European University Institute

DEPARTMENT OF ECONOMICS

Essays in Macroeconomics

Laure Simon

Thesis submitted for assessment with a view to obtaining the degree of Doctor of Economics of the European University Institute

Florence, 10 June 2020

European University Institute
Department of Economics

Essays in Macroeconomics

Laure Simon

Thesis submitted for assessment with a view to obtaining the degree of Doctor of Economics of the European University Institute

Examining Board

Prof. Evi Pappa, EUI and University Carlos III of Madrid, Supervisor
Prof. Axelle Ferrière, Paris School of Economics
Prof. Jean Imbs, NYU Abu Dhabi and Paris School of Economics
Prof. Morten Ravn, University College London

European University Institute
DEPARTMENT OF ECONOMICS

Researcher declaration to accompany the submission of written work

I, Laure Simon, certify that I am the author of the work *Essays in Macroeconomics* I have presented for examination for the PhD thesis at the European University Institute. I also certify that this is solely my own original work, other than where I have clearly indicated, in this declaration and in the thesis, that it is the work of others.

I warrant that I have obtained all the permissions required for using any material from other copyrighted publications.

I certify that this work complies with the *Code of Ethics in Academic Research* issued by the European University Institute (IUE 332/2/10 (CA 297).

Statement of inclusion of previous work (if applicable):

I confirm that chapter 2 was jointly co-authored with Alica Ida Bonk and I contributed 50% of the work.

I confirm that chapter 3 was jointly co-authored with Jean-Paul L'Huillier and Kirill Shakhnov and I contributed 33% of the work.

Signature and Date:

Laure Simon

May 12, 2020

Abstract

The first chapter uncovers a key interaction between government spending, demographics and productivity. I document that age is a key driver of consumption adjustment to government spending shocks, with significantly larger responses among young people, regardless of financial constraints. Further evidence reveals that productivity, wages and hours worked increase relatively more among young workers. I rationalize these findings with a life-cycle model where I introduce learning-by-doing. Young workers accumulate skills on-the-job at a fast rate, while the productivity of the prime-age remains stable. Then, by raising hours worked, a fiscal expansion can generate higher wage increases for young individuals, thus stimulating their consumption.

The second chapter analyzes the heterogeneous effects of government spending shocks from a gender perspective. Men typically bear the brunt of recessions due to stronger cyclicality of their employment and wages relative to women's. We study the extent to which fiscal policy may offset or worsen these asymmetric effects across genders. We find that men are hurt or benefit less than women from increases in major government spending components. This result is largely driven by negative spillovers for men working in the private sector. Furthermore, fiscal expansions cannot reconcile both policy goals: offsetting inequitable business cycle effects and closing gender gaps.

The third chapter uncovers the crucial role of the horizon in shaping the macroeconomic effects of news shocks, using a novel dataset on worldwide giant mineral discoveries. The median delay between the discovery of a mineral and its exploitation is about twice the delay reported for other commodity-discovery data considered in the literature so far, which allows to study longer-run news events. We find that macroeconomic responses to long-run discoveries are delayed. A news effect appears only two or three years before production starts, underlining an existing, but myopic, effect of these discoveries on macroeconomic expectations.

Acknowledgements

I want to express my very sincere thanks and appreciation to Evi Pappa for encouraging my research and providing personal and professional guidance with a unique combination of energetic enthusiasm, lightness, warmth, empathy and thoughtfulness; Axelle Ferriere for her persistent support, dedication, dependability, constructive criticism and friendly advice; Jean Imbs for his willingness to participate and help me in this adventure; Morten Ravn for accepting to be part of my thesis committee; Juan Dolado for providing invaluable advice and feedback in all my projects, with a touch of kindness and playfulness; my coauthors, Alica Bonk, Jean-Paul L'Huillier and Kirill Shakhnov, for sharing their knowledge, experience and many enjoyable moments, for their sincerity and support; the departmental staff, in particular Lucia Vigna, Rossella Corridori, Sarah Simonsen, Thomas Bourke, Alessandro Barucci, for their cheerfulness and superb assistance; all the professors, researchers, friends and family who offered precious encouragement and support.

Contents

Chapter 1

Fiscal Stimulus and Skill Accumulation over the Life Cycle

1.1 Introduction

There is ample evidence that changes in government expenditures affect consumers unevenly depending on characteristics such as age, income, wealth, or education.[1] Understanding the channels that drive these heterogeneous effects across households is crucial for implementing better designed and more targeted policies. It is also important for assessing how changes in the composition of the population, in terms of demographic or socio-economic characteristics, may affect the effectiveness of fiscal policy. A large body of literature argues that financial constraints play a major role in shaping these differential outcomes.[2] Specifically, financially constrained agents are more responsive to fiscal shocks as they have a higher marginal propensity to consume out of current disposable income.

In this paper, I show that age is a key determinant of consumption adjustment to government spending shocks and I propose a new transmission channel that can account for the heterogeneous effects of fiscal policy across age groups. Using household level data for the U.S. from the Consumer Expenditure Survey (CEX), I document that young individuals increase their consumption after a government spending shock while the prime-age are more negatively affected. My analysis suggests

[1]See, e.g., De Giorgi and Gambetti (2012), Anderson et al. (2016), Cloyne and Surico (2016).

[2]See, e.g., Galí et al. (2007), Jappelli and Pistaferri (2014), Kaplan and Violante (2014), Misra and Surico (2014), Brinca et al. (2016), Ferriere and Navarro (2018), Kaplan et al. (2018), Hagedorn et al. (2019).

that this result is not primarily driven by financial constraints. I also bring evidence, using micro data from the Current Population Survey (CPS), that productivity, wages and hours worked increase relatively more among young workers. I rationalize these findings with a parsimonious life-cycle model where I introduce a learning-by-doing process for young agents. Young workers accumulate skills on-the-job at a fast rate, while the productivity of their prime-age counterparts remains roughly stable. Then, by raising hours worked, a fiscal stimulus can generate higher wage increases for young individuals, thus stimulating their consumption.

I start my analysis by examining consumption responses to a fiscal expansion for different groups of households, aggregated into pseudo-cohorts according to their characteristics. Impulse responses are estimated using a VAR approach and fiscal shocks are identified with a forecast-based measure using the Survey of Professional Forecasters. Consistent with the existing literature, I find that households more likely to be credit constrained tend to raise their consumption after a government spending shock, while unconstrained households tend to lower it. However, after conditioning on age, income or wealth do not appear to be the main drivers of heterogeneity in consumption responses. Specifically, my evidence suggests that younger households strongly increase their consumption after a positive shock, while the prime-age reduce it, regardless of their income level or debt position. This finding sheds light on the importance of demographics for the transmission of fiscal policy and suggests the age-related heterogeneity in consumption responses could potentially affect the effectiveness of fiscal policy as the population ages.

I propose a new and distinct transmission channel of fiscal policy that can account for this heterogeneity across the life cycle. Fiscal policy affects more young agents because it enhances human capital accumulation. It is well documented that the age-productivity profile is steep for young workers and becomes flat for the prime-age.[3] This implies that the return to learning is high for young individuals, but falls to zero for the prime-age, which may induce age-related heterogeneity in adjustment to shocks. According to my proposed mechanism, a fiscal stimulus induces a surge in labor demand that increases hours worked. In turn, young workers raise their skill level through learning-by-doing, while the productivity of the prime-age remains roughly stable.

[3] The concave shape of the age-productivity profile, as predicted by Ben-Porath (1967) model, is documented for instance in Bowlus and Robinson (2012) who build human capital profiles using a wage-based approach. See also other studies which estimate productivity-tenure profiles using employer-employee matched data, such as Hellerstein and Neumark (1995), Hellerstein et al. (1999), Fukao et al. (2006), Hellerstein and Neumark (2007).

To investigate this channel, I explore empirically the dynamics of human capital accumulation of young and prime-age workers after a fiscal stimulus. To do so, I use micro-level data from the CPS and build a measure of age-specific productivity. I follow the wage-based approach of Bowlus and Robinson (2012) which allows me to identify the number of supplied efficiency units from the hourly wage. The impact of government spending shocks on these productivity series is then estimated using a similar VAR specification to that used in the analysis of consumption responses. I find that an increase in government spending raises significantly the productivity of young workers, while the response for prime-age workers is statistically insignificant. Turning to the responses of other labor market variables, micro evidence indicates that wages and hours worked increase relatively more for young individuals. I confirm the evidence on the effects of fiscal policy on the productivity of young workers in macro-level data using a structural vector autoregression approach for a panel of countries. After a positive shock in government spending, labor productivity significantly increases only in the group of countries with a high share of young people in the population.

To rationalize my empirical findings, I develop an overlapping generations model in the spirit of Gertler (1999) that illustrates how the different dynamics of human capital accumulation across the life-cycle can shape the heterogeneity observed. The model embeds a tractable demographic structure, with three stages of life (young, prime-age, and retiree), within a dynamic stochastic general equilibrium framework where I introduce price and wage rigidities as well as segmentation in labor markets. The key feature in this model is a learning-by-doing (LBD) mechanism for young workers, as originally developed in Chang et al. (2002) in a standard real business cycle model and subsequently extended to a New Keynesian framework by d'Alessandro et al. (2019). This mechanism implies that young workers accumulate skills as they work, increasing their labor productivity in subsequent periods. Prime-age workers, in contrast, have already reached their highest level of efficiency which remains stable.

Due to pricing frictions, a fiscal stimulus generates an increase in aggregate demand, which in turn leads to higher labor demand. Through the learning-by-doing mechanism, young workers raise their skill level. Due to wage rigidities, this translates into a greater demand of firms for young workers, thus boosting their wages. In addition, the increase in productivity, if sufficiently powerful, pushes down marginal costs, and thus expected inflation falls.[4] The Central Bank reacts by lowering the nominal interest rate, which induces a fall in the real interest rate by the Taylor principle. Therefore,

[4] The negative effect of government spending shocks on inflation was already documented in several papers. See Jørgensen and Ravn (2018) and d'Alessandro et al. (2019) for recent examples.

the fiscal stimulus operates through two channels. On the one hand, lower real interest rates stimulate consumption expenditure for all individuals via intertemporal substitution. On the other hand, the fiscal stimulus also generates asymmetric effects across age groups through redistribution, which cannot be captured in a model without heteregeneous agents. In particular, young borrowers gain from lower real interest rates, at the expense of prime-age savers, since the real value of nominal assets declines. More importantly, young workers, who primarily finance their consumption through labor income, benefit from higher wage increases. As a result, the young win, while the prime-age partly lose from the fiscal expansion.

Since most of the existing literature has focused on the importance of financial constraints to explain the heterogeneous responses to fiscal shocks, I compare the predictions of the life-cycle model with LBD to a model with young "hand-to-mouth", who fully consume their current disposable income. Both models are able to explain the increase in consumption for young individuals along with the decrease for prime-age individuals after a positive government spending shock. However, the models strongly differ regarding the effects of the fiscal expansion on wages. Specifically, in the model with LBD, the growth in wages tracks the increase in skills, and thus is more pronounced for young individuals than for the prime-age, in line with the data. In contrast, the model with young hand-to-mouth predicts that wage growth remains subdued for both young and prime-age workers.

This study provides important policy implications. Given the accelerating demographic transition towards an older population in the U.S. and other developed countries, results in this paper indicate that fiscal stimulus measures could become increasingly less effective in boosting the economy. On the other hand, policies which promote human capital formation may increase the effectiveness of fiscal policy, in particular if they are targeted at young individuals.

The remainder of the paper is structured as follows. Section 1.2 provides an overview of related literature. Section 1.3 documents the heterogeneous effects of unexpected government spending shocks on household consumption and on labor market variables across age groups. Section 1.4 introduces the life-cycle model with learning-by-doing, its parametrization and describes the transmission mechanism. Finally, Section 1.5 concludes.

1.2 Related Literature

The effects of fiscal policy shocks on household consumption are still debated, as mainstream theoretical models make different predictions and empirical evidence is mixed.[5] This paper contributes to this debate by complementing other studies which analyze the effects of government spending shocks at household level.[6] These papers typically document the importance of heterogeneity in consumption responses to fiscal policy shocks across groups of households with different characteristics, such as income, wealth and age. The most prominent explanation is the presence of liquidity constraints. Financially constrained households behave in a non-Ricardian fashion as their consumption depends on their current disposable income, while unconstrained households base their consumption decisions on their lifetime resources.[7] As a result, the former tend to increase their consumption after a government spending shock that raises labor income, while the latter reduce it due to ensuing higher taxes. In contrast, in this paper I document that age is a key determinant of household adjustment to these shocks. In particular, my results suggest that after controlling for age, financial constraints account for limited heterogeneity in consumption responses to changes in government expenditure. Furthermore, I extend my analysis to variables which have received less attention in this literature, notably labor market variables.

This paper also adds to the literature that studies the role of demographics for the transmission of fiscal policy and how it affects aggregate outcomes.[8] Basso and Rachedi (2020) exploit the heterogeneity in age structure across U.S. states to estimate the effects of government spending shocks on output and employment and how these depend on demographics. They find that higher local fiscal multipliers are associated with a higher share of young people in total population. My results, based on both micro and macro data, are broadly in line with this finding. The authors emphasize the role of credit constraints and capital-experience complementarity in explaining the link between demographics and fiscal multipliers. In this paper, I provide an alternative rationale based on the different dynamics of human capital accumulation over the life cycle and document the

[5]See Ramey (2016) for a survey of the literature.

[6]See, e.g., Johnson et al. (2006), De Giorgi and Gambetti (2012), Agarwal and Qian (2014), Jappelli and Pistaferri (2014), Misra and Surico (2014), Anderson et al. (2016), Cloyne and Surico (2016), Baugh et al. (2018), Ferriere and Navarro (2018), Zidar (2019).

[7]See, e.g., Heathcote (2005), Galí et al. (2007), Kaplan and Violante (2014), Kaplan et al. (2018) for theoretical contributions.

[8]See, e.g., Janiak and Monteiro (2016), Basso and Rachedi (2020), Fiori and Ferraro (2020) for recent contributions.

importance of this channel in shaping heterogeneous outcomes in response to government spending shocks. In particular, I argue that the evolution of skill accumulation over the life cycle is able to generate age-specific differences in labor demand, similar to capital-experience complementarity as shown in Jaimovich et al. (2013).

Lastly, this paper relates to the strand of literature which studies the interaction between fiscal policy and productivity.[9] Most of these papers document a positive response of TFP or labor productivity after an increase in government spending. A potential explanation for this finding, studied in d'Alessandro et al. (2019), is that a government spending shock induces an increase in hours worked leading to future human capital improvement. The authors build a New Keynesian model with a learning-by-doing mechanism, as originally proposed by Chang et al. (2002) in a real business cycle framework. They show that it can generate an increase in real wages, TFP and consumption in response to a government spending shock. In this paper I provide evidence of the importance of demographics in the transmission of government spending shocks to productivity. Specifically, my results suggest that the increase in productivity in response to fiscal shocks is driven by young workers. The model I develop does feature a similar skill accumulation mechanism, but I emphasize its age-dependence and introduce it within a heterogeneous agents framework with a life-cycle structure. I show that this transmission channel also leads to important redistributive effects of government spending which cannot be captured in an economy without heterogeneity.

1.3 Empirical Analysis

In this section, I investigate the presence of heterogeneous effects of government spending shocks across different groups of individuals. First, I document that age is a key driver of the heterogeneity in consumption responses to these shocks. Young households increase significantly their consumption, while the prime-age are more negatively affected. My analysis suggests that this result is not primarily driven by financial constraints. Second, I explore the role of human capital accumulation in shaping this heterogeneity. I build age-specific measures of productivity using a wage-based approach. Results indicate that the impact of government spending shocks on productivity and other labor market variables is age-dependent, with greater responsiveness among young workers.

[9] See, e.g., Evans (1992), Ramey (2011), Bachmann and Sims (2012), Aghion et al. (2014), Ben Zeev and Pappa (2015), Jørgensen and Ravn (2018), d'Alessandro et al. (2019).

1.3.1 Data Description

For government spending shocks, I use the forecast-based measure developped by Ramey (2011), based on the Survey of Professional Forecasters (SPF).[10] Government spending refers to government purchases, thus does not include transfer payments. Household level data on consumption is from the Consumer Expenditure Survey (CEX) conducted by the Bureau of Labor Statistics (BLS). Labor market data is from the Current Population Survey (CPS) Merged Outgoing Rotation Groups. The empirical analysis covers the period from 1981Q4 to 2007Q4 with quarterly data.[11]

Government spending shocks

To identify government spending shocks, I follow Anderson et al. (2016) and use a measure based on the Survey of Professional Forecasters (SPF), constructed by Ramey (2011).[12] The shock is measured as the difference between the actual real government spending growth and the forecast of government spending growth made one quarter earlier.[13] Ramey (2011) shows these shocks have good explanatory power for government spending for the recent period considered in my analysis. Since professional forecasts implicitly contain a very rich information set, including anticipated changes in fiscal policy and other economic and policy variables, using this forecast error measure allows to control for all available information and anticipated future policy actions.[14] Thus it effectively deals with the issue of fiscal foresight, namely that most government spending is anticipated by economic agents prior to implementation. Another major advantage of this approach is that the shock is directly identified using information outside the VAR. Therefore the shocks are model-independent,

[10] As showed by Ramey (2011), this measure has good explanatory power for government spending over the time period considered.

[11] The starting date of the sample is determined by the availability of SPF data. The sample is restricted to 2007 to avoid nonlinearities caused by the onset of the Great Recession and the Fed funds rate being constrained by the zero lower bound.

[12] Ramey (2011) first develops a narrative time series of estimates of changes in the expected present value of government defense spending, using information from articles in several newspaper sources such as the Business Week magazine. However, she finds that this defense news shock variable has very low predictive power if both WWII and the Korean War are excluded from the sample.

[13] Following Ramey (2011), the difference in the growth rates is preferred to the difference in the levels as the base year changed multiple times during the sample period.

[14] As showed in Coibion and Gorodnichenko (2012), although there appears to be more disagreement among households compared to other economic agents, their rate of information acquisition and processing is similar to the one of professional forecasters or firms, supporting that the forecasts from the SPF can be used as proxy for agents' expectations about government spending.

they are unaffected by potential misspecifications of the VAR or by identification assumptions. This makes it particularly appealing for estimation techniques like local projections and distributed lag models, which require a series of previously identified structural shocks.[15]

Consumption data

Household level data on consumption is from the Interview portion of the Consumer Expenditure Survey (CEX) conducted by the Bureau of Labor Statistics (BLS). The survey records information on detailed categories of consumption expenditures over the preceding quarter for all households interviewed, on demographic characteristics, as well as information on labor and financial income. Nondurable consumption expenditures are measured in log of real per capita terms.[16] I focus my analysis on the role of the following characteristics: age, defined by the age of the head of household, income level, where the measure used is household's income after taxes for the past 12 months, and housing tenure, used as proxy for households' debt position, as in Cloyne and Surico (2016). See Section 1.A for more details on this dataset and the construction of the pseudo-cohorts.

Labor market data

Individual-level data on real hourly wages and hours worked is from the Current Population Survey (CPS), which is the source of official U.S. goverment labor market statistics. The survey also contains detailed information on demographic characteristics. The extremely large sample size of this dataset allows accurate analysis at a high degree of disaggregation. Therefore I also use this dataset to construct age-specific productivity series using the wage-based approach proposed by Bowlus and Robinson (2012).[17] Individuals pursuing studies, self-employed and individuals with zero or missing wage are excluded from the sample. To construct the productivity series, I further restrict the sample to full-time male workers, i.e. who usually work at least 35 hours a week.[18] See Section 1.A for more

[15]The shocks are plotted in Figure 1.B.1.

[16]Household consumption expenditures data is divided by the number of family members and deflated by the nondurables price deflator.

[17]Bowlus and Robinson (2012) use the annual March supplement of the Current Population Survey, available since 1964, which reports households' income earned during the previous calendar year. To build productivity series at a quarterly frequency, I use instead monthly data from the CPS Merged Outgoing Rotation Group, which is available from 1979.

[18]As explained in Bowlus and Robinson (2012), females have experienced considerable changes in their labor force participation, as well as fluctuating discrimination, which raises selection issues, in particular to

details on this dataset.

1.3.2 Empirical Specification

To compute the responses to exogenous government spending shocks, I estimate the following VAR model where the shocks are explicitely treated as an exogenous variable, in line with the empirical literature which uses narrative measures of fiscal shocks.

$$X_{i,t} = \alpha_i + \beta_i \text{trend} + \gamma_i \text{qtrend} + A_i(L)X_{i,t-1} + B_i(L)u_t^G + \varepsilon_{i,t} \tag{1.1}$$

where $X_{i,t}$ is a vector of endogenous variables, α_i, β_i and γ_i control for a constant, a linear trend and a quadratic trend, $A_i(L)$ is a P-order lag polynomial, $B_i(L)$ is a (R+1)-order lag polynomial, i represents the group each household belongs to. The vector $X_{i,t}$ includes the log of real government spending per capita, the variable of interest for household group i, as well as the three month Treasury bill rate and the average marginal income tax rate from Barro and Redlick (2011) to control for monetary policy shocks and tax shocks respectively. The variable of interest alternates between the log of real nondurable consumption per capita, the log of the productivity measure, the log of real hourly wage and the log of hours worked. u_t^G denotes the series of SPF shocks. Finally, I assume that R=7 and P=2.[19] Standard errors are estimated using a wild bootstrap with 10,000 replications.

1.3.3 Heterogeneous Effects on Consumption: Age Matters

First, I examine the responses when households are grouped according to a single characteristic. A large body of literature has strongly advocated an important role for income or wealth in understanding the effects of transitory fiscal shocks on consumption behavior.[20] Households who are credit constrained or lack access to financial markets have a high marginal propensity to consume out of transitory income changes, thus raise their consumption after a government spending shock that leads to an increase in wages. In contrast, unconstrained households behave in a Ricardian fashion,

appropriately identify their flat spot regions. They are thus exluded from the sample.

[19] The results are robust to assuming longer lag structures.

[20] See, e.g., Johnson et al. (2006), Galí et al. (2007), De Giorgi and Gambetti (2012), Jappelli and Pistaferri (2014), Kaplan and Violante (2014), Misra and Surico (2014), Anderson et al. (2016), Cloyne and Surico (2016).

lowering their consumption as the net present value of their life-time resources decreases after the shock from the associated higher taxes. Figure 1.B.2, Figure 1.B.3, and Figure 1.B.4 display the responses to government spending shocks of households grouped according to their age, income level and their housing tenure respectively.[21] As can be observed, the responses are broadly in line with these predictions. Specifically, the prime-age group, highest income group and outright owners behave in a Ricardian fashion lowering their consumption, while the young group and lowest income group, more likely to be financially constrained, behave in a non-Ricardian fashion with positive consumption responses.

Next, I further investigate what drives the heterogeneity in consumption responses by splitting each income group (or housing tenure group) by age.[22,23,24] Figure 1.1 displays the responses of consumption to a government spending shock for the young (first column) and the prime-age (second column) grouped by their income level (Panel (a)) and by their housing tenure (Panel (b)), respectively. Regardless of being mortgagors or not, young households increase their consumption in response to a positive shock, while the prime-age decrease or do not adjust their consumption. The same conclusion applies when analyzing the responses of age groups by level of income. Note that the responses of young households in the low income group or in the mortgagors group are more pronounced on impact, which reflects the presence of financial constraints. The impulse responses of the ratio of consumption between young and prime-age groups split by housing tenure or income level, depicted in the rightmost column, confirm that the young tend to adjust their consumption relatively more than prime-age consumers. Therefore the consumption responses to government spending shocks display more homogeneity along the age dimension than along the income dimension or debt position.

To inspect this further, I perform a descriptive analysis of income tertiles and housing tenure groups, with a focus on the age composition of each group. The distributions of age across housing

[21] As explained in Cloyne and Surico (2016), housing tenure status is an effective proxy for household debt position. It allows to distinguish between households with mortgage debt and those without (outright owners, renters). Specifically, the authors document that for nearly half of mortgagors, their net liquid wealth represents less than 50% of their monthly income, thus they appear far more likely to be liquidity constrained.

[22] The average cell size for each group is documented in Table 1.A.1.

[23] For income level, I split households in two groups, depending on whether their after-tax income is below or above the 35th percentile. Results are robust to other income percentile splits.

[24] Given the limited number of observations per cell for young outright owners, I define a broad category of "non-mortgagors" which merges outright owners and renters. Results are robust to considering outright owners and renters separately, with less precise estimates for young outright owners.

Figure 1.1: Impulse responses to government spending shocks

(a) By age and income level

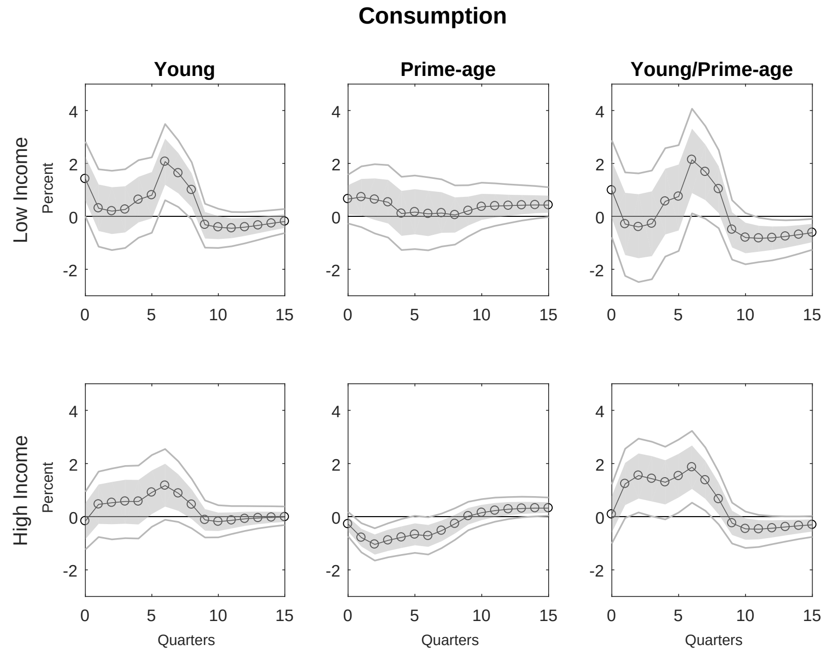

(b) By age and housing tenure

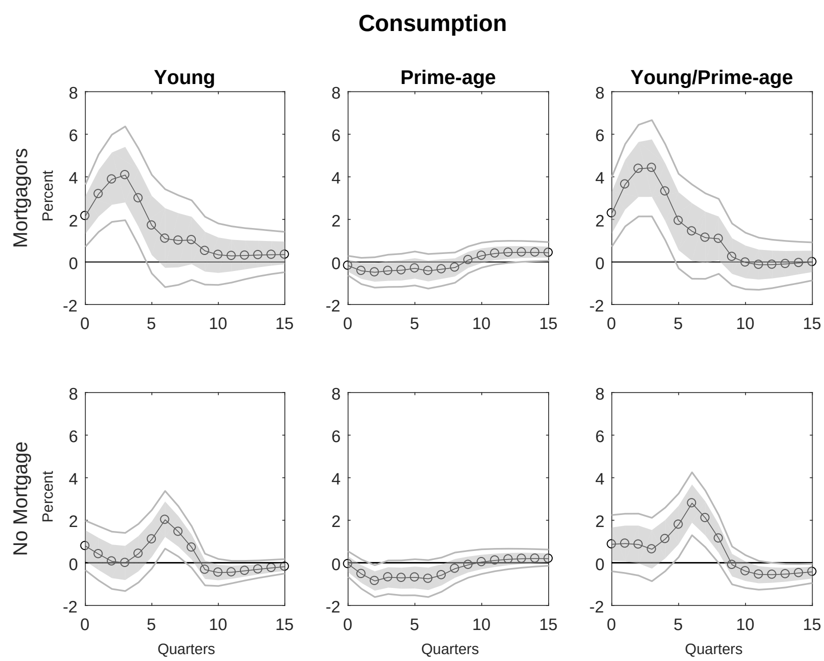

Notes: These graphs show the impulse responses of nondurable consumption for young and prime-age households by their income level in Panel (a) and by their housing tenure in Panel (b) to an exogenous government spending shock leading to an initial 1% increase in government expenditures. The impulse responses for the young (below 30) are depicted on the first column, for the prime-age (30-64) on the second column, and for the ratio between young and prime-age on the last column. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

tenure groups and income tertiles are plotted in Figure 1.B.5. Homeowners with mortgage are mainly young and working-age (25-55) while homeowners without mortgage are mainly over 55. The lowest income tertile consists mainly of young and old, while the highest tertile includes mainly working-age individuals (35-55). These statistics help explain the responses obtained when households are grouped depending on their housing tenure or their income level. For instance, positive consumption responses are found for the low income group and for renters, which are characterized by the largest share of young individuals.[25] Similarly, negative consumption responses are found for the high income group which is characterized by the largest share of prime-age households. Furthermore, the response of consumption for mortgagors is not significant as it pools together young and prime-age households who adjust their consumption in opposite ways.

To sum up, I find substantial heterogeneity in consumption responses to government spending shocks across households with different characteristics. Consistent with existing literature, households more likely to be credit constrained tend to raise their consumption after a government spending shock, while unconstrained households tend to reduce it. However, after conditioning on age, income or wealth provide less compelling explanations for the heterogeneity in consumption responses than age does. Young households increase their consumption after a government spending shock while the prime-age tend to decrease it, regardless of their income level or their housing tenure. In Section 1.3.4, I provide new evidence on a potential underlying transmission channel accounting for this heterogeneity.

1.3.4 Exploring the Role of Human Capital Accumulation

In this subsection I provide evidence that government spending shocks affects young and prime-age workers' productivity differently. My approach proceeds in two steps. First, I identify the number of supplied efficiency units from the real hourly wage for each age group. I use these efficiency time series as proxy for the productivity of workers. Then, I estimate the impact of government spending shocks on this measure of productivity for young and prime-age workers separately. I also use these efficiency series to build life-cycle productivity profiles for different cohorts.

To identify the number of efficiency units from the wage, I follow the wage-based approach of

[25] Households above 65, who represent the largest share of the low income group together with the young, increase their consumption as well in response to positive government spending shocks. The analysis of this group, mainly composed of retirees, is however outside the scope of this paper.

Bowlus and Robinson (2012), inspired by Ben-Porath (1967) model of optimal life-cycle production of human capital and its extensions.[26] As is standard in the human capital literature, the hourly wage can be defined as the product of a quantity of human capital, i.e. the number of supplied efficiency units, and its price, both unobservable. It is assumed that there are different "types" of human capital, associated with different education levels, implying different prices. At each period, the prices are assumed to be identical for all workers in a given education group, irrespectively of their age or experience. The key strategy to identify prices from observed wages, based on human capital theory, is that towards the end of working life, supplied efficiency units are constant. This implies that the change between two periods in the wage of workers on their "flat spot" only reflects changes in the price. This assumption offers a way to construct price series for each education group. Then, the worker's productivity can be calculated by dividing the real hourly wage by the price.

To construct the productivity series, I consider two education levels: low-educated are defined as workers with at most a high-school degree, and high-educated as those with some college and above. I choose the flat spot age regions for high-educated to be 48-57, and for low-educated 44-53, which is in line with Bowlus and Robinson (2012).[27] Then, to build age-specific productivity series, I compute the average productivity across the two education levels for each age group, weighted by their share in this age group. Lastly, the impact of government spending shocks on productivity is estimated for young and prime-age groups using the same specification and for the same sample period as in the previous section. I further use these productivity series to build estimates of the life-cycle human capital profiles for different cohorts.[28]

Figure 1.2 plots the estimated impulse response functions of this productivity measure, for young and prime-age workers, to a government spending shock. Productivity of the young significantly increases, with a peak of about 1.4% around 4 quarters after the shock. In contrast, the response for prime-age workers is not statistically different from zero. Thus, these results indicate that young people raise their skill level following a positive government spending shock. Figure 1.3 depicts the life-cycle profile of productivity. As predicted by the Ben-Porath model, the profile is steep for young workers, then it gradually becomes flat for prime-age workers. Interestingly, the life-cycle profile is

[26] Bowlus and Robinson (2012) extend Ben-Porath framework by incorporating two sources of cohort effects, namely selection on ability in education choices, as well as technological change.

[27] Bowlus and Robinson (2012) consider four education levels, high-school dropouts, high-school graduates, some college and college graduates, with flat spot age ranges of 44-53, 46-55, 48-57 and 50-59 respectively.

[28] To build the life-cycle profile of productivity for different cohorts, I use CPS data over the period 1979-2016.

Figure 1.2: Impulse responses of measured productivity to government spending shocks

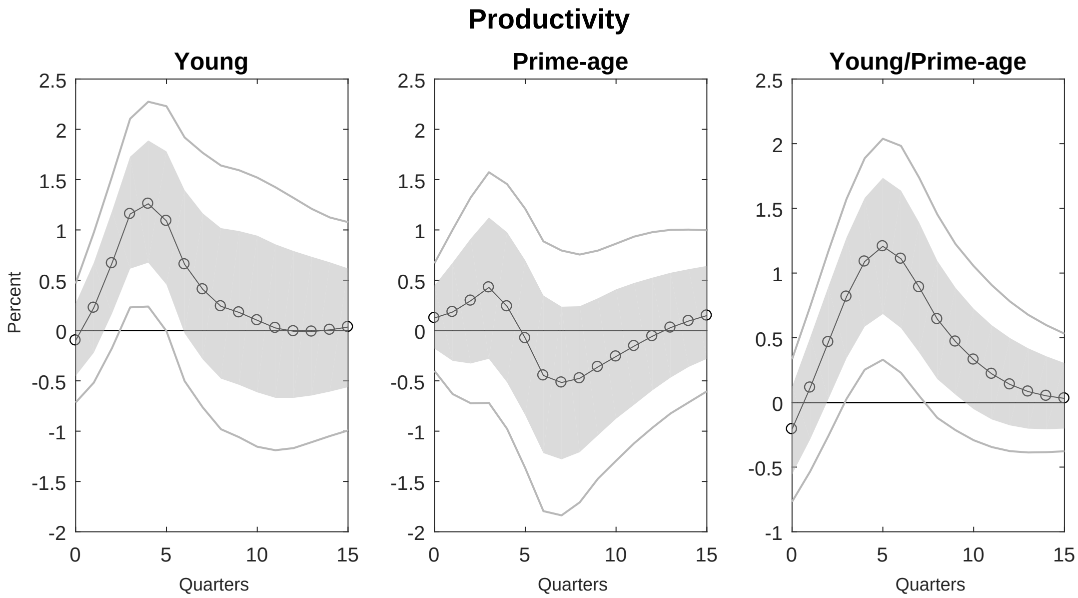

Notes: These graphs show the impulse responses of measured productivity to a 1% shock to government expenditure for young and prime-age workers, as well as for the ratio between young and prime-age. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Figure 1.3: Estimated productivity life-cycle profiles by cohort

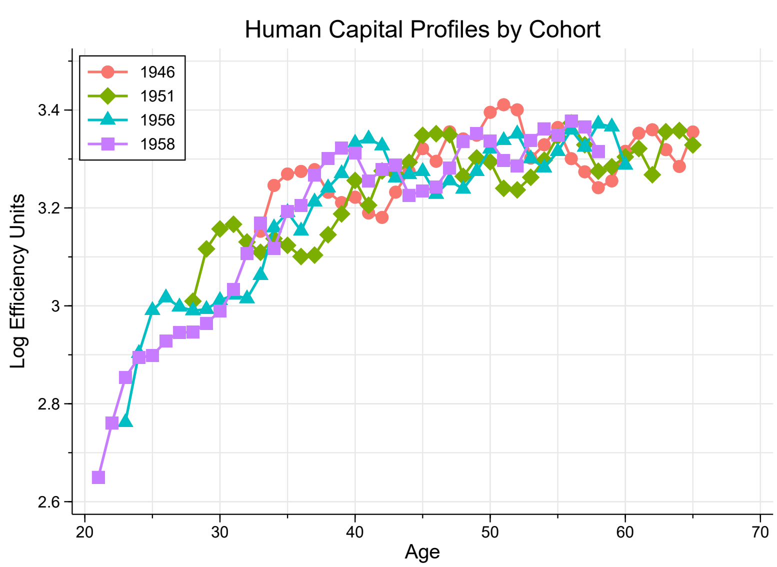

Figure 1.4: Impulse responses of wages and hours worked to government spending shocks

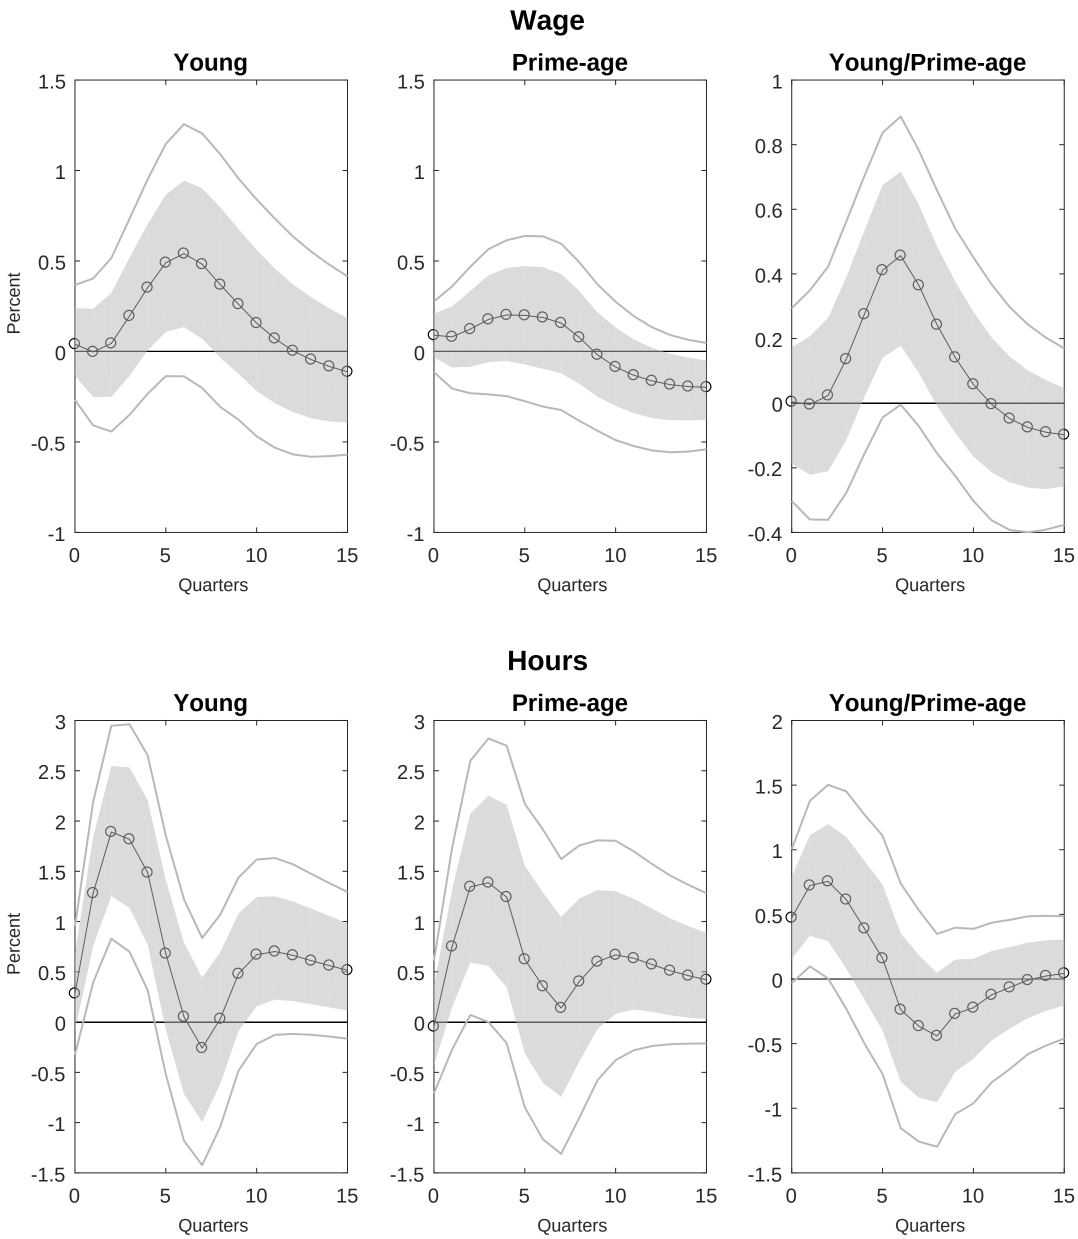

Notes: These graphs show the impulse responses of the hourly wage and hours worked for young and prime-age workers to a 1% shock to government expenditure. The impulse responses of the ratios between young and prime-age are displayed in the rightmost graphs. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

very similar across the different cohorts.

