Year : 2023

Four Essays in Monetary Economics and Heterogeneity

Mangiante Giacomo

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT D'ÉCONOMIE

Four Essays in Monetary Economics and Heterogeneity

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Docteur en Économie

par

Giacomo MANGIANTE

Directeur de thèse
Prof. Jean-Paul Renne

Co-directeur de thèse
Prof. Florin Bilbiie

Jury

Prof. Rafael Lalive, président
Prof. Aurélien Eyquem, expert interne
Prof. Yuriy Gorodnichenko, expert externe

LAUSANNE
2023

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT D'ÉCONOMIE

Four Essays in Monetary Economics and Heterogeneity

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Docteur en Économie

par

Giacomo MANGIANTE

Directeur de thèse
Prof. Jean-Paul Renne

Co-directeur de thèse
Prof. Florin Bilbiie

Jury

Prof. Rafael Lalive, président
Prof. Aurélien Eyquem, expert interne
Prof. Yuriy Gorodnichenko, expert externe

LAUSANNE
2023

UNIL | Université de Lausanne
HEC Lausanne
Le Décanat
Bâtiment Internef
CH-1015 Lausanne

I M P R I M A T U R

Sans se prononcer sur les opinions de l'auteur, la Faculté des Hautes Etudes Commerciales de l'Université de Lausanne autorise l'impression de la thèse de Monsieur Giacomo MANGIANTE, titulaire d'un bachelor en Administration des affaires de l'Université de Gênes, et d'un master en Sciences économiques de l'Université Bocconi, en vue de l'obtention du grade de docteur en Économie.

La thèse est intitulée :

FOUR ESSAYS IN MONETARY ECONOMICS AND HETEROGENEITY

Lausanne, le 24 mai 2023

La Doyenne

Marianne SCHMID MAST

Members of the doctoral committee

Professor Florin Bilbiie

Professor of Economics, University of Cambridge

Thesis cosupervisor

Professor Jean-Paul Renne

Professor of Economics, HEC, University of Lausanne

Thesis cosupervisor

Professor Aurélien Eyquem

Professor of Economics, HEC, University of Lausanne

Internal member of the doctoral committee

Professor Yuriy Gorodnichenko

Professor of Economics, UC Berkeley

External member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Giacomo MANGIANTE

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ________________________ Date: 21 April 2023

Prof. Jean-Paul RENNE
Thesis supervisor

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Giacomo MANGIANTE

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ______________________________ Date: _25.04.2023_________

Prof. Florin BILBIIE
Thesis co-supervisor

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Giacomo MANGIANTE

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ______________________________ Date: 22.04.2023

Prof. Aurélien EYQUEM
Internal member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Giacomo MANGIANTE

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: _______________ Date: April 21, 2023

Prof. Yuriy GORODNICHENKO
External member of the doctoral committee

Four Essays in Monetary Economics and Heterogeneity

By
Giacomo Mangiante

ACKNOWLEDGEMENTS

I would like to thank my supervisors Florin Bilbiie and Jean-Paul Renne for their invaluable guidance during the last five years. You have always been there to encourage me both from a professional and personal perspective. I am particularly indebted to Yuriy Gorodnichenko for hosting me during my vising period and for his constant support afterward. I am also thankful to my coauthors Jannik Hensel, Christoph Lauper, Riccardo M. Masolo, Pascal Meichtry, Luca Moretti, and Federico di Pace. I have learned a lot from each of you and I hope to keep doing so in the future. I am extremely grateful to my partner Sara. Your patience, enthusiasm, and lightheartedness have sustained me even in the most difficult moments of my journey. A huge thank goes to my family. I always felt your support and you never lost a chance to show me how proud you have been of me. Finally, I want to express my gratitude to all my colleagues and friends who shared this experience with me. Thank you Aleksandra, Andrea (x2), Brendan, Claudia, Elio, Enrico, Fabrizio, Kevin, Krishna, Ilaria, Miriam, Paola, Pauline, Resuf, and Tiziano.

Contents

Introduction

The overarching theme of my thesis is the study of the relationship between monetary policy and heterogeneity both from an empirical and theoretical perspective. Economic agents, e.g., households, firms, banks, etc., are heterogeneously exposed to the decisions of the monetary authorities which might have important distributional effects. At the same time, the ability of the central banks to achieve their mandates is significantly influenced by the heterogeneity in the economy.

In the first chapter of my thesis, I study how demographic trends influence the effectiveness of monetary policy. Almost every country in the world is expected to experience significant demographic transitions over the next decades. Lower mortality rates and longer life expectancies have already increased the share of retired people and reduced the size of the working population and monetary authorities are unlikely to be immune from the effects of these trends. However, given the slow-moving pace of the trends, central banks have so far neglected the impact that population aging might have on the pass-through of monetary policy.

I propose and quantify a novel channel to explain how demographic trends might alter the transmission of monetary policy shocks: consumption heterogeneity across age groups. I combine household-level data from the Consumer Expenditure Survey (CEX) with measures on the sectoral frequency of price adjustments, i.e., how often prices are adjusted, and I document a negative relationship between the price stickiness of the consumption bundle and age. The main driver is the higher share of services consumed by older individuals since services tend to adjust their prices only every 13 months compared to goods which adjust their prices every 3 months. I then evaluate whether this micro-level relationship translates into actual differences at the macro level in the responsiveness to shocks. Indeed, if prices are more rigid, output should respond more to monetary shocks. To test this hypothesis, I exploit the cross-sectional variation in demographic structures across the U.S. and I show that the economic activity in states with a higher old-age dependency ratio reacts more to monetary shocks in line with the prediction. Finally, I rationalize these findings using a two-sector OLG New Keynesian model. I show that population aging has significantly increased the responsiveness of output to monetary shocks and will increase it even further in the future, that consumption heterogeneity plays an important role in explaining the change in monetary policy effectiveness induced by demographic trends, and that younger households are the most exposed to these trends.

Consumption heterogeneity plays a crucial role in also evaluating the different inflation rates to which households are exposed. Households tend to purchase different consumption baskets and the price of each good has a different sensitivity to changes in the interest rate. In the second chapter of the thesis, coauthored with Christoph Lauper, we study how the distribution of individual inflation rates is affected by monetary policy shocks. We compute a measure of inflation at the household level using the expenditure data from the CEX and we identify different moments of the inflation rate distribution. In response to a contractionary monetary policy shock, the median inflation rate as well as the cross-sectional standard deviation is significantly reduced. The decrease in inflation dispersion is almost entirely driven by expenditures on *Energy*, *Water*, and *Gasoline*. Moreover, *inflation inequality*, defined as the cross-sectional standard deviation of the median inflation rates across expenditure, salary, and income deciles, decreases after a contractionary monetary shock. The reason is that households at the bottom of the distribution are on average exposed to a higher inflation rate which tends at the same time to decrease more following a monetary shock. Finally, we find that the increase in expenditure inequality in response to monetary shocks is significantly more muted once inflation heterogeneity is taken into account.

Households are not the only important economic agents for central banks. Firms are the price setters in the economy and therefore they ultimately determine inflation. In the third chapter, jointly with Federico di Pace and Riccardo Masolo, we evaluate whether firms' expectations react to monetary policy announcements. We compare the responses to the Decision Maker Panel survey filed by firms immediately before with those that filed after a Monetary Policy Committee meeting of the Bank of England (BoE). We find that firms' expectations and uncertainty about their own business do not respond the same way financial markets do. Announcements that surprised the markets have basically no effect on firms' expectations. However, announced changes in the monetary policy rate induce firms to revise their price expectations, with rate hikes resulting in a reduction in price expectations and uncertainty surrounding them, even though the markets were perfectly able to predict them.

Central banks are becoming more and more vocal in the fight against climate change. Some of them, like the ECB and BoE, have even adopted an active role in fostering the transition toward a greener economy. However, the empirical evidence on whether climate policies might be combined with the monetary authorities' objectives is still limited. In the fourth chapter, coauthored with Jannik Hensel and Luca Moretti, we study the impact of carbon pricing on firms' inflation expectations and we discuss the potential implications for what constitutes the core of most central banks' mandate: price stability. As in Känzig (2022), we identify exogenous variations in carbon price from changes in carbon futures price around regulatory events. The shock series is combined with French firm-level survey data which reports information on firms' inflation expectations, own price expected growth, and realized price growth. We document that a change in the price of carbon increases firms' inflation expectations. Moreover, firms' own expected and realized price growth respond similarly to inflation expectations. The effect on price expectations is more persistent than on actual price growth leading to positive forecast errors in the medium-/long-run. Finally, we show that a sizable share of the increase in inflation expectations is due to indirect effects. Firms rely on their own business conditions to form expectations about the aggregate price dynamics. Therefore, the expected positive growth in their own prices significantly contributes to the observed increase in inflation expectations.

Demographic Trends and the Transmission of Monetary Policy

Giacomo Mangiante*

Link to most recent version

March 2023

Abstract

This paper studies the impact of demographic trends on the effectiveness of monetary policy. I propose and quantify a novel channel to explain how population aging affects the transmission of monetary policy: older individuals devote a larger share of their consumption bundle to product categories with higher levels of price rigidity – categories that adjust their prices less often – so the aggregate frequency of price adjustment decreases as the population ages. Using micro data on consumer expenditure, I document that the main driver of the negative relationship between age and the frequency of price adjustment is the higher share of services consumed by old households. At the macro level, if prices are more rigid output should respond more to monetary shocks. To test this hypothesis, I exploit the cross-sectional variation in demographic structures among U.S. states. I show that population aging is related to a shift towards the service sector and that the economic activity in more service-intensive states reacts more to monetary shocks. I rationalize these findings using a two-sector OLG New Keynesian model. Combining the model with population projections for the U.S., I find that changes in the age distribution between 1980 and 2010 increased the contemporaneous response of output to monetary shocks by 6% and will increase it by 10% by 2050. Moreover, demographic trends explain around 10% of the observed decrease in the slope of the Phillips curve.

Keywords: Monetary policy, age structure, consumption heterogeneity, Phillips curve
JEL classification: E31, E52, J11

*University of Lausanne, Lausanne, Switzerland, email: giacomo.mangiante@unil.ch
I am extremely grateful to my advisors Florin Bilbiie and Jean-Paul Renne for their invaluable guidance and support. I would also like to thank Michele Andreolli, Adrien Auclert, Martin Bielecki, Luigi Bocola, Martin Brown, Aurélien Eyquem, Jordi Galí, Claudia Gentile, Isabel Gödl-Hanisch, Yuriy Gorodnichenko, Andrei Levchenko, Christian Keuschnigg, Aleksandra Malova, Riccardo Masolo, Emi Nakamura, Andrea Papetti, Ricardo Reis, Rana Sajedi, Benjamin Schoefer, Jón Steinsson, Andreas Tischbirek, Martin Wolf and seminar and conference participants at the Bank of Canada, the Bank of England, the Banque de France, the Bank of Italy, the Dynare Conference, the European Economic Association, the Federal Reserve Board, the Lausanne Research Days, the JME-SNB-SCG Conference, the Macro Research Cluster, the 2nd edition of QuickTalks at King's College London, the RES Symposium of Junior Researchers, the SNB Research Conference, the Swiss Finance Institute Research Days, the UC Berkeley Macro Lunch, and the Ventotene Workshop in Macroeconomics for helpful comments. All errors are my own.

1 Introduction

The world population has aged rapidly over the past half-century. In the United States, lower fertility rates and longer life expectancies have already increased the share of retired people and reduced the size of the working population. As shown in the left panel of Figure 1, the ratio of these two groups, defined as the old-age dependency ratio, has significantly grown since 1960 and it is projected to rise even further in the following decades. The U.S. is not alone in this demographic transition. Every country is expected to experience similar demographic trends as the U.S. These trends influence many central aspects of the economy and are not limited to the pension system sustainability or labor market participation. Monetary authorities are also not immune to the effects of the changes in the population distribution. Given the magnitude and the increasing pace of these trends, it is of great importance for the monetary authorities to understand the extent to which demographic trends might affect their abilities to achieve their mandates.

Figure 1: Demographic trends

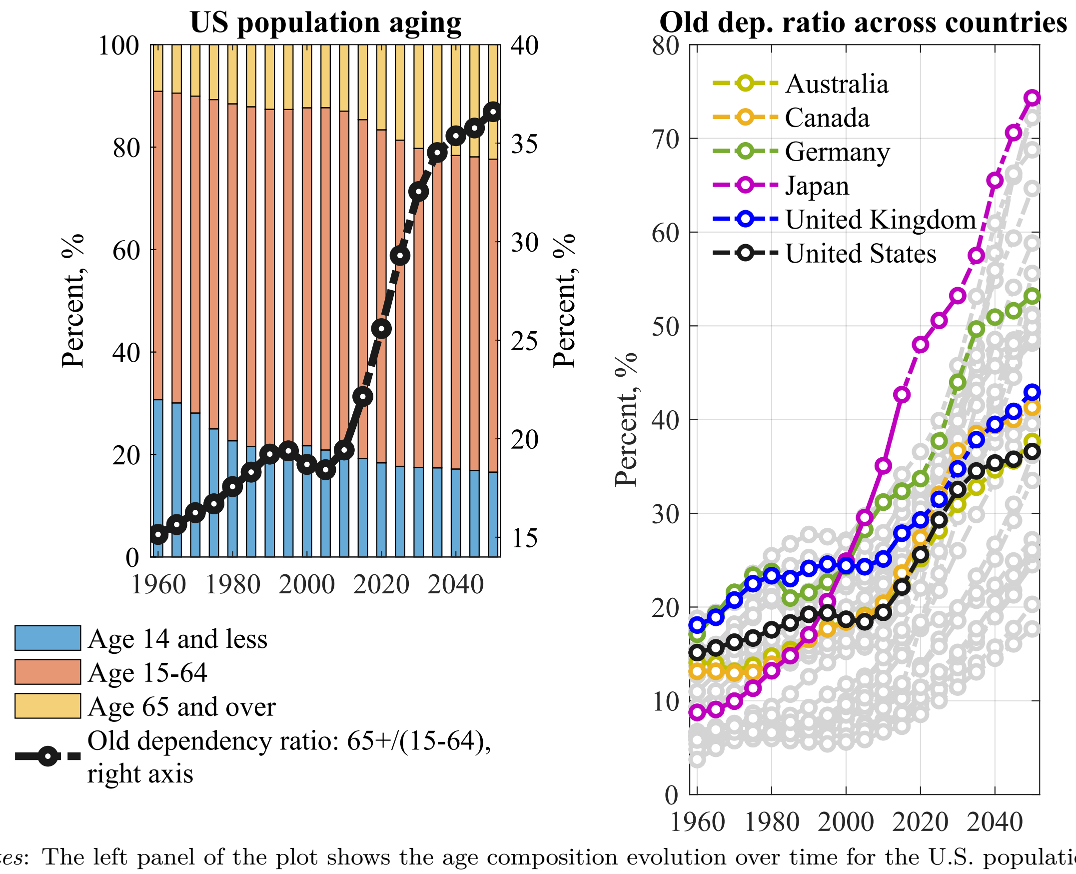

Notes: The left panel of the plot shows the age composition evolution over time for the U.S. population as well as the relative old-age dependency ratio from 1960 to 2050. The right panel compares the time series of the old-age dependency ratio across major economies. The source of the data is the World Bank Population Estimates And Projections.

This paper studies the impact of population aging on the effectiveness of monetary policy. I propose a novel channel to explain how the transmission of monetary policy might be influenced by demographic trends. Older individuals devote a larger share of their expenditures to services, and services tend to adjust their prices less often than goods. As the population ages, the relative importance of services rises leading to an increase in price stickiness. Since fewer firms can adjust their price in response to a monetary shock, output responds more strongly. Using household-level data for the U.S., I document that the negative relationship between age and the frequency of price adjustment of the consumption bundle is driven by significant differences in sectoral expenditure shares across age groups. In line with this micro evidence, I show that populating aging is accompanied by an increase in the relative size of the service sector and that the economic activity of more service-intensive U.S. states is more responsive to monetary shocks. I then use a theoretical model to quantify how much of the change in the effectiveness of U.S. monetary policy from 1980 to 2050 can be accounted for by population aging.

To study the relationship between age and price stickiness, I combine household-level data from the U.S. Consumer Expenditure Survey (CEX) for the period 1982-2018 with the sectoral frequency of price adjustment computed by Nakamura and Steinsson (2008). I find that older households spend significantly more on services. The services expenditure share of households over 80 years old is 20 percentage points higher compared to that of households in their early 30s. At the same time, services adjust their prices on average every 13 months, whereas goods every 3 months. The average frequency at which the price of the consumption bundle is adjusted is highly heterogeneous across age groups ranging from 8.2 months for young households to almost 10 months for older households. This relationship is stable over the sample period and when controlling for other households' characteristics.

Through the lens of a standard 3-equation New Keynesian model, I evaluate how changes in price stickiness affect the responsiveness of output and inflation to monetary shocks. A decrease in the frequency of price adjustment results in a more muted response of inflation since fewer firms adjust their price, but a more substantial response of output, since firms would need to adjust their production more vigorously. However, output and inflation are not equally sensitive to changes in the price stickiness parameter. The response of output is significantly influenced by the frequency of price adjustment, whereas inflation is only marginally affected. This is due to the fact that with higher price stickiness fewer firms can adjust their price every period. Inflation responds less to shocks and also becomes less sensitive to changes in the other macroeconomic variables. Since prices cannot be adjusted, firms respond by adjusting their production more. Moreover, firms anticipate that on average they might not be able to adjust their price for a longer time period. The expectations channel results in a further increase in output responsiveness. Due to the lower sensitivity of inflation to changes in the economy, the increase in output responsiveness has only a marginal impact on the responsiveness of inflation.

The theoretical framework delivers two key predictions on how monetary policy transmission is influenced by demographic trends. An increase in the share of older individuals increases the demand for services resulting in a lower frequency of price adjustment at the aggregate level. Therefore, the first prediction is a stronger response of output following a monetary shock because fewer firms can adjust their price. The second prediction is that the response of inflation in older economies is only slightly more muted because the sensitivity of inflation to changes in the economy is lower. I test these macroeconomic predictions by exploiting the cross-sectional variation in demographic structures and economic activity among U.S. states. I document that the share of services is positively related to population aging. I then compute the responses of state-level real personal income and GDP from the Bureau of Economic Analysis (BEA) as well as inflation rates from Hazell et al. (2022) to a monetary shock adopting a panel local projection approach à la Jordà (2005). Exogenous variations in interest rate are captured using the Romer and Romer (2004) monetary shocks series. By interacting the responses with state-level services intensity, I confirm that the economic activity of states with a relatively higher share of services responds more to monetary shocks. In contrast, the response of inflation is not significantly influenced by the different economic structures.

This empirical evidence motivates the last part of the paper, where I develop a two-sector overlapping generations New Keynesian model to investigate how monetary policy shock propagation is influenced by population aging. The model incorporates a rich demographic structure with age-specific mortality rates, labor productivity, and consumption preferences over the services and goods sectors. The sectors differ in their degree of price stickiness, and only the output from the goods sector can be stored and invested. I calibrate the model to match the realized and projected population distribution and the different sectoral preferences across age groups observed in the data.

The theoretical model is then used to answer the following questions: What is the relationship between monetary policy effectiveness and demographic trends? To what extent does the new channel proposed in this paper contribute to changes in this relationship? And, finally, did population aging play any role in the observed decrease in the sensitivity of inflation to changes in economic activity, i.e., on the flattening of the Phillips curve?

In line with the empirical evidence, the model implies that the change in the U.S. population distribution and mortality rate between 1980 and 2010 increased the contemporaneous response of output to monetary shocks but only marginally affected the response of inflation. Demographic trends alone increased the output response by 6% in 2010 relative to 1980 and in 2050 the response is expected to be 10% higher relative to 1980. The increase in output responsiveness is mainly driven by an increase in the sensitivity of the consumption of younger households to changes in interest rate. Moreover, I find that population aging accounts for around one-third of the overall change in monetary policy effectiveness induced by the higher share of expenditures dedicated to services and that consumption heterogeneity across age groups significantly contributed to that. Finally, through the shift in aggregate demand towards services, demographic trends explain around 10% of the decrease in the slope of the Phillips curve.

Understanding how and through which channels the shifts in demographic structure influence the transmission of monetary policy shocks is crucial for policymakers and central bankers to conduct optimal monetary policy. While in the recent literature, much attention has been dedicated to studying the effects of aging on government debt and fiscal policy, the focus on the implications for monetary policy has been limited. Most of these studies concentrate on the long-term consequences on the level of the interest rate and inflation. Indeed, given the slow-moving pace of demographic trends, the impact of population aging on the transmission of monetary policy shocks has been considered negligible. However, the results of this paper show that population aging can significantly influence the effectiveness of short-term monetary policy.

Related literature. This paper contributes to three strands of the literature. First, the results complement the large body of empirical and theoretical evidence on the relationship between monetary policy and demographic trends. As previously mentioned, most of the literature has focused on the effects on the long-term steady-state level of the interest rates and inflation[1] rather than on the short-term implications. Few exceptions include Fujiwara and Teranishi (2008), Kantur (2013), and Yoshino and Miyamoto (2017), which use a two-agents model with workers and retirees to study the effectiveness of monetary policies from

[1]See, among others, Carvalho et al. (2016), Aksoy et al. (2019), Eggertsson et al. (2019), Papetti (2019), Lis et al. (2020), Papetti (2021), Bielecki et al. (2020), Lisack et al. (2021) and Auclert et al. (2021).

a theoretical perspective. Bielecki et al. (2021) develop a life-cycle model calibrated on the Euro Area to show that demographic trends have contributed to the decline in the natural interest rate and have exacerbated the risk of hitting the lower bound and that the pressure is expected to continue. Finally, Brzoza-Brzezina and Kolasa (2021) study the importance of asset distribution across generations for the redistributive effects of monetary policy.

From an empirical point of view, Wong (2014) and Wong (2021) find that the consumption of younger households tends to respond more to monetary shocks since they refinance or enter new loans as interest rates change. Leahy and Thapar (2022) show that the responses of private employment and personal income are stronger the greater the share of the population between 40 and 65 years of age. In contrast, Kimberly et al. (2021) demonstrate that the consumption of older households is more responsive to monetary policy shocks because of their portfolio composition. Kopecky (2022) provides empirical evidence that population age structure plays an essential role in the relationship between excess money growth and inflation. Using a cointegrated VAR approach for the U.S. and Euro Area, Bobeica et al. (2017) find a positive long-run relationship between inflation and the growth rate of the working-age population. Similarly, de Albuquerque et al. (2020) document in a panel of 24 countries that the 35-64 years old group creates disinflationary pressure while very old population groups appear to contribute strongly to inflation. I contribute by proposing and analyzing a novel channel through which demographic trends might affect monetary policy effectiveness: consumption heterogeneity across age groups.

The second strand is the literature on the time-varying effects of monetary policy shocks on real activity and inflation. Reforms in the institutional structure of the credit markets (Boivin et al., 2010), stronger anchoring of expectations as well as demographic trends (Imam, 2014, Kronick and Ambler, 2019) have been proposed as potential explanations for the fact that the responses of output and inflation to shocks have changed in the last decades. I show that population aging is putting downward pressure on the aggregate frequency of price adjustment increasing output responsiveness and decreasing inflation responsiveness to shocks over time. This result is also confirmed in a cross-country comparison by Galesi and Rachedi (2018) who illustrate that the response of inflation to monetary shocks in countries with a larger share of services intermediaries is more muted but the response of output is stronger.

Finally, this paper relates to the literature that studies the flattening of the Phillips curve, i.e., the positive relationship between inflation and economic activity. The empirical disconnect between inflation and economic activity has been interpreted as potential evidence that the

Phillips curve has weakened or even disappeared (Coibion and Gorodnichenko, 2015, Blanchard et al., 2015, Laurence and Mazumder, 2011). Potential explanations include the successful anchoring of expectations (Bernanke, 2010), the increase in central bank credibility (McLeay and Tenreyro, 2019), global forces (Jorda et al., 2019), and the change in the input-output network (Rubbo, 2022). Related to this last channel, I find that demographic trends shift aggregate demand towards the services sector, which has a lower slope of the Phillips curve, resulting in an overall decrease in the sensitivity of inflation to real activity.

Road map. The remaining paper is organized as follows. Section 2 uses household-level expenditure data to document the negative relationship between age and the frequency of price adjustment. In section 3, I derive which are the theoretical predictions of a change in price stickiness using a standard 3-equation New Keynesian model. Section 4 studies the heterogeneous effects of monetary policy shocks across U.S. states according to their economic structures. In section 5, I develop the two-sector OLG NK model to assess how the transmission of monetary policy shocks in the U.S. has been influenced by demographic trends and to what extent consumption heterogeneity explains this. Finally, section 7 concludes.

2 Micro-level evidence

Using household-level data for the U.S., I document significant heterogeneity in price stickiness across the consumption bundles of different age groups. In particular, older people purchase more services rather than goods and the firms in the services sector tend to adjust less often their prices. Therefore, an increase in the share of old people puts downward pressure on the aggregate frequency of price adjustment.

2.1 Heterogeneity in the frequency of price adjustment

2.1.1 Data

I show how the frequency of price adjustment varies with household age using micro-data for the U.S. To do so, I combine data on expenditure shares from the Consumer Expenditure Survey (CEX) run by the Bureau of Labor Statistics (BLS)[2] for the 1982-2018 period with

[2]The CEX survey respondents are asked about their expenditures for the full consumption basket. The CEX is made up of two separate surveys: the Interview and the Diary. The first one covers the full range of expenditures on a quarterly basis, while the second provides more detailed information at a weekly frequency for certain product categories like food and clothing. A set of demographic characteristics are reported in both surveys. Overall, in the two modules, there are questions regarding around 600 Universal Classification Code (UCC) categories.

the item-level frequency of price adjustment data from Nakamura and Steinsson (2008), which is computed as the fraction of the number of times an item changes its price over the number of times the item is observed[3]. The expenditure data from the CEX are available at Universal Classification Code (UCC) level for about 600 categories whereas the frequency of price adjustment from Nakamura and Steinsson (2008) at the Entry Level Items (ELI) level for 272 categories. Therefore, as in Clayton et al. (2018) and Cravino et al. (2020), I implement a "many-to-one" merge from UCCs to ELIs by summing up the expenditures of all UCCs linked to the same ELI. Because a few ELIs do not find a linked UCC, e.g., rent, the final dataset covers 263 ELIs out of 272[4].

I then aggregate households into age groups based on the reference person's age, that is the age of the household head[5]. The average frequency of price changes for age group a, $\bar{\theta}^a_t = \sum_j \omega^a_{t,j}\theta_j$, is computed as the weighted average of the product-specific frequencies of price changes θ_j from Nakamura and Steinsson (2008) using as weights the age group-specific expenditure shares $\omega^a_{t,j}$ from the CEX[6]. As an alternative measure of price stickiness, I compute the mean implied duration. I define for each ELI category the mean implied duration as $d = \frac{-1}{ln(1-f)}$, where f is the frequency of price adjustment, which measures after how many months, on average, a firm in sector j adjusts its price. I then compute the mean implied duration for each age group a similarly to the frequency of price changes.

Before presenting the price stickiness heterogeneity across age groups, it is useful to see how it evolved over time and how it relates to demographic trends. The core idea of this paper is well summarized in Figure 2. On the left panel, I compare the time series from 1980 to 2018 for the U.S. old-age dependency ratio (left axis) with the scatterplot of the share of consumption devoted to services as well as the relative polynomial fit (right axis). The distinction between goods and services, which I will discuss more in detail later, is extremely important for my analysis since the share of services consumed increases over the life cycle (with the share for older households being around 20 percentage points more than for younger households) and because the two categories have remarkably different frequencies of price adjustments (goods adjust on average every 3 months whereas services every 13 months). On the right panel, I compare the same time series of the U.S. old-age dependency ratio (left axis)

[3]Figure 24 reports heterogeneity in price rigidities across 19 categories and between goods and services.

[4]See section A of the Online Appendix for more details about the data.

[5]The results are similar if it is used the average age across all household members.

[6]The implicit assumption I make is that the frequency of price adjustment at sectoral level θ_j is constant over time. This assumption is partly tested by Nakamura and Steinsson (2008) who compare the frequency of price adjustment over two different periods, 1988-1997 and 1998-2005, and they show that the parameters are rather stable over time.

Figure 2: Old-age dependency ratio, service share, and price stickiness

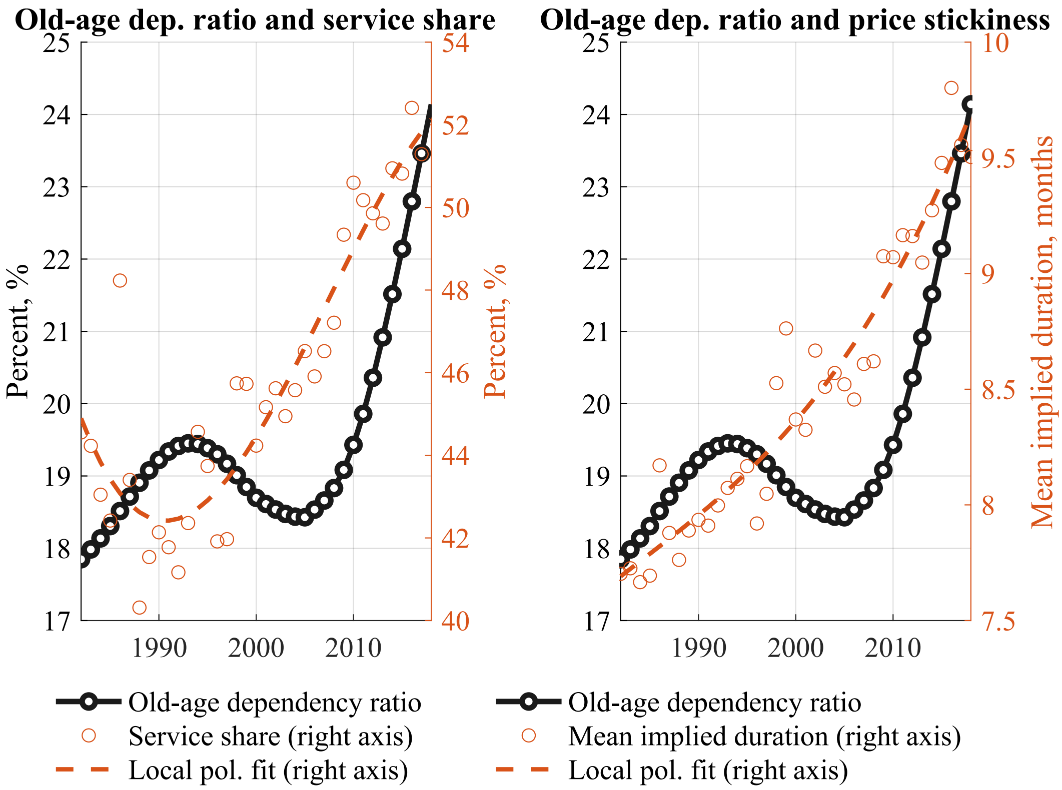

Notes: The left panel of the plot shows the evolution of the U.S. old-age dependency ratio over time (left axis) alongside the time series of the share of consumption devoted to services (right axis). The right panel compares the time series of the U.S. old-age dependency ratio with the mean implied duration of prices (right axis). The source of the data is the World Bank Population Estimates And Projections as well as the CEX data.

with the scatterplot of the mean implied duration as well as the relative polynomial fit (right axis). The old-age dependency ratio in the U.S. increased throughout the 80s and until the mid-90s. It slightly decreased in the subsequent 10 years and then it significantly rose again and is expected to keep rising in the next decades as shown in Figure 1.

The evolution of the demographic structure can be considered to some extent exogenous but, despite being rather slow-moving, it is likely to have non-negligible effects on the overall economy. In particular, as shown in Cravino et al. (2022), population aging explains around a fifth of the increase in the share of services consumed over the last 40 years which overall rose from 44% to 52%. Moreover, given that firms in the services sector adjust their prices much less frequently than firms in the goods sector, the rise in the share of services resulted in a decrease in the overall frequency of price adjustment with the mean implied duration increasing from around 8 months to 9.5 months. Therefore, since demographic trends contributed to the change in the share of services, they are also partially responsible for the observed decrease in the frequency of price adjustment. As every standard New Keynesian model predicts, the lower the frequency of price adjustment, the stronger the response of output and the more muted the response of inflation to monetary policy shocks.

2.1.2 Price stickiness across age groups

In this section, I document significant heterogeneity in price stickiness across age groups due to the different expenditure categories they consume. Figure 3 plots the weighted average frequency of price adjustment for each age group, $\bar{\theta}^a$ (left axis). There is a clear and significant negative correlation between age and the consumption bundle's price adjustment frequency. The average frequency of price adjustment for households above the age of 80 years is more than 20% lower than that of households between the ages of 15 and 25 years[7]. Figure 25 also reports the mean implied duration for each age group. The mean implied duration significantly increases over the life cycle from around 8.4 months to almost 9.8 months.

Figure 3: Frequency of price adjustment and services consumption across age groups

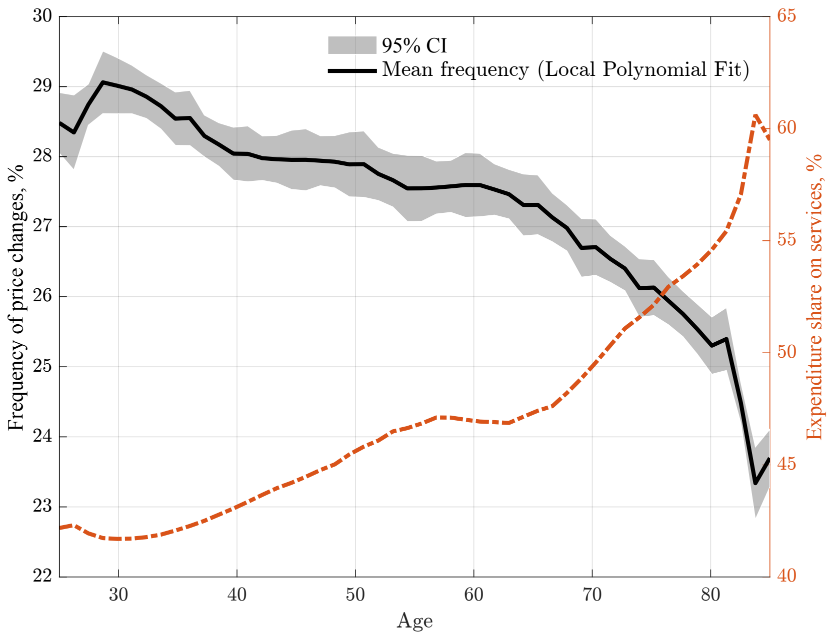

Notes: The figure plots the weighted average frequency of price adjustment (left axis) alongside the share of consumption devoted to services (right axis) across age groups. The shaded area is the 95% confidence band. The frequency of price adjustment is computed as the fraction of the number of times an item changes its price over the number of times the item is observed and expressed in percent per month. The expenditure shares are computed using data from the CEX whereas the sectoral price stickiness parameters are retrieved from Nakamura and Steinsson (2008).

The main driver behind this negative relationship is the higher share of services consumed by older households. As it can be noticed in Figure 3, the share of consumption devoted

[7]As shown in Figure 25, excluding temporary sales in the computation of the frequency of price adjustments shift the entire relationship downward but does not affect the relative relationship across age groups.

to services (right axis) increases from around 40% for younger households up to 60% for older ones[8]. Nakamura and Steinsson (2008) document that services tend to have a much higher level of price stickiness with an average price duration of 13 months compared to a 3 months duration for goods[9]. Given the heterogeneity in consumption bundles across age groups and the different frequency of price adjustments across sectors, the expenditures of older households are characterized by a much stronger price stickiness relative to young households[10].

To shed further light on which categories mainly drive the relationship between age and price stickiness, I focus now on more granular expenditure categories. Table 7 shows the expenditure shares across some age groups for twenty of the main consumption categories. In line with previous findings, the largest disparity can be observed in health expenditures where the average consumption share of households above the age of 80 years is almost 16 percentage points larger than that of households below the age of 25 years. Moreover, younger households tend to spend relatively more on categories like Education, Entertainment, and Private Transportation. In contrast, Energy and Household Furnishings and Operations constitute a larger component of the older household consumption bundle.

The left panel of Figure 4 plots the frequency of price change on the y-axis against the difference in the expenditure shares between the age groups $(75; 80]$ and $(25; 30]$ on the x-axis. A positive value means that the older group has higher expenditure shares in that category. Most of the categories gather around zero suggesting that the two age groups have similar expenditure shares. However, the categories more intensively brought by older households tend to be characterized by a lower frequency of price adjustment while the opposite holds for the categories mainly purchased by younger households. The correlation between the x-axis and y-axis variables is -0.153.

[8] I classify as Goods the following expenditure categories: Food at home, Vehicle purchasing, Gas, Entertainment equipment, Appliances, furniture and fixtures, Alcoholic beverages, Clothing and other apparel, Tobacco, Personal care goods. I classify them as Services: Health, Utilities, Car maintenance, Repairs and insurance, Food away from home, Domestic services and childcare, Education, Entertainment services, Public transport, and Personal care services.

[9] Several potential explanations have been suggested in the literature to explain the difference in price stickiness between the two sectors. For example, the production of services is much more labor-intensive than the production of goods. The high wage stickiness might then translate into a lower frequency of price adjustment for services. Alternatively, services face lower price competition since innovation is less common in the services sector than in the goods sector. Micro funding the different frequencies of price adjustment across sectors is beyond the scope of this paper and, therefore, in the theoretical model price stickiness is set exogenously.

[10] Cravino et al. (2022) test more systematically the relationship between age and the share of consumption devoted to services. The authors control for income decile dummies and region-time fixed effects and they still document large differences in service expenditures across households of different age groups.

Figure 4: Expenditure differences across age group

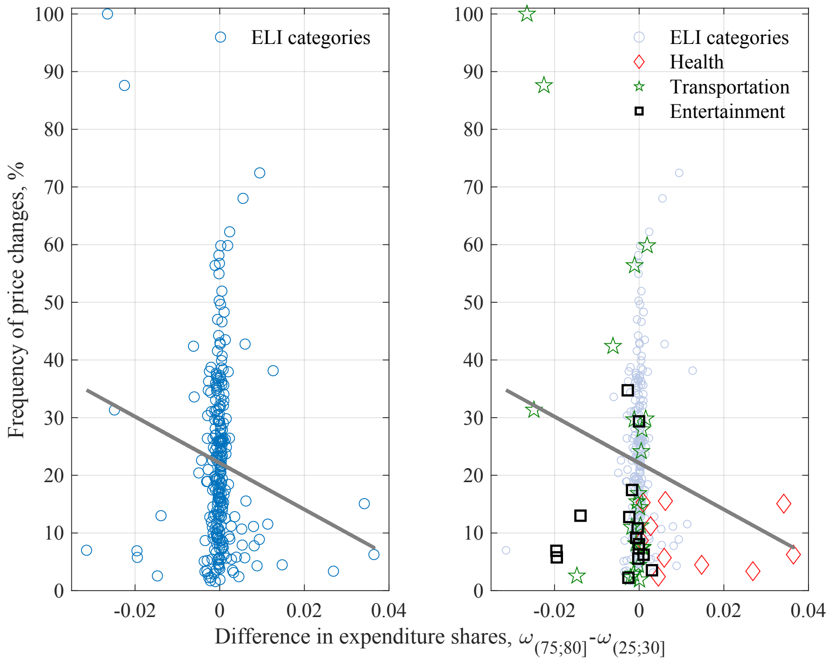

Notes: The left panel plots the frequency of price adjustment against the difference in sectoral expenditure shares for the age groups (75; 80] and (25; 30]. The right panel shows the same plot highlighting some important categories: Entertainment, Health, and Transportation. The fitted linear regression line of the data is included in both panels.

On the right panel of Figure 4, I highlight some of the categories for which expenditure heterogeneity is more evident. As previously mentioned, medical expenses are a major component of the elderly consumption bundle and at the same time, they are characterized by an extremely low frequency of price adjustment. The opposite is true for Transportation: younger households spend more on these categories and the firms in this sector are able to adjust their prices more frequently.

2.2 Decomposing the rise in services share

Since older people allocate a larger share of their consumption towards services, and since services tend to adjust their price much less frequently, an increase in the share of old people will increase the aggregate demand for services resulting in a lower frequency of price adjustment at the aggregate level.

To quantify the contribution of observed changes in the age distribution to the observed changes in services shares in the U.S. between 1982 and 2018, I carry out a shift-share decomposition similar to Cravino et al. (2022). This exercise allows us to quantify to what extent the increase in the share of services is due to the change in expenditure shares within

Table 1: Within-between decomposition, 1982 to 2018

	Services share	Contribution	Implied duration, months
Within	0.058	80 %	1.44 (+18.70 %)
Between	0.015	20 %	0.36 (+4.68 %)
Total	0.073 (44.95 % to 52.23 %)	100 %	1.80 (+23.38 %) (7.70 to 9.50)

age groups, i.e., each age group consumes more services but the share of aggregate expenditure of each age group is the same, and to what extent is due to reallocation of expenditures between groups, i.e., the share of services for each age group is unchanged but the age groups which have a higher share of services now account for a larger share of aggregate expenditure.

The share of services in aggregate consumption can be written as:

$$\alpha_t^s = \frac{\sum_a C_t^{s,a}}{\sum_a \sum_j C_t^{j,a}} = \sum_a \alpha_t^{s,a} s_t^a \tag{1}$$

where $\alpha_t^{s,a} = \frac{C_t^{s,a}}{\sum_j C_t^{j,a}}$ is the within age group share of expenditure devoted to services and $s_t^a = \frac{\sum_j C_t^{j,a}}{\sum_a \sum_j C_t^{j,a}}$ is the share of age group a in aggregate expenditure.

I can then decompose the change in services between two periods t_1 and t_2 as:

$$\Delta\alpha_t^s = \underbrace{\sum_a \Delta\alpha^{s,a}\bar{s}^a}_{\text{Within}} + \underbrace{\sum_a \bar{\alpha}^{s,a}\Delta s^{s,a}}_{\text{Between}} \tag{2}$$

with $\Delta x = x_{t_2} - x_{t_1}$ and $\bar{x} = \frac{x_{t_2} - x_{t_1}}{2}$ for any variable x. The term "*Within*" captures changes in the age-specific expenditure shares keeping age distribution fixed whereas the term "*Between*" captures changes in the share of age group a in aggregate expenditures keeping the preferences fixed.

I compute the within-between decomposition using the CEX data for 1982 and 2018 and report the results in Table 1. The services share increased by 7.3 percentage points between the two periods considered (first column) and 20% of the increase is attributed to between age group changes in expenditures (second column). The remaining 80% is due to changes in expenditure shares within groups. This result is in line with the findings in Cravino et al. (2022).

In terms of contribution to the change in price stickiness observed in Figure 2, from 1982 to 2018 the mean implied duration increased by 1.8 months, from 7.7 months to 9.5 months (third column); an increase of approximately 23%. Of this, the between-age group changes in expenditures alone account for 0.36 months.

2.3 Robustness

I first control that the negative relationship between age and frequency of price adjustment is stable over time. Figure 5 shows the same pattern for different periods. There is some marginal variation across time periods, partly due to the fact that some consumption categories are dropped and some are added, and partly due to actual changes in expenditure weights. However, the main conclusion still holds: the frequency of price adjustment decreases with age.

Figure 5: Frequency of price adjustment across age groups and time

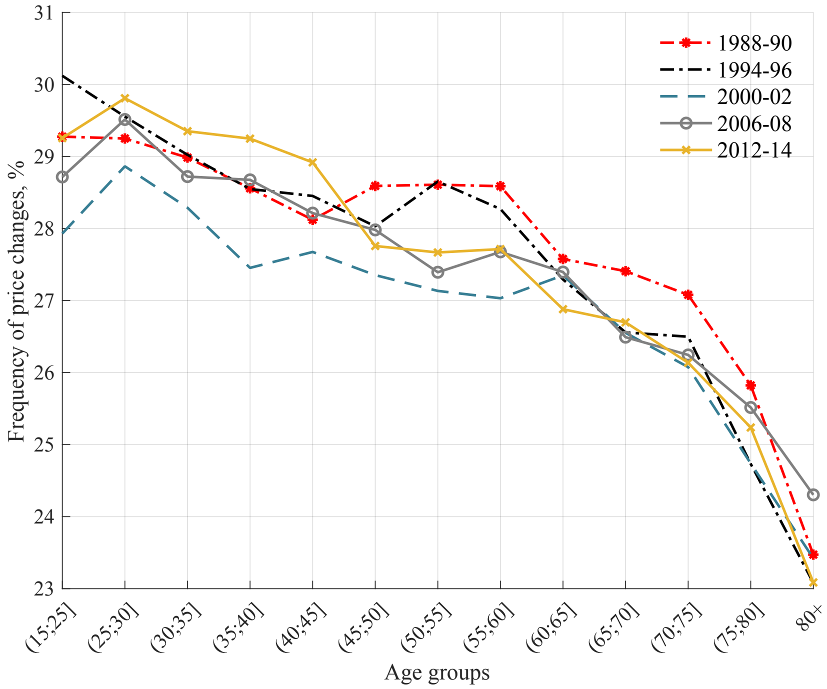

Notes: The figure plots the weighted average frequency of price adjustment at the age group level across five different time periods. The frequency of price adjustment is computed as the fraction of the number of times an item changes its price over the number of times the item is observed and expressed in percent per month. The source of the data is the CEX.

A potential source of concern regarding the findings in Figure 3 is that these patterns might be explained by demographic characteristics other than age. Indeed, Clayton et al. (2018) show that prices are more rigid in sectors selling to college-educated households whereas Cravino et al. (2020) demonstrates that price stickiness displays an inverse U-shaped distribution across income groups.

To control that these demographic characteristics do not drive the results, I compute the frequency of price adjustment across age groups conditioning on the education level of the respondents as well as on the consumption quantile to which they belong[11].

Figure 6: Frequency of price adjustment across age groups, education levels, and consumption quantiles

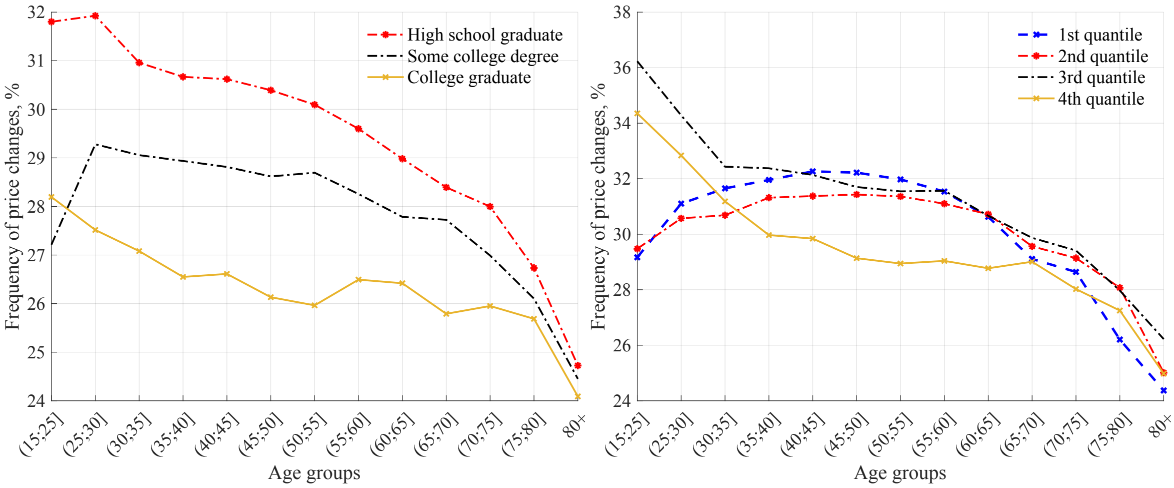

Notes: The left panel plots the weighted average frequency of price adjustment at the age groups level for three different education levels. The right panel reports the weighted average frequency of price adjustment at the age group level for different consumption quantiles. The frequency of price adjustment is computed as the fraction of the number of times an item changes its price over the number of times the item is observed and expressed in percent per month.

The left panel of Figure 6 confirms that the consumption bundles of college-educated households have a lower frequency of price adjustment as in Clayton et al. (2018). In line with the findings of Cravino et al. (2020), the right panel of Figure 6 shows that the average frequency of price adjustment tends to decrease along the consumption distribution. However, conditioning on education level as well as on consumption does not weaken the negative relationship between the frequency of price adjustment and age.

Finally, I check that no outlier in the expenditure categories is responsible for the pattern observed. Figure 7 shows that aggregating the 263 items into less and less granular groups does not remarkably affect the observed negative relationship between age and frequency of price adjustment. In particular, the classification of each expenditure category into goods or services almost entirely captures the relationship of interest

[11] Cravino et al. (2020) use the imputed income level which is available only from 2004 onward. For this reason, I use consumption level as a proxy for income. Moreover, since the households interviewed in the Interview survey are not the same ones interviewed in the Diary survey, for this robustness check I focus only on the Interview survey.

Figure 7: Frequency of price adjustment across age groups, alternative aggregation

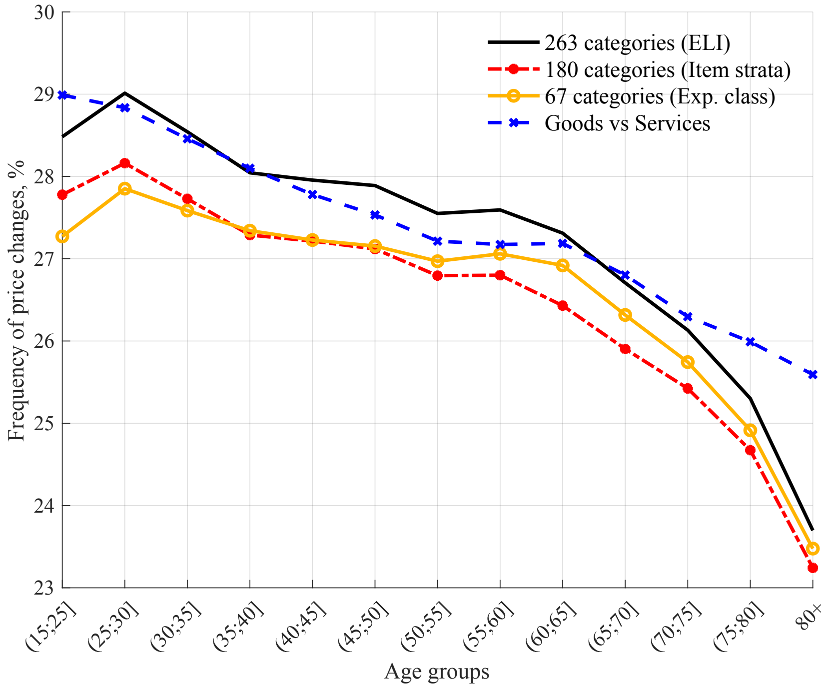

Notes: The figure plots the weighted average frequency of price adjustment across age groups when the expenditure categories are aggregated at ELI, Item Stata, and Expenditure Class level as well as Goods and Services. The frequency of price adjustment is computed as the fraction of the number of times an item changes its price over the number of times the item is observed and expressed in percent per month.

3 The 3-Equation New Keynesian model

In the previous section, I document that the frequency of price adjustment has decreased over time. One of the main reasons is that the share of services has significantly increased in the last 40 years and the firms in these sectors tend to adjust their prices less often. Part of this structural transformation can be explained by demographic trends: old households consume a larger share of services relative to young households so the share of expenditure devoted to services increases as the population ages. To evaluate how the increase in price stickiness induced by population aging might affect the propagation of monetary shocks, I start with a standard 3-equation New Keynesian model[12].

The three equations of the model are the IS curve (3), the Phillips curve (4), and the interest rate rule (5). These equations relate the output gap $\hat{x}_t$ (defined as the deviation of

[12]The derivation of the model is rather standard in the literature so I refer the interested reader to Galí (2015).

output from its flexible price counterpart), the inflation rate $\hat{\pi}_t$ and the real interest rate $\hat{r}_t$:

$$\hat{x}_t = -\frac{1}{\sigma}\left(\hat{r}_t - E_t\hat{\pi}_{t+1}\right) + E_t\hat{x}_{t+1} \tag{3}$$

$$\hat{\pi}_t = \beta E_t\hat{\pi}_{t+1} + \kappa(\sigma+\eta)\hat{x}_t \tag{4}$$

$$\hat{r}_t = \phi_\pi\hat{\pi}_t + \phi_x\hat{x}_t + \nu_t \tag{5}$$

where $\kappa \equiv \frac{(1-\theta)(1-\beta\theta)}{\theta}$ is the slope of the Phillips curve. All variables are expressed in log deviation from a zero inflation steady state. σ is the intertemporal elasticity of substitution, β is the discount factor, η is the Frisch elasticity of labor supply and θ is the fraction of firms that cannot reset their prices each period. The interest rate rule coefficients, ϕ_π and ϕ_x, capture the response of the central bank to changes in inflation and output gap respectively. We assume that the monetary policy shock ν_t follows an AR(1) process with persistence ρ:

$$\nu_t = \rho\nu_{t-1} + \varepsilon_t^\nu \tag{6}$$

It is possible to express the output gap and the inflation as a function of only the monetary policy shock and the model parameters using the method of undetermined coefficients[13]. It can be shown that:

$$\hat{x}_t = -\left(1-\beta\rho\right)\Lambda_\nu\nu_t \tag{7}$$

$$\hat{\pi}_t = -\kappa\Lambda_\nu\nu_t \tag{8}$$

where $\Lambda_\nu \equiv \frac{1}{(1-\beta\rho)[\sigma(1-\rho)+\phi_y]+\kappa(\phi_\pi-\rho)}$. If the conditions for a unique stationary equilibrium are satisfied, Λ_ν is greater than zero so both the coefficients $\left(1-\beta\rho\right)\Lambda_\nu$ and $\kappa\Lambda_\nu$ are positive. Therefore, an expansionary monetary policy shock, i.e., a decrease in ν_t, leads to a persistent increase in the output gap and inflation.

However, the two coefficients differ in magnitude as well as in terms of their sensitivity to changes in the frequency of price adjustment. To see this, I set the model parameters to their standard value in the literature[14], and I compute the contemporaneous response of the output gap and inflation to a 100 basis point expansionary shock, i.e., $\nu_t = -1$, as a function of the

[13]See Chapter 3 of Galí (2015).

[14]$\sigma = 1$ such that the utility function is in log-form, $\beta = 0.995$, $\eta = 1$, $\phi_\pi = 1.5$, $\phi_x = 0.2$ and $\rho = 0.8$.

