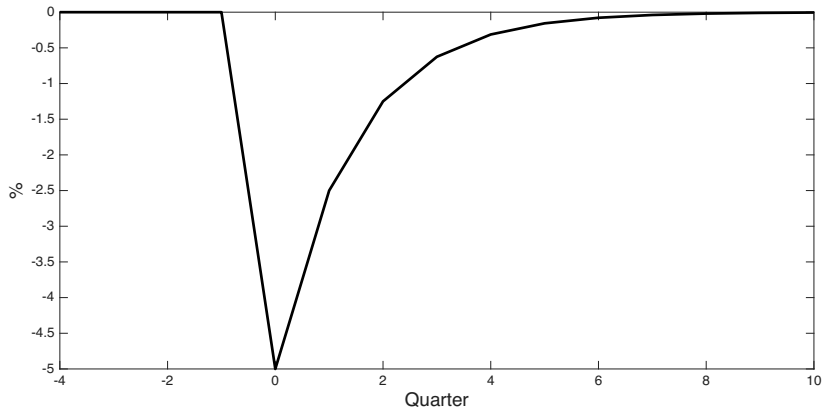


Notes: The figure shows the consumption of goods and the consumption of services responses after a monetary policy contraction of 100 basis points by wealth position under the baseline (black line) and homothetic counterfactual ($\epsilon = 1$) built as described in Section 5.2 (red line). The responses correspond to the high-service share economy. All values are expressed as percentage deviations from the steady-state values.

D.1.3 Additional Negative TFP Shock Figures

This subsection contains additional figures about the responses to the negative TFP that complement the analysis in the main text of Section 6.

Figure D.11: Aggregate TFP Shock



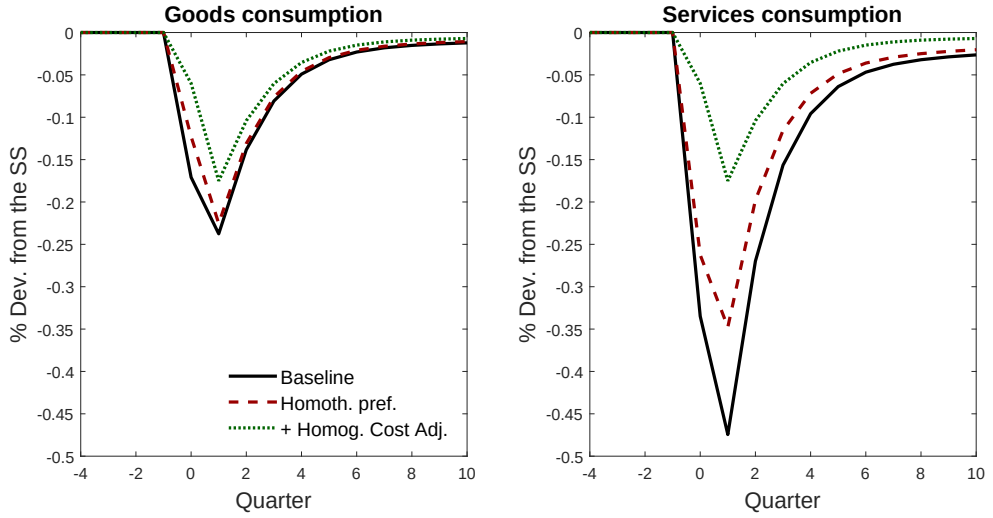
Notes: The figure shows the aggregate TFP shock path as described in Section 6. The values are in percentages.

D.2 Sectoral Response Heterogeneity Decomposition

One of the model's predictions is that a contractionary monetary policy shock alters the composition of demand. Although consumption falls in both sectors, the decline is smaller for goods than for services. As a result, the share of services in total demand decreases following the shock. This shift in demand composition is driven by sectoral differences in price rigidity and markups—captured by the slope of the NKPC in Equation (16)—as well as by differences in income elasticity, which arise from the non-homothetic nature of preferences.

In this subsection, I construct counterfactual impulse response functions to the same 100 basis point contractionary monetary policy shock. The goal is to disentangle the role of the differences in the slope of the NKPC and the role of non-homothetic preferences. The counterfactuals are constructed such that the share of services in the economy equals 68%, the value observed in 2019.

Figure D.12: Decomposition of Responses



Notes: The figure shows the homothetic and the homogeneous NKPC counterfactuals of the impulse response function of sectoral consumptions to a contractionary 100 basis point monetary contraction.

Figure D.12 shows these counterfactuals. The black full lines show the response as in Figure 7 in 2019. The first counterfactual is the case when preferences are homothetic. For that, I set ϵ equal to 1, which makes the sectoral consumption aggregator in Equation 4 become the standard CES aggregator and, in turn, the share of services across households constant. The dashed blue line plots this first counterfactual. The second counterfactual eliminates the differences in the slope of the NKPC in addition to the homothetic preferences. This means that

both sectors face the same costs of price adjustments and have the same markups. The dotted purple line shows the sectoral consumption responses in this case.

When both sectors have no differences between them, the sectoral consumption responses are the same. In response to a contractionary 100 basis point monetary contraction, services and goods consumption fall by 0.175%. When both sectors only differ on the supply side and have different NKPC slopes, the response of services consumption becomes larger than the response of goods consumption to a monetary policy shock. Services consumption response is almost two times larger and goods consumption 29% larger with respect to the homogeneous sector case. This means that services consumption falls by 0.348% and goods consumption by 0.225%. Finally, when I reintroduce the non-homothetic structure of preferences, we restore the baseline results and the differences in sectoral responses. Services respond three times more compared to the baseline case. Overall, differences in the slope of the NKPC and the non-homothetic preference structure have the same relevance in explaining the monetary response differences between the two sectors.