I further examine the effects on government spending shocks on wages and hours worked for young and prime-age agents. Impulse response functions of hours worked and hourly wages for young and prime-age workers are displayed in Figure 1.4. They show that young wages increase significantly at 68% level, with hump-shaped pattern, while the response of the prime-age group is not statistically significant. In addition, hours worked strongly increase in the short-run for young people and, to a lesser extent, for the prime-age. The impulse response functions of the ratios further confirm the greater responsiveness of productivity, wage and hours worked to positive government spending shocks for young workers.

1.3.5 Robustness and Extensions

Heterogeneous effects on consumption

Results documented in Section 1.3.3 hold when considering other proxies for financial constraints, such as educational attainment and financial market participation. As proxy for financial market participation, I define a dummy variable that takes the value one for households that report non-zero financial wealth from savings accounts, stocks, bonds, mutual funds or other financial assets, and zero otherwise.[29] In addition, the results are also robust to using a broader definition of consumption that includes purchases of small durables, imputed services from vehicles, rents, imputed rents for home owners, mortgage payments, pensions, and cash contributions. I also considered restricting the sample to employed households, for which added hours worked of head and spouse are strictly positive. The results are similar.

A concern when splitting households according to their housing tenure or income level is the possibility of endogenous changes in group composition over time. To tackle this issue, I consider another grouping estimator using the propensity score approach proposed in Attanasio et al. (2002), which allows to hold the composition of the group constant. Specifically, for each age group the predicted probability of being a mortgagor (or having a low income) is estimated based on variables that are fully predictable from one period to the next, namely age and education. Results are not overturned. I also estimated the impulse responses of the group shares to an unexpected government

[29]Using non-zero income from financial assets instead of non-zero financial wealth leads to similar results.

spending shock. The estimates suggest that changes in government spending do not lead to significant changes in group composition. All related figures can be found in Section 1.D.1.

Further robustness

I perform additional robustness checks of the main findings regarding age-dependent effects of government spending shocks on consumption, productivity, wages and hours worked. I check the sensitivity of the results to another measure of unexpected government spending shocks that explicitely controls for information provided by several macroeconomic variables and for other forecasts. Specifically, I regress the real government spending growth on the one-quarter ahead forecast of government spending growth and lags of government spending, tax revenues, output and unemployment. I also include in the regressors the one-year ahead forecast of annual output growth as government spending forecasts can be driven by expected changes in output. The shocks are defined as the estimated residuals from this regression. In addition, I check if the results are affected by the lack of control for inflation and for variables capturing the state of the economy or the degree of labor market slack. Thus I augment the VAR with additional variables, namely output growth, unemployment and inflation. Impulse responses are estimated using the same dynamic specification as in the baseline. They are qualitatively similar to benchmark results. All related figures are reported in Section 1.D.3. I also check that the results based on the age-specific productivity measures are robust to alternative specifications of the flat spot ranges (see Section 1.D.2).

1.3.6 Further Evidence Using Macro Panel Data

I provide further evidence on the age-dependent effects of government spending shocks on consumption and productivity using a structural vector autoregression approach for a panel of countries. For this analysis, I use the unique quarterly dataset compiled by Ilzetzki et al. (2013), covering government expenditure, output, consumption and other macroeconomic variables for 44 developing and developed countries from 1960 to 2007. I complement it with series on labor productivity and demographic data on the shares of young people (aged 15-29) in total population.[30] I estimate the following SVAR model:

[30] See Section 1.A for more details on the data.

$$AX_{j,t} = \sum_{k=1}^{K} C_k X_{j,t-k} + BU_{j,t} \tag{1.2}$$

where $Y_{j,t}$ is a vector of endogenous variables in country j at quarter t that consists of real government consumption expenditure, real GDP, real private consumption and labor productivity. Following Ilzetzki et al. (2013), government spending shocks are identified using Blanchard and Perotti (2002) identification strategy, and the model is estimated by panel OLS regression with fixed effects, with four lags included. This identification hinges on the assumption that there is no response of government spending to changes in other macroeconomic variables within a quarter due to decision and implementation lags.

To inspect the role of demographics, I split the panel of countries in two groups, characterized by shares of young people in total population above and below the sample mean. The VAR model is then estimated for the two groups separately in order to compare the impulse response functions of productivity and private consumption to government spending shocks in the two groups of countries.

As can be observed in Figure 1.5, countries with high share of young people display very different responses of productivity and consumption to government spending shocks compared to countries with low share of young people. Specifically, in the group of countries with high share of young people, there is a strong and significant increase in both productivity and consumption after a positive shock in government spending. In contrast, in the group with low share of young people, productivity remains virtually unchanged while consumption significantly drops after the shock. These results are in line with findings from the micro level analysis reported in Sections 1.3.3 and 1.3.4, confirming the importance of demographics in the transmission of government spending shocks to productivity and aggregate variables. Results are similar when the current account and the real exchange rate are included as controls (see Section 1.D.2).

1.4 A Life-Cycle Model with Learning-By-Doing

In this section, I build a New Keynesian DSGE model with a parsimonious life-cycle structure that rationalizes the empirical findings documented in Sections 1.3.3 and 1.3.4. The key feature is a learning-by-doing (LBD) mechanism which I introduce for young workers. The model illustrates how

Figure 1.5: Impulse responses of productivity and consumption to government spending shocks in countries with high vs. low shares of young people in total population

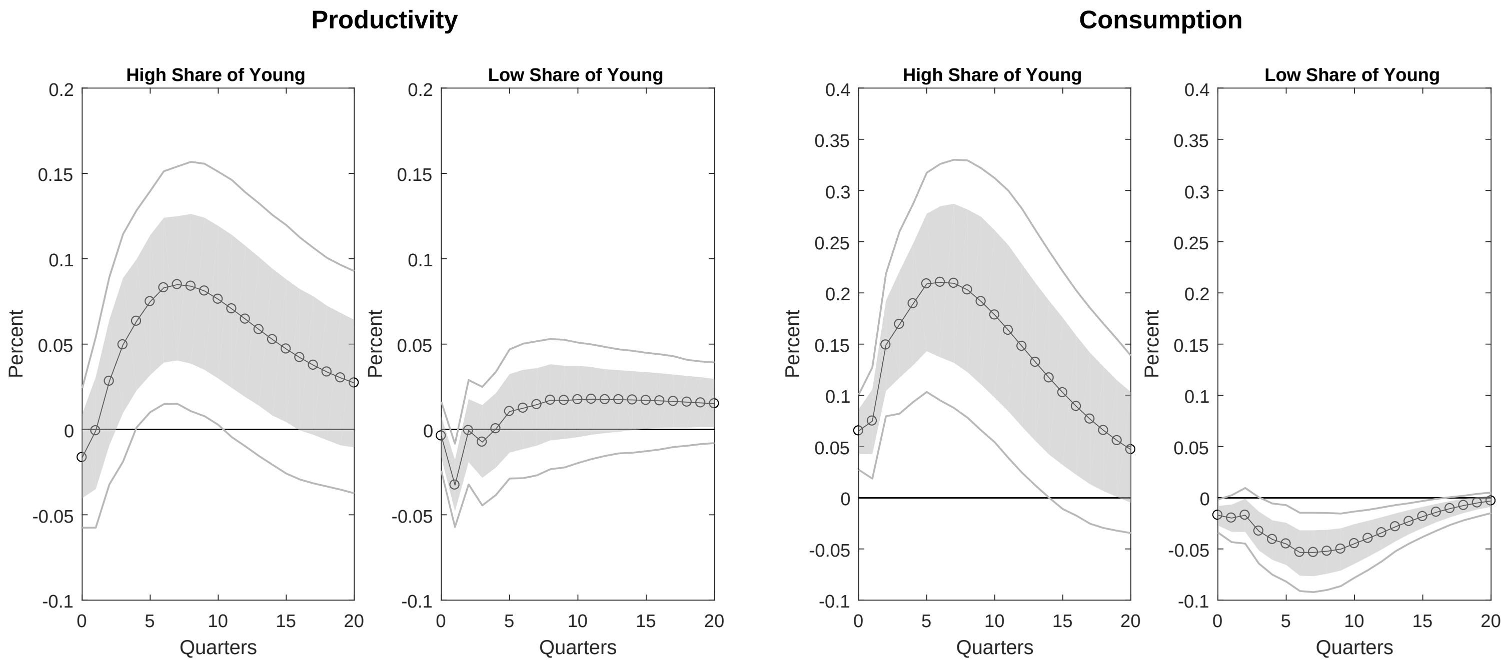

Notes: These graphs show the impulse responses of labor productivity (left panels) and private consumption (right panels) to a 1% shock to government consumption expenditure in countries with low share of young people (aged 15-29) vs. high share of young people in total population. 90% and 68% confidence intervals are shown in all cases.

different dynamics of human capital accumulation across the life cycle lead to redistribution effects of fiscal stimulus that differentially affect young and prime-age workers' incentives to consume.

1.4.1 Model

The model features a life-cycle structure in the spirit of Gertler (1999), where individuals face three stages of life: young, prime-age and retirement. Prime-age individuals cannot insure against retirement risk and retirees against longevity risk. The incompleteness of financial markets leads to life-cycle saving behavior, with prime-age workers accumulating assets to finance consumption during retirement. On the other hand, young agents borrow as they expect higher income when prime-age. I further incorporate life-cycle human capital accumulation. Young workers accumulate skills on the job while prime-age workers have reached their highest level of efficiency which remains stable. To model skill accumulation, I specify a learning-by-doing mechanism as originally proposed by Chang et al. (2002) in a real business cycle model and adapted recently to a New Keynesian framework by d'Alessandro et al. (2019).

As is standard in the New Keynesian literature, the supply side of the economy consists of a continuum of firms under monopolistic competition facing staggered price setting à la Calvo. They produce differentiated intermediate goods that are used as inputs by a perfectly competitive firm to produce a final good for private and public consumption. Wages are set by representative unions for young and prime-age workers on segmented labor markets, subject to adjustment costs. The central bank sets the nominal interest rate following a standard Taylor rule. The Treasury finances its expenditures by issuing one-period bonds and collecting lump-sum taxes from young and prime-age individuals.

Aggregation is typically challenging in this type of heterogeneous agents models, notably because the wealth distribution responds endogenously to aggregate shocks. Gertler (1999) proposes a tractable overlapping generations setup which allows to derive closed-form aggregate consumption and savings functions while preserving life-cycle behavior. However, this framework requires the use of specific nonexpected utility preferences and the assumption of risk neutrality. To make the model more flexible, I incorporate a transfer to new young and to new prime-age agents that aims at removing heterogeneity in wealth among each age group. This strategy, motivated by evidence that heterogeneity in wealth plays a smaller role than age in the transmission of government spending shocks, allows in particular to incorporate wage rigidity and to specify standard preferences. Although there is no within-group heterogeneity in this model, heterogeneity across the age dimension, i.e. between-group, is preserved. The model can be solved with standard linearization methods, using the certainty equivalence property of the first order approximation.[31]

Households and Life-cycle Structure

The economy is populated by a continuum of households, who belong to three different age groups: young (y), prime-age (p) and retiree (r). At each period, young agents face a constant probability of becoming prime-age ω_p. Similarly, prime-age households face a constant probability of becoming

[31] Details of the derivations are provided in Section 1.F.

retiree ω_r, and retirees face a time-invariant death probability ω_x. The population size is normalized to one and its composition remains constant over time.

The share of each age group in total population can be computed using the fact that the number of new prime-age agents is equal to the number of prime-age retiring, i.e. $\omega_p N_y = \omega_r N_p$, and the number of prime-age retiring is equal to the number of deaths in the economy, i.e. $\omega_r N_p = \omega_x N_r$, where N_y, N_p and N_r are the number of young, prime-age and retirees, respectively. Therefore, denoting ν_y, ν_p and ν_r the shares of each age group, we get:

$$\begin{aligned} \nu_y &= \frac{1}{1+\frac{\omega_p}{\omega_r}+\frac{\omega_p}{\omega_x}} \\ \nu_p &= \frac{1}{1+\frac{\omega_r}{\omega_p}+\frac{\omega_r}{\omega_x}} \\ \nu_r &= 1-\nu_y-\nu_p \end{aligned} \tag{1.3}$$

Individual i in age group j derives utility from consumption $C^i_{j,t}$ and disutility from hours worked $L_{j,t}$. In period t, this individual chooses consumption $C^i_{j,t}$ and asset holdings $A^i_{j,t+1}$ which solve the following optimization problem

$$\max \quad V^i_{j,t} = \left(\log(C^i_{j,t}) - \chi_j \frac{{L^i_{j,t}}^{1+\varphi_j}}{1+\varphi_j} \mathbb{I}_{\{j=y,p\}} + \beta \mathbb{E}\left(V^i_{j',t+1} | j \right) \right)$$

s.t.

$$\begin{cases} P_t C^i_{j,t} + B^i_{j,t} = A^i_{j,t} + W^i_{j,t} L^i_{j,t} - P_t \tau^i_t + \text{bq}^i_{j,t} + P_t {\tau^G_{NY,t}}^i & \text{if } j = \{y\} \\ P_t C^i_{j,t} + B^i_{j,t} = A^i_{j,t} + W^i_{j,t} L^i_{j,t} - P_t \tau^i_t + (1-\tau_d)\text{div}^i_t + P_t {\tau^G_{NP,t}}^i & \text{if } j = \{p\} \\ P_t C^i_{j,t} + B^i_{j,t} = A^i_{j,t} & \text{if } j = \{r\} \end{cases} \tag{1.4}$$

$$A^i_{j,t+1} = (R_{n,t} + \zeta \mathbb{I}_{\{j=y\}}) B^i_{j,t}$$

where β denotes the subjective discount factor, φ_j the inverse of the Frisch elasticity of labor supply, and χ_j the weight of the disutility of labor.

Households have access to bonds $B_{j,t}$ which yield a nominal return given by $R_{n,t}$, where $R_{n,t}$ is the gross nominal interest rate. To avoid overborrowing from young agents, it is assumed that they face a constant risk premium ζ. P_t is the price level. Young and prime-age individuals supply labor services to firms for a nominal wage $W_{j,t}$. Workers take wages and hours as given. Wages are fixed by unions, and hours worked are determined by intermediate goods firms' labor demand. Retirees have no labor income and consume only out of asset income. The wealth of deceased retirees is equally distributed as bequests $\text{bq}^i_{j,t}$ among young individuals. Prime-age agents earn nominal dividends $\text{div}_{j,t}$, taxed at proportional rate τ_d, from imperfectly competitive intermediate firms. Young and prime-age individuals pay the same amount of lump-sum taxes. Finally, the newborns and new prime-age individuals receive government transfers ${\tau^G_{NY,t}}^i$ and ${\tau^G_{NP,t}}^i$, respectively.[32]

Firms

The supply side of the economy is composed of a continuum of firms under monopolistic competition which produce differentiated intermediate goods, indexed by $z \in [0, 1]$, that are used as inputs by a perfectly competitive firm to produce a final good.

Final goods firm. The production of final goods by the representative firm is given by

$$Y_t = \left(\int_0^1 Y_t(z)^{\frac{\epsilon-1}{\epsilon}} dz \right)^{\frac{\epsilon}{\epsilon-1}} \tag{1.5}$$

where Y_t correspond to the quantity of the final good and $Y_t(z)$ to the quantity of intermediate good z at time t. ϵ denotes the elasticity of substitution across varieties.

Profit maximization under perfect competition yields the following set of demand schedules for intermediate goods and zero-profit condition

$$Y_t(z) = \left(\frac{P_t(z)}{P_t} \right)^{-\epsilon} Y_t \tag{1.6a}$$

$$P_t = \left(\int_0^1 P_t(z)^{1-\epsilon} dz \right)^{\frac{1}{1-\epsilon}} \tag{1.6b}$$

where $P_t(z)$ and P_t denote the price of intermediate good z and the price of the final good, respectively.

Intermediate goods firms. Each intermediate goods firm produces good z with a technology that

[32] As explained in Section 1.4.1, this assumption aims at making the model tractable.

is linear in labor

$$Y_t(z) = L_t(z) \tag{1.7}$$

Each firm hires both young and prime-age workers, aggregated into a labor input index $L_t(z)$ using CES technology.

$$L_t(z) = \left(\nu_y (X_{y,t} L_{y,t}(z))^{\frac{\eta-1}{\eta}} + \nu_p (X_{p,t} L_{p,t}(z))^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}} \tag{1.8a}$$

$$L_{y,t}(z) = \left(\frac{1}{\nu_y}\int_0^{\nu_y} L_{y,t}(z,k)^{\frac{\varepsilon_w-1}{\varepsilon_w}} dk\right)^{\frac{\varepsilon_w}{\varepsilon_w-1}} \tag{1.8b}$$

$$L_{p,t}(z) = \left(\frac{1}{\nu_p}\int_{\nu_y}^{\nu_y+\nu_p} L_{p,t}(z,k)^{\frac{\varepsilon_w-1}{\varepsilon_w}} dk\right)^{\frac{\varepsilon_w}{\varepsilon_w-1}} \tag{1.8c}$$

where η denotes the elasticity of substitution between young and prime-age labor inputs, and ε_w the elasticity of substitution between different varieties of labor.[33] $X_{y,t}$ and $X_{p,t}$ correspond to the skill level of young and prime-age workers, respectively. $L_{y,t}(z)$ is the quantity of young labor hired by the firm to produce good z, and $L_{y,t}(z,k)$ is the quantity of young labor of variety k. The same notation holds for prime-age workers. The skill level of young workers $X_{y,t}$, i.e. their productivity, evolves according to a learning-by-doing mechanism as proposed by Chang et al. (2002). Hours worked in a given period induce an increase in skills in the next period, with an elasticity μ, which persists over time at rate ϕ. On the other hand, prime-age agents are assumed to have reached their maximum level of efficiency which remains stable. Following d'Alessandro et al. (2019), firms take productivity levels as given.

$$X_{y,t} = X_{y,t-1}^{\phi} L_{y,t-1}^{\mu} \tag{1.9a}$$

$$X_{p,t} = X_{p,t-1} \tag{1.9b}$$

Each firm minimizes costs taking nominal wages $W_{y,t}(k)$ and $W_{p,t}(k)$ as given, which leads to the following set of demand schedules for young and prime-age labor inputs

$$L_{j,t}(z,k) = \left(\frac{W_{j,t}(k)}{W_{j,t}}\right)^{-\varepsilon_w} L_{j,t}(z) \qquad \text{for} \quad j \in \{y,p\} \tag{1.10}$$

[33]For $\eta = 0$, the two labor inputs are perfect complements. $\eta = 1$ corresponds to the Cobb-Douglas case. As $\eta \to \infty$, the two labor inputs become perfect substitutes.

where the wage indexes for young and prime-age workers are given by

$$W_{y,t} = \left(\frac{1}{\nu_y}\int_0^{\nu_y} W_{y,t}(k)^{1-\varepsilon_w} dk\right)^{\frac{1}{1-\varepsilon_w}} \tag{1.11a}$$

$$W_{p,t} = \left(\frac{1}{\nu_p}\int_{\nu_y}^{\nu_y+\nu_p} W_{p,t}(k)^{1-\varepsilon_w} dk\right)^{\frac{1}{1-\varepsilon_w}} \tag{1.11b}$$

Finally, taking the wage indexes $W_{y,t}$ and $W_{p,t}$ as given, each firm minimizes labor costs subject to Equation (1.8a). The optimality conditions with respect to $L_{y,t}(z)$ and $L_{p,t}(z)$ yield

$$MC_t = \left(\nu_y\left(\frac{W_{y,t}}{X_{y,t}}\right)^{1-\eta} + \nu_p\left(\frac{W_{p,t}}{X_{p,t}}\right)^{1-\eta}\right)^{\frac{1}{1-\eta}} \tag{1.12a}$$

$$\frac{L_{y,t}}{L_{p,t}} = \left(\frac{W_{y,t}}{W_{p,t}}\right)^{-\eta}\left(\frac{X_{y,t}}{X_{p,t}}\right)^{\eta-1} \tag{1.12b}$$

where MC_t is the nominal marginal cost.

Intermediate goods firms face staggered price setting à la Calvo. Each period, only a fraction $1-\theta_p$ of them are able to reset their prices. These firms maximize expected discounted real profits with respect to prices

$$\max_{P_t^*(z)} \; \mathbb{E}_t\left(\sum_{s=0}^{\infty}\theta_p^s Q_{t,t+s}\left[P_t^*(z) - MC_{t+s}\right]Y_{t+s}(z)\right) \tag{1.13}$$

subject to the final goods firm's demand constraint Equation (1.6a) for each variety z. $Q_{t,t+s}$ corresponds to the stochastic discount factor of prime-age agents between period t and $t+s$. This optimization problem implies that the optimal reset price and the dynamics of the aggregate price level are given by

$$P_t^*(z) = \frac{\epsilon}{\epsilon-1}\frac{\mathbb{E}_t\sum_{s=0}^{\infty}\theta_p^s Q_{t,t+s}P_{t+s}^{\epsilon-1}Y_{t+s}(z)MC_{t+s}}{\mathbb{E}_t\sum_{s=0}^{\infty}\theta_p^s Q_{t,t+s}P_{t+s}^{\epsilon-1}Y_{t+s}(z)} \tag{1.14a}$$

$$P_t = \left((1-\theta_p)P_t^*(z)^{1-\epsilon} + \theta_p P_{t-1}^{1-\epsilon}\right)^{\frac{1}{1-\epsilon}} \tag{1.14b}$$

Unions

To model wage stickiness, I follow the literature and assume that wages are set by unions. These unions act as monopolistic suppliers of differentiated labor services provided by workers. These labor services are bundled into a composite labor input by intermediate goods firms as specified in Section 1.4.1. To allow young and prime-age workers to get different wages, it is assumed that labor markets are segmented, so that there is one union for young workers and one for prime-age workers.

Each type of union chooses the nominal wage $W_{j,t}$ for an effective unit of labor so that $W_{j,t}(k) = W_{j,t}$ for all varieties k of workers, with $j \in \{y, p\}$, to maximize profits taking into account its members' utility and some wage adjustment costs, subject to the labor demand function for the workers it represents. The profits correspond to the difference between the wage income and the disutility of work, where $\lambda_{j,t}$ denotes the marginal utility of consumption. The wage adjustment cost is proportional to the total wage bill and is a quadratic function of the change in wages decided by the union, similar to Rotemberg (1982) for prices.[34] The adjustment cost parameter θ_w is the same for young and prime-age workers.

$$V_t^{w_j}(W_{j,t-1}(k)) = \max_{W_{j,t}(k)} \int \left(\frac{W_{j,t}(k)}{P_t} L_{j,t}(k) - \chi_j \frac{L_{j,t}(k)^{i^{1+\varphi_j}}}{1+\varphi_j} \frac{1}{\lambda_{j,t}} \right) dk$$

$$- \int \frac{\theta_w}{2} \left(\frac{W_{j,t}(k)}{W_{j,t-1}(k)} - 1 \right)^2 \frac{W_{j,t}}{P_t} L_{j,t} dk + \beta \mathbb{E} V_{t+1}^{w_j}(W_{j,t}(k))$$

subject to

$$L_{j,t}(k) = \left(\frac{W_{j,t}(k)}{W_{j,t}} \right)^{-\varepsilon_w} L_{j,t} \qquad j \in \{y, p\}$$

Solving this wage setting problem gives the following wage inflation equation for each type of worker, where $MRS_{j,t}$ is the marginal rate of substitution and $\Pi^w_{j,t} \equiv \frac{W_{j,t}}{W_{j,t-1}}$ is the wage inflation

[34] Another way to model wage stickiness would be to assume that the union is able to reset its wage rate at each period with probability $1 - \theta_w$, similar to Calvo (1983). In the case of a Calvo type of friction, the wage setting problem of the union would imply maximizing the present value of its members expected lifetime utility, which makes it difficult to adapt to this life-cycle setup. Consequently, the Rotemberg adjustment-cost version is preferred here.

rate.[35]

$$(1-\varepsilon_w)\frac{W_{j,t}}{P_t} = -\varepsilon_w MRS_{j,t} + \theta_w(\Pi^w_{j,t}-1)\Pi^w_{j,t}\frac{W_{j,t}}{P_t} - \beta\mathbb{E}_t\theta_w(\Pi^w_{j,t+1}-1)\Pi^w_{j,t+1}\frac{L_{j,t+1}}{L_{j,t}}\frac{W_{j,t+1}}{P_{t+1}} \quad j\in\{y,p\} \tag{1.16}$$

I follow Hagedorn et al. (2019) and assume that the wage adjustment process does not lead to actual costs, so as to avoid distortions due to large fluctuations in these costs after a government spending shock.[36]

Fiscal Policy

The government purchases consumption goods G_t and makes transfers to new young and new prime-age agents. These expenditures are financed by issuing debt, which consists of one-period non-contingent bonds $B_{G,t}$ yielding a nominal gross interest rate $R_{n,t}$, and by collecting lump-sum taxes T_t from young and prime-age households, as well as taxes on dividends from prime-age households.

$$B_{G,t} = R_{n,t-1}B_{G,t-1} + P_tG_t - P_tT_t + P_t(\nu_y\omega_x\tau^G_{NY,t} + \nu_p\omega_p\tau^G_{NP,t}) - \tau_d\nu_p\text{div}_t \tag{1.17}$$

The government follows a fiscal rule which dictates the response of debt and taxes to a change in government expenditures. The parameters Φ_B and Φ_G determine the response of deficits to debt and the extent of deficit financing, respectively.

$$P_tT_t = \Phi_B B_{G,t} + \Phi_G P_t G_t \tag{1.18}$$

Government expenditures evolve exogenously and follow a first order autoregressive process.

$$G_t = \rho_G G_{t-1} + \varepsilon_{G,t} \tag{1.19}$$

[35] See details in Section 1.F.

[36] As explained by Hagedorn et al. (2019), not making this assumption may lead to different results compared to using a price setting à la Calvo.

Total transfers to the new young and new prime-age agents are given by $\nu_y \omega_x \tau^G_{NY,t}$ and $\nu_p \omega_p \tau^G_{NP,t}$ respectively. These transfers are introduced to make the model tractable. In particular, these transfers are aimed at removing inequality in wealth between new young and pre-existing young, as well as between new prime-age and pre-existing prime-age. This ensures that all young agents solve the same optimization problem, and similarly for all prime-age agents. Thus the groups of young and prime-age individuals can be reduced to a representative young agent and a representative prime-age agent. Note that this assumption removes heterogeneity among young individuals and among prime-age individuals, but it preserves heterogeneity across the life cycle, and among retirees. The transfers are given by

$$\begin{aligned} P_t \tau^G_{NY,t} &= A_{y,t} \\ A_{y,t} + P_t \tau^G_{NP,t} &= A_{p,t} \end{aligned} \tag{1.20}$$

where $A_{y,t}$ and $A_{p,t}$ denote the average wealth among pre-existing young agents, and among pre-existing prime-age agents, respectively.

Monetary Policy

The nominal interest rate is set by the monetary authority and follows the Taylor rule

$$\frac{R_t}{R_{ss}} = \left(\frac{R_{t-1}}{R_{ss}} \right)^{\gamma} \left[\left(\frac{\Pi_t}{\Pi_{ss}} \right)^{\phi_\pi} \right]^{1-\gamma} \tag{1.21}$$

where R_{ss} stands for the steady-state gross nominal interest rate, Π_{ss} is the steady-state inflation. ϕ_π measure the reaction of monetary policy to current inflation. γ denotes the degree of interest rate smoothing.

Aggregation and Market Clearing

Aggregate assets of young individuals correspond to the sum of bequests left by deceased retirees and of the asset holdings of the fraction of young agents who do not become prime-age. The laws of motion of assets held by prime-age and retired agents can be defined similarly.

$$A_{y,t} = (1-\omega_p)\left((R_{n,t-1}+\zeta)B_{y,t-1}\right) + \omega_x\left(R_{n,t-1}B_{r,t-1}\right) \tag{1.22a}$$

$$A_{p,t} = (1-\omega_r)\left(R_{n,t-1}B_{p,t-1}\right) + \omega_p\left((R_{n,t-1}+\zeta)B_{y,t-1}\right) \tag{1.22b}$$

$$A_{r,t} = (1-\omega_x)\left(R_{n,t-1}B_{r,t-1}\right) + \omega_r\left(R_{n,t-1}B_{p,t-1}\right) \tag{1.22c}$$

Similarly, the aggregate levels of skills of young and prime-age workers respectively follow the laws of motion

$$X_{y,t} = (1-\omega_p)\left(X^{\phi}_{y,t-1}L^{\mu}_{y,t-1}\right) + \omega_x X_{y,0} \tag{1.23a}$$

$$X_{p,t} = (1-\omega_r)X_{p,t-1} + \omega_p\xi_p X_{y,t-1} \tag{1.23b}$$

where $X_{y,0}$ is the initial level of skills of newborns, and ξ_p aims at replicating the life-cycle productivity profile.

Finally, total consumption is given by the sum of each age group's consumption, weighted by their respective share in total population, and similarly for total taxes. Markets for bonds and goods clear.[37]

$$C_t = \nu_y C_{y,t} + \nu_p C_{p,t} + \nu_r C_{r,t} \tag{1.24a}$$

$$T_t = (\nu_y + \nu_p)\tau_t \tag{1.24b}$$

$$B_{G,t} = \nu_y B_{y,t} + \nu_p B_{p,t} + \nu_r B_{r,t} \tag{1.24c}$$

$$Y_t = C_t + G_t \tag{1.24d}$$

1.4.2 Calibration

In this section I discuss the parametrization of the model. One period corresponds to one quarter. Parameter values are summarized in Table 1.A.2.

Demographic structure. As in the empirical analysis, the young are defined as individuals aged between 15 and 29, and the prime-age between 30 and 64. This implies a probability of transition from young to prime-age $\omega_p = 0.0167$ and a probability of retirement $\omega_r = 0.0071$. The probability

[37]The goods market clearing condition is redundant by Walras' law.

of death for retirees is defined to match the average share of individuals aged 65 and above in total population over the sample, approximately 17%, which yields $\omega_x = 0.0243$.

Preferences. The disutility of labor of young and prime-age agents are fixed to match a fraction of hours worked in steady state of 0.4 for prime-age workers, and 0.35 for young workers. The value for young workers is obtained by multiplying steady hours of the prime-age by the relative employment rates of young and prime-age workers in the data.[38] The Frisch elasticities are set to $\varphi_y = 0.5$ and $\varphi_p = 0.5$, in line with conventional micro estimates.[39] The subjective discount factor is fixed to match an annualized interest rate at steady state of 2%, which leads to $\beta = 0.97$. The risk premium faced by young agents is calibrated to match the consumption ratio of prime-age relative to young individuals in the data, which is approximately equal to 1.4.

Production. The elasticity of substitution across varieties and the price stickiness parameter are calibrated to standard values used in the New Keynesian literature. Specifically, the elasticity of substitution across varieties ϵ is set to 10, which implies a price markup of 10%. The price stickiness parameter is set to $\theta_p = 0.75$, which implies that firms can reset their prices once every 4 quarters. Following Erceg et al. (2000), the elasticity of substitution across labor types is fixed to $\epsilon_w = 4$, which implies a wage markup of 1/3. To calibrate the adjustment cost on wages, I set the slope of the wage Phillips curves to 0.0066, which is the benchmark value used in Schmitt-Grohé and Uribe (2006). This implies $\theta_w \approx 500$. As regards the elasticity of substitution between age groups, I choose $\eta = 5$, which is in line with estimates reported in micro empirical studies such as Welch (1979), Card and Lemieux (2001) or Ottaviano and Peri (2012).[40] The profits are fully taxed.

Learning-by-doing. The parameters of the LBD are obtained from Chang et al. (2002) who estimate them using PSID data and find $\phi = 0.797$ and $\mu = 0.111$. The parameter ξ_p in the aggregate law of motion of skills for prime-age workers is calibrated to match the wage ratio of prime-age relative to young workers in the data, which is approximately equal to 1.4. The initial level of skills of newborn young $X_{y,0}$ is normalized to 0.5.

[38] The average employment rate of young and prime-age individuals is approximately 65% and 74% respectively over the sample period 1981-2007.

[39] See in particular Chetty et al. (2011) for a meta-analysis of existing micro and macro evidence on labor supply elasticities.

[40] Welch (1979) finds a value between 4.6 and 12.5 for the elasticity of substitution across experience groups. Card and Lemieux (2001) estimate the value of the elasticity of substitution between different age groups in the range of 4 to 6 for both low and high education workers. Estimates reported in Ottaviano and Peri (2012) imply values in the range of 3.2 to 7.7.

Fiscal and monetary policy. The government spending to output ratio is set to 0.2, consistent with the sample average. The parameters of the fiscal rule are set to $\Phi_G = 0.1$ and $\Phi_B = 0.33$, following Galí et al. (2007). The persistence of the fiscal shock is set to 0.8, as in Christiano et al. (2011). The parameters of the Taylor rule are set to $\rho = 0.85$ and $\phi_\pi = 2.4$, in line with Christiano et al. (2014).

1.4.3 Results

Figure 1.6 displays the impulse responses of key aggregate and disaggregate variables for young and prime-age workers to an expansionary government spending shock, both in the standard model without LBD (dashed black lines) and in the model with LBD (blue lines). Responses are measured in quarterly percent deviations from steady state values, except for the responses of inflation and interest rates which are measured as annualized percentage-point deviations from steady state.

Let's first consider the responses of aggregate variables (panel (a)). A positive government expenditure shock leads to a negative wealth effect as the present value of taxes paid by households increases to finance the fiscal expansion. This translates into a reduction of consumption and leisure which are normal goods. As prices are sticky, firms increase their production to meet higher demand since some of them cannot adjust their prices, hence an outward shift of the labor demand curve which raises real wages. In the standard model without LBD, productivity of workers is unresponsive to the shock, so the increase in real wages generates a surge of marginal costs. This in turn leads to higher inflation since it depends on current and expected future marginal costs.[41] The Central Bank responds by raising the nominal interest rate, which translates into higher real interest rates by the Taylor principle, encouraging households to postpone consumption. This intertemporal substitution effect thus amplifies the drop in consumption. In contrast, in the model with LBD, young workers raise their skill level as they work more, boosting the future productivity of the firms, which in turn dampens the increase in marginal costs. If wages are sticky enough, marginal costs can actually fall. This leads to a decrease in expected inflation, and through the monetary policy rule, to a decline in real interest rates which boosts consumption by intertemporal substitution. The increase in aggregate consumption and productivity, along with a reduction in inflation and the nominal interest rate, is in line with my estimates based on aggregate data, displayed in Section 1.B, and

[41] This is a key implication of the New Keynesian Phillips Curve. Iterating forward this equation yields that current inflation is determined by the discounted sum of expected future real marginal costs.

with evidence reported notably in Jørgensen and Ravn (2018) and d'Alessandro et al. (2019).[42]

Turning to disaggregate variables (panel (b)), we can observe that the responses of aggregate variables mask substantial heterogeneity between young and prime-age groups in the model with LBD. In particular, the surge of aggregate consumption appears to be mainly driven by the increase in young individuals' consumption while prime-age agents tend to reduce it, in line with empirical results. In contrast, the standard model without LBD predicts a decrease in consumption for both age groups. In the model with LBD, the fiscal expansion operates not only through intertemporal substitution effects but also through redistribution effects which shift resources from prime-age individuals (and retirees) to young households. First, the decline in real interest rates favors borrowers but penalizes savers. As a consequence, young agents increase borrowing, while prime-age agents, who rely heavily on interest income, incur losses on their financial assets. Second, heterogeneity in skill accumulation further induces redistribution effects through the impact on wages. Specifically, although prime-age workers are more productive, the relative productivity of young workers increases, through the LBD mechanism, as they supply more hours. However, their productivity growth in only partially reflected in the growth of their labor income due to wage stickiness, which implies that young workers become relatively more profitable for firms compared to prime-age workers. As a result, labor demand increases relatively more for young workers, which translates in stronger wage increases for those who primarily finance their consumption from labor income, compared to prime-age workers. Remarkably, the model with LBD generates hump-shaped responses for consumption, hours, wage and productivity for the young, in line with empirical evidence.