Figure 8: Contemporaneous response of output gap and inflation as a function of price stickiness

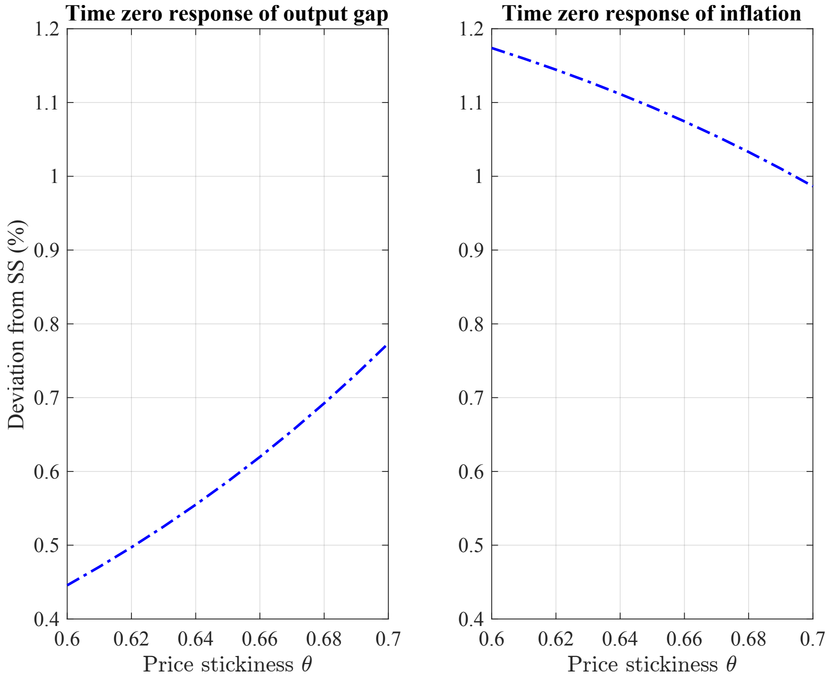

Notes: The figure plots the contemporaneous response of output gap (left panel) and inflation (right panel) to a 100 basis point decrease in interest rate as a function of the price stickiness parameter θ.

price stickiness parameter θ. From 1980 to 2020 the mean implied duration has increased from 7.5 months to almost 10 months as one can see from Figure 2, which would suggest that the price stickiness parameter has changed from 0.6 to 0.7 so I consider this interval.

The relationships between the contemporaneous responses and price rigidity are reported in Figure 8. First of all, the size of the inflation coefficient is significantly larger than the output one resulting in a stronger response of inflation to the monetary shock. Second, the relationship is upward-sloping for the output gap but downward-sloping for inflation confirming that an increase in price stickiness results in a more muted response of inflation to shocks (fewer firms can adjust their price) but stronger for output (firms need to adjust their production since they cannot adjust their prices). Third, the response of inflation is remarkably less sensitive to changes in price rigidities. Increasing the price stickiness parameter from 0.6 to 0.7 increases the time zero response of output by 75% (from 0.44% to 0.77%) whereas it decreases the response of inflation by only 20% (from 1.17% to 0.98%).

The different sensitivities of inflation and output to changes in price stickiness is due to the fact that a lower frequency of price adjustment implies that fewer firms can adjust their price every period. Therefore, following a monetary shock the response of inflation is more muted and inflation also becomes less sensitive to changes in the other macroeconomic

Figure 9: Impulse response functions from the 3-equation NK model

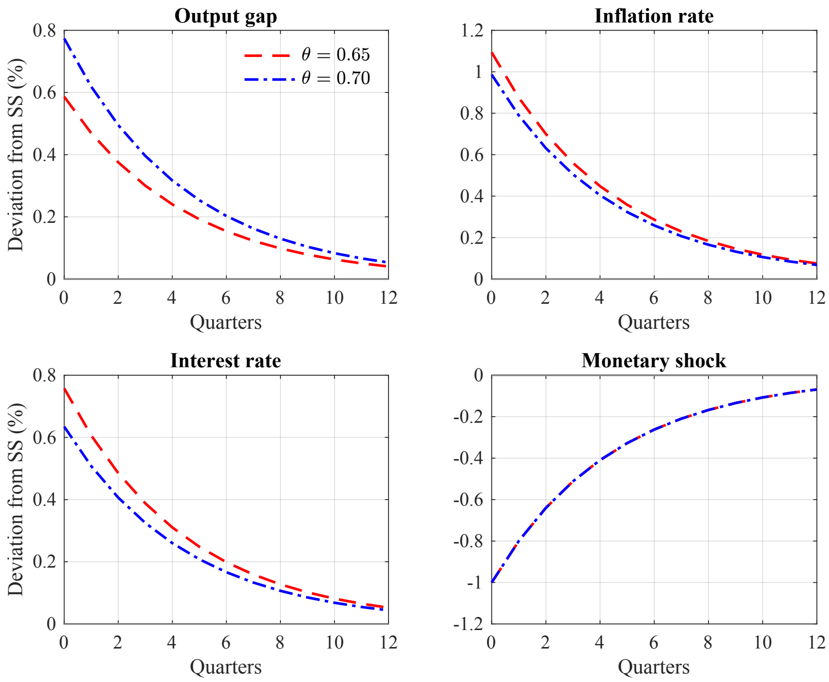

Notes: The figure plots the impulse responses of the output gap, inflation, interest rate, and the monetary policy shock from the 3-equation NK model. The red lines are relative to the model with the price stickiness parameter θ sets to 0.6 and the blue line to 0.65.

variables, i.e., $\kappa(\sigma + \eta)$ from the Phillips curve (4) is decreasing when θ increases. The firms that cannot adjust their price respond by adjusting their production more. On top of that, firms anticipate that on average they might not be able to adjust their price for a longer time period. The expectations channel increases the responsiveness of output even more. Due to the lower sensitivity of inflation to changes in the economy, the increase in output responsiveness has only a marginal impact on the responsiveness of inflation.

Figure 3 shows that the mean implied duration across age groups varies from 8.5 months to almost 10 months. So if everyone had the same consumption bundle of young households the price stickiness parameter would be 0.65 while if everyone had the same consumption bundle of old households it would be 0.7. Therefore, to get a sense of the magnitude we could expect to find empirically, I compute the impulse response functions of output, inflation, and interest rate following a decrease of 100 basis points in ε_t^{ν} when θ is set to 0.65 and 0.7. The responses are reported in Figure 9.

Following the expansionary monetary policy shock both the output gap and inflation increase. As expected from the previous analysis, the response of output is smaller in magnitude than the response of inflation. Moreover, increasing the price stickiness parameter from 0.65 to 0.7 results in a stronger response of output and in a more muted response of inflation. Finally, it is important to notice how the response of output is also much more sensitive to the change in the frequency of price adjustment relative to the response of inflation. Indeed, under these two extreme scenarios the former increases by approximately 30% whereas the latter decreases by less than 10%.

Overall the results from the standard 3-equation NK model suggest that the impact of demographic trends on the transmission of monetary shocks is asymmetric between output and inflation. The decrease in the frequency of price adjustment due to the heterogeneity in consumption bundle across age groups is expected to significantly increase the responsiveness of output and will have a more negligible effect on inflation. In the next section, I empirically test these hypotheses by exploiting the cross-sectional variation in demographic structures and economic activity across U.S. states.

4 Macro-level implications: across U.S. states comparison

In section 2 I provide evidence of a positive relationship over time between the mean implied duration, the services share, and the old-age dependency ratio which might influence the way monetary policy shocks propagate in the economy. At the aggregate level, a decrease in the frequency of price adjustment leads to a more muted response of inflation (since only a smaller fraction of firms resets their price every period) and to a stronger response of output (since firms that are unable to reset their prices need to respond by adjusting their production). As I document in section 3, these variations are not expected to be symmetrical for output and inflation. In particular, the response of output should be much more sensitive to changes in price stickiness than that of inflation.

To test the macro-level implications of the micro-level results I find, ideally, I would like to compare how economic activity reacts to shocks in periods of a high and low old-age dependency ratio. However, as shown in Figure 1, the demographic structure in the U.S. evolved slowly in the past decades so this state-dependent approach is not feasible since there is basically no variation over time. Therefore, I compensate for the lack of time variation by exploiting the cross-sectional variation in the old-age dependency ratio across U.S. states. First, I document that within-state population aging is related to a shift in economic activity toward service sectors. Second, I provide new empirical evidence that the economic activity of more service-intensive states is more responsive to monetary policy shocks in line with the predictions from the household-level data.

4.1 Data

I collect state- and country-level macroeconomic variables from different sources. The main variable of interest at the state level is the real personal income, sectoral employment, sectoral and aggregate GDP from the Bureau of Economic Analysis (BEA) as well as the annual inflation rate from Hazell et al. (2022). Whereas personal income and inflation rate are available at a quarterly frequency, the GDP is available only at an annual frequency. The country-level variables that are used as controls are collected from FRED and include the industrial production (IP), the consumer price index (CPI), the federal funds rate (FFR), the unemployment rate, and the commodity price index computed by Ramey (2016). I also include information on state population size and demographics from the U.S. Census Bureau.

Figure 10: Old-age dependency ratio across U.S. states and over time

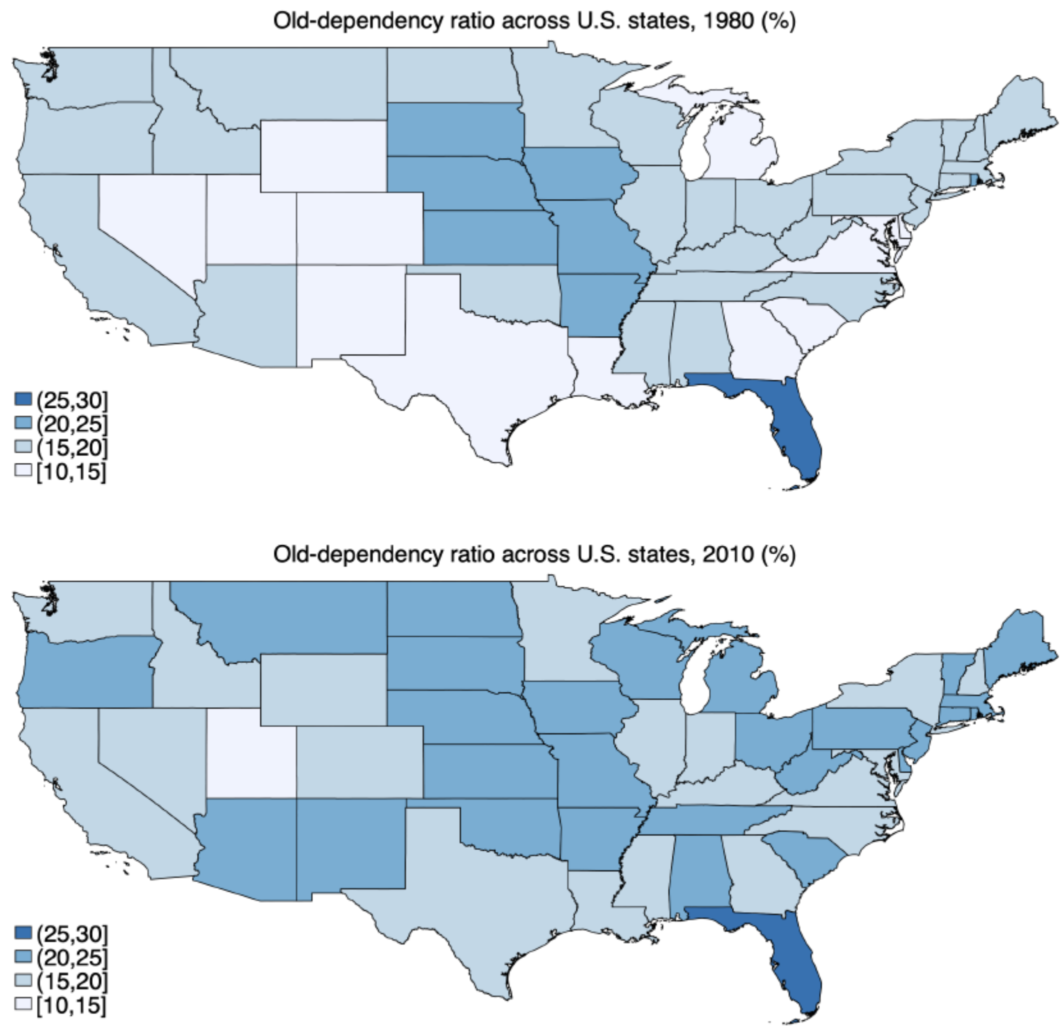

Notes: The figure shows the old-age dependency ratio across the U.S. in 1980 (top panel) and 2010 (bottom panel) using data from the Census Bureau.

Figure 10 shows the significant heterogeneity across states in terms of demographic structure for two different periods that I will use in the theoretical exercise, that is in 1980 (top panel) and in 2010 (bottom panel). These maps illustrate the substantial variation across states in both years, with the old-age dependency ratio ranging from 11% to 27% in 1980 as well as the shift in demographics across almost every state in the U.S. over the past 3 decades. From 1980 to 2010 the state-level old-age dependency ratio has increased on average by 3 percentage points so the demographic variation I exploit comes from the cross-sectional differences in the demographic structures across states rather than from the within-state variations over time.

4.2 Population aging and the service sector

I start by studying how changes in demographic structure are related to the structural transformation that shifts the economic activity from manufacturing to services. Figure 11 shows the relationship between the old-age dependency ratio and the share of GDP from services at the U.S. state level from 1965 until 2020. There is a striking positive correlation between the two variables both at the state level as well as aggregate level, i.e., the higher the state-level old-age dependency ratio the higher the share of GDP from the service sector.

Several confounding factors could explain the positive relationship documented in Figure 11. For example, as the population of a state gets older the state also becomes richer leading to a rise in the relative size of the service sector. To more formally evaluate how populating aging contributes to the increase in the service sector, I follow the same empirical specification from Cravino et al. (2022):

$$\omega_{i,t} = \alpha_i + \beta Age_{i,t} + \gamma GDP\_pc_{i,t} + \varepsilon_{i,t}, \tag{9}$$

where $\omega_{i,t}$ is the share of GDP from the service sector for state i at time t, α_i is a state fixed effect, $Age_{i,t}$ is the state-level old-age dependency ratio, and $GDP\_pc_{i,t}$ is the state level log of the GDP per capita. Standard errors are clustered at the state level.

The results from equation (9) are reported in Table 2. The positive coefficient β in Column (1) indicates that indeed aging is associated with a reallocation of economic activity towards the service sector, even after controlling for income. The finding is in line with Cravino et al. (2022) who document the same relationship for a broad set of developed countries.

Figure 11: State-level old-age dependency ratio and services share of GDP

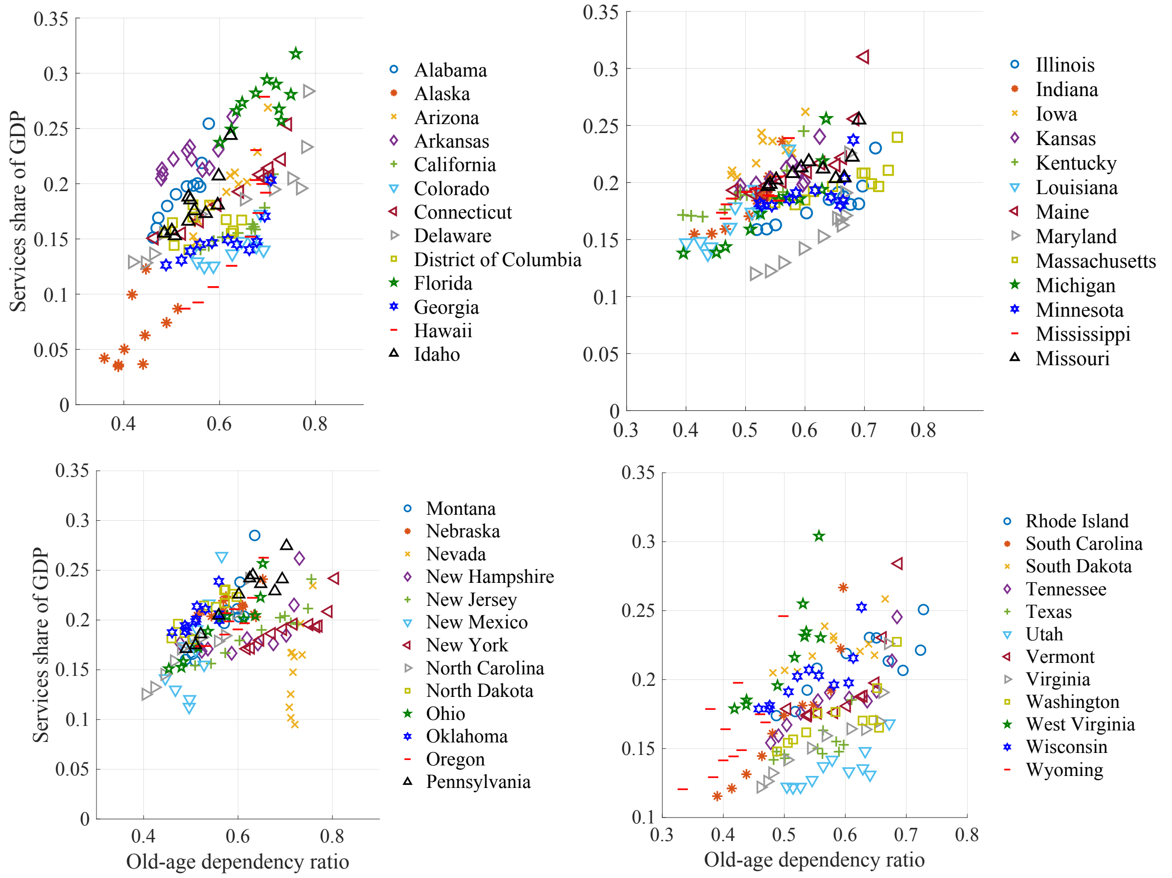

Notes: Each panel reports the scatter plot of the relationship between the old-age dependency ratio and the share of GDP from services at the U.S. state level. For ease of exposition, the states are split into four groups in alphabetical order. Each dot represents the average over a 5-year period starting from 1965 until 2020. *Source*: Bureau of Economic Analysis and U.S. Census Bureau.

The result also holds controlling for the square of the log of GDP per capita, adding a battery of extra controls[15], and using as dependent variable the share of employment in the service sector as in Columns (2) to (4) respectively. Finally, to tackle any concern about the endogeneity of the demographic structure, I follow Shimer (2001) and Rachedi and Basso (2021) and instrument the old-age dependency ratio with lagged birth rates. As baseline instrument I use the 20-year lagged birth rates but quantitatively similar results are obtained with the 25-year lagged birth rates or the 20-30 year lagged birth rates as in Rachedi and Basso (2021). The results from the IV specification are reported in Column (5)[16] and the coefficient of the variable of interest is still positive and statistically significant. Overall, these

[15] I include the state level share of male workers, white workers, college-educated workers, small firms, young firms, aggregate inflation rate as well as inflation rates in the tradable and non-tradable sector, log of housing price, establishment deaths and births.

[16] Cragg-Donald F-tests for weak instruments are reported.

Table 2: Population aging and the service sector

	(1) Services share GDP	(2) Services share GDP	(3) Services share GDP	(4) Services share empl.	(5) Services share GDP
Old-age dep. ratio	0.436***	0.394***	0.610***	0.141***	1.366***
	(0.0985)	(0.0981)	(0.123)	(0.0462)	(0.130)
Log(GDPpc)	0.0587***	0.0248*	0.148	0.167*	-0.0193***
	(0.00475)	(0.0129)	(0.200)	(0.0862)	(0.00687)
Log(GDPpc)*Log(GDPpc)		0.00604**	-0.0377	-0.0205*	0.0117***
		(0.00260)	(0.0286)	(0.0119)	(0.00107)
Observations	2595	2595	702	1149	2172
R^2	0.901	0.903	0.959	0.910	0.724
Extra controls	No	No	Yes	No	No
$C-DF$					465.700

Standard errors in parentheses

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

findings confirm that population aging is an important driver behind the observed increase in the service sector as shown in Figure 11.

4.3 Regional responses to monetary policy shocks

Having shown that demographic trends lead to an increase in the relative size of services, I now investigate how states with different levels of service intensity are heterogeneously affected by monetary policy shocks. The predictions from the theoretical model in section 3 are that, due to the higher level of price stickiness in the service sector, the economic activity in more service-intensive U.S. states should be more sensitive to changes in interest rate. At the same time, no significant differences should be found for the responses of inflation.

I compute the average state-level response to a monetary policy shock by estimating a panel local projection à la Jordà (2005):

$$y_{i,t+h} = \alpha_{i,h} + \beta_h MP_t + \theta_{i,h} X_{i,t-1} + \gamma_h X_{t-1} + \epsilon_{i,t+h} \tag{10}$$

for different horizons $h = 1, ..., 16$. As dependent variable $y_{i,t}$ I use the state-level log of real personal income, the annual inflation rate, and the log real GDP. As monetary shocks MP_t I use the narrative based Romer and Romer (2004) shocks and include state fixed effects $\alpha_{i,h}$. As state controls $X_{i,t-1}$ I use the lagged dependent variable and the log of the population size whereas as aggregate controls X_{t-1} I follow Ramey (2016) by including IP, CPI, FFR, unemployment rate, and commodity price index. To deal with the potential endogeneity, all control variables, except for the monetary policy shocks, are lagged by one period. Standard errors are clustered at the state level. The main coefficient of interest is β_h which captures the impact of monetary policy shocks on the dependent variable over the horizon h.

To evaluate how the different service intensities across U.S. states influence monetary policy effectiveness, I follow an approach similar to the one proposed by Cloyne et al. (2022) and Jamilov et al. (2023). I define a dummy variable $D^S_{i,t}$ equal one when the ratio of services to manufacturing GDP in state i at time t and zero otherwise is in the top quintile of the cross-sectional distribution and I interact the dummy with the monetary shock MP_t:

$$y_{i,t+h} = \alpha_{i,h} + \delta_{i,h} + \gamma_h D^S_{i,t} + \beta^S_h D^S_{i,t} MP_t + \theta_{i,h} X_{i,t-1} + \epsilon_{i,t+h}, \tag{11}$$

where $\delta_{i,h}$ is the time fixed effects that absorb the monetary shocks and the aggregate variables. The coefficients β^S_h capture how states are heterogeneously affected by monetary policy shocks according to their service intensity. The interaction coefficient can be interpreted as the differential response to a contractionary monetary shock of high service-intensive states (for which the ratio of services to manufacturing GDP belongs to the top 20%) relative to the baseline group (states whose ratio belongs to the bottom 80%).

I start by focusing on the impact of monetary policy shocks on the log of real personal income. The left panel of Figure 12 plots the estimated β_h coefficient at different horizons h from equation (10). The dark and light-shaded areas are respectively the 68% and 95% standard deviation confidence intervals. Following a contractionary monetary shock, that is an exogenous increase in interest rate, the real personal income decreases by 0.4% after 3 years. The magnitude and the shape of the response are in line with those found in the literature.

The right panel of Figure 12 plots the estimated β^S_h coefficients from equation (11). The negative coefficients suggest that the economic activity of more service-intensive states is more responsive to monetary shocks. The effect is statistically significant and economically meaningful. After twelve quarters these regions experience a decrease in real personal income up to 0.5% larger relative to the baseline group which is particularly sizable once compared to the overall average decrease of about 0.5%.

Figure 13 reports the same responses using the annual inflation rate as the dependent variable. Following a contractionary shock, the annual inflation rate, after an initial increase, decreases by approximately 0.4 percentage points. In line with the theoretical predictions of section 3, I find no significant differences in the inflation responses across states along the service intensity distribution. As I will show in section 5, also the more complex model I develop is able to replicate this result.

Figure 12: Impact of monetary policy on the regional real personal income

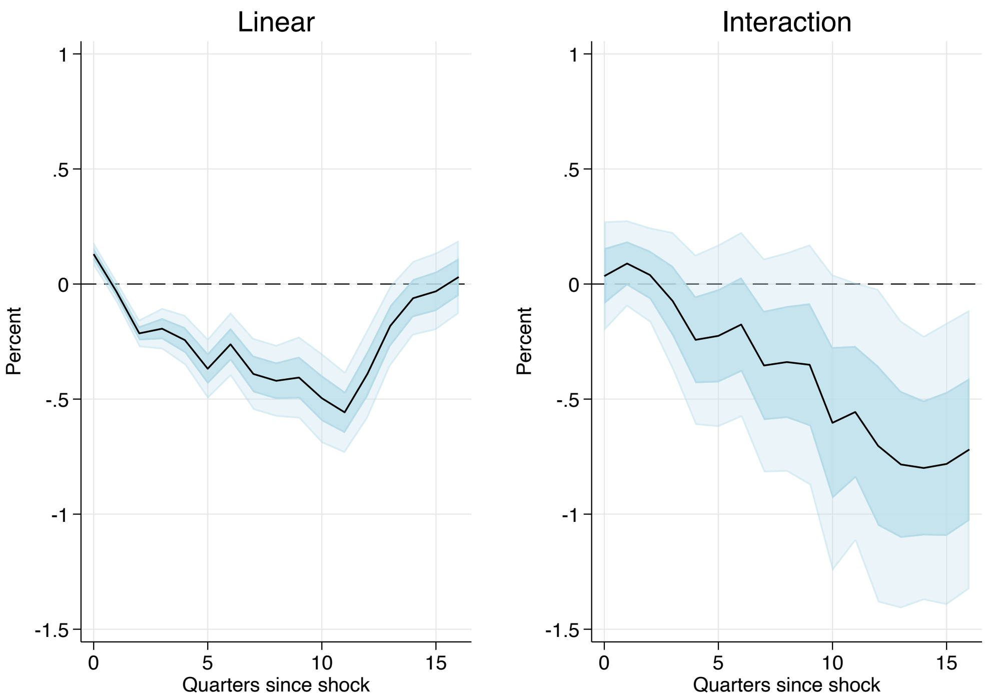

Notes: The left panel of the figure plots the response of the state-level log of real personal income to a percentage point contractionary monetary policy shock, as well as the 68% (dark shaded area) and 95% (light shaded area) confidence intervals. The horizontal axis is in quarters. The right panel reports the interaction coefficients between the monetary policy shock and the dummy identifying the top 20% of the services/manufacturing production ratio distribution.

Finally, in Figure 14 I repeat the same analysis but using the log of the real GDP at an annual frequency as the dependent variable. Real GDP decreases by around 1.2% after a monetary shock and even in this case, the states with a higher ratio of services to manufacturing GDP tend to react much more strongly.

These empirical findings confirm that economic structure plays an important role in the pass-through of monetary policy. Change in the population distribution is partially responsible for the shift toward the service sector. At the same time, the higher the relative share of services in the economy, the lower the frequency of price adjustment, and the stronger the response of output to shock. Armed with this empirical understanding, in the next section I develop a two-sector OLG-NK model to evaluate the impact of demographic trends on the transmission of monetary policy shocks and quantify the size of the new channel.

Figure 13: Impact of monetary policy on the regional annual inflation rate

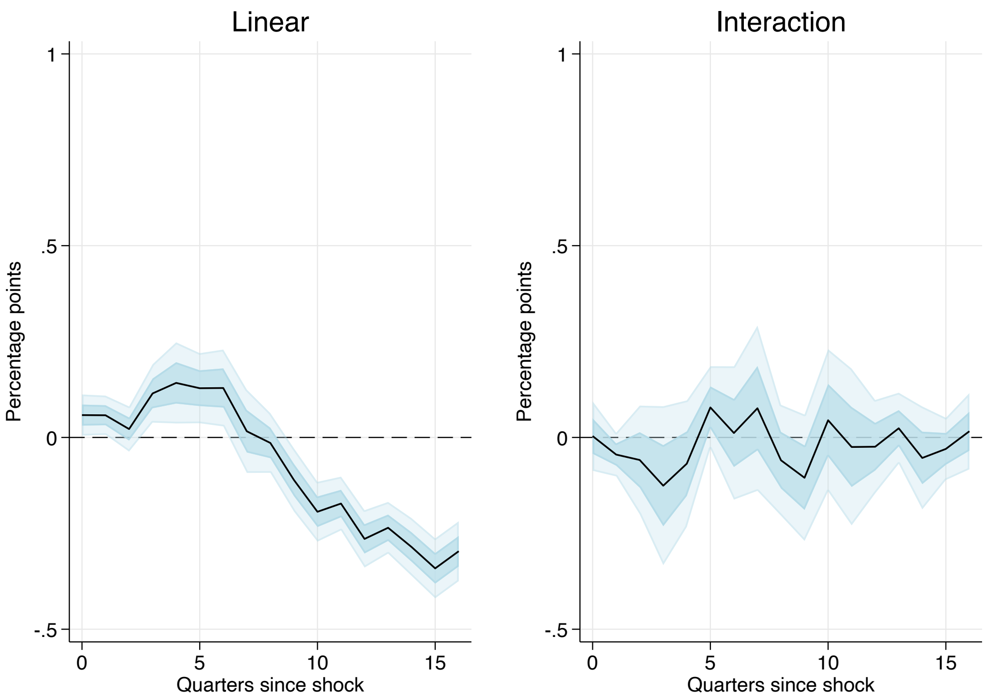

Notes: The left panel of the figure plots the response of the state-level annual inflation rate to a percentage point contractionary monetary policy shock, as well as the 68% (dark shaded area) and 95% (light shaded area) confidence intervals. The horizontal axis is in quarters. The right panel reports the interaction coefficients between the monetary policy shock and the dummy identifying the top 20% of the services/manufacturing production ratio distribution.

4.4 Robustness

To strengthen the validity of the results, I try a number of alternative specifications whose figures are reported in Appendix C. First, I repeat the same empirical analysis using different thresholds to distinguish between high and low service-intensive states. Second, as an alternative measure of monetary shocks, I also employ the high-frequency identification from Nakamura and Steinsson (2018). Third, for the dependent variable, I use the services component of the local GDP as a proxy for the non-tradable sector. The bottom line is that the basic pattern, in which the higher the share of services the stronger the effects of monetary policy, survives all of these modifications.

For the first robustness check, I consider different thresholds of the service-to-manufacturing GDP ratio distribution which I interact with the monetary shock. I consider a state high service intensive if its ratio belongs to the top quartile, one-third, and half of the distribution. I also interact the monetary policy shocks directly with the level of the ratio. The impulse

Figure 14: Impact of monetary policy on the regional real GDP

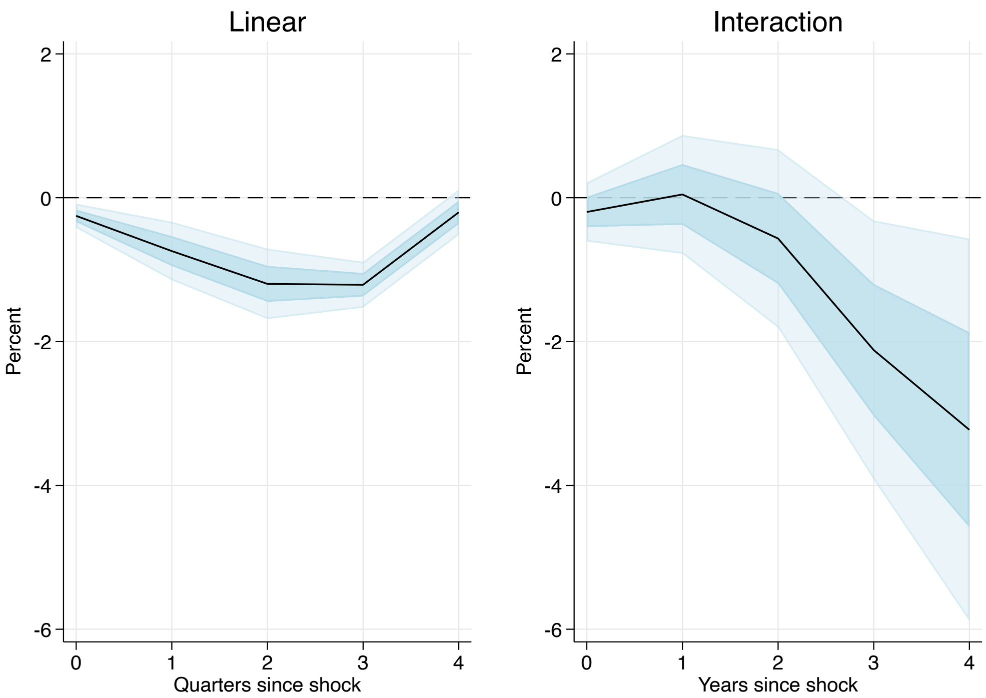

Notes: The left panel of the figure plots the response of the state-level log of the real GDP to a percentage point contractionary monetary policy shock, as well as the 68% (dark shaded area) and 95% (light shaded area) confidence intervals. The horizontal axis is in years. The right panel reports the interaction coefficients between the monetary policy shock and the dummy identifying the top 20% of the services/manufacturing production ratio distribution.

response functions are reported in Figure 26 and Figure 27 for the real personal income and inflation respectively. This alternative thresholds reinforce the conclusion that the effectiveness of monetary policy is influenced by the share of services in the economy.

One obvious question is whether the results are driven by the choice of monetary policy shocks. Therefore, as an additional estimation technique, I present the results using the high-frequency identification from Nakamura and Steinsson (2018) for the monetary shocks. The key idea of this approach is to use changes the change in the 3-month ahead Fed Funds futures within a 30-minute window surrounding scheduled Federal Reserve announcements. Since the time window is relatively small, one can consider these changes to be entirely due to the announcement itself and orthogonal to the information set of the financial market.

The results are presented in Figure 28 using as dependent variables the real personal income and annual inflation rate. All the regressions include the same controls as in the baseline specification. The responses of the interaction coefficients are comparable in shape and magnitude to the baseline specification being significantly stronger for states with a higher share of services.

Finally, spillover effects from other states might bias the results. It could be the case that the stronger response of personal income and GDP observed in more service-intensive states is actually due to an increase in the demand for tradable goods from the surrounding states rather than from the different frequencies of price adjustment across age groups. I test this hypothesis by using the services component of GDP as the dependent variable and as a proxy for the consumption of non-tradable goods: since services are usually not traded across states, differences in responses to shocks are mainly caused by local characteristics. The results are reported in Figure 29. The response of services in states with a higher service-to-manufacturing GDP ratio is significantly stronger suggesting that the main results are not driven by spillover effects.

In section B of the Online Appendix I try a number of alternative specifications which I here briefly summarize. First, I include extra state-level control in the regressions like measures of the housing market, firms' and workers' characteristics, and GDP per capita. Second, I repeat the same empirical analysis excluding the five smallest states by population, i.e., Alaska, North Dakota, Vermont, Washington D.C., and Wyoming, as well as Florida. Third, I investigate whether our results are sensitive to altering the beginning and the end of the sample as well as the number of lags. Fourth, whether spillover effects from other states might bias the results.

5 A Quantitative Life-Cycle Model

This section presents a two-sector overlapping generations (OLG) model for a closed economy with New Keynesian frictions in price settings that will be used to evaluate the impact of population aging in the U.S. on monetary shock propagation. The model presented here is an extension of the OLG models derived in Heer et al. (2017), Bielecki et al. (2020), and Bielecki et al. (2021) with one crucial modification: households of different ages have heterogeneous preferences over two sectors, services and goods, which differ in terms of the frequency of price adjustment.

5.1 Demographics

Households are born at age $j = 1$ (equivalent to real life age of 15), live for a maximum of $J = 85$ years (real-life age of 99), and survive each period with an age-specific probability s_j. The parameter $(1 - s_j)$ is then the age-specific mortality rate. The households work until they are $jw = 50$ years old (real-life age of 64) and then retire. I denote with N_j the size of cohort j relative to the overall population and so we have that $\sum_{j=1}^{J} N_j = 1$. As in Jaimovich et al. (2013) and Heer et al. (2017), the size of each age group is constant over time in order to match the empirical age-specific population shares with the model implied ones[17].

5.2 Households

The representative household of age j at time t maximizes its discounted lifetime utility (17) by choosing aggregate consumption $c_{t,j}$, the amount of hours to supply $l_{t,j}$ and the amount of assets to hold the sequent period $a_{t+1,j+1}$ subject to a budget constraint (14). The household receives a lump-sum transfer beq_t as well as an income $y_{t,j}$ composed of the net of tax labor-income $(1 - \tau_t)W_t l_{t,j} h_j$ if younger than jw years old, pension transfer from the government pen_t if older than jw years old. The transfers come from the unintentional bequests left by the households who die every period which is redistributed equally across all living agents. I express a variable in real terms by deflating it by the aggregate price index and define the relative price of the two sectors as:

$$Z_t = \frac{P_t^G}{P_t^S}. \tag{12}$$

The value function of the household of age j at time t can then be summarized as:

$$V_{t,j} = \max_{c_{t,j}, l_{t,j}, a_{t+1,j+1}} u\left(c_{t,j}, l_{t,j}\right) + \beta s_j \mathbb{E}_t V_{t+1,j+1}, \tag{13}$$

subject to the following constraints:

$$P_{t,j} c_{t,j} + P_t a_{t+1,j+1} = R_t^a P_{t-1} a_{t,j} + y_{t,j} \tag{14}$$

$$y_{t,j} = (1 - \tau_t) W_t l_{t,j} h_j \mathbf{I}_{j \le jw} + pen_t \mathbf{I}_{j > jw} + beq_t \tag{15}$$

[17] Households die every period at a rate $(1 - s_j)$ so the reader might think of an age-specific migration rate that keeps the size of each cohort constant. This assumption has a limited influence on the results since I will focus only on 3/4 years around the steady state and in such a short time span population distribution is basically constant.

$$a_{t,0} = 0 \qquad a_{t+J+1,J+1} = 0, \tag{16}$$

where R_t^a is the gross nominal rate on the real stock of assets that are managed by investment funds, W_t is the nominal wage per effective hour, h_j is the age-specific labor productivity rate, $\mathbf{I}$ is an indicator function to distinguish workers from retirees. Households are born and die without assets. Finally, the utility function takes the form:

$$u\left(c_{t,j}, l_{t,j}\right) = \Big(\frac{c_{t,j}^{1-\sigma}}{1-\sigma} - \nu \frac{l_{t,j}^{1+\eta}}{1+\eta}\Big). \tag{17}$$

The bundle of services and goods consumed by the household is given by:

$$c_{t,j} = \Big[\alpha_j^{\frac{1}{\eta}} (c_{t,j}^S)^{\frac{\eta-1}{\eta}} + (1-\alpha_j)^{\frac{1}{\eta}} (c_{t,j}^G)^{\frac{\eta-1}{\eta}}\Big]^{\frac{\eta}{\eta-1}}, \tag{18}$$

where the parameters $0 < \alpha_j < 1 \; \forall j$ capture the age-specific preferences over the services sector and will be used to match the expenditure shares observed in the data. η is the elasticity of substitution between services and goods. The price index associated with this bundle is:

$$P_{t,j} = \Big[\alpha_j^{\frac{1}{\eta}} (P_t^S)^{\frac{\eta-1}{\eta}} + (1-\alpha_j)^{\frac{1}{\eta}} (P_t^G)^{\frac{\eta-1}{\eta}}\Big]^{\frac{\eta}{\eta-1}}. \tag{19}$$

5.3 Firms

On the firms' side, there are two sectors: one that produces services and one goods. The main differences between the two sectors stem from the fact that only the output of the goods sector can be used for capital investment and they differ in their frequency of price adjustment. In line with the empirical evidence, a lower share of firms in the services sector is able to adjust prices each period. As in standard New Keynesian models, the production side in each sector is split into a competitive final goods firm and a continuum of intermediate goods firms.

Final firms. For each sector, $s \in \{S, G\}$ the final good is produced under perfect competition using a continuum of intermediate goods indexed by i with a constant-returns-to-scale technology. The final firms are price-takers and they solve the profit-maximization problem:

$$\max_{Y_{i,t}^s} P_t^s Y_t^s - \int_0^1 P_{i,t}^s Y_{i,t}^s \, dj, \tag{20}$$

subject to the CES production function where the parameter ϵ denotes the elasticity of substitution across different varieties of intermediate goods:

$$Y_t^s = \Big(\int_0^1 (Y_{i,t}^s)^{\frac{\epsilon-1}{\epsilon}} \, di \Big)^{\frac{\epsilon}{\epsilon-1}}. \tag{21}$$

The solution to the maximization problem gives the standard demand function for variety i for the production of final good s:

$$Y_{i,t}^s = \Big(\frac{P_{i,t}^s}{P_t^s} \Big)^{-\epsilon} Y_t^s. \tag{22}$$

Intermediate firms. The optimization problem of the monopolistically competitive intermediate good producer i is divided into two stages. In the first stage, for a given production function $Y_{i,t}^s$, the intermediate firm chooses the amount of inputs $L_{i,t}^s$ and $K_{i,t}^s$, taking nominal prices as given, such that costs are minimized:

$$\min_{L_{i,t}^s, K_{i,t}^s} W_t L_{i,t}^s + R_t^k K_{i,t}^s \tag{23}$$

$$s.t. \;\; Y_{i,t}^s = (K_{i,t}^s)^{\psi} (L_{i,t}^s)^{1-\psi},$$

where ψ is the capital share in the production function and R_t^k is the nominal rental rate on capital.

In the second stage, $Y_{i,t}^s$ and $P_{i,t}^s$ are determined such that the discounted real profits are maximized subject to the demand function of the final output producer. However, firms are not free to adjust their prices as they want since they face a Calvo staggered price setting mechanism: in each period, a fraction θ^S of services intermediate goods producers and a fraction θ^G of manufacturing intermediate goods producers cannot reset their prices and maintain those of the previous period. The Calvo friction parameters are constant over time and differ across sectors to match the empirical estimates on the lower frequency of price adjustment in the services sector relative to the goods sector, that is $\theta^S > \theta^G$.

The fact that a firm in sector s might not be able to adjust its price in period t with probability θ^s makes the pricing problem dynamic equal to solving:

$$\max_{P_{i,t}^s} \mathbb{E}_0 \sum_{t=0}^{\infty} \Big(\prod_{r=0}^{t} R_r^{-1} \Big) (\theta^s)^r \Big[(P_{i,t}^s - MC_{t+r}^s) \Big(\frac{P_{i,t}^s}{P_{t+r}^s} \Big)^{-\epsilon} Y_{t+r}^s \Big], \tag{24}$$

where MC_t^s is the nominal marginal cost in sector s. Since intermediate goods producers are risk-neutral they use the nominal risk-free rate to discount expected future profit flows.

5.4 Investment funds

As in Bielecki et al. (2021), the households' savings are managed by perfectly competitive and risk-neutral investment funds which transfer the earned gross return back to households every period. The portfolio managed by the investment funds consists of physical capital K_t, bonds B_t, and claims on intermediate goods-producing firms (shares) $D_{i,t}$. A representative investment fund maximizes the expected present value of future gross returns:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \Big(\prod_{r=0}^{t} R_r^{-1} \Big) \Big[[R_{t+1}^k + (1-\delta) Q_{t+1}] K_{t+1} + R_t P_t B_{t+1} + \int_0^1 [P_{t+1} F_{i,t+1} + P_{i,t+1}^d] D_{i,t+1} \, di \Big], \tag{25}$$

where δ is the depreciation rate of capital, R_t denotes the gross nominal risk-free rate, Q_{t+1} is the nominal price of a unit of capital, and $D_{i,t}$ refers to the number of shares issued by intermediate goods producing firm i which are traded at the end of period t at price $P_{i,t}^d$ and yield real dividends $F_{i,t}$. The nominal balance sheet of investment funds at the end of period t can be written as:

$$P_t A_{t+1} = Q_t (1-\delta) K_t + P_t I_t + P_t B_{t+1} + \int_0^1 P_{i,t}^d D_{i,t+1} \, di. \tag{26}$$

I_t denotes investment in physical capital which accumulates according to:

$$K_{t+1} = (1-\delta) K_t + \Big[1 - S_k \Big(\frac{I_t}{I_{t-1}} \Big) \Big] I_t, \tag{27}$$

where $S_k()$ captures investment adjustment costs which have the following functional form:

$$S_k \Big(\frac{I_t}{I_{t-1}} \Big) = \frac{S_1}{2} \Big(1 - \frac{I_t}{I_{t-1}} \Big)^2. \tag{28}$$

Finally, since I assume that all revenues are transferred back to households, the ex-post rate of return on assets R_t^a is implicitly given by:

$$R_t^a P_{t-1} A_t = [R_t^k + (1-\delta) Q_t] K_t + R_{t-1} P_{t-1} B_t + \int_0^1 [P_t F_{i,t} + P_{i,t}^d] D_{i,t} \, di. \tag{29}$$

5.5 Government

The government funds a pay-as-you-go social security system. The amount of pension benefits pen_t received by households with age above jw is given by the replacement rate $\bar{d}$ and the average net labor income $(1-\tau_t)W_t\bar{h}$. The tax rate on labor income τ_t is set such that the budget is balanced in each period:

$$pen_t = \bar{d}(1-\tau_t)W_t\bar{h} \tag{30}$$

$$\tau_t W_t \sum_{j=1}^{jw} N_j l_{t,j} h_j = pen_t \sum_{j=jw+1}^{J} N_{t,j}, \tag{31}$$

where $\bar{h} = \frac{\sum_{j=1}^{jw} h_j}{jw}$ is the average efficiency-hours worked in the working life-periods.

5.6 Monetary authority

The central bank follows the following simple Taylor-type rule:

$$\frac{R_t}{R} = \left(\frac{\pi_t}{\pi}\right)^{\phi_\pi} \left(\frac{Y_t}{Y}\right)^{\phi_y} e^{\nu_t}, \tag{32}$$

where R_t is the gross nominal interest rate, $\pi_t = \frac{P_t}{P_{t-1}}$ is the gross rate of aggregate inflation, Y_t is the aggregate output and R, π, and Y are the steady state values of the respective variable. ϕ_π and ϕ_y measure the elasticity at which the monetary authority adjusts the interest rate to changes in the current inflation rate and output and ν_t is a monetary shock following an $AR(1)$ process with persistence ρ.

Aggregate output is defined as:

$$P_t Y_t = P_t^S Y_t^S + P_t^G Y_t^G \tag{33}$$

and aggregate price level as $P_t = \left[\omega_t^{\frac{1}{\eta}} (P_t^S)^{1-\eta} + (1-\omega_t)^{\frac{1}{\eta}} (P_t^G)^{1-\eta}\right]^{\frac{1}{1-\eta}}$ where $\omega_t = \sum_j \alpha_j \chi_{t,j} \frac{P_{t,j}^{\eta-1}}{\sum_j \chi_{t,j} P_{t,j}^{\eta-1}}$ and $\chi_{t,j}$ is the share of household j expenditure in aggregate expenditures at time t. See Appendix A for the full list of model equations.

5.7 Market clearing

The market for final output in both sectors needs to clear. Only the output of the goods sector can be stored into the next period and used for capital investment while the output of the services sector needs to be consumed every period. Hence:

$$Y_t^S = C_t^S \tag{34}$$

$$Y_t^G = C_t^G + K_{t+1} - (1-\delta)K_t. \tag{35}$$

Moreover, both the labor and the capital market also need to clear:

$$L_t = L_t^S + L_t^G = \sum_{j=1}^{J} N_{t,j} l_{t,j} h_j \tag{36}$$

$$K_t = K_t^S + K_t^G = \sum_{j=1}^{J} N_{t,j} a_{t+1,j+1}. \tag{37}$$

Since bonds are traded only between (identical) investment funds they are in zero net supply, $B_t = 0$. Finally, the lump-sum transfer beq_t from the unintentional bequests is equal to:

$$beq_t = \sum_{j=1}^{J} (N_{j-1} - N_j) \frac{R_t^a}{\pi_t} a_{t,j}. \tag{38}$$

5.8 Quantitative analysis

I am interested in studying how demographic trends affect monetary policy effectiveness. Therefore, I use the model to study the transmission of monetary policy shocks around three steady states that differ only in terms of population distribution N_j, mortality rate $(1 - s_j)$, and service preferences α_j. All other parameters are fixed. I choose 1980 as the first steady state and baseline since that is when CEX data, necessary to compute the sectoral preferences across age groups, becomes available. The second steady state is 2010 and the final steady state is set at 2050 using the World Bank population projection for the U.S.

5.8.1 Calibration

The model parameters are set in two ways: externally set with the values in the literature and internally set to target data moments.

The externally set parameters are reported in Table 3. As previously mentioned, households live for a maximum of 85 ($J = 85$) years and then die with certainty. They work until they are $jw = 50$ years old (64 years old in real life) and then they retire. The elasticity of intertemporal substitution σ, the disutility of labor supply ϕ, and the inverse of the Frisch elasticity ν are set to their standard values of 1, 4, and 2 respectively. The elasticity of substitution between the

Table 3: Externally set parameters

Parameter	Value	Description
J	85	Terminal life-age (99). Death with certainty at age 100
jw	50	Terminal working-age (64)
σ	1	Elasticity of intertemporal substitution
ϕ	4	Disutility of labor supply
ν	2	Inverse of the Frisch elasticity of labor supply
η	0.4	Elasticity of substitution between services and goods from Galesi and Rachedi (2018)
ψ	0.33	Cobb-Douglas capital elasticity of output
S_1	4.39	Investment adjustment cost curvature from Bielecki et al. (2021)
$\bar{d}$	0.33	Pension replacement rate. Source: Bárány et al. (2022)
ϕ_π	1.5	Inflation coefficient in the Taylor rule
ϕ_y	0.2	Output coefficient in the Taylor rule
ρ	0.8	Monetary shock persistence
σ_{ϵ^r}	1	Std. Dev. of Monetary shock

Notes: The table reports the externally set parameters of the model.

two sectors η, which captures how easy it is for the household to switch goods and services, is from Galesi and Rachedi (2018) and set to 0.4. The investment adjustment cost curvature S_1 equals 4.39 as in Bielecki et al. (2021). The pension replacement rate $\bar{d}$ is taken from Bárány et al. (2022) and the Taylor rule coefficients are set to the standard values in the literature. Finally, the Calvo parameters for the services sector θ^S and the goods sector θ^G are set to 0.75 and 0.25 respectively as in Galesi and Rachedi (2018) in order to match the mean implied duration in months estimated by Nakamura and Steinsson (2008).

Table 4: Calibrated parameters

Parameter	Value	Description	Target
β	0.999	Discount factor	Annual interest rate between 4 and 5 %
δ	0.02	Depreciation rate	Capital-output ratio between 2 and 2.7
N_j	Panel B of Figure 15	Population shares. *Source*: UN (2017) World Population Prospects	Realised and forecasted population shares
$(1-s_j)$	Panel C of Figure 15	Survival probability. *Source*: Social Security Administration	Realised and forecasted mortality rates
α_j	Panel D of Figure 15	Share of consumption devoted to services	Age-group service preferences from CEX
h_j	Panel A of Figure 15	Age-group specific labor productivity from Fullerton (1999)	Wage profile
ϵ	6	Elasticity of demand for each intermediate good	Steady state markup of 20%
θ^S	0.75	Calvo Frequency Services. *Source*: Nakamura and Steinsson (2008)	Price adjustment every 13 months
θ^G	0.25	Calvo Frequency Goods. *Source*: Nakamura and Steinsson (2008)	Price adjustment every 3 months

Notes: The table reports the internally calibrated parameters of the model.

The internally calibrated parameters are reported in Table 4. The discount factor β and the depreciation rate δ are set to 0.999 and 0.02 respectively in order to match the annual interest rate and the capital-output ratio estimated in the early 80s. The elasticity of demand for each intermediate good ϵ is set to 6 such that the steady-state markup is equal to 20%. The age-group-specific labor productivity parameters h_j, shown in Panel A of Figure 15, are taken from Fullerton (1999) in order to match the hump-shaped distribution of labor income over the life cycle.

Figure 15: Age specific parameters

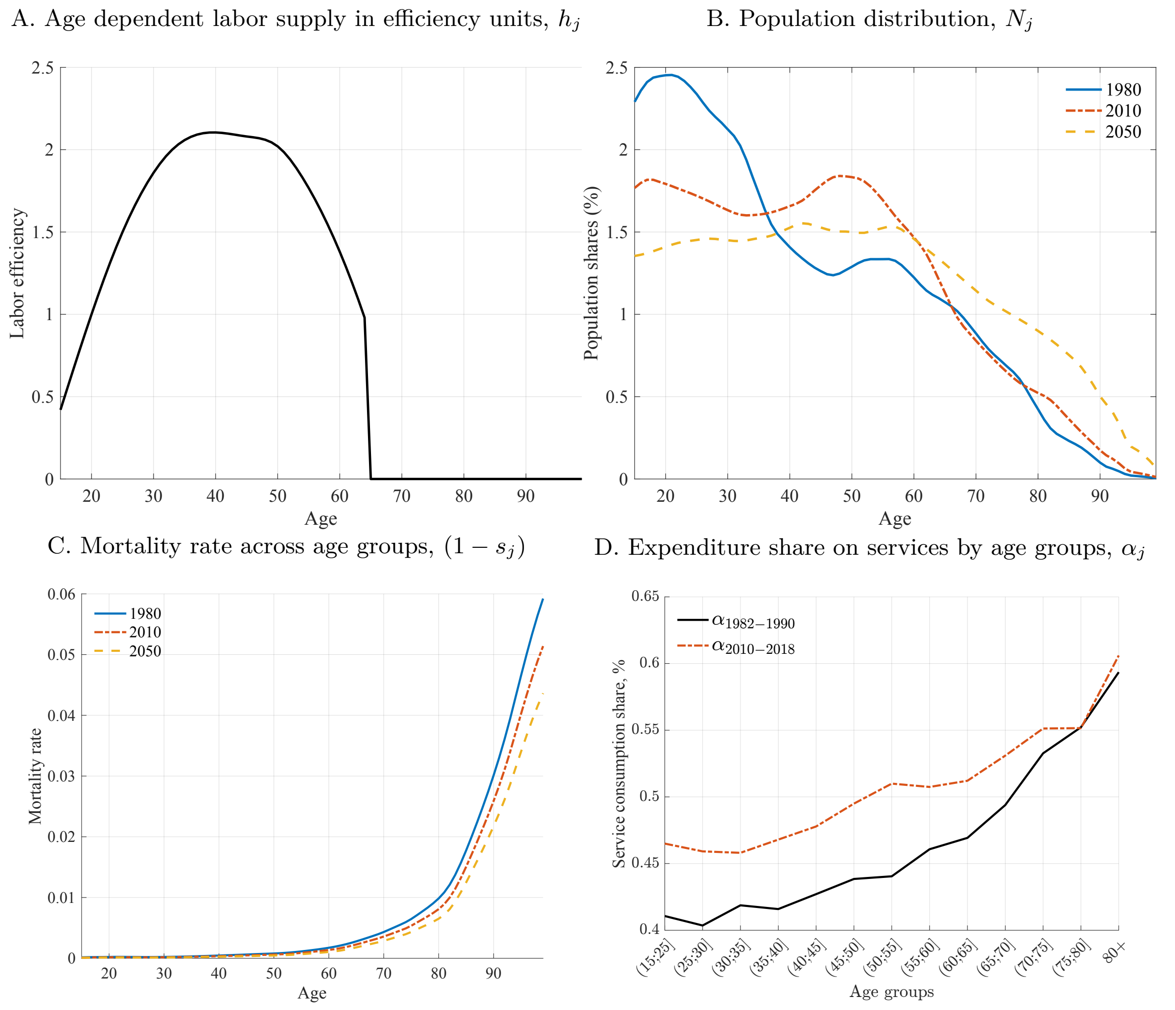

Notes: Panel A: The profile of the age-specific labor productivity is obtained by interpolating the estimates from Fullerton (1999). Panel B: The plot shows the population share distribution across age groups for 1980, 2010, and the forecasted values for 2050. *Source*: UN (2017) World Population Prospects. Panel C: The plot displays the age-group quarterly mortality rates in 1980, 2010, and the forecasted values for 2050. *Source*: Table 7 from the Cohort Life Tables for the Social Security Area. Panel D: The plot displays the average age group level expenditure shares on services across age groups over two different periods. *Source*: CEX.