To sum up, young workers benefit from redistribution effects due to changes in interest rates and wages, as well as from intertemporal substitution effects, hence the surge of their consumption. However, for prime-age individuals, the negative redistribution effects generated by the fiscal expansion tend to offset the positive effect from intertemporal substitution, leading to a small reduction of their consumption.

[42] In addition, several other papers document a fall in inflation after a positive government spending shock, such as Fatás and Mihov (2001), Mountford and Uhlig (2009), Hall and Thapar (2018).

Figure 1.6: Impulse responses of selected variables to an expansionary government spending shock in the life-cycle model with and without LBD

(a) IRFs of aggregate variables

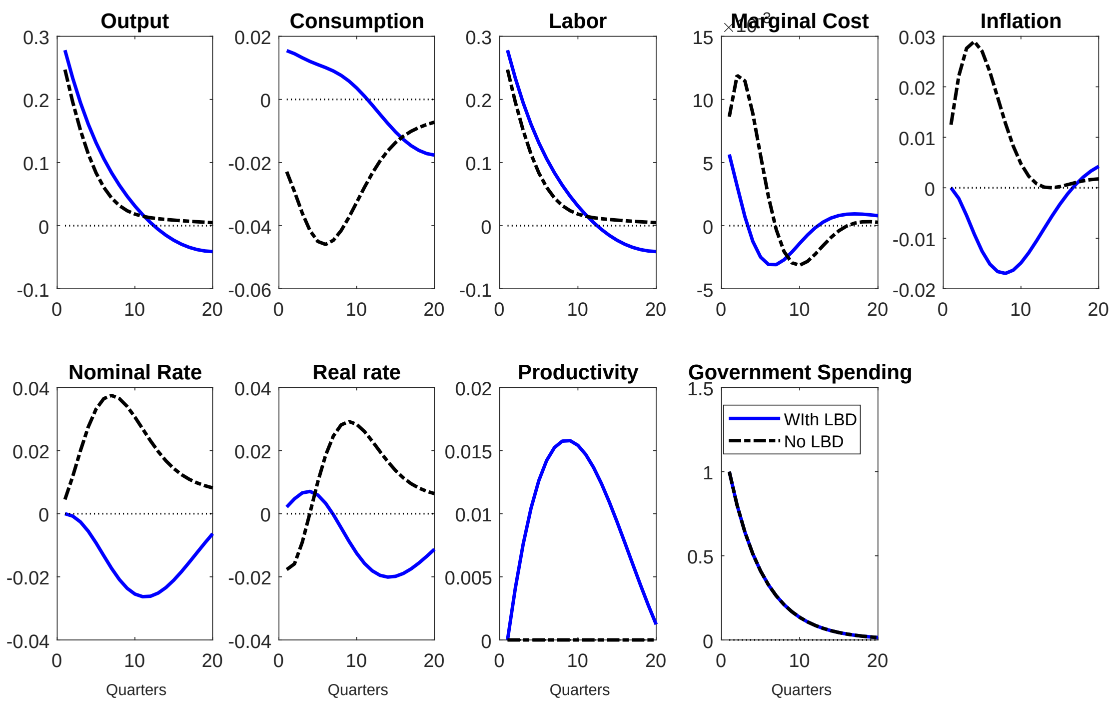

(b) IRFs of disaggregate variables

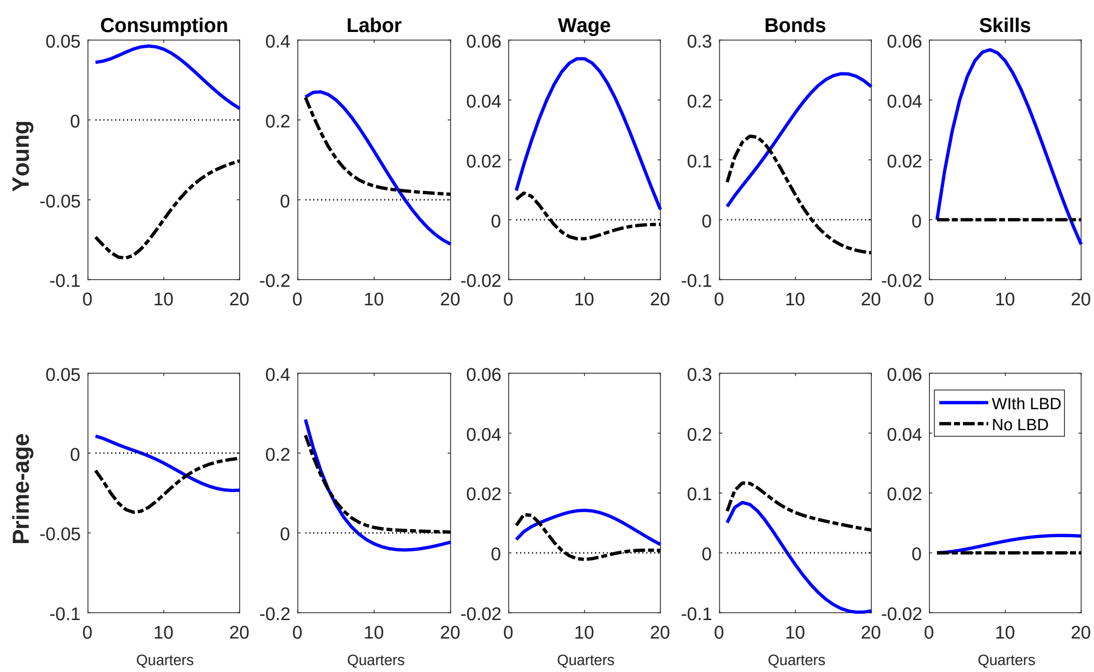

Notes: this figure shows the impulse responses of aggregate variables (Panel (a)) and disaggregate variables for young and prime-age workers (Panel (b)) to a 1% shock to government expenditure in the life-cycle model with LBD (blue solid lines) and without LBD (black dashed lines).

1.4.4 Sensitivity Analysis

In this section I provide a sensitivity analysis of my findings with respect to some key parameters in the life-cycle model with LBD. All related figures are displayed in Section 1.E.

First, it is worth noting that price and wage rigidities are critical in generating an increase in consumption for young agents. Figure 1.E.1 shows the impulse responses to a government spending shock in the baseline case with nominal frictions, compared to the case with flexible price and wages. Without price rigidities, firms adjust prices and not quantities in response to a rise in government demand. This leads to a reduced increase in hours worked, and thus in productivity and wage growth for young workers.[43] Under flexible wages, the rise in productivity of young workers is fully reflected in the increase of their real wage, which implies a surge in marginal costs. In turn, inflation and real rates strongly increase, leading to a crowding-out of consumption. However, as showed in Figure 1.E.2 and Figure 1.E.3, the response of consumption is still positive for the young when the degree of price rigidity falls to 0.4, or when the wage adjustment cost falls to 300.

Figure 1.E.4 considers variations in the value of the Taylor rule inflation parameter ϕ_π. The stronger the response of the Central Bank to inflation, the larger the drop in the real interest rate. This reinforces the positive intertemporal substitution effect and, to a lesser extent, the redistribution effect from savers to borrowers. Therefore, as the Taylor rule inflation parameter increases, the response of consumption is larger for the young, and becomes less negative for the prime-age. In particular, the prime-age response becomes slightly positive on impact for values of ϕ_π greater than 1.5, although their response at 12 quarter horizon remains negative.

Figure 1.E.5 displays the responses of consumption for young and prime-age agents at different horizons for different values of the persistence of the government spending shock. The young increase their consumption for values of the persistence ρ_G up to 0.9. For higher values, the negative wealth effect from higher taxes more than offsets the positive redistribution effect from higher wages and lower real interest rates. As a result, young agents also reduce their consumption for high values of the persistence.

Figure 1.E.6 reports the sensitivity of consumption responses for young and prime-age agents to the fiscal rule parameters at different horizons. The consumption responses among young agents

[43] Marginal costs and inflation are unchanged in the case of flexible prices.

are always positive and display little sensitivity to the fiscal rule parameters. This indicates that, everything else equal, the negative wealth effect from higher taxes is always lower than the positive redistribution effect from higher wages and lower real interest rates. Young people anticipate that they will have to pay higher taxes, either while they are still young or when they will be prime-age workers, so their consumption decision is little affected by the timing of taxes or the extent of deficit financing of government spending. The consumption response of prime-age agents is broadly negative, except for lower values of Φ_B and Φ_G, corresponding to a higher degree of deficit financing and a lower response to debt. In this case, debt is allowed to accumulate for a long time and taxes are raised very gradually. Therefore the negative wealth effect is reduced for prime-age agents as they expect the burden of taxes to fall on the next generation of workers.

Figure 1.E.7 reports the sensitivity of consumption responses for young and prime-age agents to the learning-by-doing parameters at different horizons. The figure shows that the consumption response of the young is positive for a reasonably large range of values, while it is broadly negative for the prime-age. Note that in this life-cycle model the LBD mechanism does not lead to indeterminacy issues as in a representative agent model. Indeed, d'Alessandro et al. (2019) show that, if the LBD process is too strong, an increase in the nominal interest rate can lead to a rise in inflation due to the fall in productivity which pushes up marginal costs. As a result, inflation expectations become self-fulfilling. However, in this life-cycle model, this mechanism is not powerful enough to generate instability as only a fraction of the population is learning.

In the baseline analysis, it is assumed that young and prime-age workers have the same labor supply elasticity, in order to emphasize the importance of differences in skill accumulation over the life cycle in shaping heterogeneous outcomes. Potential age differences in labor supply elasticity have received some attention in the litterature, in particular to partially account for the greater volatility of young hours relative to the prime-age observed in the data.[44] Figure 1.E.8 thus considers variations in values of labor supply elasticity for young and prime-age workers. The figure shows that baseline results are preserved and reinforced when assuming higher labor supply elasticity for young individuals.

[44]See, e.g., Ríos-Rull (1996), Jaimovich et al. (2013), Janiak and Monteiro (2016). However, direct evidence that young workers have higher labor supply elasticity is scarce.

1.4.5 Discussion: Age vs. Financial Constraints

A key mechanism to generate asymmetric consumption responses of government spending shocks across agents is the introduction of "hand-to-mouth" agents, who fully consume their current disposable income. Galí et al. (2007) show that these households, who do not borrow or save due to a lack of access to financial markets or borrowing constraints, increase their consumption after a government spending shock in presence of price stickiness and under the assumption of imperfectly competitive labor market. Indeed, the surge in labor demand puts upward pressure on real wages, which stimulates consumption of these rule-of-thumb agents. Furlanetto (2011) shows that this result is preserved and reinforced when assuming that wages are sticky. This theory could also justify why young agents, more likely to be financially constrained, tend to increase their consumption after a government spending shock.

I compare the predictions of the life-cycle model with LBD described in Section 1.4.1 to a life-cycle model where a share of young agents behave as hand-to-mouth. Figure 1.7 displays the impulse response functions of selected aggregate and disaggregate variables to a positive government spending shock in both models. The share of young hand-to-mouth is set equal to 0.5 in the model with financial constraints. At the aggregate level, both models predict an increase in labor, output and consumption. However, the model with financial constraints predicts an increase in marginal costs due to the surge in real wages which is not counteracted by a rise in productivity. This drives up expected inflation, to which the monetary authority reacts by raising interest rates. In contrast, in the model with LBD, the shift in government spending triggers an increase in TFP that puts downward pressure on marginal costs and inflation, leading instead to a reduction in the nominal interest rate, which is consistent with the existing evidence as discussed in Section 1.4.3.

Turning to disaggregated variables, consumption strongly increases for young individuals in both models, while it remains flat or decreases for the prime-age. The increase in consumption for young individuals is more pronounced on impact in the model with financial constraints, but is less persistent as the response becomes negative after a couple of quarters. In addition, both models predict a similar increase in hours worked for young and prime-age workers, albeit more persistent for young individuals in the model with LBD. As pointed out by Furlanetto (2011), wage stickiness strongly reduces the heterogeneity in the adjustment of hours worked between Ricardian and hand-to-mouth agents. However, the models differ regarding the effects of a positive government spending shock

Figure 1.7: Impulse responses of selected variables to an expansionary government spending shock in the life-cycle model with LBD vs. with financial constraints

(a) IRFs of aggregate variables

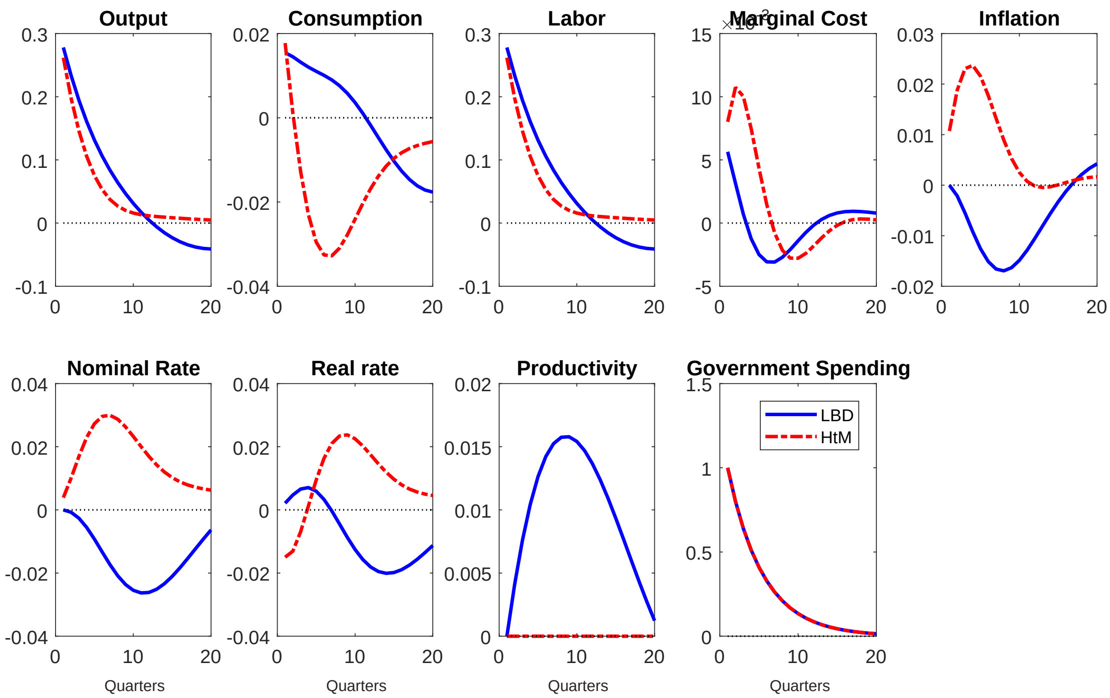

(b) IRFs of disaggregate variables

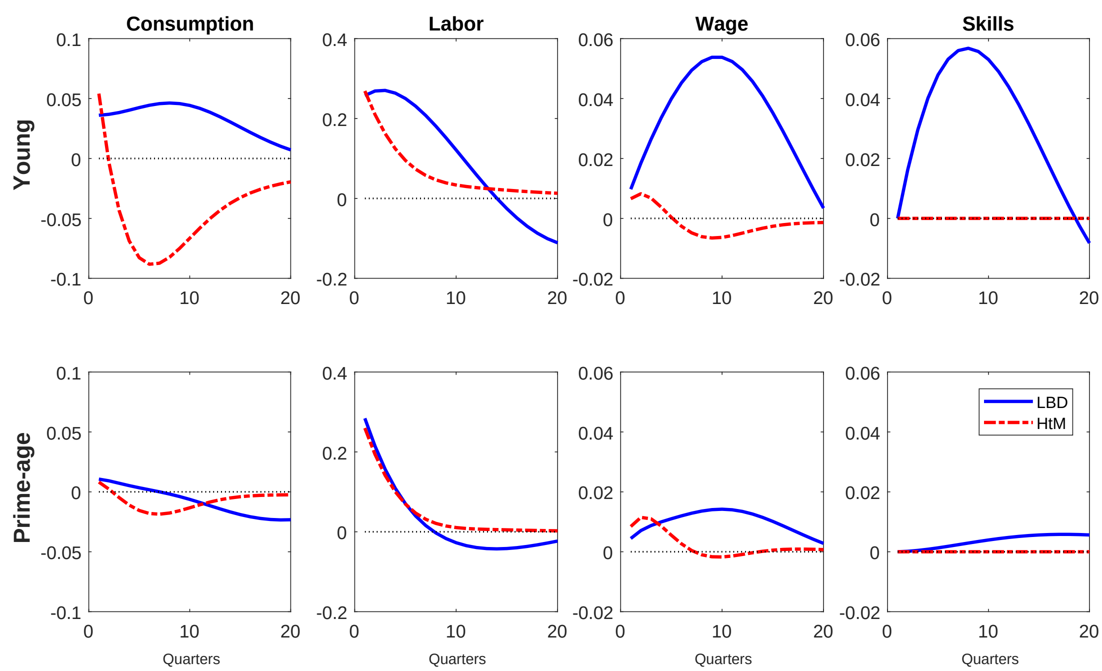

Notes: this figure shows the impulse responses of aggregate variables (Panel (a)) and disaggregate variables for young and prime-age workers (Panel (b)) to a 1% shock to government expenditure in the life-cycle model with LBD (blue solid lines) and with a share of Hand-to-Mouth young agents (red dashed lines).

on wages. The model with LBD predicts that wage growth for young agents is stronger than for the prime-age, reflecting the increase in skills, while the response of wages is nearly flat for both young and prime-age individuals in the model with financial constraints. As showed in Section 1.3.4, evidence suggests that the growth in real wages is more pronounced for young workers, as predicted by the model with LBD.

This analysis suggests that the model with LBD manages to replicate qualitatively the age-specific effects of a fiscal expansion on labor market variables observed in the data, which a model with hand-to-mouth young does not capture. However, heterogeneity in skill accumulation and heterogeneity in marginal propensity to consume over the life cycle likely interact. Quantifying the contribution of each of these channels in shaping these differential outcomes across age groups would likely yield interesting insights.

1.5 Conclusion

This paper provides a new perspective on the transmission of government spending shocks by uncovering a key interaction between fiscal policy, demographics and productivity. First, I present new evidence that age is a strong predictor of consumption adjustment to shifts in government expenditures. Young households increase their consumption after an unexpected increase in government spending, while prime-age households tend to reduce it. My analysis suggests that this result is not primarily driven by financial constraints.

The second contribution of this paper is to propose and study an alternative channel of fiscal policy transmission able to generate heterogeneous responses across age groups. Government spending enhances human capital accumulation, and thus affects more young agents who have a steep age-productivity profile. To illustrate the mechanism, I build a life-cycle New Keynesian model with learning-by-doing for young individuals. As a government spending shock stimulates hours, young workers accumulate skills faster than their prime-age counterparts. Relative labor demand for the young increases, boosting their wages. The rise in the productivity of the firm leads to a reduction in marginal costs and inflation, and thus in the real interest rate through the monetary policy rule. As a result, the fiscal stimulus tends to crowd in consumption via intertemporal substitution, but also generates redistribution effects which benefit young people. Specifically, consumption of the young is stimulated by lower real interest rates which encourage borrowing, and by higher labor income.

Finally, I provide both micro- and macro-level evidence that corroborates this mechanism.

Given the accelerating demographic transition towards an older population in the U.S. and other developed countries, results in this paper indicate that fiscal stimulus measures could become increasingly less efficient in boosting the economy. On the other hand, policies which promote human capital formation may increase the effectiveness of fiscal policy, in particular if they are targeted at young individuals.

Bibliography

Agarwal, S. and W. Qian (2014). Consumption and debt response to unanticipated income shocks: Evidence from a natural experiment in singapore. *American Economic Review 104*(12), 4205–30.

Aghion, P., D. Hemous, and E. Kharroubi (2014). Cyclical fiscal policy, credit constraints, and industry growth. *Journal of Monetary Economics 62*, 41–58.

Aguiar, M. and M. Bils (2015). Has consumption inequality mirrored income inequality? *American Economic Review 105*(9), 2725–56.

Aguiar, M. and E. Hurst (2013). Deconstructing life cycle expenditure. *Journal of Political Economy 121*(3), 437–492.

Anderson, E., A. Inoue, and B. Rossi (2016). Heterogeneous consumers and fiscal policy shocks. *Journal of Money, Credit and Banking 48*(8), 1877–1888.

Attanasio, O. P., J. Banks, and S. Tanner (2002). Asset holding and consumption volatility. *Journal of political Economy 110*(4), 771–792.

Bachmann, R. and E. R. Sims (2012). Confidence and the transmission of government spending shocks. *Journal of Monetary Economics 59*(3), 235–249.

Barro, R. J. and C. J. Redlick (2011). Macroeconomic effects from government purchases and taxes. *The Quarterly Journal of Economics 126*(1), 51–102.

Basso, H. and O. Rachedi (2020). The young, the old and the government: Demographics and fiscal multipliers. *Mimeo*.

Baugh, B., I. Ben-David, H. Park, and J. A. Parker (2018). Asymmetric consumption response of households to positive and negative anticipated cash flows. *NBER Working Paper No. 25086*.

Ben-Porath, Y. (1967). The production of human capital and the life cycle of earnings. *Journal of political economy 75*(4, Part 1), 352–365.

Ben Zeev, N. and E. Pappa (2015). Multipliers of unexpected increases in defense spending: An empirical investigation. *Journal of Economic Dynamics and Control 57*, 205–226.

Blanchard, O. and R. Perotti (2002). An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *The Quarterly Journal of Economics 117*(4), 1329–1368.

Bowlus, A. J. and C. Robinson (2012). Human capital prices, productivity, and growth. *American Economic Review 102*(7), 3483–3515.

Brinca, P., H. A. Holter, P. Krusell, and L. Malafry (2016). Fiscal multipliers in the 21st century. *Journal of Monetary Economics 77*, 53–69.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of monetary Economics 12*(3), 383–398.

Card, D. and T. Lemieux (2001). Can falling supply explain the rising return to college for younger men? a cohort-based analysis. *The Quarterly Journal of Economics 116*(2), 705–746.

Chang, Y., J. F. Gomes, and F. Schorfheide (2002). Learning-by-doing as a propagation mechanism. *American Economic Review 92*(5), 1498–1520.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review 101*(3), 471–75.

Christiano, L., M. Eichenbaum, and S. Rebelo (2011). When is the government spending multiplier large? *Journal of Political Economy 119*(1), 78–121.

Christiano, L. J., R. Motto, and M. Rostagno (2014). Risk shocks. *American Economic Review 104*(1), 27–65.

Cloyne, J. S. and P. Surico (2016). Household debt and the dynamic effects of income tax changes. *The Review of Economic Studies 84*(1), 45–81.

Coibion, O. and Y. Gorodnichenko (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy 120*(1), 116–159.

d'Alessandro, A., G. Fella, and L. Melosi (2019). Fiscal stimulus with learning-by-doing. *International Economic Review 60*(3).

De Giorgi, G. and L. Gambetti (2012). The effects of government spending on the distribution of consumption. *Mimeo*.

Deaton, A. (1985). Panel data from time series of cross-sections. *Journal of econometrics 30*(1-2), 109–126.

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics 46*(2), 281–313.

Evans, C. L. (1992). Productivity shocks and real business cycles. *Journal of Monetary Economics 29*(2), 191–208.

Fatás, A. and I. Mihov (2001). The effects of fiscal policy on consumption and employment: theory and evidence. *Mimeo*.

Ferriere, A. and G. Navarro (2018). The heterogeneous effects of government spending: It's all about taxes. *Mimeo*.

Fiori, G. and D. Ferraro (2020). Aging of the baby boomers: Demographics and propagation of tax shocks. *American Economic Journal: Macroeconomics 12*(2), 167–93.

Fukao, K., R. Kambayashi, D. Kawaguchi, H. U. Kwon, Y. G. Kim, and I. Yokoyama (2006). Deferred compensation: evidence from employer-employee matched data from japan. *Hitotsubashi Institute of Economic Research*.

Furlanetto, F. (2011). Fiscal stimulus and the role of wage rigidity. *Journal of Economic Dynamics and Control 35*(4), 512–527.

Galí, J., J. D. López-Salido, and J. Vallés (2007). Understanding the effects of government spending on consumption. *Journal of the European Economic Association 5*(1), 227–270.

Gertler, M. (1999). Government debt and social security in a life-cycle economy. *Carnegie ĲRochester Conference Series on Public Policy 50*, 61–110.

Hagedorn, M., I. Manovskii, and K. Mitman (2019). The fiscal multiplier. *NBER Working Paper No. 25571*.

Hall, M. and A. Thapar (2018). The economic effects of government spending: The importance of controlling for anticipated information. *Mimeo*.

Heathcote, J. (2005). Fiscal policy with heterogeneous agents and incomplete markets. *The Review of Economic Studies 72*(1), 161–188.

Hellerstein, J. K. and D. Neumark (1995). Are earnings profiles steeper than productivity profiles? evidence from israeli firm-level data. *Journal of human resources*, 89–112.

Hellerstein, J. K. and D. Neumark (2007). Production function and wage equation estimation with heterogeneous labor: Evidence from a new matched employer-employee data set. In *Hard-to-measure goods and services: Essays in honor of Zvi Griliches*, pp. 31–71. University of Chicago Press.

Hellerstein, J. K., D. Neumark, and K. R. Troske (1999). Wages, productivity, and worker characteristics: Evidence from plant-level production functions and wage equations. *Journal of labor economics 17*(3), 409–446.

Ilzetzki, E., E. G. Mendoza, and C. A. Végh (2013). How big (small?) are fiscal multipliers? *Journal of monetary economics 60*(2), 239–254.

Jaimovich, N., S. Pruitt, and H. E. Siu (2013). The demand for youth: explaining age differences in the volatility of hours. *American Economic Review 103*(7), 3022–44.

Janiak, A. and P. S. Monteiro (2016). Towards a quantitative theory of automatic stabilizers: the role of demographics. *Journal of Monetary Economics 78*, 35–49.

Jappelli, T. and L. Pistaferri (2014). Fiscal policy and mpc heterogeneity. *American Economic Journal: Macroeconomics 6*(4), 107–36.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review 96*(5), 1589–1610.

Jørgensen, P. L. and S. H. Ravn (2018). The inflation response to government spending shocks: A fiscal price puzzle? *Københavns Universitet*.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to HANK. *American Economic Review 108*(3), 697–743.

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Krueger, D. and F. Perri (2006). Does income inequality lead to consumption inequality? evidence and theory. *The Review of Economic Studies 73*(1), 163–193.

Lusardi, A. (1996). Permanent income, current income, and consumption: Evidence from two panel data sets. *Journal of Business & Economic Statistics 14*(1), 81–90.

Mertens, K. and J. L. Montiel Olea (2018). Marginal tax rates and income: New time series evidence. *The Quarterly Journal of Economics 133*(4), 1803–1884.

Misra, K. and P. Surico (2014). Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics 6*(4), 84–106.

Mountford, A. and H. Uhlig (2009). What are the effects of fiscal policy shocks? *Journal of applied econometrics 24*(6), 960–992.

Ottaviano, G. I. and G. Peri (2012). Rethinking the effect of immigration on wages. *Journal of the European economic association 10*(1), 152–197.

Ramey, V. A. (2011). Identifying government spending shocks: it's all in the timing. *The Quarterly Journal of Economics 126*(1), 1–50.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. In *Handbook of Macroeconomics*, Volume 2, pp. 71–162. Elsevier.

Ríos-Rull, J.-V. (1996). Life-cycle economies and aggregate fluctuations. *The Review of Economic Studies 63*(3), 465–489.

Rotemberg, J. J. (1982). Sticky prices in the united states. *Journal of Political Economy 90*(6), 1187–1211.

Schmitt-Grohé, S. and M. Uribe (2004). Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control 28*(4), 755–775.

Schmitt-Grohé, S. and M. Uribe (2006). Comparing two variants of calvo-type wage stickiness. *NBER Working Paper No. 12740*.

Welch, F. (1979). Effects of cohort size on earnings: The baby boom babies' financial bust. *Journal of Political Economy 87*(5, Part 2), S65–S97.

Zidar, O. (2019). Tax cuts for whom? heterogeneous effects of income tax changes on growth and employment. *Journal of Political Economy 127*(3), 1437–1472.

1.A Tables

Table 1.A.1: Average Cell Size by Groups - CEX dataset

Age	***15-29***	***30-64***	
	764.0	2933.5	
Housing tenure	***Mortgagors***	***Outright Owners***	***Tenants***
	1752.9	774.6	1510.4
Income level	***Low***	***High***	
	1387.5	2773.6	
Education	***Low***	***High***	
	1939.6	2221.5	
Financial market participation	***Limited***	***Not limited***	
	2863.8	1235.8	
Age and housing tenure	***15-29***	***30-64***	
Mortgagors	151.4	1516.5	
No mortgage	588.6	1331.1	
Age and income	***15-29***	***30-64***	
Low	387.1	714.6	
High	376.9	2218.8	
Age and education	***15-29***	***30-64***	
Low	304.5	1260.1	
High	459.5	1673.4	
Age and financial market participation	***15-29***	***30-64***	
Limited	578.5	2002.6	
Not limited	175.8	887.7	

Notes: This table reports the average cell size for each group of consumers in the CEX dataset, where the cell size refers to the number of households used to make one quarterly observation.

Table 1.A.2: Calibration values

Parameter	Value	Description	Target/Source
β	0.97	Discount factor	Annualized interest rate 2%
ω_p	0.0167	Probability of becoming prime-age	Young for 15 years
ω_r	0.0071	Probability of retirement	Prime-age for 35 years
ω_x	0.0243	Probability of death	Share of 65+ in population
ϵ	10	Elasticity of substitution across varieties	Price mark-up of 10%
θ_p	0.75	Probability of fixed price	Average duration 4 quarters
ζ	0.035	Risk premium for young	Consumption ratio prime-age/young
ϕ	0.111	LBD: coefficient of hours impact	Chang et al. (2002)
μ	0.797	LBD: coefficient of auto-correlation	Chang et al. (2002)
ξ_p	0.6	Efficiency parameter of prime-age	Wage ratio prime-age/young
$X_{y,0}$	0.5	Initial level of skills	Normalization
φ_y	0.5	Frisch elasticity of labor supply for young	Chetty et al. (2011)
φ_p	0.5	Frisch elasticity of labor supply for prime-age	Chetty et al. (2011)
χ_y		Disutility of labor for young	Fraction of hours worked 0.35
χ_p		Disutility of labor for prime-age	Fraction of hours worked 0.4
θ_w	500	Adjustment cost of wages parameter	Slope Phillips curve 0.006
ε_w	4	Elasticity of substitution across labor types	Erceg et al. (2000)
η	5	Elasticity of substitution between young and prime-age	Ottaviano and Peri (2012)
g_Y	0.2	Government spending to output ratio	Sample average
Φ_G	0.1	Degree of deficit financing	Galí et al. (2007)
Φ_B	0.33	Response of deficits to debt	Galí et al. (2007)
ρ_G	0.8	Persistence of government spending shock	Christiano et al. (2011)
ρ	0.85	Taylor rule: interest smoothing parameter	Christiano et al. (2014)
Φ_π	2.4	Taylor rule: coefficient of inflation	Christiano et al. (2014)

1.B Figures

Figure 1.B.1: Government spending shocks

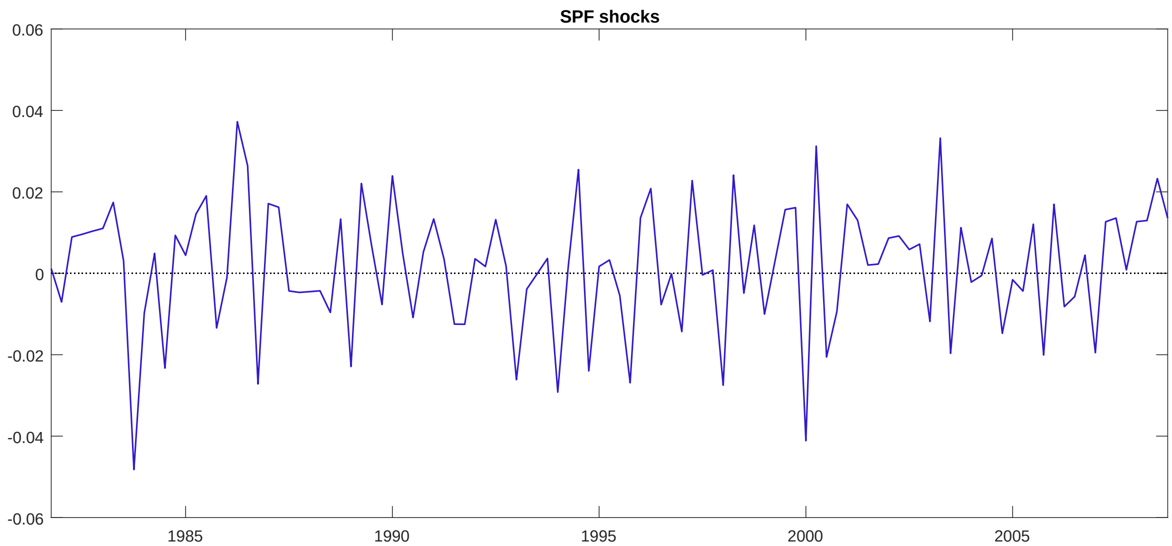

Notes: This graph displays the identified government spending shocks based on SPF forecasts.