The most important parameters for the analysis are the shares of each age group N_j, the mortality rates $(1 - s_j)$, and the shares of consumption devoted to services α_j. The U.S. population distributions for the years 1980, 2010, and 2050, reported in Panel B of Figure 15, are retrieved from the UN (2017) World Population Prospects. As one can notice, demographic trends are a complex phenomenon that cannot be entirely captured by simply considering the effects on workers and retirees. On the one hand, the share of people below 35 years old is decreasing over time whereas the share of people above 65 years old is increasing. On the other hand, the share of highly productive workers (households between 35 and 65 years old) has actually increased relative to 1980. These shifts in labor force participation might have conflicting predictions regarding the effectiveness of monetary policy if not properly included in the model.

Much more straightforward is the analysis of the changes in the U.S. mortality rates $(1 - s_j)$ reported in Panel C of Figure 15. For all the age groups considered, the survival probability has increased from 1980 to 2010 and it is expected to increase even further in 2050.

Panel D of Figure 15 shows the share of consumption α_j that each age group devotes to services. The services shares are computed from the CEX data which is available since the early 80s. Since there are no predictions regarding the state of these shares in 2050 when I evaluate how changes in preferences influence the pass-through of monetary policy I focus only on the 1980 and 2010 steady states. The share of services increases over the life cycle in line with previous findings. Since the early 80s, each age group has increased its consumption of services mainly because of income and price effects as shown in Cravino et al. (2022).

I assess the quality of the calibration of the lifecycle parameters by comparing some untargeted moments with the data. In particular, Figure 30 plots the age profile of assets implied by the model (normalized to asset holdings at age 65) with the age profile observed over different years in the Survey of Consumer Finances (normalized for the group 65-54). The model performs quite well in replicating the hump-shaped lifecycle asset profile which peaks around 60 years old. In line with the data, individuals borrow when young and dissave after they retire.

The theoretical model developed is used to answer several questions concerning the relationship between demographic trends and monetary policy effectiveness. First, I evaluate whether population aging has influenced the responsiveness of output and inflation to monetary shocks. Second, I exploit the rich demographic structure of the model to study the consumption of which age groups is the most sensitive to changes in the interest rate. Third, I show that demographic trends play a significant role in explaining the variation of the pass-through of monetary policy even when compared to other channels like the increase in the share of services consumed due to price and income effects. Fourth, I quantify the importance of the novel channel I propose, i.e., consumption heterogeneity. Fifth, I document that demographic trends partially contributed to the flattening of the Phillips curve.

5.8.2 Demographic trends and the effectiveness of monetary policy

In this section, I evaluate how demographic trends influence the way monetary policy shocks propagate in the U.S. Figure 16 reports the IRFs to an expansionary monetary shock of the main variables in the model computed using the demographic structure in 1980. The shapes and the magnitudes are in line with the literature. Following a 100 basis points expansionary monetary policy shock, i.e., an exogenous decrease in the interest rate, output, inflation, consumption, and investment increase. The central bank then responds by increasing the real interest rate to slow down economic growth until the economy returns to the initial steady state.

Figure 16: Model impulse response functions

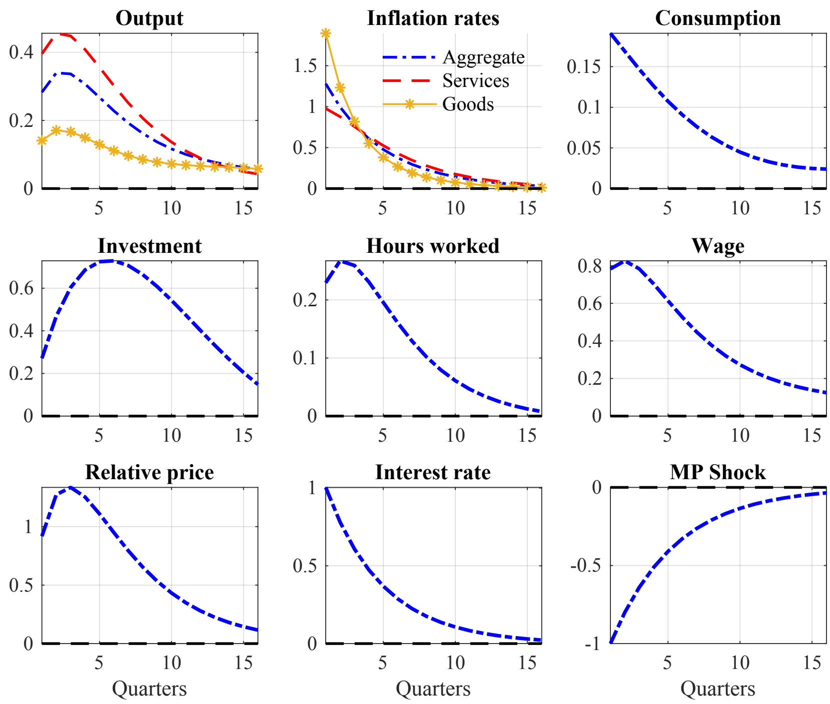

Notes: The plot reports the IRFs of several variables of interest computed using 1980 as steady states.

Of particular interest are the responses in the two top left panels. Given the different price stickiness parameters between the two sectors, the price response in the services sector is more muted relative to the response in the goods sector. Since firms in the services sector cannot adjust their prices as frequently, they respond to the shock by adjusting their production more vigorously leading to a stronger and less persistent response of the output in the services sector relative to the response in the goods sector. As previously underlined, the sensitivity of output and inflation to different price stickiness is not symmetric. Output in the services sector is significantly more responsive to shocks whereas the response of inflation in the same sector is only marginally more muted.

I now focus on the influence that demographic trends have had on monetary shock propagation in the U.S. over the last decades and the influence that they will have in the next 30 years. To evaluate this relationship, I compute the response of output to an expansionary monetary shock using the population distribution N_j and the mortality rates $(1 - s_j)$ in 1980, 2010, and 2050. All the other parameters, including the services shares α_j, are kept fixed and set to their 1980 values.

Figure 17: Model IRFs of output for different demographic structures

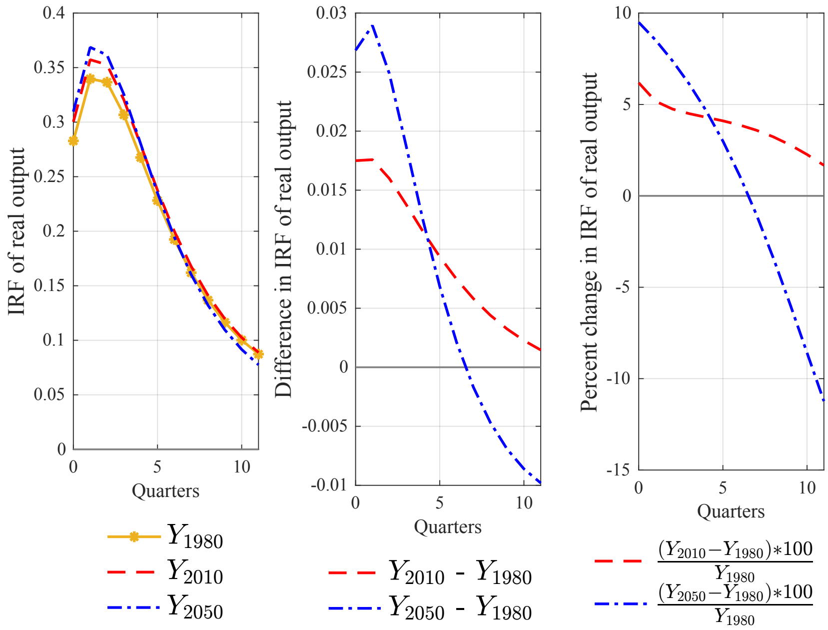

Notes: The left panel of the plot reports the IRFs of output across the three different steady states changing only the population distribution and mortality rate and keeping service preferences at the 1980 values. The middle panel shows the first differences between these IRFs, i.e., the difference between the IRF of output in 2050 and 2020 with the respect to the baseline IRF in 1980, whereas the right panel reports the percentage change in IRFs across the different steady states.

The responses are plotted in the left panel of Figure 17. Moving from 1980 to 2010 and then to 2050 results in a stronger response of output to the shock. On top of that, the responses have become less persistent over time. In the middle panel, I report the differences in output responses with respect to the baseline of 1980. By increasing the share of old people who have a higher preference for the services sector, the demographic structures of 2010 and 2050 increase the response of output by 1.8 and 2.6 percentage points respectively relative to that of 1980. The right panel shows the same results in percent deviation: simply changing the population distribution and the mortality rate over time makes the response of output 6% stronger in 2010 relative to 1980 and 10% stronger in 2050.

Figure 18: Model IRFs of inflation for different demographic structures

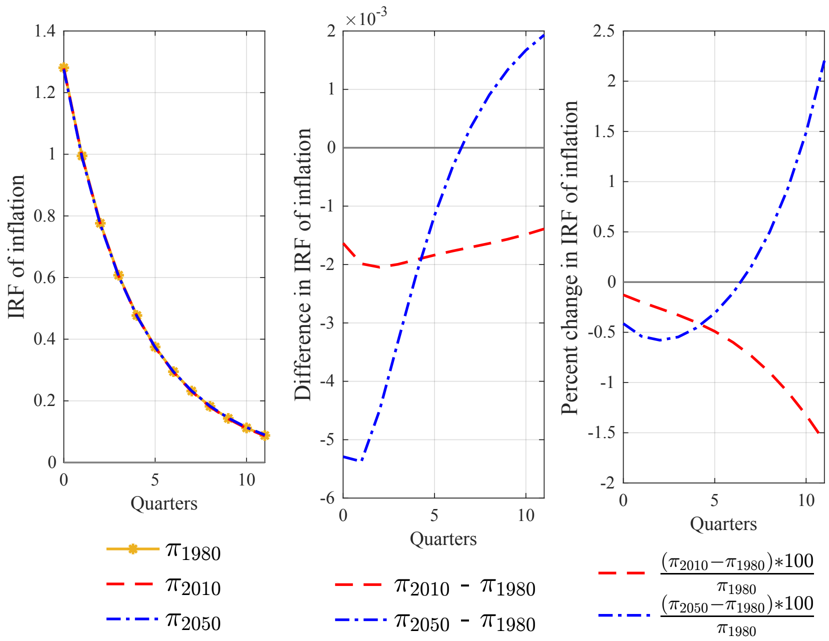

Notes: The left panel of the plot reports the IRFs of inflation across the three different steady states changing only the population distribution and mortality rate and keeping service preferences at the 1980 values. The middle panel shows the first differences between these IRFs, i.e., the difference between the IRF of inflation in 2050 and 2020 with the respect to the baseline IRF in 1980, whereas the right panel reports the percentage change in IRFs across the different steady states.

Figure 18 reports the same analysis for the responses of the aggregate inflation rate. In line with the empirical evidence found in Section 4, demographic trends have a negligible impact on the IRFs of inflation: the demographic structures of 2010 and 2050 relative to that in 1980 result in more muted responses of inflation (so the differences $\pi_{2010} - \pi_{1980}$ and $\pi_{2050} - \pi_{1980}$ are negative) but the overall decrease is less than 1% for both steady states.

5.8.3 Heterogeneous consumption responses by age

The shift in aggregate demand towards services caused by demographic trends leads to a stronger response of output following an expansionary monetary shock. However, these changes are unlikely to be homogeneous across age groups. I assess which age groups are more exposed to the structural transformation induced by population aging.

Figure 19: Heterogeneous consumption responses to expansionary monetary policy shocks, by age

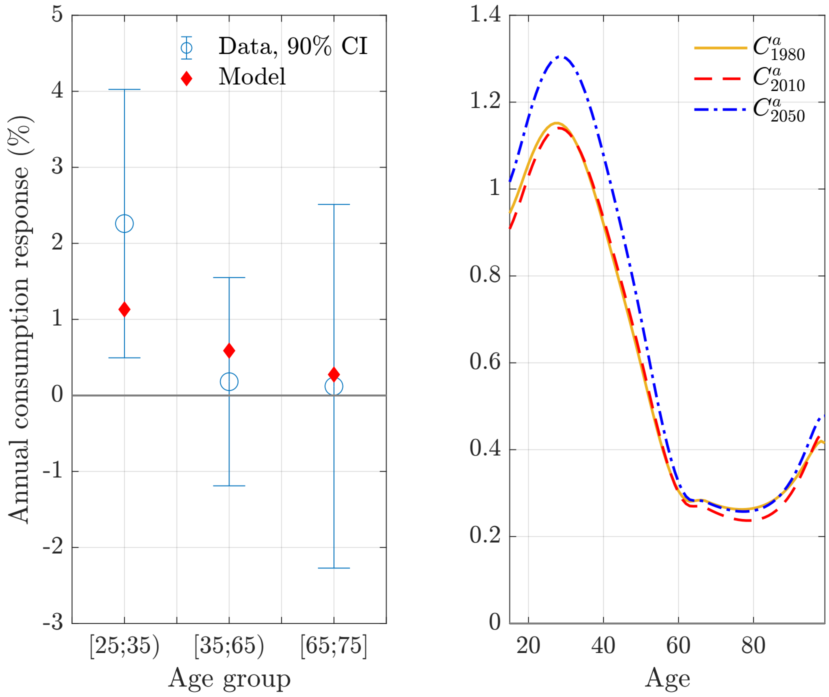

Notes: The left panel compares the annual percent change of consumption following a contractionary shock from the model (red diamond) with the empirical estimates from Wong (2021) (blue error plot with 90% confidence bands) for three age groups. The right panel reports the model-implied annual percent change of consumption across age groups following an expansionary shock using the demographic structure of 1980, 2010, and 2050.

I start by comparing the model-implied consumption responses with the empirical estimates from the literature. The left panel of Figure 19 reports the annual percent change of consumption, i.e., the sum of the responses of the first four quarters, from the model with those estimated by Wong (2021) for three age groups. The model is able to capture quite well the negative relationship between age and consumption responsiveness and the predicted responses fall within the 90% confidence bands of the empirical estimates. The consumption of young households is the most sensitive to changes in the interest rate. This is due to the fact that they just entered the labor market and they have fewer assets so whenever there is a shock to the economy they are not able to smooth consumption over time as much as older households.

The left panel of Figure 19 shows the model-implied annual percent change of consumption for each age group C^a following an expansionary monetary shock using the demographic structure of 1980, 2010, and 2050. The relationship between age and consumption responsiveness is not linear. In particular, it increases until households are 30 years old and then drastically decreases. After they turn 60, the relationship becomes rather stable with a slight increase towards the end. The nonlinearity in the relationship is due to the hump-shaped distribution of assets and labor productivity reported in Figure 30 and Panel A of Figure 15 respectively.

The change in population distribution from 1980 to 2010 has a negligible and rather homogeneous effect on consumption responsiveness across age groups. However, in 2050 demographic trends will have an extremely heterogeneous impact across age groups and the consumption of younger households will be the one most affected by demographic trends. For the age group between 25 and 35 years, old consumption will respond 15% more in 2050 than in 1980. The consumption responses of older people are basically unaffected.

The negative relationship between consumption responsiveness and age might erroneously lead to the conclusion that demographic trends will make output *less* responsive to monetary policy shocks. Indeed, ceteris paribus, by increasing the share of the less responsive members of the society, i.e., older people, demographic trends should reduce the aggregate consumption sensitivity resulting in a more muted response of output to shocks. However, as shown in Figure 19, the general equilibrium effects caused by changes in population distribution are such that the sensitivity of the young households' consumption increases even further completely offsetting the partial equilibrium effects and overall increasing output responsiveness.

5.8.4 Demographic trends vs. price and income effects

As demonstrated in Cravino et al. (2022), population aging accounts for around a fifth of the overall rise in the share of services whereas the real income growth and changes in relative prices explain another three fifth. To quantify the importance for monetary policy propagation of demographic trends relative to other channels, I compare the variation in output and inflation responsiveness under three different scenarios. In the first scenario, I isolate the demographic component by computing the percent change in the IRFs of output and inflation from 1980 to 2010 by adjusting the population distribution and mortality rates but keeping the service preferences constant as in Figures 17 and 18. The results are reported in the blue bars of Figure 20. The responses of output are shown on the left panel and the responses of inflation are on the right panel. In the second scenario, I isolate the importance of all the other channels that lead to an increase in services, e.g., price and income effects, excluding population aging by varying the service preferences from 1980 to 2010 but keeping the demographic structure of 1980 (red bars). In the third scenario, I adjust both the demographic structures and service preferences to the two steady states (black line).

Figure 20: Model IRFs under different scenarios

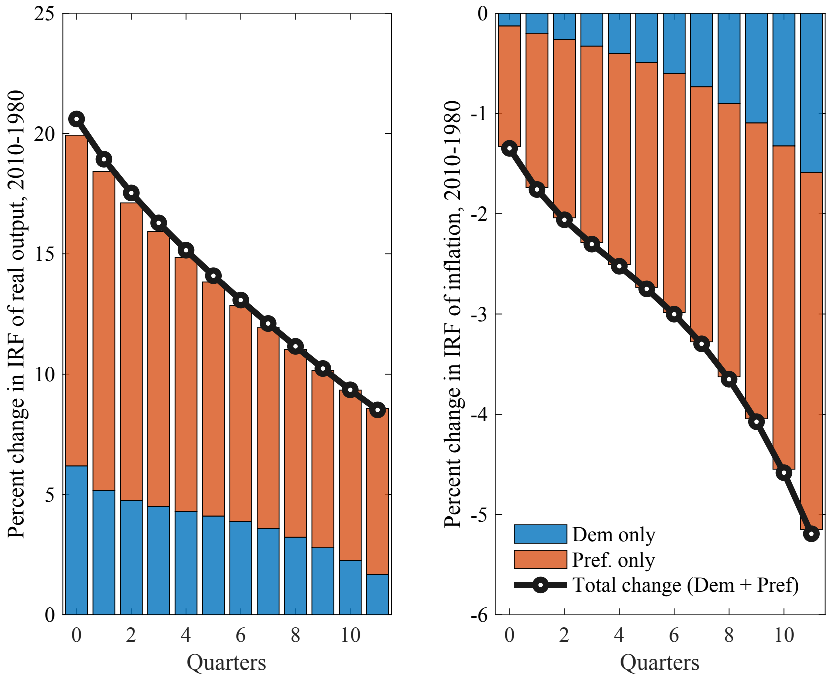

Notes: The left panel of the plot shows the percent change in impulse responses for output from 1980 to 2010 under three different scenarios: using the population distribution and mortality rates of 1980 and 2010 but services preferences kept fixed at the 1980 values (blue bars, same plot as before), using the service's preferences of 1980 and 2010 but the demographic structure of 1980 (red bars) and finally using both the demographic structures and services preferences of the two steady states (black line). The right panel shows the same percent change but for inflation.

The response of output in 2010 is 20% stronger than in 1980 when both the demographic structure and the service preferences are changed and a significant share of this increase is explained by population aging alone. The ratio between the blue bars and the black line in the left panel is approximately 30% suggesting that, even though other structural changes like income and price effects are important drivers of the change in services share, demographic trends account for a sizable extent of the overall effect.

The right panel of Figure 20 delivers a similar story for inflation. The overall percent change in IRFs is between 1.5% and 4.5% more muted in 2010 relative to 1980 and the share explained by demographic trends is between 10% and 25%

5.8.5 The importance of consumption heterogeneity

The variation in monetary policy effectiveness caused by changes in the demographic structure so far documented is the result of the interaction of several channels. The shift in population distributions across the three steady states results in different labor market participation, asset distribution, etc. To quantify the importance of the new proposed channel, i.e., consumption heterogeneity across age groups, I compare the changes in output and inflation responsiveness shown in Figure 17 and Figure 18 with a counterfactual scenario. In the baseline specification, I compute the responses to a monetary policy shock using the three different steady-state values for the demographic structure and keeping everything else fixed including the service preferences. In the counterfactual scenario, I still change the demographic parameters from 1980 to 2050 but I assume that the share of consumption devoted to services α_j is constant across age groups and equal to the weighted mean value of 1980.

Figure 21: Model IRFs of output between the baseline and the contrafactual scenario

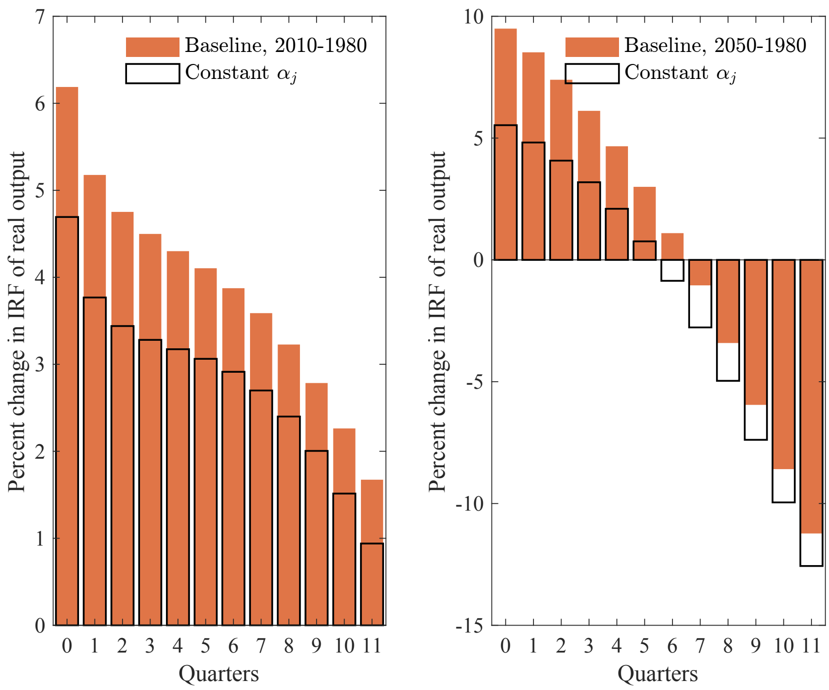

Notes: The plot compares the percent changes of output for 2010 relative to 1980 (left panel) and for 2050 relative to 1980 (right panel) for the baseline and a contrafactual scenario in which all age groups have the same sectoral preferences.

Figure 21 reports the percent changes of output under the baseline and the contrafactual scenario for 2010 relative to 1980 (left panel) and for 2050 relative to 1980 (right panel). Neglecting consumption preferences heterogeneity across age groups leads to a sizable underestimation of the effect of demographic trends on monetary policy: the percent change of the response of output on impact drops from 6.3% to 4.8% in 2010 and from 9.9% to 5.3% in 2050.

It is important to notice that demographic trends still lead to an increase in the overall effectiveness of monetary policy. This is mainly due to the other demographic channels included in the model. For instance, the aggregate share of workers decreases over time so the firms need to adjust the wage level more vigorously to shocks in order to increase the supply of hours of labor. Moreover, asset distribution becomes more negatively skewed over time because of the increase in the share of older households. As shown in the right panel of Figure 19, fewer assets are then owned by younger households who become even more financially constrained and, therefore, more sensitive to changes in the interest rate.

Figure 22: Model IRFs of inflation between the baseline and the contrafactual scenario

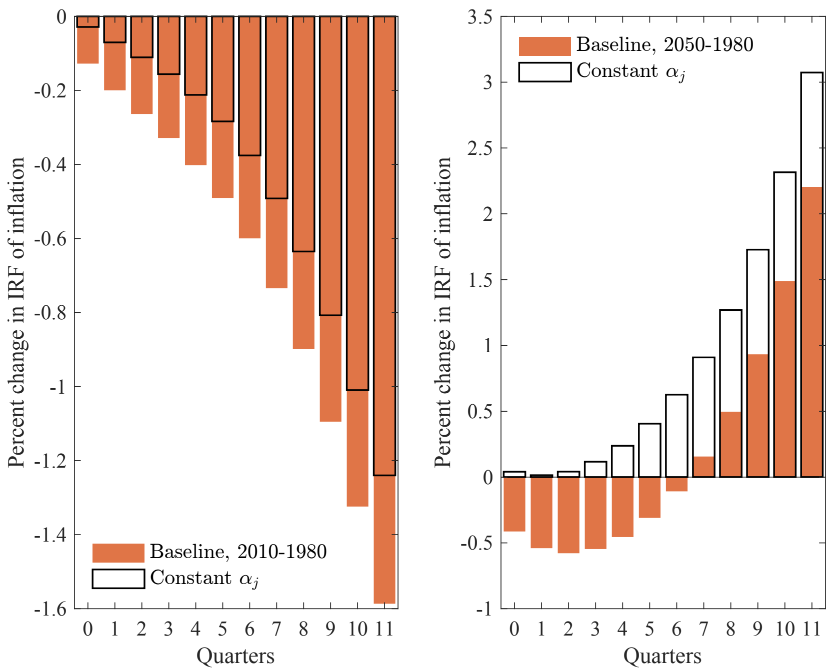

Notes: The plot compares the percent changes of inflation for 2010 relative to 1980 (left panel) and for 2050 relative to 1980 (right panel) for the baseline and a contrafactual scenario in which all age groups have the same sectoral preferences.

The same exercise is repeated for inflation and reported in Figure 22. A symmetrical effect is found here: neglecting preference heterogeneity results in an overestimation of the impact of population aging on the response of inflation. The effect is such that the percent change is smaller from 1980 to 2010 than in the baseline and becomes even positive for 2050.

5.8.6 Phillips curve

The slope of the Phillips curve, which captures the strength of the relationship between inflation and economic activity, has been found to decrease over time. This so-called "flattening" of the Phillips curve has crucial implications for policymakers and central bankers. A lower sensitivity of inflation to real activity implies that to stabilize inflation, larger movements in economic activity are needed, which in turn require larger shifts in the interest rate. This is of particular importance in times when the interest rate is close to zero.

Several explanations have been proposed to justify this phenomenon. The potential causes include the success of monetary policy in anchoring expectations (Bernanke, 2010), the increase in central bank credibility and transparency (McLeay and Tenreyro, 2019), or global forces (Jorda et al., 2019). In this paper, I argue that part of the flattening of the Phillips curve is due to the increase in the consumption share devoted to services and, therefore, to demographic trends that shift demand towards this stickier category.

The New Keynesian Phillips curve for sector s can be derived by linearizing equation (24) around a steady state with zero inflation in both sectors. Applying the canonical derivations leads to the following sectoral Phillips curves:

$$\hat{\pi}_t^S = \beta \mathbb{E}_t \hat{\pi}_{t+1}^S + \kappa^S \hat{mc}_t^S \tag{39}$$

$$\hat{\pi}_t^G = \beta \mathbb{E}_t \hat{\pi}_{t+1}^G + \kappa^G \hat{mc}_t^G, \tag{40}$$

with

$$\kappa^S = \frac{(1-\theta^S)(1-\theta^S\beta)}{\theta^S}, \qquad \kappa^G = \frac{(1-\theta^G)(1-\theta^G\beta)}{\theta^G}. \tag{41}$$

Inflation in sector $s \in \{S, G\}$ is a function of the next period expected sectoral inflation discounted by β and the sectoral marginal cost $\hat{mc}_t^s$ times the slope of the Phillips curve κ^s. Notice that since $\theta^S > \theta^G$, that is, the share of firms that cannot reset their price every period is higher in the services sector, it follows that $\kappa^S < \kappa^G$ so the inflation in the services sector has a lower sensitivity to changes in marginal cost.

Table 5: Effect of population aging on the slope of the Phillips curve

	Baseline 1980	Dem+Pref 2010	Only Dem 2010
Service weight ω	0.4498	0.4953 (+10.11 %)	0.4542 (+0.97 %)
PC slope	1.2759	1.1773 (-7.72 %)	1.2665 (-0.74 %)

Notes: The table compares the weight given to the services sector and the slope of the Phillips curve under different contrafactuals.

As shown in Appendix B, I can derive a general formula for the aggregate Phillips curve as a weighted average of the sectoral ones:

$$\hat{\pi}_t = \omega\hat{\pi}_t^S + (1-\omega)\hat{\pi}_t^G = \beta\mathbb{E}_t\hat{\pi}_{t+1} + \Big[\omega\kappa^S + (1-\omega)\kappa^G\Big](\hat{w}_t - \psi(\hat{K}_t - \hat{L}_t)) - \lambda_t, \tag{42}$$

with $\omega = \sum_j \alpha_j \chi_j \frac{P_j^{\eta-1}}{\sum_j \chi_j P_j^{\eta-1}}$, $\chi_j = \frac{N_j P_j C_j}{\sum_j N_j P_j C_j}$ and $\lambda_t = \omega\kappa^S\hat{P}_t^{S,*} + (1-\omega)\kappa^G\hat{P}_t^{G,*}$.

Aggregate inflation is then a function of the discounted next period expected inflation, the ratio between the prices of the two sectors λ_t and the price mark-up $(\hat{w}_t - \alpha(\hat{K}_t - \hat{L}_t))$. The slope of the aggregate Phillips curve is $\Big[\omega\kappa^S + (1-\omega)\kappa^G\Big]$. The weight ω used to combine the sectoral slopes can be considered as a weighted average of the age-group service preferences α_j using the share of nominal consumption of age group j as weight.

Therefore, whereas the slopes of the sectoral Phillips curves are constant over time, changes in service preferences and population distribution might affect the slope of the aggregate Phillips curve through the weight ω. The first row of Table 5 examines this relationship: the service weight ω increased by approximately 10% from 1980 to 2010 when both changes in preferences and demographic trends are taken into account (from 45% to around 50%, in line with the empirical evidence of section 2.1) and population aging alone (third column) accounts for around 10% of the overall effect.

In terms of the slope of the aggregate Phillips curve (second row), the coefficient decreased overall by around 8% (from 1.28 to 1.18) moving from 1980 to 2010, and again demographic trends explain approximately 10% of the decrease. Therefore, these results suggest that changes in service preferences and population distribution played a non-negligible role in the flattening of the Phillips curve observed in the last decades.

Overall the results so far presented suggest that the demographic trends experienced by the U.S. in the last decades and that are expected to happen in the next 30 years will significantly influence the way monetary policy shocks propagate. Different age groups are not homogeneously affected by the increase in the pass-through of monetary policy and younger households are the most exposed. Moreover, I demonstrate that population aging accounts for a sizable share of the overall change in monetary policy effectiveness and that the novel channel proposed in this paper, i.e., consumption heterogeneity across age groups, significantly contributes to the increase in output responsiveness. Finally, I show that the shift in aggregate demand towards the stickier expenditure category induced by demographic trends partially explains the flattening of the Phillips curve.

5.8.7 Sensitivity analyis

I evaluate the robustness of the theoretical results in a number of variations of the benchmark model. For each alternative specification, I compute the percent change in the IRFs of output and inflation under the different population distribution and mortality rates for 1980 and 2010. Table 6 reports the results.

First of all, I relax the assumption that the production function of the services and the goods sectors have the same labor share. As in Galesi and Rachedi (2018), the labor share of services is set equal to 0.5283 whereas the labor share of goods is set equal to 0.2927. Second, I allow the two sectors to differ in their elasticity of substitution across varieties within sectors. In particular, the elasticities are calibrated to match the estimates of Rebekka and Vermeulen (2012) on the markups of services and manufacturing in the United States. I target a markup equal to 38% in the services sector and to 28% in the goods sector.

Third, following Jones (2021) and Papetti (2019), instead of imposing a constant disutility of labor ϕ across age groups, I assume it to be equal to the cumulative density function of a normal distribution. Figure 31 shows the shape and details of the functional form and parameter values. Fourth, for the PAYGO pension system instead of the constant replacement rate $\bar{d}$ used in the baseline, I fix the contribution rate at the steady state level $\tau = 0.0653$ while the replacement rate $\bar{d}$ is adjusted such that the government budget is balanced in each period.

All these cases deliver quantitatively similar results to the baseline specification (which is reported in the first row of Table 6). This holds on impact as well as after one and two years after the monetary shocks. Overall the robustness exercise confirms that the main conclusions of the previous section are insensitive to several of the assumptions made: demographic trends from 1980 to 2010 significantly increased the responsiveness of output to shocks whereas they had a minor effect on the responsiveness of inflation.

Table 6: Response of Output and Inflation - Robustness Checks

	Output response (%)			Inflation response (%)		
	Time 0	After 1 year	After 2 years	Time 0	After 1 year	After 2 years
Baseline	6.18	4.30	3.22	-0.12	-0.40	-0.89
Different ψ	5.63	4.01	2.93	-0.07	-0.26	-0.64
Different ϵ	5.07	3.72	2.83	-0.15	-0.34	-0.63
Different ϕ	6.97	4.58	2.95	-0.12	-0.36	-0.82
Constant τ	5.79	4.03	3.02	-0.09	-0.31	-0.71
$\theta^G = \theta^S$	2.78	3.85	2.79	-0.02	-0.21	-1.09

Notes: The table reports the percent change in IRFs of output and inflation between 1980 to 2010 under alternative assumptions of the model.

The implicit assumption I have made throughout the paper is that sectoral price stickiness will not change in the future. However, several ongoing trends, e.g., automation and online shopping, are likely to affect the frequency at which prices are adjusted. Therefore, in Figure 23 I test how sensitive are the results to difference price stickiness parameters. The black dashed line reports the percentage change in output due to demographic trends from 1980 to 2050. This is the baseline response obtained by setting $\theta^S = 0.75$ and $\theta^G = 0.25$ as for the main analysis. I compare this response with the responses from all the possible combinations of θ^S and θ^G increasing and decreasing their values by 20% (gray lines). The range of percentage changes of the contemporaneous responses to an expansionary monetary shock is between 8 and 11.5%. Therefore, the specific values chosen for the price stickiness parameters do not affect the main conclusion that demographic trends will increase the responsiveness of output to monetary shocks.

Finally, one might be concerned that the different responses of output and inflation between the two sectors stem from their structural differences, e.g., the fact that only the output from the goods sector can be stored and invested, rather than from the different frequencies of price adjustments. To isolate the role played by price stickiness, the last row of Table 6 reports the percent change in the IRFs assuming that the share of firms unable to adjust their prices is the same between the two sectors, i.e., $\theta^G = \theta^S = 0.75$. The contemporaneous effect of demographic trends on the responsiveness of output and inflation is reduced by approximately one-third and one-fourth respectively. This suggests that the structural differences between the

Figure 23: Model IRFs of output for different demographic structures, alternative price stickiness parameters

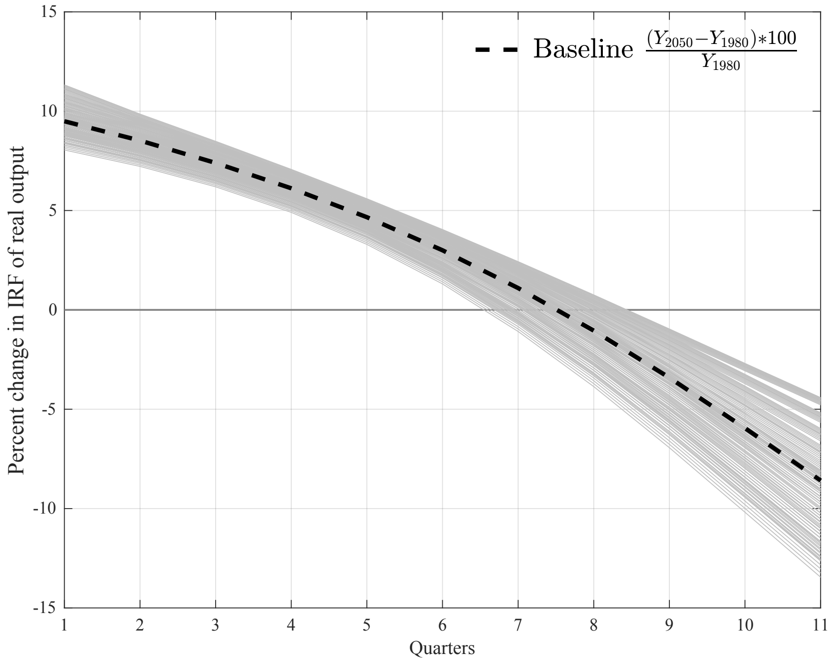

Notes: The figure reports the percentage change in IRFs changing the population distribution and mortality rate from 1980 to 2050. The black dashed line is the baseline response obtained with $\theta^S = 0.75$ and $\theta^G = 0.25$. The gray lines are the responses from all the possible combinations of θ^S and θ^G increasing and decreasing their values by 20%.

two sectors only marginally contribute to the overall change in monetary policy effectiveness caused by population aging.

6 External validity and potential alternative channels

The main contribution of the paper is to study how demographic trends can affect monetary policy effectiveness through consumption heterogeneity. I document that older people tend to consume relatively more services and, since services have a higher level of price rigidities, population aging is shifting the economy towards services leading to an overall increase in price stickiness. The lower aggregate frequency of price adjustments results in an increase in output responsiveness to monetary shocks.

Can the results provided be generalized to other countries? The household-level relationship between price stickiness is documented using a U.S. survey. The same is true for the macroeconomic analysis entirely based on economic and demographic variation among U.S. states. The U.S. has some unique characteristics that could weaken the external validity of the novel channel studied. For instance, a high share of expenditure is devoted to medical expenses.

However, I believe the main findings can be extrapolated to other countries. In a panel of 20 developed countries over the 1970-2007 period, Cravino et al. (2022) document that population aging is accompanied by the rise in the relative size of the service sector. At the same time, Galesi and Rachedi (2018) adopt a VAR specification for a panel of 25 countries to show that output reacts more to monetary policy shocks in countries that are more intensive in services intermediates. Therefore, the positive relationship between demographic trends, the importance of the service sector, and monetary policy effectiveness have already been confirmed outside the U.S. despite differences in the health system, pension scheme, etc.

Another source of concern could be that the different price stickiness in the services and the goods sectors are not the only sources of heterogeneity between the two sectors. In the sensitivity analysis of subsection 5.8.7, I evaluate whether other differences might explain the results. I allow the two sectors to differ in labor share, the elasticity of substitution across varieties within sectors (which lead to different mark-ups) and to have the same frequency of price adjustments. All the structural differences considered do not affect the main findings and confirm that heterogeneity in price stickiness is the most important channel.

There are other channels that are not currently included in the model and which I plan to study more in-depth in the future. For instance, wage rigidities might differ between services and goods. Moreover, I assume workers supply their labor to both sectors and are free to shift freely between them. Frictions in the labor market might make this adjustment costly probably leading to a stronger response in labor demand. Finally, workers of different ages might systematically differ in their skills or their preferences towards the two sectors leading to additional frictions that could strengthen or weaken the heterogeneous sectoral output response to monetary shocks. The empirical evidence suggests that overall the impact of population aging on monetary policy effectiveness is positive but the relative contribution of each channel is a priori difficult to quantify and, therefore, necessitates further analysis.

7 Conclusion

For almost every country in the world, the share of old people is projected to significantly increase and the share of the working population to decrease over the next decades. However, given the extremely slow-moving pace of this transition, limited attention has been given to the way these demographic trends might influence the effectiveness of monetary policy.

I propose and quantify a new channel through which the transmission of monetary policy shocks is affected by the demographic structure of the economy. Using household-level data for the U.S., I show that older people tend to purchase more from product categories that on average adjust their prices less often. Therefore, changes in the population distribution shift the aggregate demand towards categories with a higher level of price stickiness resulting in a stronger response of output to shocks.

To confirm the macro implications of these micro-level findings, I empirically evaluate whether the responsiveness of U.S. states' economic activity to monetary shocks is heterogeneous in their economic structure. I find that population aging leads to an increase in the relative share of services and that the real personal income and the real GDP of more service-intensive states respond significantly more to shocks. No significant differences are found for inflation.

Finally, to assess the overall effects of population aging on the pass-through of monetary policy, I develop a two-sector OLG NK model. I find that demographic trends have a sizable impact on the response of output, that the consumption of younger households is the most exposed to these trends, that the novel channel I proposed significantly contributes to this, and that the flattening of the Phillips curve is partially explained by the fact that the U.S. society is aging.

In conclusion, my research provides substantial evidence that demographic trends, despite their long-term nature, should not be overlooked by policymakers and central bankers even when it comes to short-term policy decisions like the level of the interest rate. The rise in output responsiveness documented implies that contractionary increases in interest rate to tackle the surge in inflation, like the ones we are experiencing in the post-Covid period, will result in a deeper recession compared to the same shocks a few decades ago just because of the different demographic structure. On top of that, younger households are the ones who will observe the strongest decrease in consumption so better coordination between the fiscal and monetary side is necessary to avoid the cost of the recession mainly borne by them.

References

Aksoy, Y., Basso, H. S., Smith, R. P., , and Grasl, T. (2019). "Demographic Structure and Macroeconomic Trends". *American Economic Journal: Macroeconomics 2019, 11(1): 193–222.*

Auclert, A., Malmberg, H., Martenet, F., and Rognlie, M. (2021). "Demographics, Wealth, and Global Imbalances in the Twenty-First Century". *NBER working paper 29161.*

Bárány, Z., Coeurdacier, N., and Guibaud, S. (2022). "Capital flows in an aging world". *Journal of International Economics, 103707.*

Bernanke, B. (2010). "The Economic Outlook and Monetary Policy". *Speech, Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming.*

Bielecki, M., Brzoza-Brzezin, M., and Kolasa, M. (2021). "Demographics, monetary policy, and the zero lower bound". *Journal of Money, Credit, and Banking.*

Bielecki, M., Brzoza-Brzezin, M., and Kolasab, M. (2020). "Demographics and the natural interest rate in the euro area". *European Economic Review, Volume 129.*

Blanchard, O., Cerutti, E., and Summers, L. (2015). "Inflation and Activity – Two Explorations and their Monetary Policy Implications". *NBER Working Paper No. 21726.*

Bobeica, E., Nickel, C., Lis, E., and Sun, Y. (2017). "Demographics and inflation". *ECB Working Paper No 2006.*

Boivin, J., Kiley, M. T., and Mishkin, F. S. (2010). "How Has the Monetary Transmission Mechanism Evolved Over Time?". *Handbook of Monetary Economics, Volume 3, 2010, Pages 369-422 Chapter 8.*

Brzoza-Brzezina, M. and Kolasa, M. (2021). "Intergenerational redistributive effects of monetary policy". *Journal of European Economic Association 20(2): 549-580 (2022).*

Carvalho, C., Ferrero, A., and Nechio, F. (2016). "Demographics and real interest rates: Inspecting the mechanism". *European Economic Review 88 (2016) 208–226.*

Clayton, C., Jaravel, X., and Schaab, A. (2018). "Heterogeneous Price Rigidities and Monetary Policy". *Working paper.*

Cloyne, J., Clodomiro, F., Maren, F., and Surico, P. (2022). "Monetary Policy, Corporate Finance and Investment". *Journal of the European Economic Association (forthcoming).*

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

Cravino, J., Lan, T., and Levchenko, A. (2020). "Price Stickiness Along the Income Distribution and the Effects of Monetary Policy". *Journal of Monetary Economics, 110:19-32.*

Cravino, J., Levchenko, A., and Rojas, M. (2022). "Population aging and structural transformation". *American Economic Journal: Macroeconomics, Vol. 14(4). 479-498.*

de Albuquerque, P. C., Caiado, J., and Pereira, A. (2020). "Population aging and inflation: evidence from panel cointegration". *Journal of Applied Economics 23(1), 469–484.*

Eggertsson, G. B., Mehrotra, N. R., and Robbins, J. A. (2019). "Population Aging and the Macroeconomy". *A Model of Secular Stagnation: Theory and Quantitative Evaluation.*

Fujiwara, I. and Teranishi, Y. (2008). "A dynamic new Keynesian life-cycle model: Societal aging, demographics, and monetary policy". *Journal of Economic Dynamics and Control 32 (2008) 2398–2427.*

Fullerton, H. N. (1999). "Labor force participation: 75 years of change, 1950-98 and 1998-2025". *Monthly Labor Review, 122:3–12.*

Galesi, A. and Rachedi, O. (2018). "Services Deepening and the Transmission of Monetary Policy". *Journal of the European Economic Association, Volume 17, Issue 4, August 2019, Pages 1261–1293.*

Galí, J. (2015). "Monetary Policy, Inflation, and the Business Cycle". *Princeton University Press, Second Edition.*

Hazell, J., Herreno, J., Nakamura, E., and Steinsson, J. (2022). "The Slope of the Phillips Curve: Evidence from U.S. States". *Quarterly Journal of Economics, 137(3), 1299-1344.*

Heer, B., Rohnbacher, S., and Scharrer, C. (2017). "Aging, the great moderation, and the business-cycle volatility in a life-cycle model". *Macroeconomic Dynamics, 21, 362–383.*

Imam, P. (2014). "Shock from Greying: Is the Demographic Shift Weakening Monetary Policy Effectiveness". *International Journal of Finance & Economics.*

Jaimovich, N., Pruitt, N. S., and Siu, H. (2013). "The demand for youth: Explaining age differences in the volatility of hours". *American Economic Review 103(7), 3022–3044.*

Jamilov, R., Bellifemine, M., and Couturier, A. (2023). "The Regional Keynesian Cross". *Quarterly Journal of Economics, 116, 969-1007.*

Jones, C. (2021). "Aging, secular stagnation and the business cycle". *Review of Economics and Statistics (forthcoming).*

Jordà, O. (2005). "Estimation and Inference of Impulse Responses by Local Projections". *American Economic Review, 95 (1), 161–182.*

Jorda, O., Chitra, M., Fernanda, N., and Eric, T. (2019). "Why Is Inflation Low Globally?". *FRBSF Economic Letter 2019-19.*

Kantur, Z. (2013). "Aging and Monetary Policy". *Working paper.*

Kimberly, B., Curtis, C., Lugauer, S., and Mark, N. C. (2021). "Demographics and Monetary Policy Shocks". *Journal of Money, Credit, and Banking.*

Kopecky, J. (2022). "Okay Boomer... Excess Money Growth, Inflation, and Population Aging". *Macroeconomic Dynamics.*

Kronick, J. and Ambler, S. (2019). "Do Demographics Affect Monetary Policy Transmission in Canada?". *International Journal of Finance and Economics. Vol. 24(2), pg. 787–811. April.*

Laurence, B. and Mazumder, S. (2011). "Inflation Dynamics and the Great Recession". *Brookings Papers on Economic Activity, 42(Spring): 337–381.*

Leahy, J. and Thapar, A. (2022). "Age Structure and the Impact of Monetary Policy". *American Economic Journal: Macroeconomics. 14, NO. 4, 136-73.*

Lis, E., Nickel, C., and Papetti, A. (2020). "Demographics and inflation in the euro area: a two-sector new Keynesian perspective". *Working Paper Series 2382, European Central Bank.*

Lisack, N., Sajedi, R., and Thwaites, G. (2021). "Population Aging and the Macroeconomy". *International Journal of Central Banking.*

McLeay, M. and Tenreyro, S. (2019). "Optimal inflation and the identification of the Phillips curve". *NBER Macroeconomics Annual 2019.*

Nakamura, E. and Steinsson, J. (2008). "Five facts about prices: a reevaluation of menu cost models". *The Quarterly Journal of Economics, Volume 123, Issue 4, November 2008, Pages 1415–1464.*

Nakamura, E. and Steinsson, J. (2018). "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect". *Quarterly Journal of Economics, 2018, 133(3).*

Papetti, A. (2019). "Demographics and the natural real interest rate: historical and projected paths for the euro area". *Working Paper Series 2258, European Central Bank.*

Papetti, A. (2021). "Population Aging, Relative Prices and Capital Flows across the Globe". *Bank of Italy Working Papers, No 1333.*

Rachedi, O. and Basso, H. S. (2021). "The Young, the Old, and the Government: Demographics and Fiscal Multipliers". *American Economic Journal: Macroeconomics, 2021, 13(4), 110-141.*

Ramey, V. (2016). "Macroeconomic shocks and their propagation". *In Handbook of Macroeconomics. Vol. 2, 71.*

Rebekka, C. and Vermeulen, P. (2012). "Markups in the Euro Area and the US over the Period 1981-2004: a Comparison of 50 Sector". *Empirical Economics, 42, 53-77.*

Romer, C. D. and Romer, D. H. (2004). "A new measure of monetary shocks: Derivation and implications". *American Economic Review 94(4), 1055-84.*

Rubbo, E. (2022). "Networks, Phillips Curves, and Monetary Policy". *Econometrica (forthcoming).*

Shimer, R. (2001). "The Impact of Young Workers on the Aggregate Labor Market". *Quarterly Journal of Economics, 116, 969-1007.*

Wong, A. (2014). "Population Aging and the Aggregate Effects of Monetary Policy". *MPRA Paper No. 57096, University Library of Munich, Germany.*

Wong, A. (2021). "Refinancing and The Transmission of Monetary Policy to Consumption". *R&R American Economic Review.*

Yoshino, N. and Miyamoto, H. (2017). "Declined effectiveness of fiscal and monetary policies faced with aging population in Japan". *Japan and the World Economy 42 (2017) 32–44.*

A List of model equations

This appendix presents the full set of equations of the model described in section 5. Each period t there is a distribution of households of different ages j with $j \in \{1, ..., J\}$. On the supply side, there are two sectors s, services and goods, so that $s \in \{S, G\}$. Variables are expressed in real terms as $x_t = \frac{X_t}{P_t}$, the sectoral MC_t^s are deflated by the relative P_t^s.

Households[18]:

$$P_{t,j}^* c_{t,j} + a_{t+1,j+1} = \frac{R_t^a}{\pi_t} a_{t,j} + (1-\tau_t) w_t l_{t,j} h_j \mathbf{I}_{j \le jw} + pen_t \mathbf{I}_{j > jw} + beq_t \tag{43}$$

$$a_{t,0} = 0 \qquad a_{t+J+1,J+1} = 0 \tag{44}$$

$$\nu l_{t,j}^{\eta} = \frac{(1-\tau_t) w_t h_j \mathbf{I}_{j \le jw}}{Z_t^{\omega-\alpha}} c_{t,j}^{-\sigma} \tag{45}$$

$$\frac{c_{t,j}^{-\sigma}}{P_{t,j}^*} = \beta s_j \frac{c_{t+1,j+1}^{-\sigma}}{P_{t+1,j+1}^*} \frac{R_{t+1}^a}{\pi_{t+1}} \tag{46}$$

Firms:

$$P_t^{S,*} mc_{i,t}^S = \Big(\frac{w_t}{(1-\psi)}\Big)^{1-\psi} \Big(\frac{r_t^k}{\psi}\Big)^{\psi} \tag{47}$$

$$P_t^{G,*} mc_{i,t}^G = \Big(\frac{w_t}{(1-\psi)}\Big)^{1-\psi} \Big(\frac{r_t^k}{\psi}\Big)^{\psi} \tag{48}$$

$$K_{i,t}^s = \frac{\psi w_t}{(1-\psi) r_t^k} L_{i,t}^s \tag{49}$$

$$f_t = Y - w_t L_t - r_t^k K_t \tag{50}$$

$$v_t^s = (1-\theta^s) \left(\pi_t^{s,\#}\right)^{-\epsilon} (\pi_t^s)^{\epsilon} + \theta^s (\pi_t^s)^{\epsilon} v_{t-1}^s \tag{51}$$

$$(\pi_t^s)^{1-\epsilon} = (1-\theta^s) \left(\pi_t^{s,\#}\right)^{1-\epsilon} + \theta^s \tag{52}$$

$$x_{1,t}^s = \frac{1}{R_t} Y_t^s P_t^{s,*} mc_{i,t}^s + \theta^s \beta \mathbb{E}_t \left(\pi_{t+1}^s\right)^{\epsilon} x_{1,t+1}^s \tag{53}$$

$$x_{2,t}^s = \frac{1}{R_t} Y_t^s P_t^{s,*} + \theta^s \beta \mathbb{E}_t \left(\pi_{t+1}^s\right)^{\epsilon-1} x_{2,t+1}^s \tag{54}$$

$$\pi_t^{s,\#} = \frac{\epsilon}{\epsilon-1} \pi_t^s \frac{x_{1,t}^s}{x_{2,t}^s} \tag{55}$$

[18]See Appendix B for the derivation of ω.