Figure 1.B.2: Impulse responses to government spending shocks by age groups

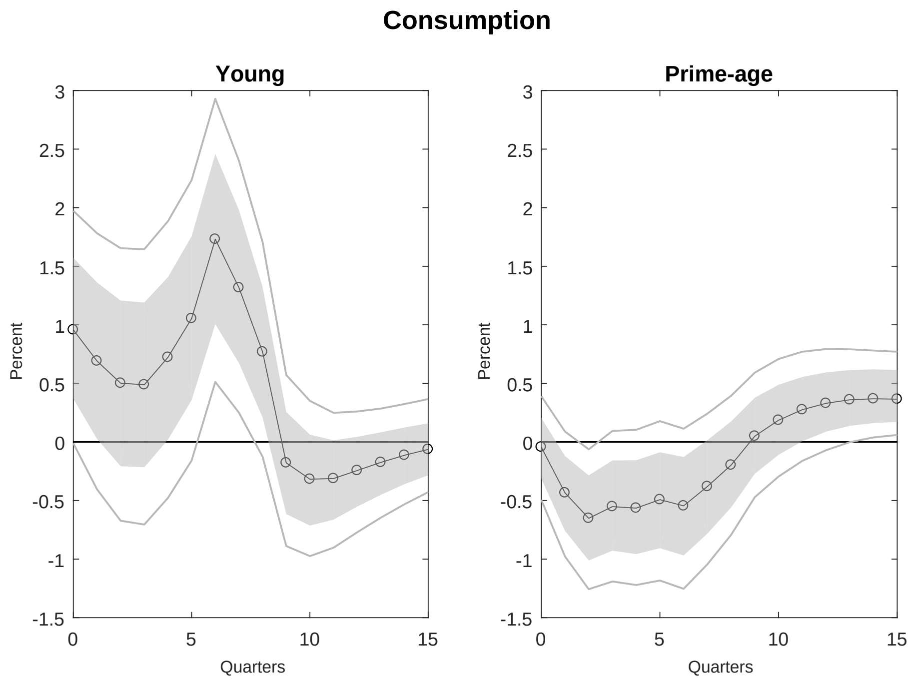

Notes: These graphs show the impulse responses of nondurable consumption for young (under 30) and prime-age (30-64) households to an exogenous government spending shock leading to an initial 1% increase in government expenditures. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Figure 1.B.3: Impulse responses to government spending shocks by income group

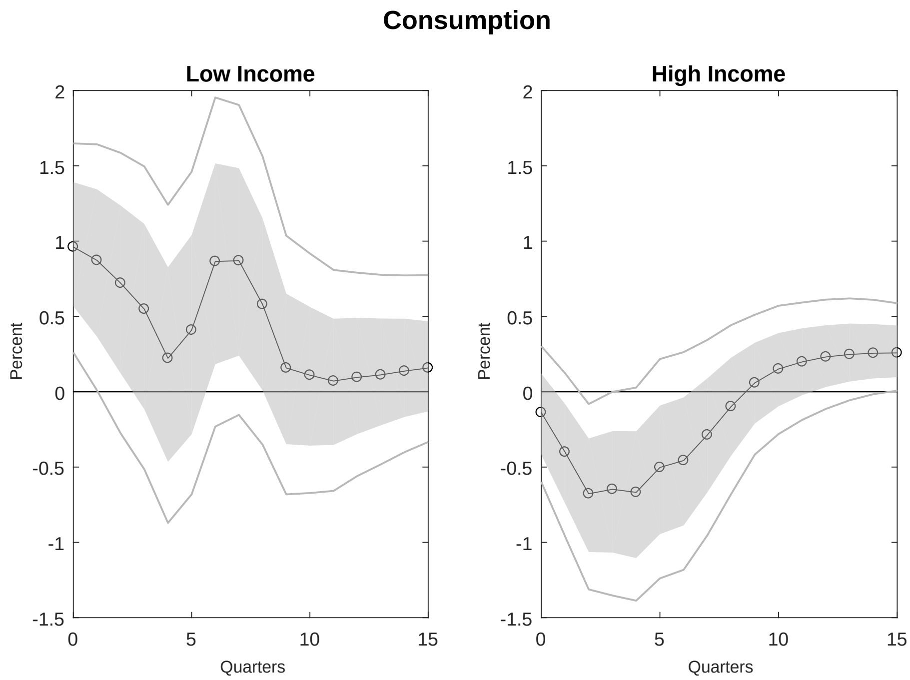

Notes: These graphs show the impulse responses of nondurable consumption for each income tertile to an exogenous government spending shock leading to an initial 1% increase in government expenditures. “Low income" denotes the group of households with after-tax income below the 35th percentile, and “High income" the group with after-tax income above the 35th percentile. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Figure 1.B.4: Impulse responses to government spending shocks by housing tenure

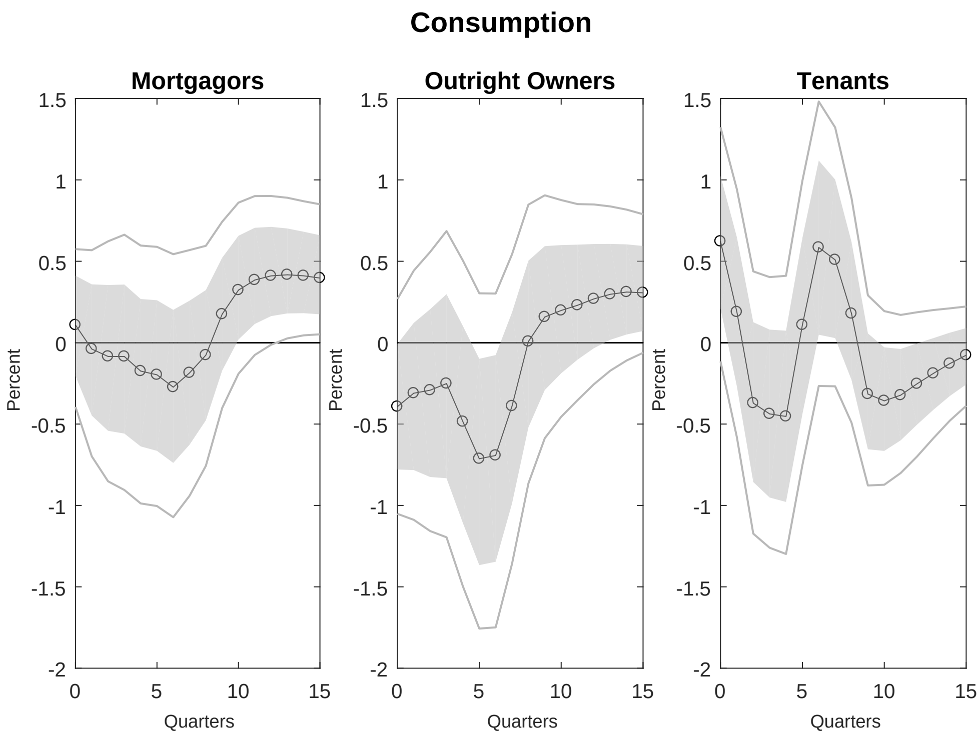

Notes: These graphs show the impulse responses of nondurable consumption for various housing tenure groups to an exogenous government spending shock leading to an initial 1% increase in government expenditures. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Figure 1.B.5: Age composition of housing tenure groups and income tertiles

(a) By housing tenure

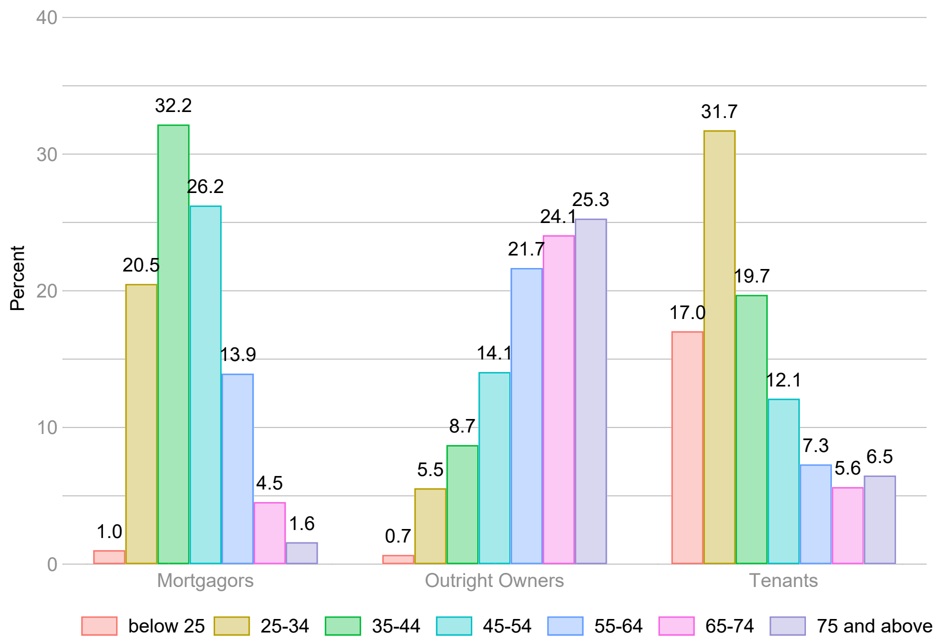

(b) By income tertile

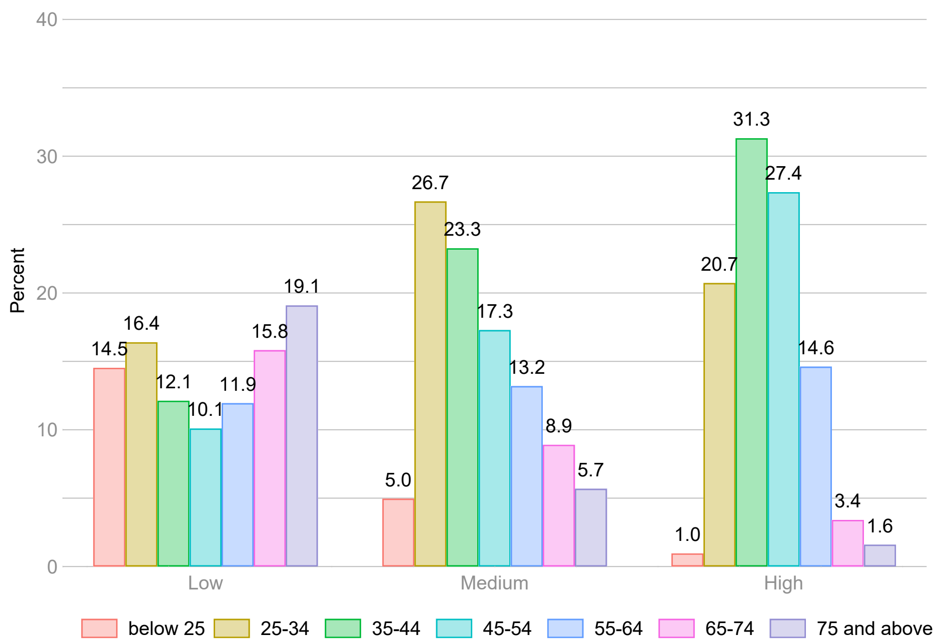

Notes: These graphs show the proportion of each age group by housing tenure (upper graph) and by income tertile (lower graph).

1.A Data

1.A.1 Aggregate U.S. Data

Time series for gross domestic product, non-durable consumption, wages, GDP price deflator, the three-month Treasury-Bill rate are available from the website of the Federal Reserve Board of St. Louis (FRED). Federal government expenditures include direct consumption and investment purchases, which excludes the imputed rent on government capital stocks. This data is from the Bureau of Economic Analysis, Table 3.2. The series for the average marginal income tax rate is taken from Mertens and Montiel Olea (2018) who update the measure of Barro and Redlick (2011) until 2012. Following Ramey (2011), the annual tax series are converted to quarterly assuming that the tax rate does not change during the fiscal year. For total factor productivity, I use the real-time, quarterly series on TFP for the U.S. business sector, adjusted for variations in factor utilization (labor effort and capital's workweek), constructed by John Fernald, which is available on the website of the Federal Reserve Bank of San Francisco. Total hours worked series is constructed as the product of average weekly hours in the nonfarm business sector and the civilian employment level, which are also available from FRED. Wages correspond to compensation per hour in the non-farm business sector. The inflation rate is constructed as the annualized rate of change of the GDP deflator. Nominal series for output, consumption, wages and government expenditures series are deflated using the GDP deflator. All quarterly series are seasonally adjusted and quantity variables are expressed in logs of per capita amounts.

1.A.2 SPF

To build a measure of government spending shocks, I follow the approach of Ramey (2011) and use data from the Survey of Professional Forecasters (SPF), which is available from the website of the Federal Reserve Bank of Philadelphia. In this survey, professional forecasters, mostly from the private sector, are asked to provide forecast values for a number of macroeconomic variables for the present quarter and up to four quarters ahead. Regarding real federal government consumption expenditures and gross investment, which is the variable of interest to build the shock, individual forecasts are available from 1981Q3 onwards. As data on macroeconomic variables are released with a lag, when the forecasts are made, the forecasters only know the value of these variables in the previous quarter, but not in the current one. As is customary, to build the shock, the difference in the growth rates is preferred to the difference in the levels as the base year changed multiple times during the sample period.

1.A.3 CEX

Household level data on consumption and hours worked is from the Interview portion of the Consumer Expenditure Survey (CEX) conducted by the Bureau of Labor Statistics (BLS). The CEX Interview Survey is a rotating panel of approximately 7,000 households, selected to be representatives of the U.S. population, who are interviewed about their expenditures for up to four consecutive quarters. The survey records information on detailed categories of consumption expenditures over the preceding quarter for all households interviewed. In addition, the survey provides detailed demographic characteristics for all household members, as well as information on their labor and financial income, which I exploit in my empirical analysis.

The household is identified with the head of the household. Following Anderson et al. (2016), all households with missing data or implausible consumption or income data are dropped, as well as households whose head is aged more than 75. My final sample contains 171,090 households over the period 1981Q3-2007Q4. Similar to Krueger and Perri (2006), nondurable consumption is defined as expenditures on food, alcoholic beverages, tobacco, apparel and services, personal care, household operations, public transportation, gas and motor oil, medical care, entertainment, reading material and education. Consumption expenditures are measured in log of real per capita terms.[45] All variables are seasonally adjusted by X-12 ARIMA.

Given the short panel dimension of the dataset, I follow the strategy described in Deaton (1985) and build pseudo-panels, which consists in aggregating individual observations into pseudo-cohorts of consumers with different characteristics and computing averages for each period. Several concerns have been raised in the literature regarding CEX data, such as the presence of measurement error and underreporting by high-income households.[46] An advantage of this approach is that it attenuates

[45]Household consumption expenditures data is divided by the number of family members and deflated by the consumer price index.

[46]See, e.g., Lusardi (1996), Aguiar and Hurst (2013) and Aguiar and Bils (2015) for a discussion of these issues.

the attrition problem and reduces measurement error since it aggregates across agents. Furthermore, income data, where measurement error is more salient, is only used to identify income groups.

1.A.4 CPS

To inspect the effects of government spending shocks on age-specific labor-market outcomes, I use data from the Uniform Extracts of the Current Population Survey (CPS) Merged Outgoing Rotation Group (MORG) from the CEPR. The CPS is the source of official US government labor market statistics. It provides information on individuals' employment status, hourly wages and weekly hours worked, as well as on their education and demographic characteristics. Interviewed households are selected to be representative of the U.S. population. About 60,000 households are interviewed for four consecutive months one year, then ignored for eight months, and interviewed again for four consecutive months. Individuals pursuing studies, self-employed and individuals with zero or missing wage are excluded from the sample. Hours worked are computed as the product of usual weekly hours and the number of persons employed. The extremely large sample size of the CPS dataset allows accurate analyses at a high degree of disaggregation. Therefore I further use this dataset to build measures of productivity for different age groups using a wage-based approach. The series are seasonally adjusted by X-12 ARIMA. My final sample contains about 5 million observations over 1981Q3-2007Q4.

1.A.5 Cross-country Panel Data

The data series on real GDP, real government consumption expenditure, real private consumption, the current account and the real effective exchange rate used in Section 1.3.6 are taken from Ilzetzki et al. (2013). These quarterly series cover the period from 1960:1 to 2009:4 for 44 developing and developed countries. I extend this dataset by collecting series on labor productivity, measured as GDP per employed person. These series are obtained from the OECD (Main Economic Indicators), Eurostat and Oxford Economics. Five countries are excluded from the sample as there is no quarterly series available for labor productivity: Botswana, Ecuador, El Salvador, Peru, Uruguay. Similar to Ilzetzki et al. (2013), the productivity series are seasonally adjusted and analyzed as deviations from their quadratic trend.

The remaining 39 countries are split in two groups according to their share of young people in total population. Population shares by age groups are computed using annual data from the World Population Prospects prepared by the Population Division of the Department of Economic and Social Affairs of the United Nations Secretariat. Table 1.A.1 lists the countries included in each group. Countries with high share of young in total population are characterized by an average share of people aged 15-29 over 1970-2010 above the sample mean of 23.6%, while the second group consists of countries with a share strictly below the sample mean. The distribution of countries in the two groups remains unchanged if the share of people aged 15-34 is considered instead of 15-29.

Table 1.A.1: Share of young people in total population across panel of 39 countries

High Share of Young		**Low Share of Young**	
Colombia	28.8%	Romania	23.1%
Malaysia	28.1%	Spain	22.8%
Mexico	28.1%	Slovenia	22.6%
Brazil	28.1%	Netherlands	22.6%
South Africa	28.0%	Lithuania	22.5%
Thailand	27.3%	Portugal	22.3%
Turkey	27.1%	Czech Republic	22.3%
Chile	26.9%	Greece	21.9%
Israel	25.1%	Finland	21.8%
Ireland	24.8%	France	21.7%
Iceland	24.7 %	Estonia	21.6%
Slovakia	24.5 %	Croatia	21.6%
Argentina	24.3%	Latvia	21.6%
Poland	24.3%	Hungary	21.6%
Canada	24.1%	Norway	21.5%
Australia	23.6%	Bulgaria	21.2%
United States	23.6%	United Kingdom	21.2%
		Belgium	21.1%
		Denmark	21.0%
		Italy	20.8%
		Germany	20.4%
		Sweden	20.1%
Mean	**26.0%**		**21.7%**

Notes: This table reports the average share of young people (aged 15-29) among total population over 1970-2010 across a panel of 39 countries. Overall sample mean is 23.6%.

1.B Aggregate Results

In this Appendix, I present time series evidence on the effects of government spending shocks on macroeconomic variables for the U.S.

I use a structural vector autoregression approach, where the vector of endogenous variables includes, in this order, TFP, the SPF shock, log real per capita quantities of government spending, GDP, nondurable consumption and total hours worked, as well as log real wages, the average marginal income tax rate, the three-month T-Bill rate and the inflation rate. The model is estimated with two lags, a constant and a quadratic trend on the same sample as in the baseline analysis.

Using a standard Choleski decomposition, shocks to government spending are identified as innovations to the SPF forecast-based measure which are orthogonal to TFP movements on impact and pre-determined with respect to remaining variables. This specification controls for the measurement error component that may induce a bias in the impulse responses of output and TFP, following Ben Zeev and Pappa (2015).[47] Furthermore, following Blanchard and Perotti (2002), this specification implies that government spending cannot react to changes in remaining variables within the same quarter due to implementation lag.

Figure 1.B.1 displays the impulse responses of these variables to a shock that raises government spending by one percent. Following a fiscal expansion, output, hours worked, wages and consumption increase significantly, while the nominal interest rate and the inflation rate drop. TFP also rises significantly after a few quarters, and the marginal income tax rate increases during the first quarters after the shock. These findings are in line with empirical estimates already reported in the litterature, notably in d'Alessandro et al. (2019) and Jørgensen and Ravn (2018).

[47] Ben Zeev and Pappa (2015) find that the positive response of output and TFP to unexpected government spending shocks could be due to correlated measurement error in the two variables, and show that forcing the fiscal shocks to be orthogonal contemporaneously to TFP fluctuations enables to properly identify the true effects of the shocks on macroeconomic variables.

Figure 1.B.1: Impulse responses of aggregate variables to government spending shocks

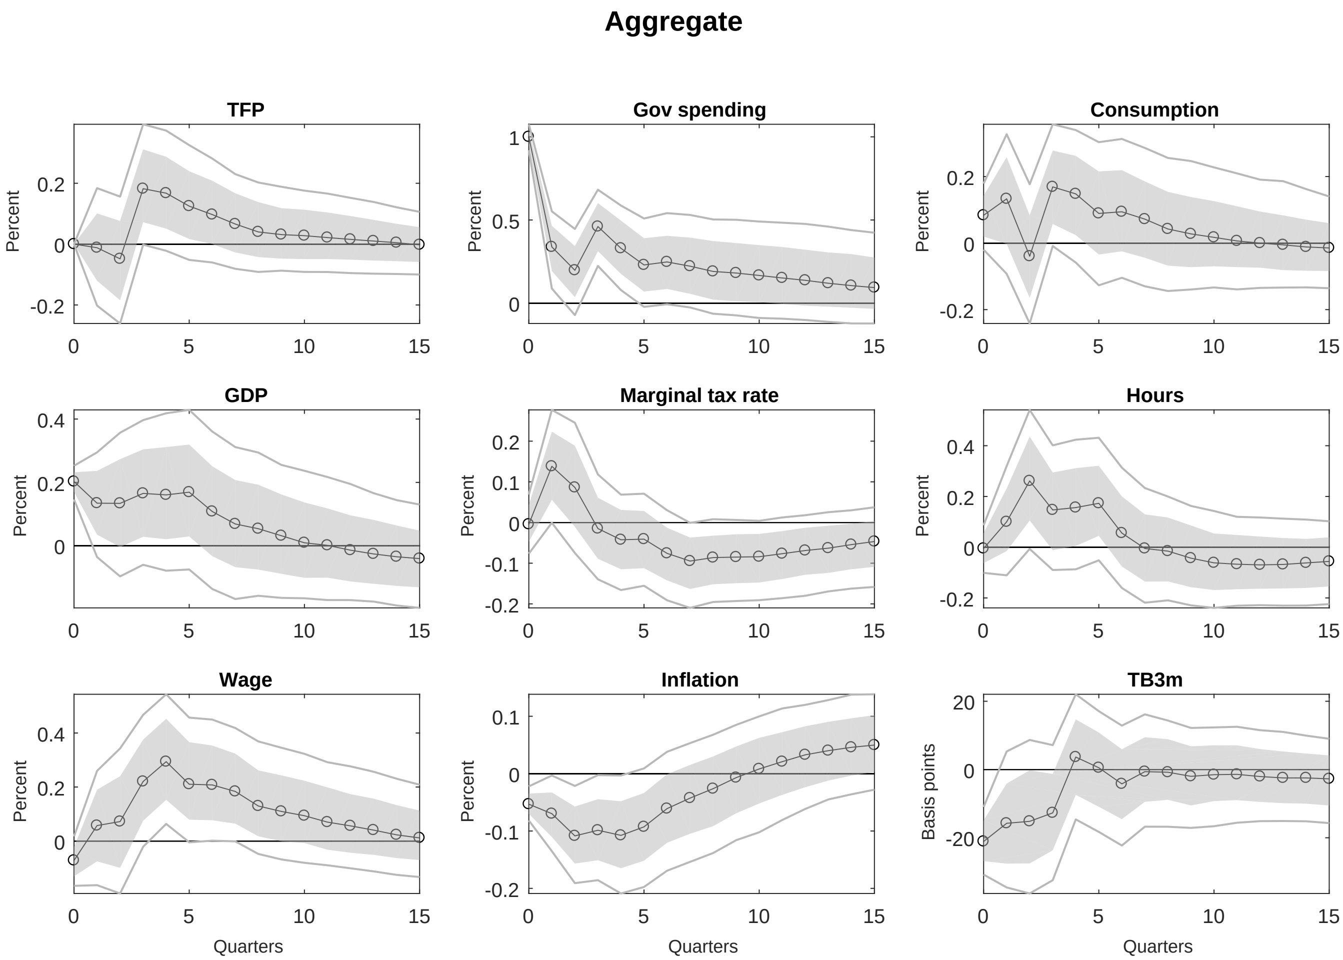

Notes: These graphs show the impulse responses of aggregate variables to an exogenous government spending shock leading to an initial 1% increase in government expenditures. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

1.C Fiscal Multipliers

Figure 1.C.1: Cumulative Fiscal Multipliers

(a) Consumption

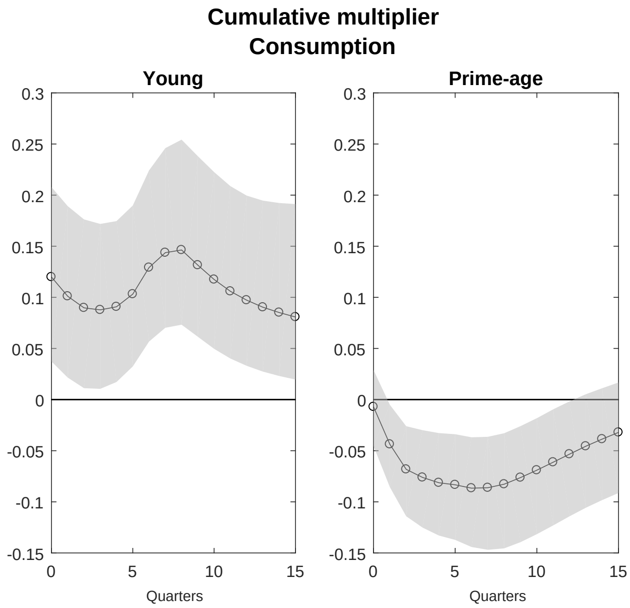

(b) Productivity

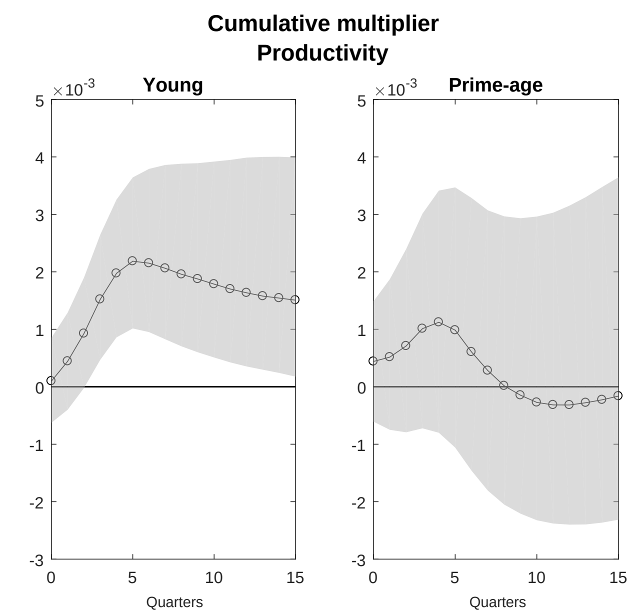

Notes: These graphs show the cumulative fiscal multipliers of nondurable consumption (Panel (a)) and productivity (Panel (b)) for young and prime-age households. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 68% confidence intervals are shown in all cases.

1.D Robustness

1.D.1 Heterogeneous Effects on Consumption: Robustness

Using other proxies for financial constraints

Figure 1.D.1: Impulse responses to government spending shocks by age and education level

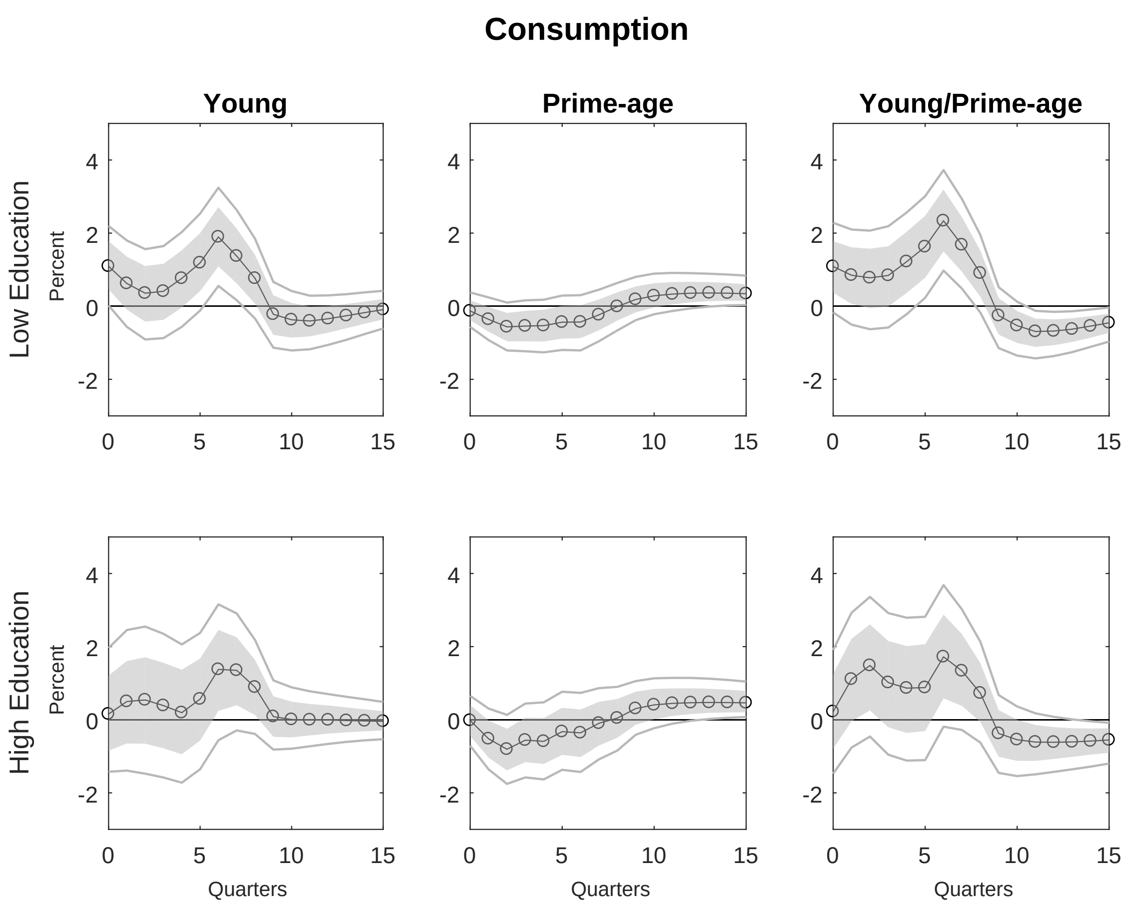

Notes: These graphs show the impulse responses of nondurable consumption for young and prime-age households, grouped by their education level ("Low": no college degree, "High": college degree) to an exogenous government spending shock leading to an initial 1% increase in government expenditures. The impulse responses for the young (below 30) are depicted on the first column, for the prime-age (30-64) on the second column, and for the ratios between young and prime-age on the last column. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Figure 1.D.2: Impulse responses to government spending shocks by age and financial market participation

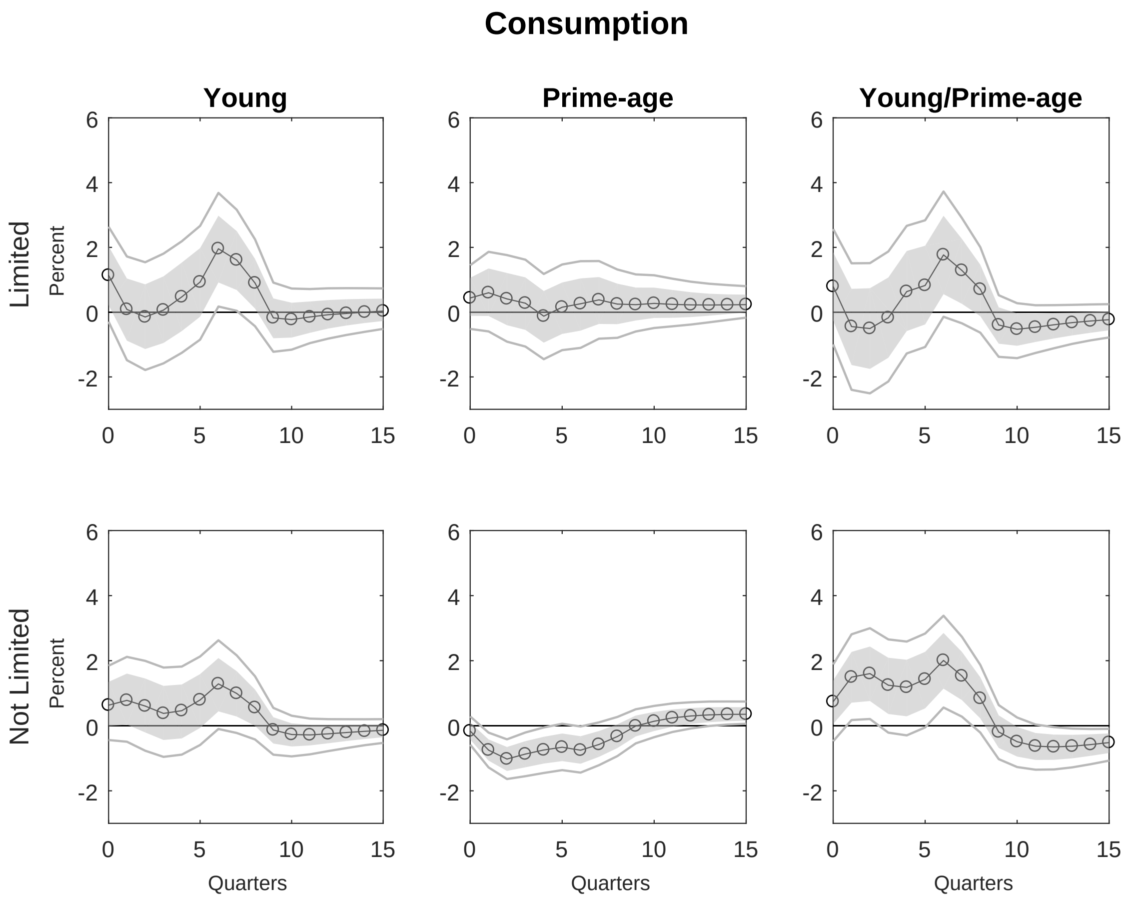

Notes: These graphs show the impulse responses of nondurable consumption for young and prime-age households, grouped by their financial market participation ("Limited": no income from financial assets, "Not Limited": non-zero income from financial assets) to an exogenous government spending shock leading to an initial 1% increase in government expenditures. The impulse responses for the young (below 30) are depicted on the first column, for the prime-age (30-64) on the second column, and for the ratios between young and prime-age on the last column. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Using a broader definition of consumption

Figure 1.D.3: Impulse responses of consumption for young and prime-age groups

(a) By income level

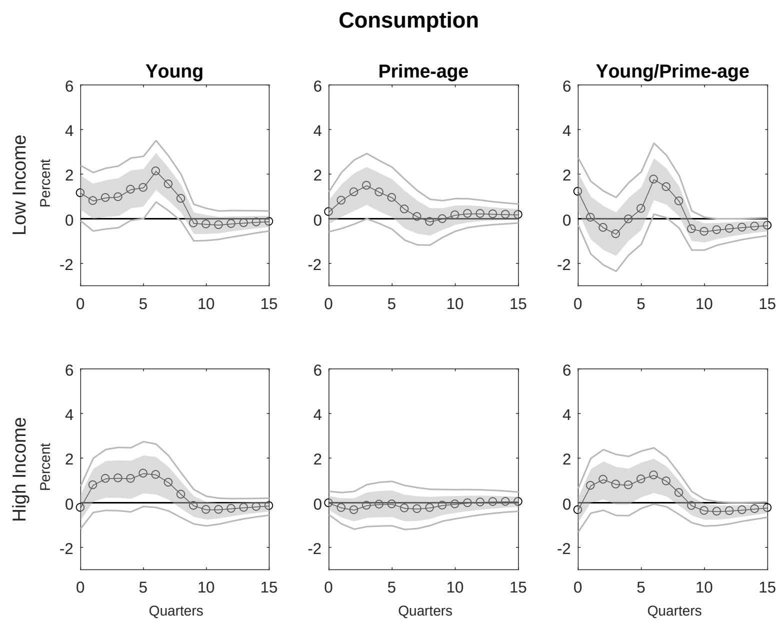

(b) By housing tenure

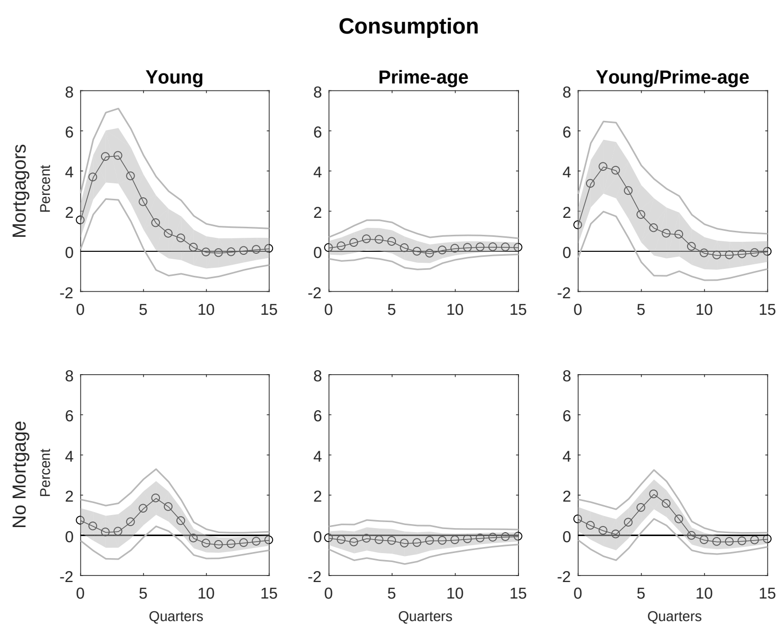

Notes: These graphs show the impulse response functions of nondurable consumption for young and prime-age households, by their income level in Panel (a) and by their housing tenure in Panel (b), to an exogenous government spending shock leading to an initial 1% increase in government expenditures. The impulse response functions of the ratios between young and prime-age are displayed in the last column. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Restricting the sample to employed households

Figure 1.D.4: Impulse responses of consumption for young and prime-age groups

(a) By income level

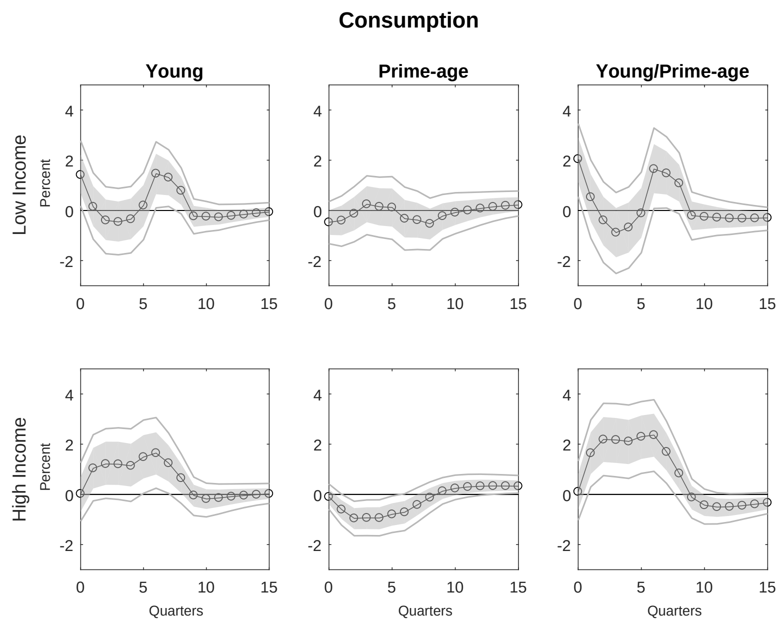

(b) By housing tenure

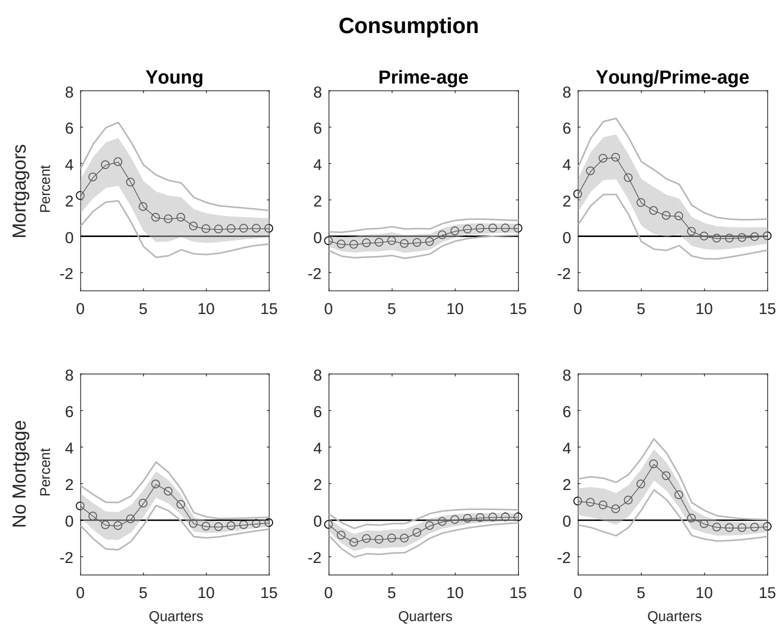

Notes: These graphs show the impulse response functions of nondurable consumption for young and prime-age households, by their income level in Panel (a) and by their housing tenure in Panel (b), to an exogenous government spending shock leading to an initial 1% increase in government expenditures. The impulse response functions of the ratios between young and prime-age are displayed in the last column. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Compositional change

Figure 1.D.5: Impulse responses of consumption for young and prime-age groups - using propensity score approach

(a) By income level

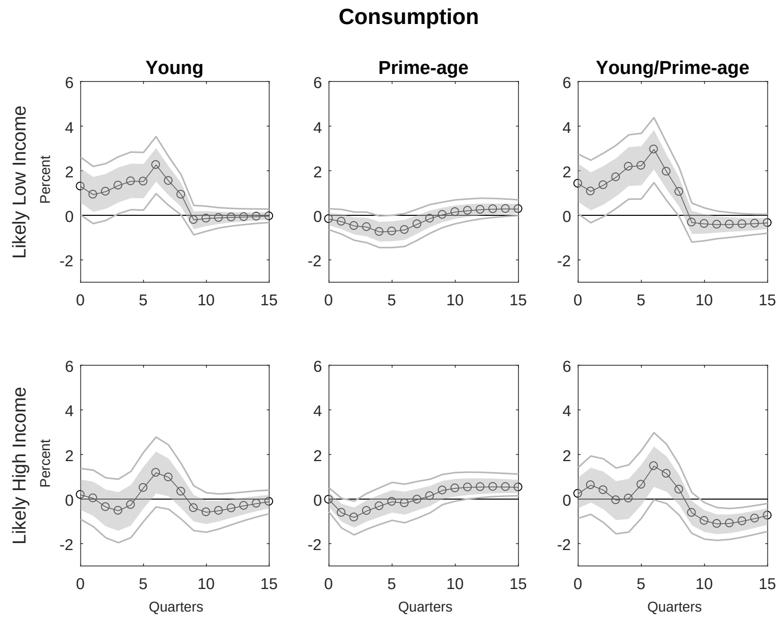

(b) By housing tenure

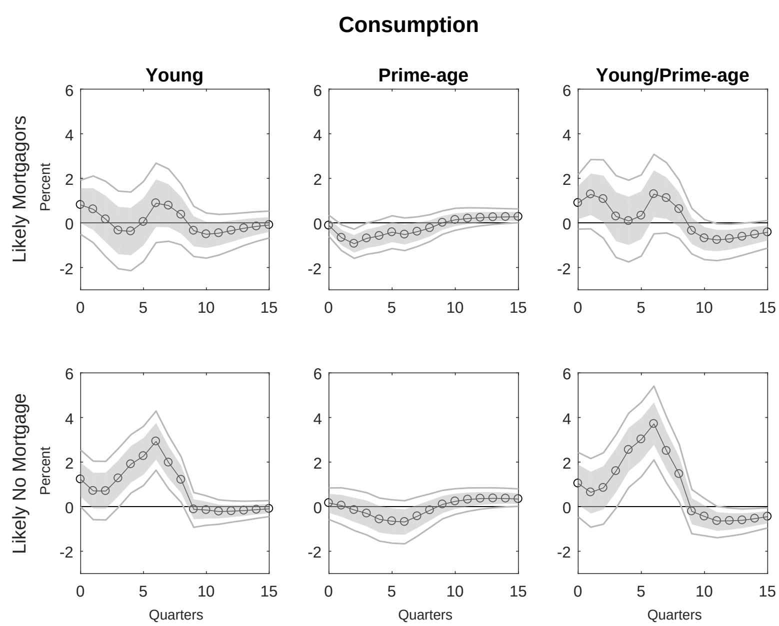

Notes: These graphs show the impulse response functions of nondurable consumption for young and prime-age households, by their income level in Panel (a) and by their housing tenure in Panel (b), to an exogenous government spending shock leading to an initial 1% increase in government expenditures. Groups are computed following Attanasio et al. (2002) propensity score approach using a fixed probability threshold. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Figure 1.D.6: Impulse responses of shares

(a) By income level

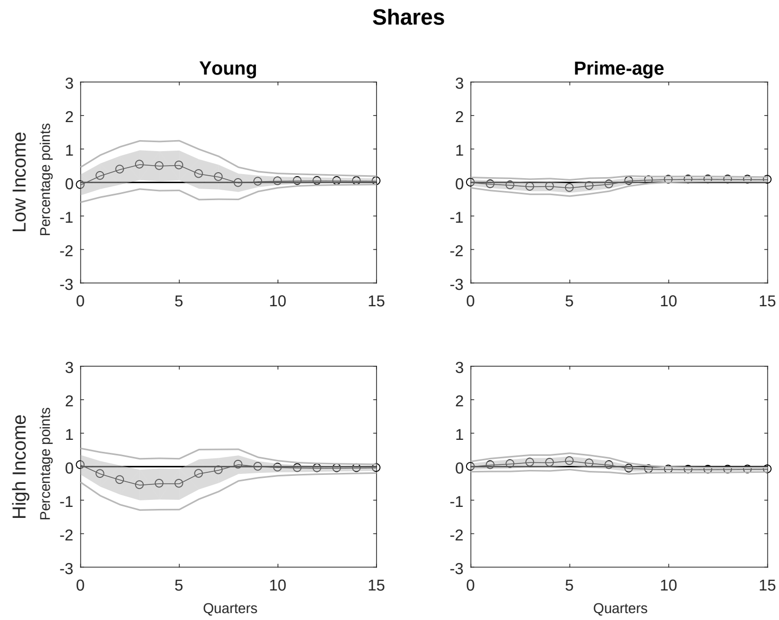

(b) By housing tenure

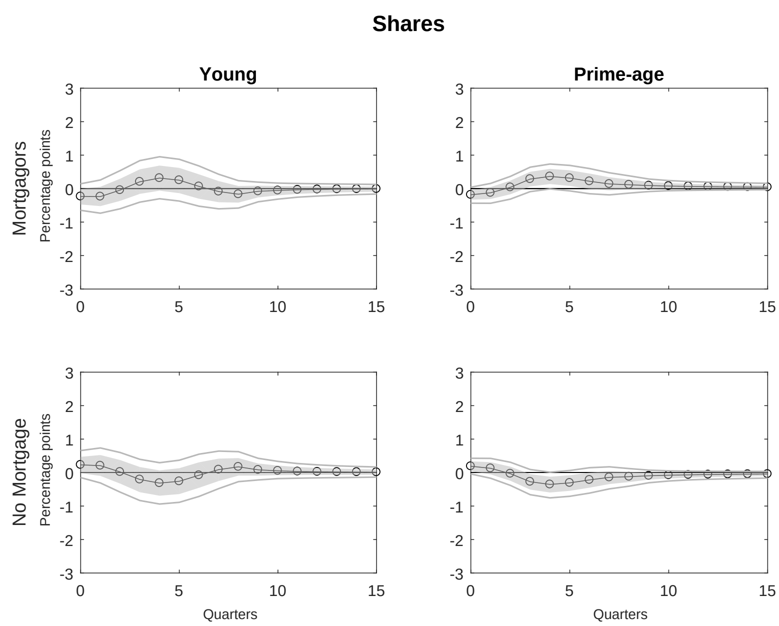

Notes: These graphs show the impulse response functions of the shares of young and prime-age households, by their income level in Panel (a) and by their housing tenure in Panel (b), to an exogenous government spending shock leading to an initial 1% increase in government expenditures. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

1.D.2 Fiscal Policy and Human Capital: Robustness

Micro Evidence

Figure 1.D.7: Impulse responses of productivity to government spending shocks - sensitivity to the flat spot range

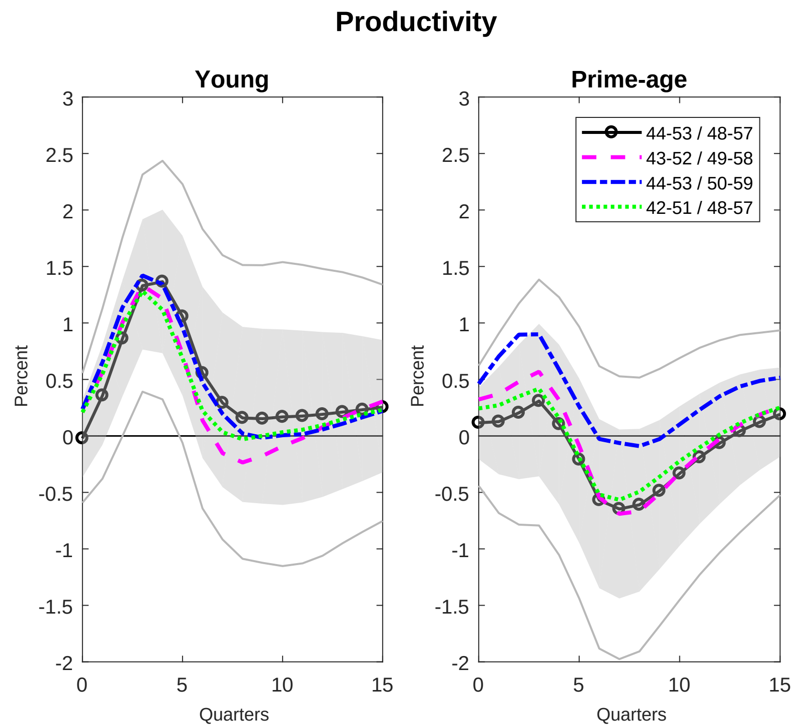

Notes: These graphs show the impulse responses of measured productivity to a 1% shock to government consumption expenditure for different values of the flat spot region for low- and high-educated groups. 90% and 68% confidence intervals are shown for the baseline impulse response function (flat spot: 44-53 / 48-57).

Panel Evidence

Additional controls

Figure 1.D.8: Impulse responses of productivity and consumption to government spending shocks in countries with high vs. low shares of young in total population - controlling for current account and real exchange rate

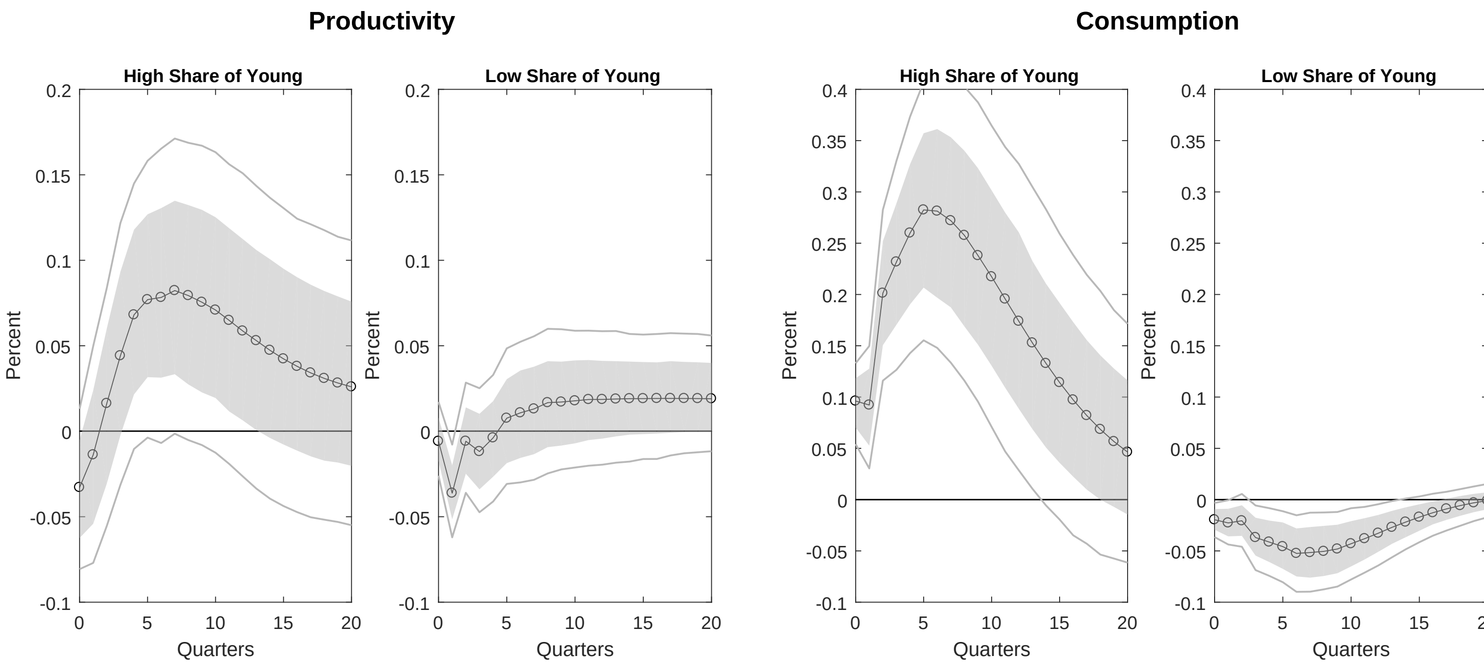

Notes: These graphs show the impulse responses of labor productivity (left panels) and private consumption (right panels) to a 1% shock to government consumption expenditure in countries with low share of young (aged 15-29) in total population vs. high share of young. 90% and 68% confidence intervals are shown in all cases.

1.D.3 Further Robustness

Other measure of government spending shocks

Figure 1.D.9: Impulse responses of nondurable consumption and hourly wages

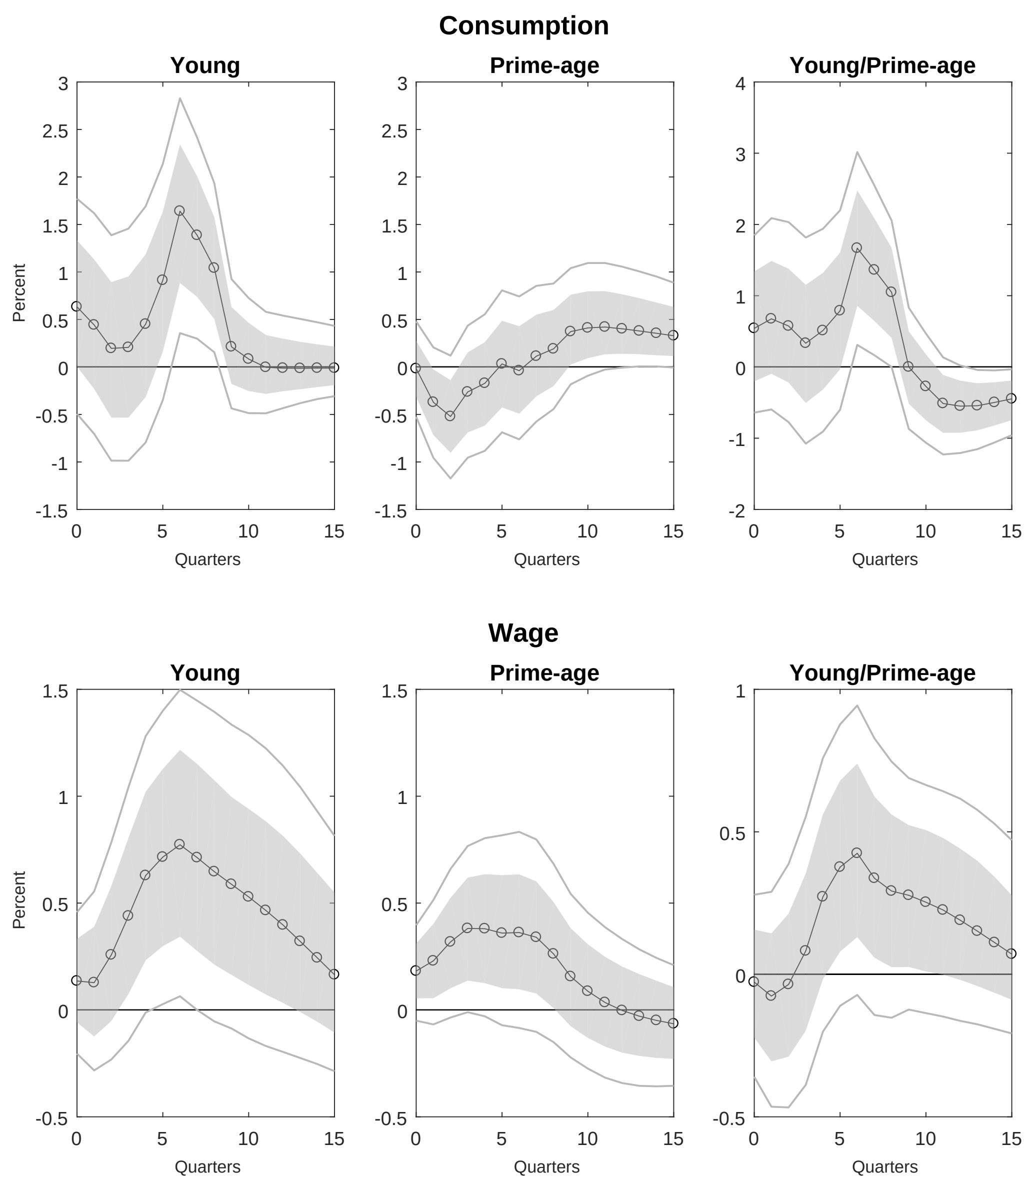

Notes: These graphs show the impulse response functions of nondurable consumption and hourly wages for young and prime-age workers, and of the ratios between young and prime-age to an exogenous government spending shock leading to an initial 1% increase in government expenditures. Several controls are added in the estimation of the shocks. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Figure 1.D.10: Impulse responses of hours worked and productivity

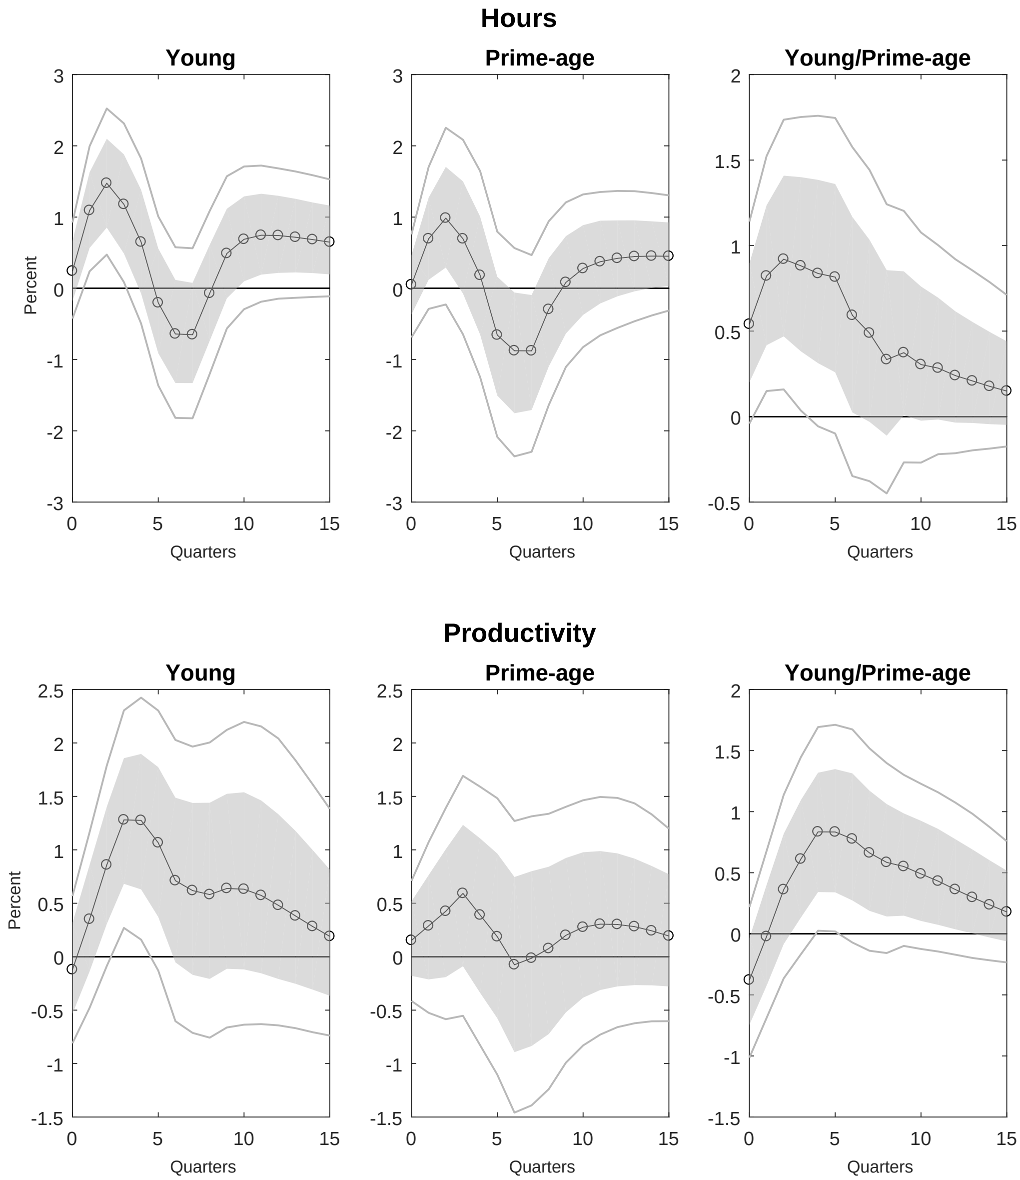

Notes: These graphs show the impulse response functions of hours worked and productivity for young and prime-age workers, and of the ratios between young and prime-age to an exogenous government spending shock leading to an initial 1% increase in government expenditures. Several controls are added in the estimation of the shocks. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Additional controls

Figure 1.D.11: Impulse responses of nondurable consumption and hourly wages

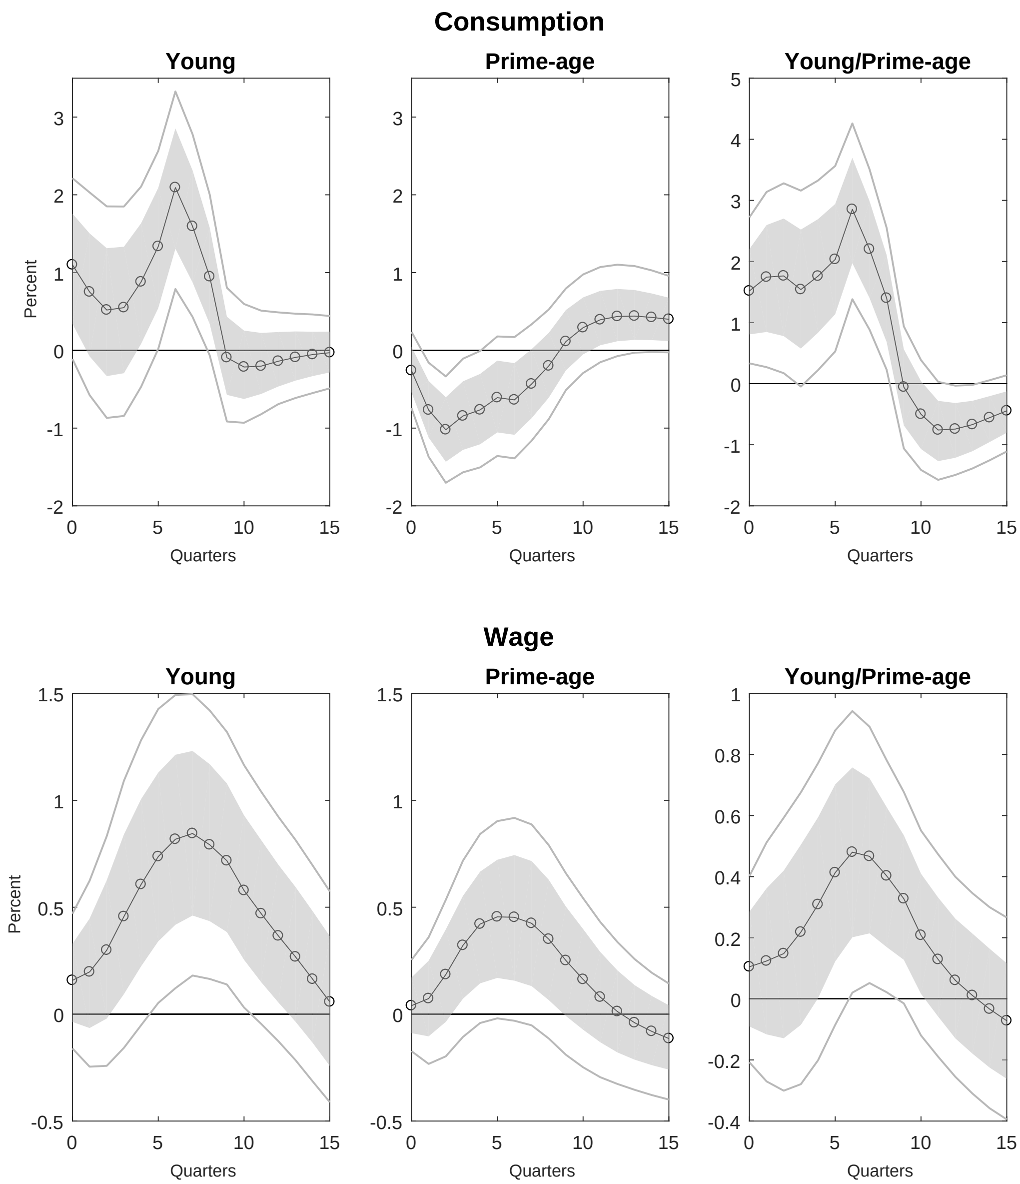

Notes: These graphs show the impulse response functions of nondurable consumption and hourly wages for young and prime-age workers, and of the ratios between young and prime-age to an exogenous government spending shock leading to an initial 1% increase in government expenditures. Several variables are added in the VAR model. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

Figure 1.D.12: Impulse responses of hours worked and productivity

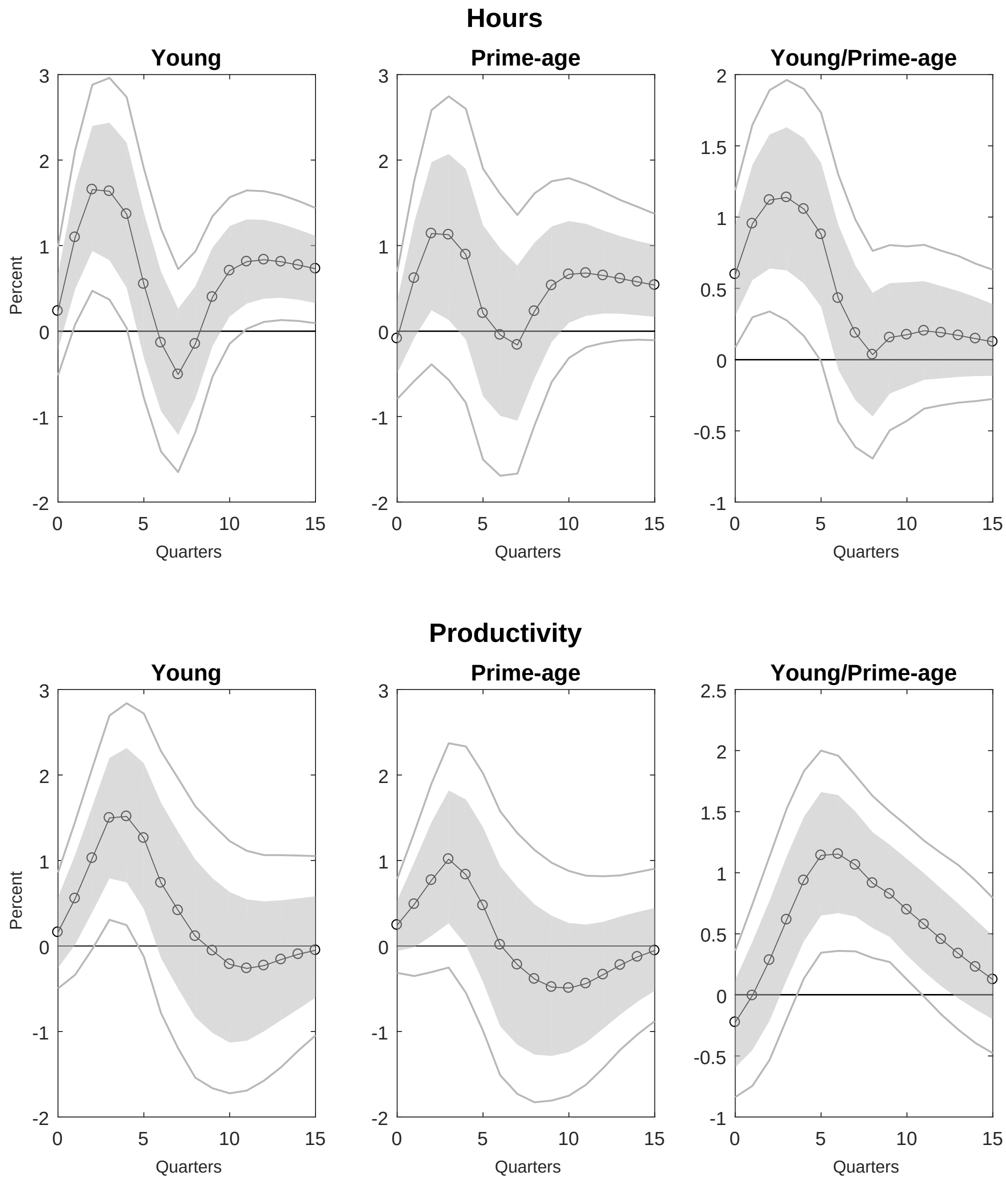

Notes: These graphs show the impulse response functions of hours worked and productivity for young and prime-age workers, and of the ratios between young and prime-age to an exogenous government spending shock leading to an initial 1% increase in government expenditures. Several variables are added in the VAR model. Except for the endpoints, the coefficients are smoothed over three consecutive periods. 90% and 68% confidence intervals are shown in all cases.

1.E Model: Sensitivity Analysis

Figure 1.E.1: Impulse responses to a government spending shock - With vs. without nominal rigidities

(a) IRFs of aggregate variables

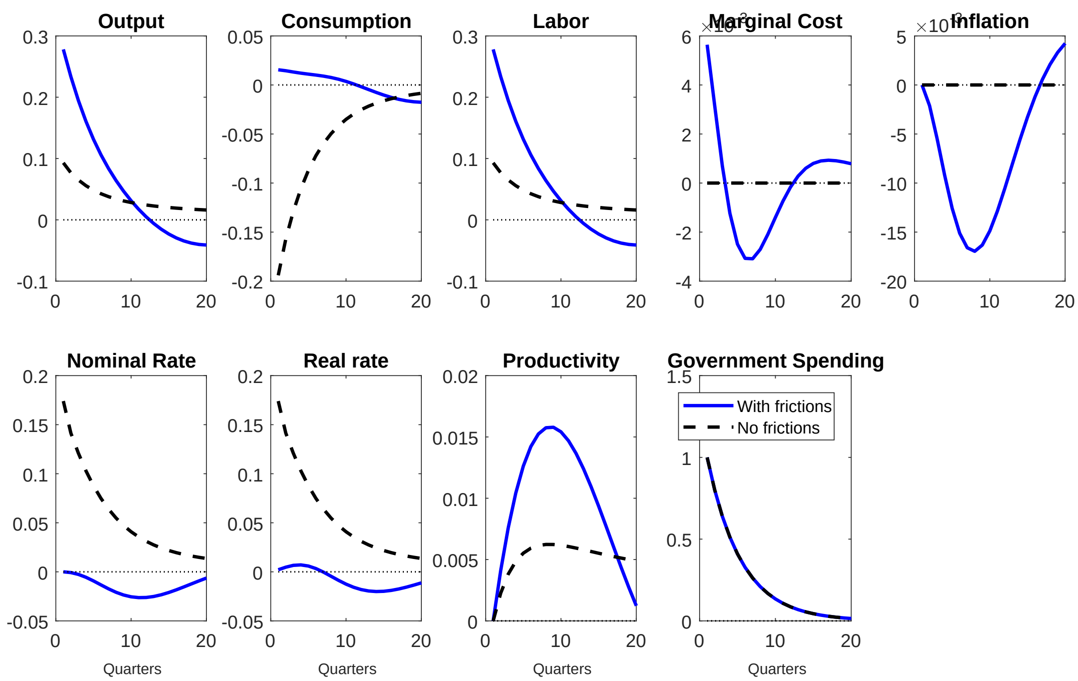

(b) IRFs of disaggregate variables

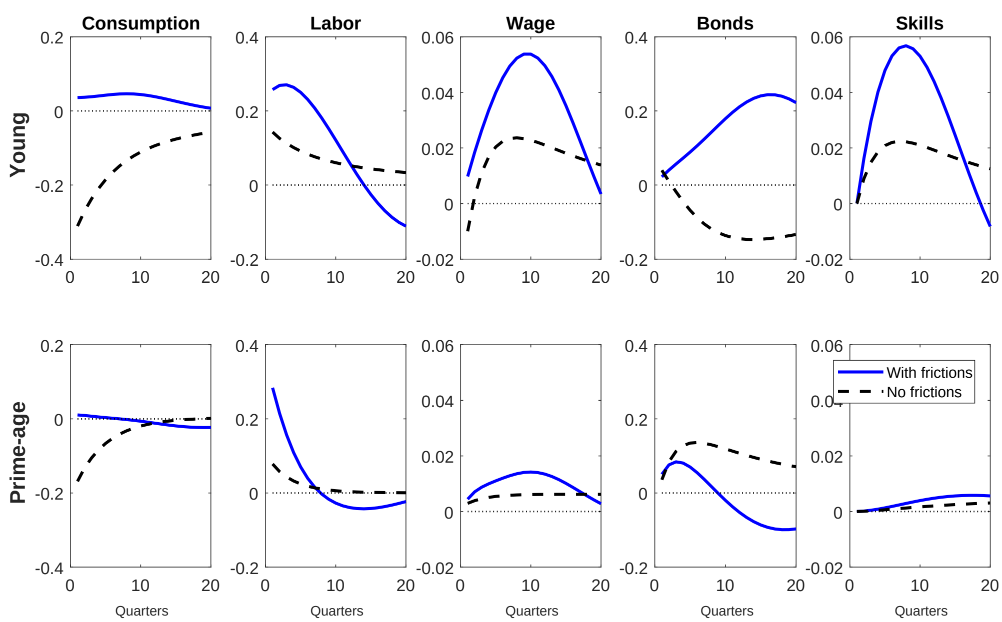

Notes: this figure shows the impulse responses of aggregate variables (Panel (a)) and disaggregate variables for young and prime-age workers (Panel (b)) to a 1% shock to government expenditure in the life-cycle model with LBD, with and without nominal rigidities.

Figure 1.E.2: Impulse responses to a government spending shock - Sensitivity to price stickiness

(a) IRFs of aggregate variables

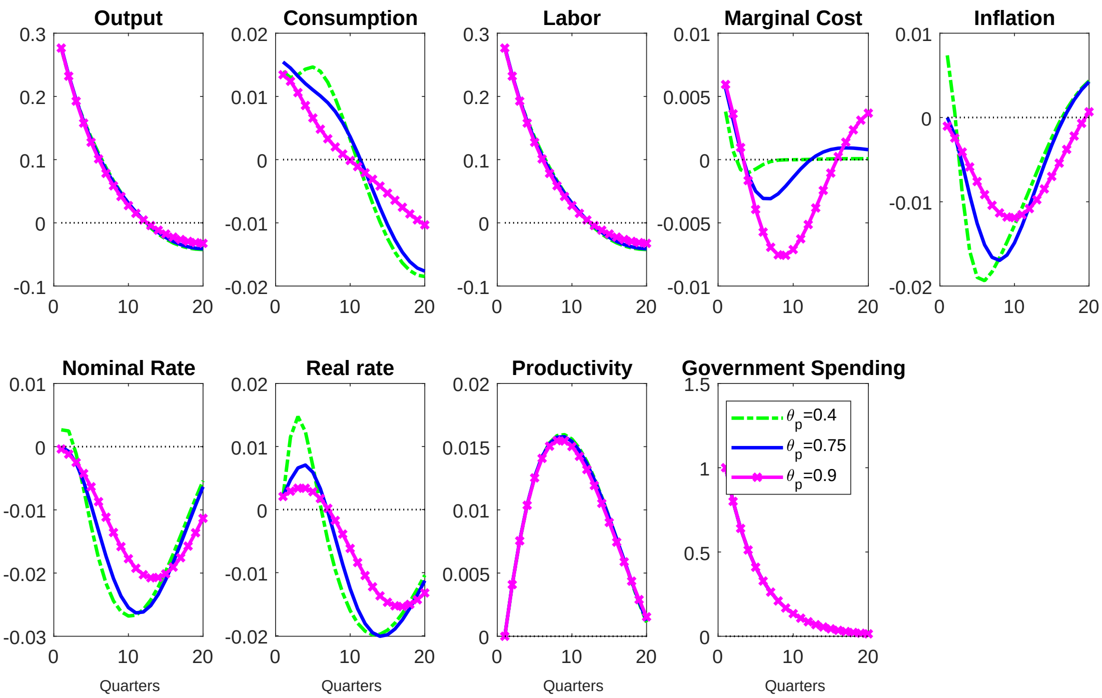

(b) IRFs of disaggregate variables

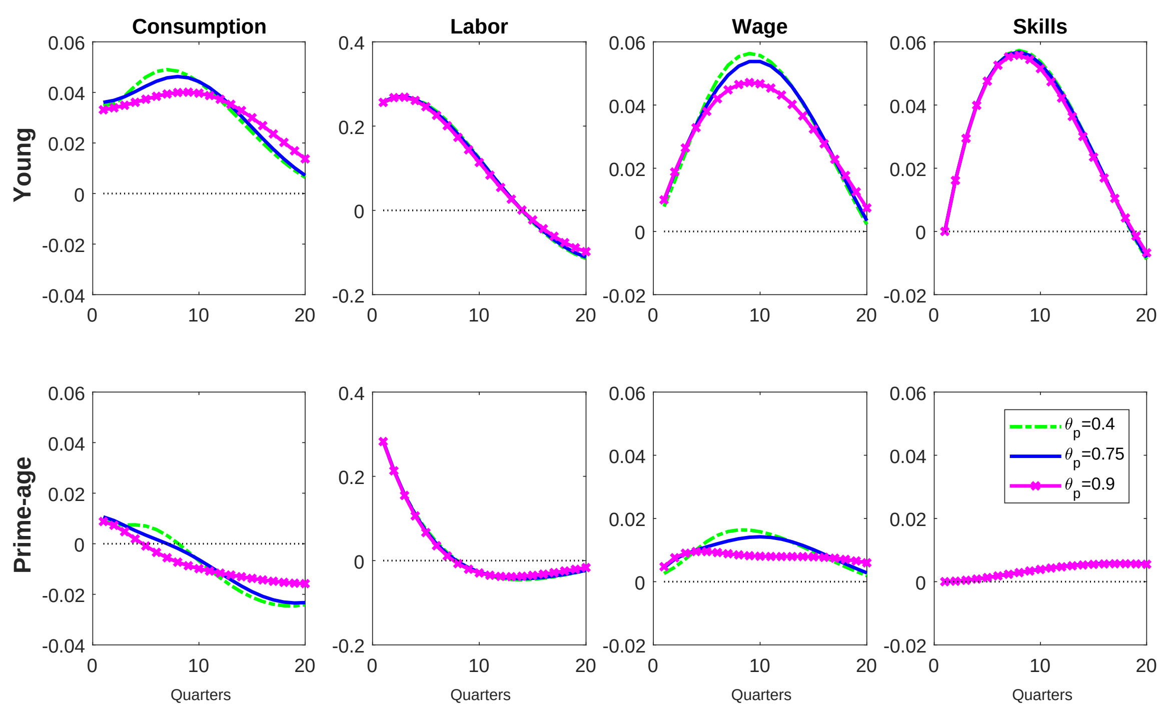

Notes: this figure shows the impulse responses of aggregate variables (Panel (a)) and disaggregate variables for young and prime-age workers (Panel (b)) to a 1% shock to government expenditure in the life-cycle model with LBD for different values of the price stickiness parameter θ_p.