Representative investment fund:

$$K_{t+1} = (1-\delta)K_t + \Big[1 - \frac{S}{2}\Big(\frac{I_t}{I_{t-1}} - 1\Big)^2\Big] I_t \tag{56}$$

$$A_{t+1} = q_t(1-\delta)K_t + I_t + p_t^d \tag{57}$$

$$\frac{R_t^a}{\pi_t} A_t = \Big[r_t^k + q_t(1-\delta)\Big]K_t + f_t + p_t^d \tag{58}$$

$$R_t q_t = \mathbb{E}_t\Big[\Big(r_{t+1}^k + q_{t+1}(1-\delta)\Big)\pi_{t+1}\Big] \tag{59}$$

$$R_t p_t^d = \mathbb{E}_t\Big[\Big(p_{t+1}^d + f_{t+1}\Big)\pi_{t+1}\Big] \tag{60}$$

$$1 = q_t\Big[1 - \frac{S}{2}\Big(\frac{I_t}{I_{t-1}} - 1\Big)^2 - S\Big(\frac{I_t}{I_{t-1}} - 1\Big)\frac{I_t}{I_{t-1}}\Big] + \mathbb{E}_t\Big[\frac{\pi_{t+1}}{R_t} q_{t+1} S\Big(\frac{I_{t+1}}{I_t} - 1\Big)\Big(\frac{I_{t+1}}{I_t}\Big)^2\Big] \tag{61}$$

Government:

$$pen_t = \bar{d}(1-\tau_t)w_t \sum_{j=0}^{jw} N_j h_j \tag{62}$$

$$\tau_t w_t \sum_{j=0}^{jw} N_j h_j = pen_t \sum_{j=jw+1}^{J} N_j \tag{63}$$

Monetary authority:

$$\frac{R_t}{R} = \Big(\frac{\pi_t}{\pi}\Big)^{\phi_\pi}\Big(\frac{Y_t}{Y}\Big)^{\phi_y} e^{\nu_t^r} \tag{64}$$

$$\nu_t^r = \rho^\nu \nu_{t-1}^r + \epsilon_t^\nu \tag{65}$$

Market clearing:

$$L_t = L_t^S + L_t^G = \sum_{j=1}^{jw} N_j h_j n_{t,j}, \quad A_t = \sum_{j=1}^{J} N_{j-1} a_{t,j}, \quad K_t = K_t^S + K_t^G \tag{66}$$

$$beq_t = \sum_{j=1}^{J} (N_{j-1} - N_j) a_{t,j} \frac{R_t^a}{\pi_t} \tag{67}$$

$$Y_t^S = (K_t^S)^\alpha (L_t^S)^{1-\alpha} / v_t^S = C_t^S \tag{68}$$

$$Y_t^G = (K_t^G)^\alpha (L_t^G)^{1-\alpha} / v_t^G = C_t^G + I_t \tag{69}$$

$$C_t = P_t^{S,*} C_t^S + P_t^{G,*} C_t^G \tag{70}$$

$$Y_t = P_t^{S,*} Y_t^S + P_t^{G,*} Y_t^G \tag{71}$$

$$C_t^S = \sum_{j=1}^{J} \alpha_j \left(P_{t,j}^{S,*}\right)^{\eta} N_j c_{t,j}, \qquad C_t^G = \sum_{j=1}^{J} \left(1-\alpha_j\right)\left(P_{t,j}^{G,*}\right)^{\eta} N_j c_{t,j} \tag{72}$$

Price dynamics

$$\frac{\pi_t^G}{\pi_t^S} = \frac{Z_t}{Z_{t-1}} \tag{73}$$

$$\pi_t = \pi_t^S \frac{\omega + (1-\omega)Z_t^{1-\eta}}{\omega + (1-\omega)Z_{t-1}^{1-\eta}} \tag{74}$$

$$P_t^{S,*} = \frac{P_t^S}{P_t} = \left[\omega + (1-\omega)Z_t^{1-\eta}\right]^{\frac{1}{\eta-1}}, \quad P_t^{G,*} = \frac{P_t^G}{P_t} = \left[\omega Z_t^{\eta-1} + (1-\omega)\right]^{\frac{1}{\eta-1}} \tag{75}$$

$$P_{t,j}^{*} = \frac{P_{t,j}}{P_t} = \left[\frac{\alpha_j + (1-\alpha_j)Z_t^{1-\eta}}{\omega + (1-\omega)Z_t^{1-\eta}}\right]^{\frac{1}{1-\eta}} \tag{76}$$

$$P_{t,j}^{S,*} = \frac{P_{t,j}}{P_t^S} = \left[\alpha_j + (1-\alpha_j)Z_t^{1-\eta}\right]^{\frac{1}{1-\eta}}, \quad P_{t,j}^{G,*} = \frac{P_{t,j}}{P_t^G} = \left[\alpha_j Z_t^{\eta-1} + (1-\alpha_j)\right]^{\frac{1}{1-\eta}} \tag{77}$$

B Aggregate Phillips curve derivation

In this section, I derive a general formula for the aggregate Phillips curve as a weighted average of the sectoral ones.

The demand functions for services and goods relative to the households maximization problem are given by:

$$c_{t,j}^S = \alpha_j \left(\frac{P_t^S}{P_{t,j}}\right)^{-\eta} c_{t,j}, \qquad c_{t,j}^G = (1-\alpha_j)\left(\frac{P_t^G}{P_{t,j}}\right)^{-\eta} c_{t,j} \tag{78}$$

where $c_{t,j}$ is the aggregate consumption of household j and $P_{t,j}$ is the price index associated with its bundle.

Adding across households, one can obtain the following expression of the sectoral aggregate demand:

$$C_t^S = \omega_t \left(\frac{P_t^S}{P_t}\right)^{-\eta} C_t, \qquad C_t^G = (1-\omega_t)\left(\frac{P_t^G}{P_t}\right)^{-\eta} C_t \tag{79}$$

where, following Cravino et al. (2020), the expenditure share is defined as $\omega_t \equiv \sum_j \alpha_j \chi_{t,j} \frac{P_{t,j}^{\eta-1}}{\sum_j \chi_{t,j} P_{t,j}^{\eta-1}}$ and $\chi_{t,j}$ is the share of household j in aggregate expenditures at time t. One can then define the aggregate price index as $P_t \equiv \left[\omega_t^{\frac{1}{\eta}}(P_t^S)^{1-\eta} + (1-\omega_t)^{\frac{1}{\eta}}(P_t^G)^{1-\eta}\right]^{\frac{1}{1-\eta}}$.

To simplify the log-linearization process, I assume that ω_t is constant and equal to its steady state value. By log-linearizing the aggregate price index I obtain:

$$\hat{p}_t = \omega \hat{p}_t^S + (1-\omega)\hat{p}_t^G \tag{80}$$

which allows obtaining an expression for the aggregate inflation rate:

$$\hat{\pi}_t = \hat{p}_t - \hat{p}_{t-1} = \omega(\hat{p}_t^S - \hat{p}_{t-1}^S) + (1-\omega)(\hat{p}_t^G - \hat{p}_{t-1}^G) = \omega\hat{\pi}_t^S + (1-\omega)\hat{\pi}_t^G \tag{81}$$

By solving the cost minimization problem of the intermediate firm i, I find the following expression for the sectoral marginal costs in real terms:

$$P_t^{S,*} mc_{i,t}^S = \Big(\frac{w_t}{(1-\psi)}\Big)^{1-\psi}\Big(\frac{r_t^k}{\psi}\Big)^{\psi} \tag{82}$$

$$P_t^{G,*} mc_{i,t}^G = \Big(\frac{w_t}{(1-\psi)}\Big)^{1-\psi}\Big(\frac{r_t^k}{\psi}\Big)^{\psi} \tag{83}$$

as well as the standard relationship between capital and labor for both sectors s:

$$K_{i,t}^s = \frac{\psi w_t}{(1-\psi) r_t^k} L_{i,t}^s \tag{84}$$

Notice that since all firms use the same capital-output ratio I can drop the subindex i. I then log-linearize the marginal cost equations for both sectors:

$$\hat{mc}_t^S = -\hat{P}_t^{S,*} + (1-\psi)\hat{w}_t + \psi\hat{r}_t^k \tag{85}$$

$$\hat{mc}_t^G = -\hat{P}_t^{G,*} + (1-\psi)\hat{w}_t + \psi\hat{r}_t^k \tag{86}$$

and by combining the two log-linearized expression of the capital-output ratios $\hat{K}_t^s - \hat{L}_t^s = \hat{w}_t - \hat{r}_t^k$, I obtain that $\hat{K}_t - \hat{L}_t = \hat{w}_t - \hat{r}_t^k$.

I can now replace the expressions of the log-linearized real marginal costs in the sectoral Phillips curve, obtained by linearizing equation (24) around a steady state with zero inflation in both sectors:

$$\hat{\pi}_t^S = \beta\mathbb{E}_t\hat{\pi}_{t+1}^S + \kappa^S \hat{mc}_t^S \tag{87}$$

$$\hat{\pi}_t^G = \beta\mathbb{E}_t\hat{\pi}_{t+1}^G + \kappa^G \hat{mc}_t^G \tag{88}$$

with

$$\kappa^S = \frac{(1-\theta^S)(1-\theta^S\beta)}{\theta^S}, \qquad \kappa^G = \frac{(1-\theta^G)(1-\theta^G\beta)}{\theta^G} \tag{89}$$

i.e.,

$$\hat{\pi}_t^S = \beta\mathbb{E}_t\hat{\pi}_{t+1}^S + \kappa^S[-\hat{P}_t^{S,*} + (1-\psi)\hat{w}_t + \psi\hat{r}_t^k] \tag{90}$$

$$\hat{\pi}_t^G = \beta \mathbb{E}_t \hat{\pi}_{t+1}^G + \kappa^G[-\hat{P}_t^{G,*} + (1-\psi)\hat{w}_t + \psi \hat{r}_t^k] \tag{91}$$

i.e.,

$$\hat{\pi}_t^S = \beta \mathbb{E}_t \hat{\pi}_{t+1}^S + \kappa^S[-\hat{P}_t^{S,*} + \hat{w}_t - \psi(\hat{w}_t - \hat{r}_t^k)] \tag{92}$$

$$\hat{\pi}_t^G = \beta \mathbb{E}_t \hat{\pi}_{t+1}^G + \kappa^G[-\hat{P}_t^{G,*} + \hat{w}_t - \psi(\hat{w}_t - \hat{r}_t^k)] \tag{93}$$

Using the fact that $\hat{w}_t - \hat{r}_t^k = \hat{K}_t - \hat{L}_t$, I find:

$$\hat{\pi}_t^S = \beta \mathbb{E}_t \hat{\pi}_{t+1}^S + \kappa^S[-\hat{P}_t^{S,*} + \hat{w}_t - \psi(\hat{K}_t - \hat{L}_t)] \tag{94}$$

$$\hat{\pi}_t^G = \beta \mathbb{E}_t \hat{\pi}_{t+1}^G + \kappa^G[-\hat{P}_t^{G,*} + \hat{w}_t - \psi(\hat{K}_t - \hat{L}_t)] \tag{95}$$

The sectoral Phillips curves can be replaced in equation (81):

$$\hat{\pi}_t = \omega \hat{\pi}_t^S + (1-\omega)\hat{\pi}_t^G = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \Big[\omega \kappa^S + (1-\omega)\kappa^G\Big](\hat{w}_t - \psi(\hat{K}_t - \hat{L}_t)) - \lambda_t \tag{96}$$

with $\lambda_t = \omega \kappa^S \hat{P}_t^{S,*} + (1-\omega)\kappa^G \hat{P}_t^{G,*}$.

C Additional figures and tables

Figure 24: Average price rigidities across expenditure categories

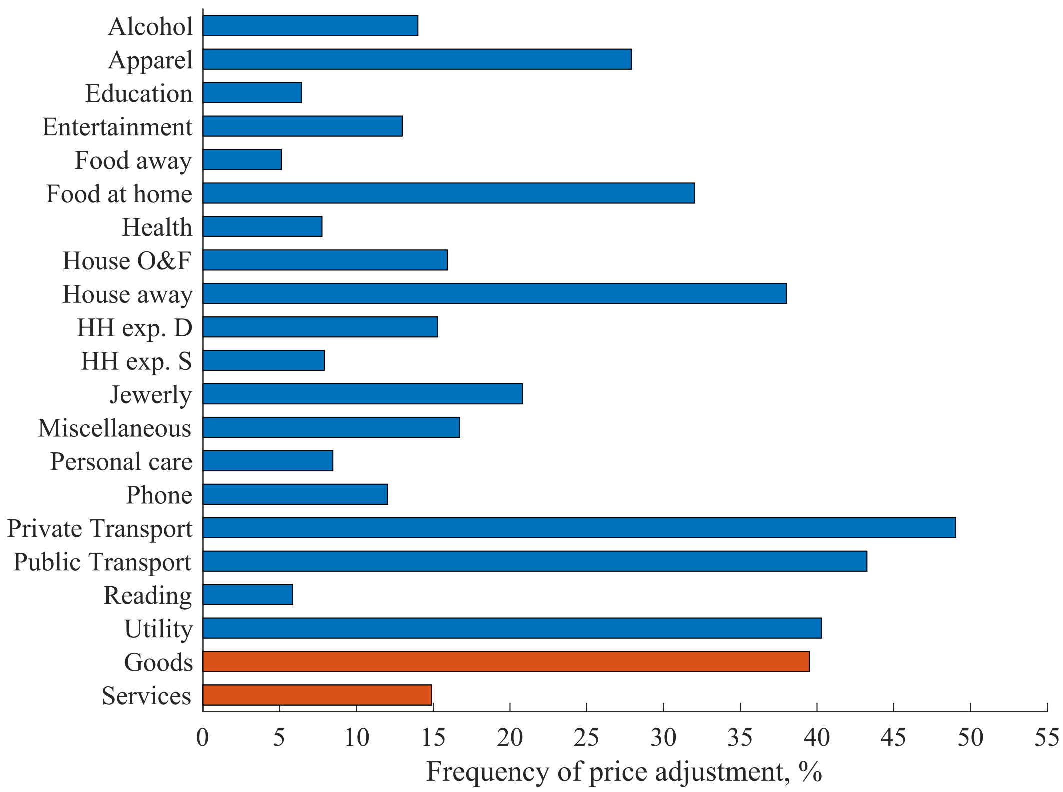

Notes: The bar plot shows the weighted average frequency of price adjustment across different expenditure categories as well as for the aggregation of the categories into Goods and Services.

Figure 25: Frequency of price adjustment and mean implied duration across age groups

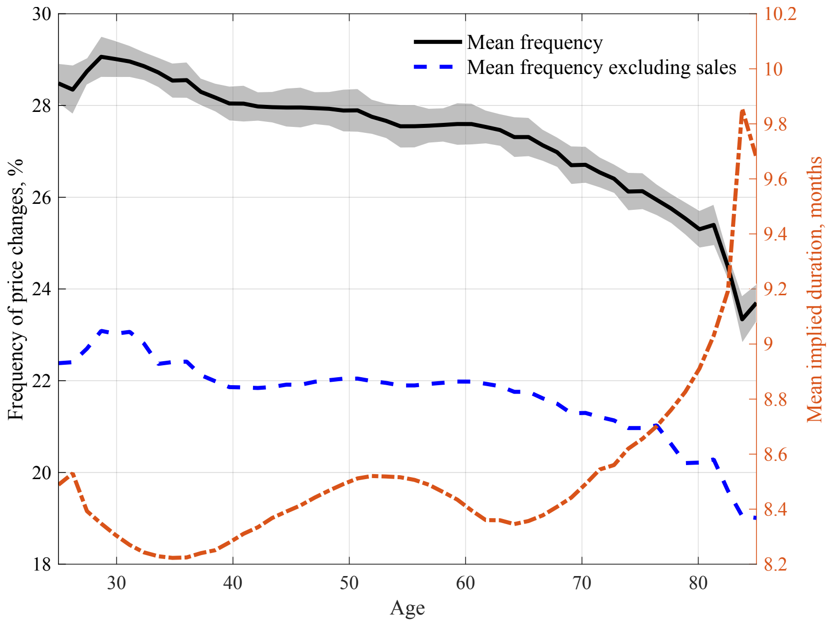

Notes: The figure plots the weighted average frequency of price adjustment with and without temporary sales (left axis) alongside the mean implied duration (right axis) across age groups. The shaded area is the 95% confidence band. The frequency of price adjustment is computed as the fraction of the number of times an item changes its price over the number of times the item is observed and expressed in percent per month. The mean implied duration captures after how many months, on average, a firm in sector j adjusts its price. The expenditure shares are computed using data from the CEX whereas the sectoral price stickiness parameters are retrieved from Nakamura and Steinsson (2008).

Figure 26: Impact of monetary policy on the regional real personal income, different thresholds

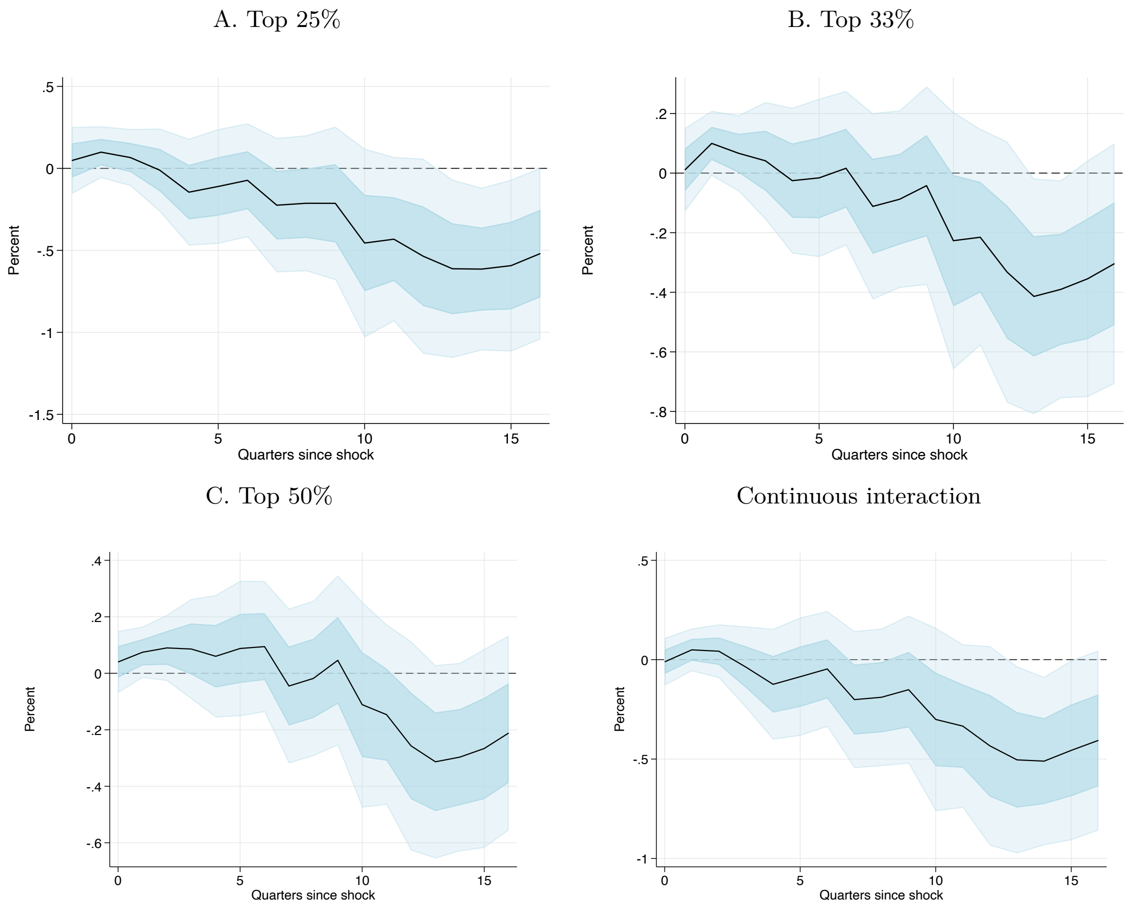

Notes: Each panel reports the interaction coefficients between the monetary policy shock and the dummy identifying the top 20% of the services/manufacturing production ratio distribution using as dependent variable either the state-level real personal income. The dark shaded area and the light shaded area represent the 68% and the 95% confidence intervals respectively. The horizontal axis is in quarters.

Figure 27: Impact of monetary policy on the regional inflation rate, different thresholds

A. Top 25%

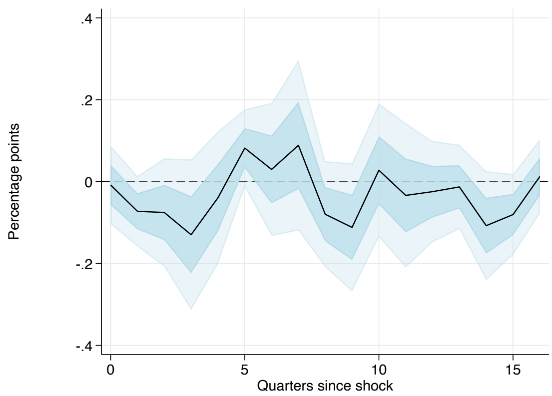

B. Top 33%

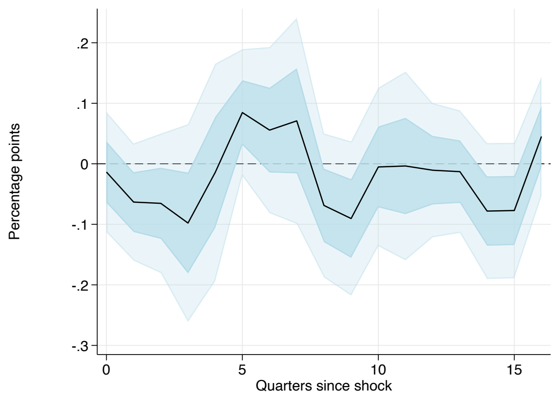

C. Top 50%

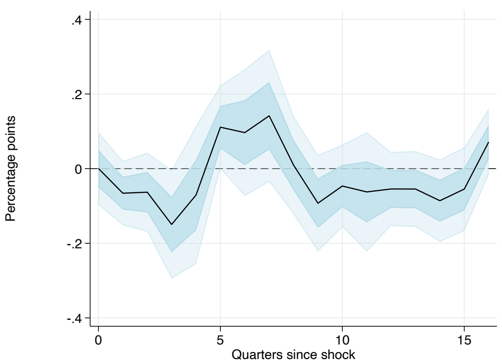

Continuous interaction

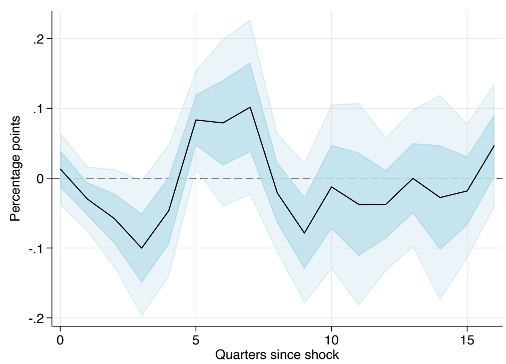

Notes: Each panel reports the interaction coefficients between the monetary policy shock and the dummy identifying the top 20% of the services/manufacturing production ratio distribution using as dependent variable either the state-level inflation rate. The dark shaded area and the light shaded area represent the 68% and the 95% confidence intervals respectively. The horizontal axis is in quarters.

Figure 28: Impact of monetary policy on regional variables, Nakamura and Steinsson (2018) monetary shocks

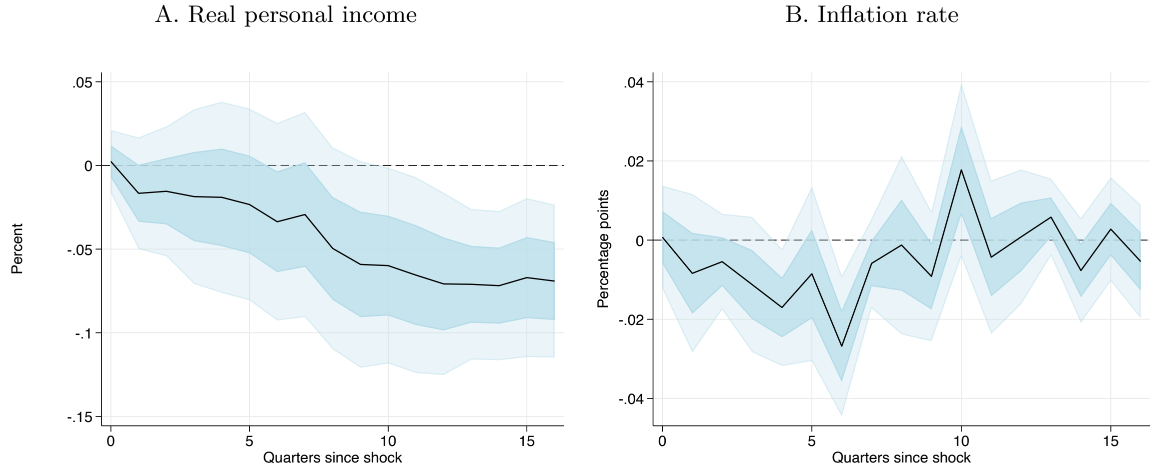

Notes: Each panel reports the interaction coefficients between the monetary policy shock and the dummy identifying the top 20% of the services/manufacturing production ratio distribution using as dependent variable either the state-level real personal income or the inflation rate. The dark shaded area and the light shaded area represent the 68% and the 95% confidence intervals respectively. The horizontal axis is in quarters.

Figure 29: Impact of monetary policy on the production of the regional services

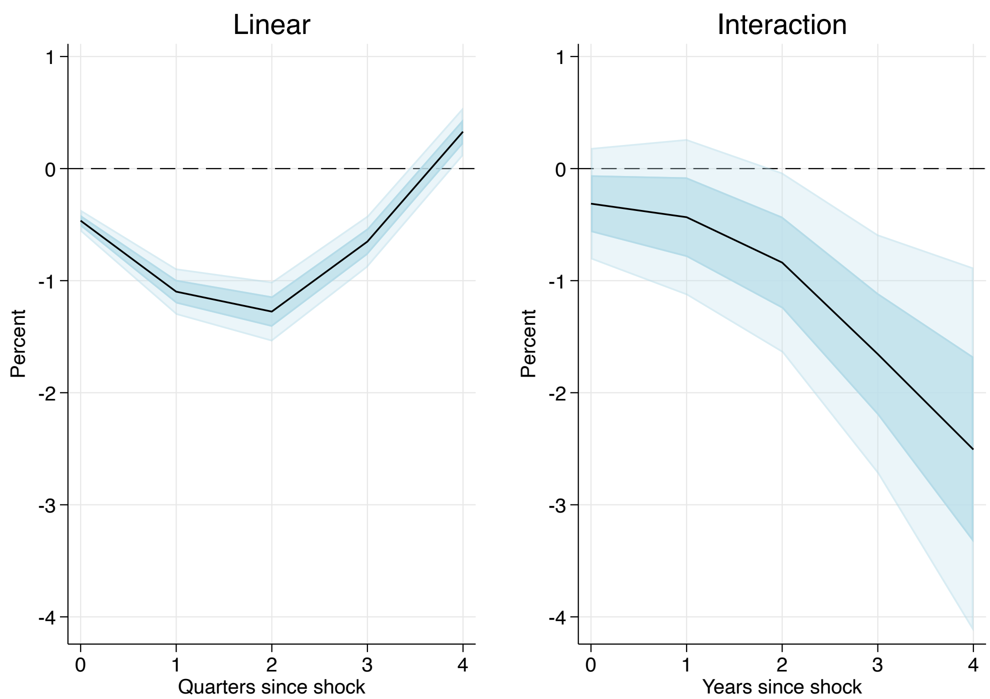

Notes: The left panel of the figure plots the response of the state-level log of the real services production to a percentage point contractionary monetary policy shock, as well as the 68% (dark shaded area) and 95% (light shaded area) confidence intervals. The horizontal axis is in years. The right panel reports the interaction coefficients between the monetary policy shock and the dummy identifying the top 20% of the services/manufacturing production ratio distribution.

Figure 30: Model vs Data

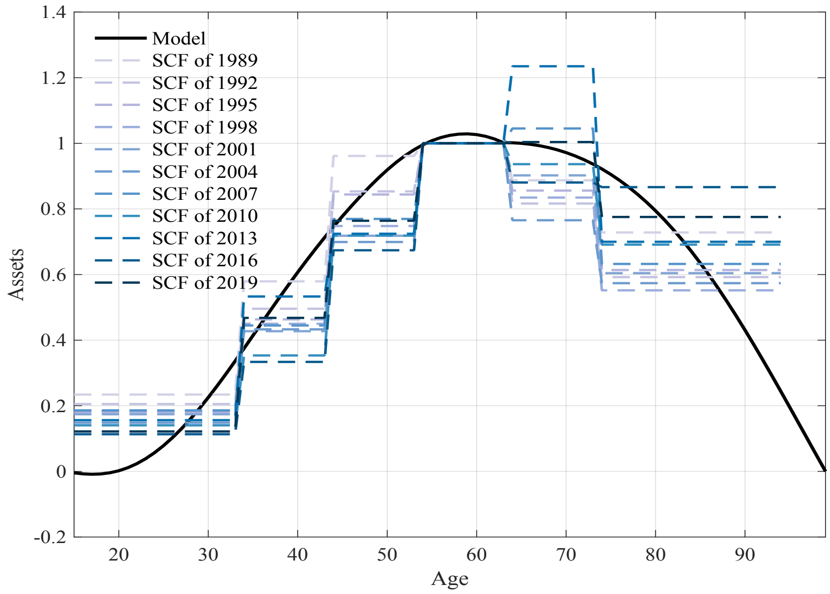

Notes: The plot compares the steady state assets profile from the model (Age 65 = 1) with the asset profile taken from the Survey of Consumer Finances for different years (Age group 55-64 = 1). *Source*: Survey of Consumer Finances.

Figure 31: Age dependent disutility of labor supply, ν_j

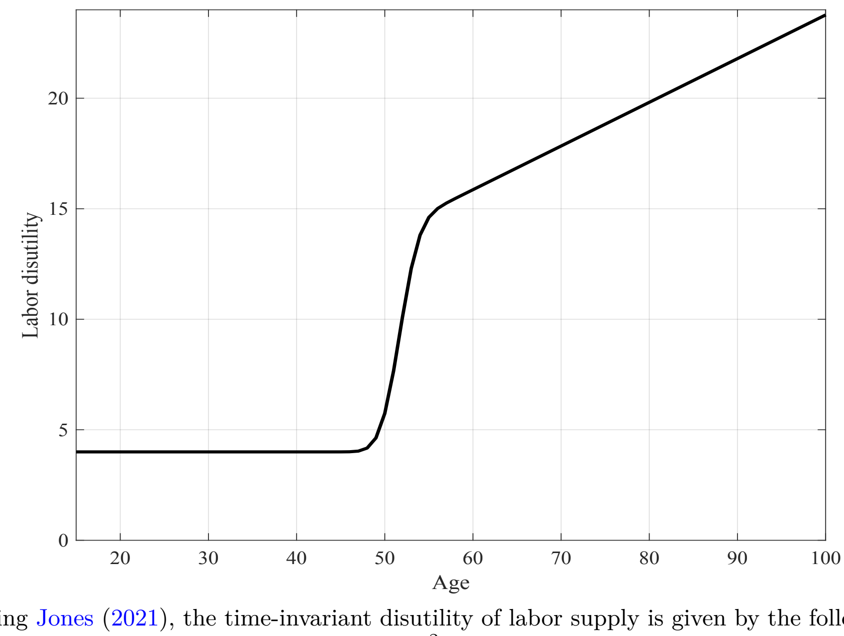

Notes: Following Jones (2021), the time-invariant disutility of labor supply is given by the following expression: $\nu_j = b_0 + (b_1 \frac{j}{J+1}) \int_{-\infty}^{J} \frac{1}{(J+1)b_3\sqrt{2\pi}} exp\left\{ \frac{1}{2} \left(\frac{j-(J+1)b_2}{(J+1)b_3} \right)^2 dj \right\}$ where the parameter values chosen are: $b_0 = 4$, $b_1 = 17$, $b_2 = 0.65$, $b_3 = 0.02$ as in Papetti (2019). $J + 1 = 86$ is the number of periods the individual can be alive since the household enters the world at age 15 and remains alive up to the maximum age of 100. Finally, the integral expression is the normal cumulative distribution function over age j with mean $b_2(J + 1)$ and standard deviation $b_3(J + 1)$.

Table 7: The table reports the expenditure shares across the major consumption categories for different age groups

	Age groups						
	25-	(30,35]	(40,45]	(50,55]	(60,65]	(70,75]	80+
Alcohol	2.1	1.4	1.2	1.2	1.2	1.1	0.6
Apparel	5.1	4.8	4.7	4.2	3.8	3.1	2.3
Education	6.7	1.5	2.4	3.9	1.0	0.6	0.4
Energy	3.8	5.0	5.4	5.5	6.0	6.7	7.9
Entertainment	5.9	7.0	7.5	6.9	6.8	6.0	4.4
Food Away	6.1	5.6	5.8	5.8	5.6	5.1	4.1
Food at Home	11.4	12.5	13.0	12.1	12.3	12.9	13.5
Medical	3.4	5.4	6.4	7.6	10.7	15.1	19.0
Household F&O	6.4	9.9	9.1	9.0	9.8	10.1	11.1
Other Lodging	1.2	1.0	1.4	2.0	1.8	2.0	0.9
Owned Dwellings	1.8	6.5	7.5	7.7	8.1	7.6	5.9
Other Expenses	0.9	1.1	1.3	1.4	1.6	1.8	2.4
Personal Care	1.9	1.9	2.0	1.9	1.9	2.0	2.1
Private Transportation	20.5	21.8	21.7	21.6	20.8	17.5	11.3
Public Transportation	1.2	1.3	1.4	1.5	1.8	1.7	1.1
Reading	0.3	0.4	0.4	0.5	0.6	0.7	0.7
Rented Dwellings	19.4	10.8	6.4	4.4	3.7	3.9	10.2
Tobacco	1.3	1.0	1.1	1.2	1.1	0.8	0.4
Water	0.6	1.1	1.2	1.2	1.3	1.5	1.7

Monetary policy shocks and inflation inequality

Christoph Lauper* Giacomo Mangiante†

Link to most recent version

March 2023

Abstract

This paper studies how monetary policy shocks influence the distribution of household-level inflation rates. We find that (i) contractionary monetary policy shocks significantly and persistently decrease inflation dispersion in the economy, and that (ii) the expenditures on *Energy*, *Water* and *Gasoline* are the main drivers behind this result. Moreover, (iii) different demographic groups are heterogeneously affected by monetary policy. Due to the different consumption baskets purchased, low- and middle-income households experience higher median inflation rates, which are at the same time more responsive to a contractionary monetary shock, leading to an overall convergence of inflation rates across income groups. The same result holds for expenditure and salary groups. These findings imply that (iv) the impact of monetary policy shocks on expenditure inequality is between 20 and 30% more muted once we control for differences in individual inflation rates. Overall, our empirical evidence highlights the importance of inflation heterogeneity in studying the distributional consequences that monetary policies can have.

Keywords: monetary policy, inflation inequality, distributional effects
JEL classification: E31, E52

*University of Lausanne, Lausanne, Switzerland, email: christoph.lauper@unil.ch

†University of Lausanne, Lausanne, Switzerland, email: giacomo.mangiante@unil.ch

We would like to thank Florin Bilbiie, Jean-Paul Renne, and Andreas Tischbirek for their valuable comments and guidance. We are also grateful to Philippe Bacchetta, Kenza Benhima, Aurélien Eyquem, Yuriy Gorodnichenko, Andrei Levchenko, Christina Patterson, Ricardo Reis, and fellow UNIL Ph.D. students for helpful suggestions that significantly improved the paper. Finally, we would like to thank the participants of the conferences organized by the International Association for Applied Econometrics, the Swiss Society of Economics and Statistics, the Society of Economics of the Household, the Ph.D. - Economics Virtual Seminar, the VMACS (Virtual Macro Seminar), the Berkeley GEMS, the Lausanne Macro Reading Group, the 9th Italian Congress of Econometrics and Empirical Economics, the 61st Annual Conference of the Italian Economic Association and the 15th Economics Graduate Students Conference for precious feedback.

1 Introduction

The relationship between monetary policy and heterogeneity has become increasingly important in macroeconomic research, both from a theoretical and empirical point of view. Changes in interest rate do not impact households homogeneously. Renters and homeowners, savers and hand-to-mouths, high-skilled and low-skilled workers are only a few examples of different demographic groups that have been found to bear the consequences of the decisions made by the monetary authorities in completely different ways. Therefore, in the last few years, both economic researchers and central bankers have shifted their focus from aggregate to more granular effects to better understand the different channels through which monetary policy can affect individual households and firms. However, the importance of inflation heterogeneity, i.e., the different inflation rates experienced by households due to the variations in the consumption baskets they purchase, for the distributional effects of monetary policy has so far received limited attention.

This paper studies how monetary policy influences the distribution of the individual inflation rates to which different households are exposed. We compute a measure of the inflation rate at the household level and we document that contractionary monetary shocks reduce the median as well as the cross-sectional standard deviation of the distribution of inflation rates. The decrease in inflation dispersion is almost entirely driven by expenditures on *Energy*, *Water*, and *Gasoline*. The inflation rates of these sectors, despite the fact that they account for a relatively small share of the aggregate consumption bundle, are extremely sensitive to changes in interest rate. We then study how the inflation rates of different demographic groups are heterogeneously affected by monetary shocks. We show that *inflation inequality*, defined as the cross-sectional standard deviation of the decile-specific inflation rates across expenditure, salary, and income deciles, decreases after a contractionary monetary shock. The reason is that households at the bottom of the distribution are exposed to a higher inflation rate which tends at the same time to decrease more following a monetary shock. Finally, we find that the increase in expenditure inequality in response to monetary shocks is significantly more muted once inflation heterogeneity is taken into account.

The first contribution of this paper is to evaluate how monetary policy influences the distribution of household-level inflation rates. To compute individual inflation rates, we combine item-level price data from the Bureau of Labor Statistics (BLS) with individual expenditure data from the Consumer Expenditure Survey (CEX) for the U.S. from 1980 onward. We evaluate how the different moments of the inflation rates distribution, i.e., the median and the standard deviation, react to monetary policy shocks by adopting a Local Projection approach à la Jordà (2005). Exogenous variations in interest rate are captured using the Romer and Romer (2004) monetary shocks series. We document that contractionary monetary policy shocks decrease the median inflation rate as well as significantly reduce the dispersion of the distribution.

The second contribution is to assess which sectors are mainly responsible for the decrease in inflation dispersion. The price indexes of different sectors have different sensitivity to monetary policy shocks. We document that *Energy*, *Water* and *Gasoline* are by far the most influenced by contractionary shocks and they explain almost entirely the response of inflation dispersion to monetary shocks even though they account for only a relatively small expenditure share.

The third contribution is to study whether the inflation rates of different demographic groups are heterogeneously affected by monetary policy. We demonstrate that contractionary shocks lead to a sizable decrease in inflation inequality. On the one hand, the inflation rates of low- and middle-income households tend to be higher than that one of high-income households. On the other hand, it is more reactive to shocks and therefore decreases relatively more after a monetary shock. The same result holds for salary and expenditure deciles, confirming the important role of endowments in the dynamics of individual inflation rates.

The fourth contribution of the paper is to evaluate how these new findings on inflation heterogeneity influence real expenditure inequality and its response to monetary shocks. We compute two measures of real expenditure at household-level: one deflating nominal expenditure by the aggregate price level (as is common in the literature, neglecting inflation heterogeneity) and one deflating each expenditure category by the relative sectoral price level. As expected, we find that assuming all households are exposed to the same inflation rate overestimates the impact of monetary policy shocks on expenditure inequality. Although the nominal expenditure of low- and middle-income households decreases more after a shock compared to that of high-income households, their inflation rates also decrease relatively more, partially offsetting this decrease in real terms. It is important to underline that real consumption heterogeneity is still found to increase after a monetary shock corroborating again the evidence of the sizable distributional effects that central banks can have on the economy.

After years of relatively stable and low price growth, inflation rates worldwide have reached historically high levels in the post-Covid period. Which are the optimal monetary policies to be implemented to tackle it are again at the center of attention for academics and policymakers. However, most of the discussion focuses on stabilizing the aggregate inflation rate. The results from this paper suggest that concentrating only on the overall inflation would miss the huge heterogeneity in inflation rates to which households are exposed.

The level as well as the sensitivity of household-level inflation rates to changes in interest rate are strongly correlated with demographic characteristics. Therefore, abstracting from also considering how the individual inflation rates adjust in response to shocks would lead to systematic biases by the monetary authorities against specific demographic groups. For instance, since low-income households experience a higher inflation rate relative to high-income households, they would benefit from a more aggressive monetary policy than the one implemented by focusing only on the aggregate inflation rate. This problem could even be exacerbated by the fact that central banks usually design their policies targeting a specific subset of the price indexes. As we document, core measures of inflation, i.e., excluding energy and food, greatly underestimate the overall level of inflation dispersion in the economy. Finally, the empirical findings we provide suggest that central banks should pay close attention to inflation heterogeneity as whether it is taken into account or not has important implications for the magnitude of the distributional effects caused by the monetary authorities' decisions.

Related literature. This paper contributes to two strands of the literature. The first one is the research agenda on inflation inequality. Households are exposed to different levels of price increases given the heterogeneous consumption baskets they consume. For the U.S., Thesia et al. (1996), Hobijn and Lagakos (2005), Leslie and Paulson (2006), Johannsen (2014), and Orchard (2022) measure inflation inequality using the CEX data which covers the full consumption basket. More recently, Kaplan and Schulhofer-Wohl (2017), Argente and Lee (2021), and Jaravel (2019) compute inflation inequality from scanner data which are available for a much more limited time period but provides information at a higher level of granularity. The differences in inflation rates across households have been found to be substantial over time as well as related to demographic characteristics. For instance, high-income households are exposed to lower inflation rates compared to low- and middle-income households. See Jaravel (2021) for a review of the growing literature on inflation inequality.

Particularly related to the results of our paper, Cravino et al. (2020) show that the inflation rate of high-income households reacts significantly less than that of middle-income households following a monetary shock. We contribute to this literature by studying how inflation dispersion across households responds to monetary policy shocks. We document that contractionary shocks decrease the cross-sectional dispersion in household inflation rates. Almost the entire effect is due to the higher sensitivity of the prices of *Energy*, *Water*, and *Gasoline* to changes in the interest rate. Combining two results from the existing literature regarding the fact that lower- and middle-income households are exposed to a higher inflation rate, as documented by Kaplan and Schulhofer-Wohl (2017) and Jaravel (2019), and that at the same time, their inflation rate decreases relatively more following a monetary shock, as shown in Cravino et al. (2020), we find that inflation inequality across income, salary, and expenditure deciles decrease in response to a monetary shock.

The second strand is the growing literature on the distributional aspects of monetary policy. With an approach analogous to the one we adopt, Coibion et al. (2017) document that consumption and income inequality in the U.S. increase following a contractionary monetary shock. Similar findings have also been found in other countries and in different time periods, e.g., Mumtaz and Theophilopoulou (2017) for the United Kingdom, Guerello (2018) and Samarina and Nguyen (2023) for the Euro Area, Furceri et al. (2018) for a panel of 32 advanced and emerging economies. A summary of the current empirical and theoretical literature on the relationship between monetary policy and inequality is provided by Colciago et al. (2019).

We show that neglecting inflation heterogeneity results in an overestimation of the impact of monetary policy shocks on expenditure inequality. In response to a contractionary monetary shock, the stronger decrease in the inflation rate of low-income households partially offset the decrease of their nominal consumption resulting in a more muted response in real terms. It follows that the distributional effects of monetary policy on expenditure inequality are more limited once inflation heterogeneity is taken into consideration.

Road map. The paper is structured as follows. Section 2 describes the dataset used, as well as the construction of individual inflation rates and dispersion measures. In Section 3 we discuss the empirical strategy and show the main results in terms of the impact of monetary policy shocks on the cross-sectional inflation distribution. Section 4 studies the heterogeneous responses across different demographic groups. Section 5 evaluates how inflation heterogeneity influences the response of real consumption inequality to monetary shocks. In section 6, we perform a battery of different robustness checks to evaluate the reliability of our findings. Section 7 concludes.

2 Individual inflation rates

In this section, we compute individual inflation rates at the household level by exploiting the differences in consumption patterns across households. There are three steps needed for the computation of any inflation rate. First, we need information on prices for different goods. Second, we need detailed information on (individual) consumer expenditure, which allows computing the share of different goods in an aggregate index and therefore provides weights[1]. Third, statistical agencies have to decide on a methodology to combine price data to get a meaningful measure of inflation. In the following, we discuss each step separately.

2.1 Inflation data

We use data from the Consumer Price Index (CPI) as computed by the BLS at a monthly frequency. In particular, we use the not-seasonally-adjusted *US City Average for all urban consumers* (CPI-U). The BLS collects price data on 211 different subgroups of goods and services, which they call item strata. This is the most disaggregated level for which it publishes information on prices. However, these item strata over the period from 1980 to today undergo regular revisions or their definition is changed. Some disappear entirely and some get newly introduced. For this reason and for data availability we need to combine these basic price indices with more aggregate ones. We follow Hobijn and Lagakos (2005) and Johannsen (2014) in creating 21 indices, for which we get consistent inflation rates during our time sample. We call the inflation rates for subgroups of the consumer basket *inflation subindices*[2]. The construction of these inflation rates is subject to a tradeoff between consistent and sufficiently long time series and finely disaggregated time series that capture as much of the difference in inflation as possible. Jaravel (2019) finds that only 20% of inflation inequality is captured when using 22 expenditure categories instead of 256 for the period from 2004 to 2015. In subsection 6.1 we show that increasing the number of categories considered from 21 to 121 significantly increases the *level* of inflation dispersion across households but does not affect its *sensitivity* to monetary policy shocks.

In Table 1 we report the mean, median, standard deviation, the 10th and the 90th percentile of the 21 inflation subindices we compute, as well as of the Official CPI-U for the period 1980-2008. The observed sectoral inflation heterogeneity will be one of the key components in explaining the evolution of inflation dispersion. Households spend different

[1] The CEX proves rich enough to provide data on expenditure, going back to 1980.

[2] The list and definitions of these subindices can be found in Appendix A.1.

shares of their overall expenditure on each category and, since these categories differ in terms of price volatility and price level, this will lead to differences in terms of experienced inflation[3]. In what follows, we have to find reliable weights with which we can combine the inflation subindices to get household-level inflation rates across all items.

2.2 Expenditure data

For the computation of expenditure weights, we use the CEX provided by the BLS. This is the same dataset that is used to compute the official CPI of the U.S. The CEX is a quarterly survey of household expenditures and is divided into a diary and an interview survey. The diary survey covers small expenditures on daily items over a period of two weeks. The interview survey is more comprehensive, with detailed questioning every three months yielding up to a year of data for a single household. Since our goal is to get inflation rates that are as comprehensive as possible, we solely rely on data from the interview survey.

There are some limitations to the CEX data. The BLS removes consumption data from the 100th percentile (it is top-coded) to ensure anonymity. Additionally, since we deal with survey data, there are likely more measurement errors in the CEX compared to other data sources[4]. However, the CEX allows us to get a comprehensive picture of virtually all consumer expenditures and it is also sufficiently large in the time dimension (starts in 1980) and along the cross-section (roughly 5000-7000 households each wave).

Like the inflation subindices, we aggregate the expenditure data into 21 groups[5], matching the classification of the CEX with the one from the price indices. In the next step, we aggregate the household-level expenses from monthly to yearly. By doing this, we get rid of seasonal patterns in expenditures, while at the same time "averaging out" extraordinary expenses and hence improving the quality of our data. With this approach, almost the entire

[3]The biggest limitation of using inflation subindices is that they are not individual prices. While we capture the inflation that is due to different consumption baskets, we are not able to capture inflation differences within a subindex. It is conceivable that taking the category *Food away* as an example, high-end restaurants have different price developments from low-end ones. This problem is circumvented with Nielsen scanner data. The dataset reports product-level information on both prices and quantities so it is more granular than the CEX data. However, two major limitations made the Nielsen data a non-viable solution for our analysis. First of all, the data covers only purchases in department stores, grocery stores, drug stores, convenience stores, and other similar retail outlets which account for approximately 15% of total household expenditures. Moreover, the dataset is available only from 2004 onward.

[4]See Bee et al. (2013) for an assessment of the quality of our consumer dataset.

[5]In computing household-level inflation rates we have to alter the *Housing* group and omit the *Vehicle* group altogether. In particular, we follow Johannsen (2014) and we use the question on rental equivalence for the owned dwelling expenditures of the homeowners. Moreover, we exclude expenditures on new and used vehicles since in a given year the purchase of a vehicle could dominate all other expenditures. When we compute the inflation rate across deciles, vehicle purchases are included since it is less likely this category can bias the decile-level inflation rates. See Appendix A.3 and Appendix B for more details.

variation in individual inflation rates comes from price changes, rather than from changes in consumption patterns. Hence, the variation in individual inflation rates is mainly driven by the dynamics of sectoral inflation rates, as opposed to being driven by changes in the consumption bundle, as we intend. The relevance of the substitution effects is studied in subsection 6.1 where we compute the expenditure shares at higher frequencies.

2.3 Computation of individual inflation rates

In the third step, we combine the expenditure data with the inflation data. For this, we compute consumption shares w_j^i for household i and item subgroup j, which are calculated by dividing the yearly consumption expenditure in a certain period by the total expenditure reported in the same period. In the baseline analysis, we use all 21 categories. We compute the individual inflation rate for household i as:

$$\pi_{t-k,t}^{i} = \sum_{j \in J} w_j^i \pi_{j,t-k,t} \tag{1}$$

where j denotes the item subgroup as defined in section 2.2. The inflation rate of the subindex for good j in period t with base period $t-k$ is denoted by $\pi_{j,t,t-k}$. We set $k = 12$, meaning year-on-year inflation rates, which removes seasonality in the inflation subindices. Additionally, we winsorize the individual inflation rates at the 1st and the 99th percentile. In the next step, we analyze the statistical properties of individual inflation rates.

2.4 Properties of individual inflation rates

We assess the validity of the measures of individual inflation computed above by comparing the official CPI inflation rate with the median of individual inflation rates in Figure 1[6]. In the same figures, we also show different percentiles of the calculated household-specific rates of inflation.

The median of the distribution of household-specific rates of inflation closely tracks the headline value of CPI inflation. Hence, our approach gives, in an aggregate world, very similar results to the official CPI inflation rate. This result shows why for many years economic models mainly focused on the representative agent: The time series of the experienced inflation for the “median household” can be considered a quite good approximation of the aggregate economy.

[6] Similar results are obtained for the mean of the distribution.

Figure 1: Official CPI inflation, cross-sectional distribution, and median individual inflation rate over time

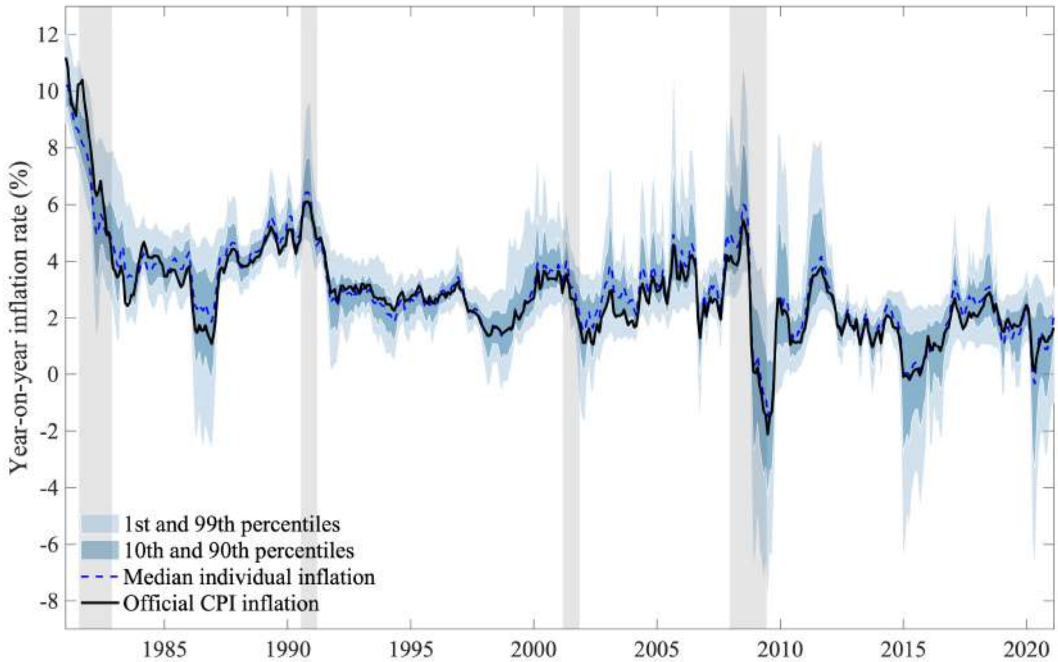

Notes: The plot shows the evolution over time of the official CPI inflation as well as the median and selected percentiles (1st, 10th, 90th, and 99th) of the winsorized cross-sectional distribution in individual inflation rates. The gray shaded areas depict U.S. recessions.

At the same time, the individual inflation rate percentiles in Figure 1 reveal how much information is lost when ignoring the heterogeneity across households. Not surprisingly, macroeconomic models have been expanded to include heterogeneity in consumption, wages, asset portfolio composition, and many more. However, most models still abstract from inflation differences and implicitly assume that households are exposed to the same inflation rate. Figure 1 strongly rejects this assumption.

2.5 Measures of dispersion

To evaluate how monetary policy shocks affect inflation dispersion in the U.S., we construct three different measures of dispersion: the cross-sectional standard deviation, the difference between the 90th percentile and the 10th percentile (depicted as 90th-10th, henceforth), and the cross-sectional interquartile range (IQR). To avoid the change in the survey composition affecting our results, we calculate the variation in the inflation dispersion measures on the households present in both periods. Therefore, when we calculate the change in the cross-sectional standard deviation from t to $t+1$, we do it only for the households which are present during both periods. Sampling weights are applied throughout the analysis.

Figure 2: Historical series of inflation dispersion measures

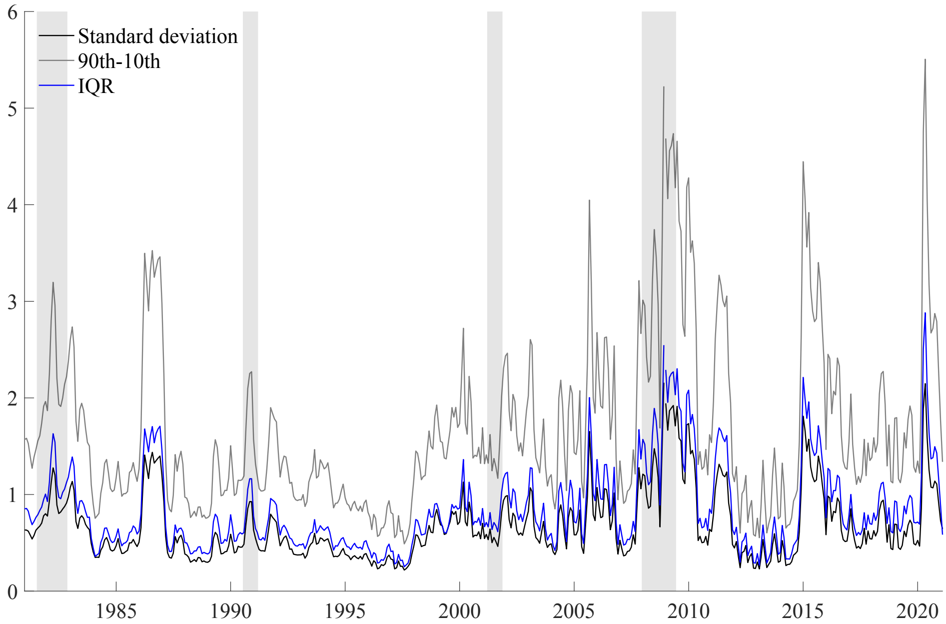

Notes: In the plot, we show the evolution of inflation dispersion measured using the cross-sectional standard deviation, the difference between the 90th and the 10th percentile of the cross-sectional distribution, and the IQR. All the series refer to the period 1981M1:2020M12. The gray shaded areas depict U.S. recessions.