Figure 1.E.3: Impulse responses to a government spending shock - Sensitivity to wage stickiness

(a) IRFs of aggregate variables

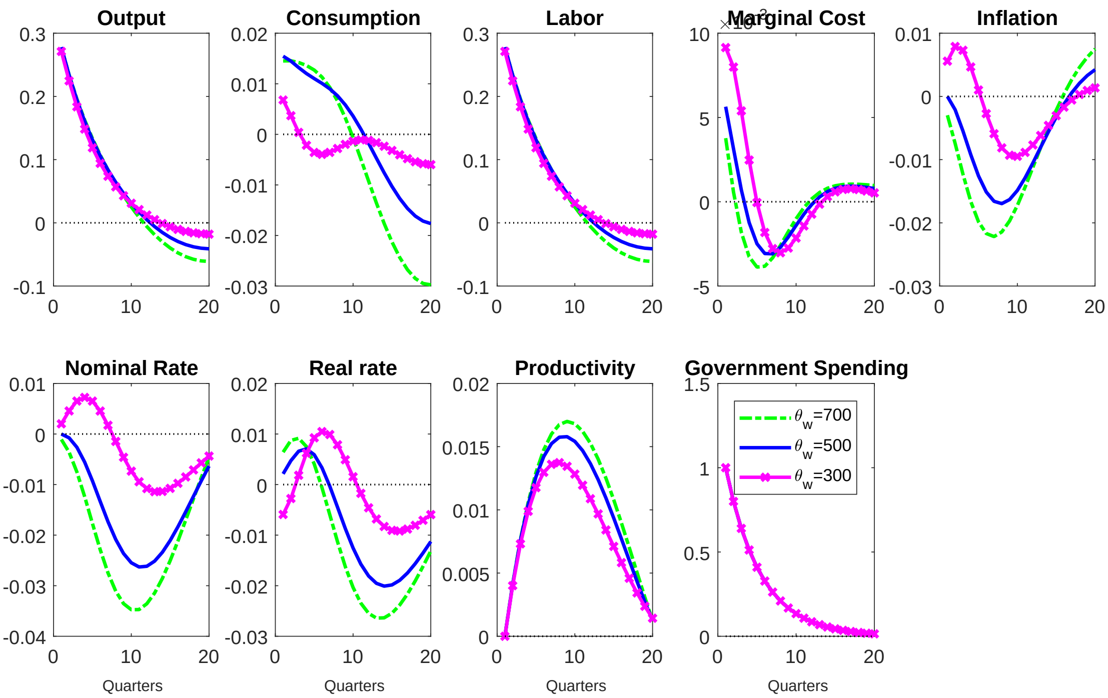

(b) IRFs of disaggregate variables

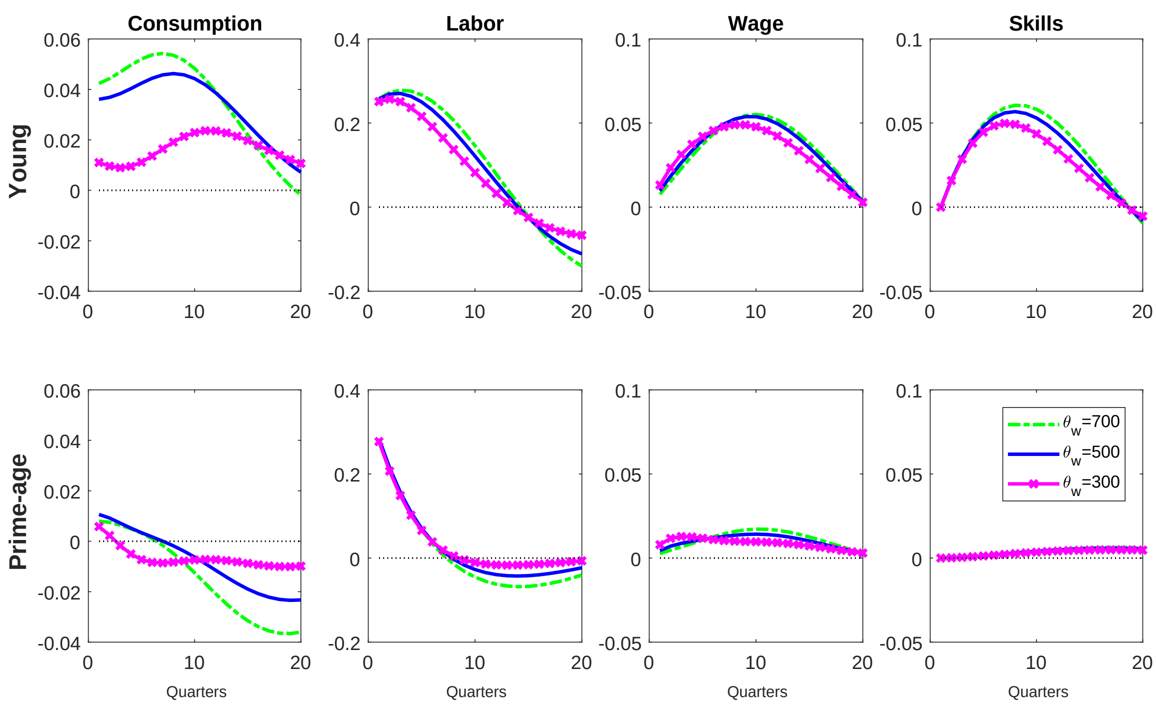

Notes: this figure shows the impulse responses of aggregate variables (Panel (a)) and disaggregate variables for young and prime-age workers (Panel (b)) to a 1% shock to government expenditure in the life-cycle model with LBD for different values of the wage adjustment cost parameter θ_w.

Figure 1.E.4: Impulse responses to a government spending shock - Sensitivity to monetary policy parameter ϕ_π

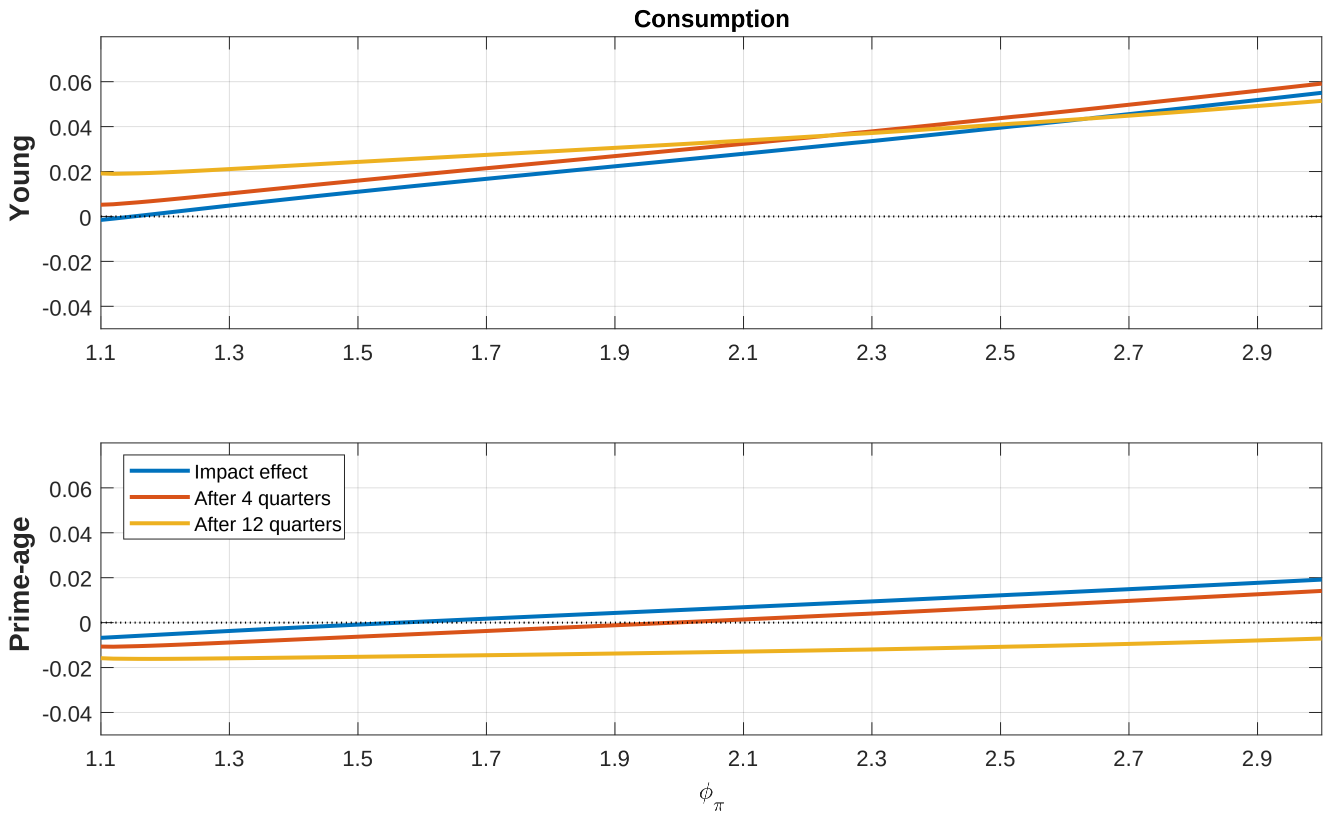

Notes: this figure shows the consumption responses for young and prime-age workers at different horizons to a 1% shock to government expenditure in the life-cycle model with LBD for different values of the Taylor rule parameter ϕ_π.

Figure 1.E.5: Impulse responses to a government spending shock - Sensitivity to shock persistence

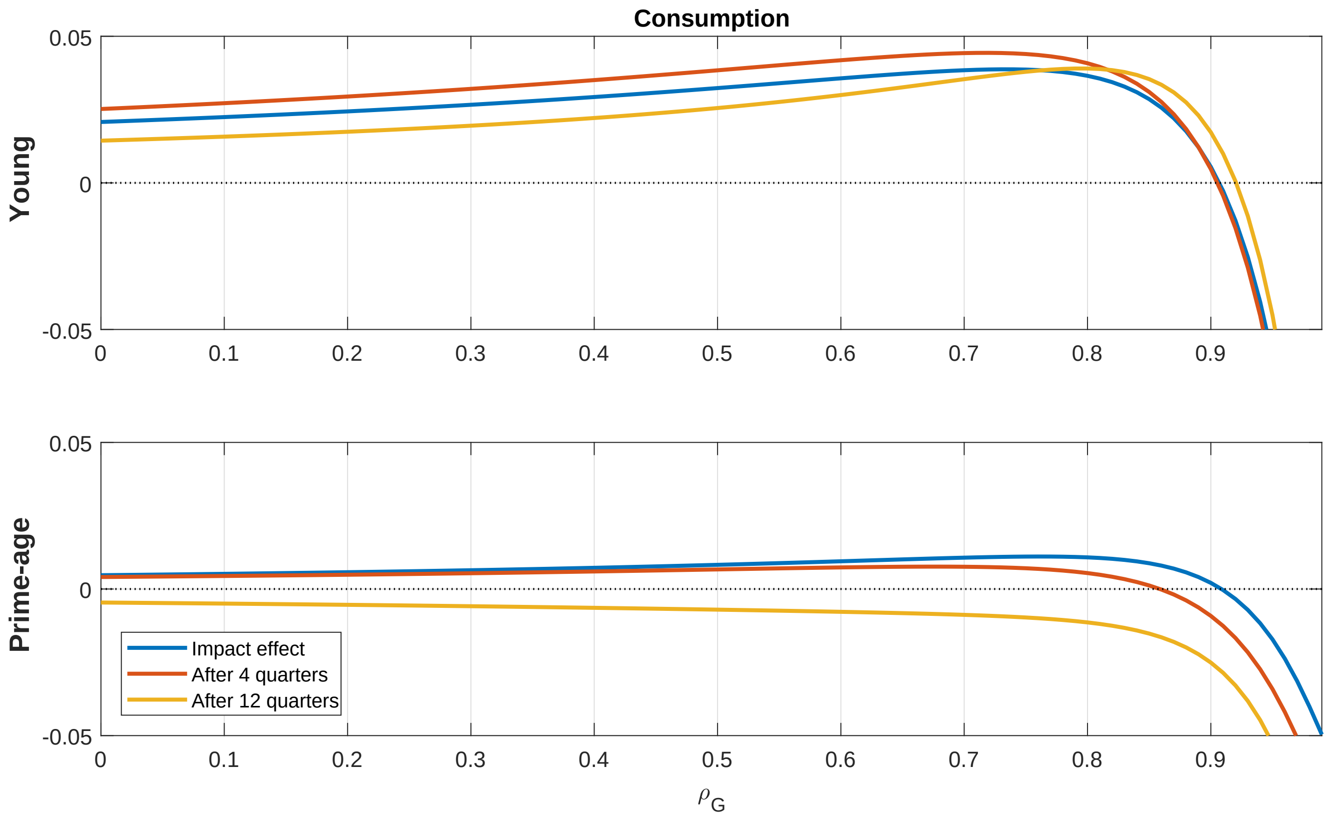

Notes: this figure shows the consumption responses for young and prime-age workers at different horizons to a 1% shock to government expenditure in the life-cycle model with LBD for different values of the persistence of the shock ρ_G.

Figure 1.E.6: Impulse responses to a government spending shock - Sensitivity to fiscal rule parameters

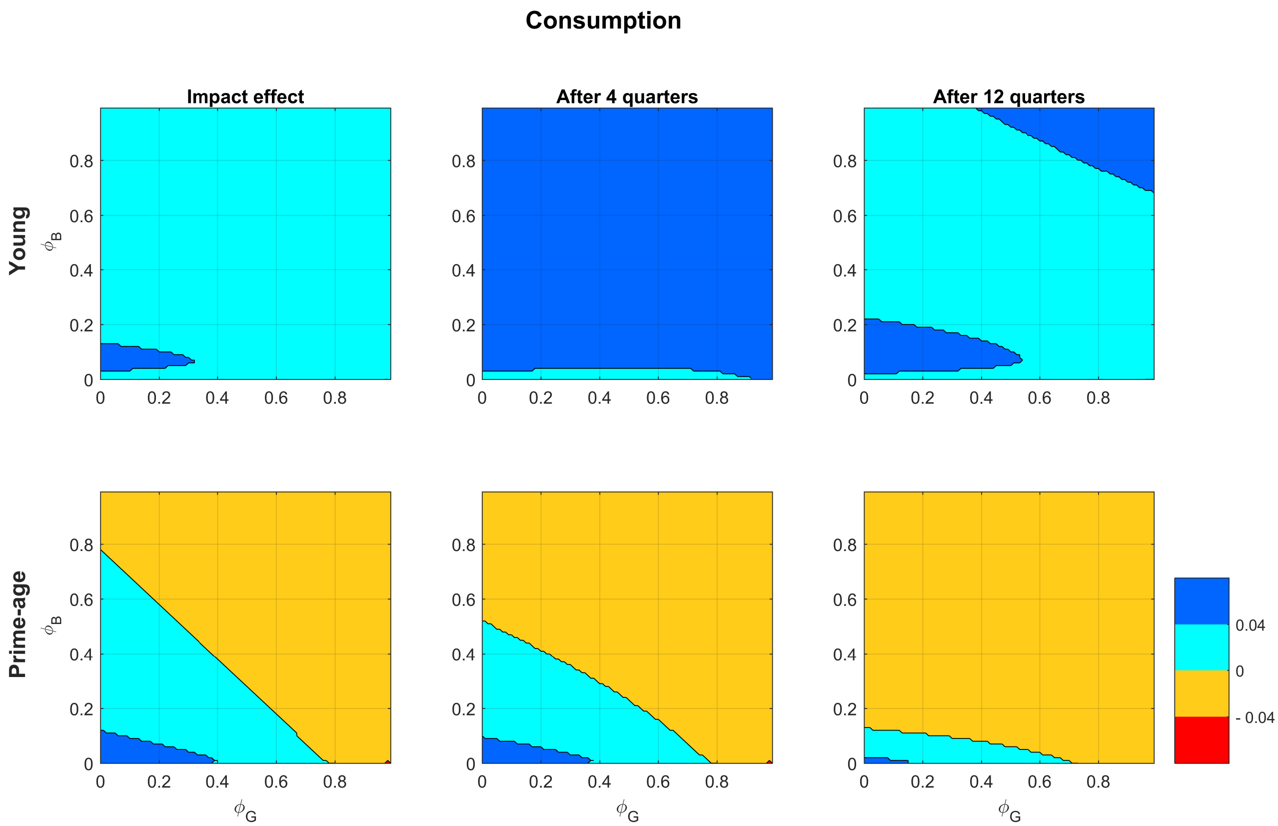

Notes: this figure shows the consumption responses for young (first row) and prime-age (second row) workers at different horizons to a 1% shock to government expenditure in the life-cycle model with LBD for different values of the fiscal rule parameters, Φ_G and Φ_B.

Figure 1.E.7: Impulse responses to a government spending shock - Sensitivity to learning-by-doing parameters

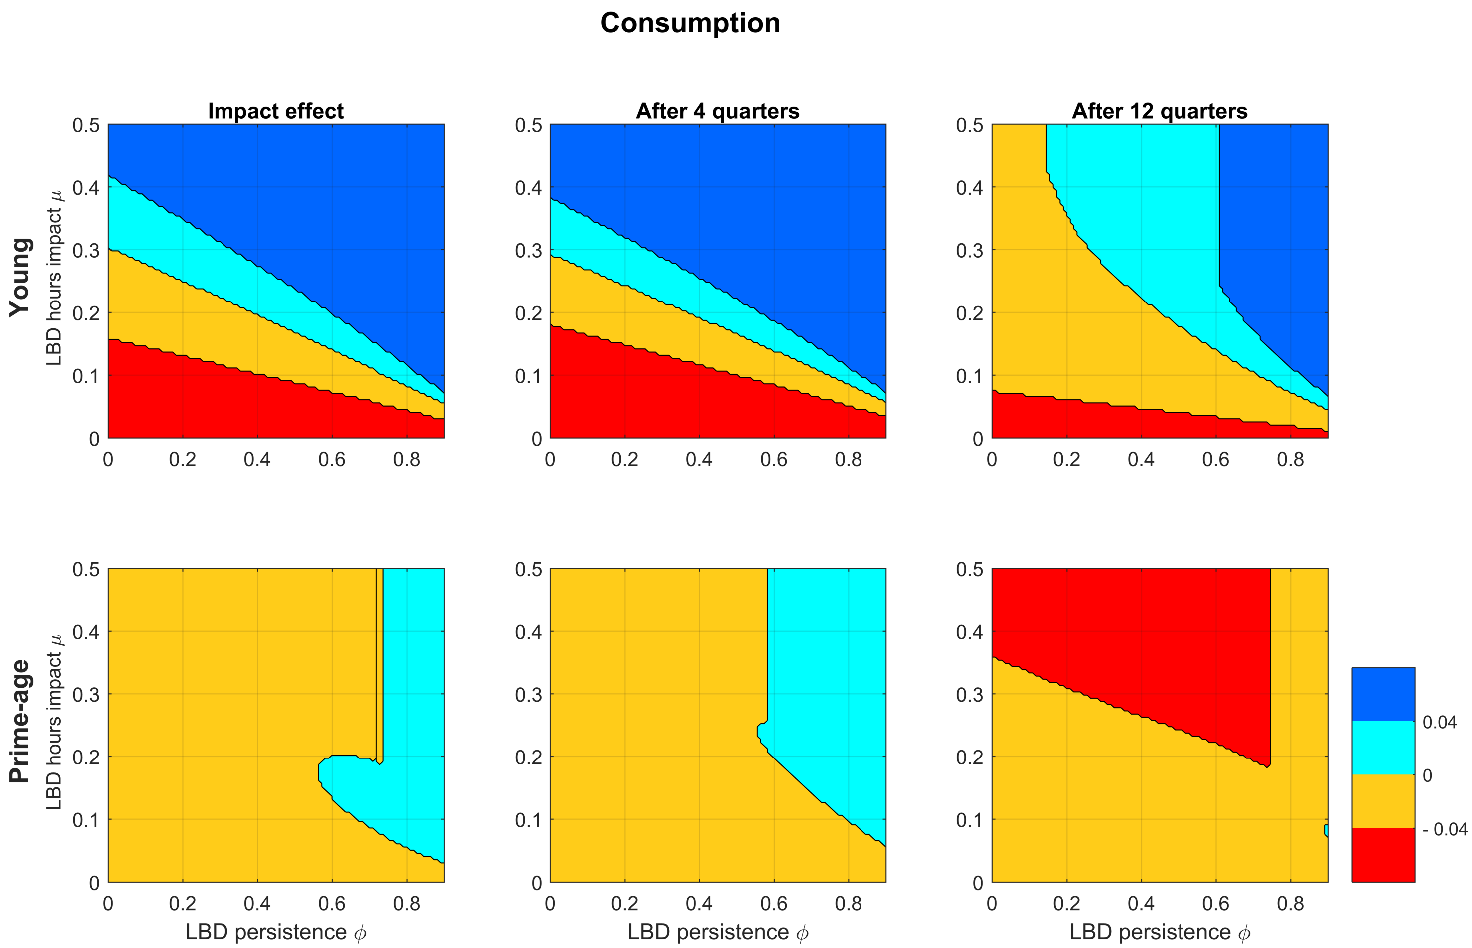

Notes: this figure shows the consumption responses for young (first row) and prime-age (second row) workers at different horizons to a 1% shock to government expenditure in the life-cycle model with LBD for different values of the parameters capturing the impact of past hours μ and the skill persistence ϕ.

Figure 1.E.8: Impulse responses to a government spending shock - Sensitivity to labor supply elasticity

(a) IRFs of aggregate variables

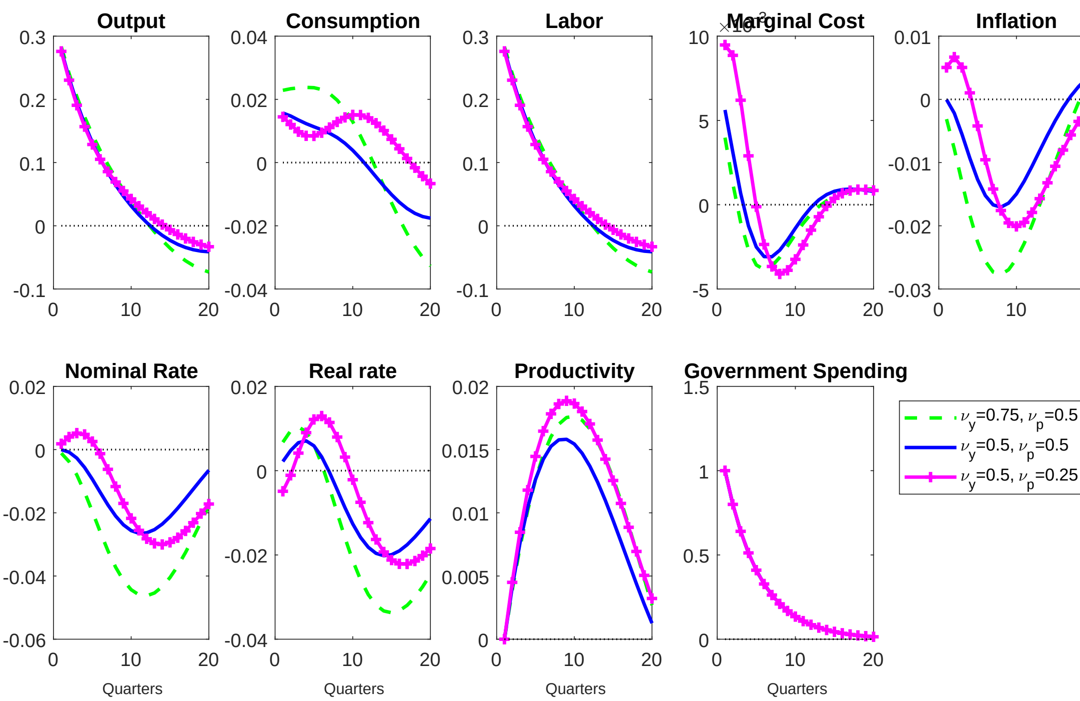

(b) IRFs of disaggregate variables

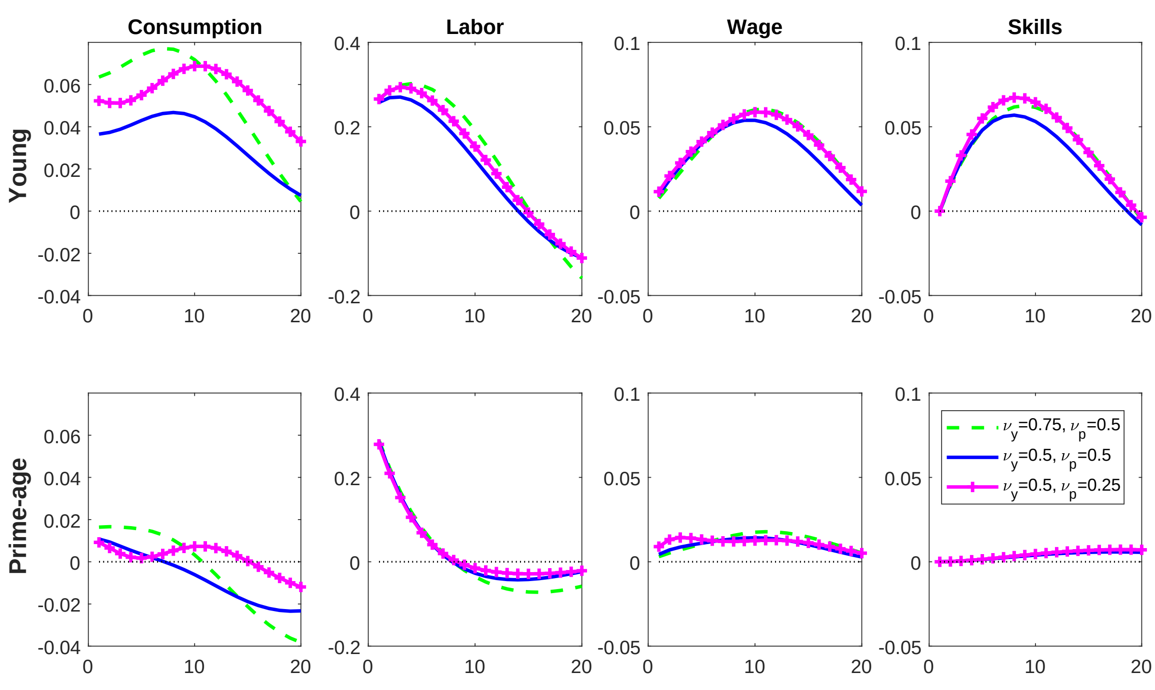

Notes: this figure shows the impulse responses of aggregate variables (Panel (a)) and disaggregate variables for young and prime-age workers (Panel (b)) to a 1% shock to government expenditure in the life-cycle model with LBD for different values of labor supply elasticity for young and prime-age workers: ν_y and ν_p respectively.

1.F Model Derivations

1.F.1 Solving the Optimization Problems of Households

In this Appendix, I derive the optimal decisions of each age group. I show that the decision rules of retirees are linear in wealth, so they can be linearly aggregated. In addition, as explained in the main text, the groups of young and prime-age individuals can be reduced to a representative young agent and a representative prime-age agent.

I solve the model using the certainty equivalence property of first-order perturbation.[48] The expectations operators are thus omitted. All the optimization problems and decisions rules are derived in real terms.

Problem of the Retiree

The optimization problem for retired agent i is given by

$$\begin{aligned} &\max_{c^i_{r,t},b^i_{r,t}} \; V^i_{r,t} = \frac{{c^i_{r,t}}^{1-\sigma}}{1-\sigma} + \beta(1-\omega_x)V^i_{r,t+1} \\ &\text{s.t.} \\ &c^i_{r,t} + b^i_{r,t} = a^i_{r,t} \\ &a^i_{r,t+1} = \frac{R_{n,t}}{\Pi_{t+1}} b^i_{r,t} \end{aligned} \tag{1.25}$$

The first-order condition with respect to consumption is given by

$$ {c^i_{r,t}}^{-\sigma} = \beta(1-\omega_x)\frac{\partial V^i_{r,t+1}}{\partial b^i_{r,t}} \tag{1.26}$$

From the envelope theorem condition we have

$$\frac{\partial V^i_{r,t+1}}{\partial b^i_{r,t}} = \frac{R_{n,t}}{\Pi_{t+1}} {c^i_{r,t+1}}^{-\sigma} \tag{1.27}$$

[48] See Schmitt-Grohé and Uribe (2004).

Combining these conditions yields the Euler equation

$$
{c_{r,t}^{i}}^{-\sigma} = \beta(1-\omega_x)\frac{R_{n,t}}{\Pi_{t+1}}{c_{r,t+1}^{i}}^{-\sigma} \tag{1.28}
$$

Next, conjecture a solution as follows, where is introduced the marginal propensity to consume $\gamma_{r,t}^{i}$

$$
c_{r,t}^{i} = \gamma_{r,t}^{i} a_{r,t}^{i} \tag{1.29}
$$

Rearranging the budget constraint we have

$$
a_{r,t+1}^{i} = \frac{R_{n,t}}{\Pi_{t+1}}(1-\gamma_{r,t}^{i})a_{r,t}^{i} \tag{1.30}
$$

Substituting $c_{r,t}^{i}$ in the Euler equation and collecting terms we get

$$
\frac{1}{\gamma_{r,t}^{i}} = 1 + (\beta(1-\omega_x))^{\frac{1}{\sigma}}\left(\frac{R_{n,t}}{\Pi_{t+1}}\right)^{\frac{1-\sigma}{\sigma}}\frac{1}{\gamma_{r,t+1}^{i}} \tag{1.31}
$$

Therefore the marginal propensity to consume is only a function of aggregate variables, thus is identical for all retired agents $\gamma_{r,t} = \gamma_{r,t}^{i} \quad \forall i$. Given the linearity of the consumption function, this implies that the aggregate consumption of retirees $c_{r,t}$ can be expressed as

$$
c_{r,t} = \gamma_{r,t} a_{r,t} \tag{1.32}
$$

where $a_{r,t}$ denotes the total wealth of retirees, which depends on the total savings of the prime-age workers who have just retired and of the retirees who are still alive.

$$
a_{r,t} = (1-\omega_x)\left(\frac{R_{n,t-1}}{\Pi_t}b_{r,t-1}\right) + \omega_r\left(\frac{R_{n,t-1}}{\Pi_t}b_{p,t-1}\right) \tag{1.33}
$$

Therefore the decision rules of retirees can be described by the following equations

$$c_{r,t} = \gamma_{r,t} a_{r,t} \tag{1.34a}$$

$$\frac{1}{\gamma_{r,t}} = 1 + (\beta(1-\omega_x))^{\frac{1}{\sigma}} \left(\frac{R_{n,t}}{\Pi_{t+1}}\right)^{\frac{1-\sigma}{\sigma}} \frac{1}{\gamma_{r,t+1}} \tag{1.34b}$$

$$a_{r,t} = (1-\omega_x)\left(\frac{R_{n,t}}{\Pi_{t+1}} b_{r,t-1}\right) + \omega_r \left(\frac{R_{n,t}}{\Pi_{t+1}} b_{p,t-1}\right) \tag{1.34c}$$

Problem of the Representative Prime-age Worker

The optimization problem of the representative prime-age worker is given by

$$\begin{aligned} \max_{c_{p,t}, b_{p,t}} \quad & V_{p,t} = \frac{{c_{p,t}}^{1-\sigma}}{1-\sigma} - \chi_j \frac{{L_{p,t}}^{1+\varphi_p}}{1+\varphi_p} + \beta\left((1-\omega_r)V_{p,t+1} + \omega_r V_{r,t+1}\right) \\ \text{s.t.} \quad & \\ & c_{p,t} + b_{p,t} = a_{p,t} + w_{p,t} L_{p,t} - \tau_t + (1-\tau_d)\mathrm{div}_{p,t} \\ & a_{p,t+1} = \frac{R_{n,t}}{\Pi_{t+1}} b_{p,t} \end{aligned} \tag{1.35}$$

The first order condition with respect to consumption is given by

$$c_{p,t}^{-\sigma} = \beta\left((1-\omega_r)\frac{\partial V_{p,t+1}}{\partial b_{p,t}} + \omega_r \frac{\partial V_{r,t+1}}{\partial b_{r,t}}\right) \tag{1.36}$$

Using the envelope theorem conditions yields the Euler equation

$$c_{p,t}^{-\sigma} = \beta \frac{R_{n,t}}{\Pi_{t+1}} \left((1-\omega_r) c_{p,t+1}^{-\sigma} + \omega_r c_{r,t+1}^{-\sigma}\right) \tag{1.37}$$

Finally, the wealth of a prime-age agent can be expressed as the total savings of the prime-age workers who do not retire and of the young agents who have just become prime-age.

$$a_{p,t} = (1-\omega_r)\left(\frac{R_{n,t-1}}{\Pi_t} b_{p,t-1}\right) + \omega_p \left(\frac{(R_{n,t-1}+\zeta)}{\Pi_t} b_{y,t-1}\right) \tag{1.38}$$

Thus the decision rules of the representative prime-age worker can be described by the following equations

$$c_{p,t} + b_{p,t} = a_{p,t} + w_{p,t}L_{p,t} - \tau_t + (1-\tau_d)\mathrm{div}_{p,t} \tag{1.39a}$$

$$c_{p,t}^{-\sigma} = \beta\frac{R_{n,t}}{\Pi_{t+1}}\left((1-\omega_r)c_{p,t+1}^{-\sigma} + \omega_r c_{r,t+1}^{-\sigma}\right) \tag{1.39b}$$

$$a_{p,t} = (1-\omega_r)\left(\frac{R_{n,t-1}}{\Pi_t}b_{p,t-1}\right) + \omega_p\left(\frac{(R_{n,t-1}+\zeta)}{\Pi_t}b_{y,t-1}\right) \tag{1.39c}$$

Problem of the Representative Young Worker

The optimization problem of the representative young worker is given by

$$\begin{aligned}
&\max_{c_{y,t},b_{y,t}} \quad V_{y,t} = \frac{{c_{y,t}}^{1-\sigma}}{1-\sigma} - \chi_y\frac{{L_{y,t}}^{1+\varphi_y}}{1+\varphi_y} + \beta\left((1-\omega_p)V_{y,t+1} + \omega_p V_{p,t+1}\right) \\
&\text{s.t.} \\
&c_{y,t} + b_{y,t} = a_{y,t} + w_{y,t}L_{y,t} - \tau_t \\
&a_{y,t+1} = \frac{(R_{n,t}+\zeta)}{\Pi_{t+1}}b_{y,t}
\end{aligned} \tag{1.40}$$

The decision rules of the representative young worker can be derived similarly to those for the prime-age worker.

$$c_{y,t} + b_{y,t} = a_{y,t} + w_{y,t}L_{y,t} - \tau_t \tag{1.41a}$$

$$c_{y,t}^{-\sigma} = \beta\frac{(R_{n,t}+\zeta)}{\Pi_{t+1}}(1-\omega_p)c_{y,t+1}^{-\sigma} + \beta\frac{R_{n,t}}{\Pi_{t+1}}\omega_p c_{p,t+1}^{-\sigma} \tag{1.41b}$$

$$a_{y,t} = (1-\omega_p)\left(\frac{(R_{n,t-1}+\zeta)}{\Pi_t}b_{y,t-1}\right) + \omega_x\left(\frac{R_{n,t-1}}{\Pi_t}b_{r,t-1}\right) \tag{1.41c}$$

1.F.2 Derivation of the Wage Phillips Curves

Each union solves the following optimization problem

$$V_t^{w_j}(W_{j,t-1}(k)) = \max_{W_{j,t}(k)} \int \left(\frac{W_{j,t}(k)}{P_t} L_{j,t}(k) - \chi_j \frac{L_{j,t}(k)^{1+\varphi_j}}{1+\varphi_j} \frac{1}{\lambda_{j,t}} \right) dk$$
$$- \int \frac{\theta_w}{2} \left(\frac{W_{j,t}(k)}{W_{j,t-1}(k)} - 1 \right)^2 \frac{W_{j,t}}{P_t} L_{j,t} dk + \beta \mathbb{E}_t V_{t+1}^{w_j}(W_{j,t}(k))$$

subject to

$$L_{j,t}(k) = \left(\frac{W_{j,t}(k)}{W_{j,t}} \right)^{-\varepsilon_w} L_{j,t} \qquad\qquad j \in \{y, p\}$$

The first order condition with respect to $W_{j,t}(k)$ gives

$$0 = (1-\varepsilon_w) \frac{1}{P_t} \left(\frac{W_{j,t}(k)}{W_{j,t}} \right)^{-\varepsilon_w} L_{j,t} + \chi_j \varepsilon_w L_{j,t}(k)^{\varphi_j} \frac{1}{\lambda_{j,t}} \left(\frac{W_{j,t}(k)}{W_{j,t}} \right)^{-\varepsilon_w - 1} \frac{L_{j,t}}{W_{j,t}} - ...$$
$$... - \theta_w \left(\frac{W_{j,t}(k)}{W_{j,t-1}(k)} - 1 \right) \frac{1}{W_{j,t-1}(k)} \frac{W_{j,t}}{P_t} L_{j,t} + \beta \mathbb{E}_t \frac{\partial V_{t+1}^{w_j}}{\partial W_{j,t}(k)} \tag{1.43}$$

From the envelope theorem

$$\frac{\partial V_{t+1}^{w_j}}{\partial W_{j,t}(k)} = \theta_w \left(\frac{W_{j,t+1}(k)}{W_{j,t}(k)} - 1 \right) \frac{W_{j,t+1}(k)}{W_{j,t}^2(k)} \frac{W_{j,t+1}}{P_{t+1}} L_{j,t+1} \tag{1.44}$$

Combining Equation (1.43) and Equation (1.44), we obtain

$$0 = (1-\varepsilon_w) \frac{1}{P_t} \left(\frac{W_{j,t}(k)}{W_{j,t}} \right)^{-\varepsilon_w} L_{j,t} + \chi_j \varepsilon_w L_{j,t}(k)^{\varphi_j} \frac{1}{\lambda_{j,t}} \left(\frac{W_{j,t}(k)}{W_{j,t}} \right)^{-\varepsilon_w - 1} \frac{L_{j,t}}{W_{j,t}} - ...$$
$$... - \theta_w \left(\frac{W_{j,t}(k)}{W_{j,t-1}(k)} - 1 \right) \frac{1}{W_{j,t-1}(k)} \frac{W_{j,t}}{P_t} L_{j,t} + \beta \mathbb{E}_t \theta_w \left(\frac{W_{j,t+1}(k)}{W_{j,t}(k)} - 1 \right) \frac{W_{j,t+1}(k)}{W_{j,t}^2(k)} \frac{W_{j,t+1}}{P_{t+1}} L_{j,t+1} \tag{1.45}$$

Using that $W_{j,t}(k) = W_{j,t}$ and $L_{j,t}(k) = L_{j,t}$, and defining the wage inflation rate $\Pi_{j,t}^w \equiv \frac{W_{j,t}}{W_{j,t-1}}$, we get

$$0 = (1-\varepsilon_w)\frac{1}{P_t}L_{j,t} + \chi_j \varepsilon_w {L_{j,t}}^{\varphi_j}\frac{1}{\lambda_{j,t}}\frac{L_{j,t}}{W_{j,t}} - \theta_w\left(\Pi^w_{j,t}-1\right)\Pi^w_{j,t}\frac{1}{P_t}L_{j,t} + ... \\ ... + \beta\mathbb{E}_t\theta_w\left(\Pi^w_{j,t+1}-1\right)\Pi^w_{j,t+1}\frac{1}{W_{j,t}}\frac{W_{j,t+1}}{P_{t+1}}L_{j,t+1} \tag{1.46}$$

Finally, after dividing by $L_{j,t}$ and multiplying by $W_{j,t}$, we get

$$(1-\varepsilon_w)\frac{W_{j,t}}{P_t} = -\varepsilon_w MRS_{j,t} + \theta_w(\Pi^w_{j,t}-1)\Pi^w_{j,t}\frac{W_{j,t}}{P_t} - \beta\mathbb{E}_t\theta_w(\Pi^w_{j,t+1}-1)\Pi^w_{j,t+1}\frac{L_{j,t+1}}{L_{j,t}}\frac{W_{j,t+1}}{P_{t+1}} \\ j \in \{y,p\} \tag{1.47}$$

Chapter 2

From He-Cession to She-Stimulus? The Impact of Fiscal Policy on Gender Gaps

2.1 Introduction

Against the backdrop of the Great Recession and the accompanying fiscal stimuli enacted in the United States and elsewhere, the debate on whether business cycles and fiscal policy affect men and women unevenly has been revived. It is well documented that business cycle fluctuations affect male workers more than female workers.[1] Men incurred 78% of the net job losses during the Great Recession, and male-to-female job losses during the five recessions between 1969 and 1991 were similar in magnitude.[2] In particular, young, less educated men as well as those working in blue-collar occupations are particularly strongly affected.[3] However, much less is known about the effects of fiscal policy on labour market outcomes of men and women, a gap which this paper attempts to fill.

[1]This can be partly ascribed to men's employment in more cyclical industries such as manufacturing and construction. See for instance Clark and Summers (1981), Solon et al. (1994), Hoynes et al. (2012).

[2]See Wall and Engemann (2009).

[3]Table 2.A.1 summarizes unemployment increases for men, women and vulnerable male subgroups during recessions that occurred since the start of our sample period. For a more detailed analysis see Bredemeier et al. (2017) and Section 2.E.

The aim of this paper is to gain a deeper understanding of the heterogeneous effects of government expenditure from a gender perspective. Our study provides policy-making insights on how fiscal policy can contribute to offsetting inequitable business cycle effects and/or to closing gender gaps, as well as on the trade-offs involved. Using micro-level data for the U.S. from the Current Population Survey (CPS), we find that men are hurt or benefit less than women from increases in the major government spending components. This result is largely driven by negative spillovers for men working in the private sector. Our study further sheds light on various policy trade-offs. In particular, fiscal expansions cannot reconcile both policy goals: offsetting inequitable business cycle effects and closing gender gaps.

To measure the effects of fiscal policy shocks on gender gaps in the labour market, we estimate several vector-autoregressive models using Bayesian estimation techniques. Following Mountford and Uhlig (2009), the fiscal shocks are identified using an agnostic sign restriction approach. The main advantage of this identification strategy is that it allows us to eliminate the confounding influence of other macroeconomic shocks, namely business cycle, monetary policy and tax revenue shocks. We then examine the impulse response functions of gender gaps in wages, hours and employment rate to fiscal policy shocks. Our study encompasses the analysis of three dimensions of heterogeneity. First, we investigate whether the effects of fiscal policy shocks vary depending on the type of public expenditure. Second, we explore how the effects vary across population groups with different characteristics such as age, education, industry and occupation. Lastly, we inspect labour market outcomes separately for men and women to understand the drivers of variations in gaps.

Our main findings can be summarized as follows. First, total government spending raises female wages and employment relatively more and hence, gender gaps narrow. However, the composition of fiscal shocks matters. Expenditure on consumption, the government wage bill and social benefits boost women's wages exclusively. In contrast, investment spending is detrimental to both genders' hourly income. As regards employment, men either suffer more or benefit less than women from fiscal expansions across all spending categories, except investment and defense. Second, splitting the working population into subgroups reveals substantial heterogeneity. In particular, we find that negative effects on male wages and employment for the total population are largely driven by men's losses in the private sector. These stand in contrast to the benefits experienced by men in the public sector and by women in both sectors. The most likely explanation is that private sector men sort into occupations which are not affected by government spending. Third, our analysis highlights that if fiscal policy can help offsetting inequitable business cycle effects, this is not compatible with promoting gender equality. Those spending components that best close gender gaps (consumption spending, the wage bill and benefits) are most harmful for young, uneducated and private-sector men. Similarly, investment spending, which fosters employment of these crisis-hit men, increases overall gender gaps.

These insights are valuable for macroeconomists and policy makers for three reasons. First, our results help to gauge to what extent government expenditure is able to "assist those most impacted by the recession" which was the explicit purpose of the American Recovery and Reinvestment Act of 2009.[4] Second, our analysis is insightful for policy makers whose goal is to promote female labor force participation and gender equality, independently of the cycle.[5] Conversely, this paper highlights the potential damaging effect that cutting government expenditure may have by widening persistent gender gaps. Hence, we underline the gender non-neutrality of budgetary decisions and substantiate the importance of implementing "gender budgeting" as suggested by the IMF (2017) and the European Parliament (2015). Third, our analysis hints at the importance of encouraging women's labour force participation as this may increase the effectiveness of fiscal policy as an aggregate stabilization tool.[6]

The remainder of this paper is structured as follows. Section 2.2 outlines the related literature and Section 2.3 provides some theoretical background. Sections 2.4 and 2.5 describe the data and the econometric approach, respectively. Results are presented in Section 2.6, followed by a discussion on the drivers of heterogeneous effects of fiscal shocks on labour market outcomes across gender, demographic groups and fiscal components in Section 2.7. Robustness checks and extensions are described in Section 2.8. Section 2.9 discusses some limitations of our study and offers directions for future research. Section 2.10 concludes. The Appendices contain some stylized facts about the

[4]This stimulus package worth $787 billion consisted of a mix of tax credits, spending on social welfare, consumption spending (mainly on education and healthcare) as well as investments in infrastructure and the energy sector.

[5]Our results suggest that the impact of fiscal policy on gender gaps can be quite persistent. In addition, government expenditure shocks show a high degree of persistence. We find estimates of the autocorrelation coefficients of the cyclical component for government spending instruments which are larger than 0.9 and highly statistically significant.

[6]Our results should only be seen as suggestive and not conclusive causal evidence since we do not have an appropriate contrapositive, i.e. we do not observe male labour market outcomes after fiscal shocks in the absence of women. Furthermore, the usefulness of fiscal policy may rely on temporarily incentivizing women to join the labour force and a permanently higher female participation may in fact diminish this "added worker" effect.

components of government expenditure, and the gender compositions across occupations and sectors. Further econometric results using a state-level analysis are presented, as well as a description of the data and of the algorithm used for the estimation of the impulse response functions. Results of the impact of a business cycle shock on gender gaps are also reported and analyzed.

2.2 Related Literature

Despite a growing interest in the evolution and the determinants of gender gaps in the labour market,[7] the literature on the impact of macroeconomic policy on gender equality is scarce. Braunstein and Heintz (2008) and Seguino and Heintz (2012) show that contractionary monetary policy may widen gender gaps in employment. Regarding the role of fiscal policy, Giavazzi and McMahon (2012) study the effects of shifts in military spending across U.S. states and find that hours worked and wages increase more for households with female heads. Bredemeier et al. (2017) document that fiscal expansions stimulate primarily female employment in the U.S. Akitoby et al. (2019) find that increases in government expenditure benefit female employment during recessions in G-7 countries. Our findings are broadly in line with these results. We contribute to this literature by exploring the effects of various subcomponents of government expenditure. We argue that who benefits from fiscal stimuli depends on the type of expenditure under consideration. We also analyse labour market outcomes separately for men and women to better understand the variations in gaps. Furthermore, our identification strategy is able to better isolate the variations in fiscal policy variables from automatic responses to other macroeconomic shocks.

Several studies have emphasized the crucial role of industry composition in shaping gender differences in labor market outcomes, including Hoynes et al. (2012), Olivetti and Petrongolo (2014) and Bredemeier et al. (2017). We explore the role of other workers' characteristics, such as age and education. Furthermore, our analysis highlights the importance of occupational sorting of men and women, in particular across public and private sectors, as a driver of heterogeneous effects of fiscal policy across genders.

Finally, this paper relates to a strand of the literature which reports heterogeneous effects of fiscal policy across industries, notably Nekarda and Ramey (2011) and Bredemeier et al. (2019),

[7] See for instance Blau and Kahn (2000), Blau and Kahn (2017), Ngai and Petrongolo (2017), Albanesi and Şahin (2018).

and across households with different characteristics, in particular Giavazzi and McMahon (2012), Misra and Surico (2014), Cloyne and Surico (2016), Anderson et al. (2016), Ferriere and Navarro (2018). We complement these studies by further documenting heterogeneity across genders and across subcomponents of government expenditure. In addition, we provide evidence of a strong public-private sector divide, as negative effects of fiscal policy on male wages and employment for the total population appear to be largely driven by negative spillovers for men working in the private sector. This finding is closely related to Kim (2018) who documents that shocks to government expenditure increase employment in the government sector at the expense of private sector jobs in the U.S.

2.3 Theoretical Background and Conjecture

Fiscal policy shocks could impact female and male labour market outcomes unequally since they face different demand and supply curves. The former is driven by sectoral and occupational sorting. Women are over-represented in services, sales, and office jobs (so-called pink-collar occupations, see Figure 2.A.6) as well as in the public sector. Men on the other hand cluster in manufacturing, installation and construction (“blue-collar" occupations) and dominate the private sector. The magnitude of their labour demand curve shifts after a fiscal expansion will depend on which type of government spending is boosted.

Investment or defense spending are targeted at blue-collar occupations. Thus, we expect the demand for male workers to increase disproportionately which should widen male-to-female employment and wage gaps. Furthermore, the resulting boost in household income may encourage some women to stay at home which would further increase the employment gap.

In contrast, expanding social benefits, which mainly consist of Medicare and Medicaid expenses, may boost demand for healthcare professionals who are predominantly female. Similarly, government consumption expenditure is expected to increase demand for pink-collar workers and the female-dominated public sector disproportionately. Hence, employment and wage gaps should narrow. The fact that women face a higher public sector wage premium (Figures 2.A.7 and 2.A.8) will reinforce this effect, i.e. increasing the public sector headcount through consumption expenditure will mechanically close the wage gap in the overall population. Furthermore, increasing demand for public sector employees may induce some private sector workers to relocate. Since women are more mobile across industries and occupations,[8] they may be the main beneficiaries of higher wages and expanded employment opportunities. Moreover, women taking up government jobs may hire (usually female) nannies and nurses to take care of children and elderly dependents which may induce second round employment effects.

As mentioned above, heterogeneous labour market effects may also be driven by differences in labour supply. Women remain the main caregivers for children and elderly dependents and are therefore more likely to stay out of the labour force or work part-time. Hence, as a result of their specialization in home production, female labour supply is more elastic (Bredemeier (2015)). Consequently, even if demand shocks were of a similar size for both genders, female employment may respond more. On the other hand, even a strong labour demand shift may have only modest effects on female wages. This may be further exacerbated since government consumption or social benefit expansions may cause a selection bias. Less productive women may be encouraged to join the labour force which may exert downward pressure on average female wages[9]. The following analysis will test the validity of our conjectures.

2.4 Data

Figure 2.A.1 schematically illustrates the composition of total government expenditure in the US, and Figure 2.A.2 shows the historical evolution of each component between 1979 and 2017.[10] While real government spending per capita has more than quadrupled since the start of our sample, the relative shares of its subcomponents remained largely unchanged. The exception are social benefits which have grown from 30% of total expenditure in 1979 to 38% in 2007 making it the largest component in the last year of our baseline sample. The second largest spending category in 2007 was consumption (net of defense) (37%) out of which 72% was allocated to the government wage bill, i.e. to paying public sector employees. Furthermore, defense and investment spending accounted for 15% and 10% of total spending, respectively.[11]

[8] See, e.g., Shin (1999).

[9] The lower wages of women who suddenly start participating in the labour force may partly be the result of women's lower labour market attachment and and more frequent employment breaks.

[10] Note that we exclude transfer payments to the rest of the world, interest payments, capital transfer payments, subsidies, social insurance funds and net purchases of nonproduced assets.

[11] Until 2017 the share of social benefits grew further reaching 46% of total government spending. Consumption (net of defense) fell slightly and amounted to 35% in 2017 but the share of the government

To investigate the impact of these different fiscal instruments on gender gaps, we use micro-level data on several labor market variables from the CEPR extracts of the Current Population Survey (CPS) Merged Outgoing Rotation Groups. The CPS is the source of official U.S. government statistics on employment, wages and unemployment, with interviewed households selected to be representative of the U.S. population. The survey records detailed information on hours worked, earnings, industry, occupation, education and unionization, as well as on demographic characteristics. About 60,000 households are interviewed for four consecutive months in one year, then ignored for eight months, and interviewed again for four consecutive months.

We build quarterly series for weekly hours worked, real hourly wages and employment rates for each gender (quarterly average of monthly observations) for full-time workers aged 16-64, i.e. who have worked at least 35 hours a week.[12] Self-employed workers are excluded from the sample. All variables are seasonally adjusted by X-12 ARIMA. On average during our sample period, full-time female workers earned 27.6% less per hour, they worked 2.5 hours less per week and their employment rate was 18.9% lower than men's (see Table 2.A.2).

A fiscal instrument may affect labor market outcomes differently across industries, occupations and demographic groups, as pointed out by Nekarda and Ramey (2011), Chodorow-Reich et al. (2012), Bredemeier et al. (2019) among others. To address this issue and explore heterogeneities among groups of workers, we build pseudo-panels by aggregating individual observations into pseudo-cohorts of workers with different characteristics and computing averages for each period.[13]

We consider several sample splits: (i) by education (whether workers have obtained a college degree or not), (ii) by age (10-year age groups), (iii) pink-collar occupations (mainly service, sales, and office) vs. blue-collar occupations (mainly production, construction, transport, and installation),[14]

wage bill remained the same. Furthermore, defense and investment spending fell to 12% and 7% of total spending, respectively.

[12] In Section 2.8, we also conduct the analysis for non-married individuals (to eliminate partner effects) and for part-time workers.

[13] This approach is described in the seminal paper by Deaton (1985).

[14] Figure 2.A.6 illustrates the female shares across occupations. To build occupational employment groups, we use the conversion factors from the U.S. Census Bureau as the occupation and industry codes in the CPS were subject to several revisions. We classify as pink-collar occupations those with a female share of more than 50% for the whole sample period (1979-2016). These are services, sales and related occupations, office and administrative support occupations. Consequently, blue-collar occupations have a female share of less than 50% and include farming, fishing, and forestry occupations, construction and extraction occupations, installation, maintenance, and repair occupations, production occupations, transportation and material moving occupations. We do not classify management, business, and financial occupations since the female

(iu) public vs. private sector employees.

We use data on fiscal variables, GDP and inflation from the U.S. Bureau of Economic Analysis, on civilian population from the U.S. Bureau of Labor Statistics, and on the federal funds rate from FRED. Details of sources and definitions of the data are provided in the Appendix Section 2.B.

2.5 Econometric Approach

2.5.1 VAR model

To measure the effects of fiscal policy shocks on gender gaps in the labour market, we estimate several structural VAR models with up to ten endogenous variables.

In our baseline specification, the vector of endogenous variables includes the log of real per capita total government expenditure, the log of real per capita net (of transfers) tax revenue, the log of real per capita GDP, the gender wage gap, a supplementary labor market gap variable, inflation and the federal funds rate. The supplementary labor market gap variable alternates between the gender gap in (i) employment rates, (ii) hours. The gender wage gap is measured as the difference between the log of real male wage and the log of real female wage. The gender hours gap is computed analogously. The gender gap in employment rates is defined as the difference between male and female rates. To control for fiscal foresight, we include eight lags of an exogenous war dummy following Ramey (2011). The VAR models are estimated with two lags, on quarterly data from 1979Q1 to 2007Q4.[15] Following Mountford and Uhlig (2009), we neither include a constant nor time trend.[16]

In a second step, we further explore the effects of different types of government expenditures on gender gaps in labor market outcomes. For that purpose, we replace total government expenditures by the fiscal instrument of interest in the vector of endogenous variables. Specifically, in order to avoid including too many variables in the VAR, we rotate the subcomponent of total expenditures, which

share ranges between 25% and 55% during the sample period. Similarly we do not classify professional and related occupations since the female share is not strictly larger than 50% in the first few years of our sample. To check robustness, we also conducted the analysis for workers split according to industry, more precisely female-dominant vs. male dominant industries. Results are qualitatively similar.

[15] The starting date of the sample is constrained by the availability of CPS micro-data, while the end date is chosen to avoid nonlinearities due to the zero lower bound period.

[16] We checked that the results are qualitatively robust when a constant and a time trend are included.

alternates between (i) government consumption expenditure, (ii) government investment expenditure, (iii) expenditure on social benefits, (iv) defense spending and (v) expenditure on the government wage bill, which is the largest constituent of consumption expenditure.[17] Furthermore, we include total expenditure net of the respective fiscal variable of interest as an anchor.

Lastly, we repeat the above analysis including male and female series of real wages (in logs) and of the supplementary labor market variable instead of gender gaps. Estimating the effects of fiscal shocks on male and female labor market outcomes separately allow us to better understand the mechanism behind changes in gender gaps and to draw finer policy conclusions.

2.5.2 Identification

Following Mountford and Uhlig (2009), Pappa (2009), Arias et al. (2014), Bermperoglou et al. (2017) among others, we identify the fiscal shocks using an agnostic sign restriction approach that sets a minimum number of restrictions on impulse responses, while controlling for other macroeconomic shocks. These identifying sign restrictions are summarized in Table 1.

The shocks are identified sequentially, as in Mountford and Uhlig (2009), Arias et al. (2014), and Bermperoglou et al. (2017). First, we identify a generic business cycle shock which leads to a positive comovement between output and government net tax revenue for four quarters. Second, we follow Bermperoglou et al. (2017) and identify a monetary policy shock by combining zero and sign restrictions. In particular, the federal funds rate should react positively and contemporaneously to output and inflation deviations only to approximate the Taylor rule. We also impose orthogonality between the monetary policy shock and the business cycle shock. Third, the government revenue shock is identified as a shock that raises net tax revenues for four quarters and that is orthogonal to the monetary and business cycle shocks. Fourth, we identify a shock to total expenditure net of the spending component of interest (G_{net}) as an increase in this variable for four quarters while imposing orthogonality to all previously identified shocks. Lastly, to identify the fiscal shocks of interest, we impose that they increase the corresponding fiscal variable for four quarters while being orthogonal to the shocks identified earlier. Orthogonality to the G_{net} shock ensures that our results are exclusively driven by the fiscal instrument of interest and not by any other expenditure component.

Following Uhlig (2005), we estimate the model using a Bayesian approach with flat priors for

[17]Investment and consumption expenditures are net of defense spending.

model coefficients and the covariance matrix of shocks (see Section 2.C). The estimations are based on 400 draws from the posterior distribution of VAR parameters and 4000 draws of orthonormal matrices. We compute the median and the 68% confidence bands of impulse responses to a shock that raises the government expenditure component of interest by 1% on impact. The analysis is conducted at the aggregate level as well as for different population groups.

Table 1: Identifying sign restrictions

Restricted variables	**Shocks**				
	ε_t^G	$\varepsilon_t^{G_{net}}$	ε_t^T	ε_t^{MP}	ε_t^{BC}
Output					+
Inflation rate					
Interest rate					
Government revenue			+		+
Government expenditure component	+				
Net government expenditure		+			

Notes: This table reports the sign restrictions on impulse responses for each identified shock. ε_t^G denotes a shock to the government expenditure component of interest, $\varepsilon_t^{G_{net}}$ a shock to net government expenditures (i.e. total expenditures net of the expenditure component of interest), ε_t^T a government revenue shock, ε_t^{MP} a monetary policy shock, and ε_t^{BC} a business cycle shock. To identify the monetary policy shock we do not impose restrictions on the IRFs but on the structural impact matrix. All restrictions apply for periods 0-3 after the shock occurred.

2.6 Results

This section reports our results for different government spending components.[18] Section 2.6.1 looks at gender gaps among all full-time workers aged 16-64, and Section 2.6.2 analyses heterogeneities across subgroups, obtained by splitting the population by age, education, occupation and sector. The purpose of this exercise is threefold. First, it helps us to gain insights into the mechanisms that drive asymmetric effects across men and women. Second, it allows for a better assessment of how to use fiscal policy to offset inequitable business cycle effects. Third, the analysis highlights trade-offs involved when attempting to close gender gaps since not all men and women react equally to fiscal stimuli. Section 2.6.3 summarizes the main takeaways. Overall, we find that gender gaps

[18]In the Appendix Section 2.E, we present responses to a business cycle shock.

close following shocks to government consumption, the government wage bill and social benefits. However, these spending components amplify the particularly adverse effects experienced by private sector men during recessions.

2.6.1 Effects of fiscal policy for the total population

An increase in total government spending decreases the gender wage gap if all workers are considered (Figure 2.B.1, row 1 & column 1). This result is consistent with Bredemeier et al. (2017) and is mainly driven by a decrease in male wages (Figure 2.B.2, row 1 & column 2). Moreover, a fiscal expansion closes the gap in male-to-female employment due to an increase in both variables for women. The gap in hours worked decreases but not significantly since both genders work less hours.[19]

A shock to government consumption expenditure leads to similar movements in gender differentials (Figure 2.B.1, row 2). However, the reduction in the wage gap is driven by a marked increase in female wages rather than a fall in male wages. Another difference is that employment fall for both genders, slightly more so for men.

Figure 2.B.1 (row 3) shows that these results are strongly driven by the government wage bill, which is the largest sub-component of consumption spending. Increasing the funds available for public sector employees, either through raising public sector wages and/or through enlarging the public sector headcount, benefits women in terms of both wages and employment.[20] In contrast, men's wages drop significantly after two quarters.

Investment spending (Figure 2.B.1, row 4) is the only component that significantly widens gaps in wages and other labour market variables except hours worked. However, both genders react similarly to an investment shock, i.e. while wages and hours drop, LFP and employment increase for approximately one year and decrease thereafter. Increases in gaps are driven by slightly stronger movements for men in each of these variables.

Furthermore, we find that expenditure on benefits has a strong decreasing effect on the wage gap

[19]Since we are considering only full-time employees in our baseline analysis, changes in hours worked after a fiscal shock are small in magnitude.

[20]However, female employment increases are not robust when shocks are identified using a Cholesky ordering.

(Figure 2.B.1, row 5) since female hourly earnings increase and male earnings decrease. Employment rates of both sexes drop, slightly more so for men, which leads to a decrease in LFP and employment gaps.

Defense spending (Figure 2.B.1, row 6) stimulates both genders' wages but the result is not significant when using a Cholesky decomposition instead of sign restrictions to identify shocks. LFP and employment are boosted to a similar extent which results in no significant changes in the associated gaps. Hours worked, however, fall slightly more for men, leading to a decline in the hours gap.

To summarize, government spending has heterogeneous effects across genders (see Table 2.A.3 for an overview).[21] In particular, expenditure on consumption, the wage bill and social benefits boost women's wages exclusively leading to a drop in the gender wage gap. In contrast, shocks to these spending components tend to move hours worked, LFP and employment of men and women in the same direction. Yet, gaps in these labour market variables fall since men are less favourably affected.[22] In the next section, we further split male and female workers and analyse responses by demographic groups which will allow us to interpret the above results.

2.6.2 Effects of fiscal policy across population subgroups

In this section, we analyse labour market responses to fiscal shocks by subgroups for both genders. For each spending component, we highlight which groups drive the results for the total population. From our analysis in Section 2.E of the Appendix as well as from Wall and Engemann (2009), we know that especially men that are younger, less educated and those that work in blue-collar occupations and/or in the private sector are adversely affected by downturns. Our analysis reveals that the same groups of men are hurt by most fiscal shocks.

The decrease in male wages that is caused by a shock to **total government expenditure** is mainly driven by falling wages of low-educated and younger men. The negative effect is strongest and most persistent for the age groups 16-24 and 25-34 and fades away as age increases. The top

[21] We also investigated potential reverse causality between gender gaps and fiscal policy. We found clear-cut evidence of a Granger-causality from government expenditures to gender gaps, but little evidence of causality in the opposite direction. Results are available upon request.

[22] The robustness of these results is corroborated by using US state-level data in the Appendix Section 2.D.

right-hand plots of Figures 2.B.3 and 2.B.4 depict these results.[23] Hence, spending hikes squeeze the wages of the same individuals that are hurt during recessions. However, the drop in wages seems to shelter less educated men from the downward adjustment in LFP and employment which is experienced by college-educated men. The compositional shift towards a less educated male workforce, may have contributed to the decline in average male wages. Furthermore, government spending seems to crowd out demand for men in the private sector: Their employment and wages drop while male public employees experience a significant increase in both variables (see Figure 2.B.5). Employment responses may be due to private sector men leaving the labour force or moving to the public sector.[24] Within the private sector, women are hurt less than men in terms of wages and their employment rate even increases on impact, although not significantly. Interestingly, men's wages decline in both pink and blue collar occupations whereas women in both groups benefit (Figure 2.B.6).[25] As Section 2.7 will discuss, heterogeneous responses across genders within the private sector and pink-collar occupations are driven by occupational sorting.

As highlighted in Section 2.6.1, boosting **spending on consumption** significantly increases female wages but leaves male wages unchanged. Women across all demographic groups benefit from higher wages (Figures 2.B.8 to 2.B.11). However, the effect is most pronounced and persistent for highly educated, middle aged (45-54) women as well as those that are employed in pink-collar occupations and the public sector. Young, less educated and private-sector women are the groups that drive the drop in employment observed for the total population. These subgroups also experience particularly large wage and employment losses among men, whereas wages of public sector men increase. In contrast, women's wages rise in both sectors and their private-sector employment decline is weaker, leading to a significant drop in the private sector employment gap.

A shock to the **government wage bill** mirrors most of the above results - which is not surprising given that it accounts for 64% of consumption spending on average. The drop in male private sector wages is now even more pronounced (see Figure 2.B.15). Once more, wages of less educated and younger men (aged 35 and under) decline and this adjustment seems to shelter them from employment losses experienced by more educated and older men (Figures 2.B.13 and 2.B.14). Hence, part of the decline in overall male wages may again be explained through negative selection, i.e. by

[23] LFP mirrors the responses for employment and hours show little variation since we only look at full-time workers. Thus, we exclude these two variables from our plots.

[24] Since the CPS is not a panel dataset, we cannot track movements between sectors.

[25] This result is robust to splitting the population into those that work in female vs. male dominated industries.

employment exits of educated and older men which tend to have higher earnings. In contrast, all women except those in blue-collar occupations (Figure 2.B.16) see their wages increase after a boost in the government wage bill. Unlike after a shock to consumption, female employment in the private sector now increases significantly making the fall in the private sector employment gap more pronounced (Figure 2.B.15). Interestingly, the employment boost is mostly experienced by women of prime child-bearing age (25-34) which maybe explained by the creation of public sector jobs that provide high levels of protection and better provisions for maternity leave.

Turning to **investment spending**, both men and women experience wage declines, especially in the private sector, but their LFP and employment increase (Figure 2.B.20). For men wage declines seem to be partly driven by negative selection: employment is boosted especially among men that are younger (aged 16-24), have less education and work in blue-collar occupations (Figures 2.B.18, 2.B.19 and 2.B.21). In contrast, employment and wages of more educated men fall. Unlike the previous fiscal components, investment spending initially increases male employment in the private sector but the effect becomes significantly negative after two quarters. Women's employment is expanded across all population groups with the exception of women above 45 whose responses are not significant.

Spending on **social benefits** seem to push both men and women to leave their job possibly through negative incentive effects. The response is especially strong for young workers (under 35), for those in the private sector as well as for blue-collar men (Figures 2.B.24 to 2.B.26). Older workers and public sector employees are initially sheltered from an employment decline and for public sector women the rate even increases significantly during the second quarter. Furthermore, the increase in women's wages for the total population is driven by female wage gains in pink collar occupations and the public sector.[26] Furthermore, significant wage increases only incur to women that attended college (Figure 2.B.23) and to females aged between 35 and 54. Male wages decrease across all age, education and occupational group. Losses are particularly severe in the private sector and among older men (aged 55 to 64). In order to interpret these results, we investigate which purposes social benefits are aimed at in the next section.

Wages for the total population modestly increase after a **defense spending** shock but the result is not significant when using a Cholesky decomposition instead of sign restrictions. However, under both identification schemes employment boosts affect mainly less educated men, blue-collar males as well as men working in the public sector (Figures 2.B.28, 2.B.30 and 2.B.31). In contrast, the

[26] However, the increase in female wages is not robust to identifying shocks by Cholesky decomposition.

employment rate of highly educated and older men decreases (Figure 2.B.29). Female employment increases independently of education or sectoral affiliation. Overall, the employment and wage effects after a defense spending shock are modest.

2.6.3 Main takeaways

From the above analysis we gain three important insights. First, there is no fiscal panacea that is able to improve the labour market outcomes of all subgroups adversely affected by downturns. Table 2.A.4 summarizes our results for the particularly crisis-hit male subgroups.[27] Men who are young, less educated and/or work in the private sector experience either no effect or see their labor market outcomes worsen after shocks to total spending, government consumption, the government wage bill and social benefits. Investment spending poses a trade-off of employment gains against wage losses for vulnerable male subgroups. From their perspective, defense spending hikes are most favourable.

Second, fiscal expansions are unable to reconcile both policy goals: offsetting inequitable business cycle effects and closing gender gaps. Those spending components that best close gender gaps (consumption spending, the wage bill and benefits) are most harmful for young, uneducated and private-sector men. Conversely, investment spending which fosters employment of crisis-hit men, increases overall gender gaps. Defense spending shocks do not significantly increase gender gaps and the positive effect on male labour market outcomes are modest.

Third, negative effects on male wages and employment for the total population are largely driven by men's losses in the private sector (see Table 2.A.6).[28] In contrast to men, women's improvements in labor market outcomes are not limited to the public sector. Hence, gender gaps in the private sector tend to narrow after a fiscal expansion.

The next section discusses potential underlying mechanisms driving the heterogeneity observed across genders and population subgroups for different fiscal components.

[27]For comparison, Table 2.A.5 illustrates the results obtained for women.

[28]This is in line with Kim (2018) who also finds that spending shocks increase employment in the government sector at the expense of private sector jobs. However, the paper does not distinguish between genders or spending components.

2.7 Discussion

In this section we analyze why the effect of fiscal shocks on labour market outcomes differs by gender, subgroup and fiscal component. Our results are mainly driven by spending components being targeted at specific occupations and sectors which differ in their gender composition. As a result, men and women face different shifts in labour demand. Understanding the mechanisms that drive heterogeneous responses allows us to derive more nuanced policy implications.

2.7.1 The public-private sector divide

The above analysis suggests that, independently of gender, public sector employees benefit more from spending shocks than those in the private sector. Narrowing gender gaps in the overall population is then a partly mechanical outcome that is driven by two factors. On the one hand, it is due to a higher public sector wage premium for women than for men and hence a lower GWG among government employees (see Figures 2.A.7 and 2.A.8).[29] On the other hand, it is driven by women's over-representation in the public compared to the private sector with average female shares of 44% and 35%, respectively, during our sample period.[30] Thus, a demand driven increase in the public sector headcount or wages following a fiscal shock automatically closes gender gaps. This result is reinforced by private sector men experiencing wage and employment losses while women in both the public and the private sector benefit from most government spending components (see Table 2.A.6). Since private sector men and women are similar in terms of education and age (see Figure 2.A.9), the results must be driven by private sector men sorting into occupations that are not affected by government spending on consumption or benefits.

To illustrate this, Figure 2.A.3 decomposes spending components into their functions. Education and general public service account for more than half of government consumption spending.[31] Since the majority of public sector women work in education and administration (see Figure 2.A.12) it is not surprising that they are beneficiaries of shocks to government consumption. Similarly, private

[29] In 2015, full-time median earnings of women amounted to only 75.8% of male earnings in the private sector but 84% in state and local government jobs and 88.5% at the federal government level.

[30] The stronger appeal of the public sector for women is likely due to high levels of protection, time flexibility during motherhood (Kolberg (1991)) and the public sector wage premium.

[31] General public service includes tax collection and financial management, interest payments, executive and legislative spending.

sector women benefit from fiscal expansions as more than a quarter of them work in administrative jobs. Hence, demand for their labour may be boosted through public outsourcing or because of a shortage of administrative female staff that may have moved to the public sector. Both cases may occur after government wage bill or defense shocks, which is reflected in an increase in female wages and employment in the private sector.[32] Additionally, private sector women may encounter positive labour demand spillovers: women taking up public sector jobs may hire caregivers for their children and elderly dependants. In fact, almost a quarter of private sector women work in healthcare, personal care and education. This may also explain why they are sheltered from wage drops after a shock to social benefits, almost all of which is targeted at health and income security.[33]

Public sector men also see their wages and employment increase after a fiscal shock because they cluster in administrative and education occupations as well as in protective services. The latter group should positively respond to spending on “public order and safety" which includes the police, fire fighters, lawyers and prison officers. However, unlike women, private sector males do not mirror the occupational structure of their public sector counterparts. Almost half of them work in either construction, installation, production or transportation occupations which are not targeted by the main functional components of consumption or social benefits spending. The only exception is expenditure on economic affairs which is mainly dedicated to transportation. Indeed, this explains why a shock to investment spending, of which 42% is channelled towards economic affairs, initially increases employment of private sector men.

2.7.2 Differences across blue- and pink-collar occupation groups

Occupational sorting of genders may also explain heterogeneous responses across genders in blue- and pink-collar occupations. More than 60% of pink-collar women work in either education, healthcare, personal care or administration. They benefit particularly from higher wages after

[32]The fact that women's private-sector employment falls after a shock to government consumption is in accordance with Bermperoglou et al. (2017). The non-wage bill component of consumption spending, which is dedicated to the purchase of intermediate goods and services, may not be complementary to private consumption. As a result, private-public wage spillovers reduce private labour demand without being overrun by complementarity effects. The latter force may be stronger after government wage bill shocks since spending on education boosts private consumption and thus employment.

[33]The latter includes old age, survivors and disability insurance payments which could boost demand for predominantly female personal care workers. According to Figure 2.A.5, medical care and social security are not only the largest but also most volatile subcomponents of social benefits and therefore drive observed shocks.

shocks to consumption, the government wage bill and benefits, as explained above. This result combined with the fact that women are overrepresented in pink-collar jobs rationalizes aggregate female responses. Pink-collar men, who mainly work in sales jobs (22.1%), are less favourably affected since they are not directly targeted by these shocks. Similarly, shocks to investment and defense spending significantly stimulate demand for blue-collar men since 97% of them work in either construction, production, installation and transportation. Hence, the group of private sector men overlaps with blue-collar men in terms of occupations. Consequently, responses to fiscal shocks are similar.

2.7.3 Differences across age and education groups

The heterogeneous responses observed across age and education groups are closely related to the demographic structures within the private and the public sector. The fact that young and less educated men are often worse off reflects that private sector men rarely benefit (see Table 2.A.4). In fact, men's age and education averages are lower in the private than in the public sector. As shown in Figures 2.A.9 and 2.A.10, on average during our sample period only 21.3% of men in the private sector hold a college degree compared to 41.9% in the public sector. Moreover, the age average for men is 38 in the private sector and 41 in the public sector.[34] Hence, men that loose out most are likely to belong to all three groups: the private sector, the less educated and the younger group.[35] Overlapping characteristics are also observed for women. Those in the public sector tend to be older and more educated than in the private sector. As a result, after a shock to consumption spending, the government wage bill and benefits, only college-educated women and those aged 45+ experience persistent wage gains.

Our analysis unveiled another striking difference across education groups. While college-educated men experience drops in employment but not in wages, the opposite is true for less educated men. A possible explanation is that the former group has higher reservation wages due to more generous

[34]Similarly, only 6.4% of blue collar men blue-collar compared to 45.4% of pink-collar men hold a college degree. Consequently, beneficiaries of an investment shock and those that loose after wage bill and benefit shocks are likely to belong to both groups, blue-collar occupations and the less educated.

[35]We cannot determine the direction of causality with certainty. In other words, are young and uneducated men more vulnerable to fiscal shocks and as a result we observe losses in the private sector? Or, alternatively, are private sector male jobs particularly crowded out and consequently we observe losses among the young and less educated? However, our analysis in the previous section provides evidence in favour of the second channel.

unemployment benefits and shorter expected unemployment spells. As a result, they may prefer to (temporarily) leave their job as they see the demand for their work decrease.[36]

2.7.4 Policy trade-offs and further insights

From the above analysis, it becomes clear that policymakers attempting to find the right fiscal mix face various trade-offs. First, closing gender gaps carries the cost of worsening labour market outcomes of crisis-stricken men.[37] Furthermore, narrowing gaps do not necessarily imply that women benefit- they may simply be less worse off such as after a shock to consumption spending when both genders' employment decreases.