Figure 2 shows the historical evolution of the three measures of dispersion, together with U.S. recessions. The three variables are highly correlated, suggesting that a normal distribution approximates the computed individual inflation rates very well. Despite using a different time period and alternative CPI categories, the time series are comparable in magnitude to those found by Johannsen (2014). As one can notice, inflation dispersion tends to increase during U.S. recessions suggesting a sort of correlation with the business cycle in the economy.

3 The effects on inflation dispersion

In this section, we present the results of our empirical analysis. We first study whether and to what extent monetary policy shocks influence aggregate inflation dispersion. We then investigate which expenditure categories drive the main results of our analysis.

3.1 Methodology

In the baseline specification, we adopt the Local Projection (LP) method developed by Jordà (2005). As in Cravino et al. (2020), we estimate a series of regressions for the dependent variable over different horizons on the monetary policy shock in period t and controlling for the lags of the shock as well as of the dependent variable:

$$x_{t+h} - x_t = c_h + \beta_h e_t^{RR} + \sum_{j=1}^{J} \theta_{h,j}(x_{t+1-j} - x_{t-j}) + \sum_{i=1}^{I} \gamma_{h,i} e_{t-i} + \epsilon_{t+h}, \tag{2}$$

where x is the variable of interest and the monetary policy shocks are denoted by e_t^{RR}. In line with the literature, we include 48 lags of the shocks and 6 lags of the dependent variable as control. The coefficient β_h for $h = 1, ..., H$ gives the response of the dependent variable at time t+h to a monetary policy shock at time t[7]. The impulse responses are computed over a horizon of 48 months using data from 1980M1 to 2008M12. Standard errors are corrected as in Newey and West (1987). For each impulse response, we present the one and 1.65 standard deviation confidence intervals. Unanticipated changes in the short-term interest rate are identified using the monetary policy shock series devised by Romer and Romer (2004, henceforth called R&R shocks) and extended by Coibion et al. (2017)[8].

The R&R shocks stop before 2009 so the zero lower bound period is excluded. In Appendix D we perform some additional analysis using as an alternative measure of monetary shocks the proxy from Bauer and Swanson (2022) which spans from 1988 to 2019. The main results of the paper hold considering the most recent period as well.

3.2 Analysis

We evaluate the overall effects of a contractionary monetary policy shock on inflation dispersion by estimating equation (2) using the cross-sectional standard deviation as the baseline measure of inflation dispersion[9]. The results are reported in Figure 3. The top panel shows the responses of the annual inflation rate computed by the BLS (blue line) as well as of the median inflation rate across households (black line): following a contractionary shock, the annual rate decreases by approximately 1.5 percentage points, a magnitude in line with the literature. As one might have expected looking at Figure 1, the response of the median inflation rate closely matches the response of aggregate inflation.

[7] As an alternative specification, we also use the R&R shocks as an instrument for the change in interest rate (IV-LP) instead of directly inserting them in the LP and the results remain basically unchanged.

[8] Coibion (2012) shows how the Romer and Romer (2004) approach might be particularly sensitive to the period in which the Federal Reserve abandoned targeting the federal fund rate between 1979 and 1982. Therefore, in Section 6 we redo the analysis starting the sample in 1985, and showing that our results are not driven by these large monetary policy shocks in the early 80s.

[9] The responses for the difference between the 90th and the 10th percentile of the cross-sectional distribution and the IQR are reported in Figure 15. Given the very high correlation among dispersion measures, the IRFs display similar patterns differing mainly in the magnitude of the response.

Figure 3: Impulse responses of the year-on-year inflation rate as well as the median and the standard deviation of the individual inflation rate distribution

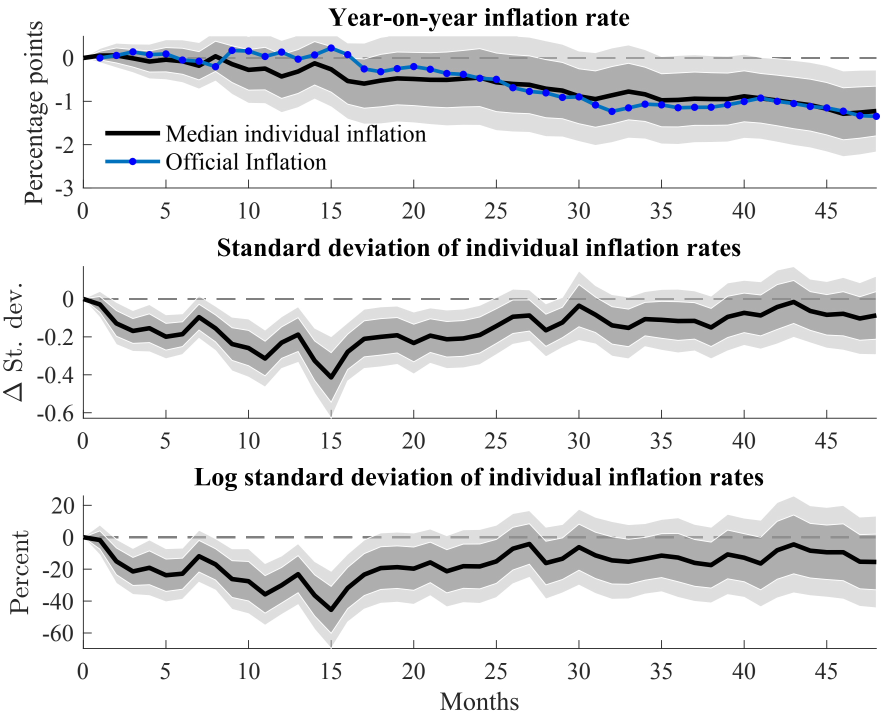

Notes: In the top panel the figure plots the impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the official annual inflation rate (black line) and the median inflation rate (blue line) of the individual inflation rate distribution. The middle panel reports the impulse response using as the dependent variable the dispersion in inflation, measured by the cross-sectional standard deviation and the bottom panel the log of the dispersion measure such that it can be interpreted as a percent change relative to the steady state. The horizontal axis is in months. Impulse responses are computed at a monthly frequency using data for the period 1980M1:2008M12.

In the middle panel, we show the impulse response of our dispersion measure. Inflation dispersion decreases after a contractionary monetary policy shock and remains persistently below zero. Looking at the one and 1.65 standard deviation confidence intervals we can easily reject the null hypothesis that the coefficients are equal to zero for the horizon considered. Therefore, the impulse response strongly suggests that monetary policy shocks lead to a decrease in the inflation dispersion in the economy.

To quantify the magnitude of the decrease in the inflation dispersion, the bottom panel computes the same impulse response but uses the log of the dispersion measure as the dependent variable, such that the magnitude can be interpreted as a percentage change relative to the steady state. Following a contractionary shock, we find that the cross-sectional standard deviation of inflation rates at the household level decreases by around 40% after 15 months and approximately 20% at the end of the horizon considered. The average inflation rate over the same time period is about 3.75% so a decrease of 1.5 percentage points corresponds to a decrease in 60% of the average value.

3.3 Sectoral contribution

The individual inflation rates are constructed assuming there is no substitution across categories in response to a monetary shock[10]. Therefore, the decrease in inflation dispersion is entirely due to the fact that the inflation of different sectors is heterogeneously sensitive to exogenous changes in the interest rate. To evaluate which sectors are mainly responsible for the results documented in the previous sections, we compute the response of several sectoral inflation rates to a contractionary shock. The results are reported in Figure 4.

Figure 4: Sectoral inflation rates impulse responses

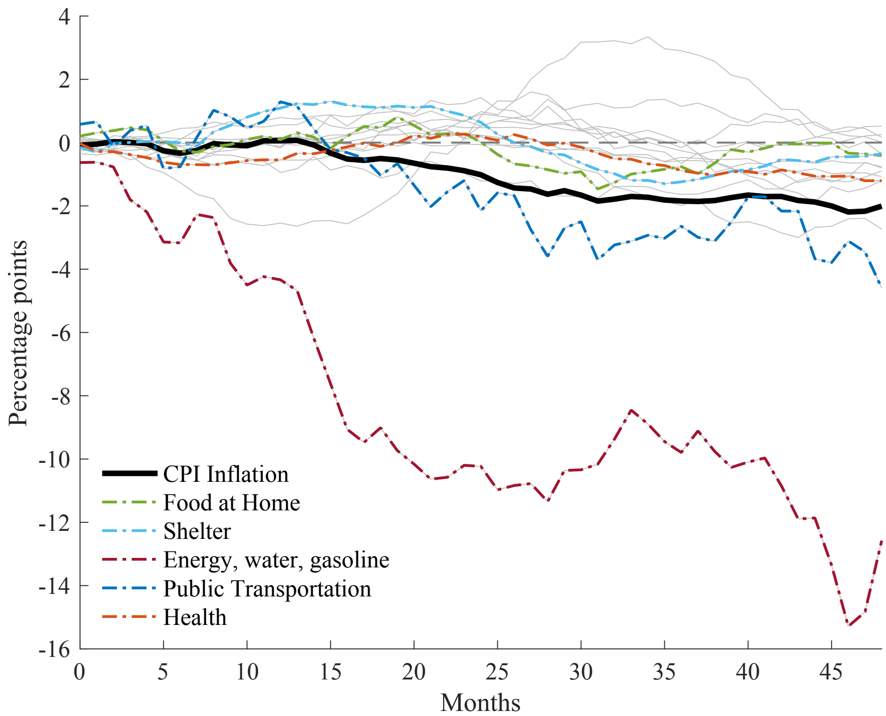

Notes: The figure plots the impulse responses of some of the different sectoral inflation rates that compose the Official CPI inflation (thick black line) to a one percentage point contractionary monetary policy shock. Impulse responses are computed at a monthly frequency using data relative to the period 1980M1:2008M12

The impact of monetary shocks on the inflation rates is extremely heterogeneous across sectors in line with the empirical evidence from Boivin et al. (2009) and Duarte and Dias (2019). Comparing the sectoral responses to the response of aggregate CPI it emerges that the majority of inflation rates at the sectoral level are only marginally affected by monetary

[10] Assumption which we relaxed in subsection 6.1.

policy shocks. In contrast, the inflation rates of *Public Transportation* and *Energy*, *Water* and *Gasoline* are significantly more responsive. This result is in line with Ider et al. (2023) which estimate a Bayesian Proxy SVAR model for the U.S. (1990-2019) and the Euro Area (1999-2019) and document that the response of the energy component of inflation to a monetary shock is ten times larger compared to the response of the headline consumer price index. Why the price indexes of some categories are more sensitive than others to monetary shocks is beyond the scope of this paper but we can expect it to be related to several factors like the different levels of price stickiness, labor intensity, etc.

Having shown that the sectoral inflation rates heterogeneously respond to monetary shocks, we now assess the contribution of the different sectors to the decrease in inflation dispersion. We start by computing inflation rates at the household level considering only a subset of the overall consumption basket. In particular, we classify each category into *non-durables*, *durables*, or *services*. As before, we then derive the response of the inflation dispersion across households for these three sub-categories, defined as the cross-sectional standard deviation, to a contractionary monetary shock.

The results are reported in Figure 5. The inflation dispersions of the three sub-categories decrease after a contractionary shock. However, they remarkably differ in the magnitude of their responses. The standard deviation of *non-durables* categories is more reactive whether the standard deviations of *durables* and *services* are less responsive to the shock and barely significant. The observed differences in the responses clearly suggest that the main drivers of the decrease in inflation dispersion can be found within the *non-durables* categories.

Therefore, we compute the same cross-sectional standard deviation of individual inflation rates but excluding one important expenditure category at a time. The results of this exercise are shown in Figure 6. As one can notice, most expenditure categories like *Housing*, *Health expenditure* and *Transportation*[11] have only a marginal effect on our main results despite accounting for a significant share of the household consumption bundles[12].

The middle left plot reports the inflation dispersion response when we exclude the categories *Energy*, *Water*, and *Gasoline*. This new specification is close to the definition of Core CPI that the Federal Reserve Bank uses to decide which monetary policy to adopt. Not surprisingly given the results shown in Figure 4, removing three of the most volatile categories cancels out the response of inflation dispersion almost entirely.

[11] *Housing* is defined as the sum of *Rented Dwellings*, *Owned Dwellings* and *Other Lodging*. *Transportation* is equal to the sum of *Public Transportation* and *Other Vehicle Expenses*.

[12] We report the average expenditure weights across different deciles for income, salary, and expenditures in Table 2.

Figure 5: Impulse responses of inflation dispersion for different sub-categories of expenditure

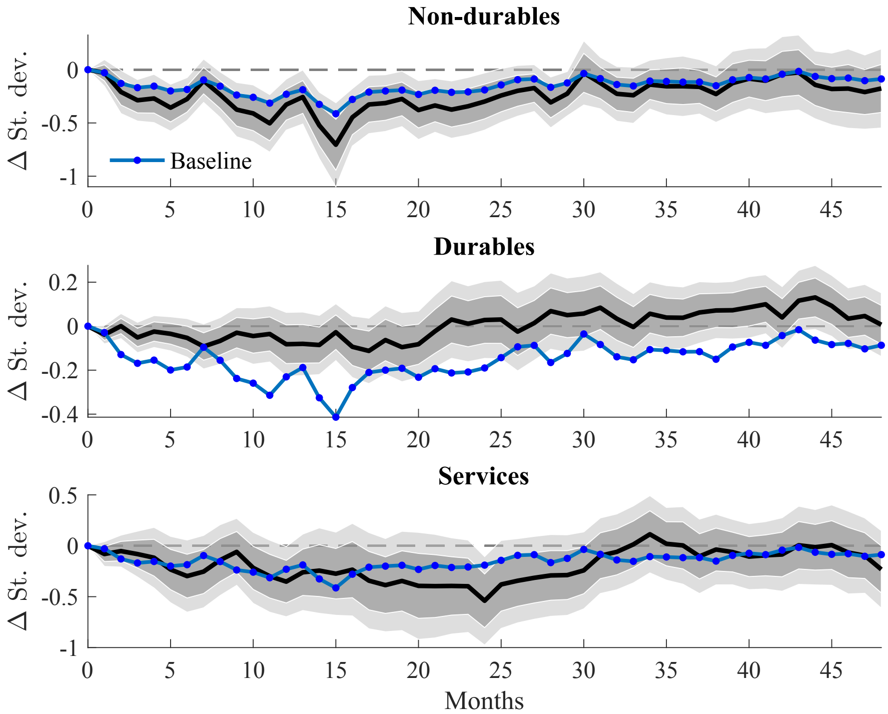

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the dispersion in inflation, measured by the cross-sectional standard deviation. The top panel uses the standard deviation in inflation rates for non-durable categories, the middle panel for durables, and the bottom panel for services. The solid blue line refers to the baseline impulse response obtained using the baseline categories. The horizontal axis is in months. Impulse responses are computed at a monthly frequency using data for the period 1980M1:2008M12.

To summarize, there is large heterogeneity in the contribution that each sector has to inflation dispersion. Many categories, even though being characterized by large expenditure share, have only a negligible impact. Most of the observed effects are due to the categories *Energy*, *Water*, and *Gasoline*. This empirical evidence suggests that central banks should not neglect the importance of these small and extremely volatile categories in setting their policy rate since most of the variation in inflation dispersion comes actually from them.

4 Heterogeneity across demographic groups

Having shown that monetary policy shocks decrease inflation dispersion in the economy, we now evaluate whether the inflation rate of some demographic groups is more sensitive to contractionary shocks relative to other groups and how this affects the cross-sectional inflation dispersion. We focus in particular on three demographic groups: income, salary, and expenditure deciles.

Figure 6: Impulse responses of inflation dispersion excluding different categories of expenditure

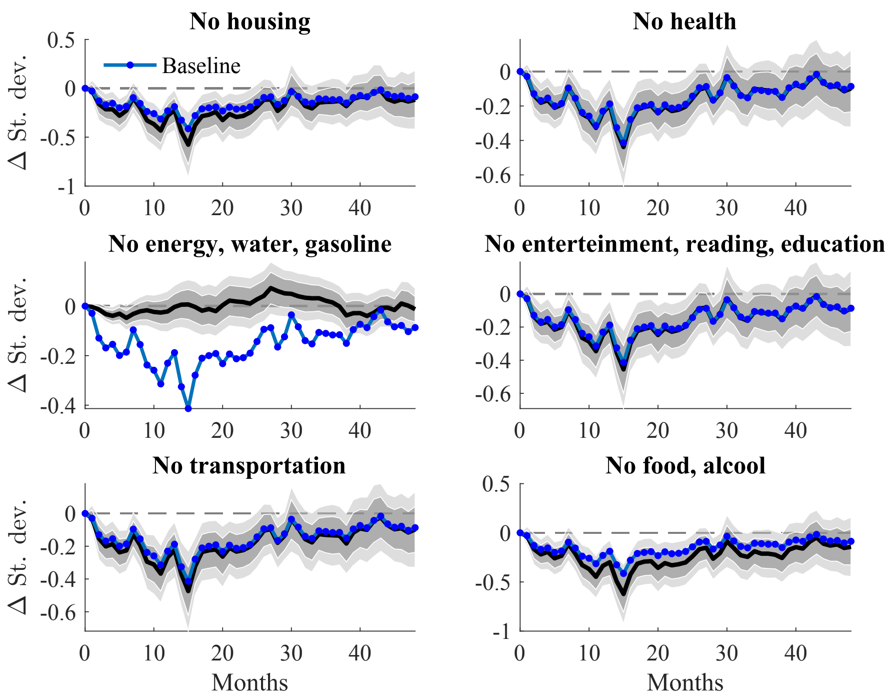

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the dispersion in inflation, measured by the cross-sectional standard deviation. Each panel uses the standard deviation in inflation rates computing excluding expenditure categories from the consumption bundle of the households. The solid blue line refers to the baseline impulse response obtained using the baseline categories. The horizontal axis is in months. Impulse responses are computed at a monthly frequency using data for the period 1980M1:2008M12.

4.1 Expenditure weights

Heterogeneity in inflation rates comes from the fact that households consume different consumption baskets. As in Cravino et al. (2020), we derive the time-varying decile-specific expenditure weights following the procedure used by the BLS to compute the aggregate CPI which we describe in detail in Appendix B[13]. We report in Table 2 the expenditure weights of the first, fifth, and tenth deciles of income, salary, and expenditure deciles for each of the 21 categories for the period 1980-2008.

Several interesting facts can be noticed: First, the pattern across deciles is quite similar for income, salary, and expenditures. This already anticipates that the decile-level inflation rates of these three categories will react in a consistent way to monetary policy shocks. Second, although the weight for most of the categories either decreases or increases from the first to the tenth deciles, some categories display a U-shape pattern, e.g., *Gasoline* and *Medical*

[13] Appendix D shows that the results are not particularly affected by considering the simple median inflation rate for each decile.

expenses. This is consistent with the findings of Cravino et al. (2020) who document that the highest price volatility is experienced by middle-income households. Finally, looking at the differences in weights across deciles, we can already anticipate the inflation rate of which deciles will be more sensitive to monetary shocks. In the previous section, we demonstrate that most of the variation in inflation dispersion comes from *Gasoline* and *Energy* and low- and middle-income households devote a significantly higher share of their income to these categories relatively to high-income households.

4.2 Impulse responses by demographic groups

We study how the inflation rates of different demographic groups react to monetary policy shocks. We start by estimating the LP with R&R shocks using as the dependent variable the cross-sectional standard deviation of the decile-specific inflation rates across income, salary, and expenditure deciles which we define as *inflation inequality*[14]. As one can see from Figure 7, following a contractionary monetary policy shock inflation inequality for the three groups significantly decreases.

To better understand the main drivers of this result, we compare the median inflation rates of the different income, salary, and expenditure deciles with their impulse responses over time. The black lines in Figure 8 report the cross-sectional distribution of the impulse responses for the inflation rate of the different income (left panel), salary (middle panel), and expenditure deciles (right panel) 24 and 48 months after a one-percentage-point contractionary monetary policy shock.

Similar to what Cravino et al. (2020) find for income, the annual inflation rate of the households at the top of the income distribution reacts substantially less to monetary policy shocks than the one of those in the middle. The difference between middle- and high-expenditure households is economically sizable and statistically significant as tested in Appendix C. After 24 months, the annual inflation rate of the households in the top decile responds to around 40% less than the inflation rate of the households in the fifth decile. After 48 months, the difference is still around 25%.

How does this relate to inflation inequality? We report in the same panels the median inflation rates across deciles relative to the time period considered (red line, left axis)[15]. One can notice how the higher the decile the lower the median inflation rate. This result is

[14] Appendix B explains in detail how the median inflation rates are computed following the same approach adopted by the BLS.

[15] Plotting the cumulative difference in inflation rates across deciles delivers similar results.

Figure 7: Impulse responses of inflation dispersion across income, salary, and expenditure deciles

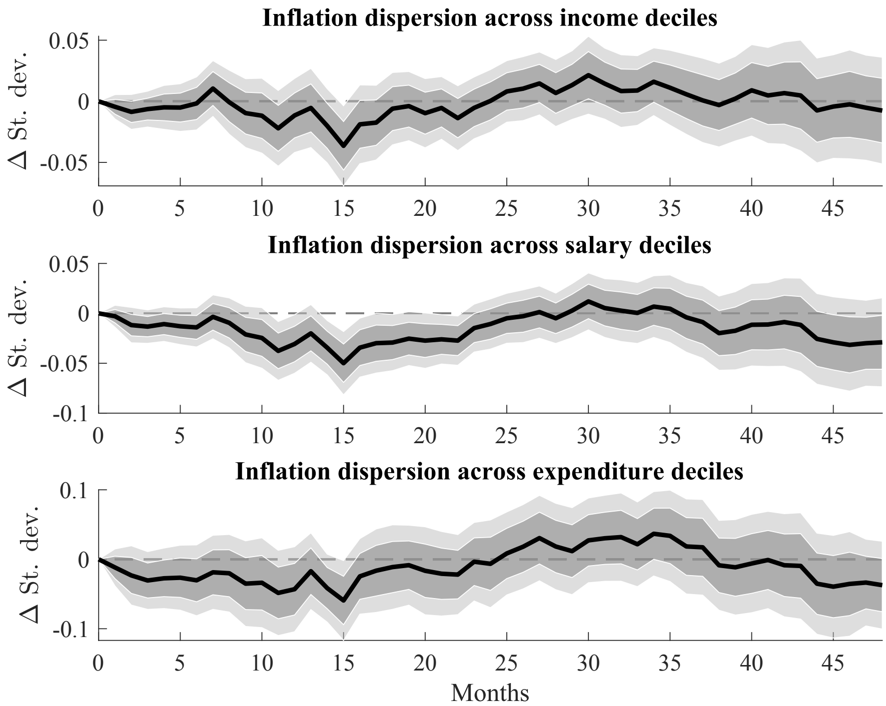

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for inflation inequality across income (top), salary (middle), and expenditure deciles (bottom). Inflation inequality is measured using the cross-sectional standard deviation of the decile-specific inflation rates. The horizontal axis is in months. Impulse responses are computed at a monthly frequency using data for the period 1980M1:2008M12.

consistent with the evidence provided by Jaravel (2019) and Kaplan and Schulhofer-Wohl (2017) using the Nielsen scanner data.

On the one hand, given their consumption bundle, high-income households experience a lower median inflation rate than the households on the left side of the distribution. On the other hand, their inflation rate reacts significantly less to monetary policy shocks. These two results combined imply that following a contractionary shock, we observe a convergence of individual inflation rates across the distribution leading to a lower inflation inequality as documented in Figure 7. Similar results can be found focusing on salary and expenditure deciles as shown in the middle and right panels of Figure 8.

Our empirical analysis strongly suggests that monetary policy shocks can have significant and non-negligible distributional effects on the economy. The median inflation rate of higher-income households is lower relative to low- and middle-income deciles. At the same time, their inflation rate is less reactive to unexpected changes in the interest rate. This results in a decrease in inflation inequality following a contractionary shock.

Figure 8: Impulse responses of the decile-specific inflation rate across income, salary, and expenditure deciles

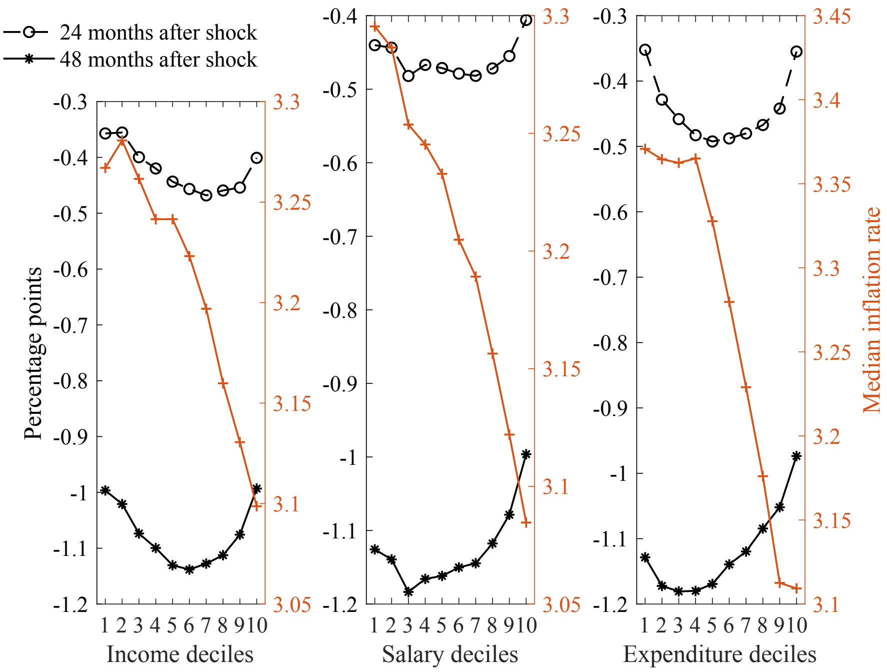

Notes: The figure reports the cross-sectional distribution of the decile-specific inflation rate responses of the different income (left panel), salary (middle panel), and expenditure deciles (right panel) 24 and 48 months after a one-percentage-point contractionary monetary policy shock. The red lines refer to the median inflation rate across deciles (left axis). Impulse responses are computed at a monthly frequency using data for the period 1980M1:2008M12.

5 Real expenditure inequality

Does the identified inflation inequality have any effect on the estimated impact of monetary shocks on real expenditure inequality? To answer this question, we follow Coibion et al. (2017) as close as possible and compute a broad measure of household expenditure which includes non-durables, durables, and services[16]. Few expenses are excluded since the relative sub-category price index is not easily identifiable, e.g., occupational expenses, mortgage, and property taxes.

To evaluate the role played by inflation inequality, we create two different series for real expenditure. In line with the literature, one is created by deflating each category by the aggregate CPI-U. The other one is obtained by deflating each item group by its relative price index. We then aggregate the expenditures at quarterly levels to reduce sampling error and to avoid having unusual purchases bias the analysis. We also winsorize at the bottom and

[16]In particular, the categories considered are: Food at Home, Food Away, Alcohol at Home, Alcohol Away, Apparel, Gasoline, Personal Care (services and durables), Reading, Tobacco, Household Furnishings and Operations, Energy, Water, Other Lodging, Public Transportation, House expenditures (services and durables), Rental expenditures (services and durables), Rent paid, Heath insurance, Health expenditures (services and durables), Education, Vehicles purchase, Vehicle expenditures (services and durables), Miscellaneous.

top 1 percent of the distribution. Expenditure inequality across households is computed as the cross-sectional standard deviation of log levels, the Gini coefficient of levels, and the difference between the 90th percentile and the 10th percentile of log levels. Finally, all series are seasonally adjusted.

Inequality is defined as $Ineq_t^{IH}$ and $Ineq_t^{NoIH}$ respectively for when inflation heterogeneity is taken into account by deflating each category by the relative price index and for when it is neglected. As an example, the standard deviations at time t across households i are equal to $Std(logC_{i,t}^{IH})$ and $Std(logC_{i,t}^{NoIH})$ with:

$$C_{i,t}^{IH} = \sum_{j \in J} \frac{C_{i,j,t}}{P_{j,t}}, \qquad C_{i,t}^{NoIH} = \sum_{j \in J} \frac{C_{i,j,t}}{P_t}, \tag{3}$$

where $C_{i,j,t}$ is the nominal consumption of household i relative to category j at time t, $P_{j,t}$ is the price index of the category j at time t and P_t is the aggregate price index.

To make our results as comparable as possible, we use the same econometric procedure adopted by Coibion et al. (2017), i.e., local projection with Romer and Romer (2004) shocks at a quarterly frequency, over the same time period, 1980Q1:2008Q4[17]. Since the series is quarterly, we include as controls 20 lags for the shocks and 2 lags for the dependent variable and we compute the impulse responses over 20 quarters.

Figure 9 plots the results. The black solid lines report the impulse responses of the three measures of expenditure inequality obtained by deflating the expenditure categories by the aggregate CPI. The shape and the magnitude of the responses are very close to those obtained by Coibion et al. (2017). After a contractionary monetary policy shock, expenditure inequality persistently and significantly increases.

However, neglecting inflation heterogeneity across consumption baskets leads to an overestimation of the overall effect. As shown by the red solid lines which report the responses of the expenditure inequality measures obtained by deflating each category by their respective price index, when the expenditure categories are properly deflated, the estimated effect of monetary policy on inequality is approximately 20% lower for standard deviation and 30% for the Gini coefficient and the 90th-10th percentile difference. It is worth mentioning that the estimated coefficients are still positive and significant which implies that monetary policy still has redistributive effects on the economy.

[17]Similar results are obtained adopting our empirical.

Figure 9: Impulse responses of expenditure inequality

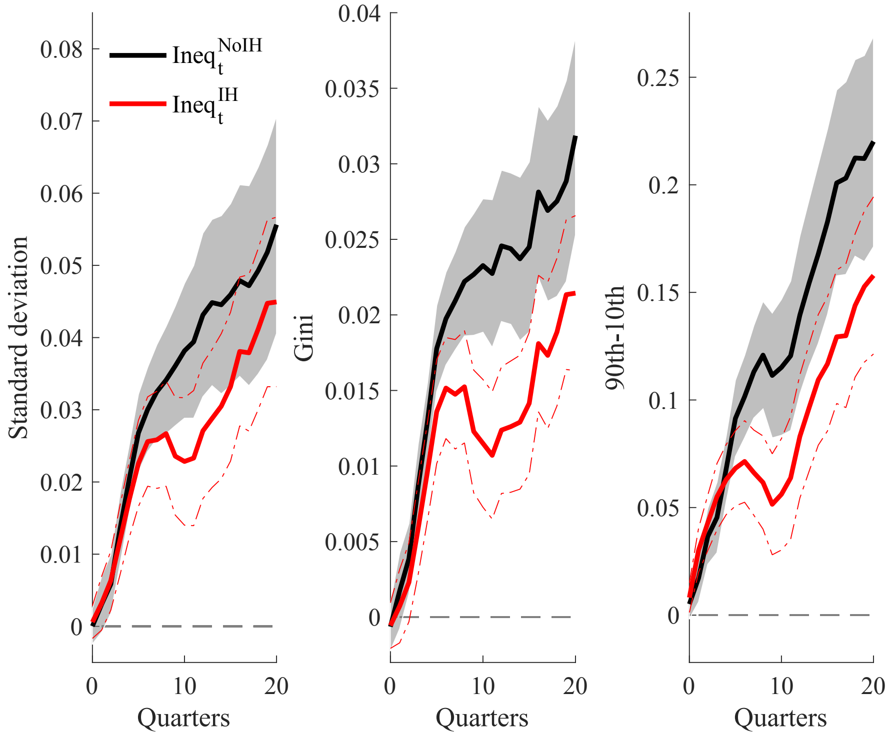

Notes: The figure plots impulse responses to a one percentage point contractionary monetary policy shock, as well as one standard deviation confidence intervals for expenditures inequality. The horizontal axis is in quarters and inequality is measured using the cross-sectional standard deviation (left), Gini coefficient (middle), and the log difference between the 90th and 10th percentiles of the cross-sectional distribution (right). The black solid line and the dark grey shaded areas depict the impulse response obtained by deflating the expenditure categories by the aggregate CPI, the red solid line and the dashed red lines refer to the impulse obtained by deflating each category by their respective price index. Impulse responses are computed at the quarterly frequency using data for the period 1980Q1:2008Q4.

This result can be explained by combining the new empirical evidence from the previous sections. Along the income distribution, a contractionary monetary shock has heterogeneous effects on nominal consumption. The nominal consumption of low- and middle-income households decreases more than that of high-income households because they are more sensitive to the monetary policy shock, e.g., they are financially constrained, they are more likely to lose their job in an economic downturn, etc. However, at the same time, the cost of their consumption basket decreases more strongly as well. Hence, the overall effect on expenditure is partially offset in real terms. This results in a more muted, but still positive and significant, increase in real expenditure inequality.

6 Robustness

To strengthen the validity of our findings in the previous sections, we show that our results are robust across a wide range of alternative specifications. First, we evaluate the importance of substitution effects. Second, we assess the sensitivity of our results to different lag specifications. Third, we perform the same analysis starting our sample in 1985M1 to control for the Volcker disinflation period. More robustness checks can be found in Appendix D. The figures are reported in Appendix E.

6.1 Substitution effects

Throughout the paper, we conduct our analysis under the assumption that differences in inflation dispersion are mainly driven by changes in prices and that variations in expenditure shares play only a marginal role. Both the inflation rate at household-level as well as at the decile level are computed using expenditure weights aggregated over multiple time periods to control for seasonal effects as well as to avoid unusual purchases by the households biasing our results. The weights for the household-level inflation rate rely on the entire time series of expenditure (maximum 12 months) whereas the weights at the decile level are computed following the BLS which updates its expenditure weight reference period approximately every ten years, and since 2002, every two years (more details can be found in Appendix B).

Cravino et al. (2020) tested whether substitution effects are important for the CEX by using the difference between the Laspeyres and Paasche price index as a proxy for the substitution bias from 1987 to 2004. These authors showed that the difference between the two indices is negligible over time demonstrating that the substitution bias must be very small.

Furthermore, using the Nielsen data, Jaravel (2019) evaluates whether the observed inflation heterogeneity along the income distribution stems from the fact that high-income households purchase different goods or whether they pay more for the same goods, for instance, because they buy from different shops. The inflation difference is then decomposed into a *between* and a *within* component. The former corresponds to the inflation difference that we would observe if households differ only in terms of the expenditure shares across categories and if they experience the same within-category inflation. Vice versa, the latter refers to the difference that would arise in case of households experience the same within-category inflation, but have different expenditure shares. The between component accounts for more than 70% of the inflation difference.

Given the importance for our results of the assumption that inflation dispersion is mainly driven by changes in prices rather than in expenditure shares, we also test whether substitution effects are a potential source of bias. We do this through two robustness checks: First, we assess if the granularity of the expenditure categories we choose plays any role. Second, we compute our measures of inflation inequality across deciles by using annual, quarterly, and monthly expenditure shares instead of using multiple years of consumption data like the BLS.

Following the literature, in computing the individual inflation rates we adopt a rather conservative aggregation in the number of categories considered. Not only do we have data for *Food and Beverage*, the most aggregate item category, but also have data for the sub-category *Eggs*, the most disaggregate. In choosing the baseline aggregation, we face a trade-off between using as disaggregate data as possible to fully capture inflation dispersion and the quality of the price index. Not all price series are available since the early 80s and this is true, especially for the most disaggregate goods and services.

We show that the main results are basically unaffected by increasing or decreasing the number of categories considered. We compute the household-level inflation dispersion using 14, 31, and 121 expenditure categories[18]. The evolution over time of the dispersion measures is reported in Figure 16.

The number of categories considered significantly affects the overall level of inflation dispersion. Relatively to the baseline inflation dispersion with 21 categories, the magnitude is slightly smaller with 14 categories and is slightly larger with 31. With 121 categories the cross-sectional standard deviation is almost twice as high compared to the baseline. However, the measures of inequality are extremely positively correlated. The correlation with the baseline specification is 0.97, 0.98, and 0.86 for the measures with 14, 31, and 121 categories respectively.

In Figure 17 we compare the response from our baseline specification with 21 categories (blue line) against the three alternative aggregations. When using price indices at a slightly more granular level (middle panel, 31 categories) or an even more conservative number of categories (top panel, 14 categories), the magnitude and the shape of the responses are

[18]For this last specification some of the price indexes were available later than 1980 so it is an unbalanced panel. The 14 categories are Food, Alcohol, Housing, Apparel, Gasoline, Other Vehicle Expenses, Public Transportation, Medical, Entertainment, Personal Care, Reading, Education, Tobacco, and Other Expenses. The 31 categories are Food at Home, Food Away from Home, Alcohol, Rental expenditures (durables), Rental expenditures (services), Rent Paid, Rent Equivalent, House Expenditures (durables), House Expenditures (services), Other House related expenses, Other Lodging, Energy, Water, Phone, Household Furnishings and Operations, Jewelry, Clothing (durables), Clothing (services), Gasoline, Vehicle Expenditure (durables), Vehicle Expenditure (services), Public Transportation, Medical, Entertainment, Personal Care (durables), Personal Care (Services), Reading, Education, Tobacco, and Other Expenses.

basically the same as that obtained in our baseline specification. Considering 121 categories the response is still significantly and persistently negative following a contractionary shock. The magnitude of the response is almost twice as much as the one of the baseline response but since the size of the inflation dispersion measure has doubled as well, in percentage terms the results are similar. This suggests that the number of categories considered in computing individual inflation rates is important for measuring the *level* of inflation inequality but not its *sensitivity* to monetary policy shocks.

As a second test for the role of substitution effects, we compute the expenditure weights for the decile-level inflation rates at annual, quarterly, and monthly frequencies. It is important to notice that by allowing the weights to vary at a much higher frequency than the biannual frequency adopted by the BLS in the last decades, our dispersion measures will not only capture potential adjustments in the consumption bundles due to the shocks but also measurement errors and unusual purchases will account for a larger share.

We report in Figure 18 the response of the cross-sectional standard deviation of the median inflation rates across income deciles as well as the one standard deviation confidence interval (black line and gray area). For comparison, the blue lines refer to the impulse response of the cross-sectional standard deviation as well as the relative confidence interval computed following the BLS methodology as shown in Figure 7.

Not surprisingly, moving from annual to quarterly and especially to monthly weights makes the responses more volatile. The responses with time-varying weights are clearly still negative and significant: inflation inequality across expenditure deciles remarkably decreases after a monetary shock. The magnitude is even more negative relative to the baseline. This might suggest that substitution effects move in the same direction as our inflation heterogeneity channel: following a contractionary shock, inflation rates of the expenditure categories purchased by low- and middle-income households decrease more strongly than the other categories so their overall inflation rates react more. Moreover, the same households might even increase their consumption of these categories since they are now relatively cheaper, leading to second-order effects. Similar evidence is found for the dispersions in median inflation across the salary and expenditure deciles whose responses are reported in Figure 19 and Figure 20 respectively.

Since we cannot further disentangle substitution effects from measurement errors in the survey or unrepresentative purchases made by households, we prefer to interpret these results with caution. Overall these findings confirm that substitution effects do not cancel out the impact of contractionary shocks on inflation dispersion and that heterogeneity in prices across, rather than within, expenditure categories is the main driver of our results.

6.2 Different lag specification

We re-estimate equation (2) with an alternative lag specification. In Figure 21 we run the LP regression including 36 and 60 lags for the monetary policy shocks as well as 4 and 8 lags for the cross-sectional standard deviation of the individual inflation. Similar results are also obtained for the other measures of dispersion. Increasing or reducing the number of lags has little to no effect on the impulse responses: after a contractionary monetary policy shock, inflation dispersion significantly decreases.

6.3 Volcker disinflation

Coibion (2012) shows how few episodes in the early 80s can be the main drivers of the impulse responses computed using LP with R&R shocks. Since then, it has been common practice for researchers to test their results excluding the period between 1979 and 1982 in which the Federal Reserve abandoned targeting the federal fund rate. Figure 22 reports the IRFs obtained using the baseline specification but starting the sample in 1985M1. In this case, the results are also robust.

7 Conclusion

Central bankers and policymakers are more and more strongly advocating the importance of the conduct of a more inclusive monetary policy where the potential negative spillovers deriving from the monetary authorities' decisions are taken into account. Similarly, macroeconomic research has shifted its focus from the aggregate effects of monetary shocks towards the different channels through which households and firms might be heterogeneously affected by it. Our results suggest that the inflation heterogeneity that arises from the different consumption baskets the agents purchase is of pivotal importance for understanding the distributional consequences of monetary policy.

In this paper, we study how monetary policy shocks affect the distribution of household-level inflation rates. We rely on individual expenditure data from the CEX and combine it with category-level inflation rates from the BLS to obtain household-level inflation rates. We compute different moments of the individual inflation rates distribution and we evaluate how monetary policy shocks influence the median and the cross-sectional standard deviation of the distribution. Inflation dispersion across households significantly and persistently decreases in response to a contractionary monetary policy shock. *Energy*, *Water* and *Gasoline* are found to explain almost entirely the observed effects despite accounting for a relatively small expenditure share.

We also evaluate how the inflation rate of different demographic groups is heterogeneously affected by monetary policy. We find that the inflation rates of low- and middle-income households are significantly more reactive to monetary shocks than that of high-income households. Since at the same time, they experience a higher median inflation rate, contractionary shocks lead to an overall convergence of inflation inequality across income groups. The same is true for expenditure and salary deciles.

Finally, we demonstrate that assuming that households are exposed to the same inflation rate results in an overestimation of the impact of monetary shocks on expenditure inequality. Following a contractionary shock, low-income households experience a stronger decrease in nominal consumption relative to high-income households. However, the price of their consumption bundles decreases relatively more as well partially offsetting the effect in real terms. Accounting for inflation heterogeneity reduces the estimated response of expenditure inequality to monetary shocks by around 20-30% depending on the measure of inequality considered.

In conclusion, our research provides substantial evidence that designing optimal monetary policies as well as studying their distributional effects cannot abstract from also considering the different inflation rates to which agents are exposed. Indeed, the economic agents experience significantly different inflation rates both in the long run as well as in response to shocks. Inflation heterogeneity in the economy is sizable and related to demographic characteristics. Therefore, focusing only on aggregate inflation or measures of inflation that exclude important components might lead to the implementation of systematically suboptimal policies for specific demographic groups. Finally, taking into account inflation heterogeneity is particularly relevant when it comes to assessing the impact of monetary policy on other forms of inequalities.

References

Argente, D. and Lee, M. (2021). "Cost of Living Inequality during the Great Recession". *Journal of European Economic Association.*

Bauer, M. D. and Swanson, E. T. (2022). "A Reassessment of Monetary Policy Surprises and High-Frequency Identification". *NBER Working Papers 29939, National Bureau of Economic Research.*

Bee, A., Meyer, B., and Sullivan, J. (2013). "The Validity of Consumption Data: Are the Consumer Expenditure Interview and Diary Surveys Informative?". In *Improving the Measurement of Consumer Expenditures*, NBER Chapters, pages 204–240. National Bureau of Economic Research, Inc.

Boivin, J., Giannoni, M., and Mihov, I. (2009). "Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data". *American Economic Review, March 2009, 99 (1), 350–84.*

Coibion, O. (2012). "Are the Effects of Monetary Policy Shocks Big or Small?". *American Economic Journal: Macroeconomics 2012, 4(2): 1–32.*

Coibion, O., Gorodnichenko, Y., Kueng, L., and Silvia, J. (2017). "Innocent Bystanders? Monetary policy and inequality". *Journal of Monetary Economics, 2017, vol. 88, issue C, 70-89.*

Colciago, A., Samarina, A., and de Haan, J. (2019). "Central bank policies and income and wealth inequality: A survey". *Journal of Economic Surveys, 33(4), 1199– 1231.*

Cravino, J., Lan, T., and Levchenko, A. (2020). "Price Stickiness Along the Income Distribution and the Effects of Monetary Policy". *Journal of Monetary Economics, 110:19-32.*

Duarte, J. B. and Dias, D. A. (2019). "Monetary Policy, Housing Rents, and Inflation Dynamics". *Journal of Applied Econometrics.*

Fisher, D. J., Johnson, S. D., and Smeeding, M. T. (2013). "Measuring the Trends in Inequality of Individuals and Families: Income and Consumption". *American Economic Review 103(3), 184-88.*

Furceri, D., Loungani, P., and Zdzienicka, A. (2018). "The effects of monetary policy shocks on inequality". *Journal of International Money and Finance, Volume 85, July 2018, Pages 168-186.*

Guerello, C. (2018). "Conventional and unconventional monetary policy vs. households income distribution: An empirical analysis for the Euro Area". *Journal of International Money and Finance, 2018, vol. 85, issue C, 187-214.*

Hobijn, B. and Lagakos, D. (2005). "Inflation inequality in the United States". *Review of Income and Wealth.*

Ider, G., Kriwoluzky, A., Kurcz, F., and Schumann, B. (2023). "The Energy-Price Channel of (European) Monetary Policy". *DIW Berlin Discussion Paper No. 2033.*

Jaravel, X. (2019). "The unequal gains from product innovations: evidence from the U.S. retail sector". *The Quarterly Journal of Economics, Volume 134, Issue 2, Pages 715-783.*

Jaravel, X. (2021). Inflation inequality: Measurement, causes, and policy implications. *Annual Review of Economics*, 13:599–629.

Johannsen, B. K. (2014). "Inflation Experience and Inflation Expectations: Dispersion and Disagreement Within Demographic Groups". *FEDS Working Paper No. 2014-89.*

Jordà, O. (2005). "Estimation and Inference of Impulse Responses by Local Projections". *American Economic Review, 95 (1), 161–182.*

Kaplan, G. and Schulhofer-Wohl, S. (2017). "Inflation at the household level". *Journal of Monetary Economics 91, 19-38.*

Leslie, M. and Paulson, A. (2006). "Constructing the Chicago Fed Income Based Economic Index–Consumer Price Index: Inflation Experiences by Demographic Group: 1983–2005". *Federal Reserve Bank of Chicago Working Paper, 2006.*

Mumtaz, H. and Theophilopoulou, A. (2017). "The impact of monetary policy on inequality in the UK. An empirical analysis". *European Economic Review, Volume 98, Pages 410-423.*

Newey, W. and West, K. (1987). "A Simple, Positive-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix". *Econometrica 55:703–708.*

Orchard, J. (2022). "Cyclical Demand Shifts and Cost of Living Inequality". *Working paper.*

Romer, C. D. and Romer, D. H. (2004). "A new measure of monetary shocks: Derivation and implications". *American Economic Review 94(4), 1055-84.*

Samarina, A. and Nguyen, A. (2023). "Does monetary policy affect income inequality in the euro area?". *Journal of Money, Credit and Banking.*

Thesia, G. I., David, J. S., and Kokoski, M. F. (1996). "An Experimental Consumer Price Index for the Poor". *Monthly Labor Review, 119 (1996), 32.*

A Data sources

In this section, we document in greater detail the data sources used and the properties of the underlying data.

A.1 Price Indices

Since individual inflation rates are a weighted average of sectoral price indices, Table 1 displays the CPI subindices used, as well as their respective statistical properties.

Table 1: Item-level CPI statistics

CPI series (Item Code)[19]	Mean	Median	Standard Deviation	p10	p90
Food at Home (SAF11)	3.05	2.72	1.84	1.01	5.60
Food Away from Home (SEFV)	3.36	3.05	1.41	1.99	4.61
Alcoholic Beverages (SAF116)	3.24	2.73	1.84	1.64	5.29
Rented Dwellings (SEHA)	3.94	3.60	1.53	2.46	6.15
Owned Dwellings (SEHC)	3.65	3.33	1.01	2.42	5.13
Other Lodging (MUUR0000SE2102-SEHB)	5.15	4.65	3.50	1.51	9.69
Energy (SAH21)	3.29	2.41	5.74	-3.19	10.82
Water (SEHG01)	5.34	5.23	2.38	2.82	7.79
Phone (SAE2)	-1.06	-1.08	1.70	-3.31	1.13
Household F&O[21] (SAH3)	1.43	1.34	1.77	-0.39	2.70
Apparel (SAA)	1.00	0.82	2.32	-1.83	4.49
Gasoline (SETB)	3.31	2.93	13.79	-13.63	20.98
Other Vehicle Expenses (SETC-SETD-SETE-SETF)	3.02	2.34	2.10	0.79	6.75
Public Transportation (SETG)	4.47	4.06	5.08	-0.93	9.54
Medical care (SAM)	5.72	4.82	2.21	3.45	9.01
Entertainment (SAR)	1.47	1.34	0.74	0.59	2.64
Personal Care (SAG1)	3.23	2.79	1.57	1.87	5.01

[20]The official series ID, as defined by the BLS, is a combination of "CUUR0000", which stands for the unadjusted CPI-U inflation rate for the whole US, and the Item Code, as shown in the table.

[21]Household Furnishings and Operations

Table 1: Item-level CPI statistics (continued)

CPI series (Item Code)[20]	Mean	Median	Standard Deviation	p10	p90
Reading (SERG)	3.64	3.36	2.50	0.86	7.01
Education (SAE)	2.40	2.40	0.96	1.10	3.70
Tobacco (SEGA)	7.56	7.11	6.08	2.27	12.75
Other Expenses (SEGD)	5.73	4.93	2.84	3.29	11.48
CPI-U (SA0)	**3.42**	**3.04**	**1.72**	**1.68**	**5.01**

A.2 Consumer expenditure survey data

In this section, we provide further details about the construction of the dataset we use in the empirical analysis. We download the raw data for the period 1980-2005 from the ASCII files available from the Inter-university Consortium for Political and Social Research (ICPSR) whereas from the year 2006 onward we use the data provided by the BLS. For each quarter, the Interview Survey is structured as follows: the expenditure data is recovered from the disaggregated MTAB files, income data is derived from the FMLY files and additional information regarding the households can be found in the MEMB files.

In line with the literature, we aggregate together expenditures about the same month which is reported in different interviews. Then, we drop households that report zero expenditure on food as well as those which report negative expenditure for categories that cannot be negative according to the data codebook, such as expenditure for elderly care. Respondents younger than 25 years and older than 75 are excluded. To correct for sample breaks caused by slight changes in the questionnaire (food at home (1982Q1-88Q1), food away from home (2007Q2), and personal care services (2001Q2)) we regress each expenditure series on a time trend and indicators for the corresponding sample breaks and then subtract the effect of the dummies from the original series. For all these transformations, we rely heavily on Coibion et al. (2017).

Finally, the CEX data started to include the imputed income in 2004. To impute income data before that year, we follow the approach adopted by Fisher et al. (2013) and Coibion et al. (2017): for households recording a bracketed range, we use the median point of the bracket. Furthermore, we estimate the remaining income observations by regressing income on a set of observable characteristics such as age, age squared, the reference person's gender, race, education, number of weeks worked full or part-time in the last 12 months, unadjusted family size, the number of children under 18, the number of people over 64, the number of earners at the annual level and with sampling weights as well as using fixed effects for the income reporting date. To account for the sampling uncertainty, we add residuals drawn randomly with replacement from the sampling distribution to the predicted values. We then trim values above the top-coding threshold at the top coding value.

We then calculate expenditure shares from the cleaned expenditure data, which constitute the weights used to calculate individual inflation rates. We find substantial variation in the weights that can be explained to a large part by either income, salary, or expenditure deciles. Table 2 shows the weights for the 1st, 5th, and 10th deciles.

Table 2: Expenditure weights for the first, fifth and tenth decile of income, salary, and expenditure

	Income deciles			Salary deciles			Expenditure deciles		
	1st	5th	10th	1st	5th	10th	1st	5th	10th
Food at Home	18.7	14.2	11.1	16.5	14.0	11.1	22.0	14.3	9.9
Food Away	7.2	7.5	7.3	7.7	7.6	7.2	8.0	7.3	6.9
Alcohol	1.0	1.1	1.2	1.1	1.2	1.2	1.1	1.1	1.1
Rented Dwellings	15.6	12.4	6.0	13.7	12.4	6.0	21.8	10.6	5.9
Owned Dwellings	15.4	17.1	22.6	14.5	16.8	22.8	6.5	19.3	22.6
Other Lodging	0.5	0.6	1.4	0.7	0.6	1.3	0.3	0.6	1.5
Energy	6.2	5.4	4.3	5.7	5.2	4.3	6.6	5.6	3.7
Water	0.9	1.0	0.9	0.9	0.9	0.9	0.9	1.0	0.8
Phone	3.4	3.0	2.3	3.2	3.0	2.3	3.8	3.0	2.1
Household F&O[22]	3.3	4.5	7.0	3.9	4.7	7.0	2.5	4.4	8.1
Apparel	4.0	4.3	5.6	4.4	4.6	5.7	3.7	4.2	5.7
Gasoline	4.2	5.3	4.4	5.0	5.6	4.5	4.3	5.4	3.8
Other Vehicle Expenses	4.3	6.8	7.2	5.5	7.2	7.3	3.7	6.8	7.0
Public Transportation	1.0	1.0	1.7	1.1	1.0	1.6	1.0	0.9	1.8
Medical	5.0	6.2	5.0	5.4	5.2	4.6	4.9	6.2	5.6
Entertainment	3.8	4.9	6.5	4.5	5.2	6.5	3.4	4.8	7.0
Personal Care	0.9	1.0	1.0	1.0	1.0	1.0	0.9	1.0	0.9
Reading	0.4	0.5	0.6	0.5	0.5	0.6	0.4	0.5	0.6
Education	1.6	0.8	2.3	2.1	1.0	2.4	1.3	0.9	2.9
Tobacco	1.7	1.4	0.6	1.7	1.4	0.6	2.3	1.3	0.6
Other Expenses	0.8	1.1	1.1	0.9	1.1	0.9	0.6	1.0	1.5

A.3 Matching of expenditure and inflation data

We match the expenditure categories with the respective price indices. Following Hobijn and Lagakos (2005), for the category *Other Vehicle Expenses* which does not have a perfect match with the available CPI sub-categories, we create the CPI index by combining the series that match this category (that is, SETC, SETD, SETE, and SETF). As sectoral weights, we use the average over the time period considered of the official weights provided by the BLS, as displayed in the table "Relative Importance in the CPI". Finally, since *Other Lodging* changed

[22] Household Furnishings and Operations

the name, we use *Lodging away from home* until 1997 (MUUR0000SE2102) and *Lodging while out of town* (SEHB) until the end of the sample. In all cases, the CPI series we use are the not-seasonally-adjusted *US City Average for all urban consumers* series.

Table 3: Matching between CEX expenditure category and CPI

BLS Expenditure Category	CPI Series (Item Code)
Food at Home	SAF11
Food Away from Home	SEFV
Alcohol	SAF116
Owned Dwellings	SEHC
Rented Dwellings	SEHA
Other Lodging	MUUR0000SE2102-SEHB
Energy	SAH21
Water	SEHG01
Phone	SAE2
Household Furnishings and Operations	SAH3
Apparel	SAA
Gasoline	SETB
Other Vehicle Expenses	SETC-SETD-SETE-SETF
Public Transportation	SETG
Medical	SAM
Entertainment	SAR
Personal Care	SAG1
Reading	SERG
Education	SAE
Tobacco	SEGA
Other Expenses	SEGD

B Decile-level expenditure weights

Before computing the decile-level expenditure weights, some adjustments need to be performed. In line with the literature and the BLS procedure, the expenditure weight for the owners' equivalent rent of primary residence is based on the following CEX question: "If someone were to rent your home today, how much do you think it would rent for monthly, unfurnished and without utilities?" The homeowners' answer to this question is stored in the variable RENTEQVX in the characteristics files.

Moreover, as we mention in the main text, vehicle purchases are likely to bias the estimated expenditure shares. Indeed, they are large in size and not representative of the usual household consumption bundle. Therefore, in line with Johannsen (2014), we drop this category when computing household-level inflation rates. Following Cravino et al. (2020), we include expenditures on used cars and trucks when computing the decile-level inflation but we reduce these spendings to half to reflect only the dealer value added.

Households are also interviewed a different number of times and for at most four consecutive quarters, which corresponds to twelve months' worth of spending information. However, this does not necessarily match the calendar year. To control for this, we compute the decile-based inflation rate closely following the BLS procedure as in Cravino et al. (2020). First, we sort households into deciles based on their annual income, salary, median, and mean expenditure. We then compute the average expenditure for each item category at every decile in the calendar year. For instance, a respondent interviewed in February will report personal consumption for January, but also for November and December of the previous year. Similar to what the BLS does for the computation of the official CPI, to account for the relative contribution of each household to the decile-mean value of a calendar year, we weight the consumption by the number of months a household reports expenditures during a calendar year (the BLS calls this variable MO_SCOPE).

We can then use the formula below to compute the average expenditure for each category j at each decile d. First, for household i at decile d, we aggregate over all the expenditures on good j during the calendar year. Second, the household total expenditure is weighted by the sampling weights, fwt, provided by BLS to make the survey sample representative of the U.S. population. Then, the weighted household expenditure is summed up at the decile level. Finally, to obtain the monthly average income spent on good j by decile d, we divide the annual weighted household expenditure for category j by the weighted number of months household at decile d reported expenditure during the calendar year. To annualize the average category expenditure at the decile level, it is sufficient to multiply the monthly average expenditure by twelve:

$$X_j^d = \frac{\sum_i fwt_i^d \sum_t c_{i,j,t}^d}{\sum_i fwt_i^d \; MO\_SCOPE_i^d} \times 12 \tag{4}$$

where fwt_i^d is the frequency weight for household i at decile d, $c_{i,j,t}^d$ refers to the annual consumption on category j by household i at decile d and $MO\_SCOPE_i^d$ identify the number of months per year household i reported its expenditure. The decile-level expenditure weight for category d can then be computed as:

$$w_j^d = \frac{X_j^d}{\sum_j X_j^d} \tag{5}$$

C Differences in responses across deciles

We evaluate whether the responses of the decile-level median inflation rates to a monetary policy shock are statistically different from each other. To do so, we estimate equation (2) using as dependent variable the difference between the inflation rate of the 10th and 1st decile of each group and the inflation rate of the 5th decile. The first column of Figure 10 reports the responses of the difference in median inflation rate for the 10th and the 5th decile, and the second column for the 1st and the 5th decile. The first row shows the responses for the differences across expenditure deciles, the second row for salary deciles, and the last row for income deciles.

As it can be noticed in Figure 8, both the median inflation rates of the 10th as well as of the 1st deciles of income, salary, and expenditures react much less to a monetary policy shock than the 5th deciles resulting in a positive and significant response of their differences. The U-shaped response across deciles is in line with what was found by Cravino et al. (2020) who document that the price volatility along the income distribution is hump-shaped with the households at the top of the distribution experiencing the lowest volatility (resulting in the flattest impulse response) and middle-income households being exposed to slightly more price volatility than lower-income households.

D Further robustness checks

As a further robustness check, Figure 11 reports the impulse responses excluding all U.S. recession periods from the analysis (1981M07:1982M11, 1990M07:1991M03, 2001M03:2001M11). The results remain qualitatively unchanged with respect to the baseline specification.

As a second set of checks, we assess whether our results are specific to the shock series we chose, i.e., Romer and Romer, 2004. The alternative measure of monetary shocks we use is the high-frequency proxy proposed in Bauer and Swanson (2022). The proxy is computed from

Figure 10: Differences in impulse responses across deciles

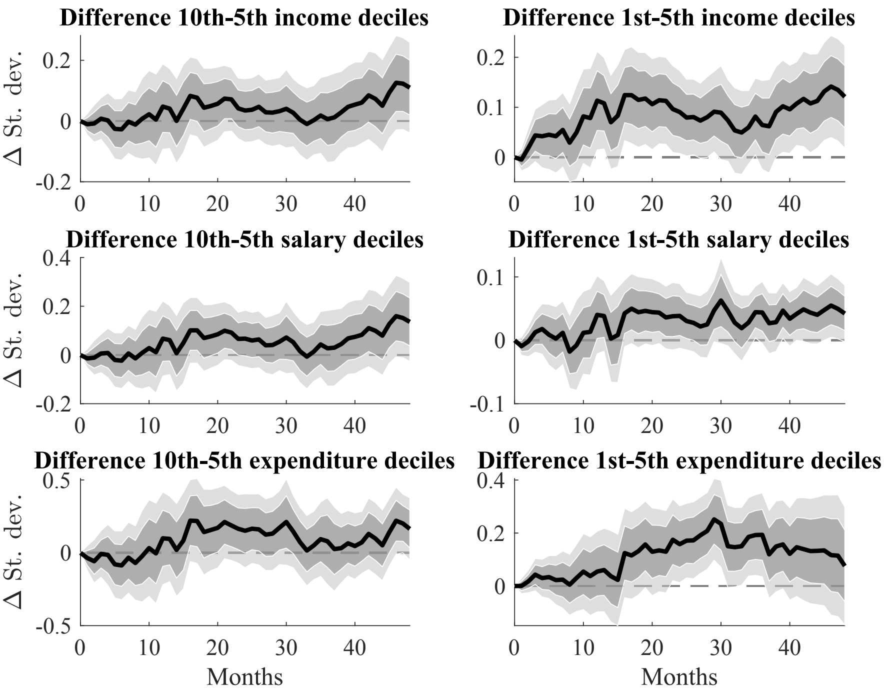

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the difference in decile-specific inflation rates across deciles of the demographic groups. The first column reports the responses of the difference in inflation rate for the 10th and the 5th decile, and the second column for the 1st and the 5th decile. The first row shows the responses for the difference across expenditure deciles, the second row for salary deciles, and the last row for income deciles. The horizontal axis is in months. Impulse responses are computed at a monthly frequency using data for the period 1980M1:2008M12.

changes in future prices in a narrow window around FOMC announcements and orthogonalized with respect to the public information about the economic and inflation outlook. The shock series is available from 1988 to 2019.

The results are presented in Figure 12. The top panel reports the response of the cross-sectional standard deviation to a contractionary shock and the bottom panel shows the response of inflation inequality across expenditure deciles. All the regressions include the same controls as in the baseline specification. In response to contractionary monetary policy shocks inflation dispersion as well as inequality decrease. Overall, the results from alternative monetary policy shocks confirm our main findings and point towards a distributional role played by monetary policy in terms of inflation dispersion.

Moreover, one might be concerned that part of the inflation heterogeneity we measured is driven by differences in consumption patterns across U.S. states rather than along the income distribution. Since the BLS does not provide price indices at the state level, but only at the division level (Northeast, Midwest, South, and West), we compute the cross-sectional

Figure 11: Impulse responses of inflation dispersion (without recession periods)

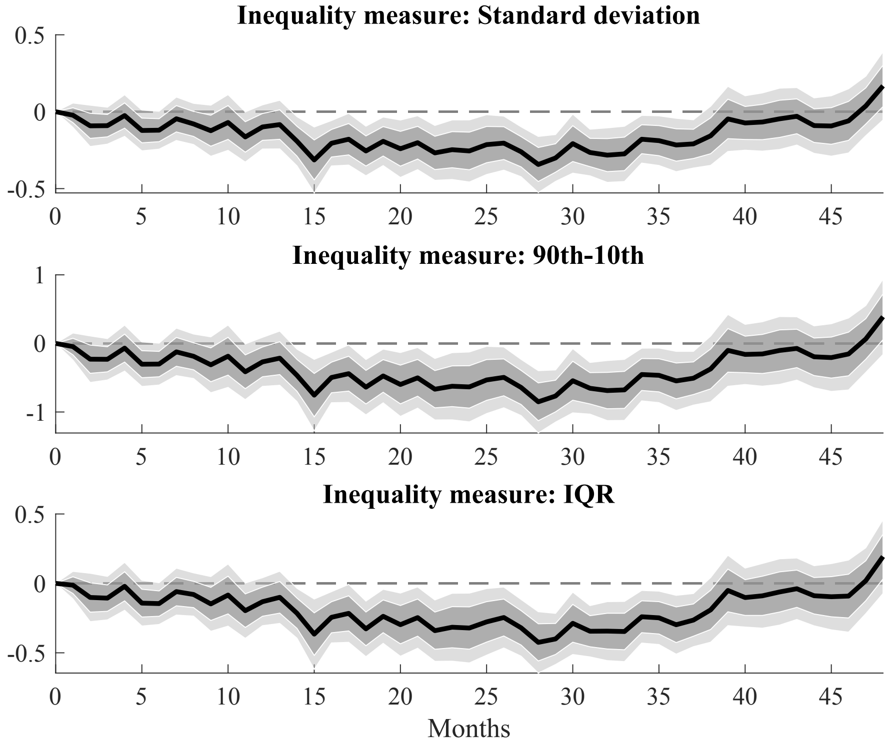

Notes: The figure plots impulse responses to a one percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the respective inflation dispersion measures. The horizontal axis is in months. Dispersion is measured using the cross-sectional standard deviation (top), the difference between the 90th and the 10th percentile of the cross-sectional distribution (middle), and the IQR (bottom). Impulse responses are computed at a monthly frequency using data relative to the period 1980M1:2008M12

standard deviation of inflation for the four divisions using expenditure weights as well as price indices at division level[23].

The responses across U.S. divisions are reported in Figure 13. There are some regional differences in the shape of the responses of inflation dispersion to contractionary shocks. However, the magnitude and significance of the results are comparable to the baseline specification. The decrease is more muted only for the West division.

Finally, in the main analysis, the decile-specific inflation rates are computed following the BLS procedure. The advantage of this approach is that for each decile all the individual expenditure information is combined to form the expenditure weights. In this way, outliers are less likely to bias the analysis. An alternative approach to the BLS methodology would be to simply consider the median of the individual inflation rates within each decile.

[23] A more limited number of price indices are available at the division level. Therefore, we used the following expenditure categories: Food at Home, Food Away from Home, Alcohol, Rented Dwellings, Owned Dwellings, Household Furnishings and Operations, Utility, Apparel, Private Transportation, Public Transportation, Gasoline, Medical, Education, and Miscellaneous.

Figure 12: Impulse responses of inflation dispersion and inequality, Bauer and Swanson (2022) monetary shocks

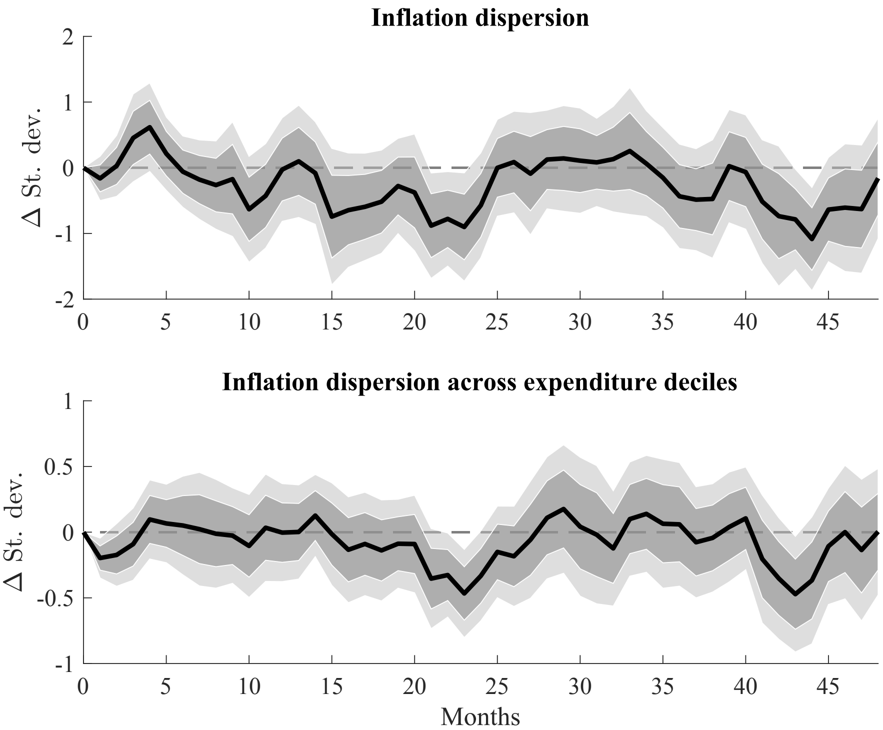

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as the 1.65 standard deviation confidence intervals. The horizontal axis is in months. Impulse responses are computed at a monthly frequency using data for the period 1988M2:2019M12.

In Figure 14 we report the responses of inflation inequality for income, salary, and expenditures to a contractionary monetary shock. Inflation inequality is measured as the standard deviation of the median inflation rates across deciles. Following a monetary shock the inflation inequality responses are still negative and statistically significant confirming the baseline results.

Figure 13: Impulse responses of inflation dispersion across U.S. divisions

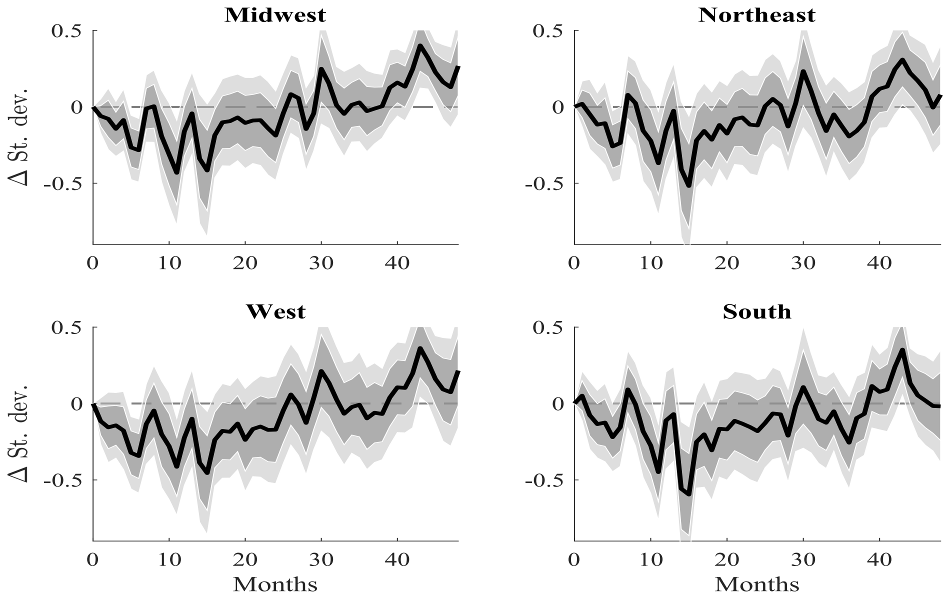

Notes: The figure plots impulse responses to a one percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the inflation dispersion measured as the cross-sectional standard deviation for the four US regions. Impulse responses are computed at a monthly frequency using data relative to the period 1980M1:2008M12.

Figure 14: Impulse responses of the dispersion across the median inflation rates for income, salary, and expenditure deciles

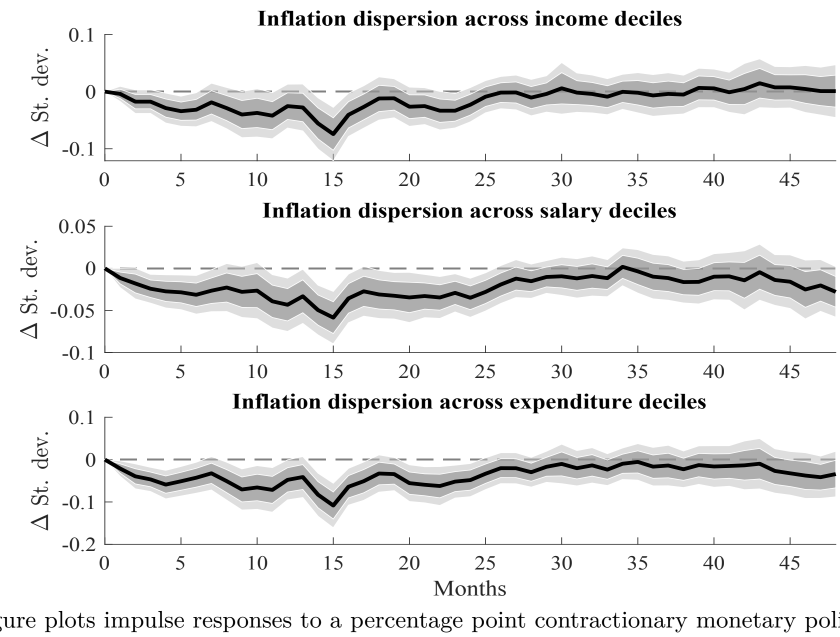

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for inflation inequality across income (top), salary (middle), and expenditure deciles (bottom). Inflation inequality is measured using the cross-sectional standard deviation of the median inflation rate for each decile. The horizontal axis is in months. Impulse responses are computed at a monthly frequency using data for the period 1980M1:2008M12.

E Robustness plots

Figure 15: Impulse responses of inflation dispersion

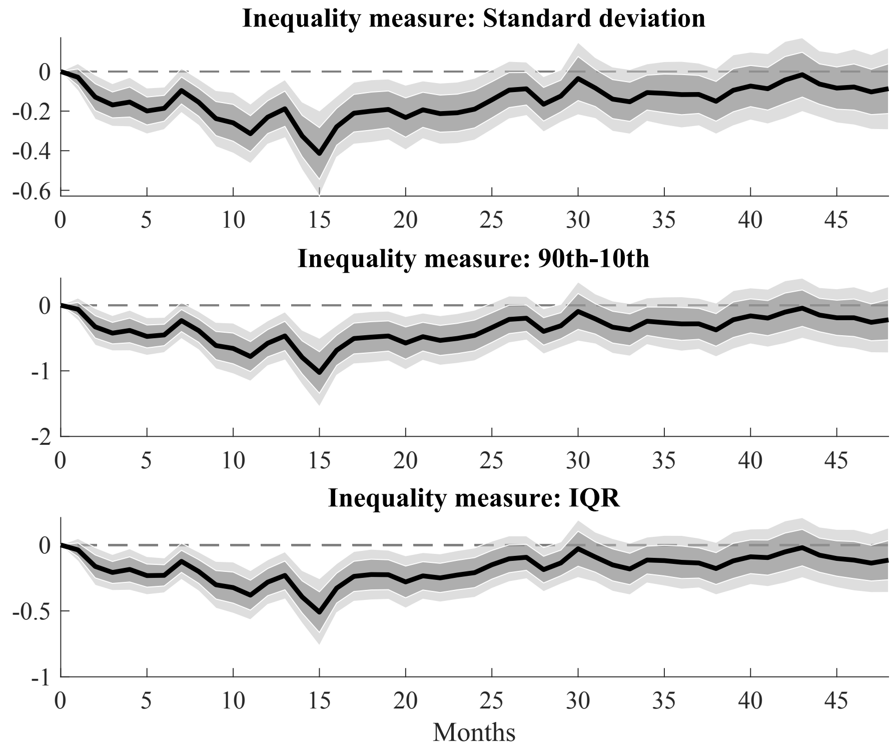

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the respective inflation dispersion measures. The horizontal axis is in months. Dispersion is measured using the cross-sectional standard deviation (top), the difference between the 90th and the 10th percentile of the cross-sectional distribution (middle), and the IQR (bottom). Impulse responses are computed at a monthly frequency using data for the period 1980M1:2008M12.

Figure 16: Historical series of inflation dispersion measures

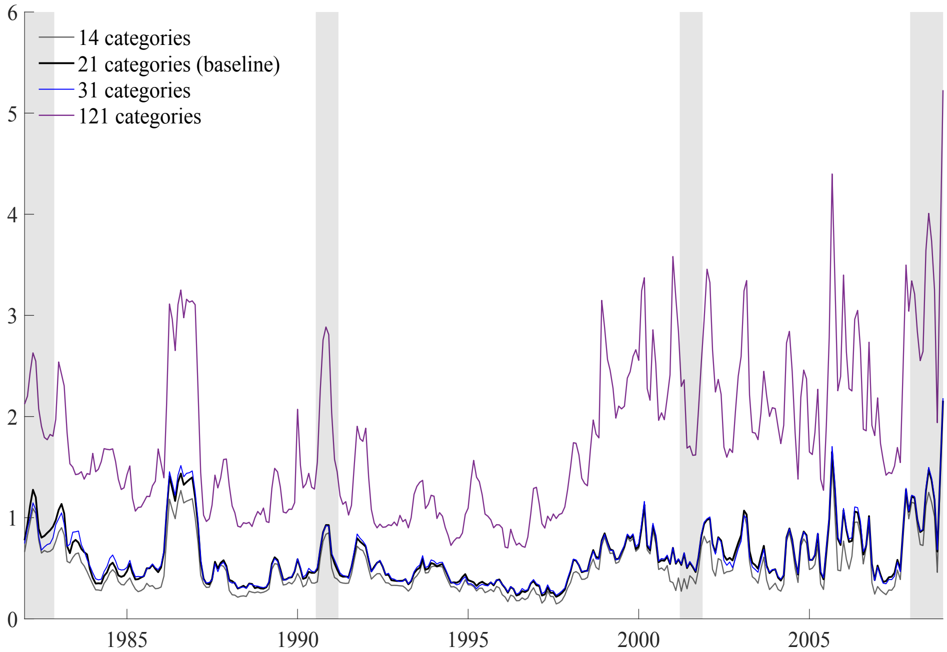

Notes: The plot shows the evolution of inflation dispersion measured using the cross-sectional standard deviation computed using 14, 21, 31, and 121 expenditure categories. All the series refer to the period 1981M1:2009M12. The gray shaded areas depict U.S. recessions.

Figure 17: Impulse responses of the cross-sectional standard deviation of inflation (alternative aggregations)

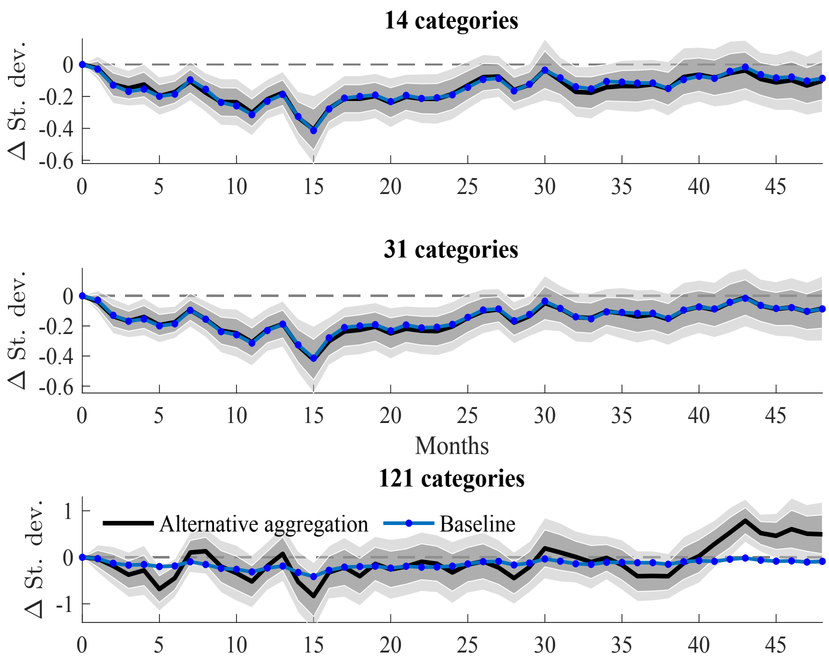

Notes: The figure plots impulse responses of alternatively aggregated inflation rates to a one percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the respective inflation dispersion measures. The solid blue line refers to the impulse response obtained using the baseline categories. Impulse responses are computed at a monthly frequency using data relative to the period 1980M1:2008M12.

Figure 18: Impulse responses of inflation inequality across income deciles with time-varying weights

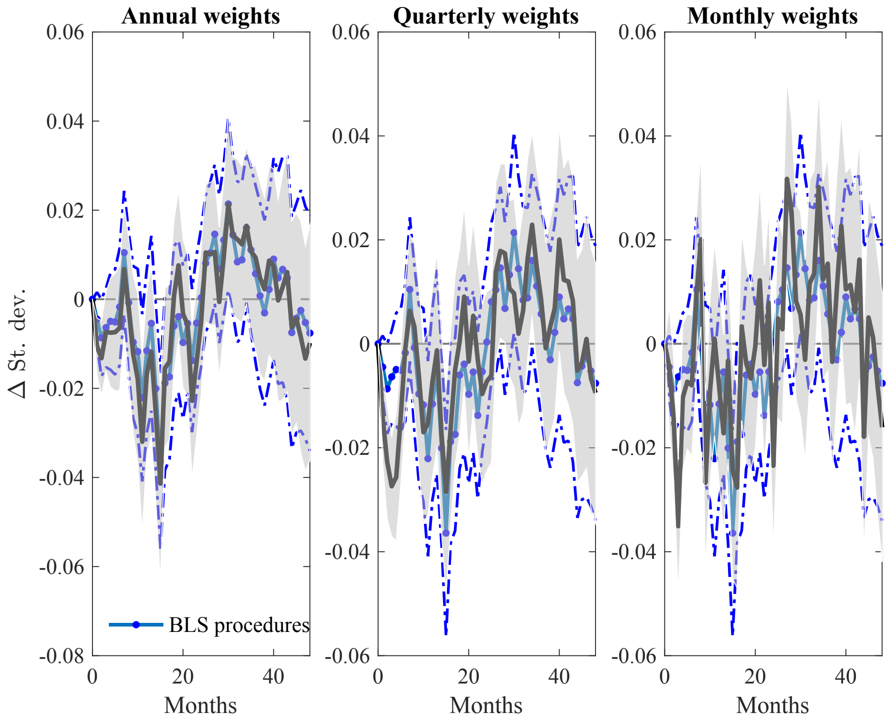

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock (black line) as well as one standard deviation confidence interval (gray area) for inflation inequality across income deciles. Inflation inequality is measured using the cross-sectional standard deviation of the decile-specific inflation rates. The expenditure weights are time-varying and computed at annual (left panel), quarterly (middle panel), and monthly (right panel) frequencies. The solid blue line refers to the baseline impulse response obtained following the BLS methodology for the expenditure weights, the blue dashed lines are the one standard deviation confidence interval. The horizontal axis is in months. The top panel uses the standard deviation in inflation rates for non-durable categories, the middle panel for durables, and the bottom panel for services. Impulse responses are computed at a monthly frequency using data for the period 1980M1:2008M12.

Figure 19: Impulse responses of inflation inequality across salary deciles with time-varying weights

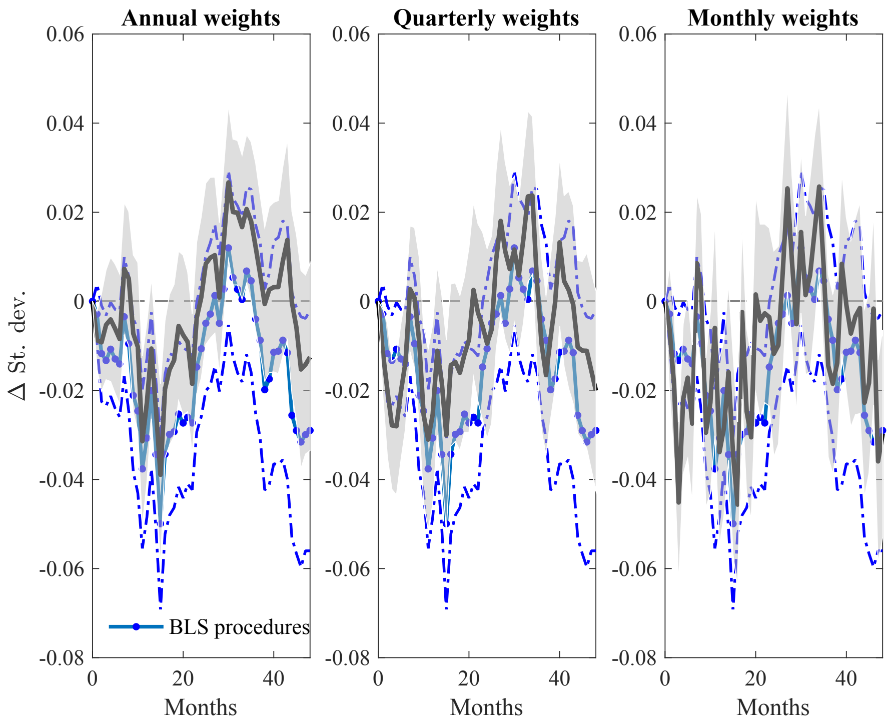

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock (black line) as well as one standard deviation confidence interval (gray area) for inflation inequality across salary deciles. Inflation inequality is measured using the cross-sectional standard deviation of the decile-specific inflation rates. The expenditure weights are time-varying and computed at annual (left panel), quarterly (middle panel), and monthly (right panel) frequencies. The solid blue line refers to the baseline impulse response obtained following the BLS methodology for the expenditure weights, the blue dashed lines are the one standard deviation confidence interval. The horizontal axis is in months. The top panel uses the standard deviation in inflation rates for non-durable categories, the middle panel for durables, and the bottom panel for services. Impulse responses are computed at a monthly frequency using data for the period 1980M1:2008M12.

Figure 20: Impulse responses of inflation inequality across expenditure deciles with time-varying weights

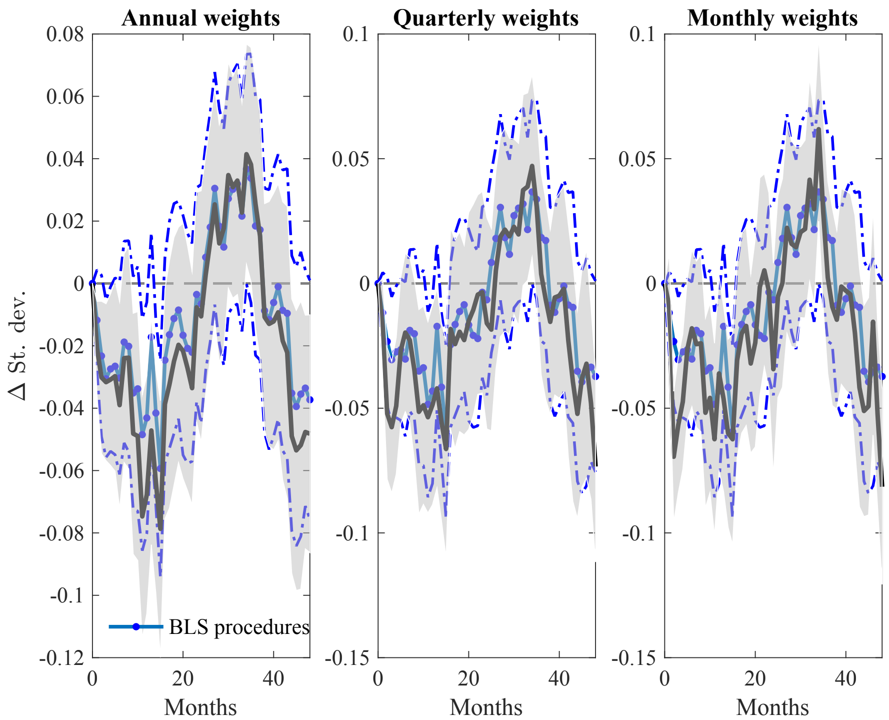

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock (black line) as well as one standard deviation confidence interval (gray area) for inflation inequality across expenditure deciles. Inflation inequality is measured using the cross-sectional standard deviation of the decile-specific inflation rates. The expenditure weights are time-varying and computed at annual (left panel), quarterly (middle panel), and monthly (right panel) frequencies. The solid blue line refers to the baseline impulse response obtained following the BLS methodology for the expenditure weights, the blue dashed lines are the one standard deviation confidence interval. The horizontal axis is in months. The top panel uses the standard deviation in inflation rates for non-durable categories, the middle panel for durables, and the bottom panel for services. Impulse responses are computed at a monthly frequency using data for the period 1980M1:2008M12.

Figure 21: Impulse responses of inflation dispersion for different lag specifications

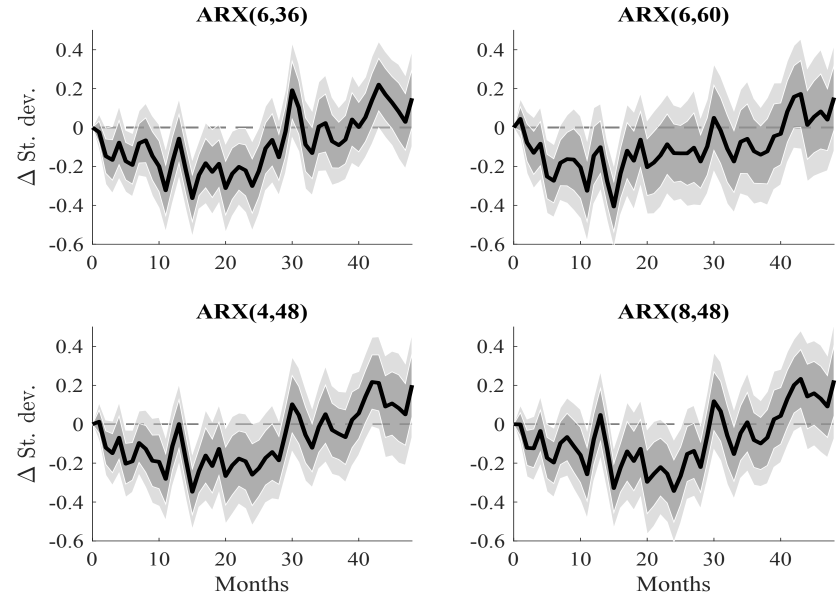

Notes: The figure plots the impulse responses to a one percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals of the cross-sectional standard deviation. The horizontal axis is in months. In an $ARX(p, r)$-model, we control for p lags of the dependent variable, and for r lags of the shock variable. Impulse responses are computed at a monthly frequency using data relative to the period 1980M1:2008M12.

Figure 22: Impulse responses of inflation dispersion (without Volcker period)

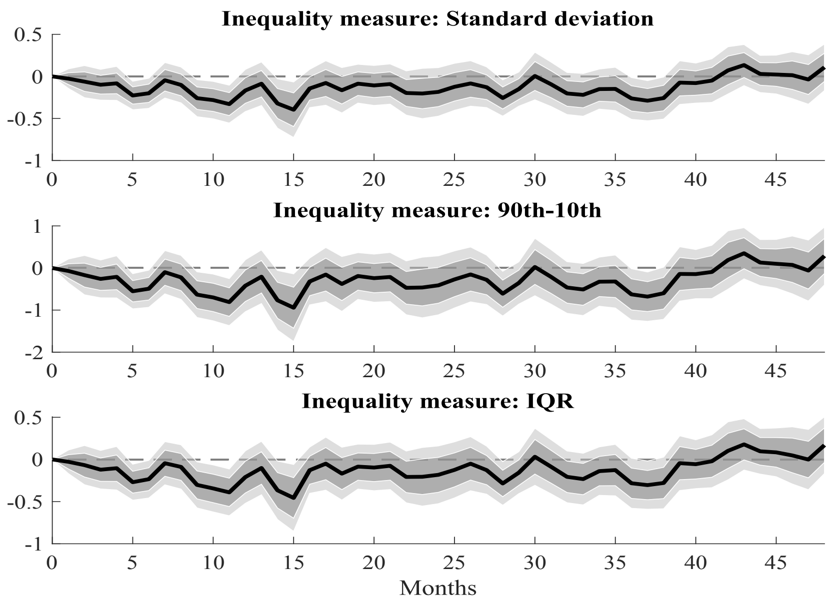

Notes: The figure plots impulse responses to a one percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the respective inflation dispersion measures. The horizontal axis is in months. Dispersion is measured using the cross-sectional standard deviation (top), the difference between the 90th and the 10th percentile of the cross-sectional distribution (middle), and the IQR (bottom). Impulse responses are computed at a monthly frequency using data relative to the period 1985M1:2008M12 in order to exclude the Volcker disinflation period.

Do firm expectations respond to Monetary Policy announcements?

Federico Di Pace* Giacomo Mangiante† Riccardo M. Masolo‡

March 2023

Abstract

We study whether firms' expectations react to the Bank of England's monetary policy announcements by comparing the responses to the Decision Maker Panel (DMP) survey filed immediately before and after a Monetary Policy Committee (MPC) meeting. On the one hand, we find that firms' expectations and uncertainty about their own business for the most part do not respond to high-frequency monetary policy *surprises*. On the other hand, announced *changes* in the monetary policy rate induce firms to revise their price expectations, with rate hikes resulting in a reduction in price expectations and the uncertainty surrounding them.

Keywords: Central bank communication, firm expectations, high-frequency identification, survey data.
JEL classification: D84, E52, E58.

*Bank of England, email: federico.dipace@bankofengland.co.uk
†University of Lausanne, Lausanne, Switzerland, email: giacomo.mangiante@unil.ch
‡Catholic University of the Sacred Heart, Milan, email: riccardomaria.masolo@unicatt.it
We are particularly indebted to Robin Braun for providing us with the monetary policy shocks series and to Tommaso Aquilante. We would also like to thank Florin Bilbiie, Carola Binder, Nick Bloom, Michael Ehrmann, Fiorella De Fiore, Yuriy Gorodnichenko, Mikosch Heiner, Marco Lombardi, Catherine L. Mann, Sebastian Rast, Ricardo Reis, Jean-Paul Renne, Johannes Schuffels, Boromeus Wanengkirtyo, and the seminar and conference participants at the Lausanne Ph.D. Macro Workshop, the Workshop on Empirical Monetary Economics, and the Bank of England for their helpful comments.
Disclaimer: The views expressed here are those of the authors alone and do not necessarily represent the views of the Bank of England or any of its committees.

1 Introduction

Understanding how expectations respond to announcements by the monetary authorities is pivotal for the transmission of monetary policy. Recently, the profession has devoted greater attention to distinguishing between different economic agents, e.g., market participants, households, firms, and professionals (Reis, 2020; Olivier et al., 2020). Financial markets react swiftly to monetary policy communications. This has been documented by Kenneth (2001) and the vast literature that ensued. However, traders are not the only decision-makers in the economy. Responses by households and firms are arguably more relevant to assessing the macroeconomic effects of policy intervention. Firms set prices and ultimately determine inflation.

In this paper, we study whether firm expectations respond to the Bank of England (BoE)'s monetary policy announcements and, if so, how. We contrast U.K. firms' survey responses filed in the days immediately preceding an MPC meeting to those submitted in its aftermath. We find that announcements of interest rate *changes* cause a significant movement in price expectations, and the uncertainty surrounding them. An interest rate hike leads to a reduction in price expectations, in line with what economic theory predicts in response to a contractionary monetary policy shock. That is, in general, not the case if we use the conventional definition of policy surprise, defined as the change in the 3-month Sterling future price in a short window around the policy announcement – Cesa-Bianchi et al. (2020) and references therein.

In sum, firms do not respond to monetary announcements the same way financial markets do. Announcements that resulted in sizeable monetary policy *surprises* but did not involve any interest rate adjustment do not affect firms' expectations. At the same time, firms' expectations do react to policy rate *changes*, even though they may well have been anticipated by financial markets.

To evaluate whether firms' expectations respond to monetary policy announcements we use the U.K. survey Decision Maker Panel (DMP). The DMP has data on firms' expectations and uncertainty about their own future sales, price, employment, and investment (Bloom et al., 2018). To isolate the effects of the monetary announcements, we exploit the date on which different firms filed their answers. By comparing the responses of those firms which responded immediately before to those which responded after a Monetary Policy Committee (MPC) meeting of the Bank of England, we can test whether the expectations are influenced by monetary policy decisions. We find that firms' expectations do not significantly react to monetary policy surprises. The results hold controlling for the size of the shock and other firms' observable characteristics.

However, firm expectations and uncertainty strongly respond to the MPC announcements of interest rate *changes*. This is consistent with Google Trends data showing that MPC announcements of changes in rates associate with spikes in attention by the general public, while that is not the case when no change in policy rates is announced. The decrease in interest rate announced on the 11th of March 2020 resulted in a sizeable increase in price expectations of around 1 percentage point (pp) whereas the policy tightening of the 16th of December 2021 and the 17th of March 2022 led to a decrease in price expectations of 1.6 and 1.8 pp respectively in line with the economic theory. All three changes caused a reduction in the level of price uncertainty.

Related Literature. This paper contributes to two strands of the literature. First, the results complement the large body of empirical evidence on the effects of monetary policy announcements on expectations that rely on event studies. Lamla and Vinogradov (2019), Rast (2021), Fiore et al. (2022) and Binder et al. (2022) focus on the response of the households' expectations. Lamla and Vinogradov (2019) run their own survey around Federal Open Market Committee (FOMC) meetings and document that the announcements have no measurable effect on average beliefs but make people more likely to receive news about the central bank's policy. Rast (2021) uses the GfK survey and finds that policy rate announcements lead to significant adjustments in household inflation expectations, unlike those about forward guidance and quantitative easing. Fiore et al. (2022) rely on the responses from the New York Fed's Survey of Consumer Expectations before and after FOMC meetings and find that only the expectations about interest rates are affected. Binder et al. (2022) use the same survey to evaluate how household inflation expectations respond to FOMC announcements, macroeconomic data releases, and news related to politics and the Covid-19 pandemic.

Similarly, Lewis et al. (2019) use daily survey data from Gallup to assess how households' beliefs about economic conditions are influenced by monetary policy: changes in the federal funds target rate have a significant and instantaneous effect on economic confidence. Singh and Mitra (2023) use the same data and find that household expectations respond primarily to announcements regarding the unemployment rate. Claus and Nguyen (2020) apply a latent factor model to consumer survey data from the Australian CASiE survey to document that expectations about economic conditions, unemployment, and readiness to spend adjust in the direction predicted by standard models following a monetary policy shock.

More closely related to our paper, Enders et al. (2019) study whether the expectations of the German firms participating in the Ifo Business Survey Industry respond to policy surprises. They find that the responses of production and price expectations are highly non-linear in the size of the monetary surprise and that many of the ECB's announcements of non-conventional policies did not shift expectations significantly. Bottone and Rosolia (2019) use the Bank of Italy's quarterly Survey of Inflation and Growth Expectations and document that firms' pricing plans are not affected by monetary policy shocks. Pinter and Kočenda (2022) show that French firms' and households' expectations react to central bank announcements only once the media response to the announcement is taken into account. In line with this evidence, we confirm that on average monetary announcements have limited influence on firm expectations. However, we also show that changes in the monetary policy rate are able to significantly alter the expectations in the direction predicted by economic theory and reduce the uncertainty around them.

We also contribute to the literature on how different economic agents form their expectations and respond to shocks. Peter et al. (2022a) study how agents expect different shocks to transmit to the macroeconomy. They provide 6,500 U.S. households and 1,500 experts with identical information about the parameters of the shocks and document that their beliefs about the directional effects and the propagation channels of shocks are widely dispersed. Similarly, Peter et al. (2022b) find that households and experts explain macroeconomic phenomena in a completely different way. In surveys with more than 10,000 US households and 100 academic experts, the authors document that households' narratives are strongly heterogeneous and coarser than experts' narratives as well as focus more on the supply side than on the demand side. Reis (2020) focuses on the discrepancy between market prices and people's long-run inflation expectations. This discrepancy is found to have large business-cycle fluctuations, to be related to monetary policy, and to be driven by disagreement across groups in the population. Candia et al. (2022) compare U.S. firms' inflation expectations to those of households and professional forecasters and show that U.S. managers have far from anchored expectations and that they are poorly informed about recent aggregate inflation dynamics or monetary policy. Similarly, Link et al. (2023) find sizable differences in information frictions between firms and households with firms' expectations about macroeconomic variables being closer to expert forecasts and less dispersed than households'. The heterogeneous response of households and firms to shocks is evaluated by Mikosch et al. (2022) as well. They show that an exogenous increase in the perceived uncertainty of the exchange rate leads to an increase in firms' demand for a report about exchange rate developments, but not households'.

We expand this literature by showing that financial markets and firms do not respond the same way to monetary policy announcements. Unlike markets, firms' expectations react significantly and sizeably only to MPC announcements of interest rate *changes*.

The remainder of the paper is organized as follows. Section 2 details the data which we use in the paper. In section 4 we describe the empirical specification that we adopt to evaluate the effects of the monetary announcements on firms' expectations. Section 5 reports the main results of the analysis. Finally, section 6 concludes.

2 Data

Decision Maker Panel. Most of the analysis relies on firm survey data from the Decision Maker Panel – Bloom et al. (2019a), Bloom et al. (2019b) and Altig et al. (2020). The DMP was launched in August 2016 by the Bank of England, the University of Nottingham, and Stanford University and it is now one of the largest regular business surveys, with a panel of 8,000 firms and around 3,000 responding in any given month.[1] The respondents of the survey are the Financial Officers from small, medium, and large U.K. companies operating in a broad range of sectors, and the survey is designed to be representative of the population of U.K. businesses.

We focus on the questions regarding subjective expectations about future growth in prices, sales, employment, and investment:

1. *Looking ahead, from now to 12 months from now, what approximate % change in your AVERAGE PRICE (EMPLOYMENT) would you assign to each of the following scenarios? (with five scenarios: lowest, low, middle, high, and highest provided)*

2. *Please assign a percentage likelihood (probability) to the % changes in your AVERAGE PRICES (EMPLOYMENT) you entered*

3. *Looking a year ahead from the last quarter, by what % amount do you expect your SALES REVENUE (CAPITAL EXPENDITURE) to have changed in each of the following*

[1] More information about the representativeness of the data and the structure of the survey can be found `here`.

scenarios?" (with five scenarios provided; i) lowest, ii) low, iii) middle, iv) high, v) highest)

4. *Please assign a percentage likelihood (probability) to the% changes in your SALES REVENUE (CAPITAL EXPENDITURE) you entered*

For each firm, we have its lowest, low, medium, high, and highest expectations and the probabilities associated with them (Altig et al., 2019). We can thus compute the firm's expected value and the uncertainty surrounding it.

Each respondent i supplies future growth rates, which we refer to as $\Delta z_{i,t}$. For each variable she provides five values, $\Delta z_{i,t,j}$ at support points $j = 1, 2, 3, 4, 5$, and the associated probabilities, $p_{i,t,j}$. We calculate the respondent's mean expectation of the growth rate, for each period, as:

$$\mathrm{Mean_{i,t}}(\Delta z_{i,t,j}) = \sum_{j=1}^{N} p_{i,t,j} \cdot \Delta \mathrm{z}_{i,t,j}, \tag{1}$$

and the relative subjective uncertainty as the standard deviation,

$$SD_{i,t}(\Delta z_{i,t,j}) = \left[\sum_{j=1}^{N} p_{i,t,j} \left(\Delta \mathrm{z}_{i,t,j} - \ \mathrm{Mean}_{i,t}(\Delta \mathrm{z}_{i,t,j}) \right)^2 \right]^{1/2}. \tag{2}$$

Finally, the DMP survey provides a rich set of firm-level characteristics, like their size and sector, which we use as controls. In Table 1 we provide some descriptive statistics for the main variables of interest of our analysis.

Table 1: Descriptive statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Expected price growth	2.620	4.736	-8.534	12.435	5,263
Price uncertainty	2.225	3.869	0.692	4.527	5,263
Expected sales growth	8.546	5.921	-2.456	16.387	6,254
Sales uncertainty	6.879	2.278	2.937	11.749	6,254
Expected emp. growth	0.835	1.267	-1.017	3.311	6,898
Emp. uncertainty	6.687	2.483	2.472	10.843	6,898
Expected inv. growth	16.16	8.541	-3.006	27.498	6,615
Inv. uncertainty	23.45	6.837	3.923	34.537	6,615

Notes: The table reports descriptive statistics for the main variables of interest of the DMP in the days around the monetary policy announcements. We report information on the expected growth in prices, sales, employment, and investment and the relative uncertainty.

Monetary policy surprises. To measure the unexpected component of the monetary announcements, we use the surprises computed by Cesa-Bianchi et al. (2020) and extended

Figure 1: Bank of England Base Rate and monetary policy surprises

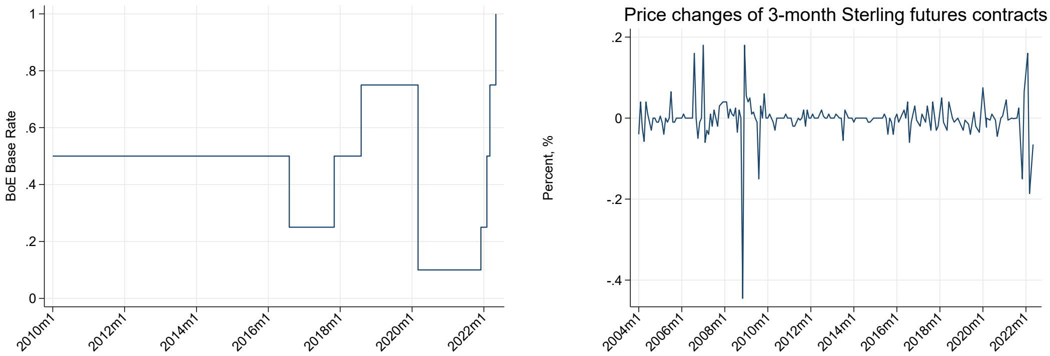

Notes: The left panel plots the Bank of England Base Rate over time. The vertical axis is in annual percentage points. The right panel reports monetary policy surprises, computed as the changes in the second front contract of the 3-month Sterling future, the 3-to-6 month ahead expectation about the 3-month Libor, in a 30-minute window around monetary policy events.

by Braun et al. (2022), based on the high-frequency identification approach developed by Gurkaynak et al. (2005). The surprises are the changes in the price of 3-month Sterling futures contracts expiring 2 quarters ahead in a 30-minute window around the announcements of the Monetary Policy Committee of the Bank of England.[2]

In the left panel of Figure 1 we plot the time series of the BoE Bank Rate. Despite the DMP survey only starting in 2016, we are able to capture important monetary events. Since 2016 the Bank Rate has been adjusted several times to respond to different events related to the Brexit referendum, Covid, and the recent increase in the inflation rate. This is reflected in the evolution over time of the monetary policy shocks, reported in the right panel of Figure 1. From 2016 onward the magnitude and the volatility of the surprises in the Bank Rate have correspondingly increased, relative to the early 2010s.

Google Trends. Google trends measure the search interest for certain topics/keywords. We use Google Trends data to assess whether the general public interest in the activities of the BoE increases in correspondence with monetary policy announcements.

[2]The list of MPC meetings for which at least one of the expectations measures is available is reported in Table (9) of the Appendix. From 2021 onward the Libor-based futures are not available anymore, so Sonia-based futures are used instead.

3 A Model of Expectation Formation

We consider a simple model of expectation formation to guide our intuition on how firms might respond to monetary announcements. The model closely follow the derivations in Singh and Mitra (2023). Suppose there is an announcement X_t which has two components: anticipated A_t and unanticipated U_t:

$$X_t = A_t + U_t. \tag{3}$$

The anticipated component is the part of the announcement that fully rational agents can forecast with the information they have. The unanticipated component is the part of the announcement that is not predictable given the information set.

We define with Y_t some fundamentals about the economy based on which firms form their future expectations. After the announcement, the fundamentals can be written as:

$$Y_t = p \cdot g\left(X_t\right) + (1-p) \cdot h\left(\psi_t\right), \tag{4}$$

where with ψ_t we denote all information aside from the announcement that is available to agents for forming expectations. The firms give a weight p to the information of the announcement in their expectations process. The expectation of the fundamental before the announcement can be written as:

$$E_{t-1}\left(Y_t\right) = E_{t-1}\left[p \cdot g\left(X_t\right) + (1-p) \cdot h\left(\psi_t\right)\right] = p \cdot E_{t-1}\left[g\left(X_t\right)\right] + (1-p) \cdot E_{t-1}\left[h\left(\psi_t\right)\right], \tag{5}$$

Therefore, the change in the firms' fundamentals before and after the announcement is equal to:

$$Y_t - E_{t-1}\left[Y_t\right] = p \cdot \left[g\left(X_t\right) - E_{t-1}\left[g\left(X_t\right)\right]\right] + (1-p) \cdot \left[h\left(\psi_t\right) - E_{t-1}\left[h\left(\psi_t\right)\right]\right] \tag{6}$$

If we consider a small enough time window around a monetary announcement we can assume that the only new information firms receive in that time frame is that provided in the announcement itself. It follows that $h\left(\psi_t\right) - E_{t-1}\left[h\left(\psi_t\right)\right] = 0$. The impact of the announcement simplifies to:

$$Y_t - E_{t-1}\left[Y_t\right] = p \cdot \left[g\left(X_t\right) - E_{t-1}\left[g\left(X_t\right)\right]\right]. \tag{7}$$

Under the assumption that g is the identity function, we have:

$$Y_t - E_{t-1}[Y_t] = p \cdot [X_t - E_{t-1}(X_t)] \tag{8}$$

Thus, the impact of a monetary announcement on firms' expectations is a function of the difference between the information released in the announcement and the forecast of it.

We consider two extreme cases of how firms form their expectations about the announcement: sophisticated and naive firms. On the one hand, sophisticated firms are fully rational and use the entire available information set to make their forecast. On the other hand, naive firms do not use any new information between two announcements to update their beliefs.

The forecast of the sophisticated firms is given by:

$$\mathbb{E}^S(X_t) = A_t, \tag{9}$$

which means they are able to perfectly anticipate the predictable part of the announcement. The change in fundamentals due to the announcement boils down to:

$$\begin{aligned} Y_t - E_{t-1}(Y_t) &= p^S \cdot \left[X_t - E^S_{t-1}(X_t)\right] \\ &= p^S \cdot [X_t - A_t] \\ &= p^S \cdot U_t \end{aligned} \tag{10}$$

where p^S is the weight sophisticated firms give to the announcement. Since firms are fully rational, only the unpredictable component of the announcement, i.e., U_t, will affect their expectations.

Naive firms do not update their beliefs between announcements. Therefore, we assume they are entirely backward-looking and their expectation at time $t-1$ of X_t is simply the value of X from the previous announcement:

$$E^N_{t-1}(X_t) = X_{t-1}. \tag{11}$$

The change in expectations due to the announcement becomes:

$$\begin{aligned} Y_t - E_{t-1}(Y_t) &= p^N \cdot \left[X_t - E^N_{t-1}(X_t)\right], \\ &= p^N \cdot [X_t - X_{t-1}], \end{aligned} \tag{12}$$

where p^N is the weight naive firms give to the announcement. Hence, naive firms update their expectations only if the value of X changes between two announcements.

4 Estimation Strategy and Identification

We estimate the treatment effect of monetary policy announcements by comparing the expectations of the survey respondents right before the MPC announcement with those right after, along the lines of Rast (2021) and Lamla and Vinogradov (2019).

DMP surveys are conducted monthly over a period of 2 to 3 weeks. Firms can respond at any time during that period. We focus on the monetary policy announcements that take place during the time window in which the DMP survey is administered and contrast the responses of firms that submitted their responses right before the announcement to those that did in the aftermath.[3]

In particular, we estimate the following regression specification:

$$y_{i,t} = \alpha + \beta D_{i,t}\mathbf{s}_t + X_{i,t} + \epsilon_{i,t}, \tag{13}$$

where $D_{i,t}$ is a dummy equal to 1 if the firm responds after the announcement (as a baseline we use a symmetric time window around the MPC announcements of 5 days), s_t represents the shock coming from new information in the announcement, and $X_{i,t}$ is the matrix of control variables, which includes sector and wave fixed effects. By wave, we refer to the monthly administering of the survey. A wave is completed within a month, so we could equivalently label it as monthly fixed effect. Standard errors are clustered at the wave level.[4]

The shock s_t to information coming from monetary announcements will vary depending on whether firms are sophisticated or naive. In the former case, only unanticipated changes in the announcement can influence expectations. Therefore, as a measure of the unpredictable component of the announcement we consider the financial market surprises. In the latter case, firms are entirely backward-looking so the shock is equal to the change in the policy rate relative to the previous announcement.

In Figure 2, we report the total number of respondents for each day of the month. The majority of firms submit their responses in the second week of the month, while only a few

[3] We exclude the responses that have been filed on the days of an announcement as we do not observe the exact time of the submission.

[4] In Section 5.2.1 we interact the term $D_{i,t}s_t$ with firm-level characteristics that may influence the reaction to monetary policy news, to isolate elements of heterogeneity in firm responses.

Figure 2: Total number of firms submitting their answers to the DMP survey by day of the month

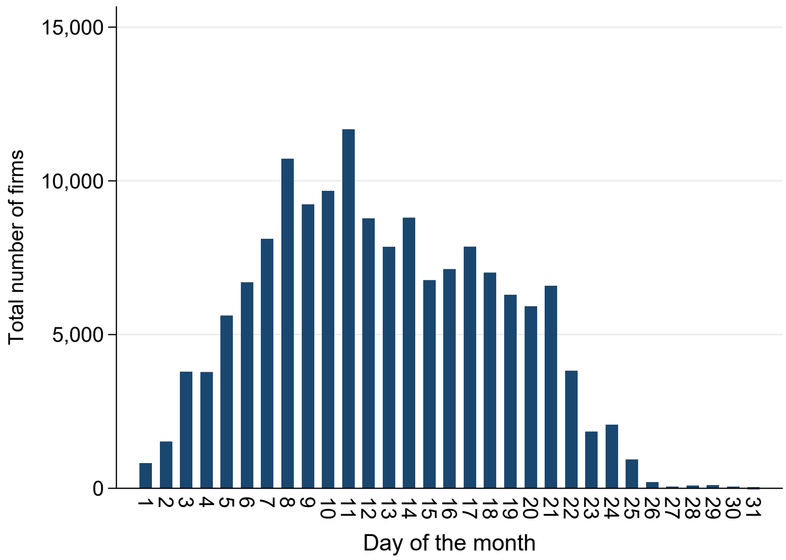

file their answers during the last week. So if an announcement is made towards the end of the month, we may not be able to include it in our analysis as no firm's observations fall within the 5-day window around the announcement.

More important for our identification strategy is that the date on which firms file their responses does not depend systematically on firm characteristics or on the timing of the policy announcement. In Section 5.2 we test this assumption and find that the probability of answering the survey before or after the MPC announcements is unrelated to firms' observable characteristics.

5 Baseline Results

In this section, we report the main results of our empirical analysis. First, we show that monetary surprises have no impact on firms' expectations about their own business. Second, we find that the general public attention to the BoE's activities spikes during MPC meetings if a change in interest rate is announced. Third, we demonstrate that firms' expectations and uncertainty significantly respond to the announcements of interest rate changes. Overall, the results suggest that the firms react to monetary policy announcements only if they involve an interest rate adjustment.

5.1 Do firms respond to monetary policy surprises?

If firms are sophisticated, only the unanticipated component of the announcements will shift their expectations. Therefore, we start by evaluating whether firms' expectations react to monetary *surprises* as defined by the reaction of financial markets. We estimate equation (13) with firm expectations and uncertainty as the dependent variable. The monetary policy shocks are the surprises computed by Cesa-Bianchi et al. (2020) and extended to 2022 by Braun et al. (2022).

Table 2 reports the coefficients of interest from equation (13), considering the 12-month ahead price and employment growth, the 4-quarter ahead of sales and investment growth as well as their relative subjective uncertainty, as the dependent variable.

Table 2: MPC announcements and firm expectations and uncertainty

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
Surprise x Dummy mpc	4.409	2.684	5.109	1.076	1.265	1.158	-13.51	28.94
	(2.988)	(2.370)	(4.903)	(5.776)	(3.496)	(1.947)	(27.09)	(21.29)
Constant	2.954***	2.239***	8.456***	6.946***	0.835***	6.751***	16.18***	48.00***
	(0.00956)	(0.00759)	(0.0166)	(0.0195)	(0.0155)	(0.00862)	(0.0968)	(0.0761)
Observations	5263	5263	6254	6254	6898	6898	6615	6615
R^2	0.070	0.038	0.055	0.041	0.059	0.035	0.034	0.019
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
Number of meetings	34	34	35	35	35	35	35	35

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The coefficients measure by how much the expectations (and uncertainty levels) of firms that respond to the survey after a hypothetical monetary policy surprise of 1 percentage point (pp) differ from those that filed their answers beforehand. None is significant. Moreover, it is important to note that the standard deviation of the surprises is around 0.05 pp so a one standard deviation shock increases price expectations by 0.23 pp. Therefore, the effects are not only statistically but also economically insignificant.

However, it could be that firms respond differently to different monetary policy shocks as documented by Enders et al. (2019). If that were the case, a linear model would not be appropriate as it would conflate different responses into a single coefficient. We thus sort the monetary surprises from 2016m8 (when the DMP survey becomes available) onward according to their size and we break them into terciles, $b \in \{\text{bottom, middle, top}\}$. It follows that the bottom (top) bin includes only large expansionary (contractionary) surprises, i.e., negative (positive) monetary policy shocks. We then estimate the following model:

$$y_{i,t} = \alpha + \sum_{b=1}^{3} \beta_b D_{i,t} \mathbf{s}_{t,b} + X_{i,t} + \epsilon_{i,t} \tag{14}$$

where $\mathbf{s}_{t,b}$ assumes the value of the monetary surprise in case it falls into bin b and zero otherwise. Table 3 reports our estimates.

Table 3: MPC announcements and firm expectations and uncertainty, bins approach

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
Surprise bottom 1/3 x Dummy mpc	6.398***	4.523***	5.836	0.438	4.344*	0.514	14.92	55.10***
	(0.708)	(1.289)	(4.937)	(5.586)	(2.161)	(1.093)	(31.74)	(12.27)
Surprise middle 1/3 x Dummy mpc	-70.33	39.79	-318.0	-154.3	-59.49***	-33.05	-1980.3	-468.5
	(71.06)	(34.97)	(268.9)	(97.17)	(21.05)	(200.0)	(1349.0)	(639.5)
Surprise top 1/3 x Dummy mpc	-4.982	-6.865***	3.786	5.873	-12.99	4.557	-120.4**	-80.44**
	(11.87)	(2.151)	(12.12)	(22.77)	(11.78)	(9.276)	(50.58)	(37.08)
Constant	3.113***	2.416***	8.441***	6.841***	1.074***	6.688***	17.87***	49.91***
	(0.173)	(0.0391)	(0.208)	(0.351)	(0.165)	(0.131)	(0.904)	(0.567)
Observations	5263	5263	6254	6254	6898	6898	6615	6615
R^2	0.070	0.039	0.055	0.041	0.059	0.035	0.034	0.019
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
Number of meetings	34	34	35	35	35	35	35	35

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Most of the coefficients are again insignificant. One exception is the positive price response pertaining to the bottom tercile of the surprises. As we discuss more in detail below, the result is driven by the announcement of the 17th of March 2022. The BoE increased its interest rate leading to a decrease in firms' price expectations. However, the interest rate increase was smaller than what was anticipated by the financial markets resulting in a negative surprise. The coefficient becomes insignificant once this announcement is removed from the sample[5].

Overall, it appears that firms do not respond to monetary policy surprises. The results suggest that the firms' expectations formation process is not as sophisticated as the financial markets. We next evaluate whether firms are more naive in their expectations and closer to the general public. Therefore, we turn to Google Trends (GT) data to investigate what type of announcements, if any, captures the general public attention the most.

The top left panel of Figure 3 shows the Google searches for "BoE Bank Rate" (right axis) and the actual Bank Rate at monthly frequency (left axis) – the top right panel zooms into March 2020, the one month in which two separate policy announcements were made.[6] Google searches for the term "BoE Bank Rate" spike when an MPC announcement corresponds to an interest rate *change*. The interest in the activity of the central bank rises on the exact day of the announcement, suggesting that the communication is effective at capturing the attention of the general public.

The bottom panel of Figure 3 overlays monetary policy surprises (Cesa-Bianchi et al., 2020) (left axis) to the Google Trend series (right axis) instead. As one can notice, there is no clear relationship between the size of the shocks and the level of attention by the general public. Some announcements which involved a change in interest rate, e.g., on the 17th of March 2022,

[5]See Table 5.

[6]The time series are standardized such that the maximum value for the time period considered is equal to 100.

Figure 3: BoE Bank Rate, Google searches, and monetary policy surprises

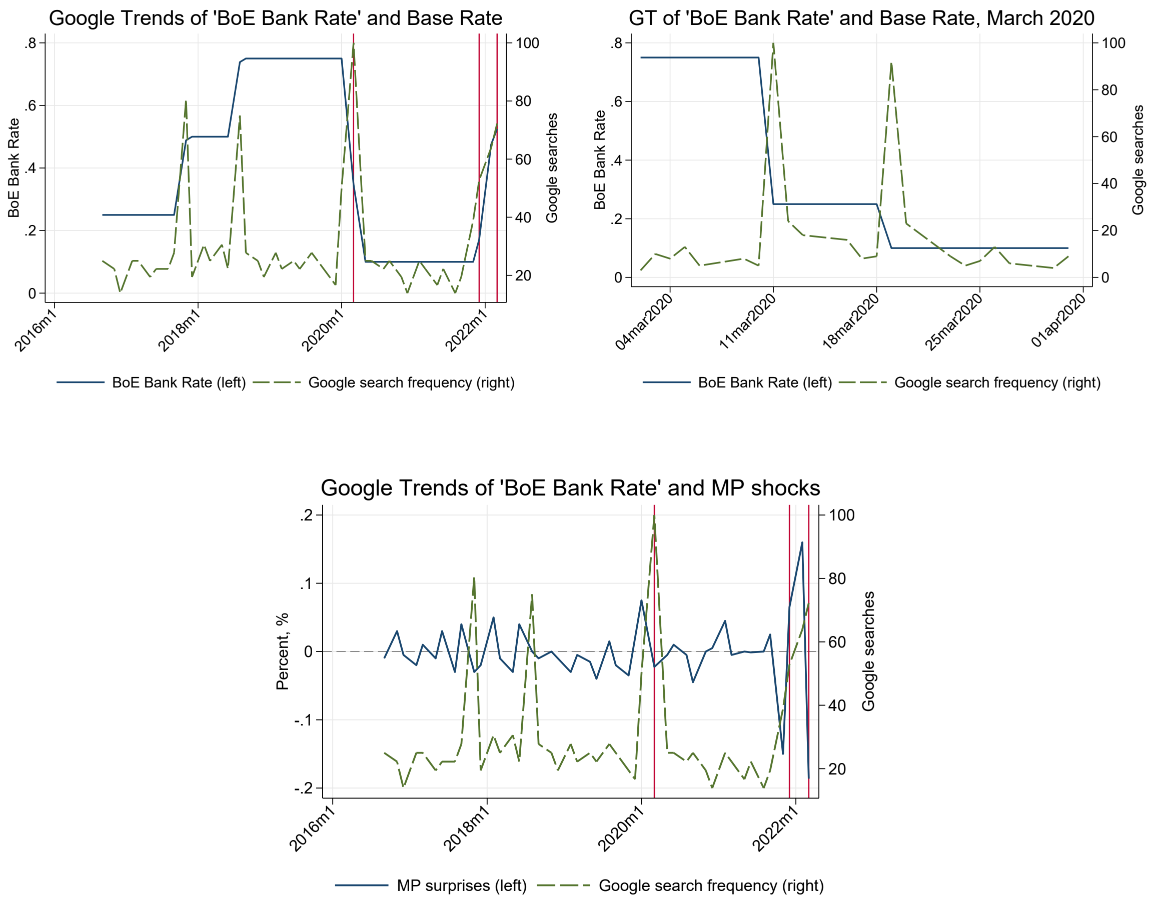

Notes: The top left panel shows the time series of the Bank of England Base Rate and the relative Google search index at a monthly frequency. The top right panel reports the same variables at daily frequency for March 2020. The bottom right panel confronts Google searches with monetary surprises. The red vertical lines are the MPC meetings which were scheduled in the middle of the month and are then part of the analysis.

were almost perfectly anticipated by the markets resulting in almost zero surprises despite sizeably increasing the general public attention. Therefore, news of rate changes reaches the public and stirs its interest, irrespective of whether the rate moves represent a surprise to financial market participants. Conversely, market-based surprises do not make the news and do not capture the attention of the general public.

Google trend data suggest that interest rate *changes* may be more *salient* than high-frequency surprises. To test whether firms are more naive in their expectation formation, we thus estimate a series of regressions that study the response of expectations and uncertainty to the MPC announcements which involved a rate change. Since the DMP survey has been launched, the BoE has changed the policy rate 7 times. Out of these 7 times, 3 times were announced in the middle of the month when the DMP survey is administered allowing us to compare how firms' expectations adjusted in response. The announcements happened on the 11th of March 2020, the 16th of December 2021, and the 17th of March 2022. The vertical red lines in Figure 3 indicate the three meetings.

As reported on the top right panel of Figure 3, on the 11th of March 2020 the BoE announced a sizeable decrease in the Base Rate from 0.75 percentage points to 0.25 to stimulate the economy in the Covid period. On the 16th of December 2021, the monetary authority adopted a more active stance against the surge in the inflation rate by increasing the interest rate from 0.1 percentage point to 0.25 and on the 17th of March 2022 from 0.5 to 0.75. Interestingly, the first interest-rate hike corresponds to a large positive market-based surprise: markets had not fully priced in the interest-rate increase. The second one to a negative surprise of almost 20bps: market participants were expecting a larger increase in the policy rate. The three announcements were extensively covered in the media given the important economic challenges they were responding to[7].

Table 4: Individual MPC announcements and firm expectations and uncertainty

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
Meeting 11/03/2020 x Dummy mpc	1.044***	-0.594***	-0.837***	-1.943***	-1.763***	0.926***	0.774	2.386**
	(0.00959)	(0.0255)	(0.0696)	(0.138)	(0.0255)	(0.0538)	(0.951)	(0.381)
Meeting 16/12/2021 x Dummy mpc	-1.594***	-0.301***	0.918	2.773***	-0.0799	-0.551*	4.126	-1.908*
	(0.0595)	(0.0106)	(0.320)	(0.220)	(0.245)	(0.159)	(2.967)	(0.491)
Meeting 17/03/2022 x Dummy mpc	-1.840***	-0.617***	-2.077*	-1.498**	-1.013*	0.529**	5.410	-7.628***
	(0.0985)	(0.0587)	(0.490)	(0.228)	(0.302)	(0.0652)	(3.548)	(0.319)
Constant	3.721***	2.654***	7.315***	7.916***	2.108***	6.036***	2.037**	46.61***
	(0.0222)	(0.0138)	(0.0430)	(0.0109)	(0.0861)	(0.00774)	(0.377)	(0.136)
Observations	402	402	486	486	553	553	540	540
R^2	0.124	0.059	0.053	0.045	0.089	0.062	0.037	0.039
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4 presents the results of the expectations and uncertainty variables regressed on announcement-specific treatment indicators. The coefficients can be interpreted as the mean difference (controlling for wave and announcement fixed effects) between the expectations of those surveyed before the respective MPC announcement and those surveyed afterward.

The monetary policy loosening of March 2020 leads to an increase in price expectations, and the two tightening episodes to a reduction. Estimates are statistically significant across the board and economically meaningful too. For instance, before the MPC announcement to increase the interest rate by 0.25 percentage points (pp) on the 17th of March 2022, the average expectation for price growth was around 6.1 pp, after the announcement it fell to 4.3 pp. All

[7]Since around those dates many other central banks were taking similar decisions for this exercise we narrow the window considered from 5 days to 2 to reduce the probability to capture the impact of events other than the BoE announcements.

changes cause a reduction in the level of price uncertainty. The average uncertainty before the announcement on the 17th of March was around 3.25 percent and fell to 2.2 afterward.

The expectations on sales and employment were negatively adjusted in response to the interest rate hike on the 17th of March 2022. The announcement of December 2021 had negligible effects on sales, employment, and investment expectations. Finally, the interest rate cut of March 2020 is followed by a reduction in expected sales and employment. The contractionary response of the real variables could be explained by the fact that the announcement happened amidst daily news concerning the pandemic, which induced firms to reduce their expectations for sales and employment downward.

Our evidence suggests that firms are naive in forming their expectations. Interest rate *changes* garner a high level of attention from the general public and firms adjust their price expectations in their aftermath. The expectations are adjusted in a consistent manner and in the direction we would expect if we considered the rate *change* to be a close proxy to the shock. The combination of the central bank intervention and the media coverage tends to reduce the level of price uncertainty. This is true irrespective of the sign and magnitude of the high-frequency market surprise.

The most likely cause of the lack of comovement between market-based monetary policy surprises and firm expectations is that financial market surprises are a poor proxy for what constitutes a surprise to firms. This is reinforced by the finding that the results in Table 3 hinge on the large negative market surprise relative to the March 2022 policy announcement. In Table 5 we repeat the same estimation excluding the three meetings we consider in Table 4. The significant response of price expectations to large expansionary shocks disappears once we remove the large negative market surprise recorded in March 2022. It should be noted that nothing in the policy discussion surrounding the March 2022 meeting leads us to believe there was a significant *forward guidance* component to the policy decision. We ascribe the negative market surprise to markets expecting a larger increase in policy rates. On the other hand, the response of firm price expectations is consistent with firms interpreting the rate change as a tightening in the monetary policy stance.

Table 5: MPC announcements and firm expectations and uncertainty, bins approach excluding meetings with policy changes

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
Surprise bottom 1/3 x Dummy mpc	5.914	-1.203	-26.59	-36.23	5.694	6.845	130.9	101.1
	(6.952)	(10.44)	(28.18)	(22.72)	(19.57)	(8.942)	(213.9)	(76.22)
Surprise middle 1/3 x Dummy mpc	-70.06	40.07	-314.7	-153.6	-57.70**	-33.30	-1998.5	-466.9
	(69.41)	(35.00)	(266.9)	(97.74)	(22.07)	(200.5)	(1354.1)	(644.8)
Surprise top 1/3 x Dummy mpc	12.35*	-4.826	1.710	-32.94***	-25.20	17.47	-164.7**	-80.95
	(6.997)	(3.427)	(24.88)	(11.08)	(20.09)	(12.48)	(76.54)	(68.54)
Constant	2.726***	2.256***	7.818***	6.769***	1.127**	6.659***	20.57***	50.78***
	(0.153)	(0.177)	(0.593)	(0.404)	(0.416)	(0.227)	(3.618)	(1.542)
Observations	5010	5010	5953	5953	6581	6581	6347	6347
R^2	0.059	0.035	0.057	0.044	0.057	0.035	0.035	0.018
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
Number of meetings	32	32	33	33	33	33	33	33

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5.2 Robustness checks

5.2.1 Firm characteristics

We investigate whether firms heterogeneously respond to monetary policy news based on their observable characteristics. In other words, we test whether the insignificant response to market-based surprises is the result of compensating effects across different firm types.

We focus on the firms' size, i.e., the number of employees, their sector, and their age. We classify a firm as small if it has less than 50 employees, medium if between 50 and 250, and large if it employs more than 250 people. We then estimate equation (13) interacting the categorical variable of the firm size with the post-announcement dummy and the monetary surprises.

The results are reported in Table 6, where the omitted group is that of small firms. While there are systematic differences in the level of expectations across groups, there are no significant differences in the responses of price expectations.[8] The response of firm expectations to monetary policy announcements is largely independent of their size.

Table 6: MPC announcements and firm expectations and uncertainty, by size

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
Medium (50-250)	-0.524**	-0.484***	-0.808	-1.931***	-0.204	-3.158***	-5.922**	-8.834***
	(0.216)	(0.0872)	(0.556)	(0.393)	(0.296)	(0.272)	(2.652)	(0.918)
Large (above 250)	-0.853***	-0.953***	-2.905***	-3.727***	-0.572*	-5.117***	-20.58***	-24.85***
	(0.171)	(0.0984)	(0.739)	(0.348)	(0.299)	(0.245)	(3.225)	(0.870)
Surprise x Dummy mpc	5.991	0.176	3.564	2.647	-1.620	4.097	-43.69	12.87
	(4.392)	(3.055)	(6.932)	(6.817)	(3.444)	(3.906)	(31.51)	(37.77)
Medium (50-250) × Surprise x Dummy mpc	0.825	0.448	4.117	-0.773	2.781	-7.351	-38.14	-21.18
	(2.298)	(0.607)	(6.280)	(7.302)	(2.354)	(4.406)	(40.98)	(16.91)
Large (above 250) × Surprise x Dummy mpc	1.340	0.739	-6.683	-3.826**	4.462**	-5.306	7.721	-21.25**
	(1.366)	(1.461)	(8.191)	(1.667)	(2.013)	(4.014)	(23.49)	(9.525)
Constant	3.451***	2.722***	9.547***	8.837***	1.068***	9.504***	24.37***	58.83***
	(0.139)	(0.0624)	(0.386)	(0.256)	(0.195)	(0.181)	(1.739)	(0.555)
Observations	4504	4504	5361	5361	6898	6898	5793	5793
R^2	0.080	0.057	0.064	0.063	0.060	0.092	0.041	0.096
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
Number of meetings	34	34	35	35	35	35	35	35

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[8] Large firms display some peculiarity in the responses of sales and employment.

Table 7 shows the results of the same regression but interacting the announcement dummy and the monetary surprises with a dummy equal to one if the firm belongs to the financial sector. Monetary surprises have a differential effect on employment prospects only. The coefficients on all the other interactions are not significant.

Table 7: MPC announcements and firm expectations and uncertainty, finance sector

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
Finance	-1.049***	0.173	0.758	1.790***	1.615***	1.065***	-11.69**	-1.546
	(0.289)	(0.132)	(0.701)	(0.572)	(0.541)	(0.327)	(4.463)	(1.567)
Surprise x Dummy mpc	4.613	2.663	5.545	0.0982	3.551	0.665	-8.954	28.73
	(3.327)	(2.663)	(5.608)	(6.132)	(3.552)	(2.102)	(30.41)	(22.30)
Finance × Surprise x Dummy mpc	-2.209	-1.971	-12.75	2.690	-13.79***	4.777	13.49	16.36
	(5.956)	(3.454)	(11.45)	(14.31)	(3.356)	(3.729)	(61.27)	(15.62)
Constant	3.018***	2.228***	8.395***	6.804***	0.722***	6.674***	16.89***	48.10***
	(0.0216)	(0.00923)	(0.0570)	(0.0456)	(0.0450)	(0.0255)	(0.278)	(0.126)
Observations	5263	5263	6254	6254	6898	6898	6615	6615
R^2	0.056	0.021	0.047	0.027	0.051	0.017	0.032	0.011
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
Number of meetings	34	34	35	35	35	35	35	35

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Finally, we assess whether firm age drives differences in the responses of firm expectations. We interact the post-dummy variable and the monetary shocks with the age of the firm obtained from Bureau van Dijk (BvD). The coefficients from Table 8 suggest that age does not influence firms' responsiveness to the average monetary announcement. Overall, this section documents that the observable characteristics considered, i.e., size, sector, and age, play at best a minor role in explaining the responsiveness of firm expectations to monetary surprises.

Table 8: MPC announcements and firm expectations and uncertainty, by age

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
Log Age	-0.0731	-0.284***	-2.258***	-0.892***	-1.190***	-0.910***	-3.506**	-0.854*
	(0.0945)	(0.0393)	(0.256)	(0.145)	(0.144)	(0.110)	(1.592)	(0.492)
Surprise x Dummy mpc	4.037	6.666	10.43	8.144	15.30*	-1.376	-55.78	11.12
	(8.198)	(4.507)	(8.961)	(9.975)	(8.566)	(3.451)	(62.38)	(50.22)
Log Age × Surprise x Dummy mpc	0.112	-1.183	-1.659	-2.094	-4.186	0.938	16.59	5.212
	(2.526)	(0.899)	(2.750)	(2.247)	(2.553)	(0.712)	(17.29)	(12.86)
Constant	3.198***	3.138***	15.54***	9.749***	4.594***	9.559***	27.02***	50.74***
	(0.293)	(0.122)	(0.798)	(0.452)	(0.448)	(0.342)	(4.977)	(1.533)
Observations	5186	5186	6154	6154	6745	6745	6521	6521
R^2	0.069	0.046	0.065	0.048	0.069	0.043	0.034	0.019
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
Number of meetings	34	34	35	35	35	35	35	35

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5.2.2 Sampling and specification

The identification strategy we adopt crucially relies on firms responding randomly to the survey within the month. As in Bottone and Rosolia (2019), we test this assumption by plotting the predicted probability of answering before or after the MPC announcement in Figure 4. The predicted probabilities are estimated with a probit model for the event of returning the questionnaire after the monetary event on dummies for industry and size class. The two distributions are essentially identical. This suggests that the decision to submit the survey responses before or after the announcements is unrelated to the observable characteristics considered.

We then evaluate whether controlling for the lagged value of the dependent variable affects our results.[9] We report the results using the monetary surprises from Cesa-Bianchi et al. (2020) in Table 10. The time dimension of the panel is short, so controlling for a lag of the dependent variable reduces the number of observations by almost half. Most coefficients on the monetary policy surprise remain not significant, with the notable exception of that on price expectations. Even in this case, however, the result is entirely driven by the MPC meeting on the 17th of March 2022, when firms decreased their price expectations after the announcement of a rate increase that markets perceived as an expansionary shock instead. Excluding this meeting from the analysis makes the coefficient not significant.

Our empirical strategy also rests on the window size around the announcement. The smaller the window, the less likely economic news other than monetary policy announcements can pollute our estimates. At the same time, a shorter window reduces the number of respondents and increases noise.

In our baseline, the treatment group is represented by firms filing survey responses up to 5 days after the announcements. Correspondingly, firms in the control group will have filed their responses 5 days prior to the announcement. The window sizes vary quite significantly across papers. Enders et al. (2019) opt for a 2-day window, while Fiore et al. (2022) consider a 21-day window.

In our robustness checks, we consider a 2-day and a 10-day window. Reducing the size of the window from 5 days to 2 days excludes 3 monetary events which fall outside the new interval considered.[10] However, as it can be seen in Table (11), it has basically no effects on the estimated coefficients. Increasing the size of the window from 5 days to 10 days includes 8 more events that now fall inside the interval considered. Table (12) reports the results. The only remarkable result of these two estimations is the significant coefficient on price changes. But just as above, this hinges on the negative market surprise from March 2022.

Some monetary announcements occur at the beginning or at the end of the month. It might be the case that they are still included in the analysis although the size of the control and treated groups respectively is quite small. As a further robustness check, we now consider only the monetary events which happen in the middle of the month, i.e., from the 10th to

[9] We lag the dependent variable by 3 months as firms are surveyed once every 3 months.

[10] A 2-day window will rule out announcements that occur 3 or more days before the survey begins to be administered for the month.

the 20th, and for which we have both a large control and treated group. The number of announcements decreases from 34 to 19 for price expectations and 35 to 20 for the other variables. The results are shown in Table (13). The new specification does not remarkably alter the findings of the main analysis.

Finally, one might be concerned that other important announcements/releases might happen near the monetary events confounding the results. In particular, the U.K. employment rate and CPI releases by the Office for National Statistics (ONS) might affect firms' prices and employment expectations. Therefore, we create a time series of the dates when these two documents are published and we test whether firms that filled the survey after the release have different expectations from those which filled the survey before the release.

Table (14) and Table (15) report the results from the baseline specification using as control variable a dummy that equals 1 if the firm responded to the survey within 5 days after the release and 0 if the firm responded within 5 days before the release. The size of the coefficients is extremely close to zero suggesting that the ONS releases have basically no effect on firms' expectations.

6 Conclusion

To what extent central banks' announcements are able to affect expectations is critical to the transmission of monetary policy to inflation. The ability to influence expectations, and ultimately actual decisions, is considered one of the most important policy tools available to monetary authorities. However, the empirical evidence on the effects of real-world announcements on expectations is still limited.

In this paper, we study whether firms' expectations respond to monetary policy announcements from the BoE. We do so by comparing the responses to the DMP survey filled before with those after an MPC meeting. Similarly to what is documented by most of the existing literature, we show that firms' expectations do not sizeably respond to monetary policy announcements when we consider high-frequency surprises as a proxy for monetary policy. At the same time, if we focus on meetings in which rates are *changed*, then we find that firm price expectations do respond to announcements.

Our findings suggest that central banks' announcements can influence expectations. However, different economic agents might pay attention to different kinds of announcements and heterogeneously respond to them. Market-based monetary policy shocks seem to not be the best proxy for what represents a monetary policy surprise to firms. Therefore, it is crucial that monetary authorities take this into account when designing their communication strategy.

References

Altig, D., Baker, S., Barrero, J. M., Bloom, N., Bunn, P., Chen, S., Davis, S. J., Leather, J., Meyer, B., Mihaylov, E., Mizen, P., Parker, N., Renault, T., Smietanka, P., and Thwaites, G. (2020). Economic uncertainty before and during the covid-19 pandemic. *Journal of Public Economics 191 (2020) 104274.*

Altig, D., Barrero, J. M., Bloom, N., Davis, S., Meyer, B., and Parker, N. B. (2019). Surveying business uncertainty. *FRB Atlanta Working Paper.*

Binder, C., Campbell, J., and Ryngaert, J. (2022). Inflation expectations: Daily dynamics. *Working Paper.*

Bloom, Nicholas, Philip, B., Scarlet, C., Paul, M., Pawel, S., Greg, T., and Garry, Y. (2018). "Brexit and Uncertainty: Insights from the Decision Maker Panel". *Fiscal Studies 4, 555–580.*

Bloom, N., Bunn, P., Chen, S., Mizen, P., Smietanka, P., and Thwaites, G. (2019a). The impact of brexit on uk firms. *NBER Working Paper 26218.*

Bloom, N., Bunn, P., Mizen, P., Smietanka, P., and Thwaites, G. (2019b). The impact of covid-19 on productivity. *NBER Working Paper 28233.*

Bottone, M. and Rosolia, A. (2019). Monetary policy, firms' inflation expectations and prices: causal evidence from firm-level data. *Bank of Italy working paper 1218.*

Braun, R., Miranda-Agrippino, S., and Saha, T. (2022). A new dataset of high-frequency monetary policy surprises for the uk. *Mimeo, Bank of England.*

Candia, B., Coibion, O., and Gorodnichenko, Y. (2022). The inflation expectations of u.s. firms: Evidence from a new survey. *Working Paper.*

Cesa-Bianchi, A., Thwaites, G., and Vicondoa, A. (2020). "Monetary policy transmission in the United Kingdom: A high frequency identification approach". *European Economic Review 123 (2020) 103375.*

Claus, E. and Nguyen, V. H. (2020). Monetary policy shocks from the consumer perspective. *Journal of Monetary Economics 114:159-173.*

Enders, Z., Hunnekes, F., and Muller, G. J. (2019). Monetary policy announcements and expectations: Evidence from german firms. *Journal of Monetary Economics 108 (2019) 45–63.*

Fiore, F. D., Lombardi, M., and Schuffels, J. (2022). Are households indifferent to monetary policy announcements? *CEPR Working Paper.*

Gurkaynak, R., Sack, B., and Swanson, E. (2005). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *International Journal of Central Banking, 1 (1), 55–93.*

Kenneth, K. (2001). "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market". *Journal of Monetary Economics.*

Lamla, M. J. and Vinogradov, D. V. (2019). "Central bank announcements: Big news for little people?". *Journal of Monetary Economics 108 (2019) 21–38.*

Lewis, D., Makridis, C., and Karel, M. (2019). Do monetary policy announcements shift household expectations? *FRB of New York Staff Report (897).*

Link, S., Peichl, A., Roth, C., and Wohlfart, J. (2023). Information frictions among firms and households. *Journal of Monetary Economics.*

Mikosch, H., Roth, C., Sarferaz, S., and Wohlfart, J. (2022). Uncertainty and information acquisition: Evidence from firms and households. *Working Paper.*

Olivier, C., Yuriy, G., Saten, K., and Mathieu, P. (2020). "Inflation Expectations as a Policy Tool?". *Journal of International Economics.*

Peter, A., Christopher, R., Carlo, P., and Johannes, W. (2022a). Subjective models of the macroeconomy: Evidence from experts and a representative sample. *The Review of Economic Studies.*

Peter, A., Christopher, R., Haaland, I., and Wohlfart, J. (2022b). Narratives about the macroeconomy. *R & R, Review of Economic Studies.*

Pinter, J. and Kočenda, E. (2022). Media treatment of monetary policy surprises and their impact on firms' and consumers' expectations. *Forthcoming, Journal of Money Credit and Banking.*

Rast, S. (2021). Central bank communication with the general public: survey evidence from germany. *Working paper.*

Reis, R. (2020). "The People versus the Markets: A Parsimonious Model of Inflation Expectations". *Working paper.*

Singh, A. and Mitra, A. (2023). "What Determines Household Expectations?". *Working paper.*

Table 9: List of MPC meetings

Year	Month	Day	Observations
2016	11	3	74
2016	12	15	202
2017	2	2	116
2017	3	16	175
2017	5	11	234
2017	6	15	234
2017	9	14	313
2017	12	14	289
2018	2	8	271
2018	3	22	208
2018	5	10	331
2018	6	21	283
2018	9	13	595
2018	12	20	433
2019	2	7	585
2019	3	21	377
2019	6	20	333
2019	9	19	322
2019	11	7	637
2019	12	19	363
2020	5	7	348
2020	6	18	328
2020	8	6	604
2020	9	17	281
2020	11	5	546
2020	12	17	359
2021	2	4	577
2021	3	18	410
2021	5	6	602
2021	6	24	217
2021	8	5	558
2021	9	23	219
2021	11	4	522
2021	12	16	321
2022	2	3	529
2022	3	17	347

A Robustness checks

Figure 4: Predicted probability of responding after the announcement, control vs treated

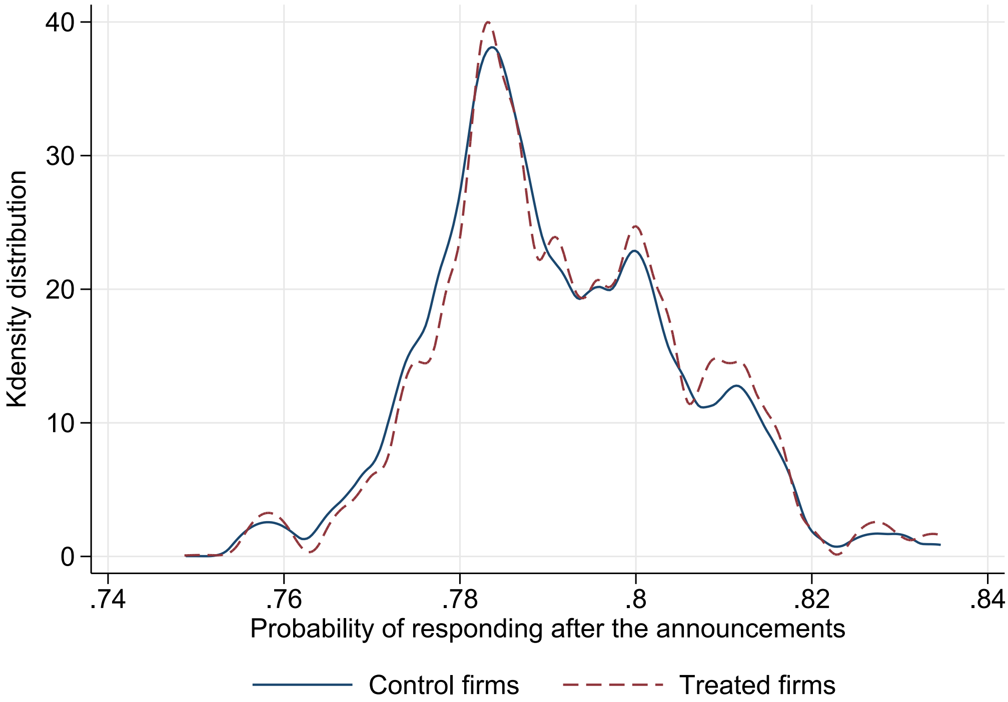

Notes: The plot shows the predicted probabilities of responding to the survey before or after a monetary announcement for the control and treated firms, i.e., those which actually filed the survey before and after the announcements. The predicted probabilities are estimated with a probit model for the event of returning the questionnaire after the monetary event on dummies for industry and size class.

Table 10: MPC announcements and firms' expectations and uncertainty controlling for lag

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
Surprise x Dummy mpc	6.168***	1.344	26.98*	-5.588	0.902	-2.907	-9.246	-9.383
	(1.737)	(1.156)	(13.77)	(5.130)	(3.866)	(2.248)	(48.90)	(13.31)
Lag expected price growth	0.485***							
	(0.0265)							
Lag price uncertainty		0.575***						
		(0.0496)						
Lag expected sales growth			0.349***					
			(0.0504)					
Lag sales uncertainty				0.623***				
				(0.0715)				
Lag expected emp growth					0.497***			
					(0.0208)			
Lag emp uncertainty						0.631***		
						(0.0392)		
Constant	1.616***	0.925***	5.424***	2.483***	0.857***	2.222***	12.29***	28.54***
	(0.0699)	(0.108)	(0.372)	(0.491)	(0.0203)	(0.264)	(0.290)	(0.926)
Observations	3188	3188	3970	3970	4652	4652	4546	4546
R^2	0.311	0.379	0.168	0.375	0.308	0.431	0.107	0.193
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
Number of meetings	34	34	35	35	35	35	35	35

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 11: MPC announcements and firms' expectations and uncertainty, 2-day window

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
Surprise x Dummy mpc	6.592*	2.487	4.924	7.517**	4.322*	-1.908	-10.44	16.18
	(3.710)	(1.510)	(8.635)	(3.638)	(2.549)	(2.079)	(37.86)	(25.33)
Constant	3.045***	2.289***	8.832***	7.142***	1.120***	6.615***	15.50***	47.64***
	(0.0161)	(0.00656)	(0.0412)	(0.0173)	(0.0103)	(0.00841)	(0.157)	(0.105)
Observations	3033	3033	3603	3603	4079	4079	3875	3875
R^2	0.072	0.039	0.052	0.042	0.048	0.041	0.035	0.022
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
Window	2 days	2 days	2 days	2 days	2 days	2 days	2 days	2 days
Number of meetings	32	32	33	33	33	33	33	33

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 12: MPC announcements and firms' expectations and uncertainty, 10-day window

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
Surprise x Dummy mpc	7.490***	3.951***	3.533	2.326	2.615	1.758	9.830	23.07
	(2.731)	(1.397)	(5.399)	(3.421)	(2.748)	(3.028)	(36.62)	(23.95)
Constant	2.917***	2.209***	8.128***	6.757***	0.812***	6.782***	15.57***	48.56***
	(0.00387)	(0.00198)	(0.00824)	(0.00522)	(0.00388)	(0.00427)	(0.0541)	(0.0354)
Observations	13402	13402	15833	15833	18067	18067	16927	16927
R^2	0.073	0.033	0.047	0.044	0.060	0.030	0.025	0.015
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
Window	10 days	10 days	10 days	10 days	10 days	10 days	10 days	10 days
Number of meetings	43	43	43	43	44	44	43	43

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 13: MPC announcements in the middle of the month and firms' expectations and uncertainty

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
Surprise x Dummy mpc	5.029	3.546*	7.883***	4.149	2.522	0.600	-8.233	26.92
	(3.008)	(2.043)	(2.623)	(5.025)	(3.399)	(1.657)	(22.51)	(20.82)
Constant	2.772***	2.205***	7.935***	6.844***	0.403***	6.621***	15.58***	48.10***
	(0.0259)	(0.0176)	(0.0252)	(0.0482)	(0.0334)	(0.0163)	(0.191)	(0.177)
Observations	3052	3052	3588	3588	4018	4018	3796	3796
R^2	0.067	0.046	0.064	0.048	0.074	0.039	0.048	0.025
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
Number of meetings	19	19	20	20	20	20	20	20

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 14: ONS employment releases and firms' expectations and uncertainty

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
ONS emp. dummy	0.00597	-0.0214	0.108	0.0145	0.00683	0.00193	-1.354	-1.356***
	(0.0681)	(0.0420)	(0.227)	(0.129)	(0.137)	(0.140)	(0.925)	(0.464)
Constant	2.754***	2.183***	7.391***	6.632***	0.561***	6.853***	14.43***	47.37***
	(0.0298)	(0.0184)	(0.0997)	(0.0566)	(0.0592)	(0.0606)	(0.406)	(0.203)
Observations	16928	16928	20035	20035	22364	22364	21850	21850
R^2	0.065	0.031	0.051	0.042	0.064	0.031	0.024	0.011
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
ONS emp. releases	65	65	65	65	65	65	65	65

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 15: ONS inflation releases and firms' expectations and uncertainty

	(1) Expected price growth	(2) Price uncertainty	(3) Expected sales growth	(4) Sales uncertainty	(5) Expected emp growth	(6) Emp. uncertainty	(7) Expected inv. growth	(8) Inv. uncertainty
ONS infl. dummy	-0.0208	-0.0384	0.0936	-0.0392	-0.0463	0.155	1.646	-1.180**
	(0.0721)	(0.0393)	(0.236)	(0.143)	(0.117)	(0.108)	(1.440)	(0.532)
Constant	2.745***	2.149***	7.072***	6.420***	0.801***	6.483***	11.30***	46.73***
	(0.0378)	(0.0206)	(0.124)	(0.0753)	(0.0603)	(0.0560)	(0.768)	(0.284)
Observations	14471	14471	17081	17081	18937	18937	18232	18232
R^2	0.061	0.034	0.043	0.041	0.056	0.024	0.020	0.015
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES
Wave FE	YES	YES	YES	YES	YES	YES	YES	YES
ONS infl. releases	68	68	68	68	68	68	68	68

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Carbon Pricing and Inflation Expectations: Evidence from France

Jannik Hensel* Giacomo Mangiante† Luca Moretti‡

Link to most recent version

April 2023

Abstract

This paper studies the impact of carbon pricing on firms' inflation expectations and discusses the potential implications for what constitutes the core of most central banks' mandate: price stability. Carbon policy shocks are identified from high-frequency changes in carbon futures price around regulatory events. The shock series is combined with French firm-level survey data. We document that a change in the price of carbon increases firms' inflation expectations. We then investigate how firms' business conditions are affected by carbon policy shocks and we find that firms' own expected and realized price growth respond similarly to inflation expectations. The effect on price expectations is more persistent than on actual price growth leading to positive forecast errors in the medium-/long-run. We also show that a sizable share of the increase in inflation expectations is due to indirect effects. Firms rely on their own business conditions to form expectations about the aggregate price dynamics. Therefore, the expected positive growth in their own prices significantly contributes to the observed increase in inflation expectations. Finally, we study how firms' responses are heterogeneously influenced by the shocks based on the share of input costs devoted to energy expenditures. We find that high energy-intensive firms tend to overreact relatively more in terms of their own price expectations compared to the actual price change the shocks induce.

Keywords: Climate policies, Carbon pricing, Inflation expectations, Monetary policy, Survey data.

JEL classification: E31, E52, E58, Q43, Q54

*University of Zurich, Zurich, Switzerland. Email: jannik.hensel@econ.uzh.ch

†University of Lausanne, Lausanne, Switzerland. Email: giacomo.mangiante@unil.ch

‡University of Zurich, Zurich, Switzerland. Email: luca.moretti@econ.uzh.ch

We would like to thank Florin Bilbiie, David Dorn, Nir Jaimovic, Jean-Paul Renne, Florian Scheuer, and Josef Zweimüller for their valuable comments and guidance. We are particularly indebted to Diego R. Känzig for providing the data on carbon policy surprises. We are also grateful to Aurélien Eyquem, Yuriy Gorodnichenko, David Hémous, Ralph Ossa, Evi Pappa, Christopher Roth, and seminar and conference participants at the UZH Macro Lunch for helpful comments. Funding for this project was provided by UZH's University Research Priority Programs Equality of Opportunity. Access to some confidential data, on which this work is based, has been made possible within a secure environment offered by CASD – Centre d'accès sécurisé aux données (Ref. 10.34724/CASD)

1 Introduction

"In short, climate change has consequences for us as a central bank pursuing our primary mandate of price stability, and our other areas of competence, including financial stability and banking supervision"

Christine Lagarde at the International Climate Change Conference (2021)

"Well, the world is running out of time to deal with the climate crisis, and the Fed has an important role to play here, and I hope the Fed will step up"

Elizabeth Warren at Jerome H. Powell's Nomination Hearing (2022)

"Given that the ECB's primary mandate is to preserve price stability, understanding the relationship between the transition to a greener economy and the price of energy is crucial"

Fabio Panetta, Member of the Executive Board of the ECB, at the Italian Banking Association (2022)

Central banks across the world have become more and more vocal about their commitment to climate change and are also facing additional pressure from policymakers to use their available toolset in such directions. Several monetary authorities have acknowledged the potential risks climate change, and in particular, the policies adopted to tackle it, pose for economic and financial stability and some of them have already adopted a more proactive role, e.g., ECB (2021). However, the empirical evidence is still limited and sometimes conflicting in their conclusion.

This paper studies the potential implications that carbon pricing has for price stability. Carbon pricing is seen as one of the most important policy tools to reduce emissions and, therefore, to mitigate the long-term shifts in temperatures and weather patterns. However, carbon pricing potentially threatens price stability which is at the core of almost every modern central bank's mandate. We document that increases in carbon prices indeed result in a rise in firms' inflation expectations as well as their own expected and realized price growth. Moreover, in the long run the effect on expectations is more persistent than on actual price changes leading to positive forecast errors. Moreover, we show that there is a direct effect of the shocks on aggregate inflation expectation but also that a significant share of the overall effect is due to indirect effects through changes in firms' own business conditions. Finally, we find that the higher the share of input costs devoted to electricity the more firms' own price expectations overreact to changes in the carbon price.

We measure exogenous changes in the carbon price using the carbon policy shock series developed in Känzig (2023). The author identifies 126 regulatory events during the period from 2005 to 2019 that influenced the supply of emission allowances in the European Union Emissions Trading System (EU ETS). The series of carbon policy surprises is computed from the change in the carbon futures price in a tight time window around the regulatory news. The surprises are then aggregated at a monthly level and used as an instrument in a proxy VAR to estimate the dynamic causal effects on the aggregate economy. The carbon policy shock series is identified from the residuals of this specification.

To evaluate how firms' inflation expectations are affected by carbon pricing policies, we combine the carbon policy shock series with French firm-level survey data. The survey, known as the Enquête de Conjoncture dans l'Industrie (ECI; "Survey of Economic Conditions in the Industry"), reports at quarterly frequency firms' inflation expectations, the expected own price growth over the next three months, and the actual price growth over the last three months. The survey is restricted to firms in the industry sector. The empirical specification we adopt is a panel local projection à la Jordà (2005).

We document that firms' inflation expectations significantly respond to carbon policy shocks. A similar effect is found for firms' own expected price growth. The responses of expected and realized price growth closely follow each other confirming that expectations translate into actual decisions. However, price forecast errors, defined as realized minus expected price growth, respond positively in the medium-/long-run suggesting that the impact of carbon policy shocks is more persistent on expectations than it is on actual price growth. We then decompose the positive response of inflation expectations into its overall and direct effect, i.e., the component of the response due to extrapolation from the firms' own business conditions. We find that the indirect effects are almost as important as the direct ones. Finally, we combine administrative balance-sheet data with the EACEI survey ("Survey on energy consumption in industries") to compute a measure of firm-level energy intensity. We document that the price expectations of high energy-intensive firms tend to overreact relatively more to carbon policy shocks.

The ability of a central bank to stabilize price growth crucially relies on its ability to control price expectations. At the same time, monetary authorities are becoming active players in tackling climate change. The findings of this paper suggest that carbon pricing is perceived by firms as inflationary. However, this does not necessarily imply that the pathway to a greener economy will cause a persistent rise in inflation. Higher taxes on fossil fuels and subsidies on green energy will impact their relative prices as well as their demand and supply. Ultimately, the overall effect on inflation will depend on the policy mix adopted.

Related literature. This paper contributes to three strands of the literature. First, the results complement the large body of empirical evidence on the effects of carbon pricing on the economy. The effectiveness of carbon pricing for emission reductions is well supported by empirical evidence (Ralf et al., 2014, Andersson, 2019). However, the impact on macroeconomic variables is still subject to debate.

Metcalf (2019) and Bernard and Kichian (2021) focus on the consequences of the British Columbia carbon tax documenting no significant impacts on GDP. Similarly, Metcalf and Stock (2020b) and Metcalf and Stock (2020a) do not find any adverse effects of carbon taxes in European countries on employment and GDP growth. Konradt and di Mauro (2023) study the potential inflationary pressure of carbon taxes in Europe and Canada and conclude that they are negligible. Moessner (2022) uses a dynamic panel estimation of New-Keynesian Phillips curves for 35 OECD economies from 1995 to 2020 and shows that an increase in prices of ETS by \$10 per ton of $CO2$ equivalents leads to an increase in energy CPI inflation by 0.8 percentage points and headline inflation by 0.08. For the California cap-and-trade market, Benmir and Roman (2022) find that carbon pricing shocks have sizable effects on the economy and result in an increase in the price of energy with negative consequences for the real economy.

The impact of carbon policies goes beyond their macroeconomic impact. The carbon policy shocks used in this paper are developed by Känzig (2023) who shows that exogenous variation in the carbon price due to regulatory events leads to an increase in inflation and a decrease in economic activity. Households along the income distribution are heterogeneously affected by the shocks mainly because of general equilibrium forces. Mangiante (2022) uses the same carbon policy shocks and documents that the real activity of poorer Euro Area countries is the most sensitive to changes in carbon price. Finally, Berthold et al. (2022) show that more carbon-intensive countries are generally more affected, $CO2$ intensive sectors do not respond differently than the green sector but within a sector, brown firms tend to suffer more. We contribute by focusing on firm-level effects. Using survey data from France we evaluate how firms' aggregate and own price expectations respond to changes in carbon price.

Second, we contribute to the literature that studies the implications of climate change and its mitigation policies for central banks. Both monetary authorities and academics are thoroughly assessing to what extent and through which channels climate change is a threat to the central banks' objective[1]. Batten et al. (2020) provide a comprehensive summary of the risks from climate change that could affect the macroeconomy and price stability.

For example, environmental disasters have been found to have large inflationary effects in emerging countries. Heinen et al. (2018) find that hurricane and flood destruction lead to an increase in consumer prices in Caribbean islands. A similar result is found by Parker (2018), who also documents heterogeneous effects across disaster types. Storms only temporarily increase food price inflation, floods also typically have a short-run impact on inflation whereas earthquakes reduce inflation excluding food, housing, and energy. Using panel local projections for 48 advanced and emerging market economies (EMEs), Faccia et al. (2021) show that hot summers increase food price inflation in the near term, especially in EMEs.

Climate change is not only a major source of concern for the central banks of developing countries. The issue is also on top of the agenda for the European Central Bank (ECB, 2021) and the members of the Executive Board (Schnabel, 2022). Moreover, modern central banks have seen an increase in public pressure to proactively contribute to the transition towards a low-carbon economy (Schoenmaker, 2021, Monnin, 2018, de Grauwe, 2019, Honohan, 2019, Lagarde, 2021, Schnabel, 2021).

We extend this literature by assessing whether carbon pricing, one of the main climate policies currently adopted, can affect price stability. We show that changes in carbon price are perceived by firms as inflationary. On top of that, firms extrapolate from the anticipated path of their own prices in forming aggregate expectations. This results in an even stronger increase in inflation expectations. Overall, our findings suggest that this climate policy potentially reduces price stability which is at the core of many central banks' mandates.

Third, this work feeds into the broader literature on inflation expectations formation. How households form their expectations about aggregate future price dynamics has been thoroughly studied in the last years[2]. The evidence on firms' inflation expectations is more scarce mainly due to limited data availability.

The empirical evidence so far suggests that firms are more similar to households than professional forecasters in forming their aggregate expectations. For the U.S., Coibion et al.

[1]See, among others, of England (2015), Carney (2015), Batten et al. (2016), of England (2018), NGFS (2020), NGFS (2021), Boneva et al. (2021)

[2]See, among others, Coibion and Gorodnichenko (2012), Coibion and Gorodnichenko (2015a), Axelrod et al. (2018), Coibion et al. (2019)

(2020b) report that disagreement in firms' inflation expectations is closer to the high levels observed for households rather than one of the professional forecasters. Candia et al. (2021) show that the inflation expectations of U.S. managers, much like those of households, are far from anchored and that the managers are largely uninformed about recent aggregate inflation dynamics or monetary policy. Kumar et al. (2015) find that firm managers in New Zealand rely on their shopping experiences as the primary determinant of their inflation expectations. Using the same survey of French manufacturing firms of this paper, Andrade et al. (2022) document that firms exploit the local prices they observe to make inferences about aggregate price dynamics despite the changes in local prices having no aggregate effects. J. et al. (2023) use data on growth expectations of German firms from the ifo Business Tendency Survey to show that firms rely on local information regarding their county, industry growth, and individual business situation when forming expectations about aggregate growth.

Households' inflation expectations have been found to be particularly sensitive to changes in gas prices[3]. This is due to the fact that gasoline is a frequently-purchased (salient) good. Households can easily observe any price changes and, given its high volatility, they tend to overestimate its importance for aggregate inflation. We extend these results to firms. We document firms' expectations strongly react to changes in carbon price and that firms rely on their own business conditions to infer the future aggregate price path.

Understanding how expectations are formed is of pivotal importance since changes in expectations affect agents' decisions and consequently their outcomes. In a series of randomized controlled trials, Coibion et al. (2019) and Coibion et al. (2020a) induce an exogenous variation in inflation expectations by providing the survey participants with different forms of information regarding inflation. The authors document that this exogenous variation has subsequent effects on household spending. With a similar empirical strategy for a survey of Italian firms, Grasso and Ropele (2018) and Coibion et al. (2020c) find that higher expected inflation is positively correlated with firms' willingness to invest, leads them to raise their prices, increase demand for credit, and reduce their employment and capital. We show that the increase in expected price growth due to changes in carbon price is closely followed by an increase in actual price growth.

Road map. The remaining paper is organized as follows. Section 2 describes the data used in this paper. In Section 3, we show the impact of carbon policy shocks on aggregate prices. Section 4 reports the results of the main analysis on firm-level data. In Section 5,

[3]See Coibion and Gorodnichenko (2015b), Cavallo et al. (2017), and D'Acunto et al. (2021)

we perform a battery of robustness checks to strengthen the validity of the baseline results. Finally, Section 6 concludes.

2 Data

2.1 Firm Level Data

The main data set used for this project is the French Outlook Survey ("ECI: Enquête de conjoncture dans l'industrie")[4]. The survey is conducted by the French economic statistics institute (INSEE) and researchers can access it after approval from the INSEE via restricted access to a secure data hub (Secure Data Access Center–CASD). It covers firms belonging to the manufacturing sector.

The survey is conducted monthly since 1992, and additional questions are asked quarterly (January, April, July, and October). Each quarter, on average 2,500 firms respond to the survey and over the sample period considered approximately 9,700 unique firms participated. The panel dimension is particularly rich since on average a firm is part of the sample for 27 quarters. Overall, for the period of interest from 1999 to 2019 our data set contains approximately 300,000 individual product-specific observations (time x firms x product) and 230,000 firm-level observations (time x firm).

The company executives are asked via postal mail or the Internet both qualitative as well as quantitative questions regarding their expectations for a variety of business-related issues such as prices, employment, production, wages, factors constraining production, and the economic outlook. Importantly, this survey also distinguishes between firm-specific questions and questions regarding aggregate measures. The most important dimension for this paper is the information about prices.

Monthly, the firms are asked about their qualitative assessment of the 3-month ahead inflation expectation (either increasing, flat, or decreasing) as well as their expectation for their own prices differentiated by individual products. Additionally, they are asked quarterly for quantitative 3-months ahead price forecasts for their own prices, as well as the quantitative price changes in the last 3 months. As shown in Figure 1, the expected price changes are positively correlated with the actual price changes in the following quarter. This suggests that the forecasts provided are of high quality as the higher the expected price growth the higher the realized price increase observed.

[4] A detailed description of the methodology of the survey can be found *here*.

Figure 1: Past and expected future price changes

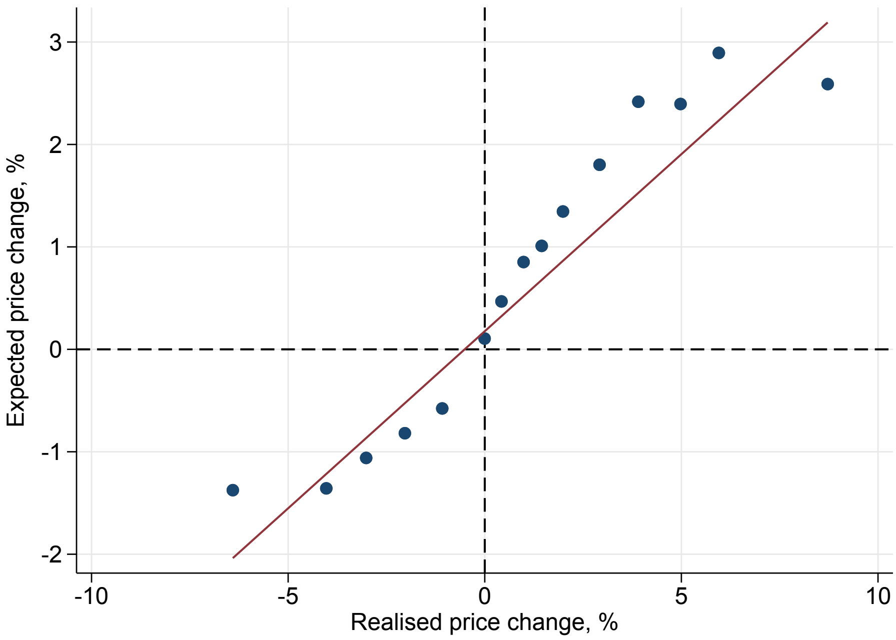

Notes: This figure shows the relationship between firms' own expected price change and the realized price growth.

Table 1 provides some descriptive statistics for the main variables of interest at quarterly frequency. The qualitative responses, i.e., the 3-month ahead expected inflation, own price growth and the realized price growth over the past 3 months which take value $\{-1, 0, 1\}$ depending on whether firms expect the variable to decrease, stay the same or increase. Andrade et al. (2022) already show that the time series of the average realized price change matches quite well the evolution of the official PPI inflation rate for France again confirming the high quality of the data. Moreover, the firms display significant heterogeneity in their forecasts of the aggregate as well as their own price growth.

2.2 Carbon Policy Shock Series

The carbon policy shocks are computed following the procedure developed by Känzig (2023) which we briefly summarize below. The main idea is similar to what has been done for monetary policy shocks (see, among others, Gürkaynak et al., 2005 and Nakamura and Steinsson, 2018). Monetary surprises are identified from changes in high-frequency asset prices around monetary policy announcements. By considering a tight window around the events,

Table 1: Descriptive statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Expected inflation	0.097	0.635	-1	1	204,936
Realised price gr.	0.038	0.509	-1	1	278,261
Expected price gr.	0.068	0.521	-1	1	249,985
Realised price gr. (Quant.)	0.11	1.97	-10	10	267,452
Expected price gr. (Quant.)	0.222	1.523	-7	8	236,393

Notes: The table reports descriptive statistics from the ECI survey on French firms for the period 1999 to 2019. The data are at quarterly frequency for the 3-month ahead inflation expectations, price growth expectations (both qualitative and quantitative) and realized price growth over the past 3 months (both qualitative and quantitative). The qualitative responses are coded as a +1 if the firm expects the variable to increase, 0 if stays the same and -1 if decreases.

the change in price can be considered unexpected and exogenous. The same methodology is applied to variations in carbon future price around regulatory events.

The European carbon market, established in 2005, operates under the cap and trade principle: a cap is set on the overall amount of certain greenhouse gases that can be emitted and, within the cap, emission allowances are auctioned off and traded in different organized markets.

Känzig (2023) identify 126 events from 2005 to 2019 concerning the overall cap in the European Union Emissions Trading System (EU ETS), the free allocation of allowances, the auctioning of allowances as well as the use of international credits. Carbon policy surprises are then computed from the changes in the futures price of the EU emission allowances (EUA) in the ICE since it is the most liquid market. In particular, the surprises are defined as the EUR change in carbon prices relative to the prevailing wholesale electricity price on the day before the event[5]. The daily surprises are then aggregated into a monthly series by summing over the daily surprises in a given month. In months without any regulatory events, the series takes zero value. The carbon policy surprise series are shown in Figure 2.

The carbon policy surprise series can be considered only a partial measure of the shock of interest due to measurement errors. To isolate the carbon policy shocks, the surprises are used as an external instrument in a VAR model with eight variables spanning the period from January 1999 to December 2019: the energy component of the HICP, total GHG emissions, the headline HICP, industrial production, the unemployment rate, the policy rate, a stock market index, as well as the real effective exchange rate. Apart from the unemployment and the policy rate, the other variables are in log levels and six lags of all variables are included.

[5]As alternative measures we also use the difference in the settlement price and its percentage change. The main results are not significantly affected by the choice of the surprise measure.

Figure 2: The carbon policy surprise series

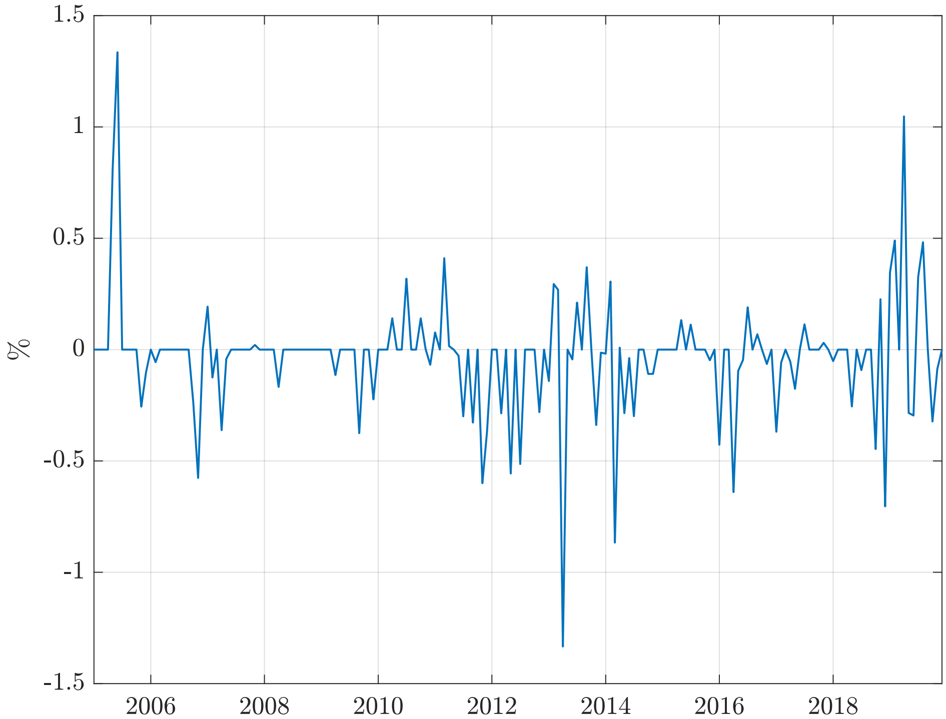

Notes: This figure shows the carbon policy surprise series, constructed by measuring the percentage change (blue solid line, left axis) as well as the change (red dashed line, right axis) of the EUA futures price around regulatory policy events.

The carbon policy shocks are then extracted from the residuals of the monthly VAR (see Stock and Watson, 2018) and are normalized to increase the energy component of the HICP by one percent on impact.

3 French Macroeconomic Variables and Carbon Policy Shocks

The Proxy-VAR used to obtain the carbon policy shock series includes macroeconomic variables for the EA-19 members. Before evaluating how carbon policy shocks affect French firms' expectations, it is important to assess the aggregate effects that these shocks have in France. To do so, we estimate the following local projection à la Jordà (2005):

$$y_{t+h} = \alpha_h + \beta_h CPShock_t + \sum_{p=1}^{P} \theta_h^p y_{t-p} + \epsilon_{t+h}, \tag{1}$$

for $h = 1, ..., 16$. y_{t+h} is the dependent variable at time $t+h$ and $CPShock_t$ are the carbon policy shocks at time t extracted from the Proxy-VAR. In the baseline specification, we include three lags of the dependent variable and we correct for autocorrelation using Newey and West

(1987) standard errors. The main dependent variables are the log of the Consumer Price Index (CPI) as well as of Producer Price Index (PPI). The coefficient of interest is β_h which captures the response of the dependent variable to a carbon policy shock for each horizon h.

Figure 3: Macro responses to carbon policy shocks

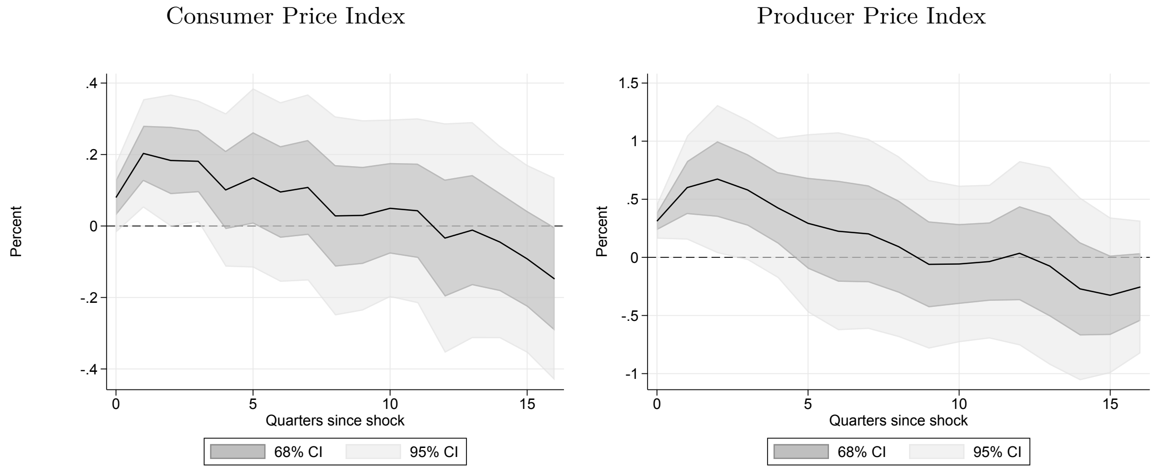

Notes: The figure plots the response to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, for the CPI (left panel) and the PPI (right panel). The dashed lines are the point estimate and the shaded areas are the 95 percent confidence bands, respectively. The horizontal axis is in quarters.

The responses of CPI and PPI to a climate policy shock are reported in Figure 3. Following a carbon policy shock that results in a one percent increase of the HICP energy component on impact, both price series significantly and persistently increase. The shock increases CPI by around 0.1 percent and PPI by 0.3 percent on impact before they slowly converge back to zero after 7/8 quarters. The inflationary effects are both statistically and economically meaningful.

In line with the findings from Känzig (2023) for the EA-19 members, the results confirm that carbon policy shocks have sizable effects at the macro level for France. We can now study whether French firms' price expectations are affected by changes in carbon price.

4 Firms' Expectations and Carbon Policy Shocks

We have shown that aggregate prices increase following a carbon policy shock. We now shift our focus from macro- to firm-level variables. Firms are asked every month about what they expect to happen to aggregate prices as well as their own prices over the next 3 months. Moreover, once every quarter firms also report the actual price change they experienced over the past 3 months. To make our results comparable we consider all the variables at quarterly frequency. The high frequency of the data and the long panel structure make it an ideal survey to study how firms' expectations are affected by changes in carbon price.

We estimate the average firm-level response to a carbon policy shock following the approach used by Andrade et al. (2022):

$$\sum_{k=0}^{h-1} \mathbb{I}\left\{E^{i}_{t+k} y^{i,j}_{t+k+1}\right\} = \alpha^{i}_{h} + \beta_h CPShock_t + \sum_{p=1}^{P} \theta^{p}_{h} X^{i,j}_{t-p} + \varepsilon^{i,j}_{t,h}, \tag{2}$$

for $h = 1, ..., 16$. $E^{i}_{t+k} y^{i,j}_{t+k+1}$ is the dependent variable, e.g., own price expectations or realized price growth, at time $t + k$ of firm i regarding its own product j. Since each firm gives a single answer to the question about the expected aggregate price change, when using inflation expectations as dependent variable the index j can be dropped and the dependent variable is equal to $E^{i}_{t+k} y^{agg}_{t+k+1}$. $\mathbb{I}\{\}$ takes value $\{-1, 0, 1\}$ depending on whether firms expect the dependent variable to decrease, stay the same or increase. α^{i}_{h} are firm fixed effect, $X^{i,j}_{t-p}$ is a matrix of controls and P is the number of lagged values[6]. Finally, standard errors are clustered at the firm level.

It is important to notice that the expectations of aggregate inflation and own price growth at monthly frequency are only qualitative. Therefore, the cumulative summation on the left-hand side can be interpreted as of the degree to which expectations respond to changes in carbon price. Due to the qualitative nature of the survey question the magnitude of the coefficient β_h does not have a direct interpretation but simply captures the share of firms that expect the dependent variable to decrease, stay the same or increase.

First, we evaluate how firms' inflation expectations are affected by carbon policy shocks. Second, we focus on firms' own price expectations, Third, we compare the effects on own price expectations with the realized price growth. Fourth, we study the price forecast errors response to assess whether firms' expectations over- or under-react to changes in carbon price. Fifth, we decompose the overall impact of carbon shocks on inflation expectations into its direct and indirect effects. Sixth, we assess whether firms heterogeneously respond based on their energy intensity level.

Figure 4: Impact of carbon policy shocks on firms' inflation expectations

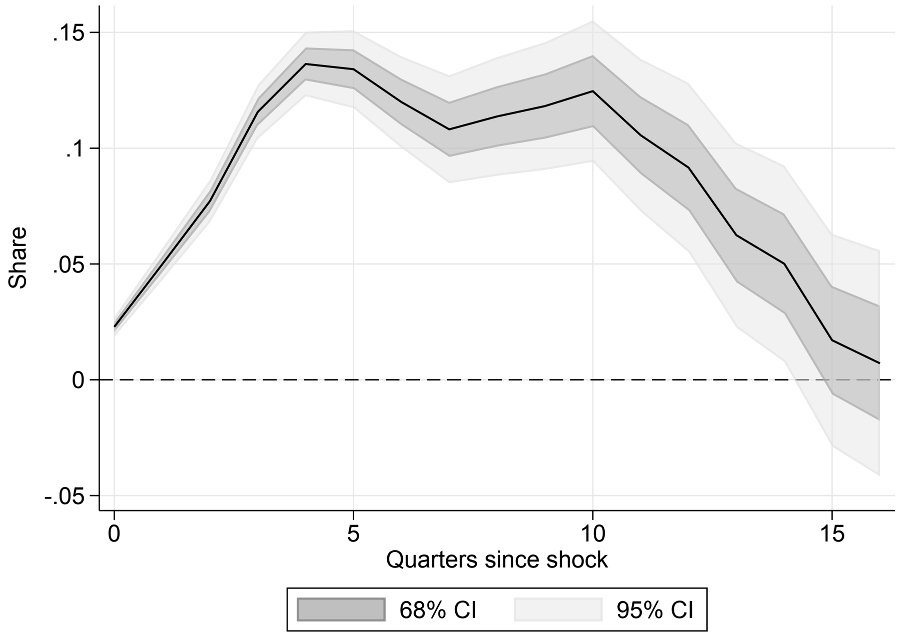

Notes: The figure plots the cumulative response to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, for the firm level inflation expectations. Inflation expectations take values {-1, 0, 1} for aggregate prices expected to decrease, stay the same or decrease. The dashed lines are the point estimate and the shaded areas are the 95 percent confidence bands, respectively. The horizontal axis is in quarters.

4.1 Inflation Expectations

The cumulative response of firms' inflation expectations is shown in Figure 4. The Figure reports the coefficients $\{\beta_h\}$ from equation (2). A carbon policy shock leads to a sizable and persistent increase in aggregate inflation expectations.

The increase in inflation expectations suggests that carbon policy might decrease price stability. Aggregate price expectations are one of the main determinants of actual inflation. Therefore, the rise in inflation expectations caused by changes in carbon price might lead to inflationary pressure on the economy. On top of that, even though the survey asks only about the 3-month inflation expectation, medium- and long-term expectations, which are the targets of the central banks, are well known to be sensitive to variations in short-term expectations (Lyziak and Paloviita, 2016). However, it is important to underline that this finding does not imply that the green transition is necessarily at odds with price stability. Changing relative prices is a desired feature of the policy. Imposing a tax on carbon is only one of the tools

[6] In the baseline specification we control for 3 lags of the dependent variable for quarterly data and for 7 lags for the monthly data.

currently available to tackle climate change and if properly complement with other policies the transition towards a greener economy and stability of prices can coexist.

4.2 Own Price Expectations and realized Price Growth

To form expectations about the evolution of aggregate prices, economic agents usually rely on personal experience even when this information is orthogonal to aggregate dynamics. For example, using the same survey of this paper, Andrade et al. (2022) show that firms' inflation expectations significantly respond to changes in industry-specific inflation rates. Therefore, changes in carbon price might not only directly increase inflation expectations but also have indirect effects due to the impact on firms' own business conditions. We study this potential channel by evaluating the response of firms' own price expectations and realized price growth to carbon policy shocks.

Figure 5: Impact of carbon policy shocks on firms' own price expectations and realized price growth (qualitative)

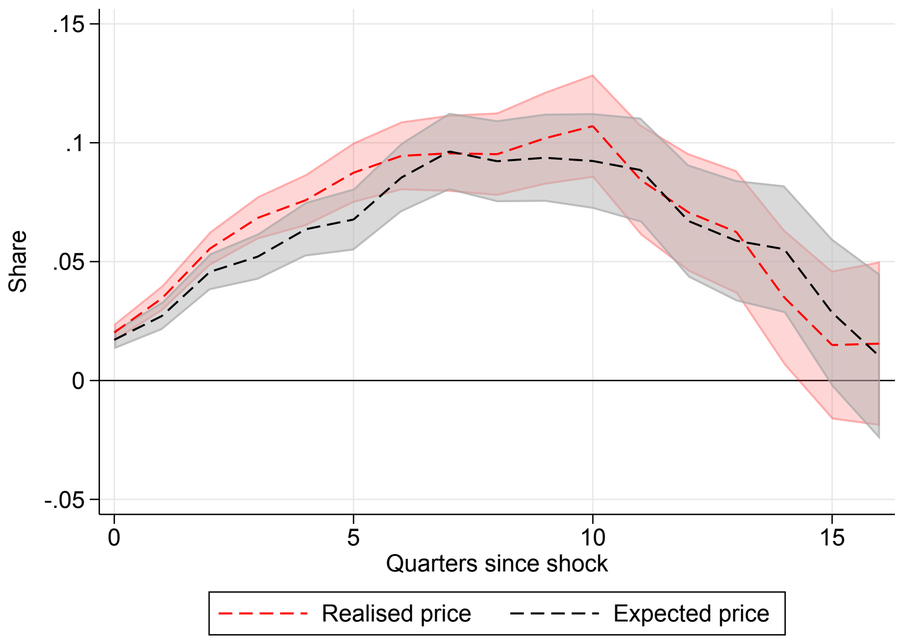

Notes: The figure plots the cumulative response to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, for the firms' own price expectations as well as the realized price growth. Price expectations take values {-1, 0, 1} for prices expected to decrease, stay the same or decrease. Realized prices take values {-1, 0, 1} based on whether prices decreased, stayed the same or decreased. The dashed lines are the point estimate and the shaded areas are the 95 percent confidence bands, respectively. The horizontal axis is in quarters.

We estimate equation (2) using the firms' expected and realized price growth from the qualitative responses as dependent variable. The cumulative responses are reported in Figure 5.

The shape and magnitude of the responses are comparable to the one of inflation expectations. It follows that changes in carbon price lead to a rise in both aggregate and firm-specific price expectations and the effect are extremely persistent over time.

Several conclusions can be drawn from the responses of expected and realized price growth. First of all, on top of the macro level, carbon policy shocks have inflationary effects at the firm level as well. The realized price growth increases in response to a change in carbon price. Second, the strong co-movement between the two responses strengthens even further the quality of the data in the survey. Firms realized price growth closely follows the expected prices confirming that the expectations they provided are on average quite precise. Third, firms' expectations are an important driver of their actual decisions: when their own price expectations increase in response to a shock, firms tend to actually raise their prices.

Figure 6: Impact of carbon policy shocks on firms' own price expectations and realized price growth (quantitative)

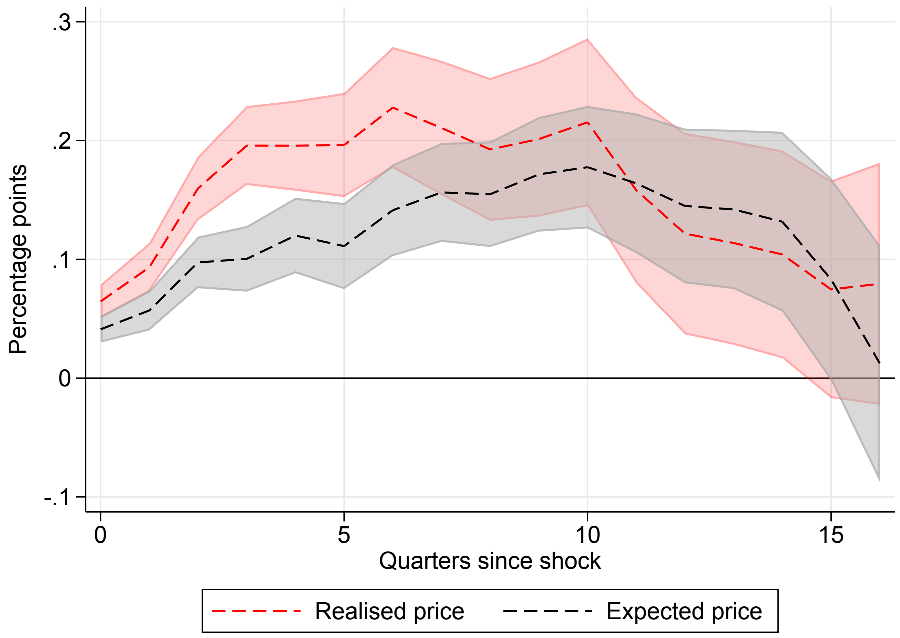

Notes: The figure plots the cumulative response to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, for the firms' own price expectations as well as the realized price growth. Price expectations and realized prices are measured as percent deviation. The dashed lines are the point estimate and the shaded areas are the 95 percent confidence bands, respectively. The horizontal axis is in quarters.

One could worry that while the share of firms expecting to raise prices and actually raising them are very similar, the actual price changes might differ significantly in magnitude. We report the responses of the quantitative variables in Figure 6. A carbon policy shock rescaled such that energy price increases by 1 percent induces an increase in expected and realized price growth of 0.05 percentage points on impact. The magnitude is comparable to the aggregate price responses we document in Section 3. The cumulative responses persistently rise up to around 0.2 percentage points after 10 quarters and then they start decreasing. The effects of carbon policy shocks on the quantitative responses show that the shocks have a significant as well as economically meaningful impact on expected and realized prices.

4.3 Price Forecast Errors

In the previous section, we have documented that the average response of the firm-level expected and realized price growth closely follow each other. However, the similar responses do not exclude that firms' expectations about the evolution of their own price either under- or over-react to carbon policy shocks when compared to the actual realization. We evaluate whether this is the case by computing the response to a carbon policy shock of price forecast errors which is defined for the quantitative responses as the difference between the expected and the realized price growth.

Figure 7: Impact of carbon policy shocks on firms' own price forecast errors

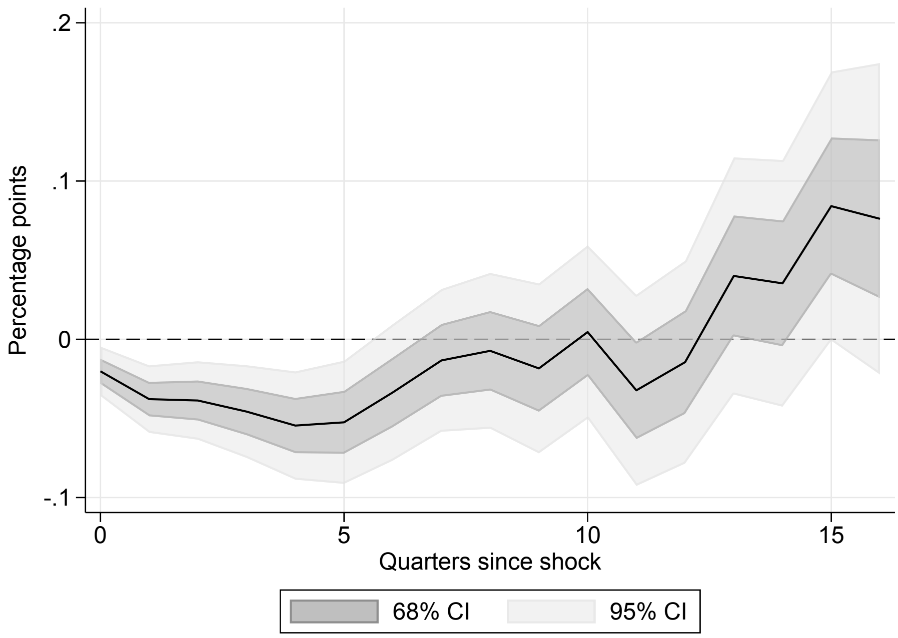

Notes: The figure plots the cumulative response to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, for the firms' own price forecast errors. Price forecast errors are measured as the difference between the expected and the realized price growth. The dashed lines are the point estimate and the shaded areas are the 95 percent confidence bands, respectively. The horizontal axis is in quarters.

The response of price forecast errors is reported in Figure 7. As one can notice, the response is initially negative. For the first five quarters, the impact of carbon policy shocks on firms' own price expectations is slightly more muted than the actual price changes they induce. For the following quarters, the forecast errors are not statistically different from zero but at the end of the time horizon considered the response turns positive. Therefore, the impact of carbon policy shocks on price expectations is more persistent than on actual price growth. With regards to the results in Andrade et al. (2022) this suggests that after this particular aggregate shock firms don't underestimate how much they are ultimately going to raise prices.

4.4 Direct vs Indirect Effects of Carbon Pricing

Carbon policy shocks have been found to sizably increase inflation expectations. Moreover, the shocks affect the industry- and firm-specific factors leading to an increase in the firms' own price expectations. Since firms tend to extrapolate from their own business conditions in forming aggregate expectations, one might expect that these indirect effects push inflation expectations even higher.

To distinguish the contribution of direct and indirect effects empirically, we follow a similar procedure as in Holm et al. (2021). We estimate two separate types of inflation expectations responses to carbon policy shocks. The first one is the baseline equation (2) which includes both direct and indirect effects. The second one is based on the same specification but also controls for the future path of the firms' expected own price[7] over the respective impulse response horizon. The estimated coefficients from the second specification capture the direct effect of changes in carbon price on inflation expectations at horizon h holding firms' expected future business conditions constant over the same time period.

The results are reported in Figure 8. The red line shows the estimated impulse response of inflation expectations without the business controls and the black dashed line shows the one with controls. The contemporaneous impact is entirely driven by direct effects. This is not surprising since the consequences of the shocks need a few months before actually materializing. After that, the two responses start to significantly diverge and the size of the overall response is around 40% larger than the size of the direct response. Therefore, a significant share of the overall impact on inflation expectations is due to indirect effects on firm-specific business conditions. The result is particularly concerning for central banks

[7] Controlling as well for the future path of production leads to similar results.

Figure 8: Direct and indirect effects of carbon policy shocks on firms' inflation expectations

Notes: The figure plots the cumulative response to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, for the firms' inflation expectations. The black line shows the estimated impulse responses controlling for price expectations, the red dashed line shows the responses without controls. The shaded areas are the 95 percent confidence bands. The horizontal axis is in quarters.

because it increases the risk that high inflation becomes entrenched even after the original shock has faded away and it would make price stability more difficult to achieve.

4.5 Heterogeneity

Firms are not homogeneously exposed to changes in energy costs. Sectoral and individual characteristics could significantly influence the propagation of an increase in carbon price to firms' expectations. For instance, one might expect that the higher the input costs devoted to energy the higher is the firm sensitivity to carbon policy shocks.

To evaluate how different degrees of energy intensity affect the propagation of shocks to expectations, we match the French survey with two additional data sources. First, the administrative balance sheet data covers the universe of French firms and provides us with information at annual frequency on the total value of the firm input costs. Second, the EACEI survey ("Survey on energy consumption in industries") reports the total expenditures by energy type. We can define different measures of energy intensity at the firm level. As a baseline measure, we compute the ratio between electricity and total input costs.

We then extend our baseline specification of equation (2) by introducing a categorical variable E_t^i which identifies different quartiles of the energy intensity distribution and which we interact with the carbon policy shock $CPShock_t$:

$$\sum_{k=0}^{h} \mathbb{I}\left\{E_{t+k}^i y_{t+k+1}^{i,j}\right\} = \alpha_h^i + \delta_h^i + \gamma_h E_t^i + \beta_h^E E_t^i CPShock_t + \sum_{p=1}^{P} \theta_h^p X_{t-p}^{i,j} + \varepsilon_{t,h}^{i,j}, \qquad (3)$$

where δ_h^i is the time fixed effects that absorb the carbon policy shocks and the aggregate variables. The coefficient β_h^E captures how firms are heterogeneously affected by the shocks according to their level of energy intensity. The interaction coefficients can be interpreted as the differential response to a carbon policy shock of the different quartiles in energy intensity relative to the baseline group (firms for which the ratio of electricity to total input costs belongs to the bottom 25%). To avoid endogeneity concerns, the categorical variable E_t^i is defined using data lagged one year. However, using contemporaneous data does not materially affect our results.

Figure 9: Impact of carbon policy shocks on firms' inflation expectations by energy intensity

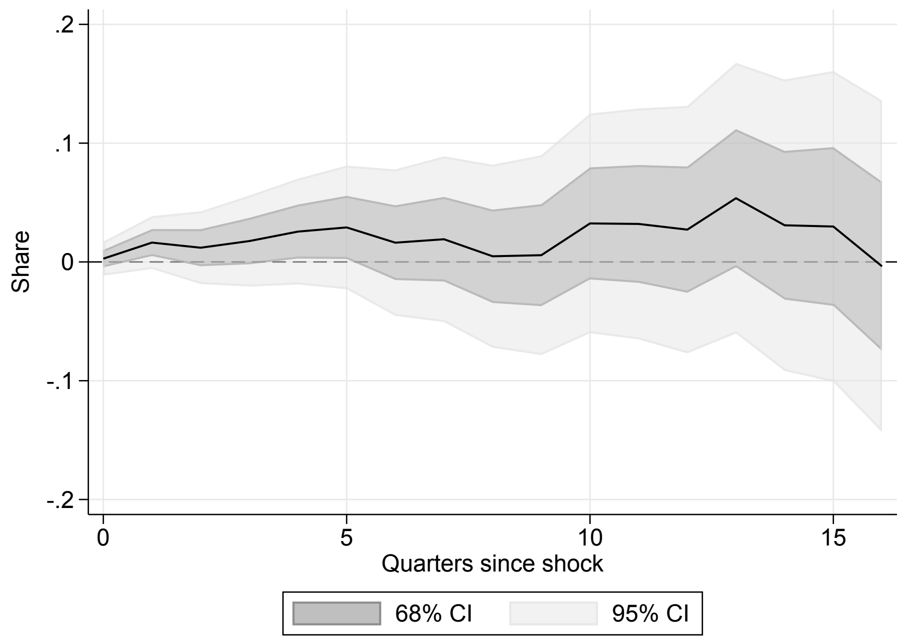

Notes: The figure plots the cumulative response to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, for the firm level inflation expectations interacted with a measure of energy intensity. Energy intensity is measured as the ratio between electricity and total input costs. Inflation expectations take values {-1, 0, 1} for aggregate prices expected to decrease, stay the same or decrease. The dashed lines are the point estimate and the shaded areas are the 95 percent confidence bands, respectively. The horizontal axis is in quarters.

We start by focusing on the impact of carbon policy shocks on inflation expectations. Figure 9 plots the coefficient β_h^E of the interaction between the shock and the top of the energy intensity quartiles. Changes in carbon price seem to influence firms' inflation expectations homogeneously along the energy intensity distribution. We do not find any statistically significant differences in the responses of the firms belonging to the top quartile relative to those at the bottom.

Figure 10: Impact of carbon policy shocks on firms' own price forecast errors by energy intensity

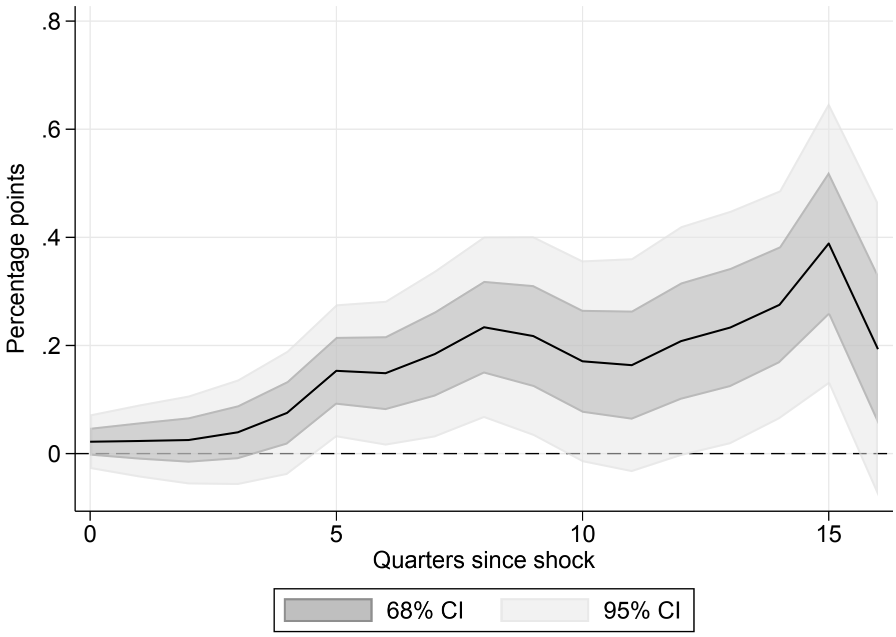

Notes: The figure plots the cumulative response to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, for the firms' own price forecast errors interacted with a measure of energy intensity. Energy intensity is measured as the ratio between electricity and total input costs. Price forecast errors are measured as the difference between the expected and the realized price growth. The dashed lines are the point estimate and the shaded areas are the 95 percent confidence bands, respectively. The horizontal axis is in quarters.

We now shift our attention to firms' own price dynamics. In Figure 10 we report the same interaction coefficient using the price forecast errors as dependent variable, i.e., expected minus realized price growth. The coefficients are positive and significant for almost the entire time horizon considered. This suggests that high energy-intensive firms tend to overreact relatively more in response to a carbon policy shock. Their own price expectations increase more compared to the actual price variation the change in carbon price induces resulting in larger price forecast errors. The effects are also economically important. Following a carbon policy shock, the forecast errors of the firms in the top quartile of the energy-intensity distribution are 0.2 percentage points higher compared to those at the bottom. Therefore, the higher the input costs devoted to energy to more firms overreact to changes in energy costs.

In conclusion, we have documented that changes in carbon policy shocks have a sizable and positive effect on inflation expectations. Moreover, firm-specific business conditions are also significantly affected leading to an increase in firms' own expected and realized price growth. In the medium-/long-run the effect on price expectations is more persistent than on the actual price growth. Moreover, the indirect effects of carbon policy shocks through changes in the firms' business conditions play a major role in the response of inflation expectations. Finally, the own price expectations of high energy-intensive firms tend to overreact to carbon policy shocks.

5 Robustness

In this section, we perform some robustness checks to strengthen the validity of the main results. First, we add extra controls to the regressions. Second, we compute the response of firms' own price expectations to carbon policy shocks only for the main product produced by the firm. The plots are reported in Appendix A.

5.1 Extra Controls

As a first robustness check, we extend the baseline specification with additional control variables. We compute the cumulative response of firms' inflation expectations to a carbon policy shock controlling as well for expected aggregate production, expected own price and production, turnovers and their respective lags. The results are shown in Figure 11.

The inclusion of controls for aggregate expectations, firms' own business conditions, and size has a negligible effect on the estimated coefficients. The magnitude and the shape of the response of inflation inequality are consistent with the baseline result. Firms' inflation expectations increase following a change in carbon price.

5.2 Price Expectations of the Main Product

In the survey, firms report the expected price growth over the next 3 months for each of their own products. In the baseline regressions, we include all these expectations. It might be the case though that firms do not pay attention homogeneously to the business conditions of each one of their products but might prioritize the most important products.

We compute the response of firms' own price expectations to a carbon policy shock only considering the product with the highest turnover. The cumulative responses are reported in Figure 12. The results are basically unaffected. Following a change in carbon price, firms' own price expectations significantly increase.

6 Conclusion

Mitigating the negative consequences of climate change is one of the most important challenges of our generation. From governments to research institutions, from households to firms, every agent in the economy is called to contribute to the reduction of greenhouse gas emissions. Monetary authorities around the world are adopting a more and more proactive role when it comes to supporting climate policies.

In this paper, we document that carbon pricing persistently increases firms' inflation expectations. This is done by combining the carbon policy shocks developed by Känzig (2023) with French firm-level survey data. We find that firms' inflation expectations are particularly sensitive to changes in carbon price. Moreover, these shocks result in an increase in firms' own price expectations as well as the ex-post realized price growth. The effect on expectations is more persistent than on actual price growth leading to positive price forecast errors in the medium-/long-run. Moreover, a significant part of the observed increase in inflation expectations is due to indirect effects, i.e., firms extrapolate from their own business conditions in forming aggregate expectations. Finally, firms that devote a higher share of input costs to energy expenditures tend to overestimate the impact that these shocks have on their prices.

Increases in the price of carbon are perceived by firms as inflationary. The empirical findings we provide suggest that carbon taxes, if not properly complemented with other green policies, might potentially be at odds with the core of the central banks' mandate, i.e., price stability. Higher short-term inflation expectations lead to higher actual prices which are likely to persist over time and propagate to longer-term inflation expectations with the risk of de-anchoring them from the inflation target. Therefore, policymakers and central bankers should carefully consider the optimal policy mix to advance the green transition without inhibiting the monetary authorities' ability to stabilize prices.

References

Andersson, J. J. (2019). "Carbon Taxes and CO2 Emissions: Sweden as a Case Study". *American Economic Journal: Economic Policy, 11(4): 1–30.*

Andrade, P., Coibion, O., Gautier, E., and Gorodnichenko, Y. (2022). "No firm is an island? How industry conditions shape firms' expectations". *Journal of Monetary Economics 125 (2022) 40–56.*

Axelrod, S., Lebow, D., and Peneva, E. (2018). "Perceptions and Expectations of Inflation by U.S. Households". *Finance and Economics Discussion Series 2018-073. Washington: Board of Governors of the Federal Reserve System.*

Batten, S., R., and Tanaka, M. (2016). "Let's talk about the weather: The impact of climate change on central banks". *Bank of England Staff Working Paper, No. 603.*

Batten, S., Sowerbutts, R., and Tanaka, M. (2020). "Climate Change: Macroeconomic Impact and Implications for Monetary Policy". *Ecological, Societal, and Technological Risks and the Financial Sector.*

Benmir, G. and Roman, J. (2022). "The Distributional Costs of Net-Zero: A HANK Perspective". *Working paper.*

Bernard, J.-T. and Kichian, M. (2021). "The Impact of a Revenue-Neutral Carbon Tax on GDP Dynamics: The Case of British Columbia". *The Energy Journal, 42(3).*

Berthold, B., Cesa-Bianchi, A., and Pace, F. D. (2022). "The economic effects of climate policies: an empirical investigation". *Forthcoming.*

Boneva, L., Ferrucci, G., and Mongelli, F. P. (2021). "To be or not to be "green": how can monetary policy react to climate change?". *ECB Working Paper, No. 285.*

Candia, B., Coibion, O., and Gorodnichenko, Y. (2021). "The Inflation Expectations of U.S. Firms: Evidence from a new survey". *NBER Working Paper 28836.*

Carney, M. (2015). "Breaking the tragedy of the horizon – climate change and financial stability". *Speech at Lloyd's of London, London, 29 September.*

Cavallo, A., Cruces, G., and Perez-Truglia, R. (2017). "Inflation expectations, learning and supermarket prices: evidence from survey experiments". *American Economic Journal: Macroeconomics, 9 (3) (2017), pp. 1-35.*

Coibion, O., Georgarakos, D., Gorodnichenko, Y., and van Rooij, M. (2020a). “How does consumption respond to news about inflation? field evidence from a randomized control trial”. *NBER Working Paper 26106.*

Coibion, O. and Gorodnichenko, Y. (2012). "What can survey forecasts tell Us about information rigidities?". *Journal of Political Economy, 120 (1) (2012), pp. 116-159.*

Coibion, O. and Gorodnichenko, Y. (2015a). “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts”. *American Economic Review 105(2015), 2644–2678.*

Coibion, O. and Gorodnichenko, Y. (2015b). “Information rigidity and the expectations formation process: A simple framework and new facts”. *American Economic Review, 105(8): 2644(78).*

Coibion, O., Gorodnichenko, Y., Kumar, S., and Pedemonte, M. (2020b). "Inflation expectations as a policy tool?". *Journal of International Economics, 124 (2020).*

Coibion, O., Gorodnichenko, Y., and Ropele, T. (2020c). "Inflation Expectations and Firm Decisions: New Causal Evidence". *Quarterly Journal of Economics 135(1): 165–219.*

Coibion, O., Gorodnichenko, Y., and Weber, M. (2019). “Monetary Policy Communications and their Effects on Household Inflation Expectations”. *NBER Working Paper 25482.*

D’Acunto, F., Malmendier, U., Ospina, J., and Weber, M. (2021). "Exposure to daily price changes and inflation expectations". *Journal of Political Economy, 129 (5) (2021), pp. 1615-1639.*

de Grauwe, P. (2019). "Green money without inflation". *Ivory Tower blog post, 26 February.*

ECB (2021). "New monetary policy strategy". *https://www.ecb.europa.eu/home/search/review/html/index.en.html.*

Faccia, D., Parker, M., and Stracca, L. (2021). "Feeling the heat: extreme temperatures and price stability". *ECB Working Paper 2626.*

Grasso, A. and Ropele, T. (2018). "Firms’ Inflation Expectations and Investment Plans". *Bank of Italy Working Paper No. 1203.*

Gürkaynak, R., Sack, B., and Swanson, E. T. (2005). “Do actions speak louder than words? The response of asset prices to monetary policy actions and statements”. *International Journal of Central Banking, 1: 55–93.*

Heinen, A., Khadan, J., and Strobl, E. (2018). "The inflationary costs of extreme weather in developing countries". *The Economic Journal 129(1).*

Holm, M. B., Paul, P., and Tischbirek, A. (2021). "The Transmission of Monetary Policy under the Microscope". *Journal of Political Economy, 129(10).*

Honohan, P. (2019). "Should monetary policy take inequality and climate change into account?". *PIIE Working Paper, No 19-18.*

J., D., L., M., and K., W. (2023). "Local Information and Firm Expectations about Aggregates". *Journal of Monetary Economics (forthcoming).*

Jordà, O. (2005). "Estimation and Inference of Impulse Responses by Local Projections". *American Economic Review, 95 (1), 161–182.*

Känzig, D. R. (2023). "The unequal economic consequences of carbon pricing". *Working paper.*

Konradt, M. and di Mauro, B. W. (2023). "Carbon Taxation and Inflation: Evidence from the European and Canadian Experience". *Journal of the European Economic Association.*

Kumar, S., Afrouzi, H., Coibion, O., and Gorodnichenko, Y. (2015). "Inflation targeting does not anchor inflation expectations: evidence from firms in New Zealand". *Brook. Pap. Econ. Activ., 46 (2 (Fall)) (2015), pp. 151-225.*

Lagarde, C. (2021). *Keynote speech at the ILF conference on Green Banking and Green Central Banking, 25 January.*

Lyziak, T. and Paloviita, M. (2016). "Anchoring of inflation expectations in the euro area: recent evidence based on survey data". *European Journal of Political Economy.*

Mangiante, G. (2022). "The Geographic Effects of Carbon Pricing". *Working paper.*

Metcalf, G. E. (2019). "On the economics of a carbon tax for the United States". *Brookings Papers on Economic Activity, 2019(1): 405–484.*

Metcalf, G. E. and Stock, J. H. (2020a). "Measuring the Macroeconomic Impact of Carbon Taxes". *AEA Papers and Proceedings, 110: 101–06.*

Metcalf, G. E. and Stock, J. H. (2020b). "The Macroeconomic Impact of Europe's Carbon Taxes". *NBER Working Paper.*

Moessner, R. (2022). "Effects of Carbon Pricing on Inflation". *CESifo Working Paper, 9563.*

Monnin, P. (2018). "Central banks should reflect climate risks in monetary policy operations". *SUERF Policy Note, No 41.*

Nakamura, E. and Steinsson, J. (2018). "High-frequency identification of monetary non-neutrality: The information effect". *The Quarterly Journal of Economics, 133(3): 1283–1330.*

Newey, W. and West, K. (1987). "A Simple, Positive-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix". *Econometrica 55:703–708.*

NGFS (2020). "Climate change and monetary policy: initial takeaways". *Technical Document, Network for Greening the Financial System, Paris.*

NGFS (2021). "Adapting central bank operations to a hotter world: reviewing some options". *Technical document, Network for Greening the Financial System.*

of England, B. (2015). "The impact of climate change on the UK insurance sector: A climate change adaptation report by the Prudential Regulation Authority".

of England, B. (2018). "The impact of adverse weather. Inflation report".

Parker, M. (2018). "The impact of disasters on inflation". *Economics of Disasters and Climate Change, 2(1), 21–48.*

Ralf, M., Preux, L. B. D., and Wagner, U. J. (2014). "The impact of a carbon tax on manufacturing: Evidence from microdata". *Journal of Public Economics, 117: 1–14.*

Schnabel, I. (2021). "From green neglect to green dominance?". *Speech given at the "Greening monetary policy – Central banking and climate change" held at the Cleveland Fed conversations on central banking, 3 March.*

Schnabel, I. (2022). "Looking through higher energy prices? Monetary policy and the green transition". *Speech at the American Finance Association 2022 Virtual Annual Meeting.*

Schoenmaker, D. (2021). "Greening monetary policy". *Climate Policy, Vol. 21, Issue 4, pp. 581-592.*

Stock, J. H. and Watson, M. W. (2018). "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments". *The Economic Journal, Volume 128 Issue 610.*

A Additional figures and tables

Figure 11: Impact of carbon policy shocks on firms' inflation expectations, extra controls

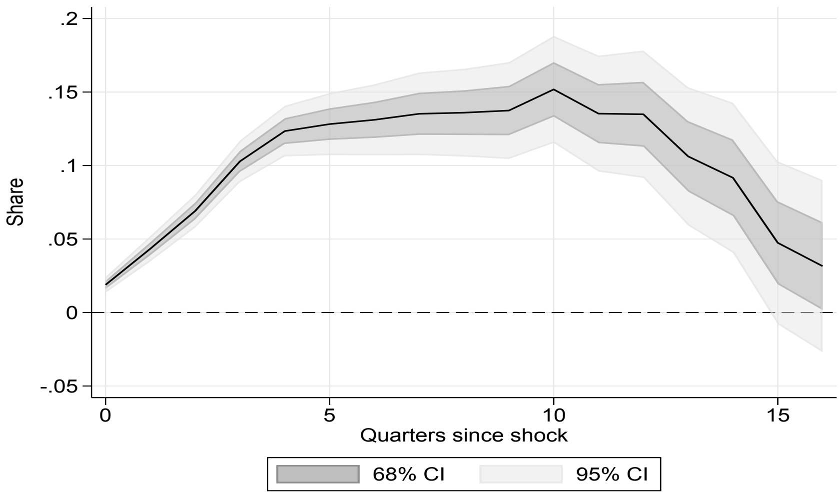

Notes: The figure plots the cumulative response to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, for the firm level inflation expectations. Inflation expectations take values {-1, 0, 1} for aggregate prices expected to decrease, stay the same or decrease. The dashed lines are the point estimate and the shaded areas are the 95 percent confidence bands, respectively. The horizontal axis is in quarters.

Figure 12: Impact of carbon policy shocks on firms' own price expectations of their main product

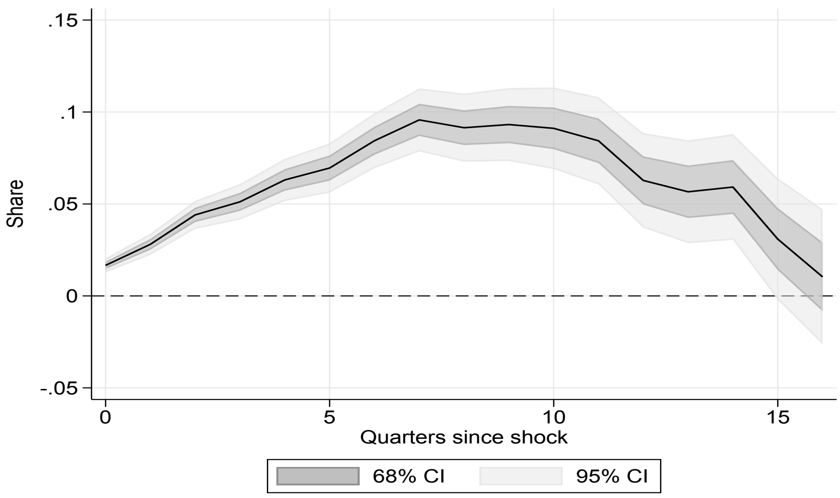

Notes: The figure plots the cumulative response to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, for the firms' own price expectations of their main product. Price expectations take values {-1, 0, 1} for prices expected to decrease, stay the same or decrease. The dashed lines are the point estimate and the shaded areas are the 95 percent confidence bands, respectively. The horizontal axis is in quarters.