Second, there are trade-offs across population subgroups. One of them lies in uneven effects on private compared to public sector men that we discussed above. Another trade-off exists between education groups. Boosting the government wage bill closes employment gaps among the highly educated but widens them among the less educated. Similarly, while defense spending is able to modestly increase employment and wages of particularly vulnerable groups, this comes at the expense of college-educated men's employment.

Third, there is a trade-off between boosting employment or wages which seems to be partly caused by selection. For example, by increasing demand for public sector employees, shocks to government consumption and benefits select older and supposedly more experienced women into the workforce. Instead, younger women are pushed out of the labour market. As a result, female wages are boosted but women's overall employment and LFP fall. Moreover, after an investment spending shock college-educated men are substituted by less educated men which lowers the average wage but initially increases employment.

All in all, the adequacy of using a particular fiscal instrument depends on the policymaker's objectives which need to be carefully weighed taking the above trade-offs into consideration. Our analysis highlights that the effects of government expenditure crucially depend on the functions

[36] This result seems puzzling at first since less educated labour supply is thought be more elastic. However, at the extensive margin, Prasad (2003) finds that college-educated workers in Germany have reservation wages that are 15% higher than for those with only a high school degree. Brown and Taylor (2009) confirm this result qualitatively for the UK.

[37] The ideal case would be if both women and young, uneducated, private-sector men benefitted and if women benefited slightly more so as to close gaps. However, none of the spending components considered can achieve this outcome.

targeted. Thus, increasing the share of consumption spending dedicated to economic affairs (i.e. transportation) may still close gender gaps without neglecting crisis-stricken men. Although the effects of functional spending components on different occupations needs to be investigated further, our analysis points at the importance of fostering cross-occupational mobility. If men were able to easily move to less cyclical jobs when the need arises, they could be better sheltered from adverse business cycle and fiscal policy effects. In addition, our findings indicate potentially large cost of austerity for women. Cuts to social benefits should harm women's wages and curbing the government wage bill may decrease both female wages and employment.[38] Another insight we get is that the presence of women in the labour force may enhance the usefulness of fiscal policy in stabilizing economic aggregates. This is because women's employment and wages show more positive responses to the major spending components.

2.8 Robustness Checks and Extensions

In this section, we consider several robustness checks and extenions of our main results. These include: (i) using an alternative identification scheme, namely a recursive (Cholesky) identification, (ii) restricting the sample to part-time workers, (iii) restricting the sample to unmarried workers.

2.8.1 Alternative Identification Scheme

We test the robustness of our main results by using a standard recursive (Cholesky) ordering to identify fiscal shocks. As in our baseline analysis, we estimate the impact of total government expenditures shocks and fiscal subcomponents, looking at the responses of gender gaps and male and female labor market variables separately. The ordering of endogenous variables is the same as in the baseline VAR specification.

The impulse responses for total population of gender gaps to all government expenditure shocks are displayed in Figure 2.B.33. Impulse responses of male and female labor market variables are shown in Figure 2.B.34. This alternative identification approach yields overall qualitatively and

[38] Perugini et al. (2019) find evidence for widening gender wage gaps for the EU-28 countries after both tax-based and expenditure-based measures implemented between 2010 and 2013. However, the authors do not distinguish between spending components and do not consider gender employment gaps.

quantitatively similar results. Results for population subgroups, available upon request, are also qualitatively similar to baseline results.

2.8.2 Part-time Workers

Conducting our analysis for part-time workers (<35 hours/week) reveals substantial differences compared to our baseline sample consisting of only full-time workers. A shock to total government spending, depresses wages and employment of both genders and is no longer limited to men (see Figure 2.B.36). However, the GWG still closes since the wage drop is stronger for part-time men (see Figure 2.B.35). The vulnerable male group, i.e. young, less educated, blue-collar workers all see their employment decrease and are hence even worse off than vulnerable men with full-time jobs (compare to Table 2.A.4). The fall in female wages is driven by the fact that the government wage bill and benefits no longer have a significantly positive effect. A potential explanation is that women working part-time are less likely to work in administrative jobs or education (see Figure 2.A.13) so that the benefits of fiscal expansion remain limited to full-time female workers.

Another striking contrast is that a shock to benefits increases both male and female part-time employment. The effect is stronger for women which is why the gender employment gap still closes. Since we observed a negative effect on full-time workers, it is possible that increasing social benefits encourage employees to switch from full- to part-time employment.

Moreover, the effect on wages of part-time male workers resembles those of women. Both genders' hourly earnings increase after a consumption spending shock and with a lag also after a benefits shock. This might be explained by the fact that the occupational structure of part-time men is more similar to women's since they are less likely to work in management, and more likely to be employed in food services, administration, cleaning and maintenance (see Figure 2.A.14).

2.8.3 Unmarried Workers

In order to verify robustness and exclude partner effects we re-run all estimations for full-time non-married workers. Our results are unchanged (Figures 2.B.37 and 2.B.38), except in three cases.

First, excluding married individuals leads to a significant increase in female wages, LFP and employment after a shock to total government expenditure. This may be due to the fact that single women are overrepresented in pink-collar jobs which are more positively affected.

Second, point estimates for both genders' employment after shocks to consumption and benefits are more negative. This could point at mild partner effects, i.e. if married individuals are included employment losses are muted due to spouses entering the labour force after the partner's job loss. Alternatively, the result may also be due to married workers enjoying greater job security.

Third, female employment no longer increases significantly after a wage bill shock, which seems to be driven by private sector women. These are less likely to work in administration, healthcare and education if they are married (Figure 2.A.15). Thus, they benefit less from fiscal expansions. This highlights, once more, that the effectiveness of fiscal policy in inducing aggregate wage changes hinges on its spillovers into the private sector. These, in turn, crucially depend on occupational sorting.

2.9 Limitations and Future Research

While we provide important insights into the impact of fiscal policy on gender gaps, several issues remain unaddressed. Future research should explore whether fiscal policy has asymmetric effects on men and women depending on the state of the cycle. This would allow for a better assessment of how to offset inequitable effects of economic slumps. Fiscal components may target different functions along the business cycle. For example, the fraction of social benefits that is spent on unemployment insurance, food stamps and refundable tax credits expands during crises. Consequently, the impact of this fiscal component on labour market outcomes of men and women might change. Another non-linearity worth investigating is whether the effects of austerity on gender gaps differ from those of fiscal expansions. Knowledge of which spending components to cut in times of tight budgets without widening gender gaps is still wanting. Furthermore, investigating the gender implications of financing fiscal expansions through tax or deficit increases should be a relevant policy issue.

In order to better understand the mechanisms behind our results, we need further evidence on the effects of functional spending components on occupations. Such an analysis should examine spillovers of e.g. health and education spending on unrelated professions. Additionally, movements between occupations and sectors as well as into and out of the labour force could be explored. The CPS data used in this paper is not suitable for this endeavour since it is not a panel dataset. Therefore, occupational and industry switches can only be observed over short time horizons. In addition, our VAR-based analysis should be complemented with a micro-econometric approach that allows us to control for several individual characteristics simultaneously. Demographic trends as well as selection biases could thus be accounted for.

2.10 Conclusions

In this paper we analyzed the labour market effects of fiscal policy from a gender perspective. Overall, increases in total government spending affect female wages and employment more favourably than male's. However, the effect varies by demographic subgroup and fiscal components. Thus, policy makers may alter the composition of expenditures according to their objectives. If their goal is to decrease gender wage and employment gaps, expanding the government wage bill and social benefits is most appropriate. However, if the fiscal authority aims at assisting young, less educated and blue-collar men who are most affected by negative business cycle shocks, defense and investment spending are the preferable fiscal tools. Hence, these two goals are incompatible.

Our analysis provided several other policy implications. First, heterogeneous results across gender are mainly driven by men's losses in the private sector as they tend to sort into occupations that are not affected by most spending components. To remedy this, barriers to cross-occupational mobility should be minimized. Second, heterogeneous effects across age and education subgroups create a trade-off between boosting aggregate employment or wages. For example, shocks to social benefits push young and less educated women out of the labour force, which increases the average female wage but reduces employment. These asymmetries should be kept in mind when designing equitable spending programmes, and the re-entry into the labour force of less favourably affected groups should be actively encouraged. Third, our findings indicate potentially large costs of austerity for women's wages and employment, especially in case of cuts to the government wage bill. Finally, since women respond more positively to major spending components, encouraging their labour force participation may enhance the efficiency of fiscal policy as a stabilization tool.

Bibliography

Akitoby, M. B., M. J. Honda, and H. Miyamoto (2019). Countercyclical Fiscal Policy and Gender Employment: Evidence from the G-7 Countries. *IMF Working Paper Series WP/19/4*.

Albanesi, S. and A. Şahin (2018). The gender unemployment gap. *Review of Economic Dynamics 30*, 47–67.

Altonji, J. G. and R. M. Blank (1999). Race and gender in the labor market. *Handbook of labor economics 3*, 3143–3259.

Anderson, E., A. Inoue, and B. Rossi (2016). Heterogeneous consumers and fiscal policy shocks. *Journal of Money, Credit and Banking 48*(8), 1877–1888.

Arias, J., J. Rubio-Ramirez, and D. Waggoner (2014). Inference based on svar identified with sign and zero restrictions: theory and applications. *FRB Atlanta Working paper, Federal Reserve Bank of Atlanta 2014-1*.

Bermperoglou, D., E. Pappa, and E. Vella (2017). The government wage bill and private activity. *Journal of Economic Dynamics and Control 79*, 21–47.

Blau, F. D. and L. M. Kahn (2000). Gender differences in pay. *Journal of Economic perspectives 14*(4), 75–99.

Blau, F. D. and L. M. Kahn (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature 55*(3), 789–865.

Braunstein, E. and J. Heintz (2008). Gender bias and central bank policy: Employment and inflation reduction. *International Review of Applied Economics 22*(2), 173–186.

Bredemeier, C. (2015). Household Specialization and the Labor-Supply Elasticities of Women and Men. *University of Cologne Working Paper Series in Economics 81*.

Bredemeier, C., F. Juessen, and R. Winkler (2017). Man-cessions, fiscal policy, and the gender composition of employment. *Economics Letters 158*, 73–76.

Bredemeier, C., F. Juessen, and R. Winkler (2019). Fiscal policy and occupational employment dynamics. *Journal of Money, Credit, and Banking*.

Brown, S. and K. Taylor (2009). Reservation Wages, Expected Wages and the Duration of Unemployment: Evidence from British Panel Data. *IZA Discussion Paper Series* (3981).

Chodorow-Reich, G., L. Feiveson, Z. Liscow, and W. G. Woolston (2012). Does state fiscal relief during recessions increase employment? evidence from the american recovery and reinvestment act. *American Economic Journal: Economic Policy 4*(3), 118–45.

Clark, K. B. and L. H. Summers (1981). Demographic differences in cyclical employment variation. *The Journal of Human Resources 16*(1), 61–79.

Cloyne, J. S. and P. Surico (2016). Household debt and the dynamic effects of income tax changes. *The Review of Economic Studies 84*(1), 45–81.

Deaton, A. (1985). Panel data from time-series of cross-sections. *Journal of Econometrics 30*, 109–126.

European Parliament (2015). Gender-responsive budgeting: Innovative approaches to budgeting. *Briefing, European Parliamentary Research Service*.

Ferriere, A. and G. Navarro (2018). The heterogeneous effects of government spending: It's all about taxes. *Mimeo*.

Giavazzi, F. and M. McMahon (2012). The household effects of government spending. In *Fiscal policy after the financial crisis*, pp. 103–141. University of Chicago Press.

Hoynes, H., D. L. Miller, and J. Schaller (2012). Who Suffers During Recessions ? *Journal of Economic Perspectives 26*(3), 27–47.

IMF (2017). Gender budgeting in G7 countries. *IMF Policy Papers*.

Kim, H. (2018). Fiscal Policy, Wages, and Jobs in the U.S. *Auburn Economics Working Paper* (2018-02).

Kolberg, J. E. (1991). The Gender Dimension of the Welfare State. *International Journal of Sociology 21*, 119–148.

Misra, K. and P. Surico (2014). Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics 6*(4), 84–106.

Mountford, A. and H. Uhlig (2009). What are the effects of fiscal policy shocks? *Journal of applied econometrics 24*(6), 960–992.

Nekarda, C. J. and V. A. Ramey (2011). Industry evidence on the effects of government spending. *American Economic Journal: Macroeconomics 3*(1), 36–59.

Ngai, L. R. and B. Petrongolo (2017). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics 9*(4), 1–44.

Olivetti, C. and B. Petrongolo (2014). Gender gaps across countries and skills: Demand, supply and the industry structure. *Review of Economic Dynamics 17*(4), 842–859.

Pappa, E. (2009). The effects of fiscal shocks on employment and the real wage. *International Economic Review 50*(1), 217–244.

Perugini, C., J. E„arkoviДЉ RakiДЉ, and M. VladisavljeviДЉ (2019). Austerity and Gender Inequalities in Europe in Times of Crisis. *Cambridge Journal of Economics 43*(3), 733–767.

Prasad, E. S. (2003). What Determines Reservation Wages of Unemployed Workers? New Evidence from German Micro Data. *IMF Working Paper Series* (WP/03/4).

Ramey, V. A. (2011). Identifying government spending shocks: It's all in the timing. *The Quarterly Journal of Economics 126*(1), 1–50.

Seguino, S. and J. Heintz (2012). Monetary Tightening and the Dynamics of US Race and Gender Stratification. *American Journal of Economics and Sociology 71*(3).

Shin, D. (1999). An Equilibrium Theory of Wage and Employment Cyclicality by Gender and by Industry. *Southern Economic Journal 65*(3), 451–471.

Solon, G., R. Barsky, and J. A. Parker (1994). Measuring the cyclicality of real wages: how important is composition bias? *The quarterly journal of economics 109*(1), 1–25.

Uhlig, H. (2005). What are the effects of monetary policy on output? results from an agnostic identification procedure. *Journal of Monetary Economics 52*(2), 381–419.

Wall, H. J. and K. M. Engemann (2009). The Effects of Recessions Across Demographic Groups. *Federal Reserve Bank of St. Louis Working Paper Series* (2009-052A).

2.A Tables

Population group	**Average increase**	**1980Q1-1980Q3**	**1981Q3-1982Q4**	**1990Q3-1991Q1**	**2001Q1-2001Q4**	**2007Q4-2009Q2**
Women	29.91%	9.36%	29.95%	9.29%	27.55%	73.40%
Men	48.48%	30.63%	53.79%	17.96%	27.36%	112.68%
Men, aged 25-34	61.69%	48.68%	68.52%	13.69%	31.45%	146.09%
Men, no college degree	47.05%	31.83%	55.73%	19.00%	20.51%	108.619%
Men, blue-collar	57.65%	38.21%	59.45%	22.66%	23.01%	144.94%

Notes: Increases in unemployment are measured from peak to trough. Micro-data obtained from CEPR-CPS. Recession dates from NBER. Red cells indicate that the unemployment increase for the male subgroup under consideration was larger than the male average.

Table 2.A.1: Increase in unemployment during recessions for full-time workers

Variable	**Mean**	**Std.Dev.**	**Min**	**Max**	**Obs.**
Wages					
Women's hourly real wage	18.189	1.954	15.103	21.342	116
Men's hourly real wage	23.868	1.267	22.305	26.26	116
Gender wage gap, $ln(w^{male}) - ln(w^{fem})$	0.276	0.066	0.191	0.416	116
Employment					
Women's employment rate	0.406	0.022	0.364	0.443	116
Men's employment rate	0.595	0.023	0.545	0.649	116
Gender employment gap	0.189	0.028	0.157	0.276	116
Hours worked					
Women's weekly hours worked	41.000	0.365	40.295	41.426	116
Men's weekly hours worked	43.534	0.366	42.768	44.218	116
Hours gap	0.060	0.004	0.051	0.069	116

Notes: Statistics for 1979Q1 through 2007Q4. Data obtained from CEPR-CPS.

Table 2.A.2: Summary statistics of labour market variables for full-time workers by gender

Spending component	Women		Men		Gender gaps	
	wage	**empl**	**wage**	**empl**	**wage**	**empl**
Total			-		↓	↓
Consumption	+	-		-	↓	↓
Wage bill	+	+	-		↓	↓
Social benefits	+	-	-	-	↓	↓
Investment	-	+	-	+	↑	↑
Defense	+	+	+			

Notes: + indicates that the response was positive within the first five quarters after shock, − denotes a negative response. A blank cell indicates that the response was not significantly different from zero during the first five quarters. Downward arrows symbolize a narrowing gender gap and upward arrows represent a widening gender gap.

Table 2.A.3: Effects of fiscal shocks on women, men and gender gaps

Men	Aged 16-24		Less educated		Blue-collar		Private sector	
Spending component	**wage**	**empl**	**wage**	**empl**	**wage**	**empl**	**wage**	**empl**
Total	-		-		-	+	-	-
Consumption	-	-	-	-	-		-	-
Wage bill	-		-		-		-	-
Social benefits	-		-	-	-	-	-	-
Investment	+	+		+	+/-	+	-	+/-
Defense		+			+	+		

Notes: + indicates that the response was positive within the first five quarters after shock, − denotes a negative response and +/− shows that the sign changed from positive to negative. A blank cell indicates that the response was not significantly different from zero during the first five quarters.

Table 2.A.4: Effects of fiscal shocks on male subgroups most affected by business cycles

Women	Aged 16-24		Less educated		Blue-collar		Private sector	
Spending component	**wage**	**empl**	**wage**	**empl**	**wage**	**empl**	**wage**	**empl**
Total			+	+	+	+		
Consumption			+	-	+	+	+	-
Wage bill	+		+	+		+	+	+
Social benefits	-	-	+	-	-	-		-
Investment	-	+		+	-	+	-	+
Defense	+	+	+		+	+	+	+

Notes: + indicates that the response was positive within the first five quarters after shock, − denotes a negative response. A blank cell indicates that the response was not significantly different from zero during the first five quarters.

Table 2.A.5: Effects of fiscal shocks on female subgroups

	Private sector				Public sector			
	Women		Men		Women		Men	
Spending component	**wage**	**empl**	**wage**	**empl**	**wage**	**empl**	**wage**	**empl**
Total			-	-	+	+	+	+
Consumption	+	-	-	-	+		+	
Wage bill	+	+	-	-	+		+	
Social benefits		-	-	-	+	+	-	-
Investment	-	+	-	+/-	-	+	-	+
Defense	+	+			+	+	+	+

Notes: + indicates that the response was positive within the first five quarters after shock, − denotes a negative response and +/− shows that the sign changed from positive to negative. A blank cell indicates that the response was not significantly different from zero during the first five quarters.

Table 2.A.6: Comparison of the effects of fiscal shocks between private and public sector workers

2.B Figures

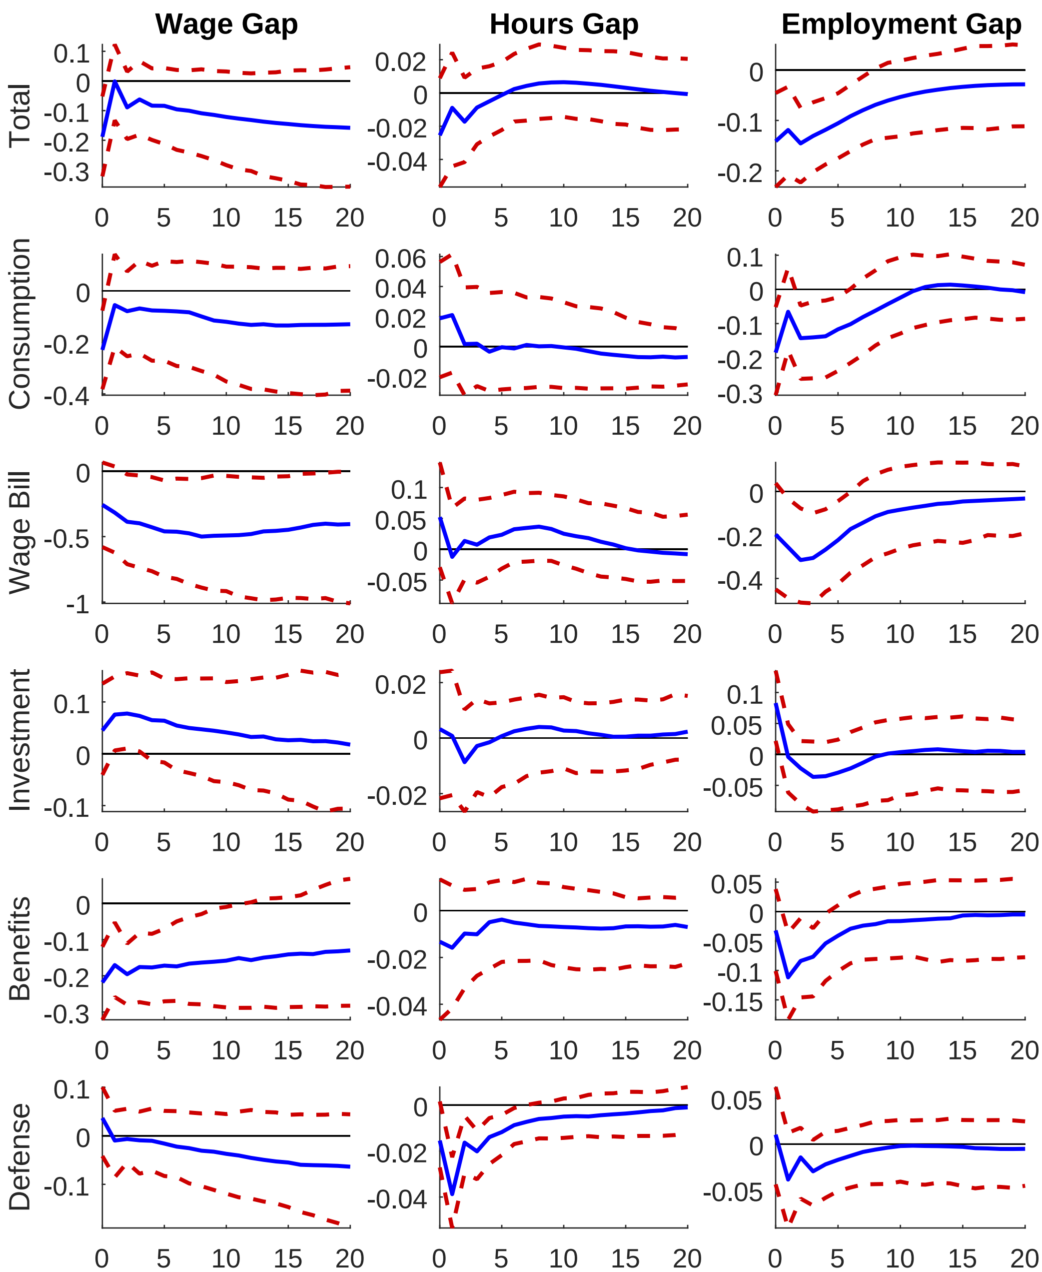

Figure 2.B.1: IRFs of gender gaps to shocks in different spending components

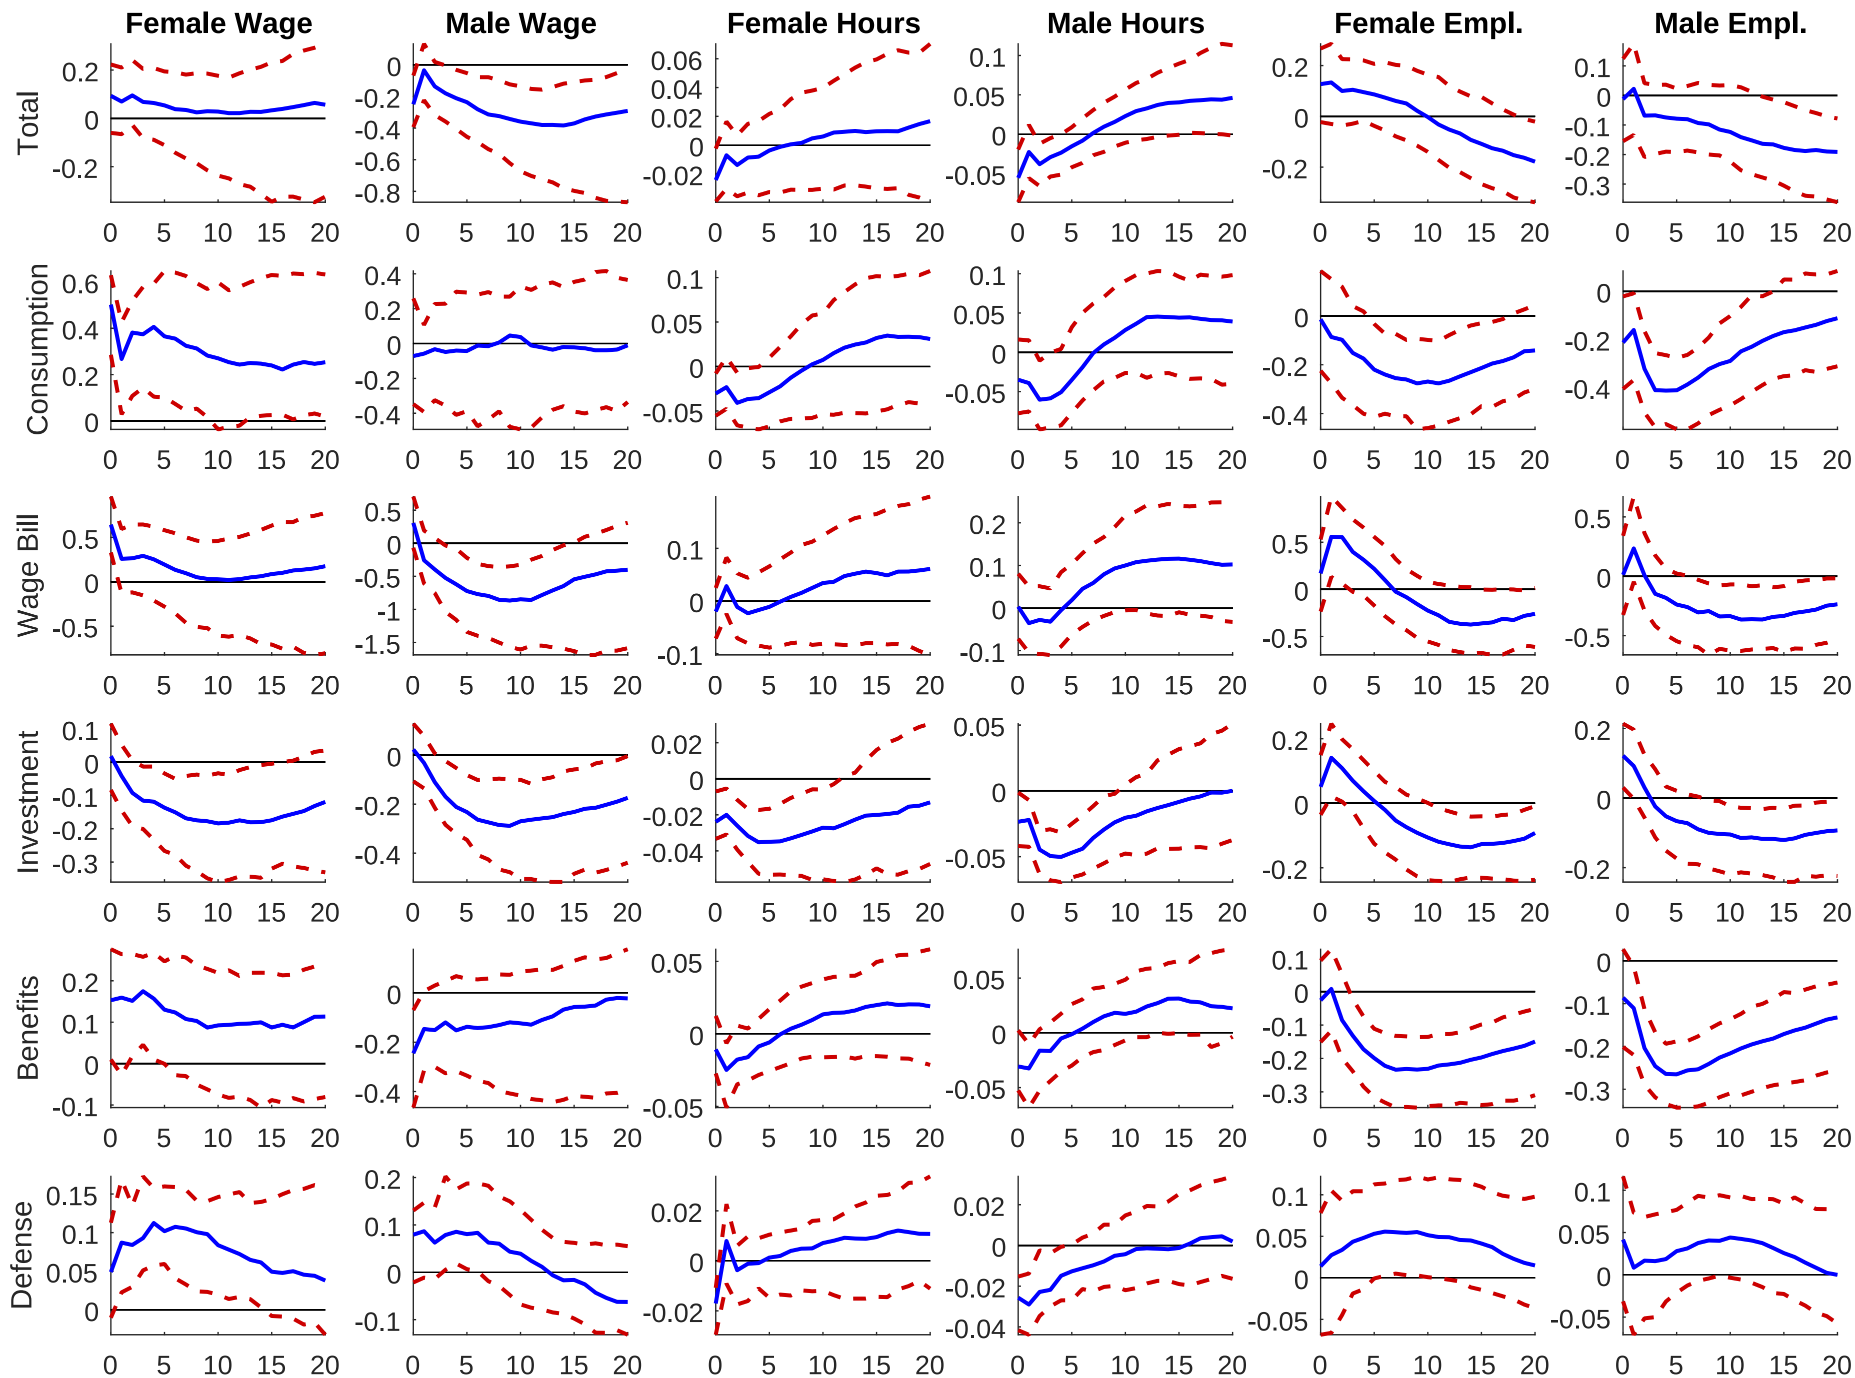

Figure 2.B.2: IRFs of labour market outcomes by gender to shocks in different spending components

Total Government Expenditure

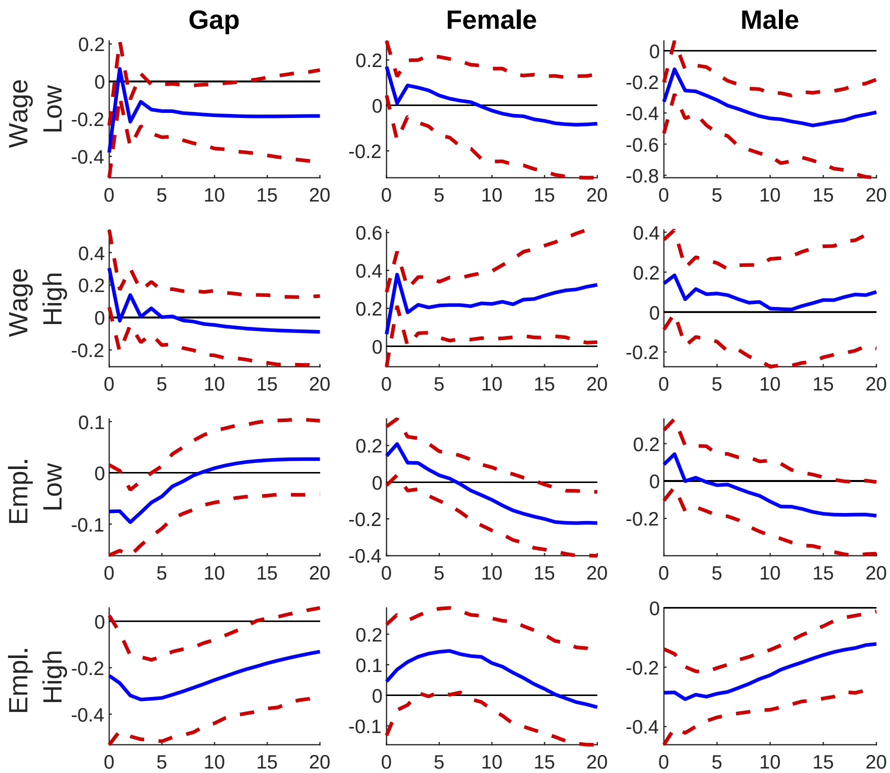

Figure 2.B.3: IRFs of gender gaps to a shock in total government spending - no college degree (low) vs. college degree (high)

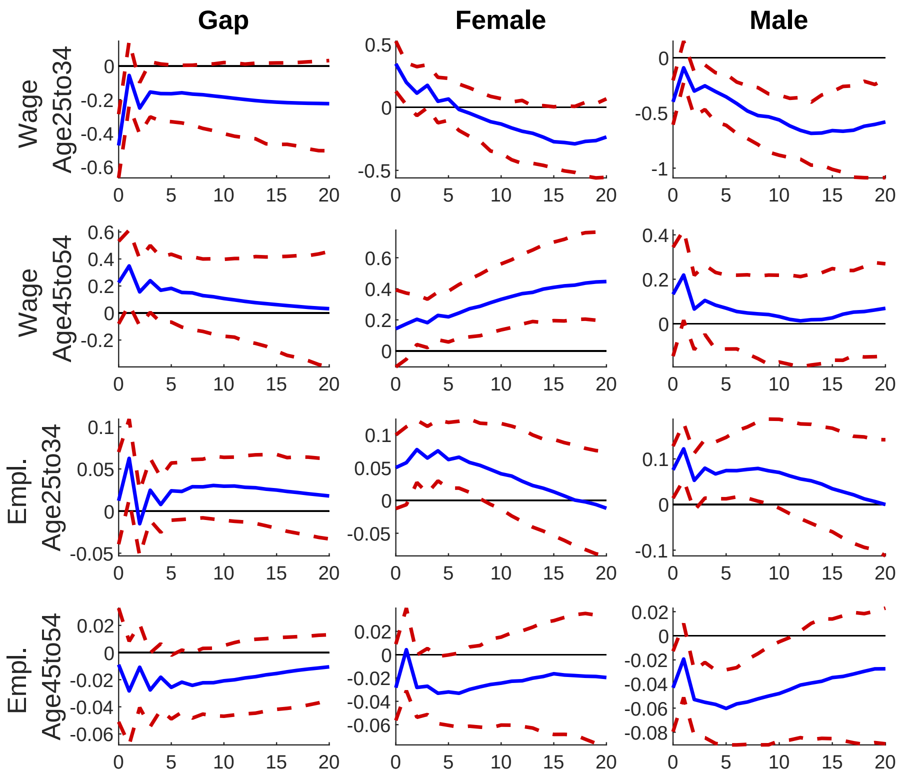

Figure 2.B.4: IRFs of gender gaps to a shock in total government spending - workers aged 25-34 vs. 45-54

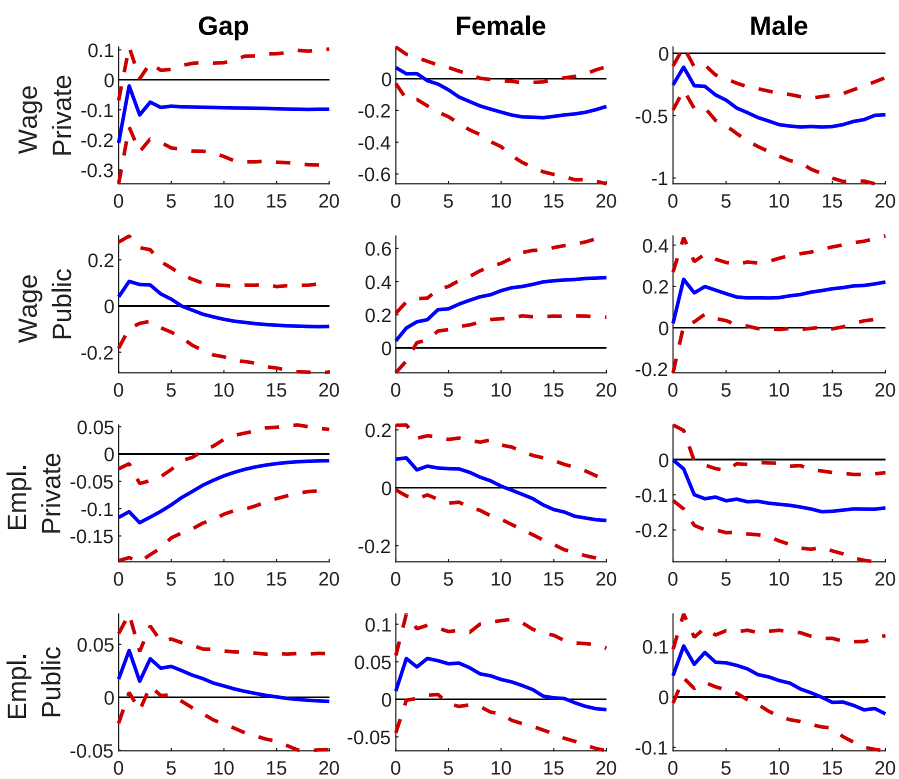

Figure 2.B.5: IRFs of gender gaps to a shock in total government spending - public vs. private sector

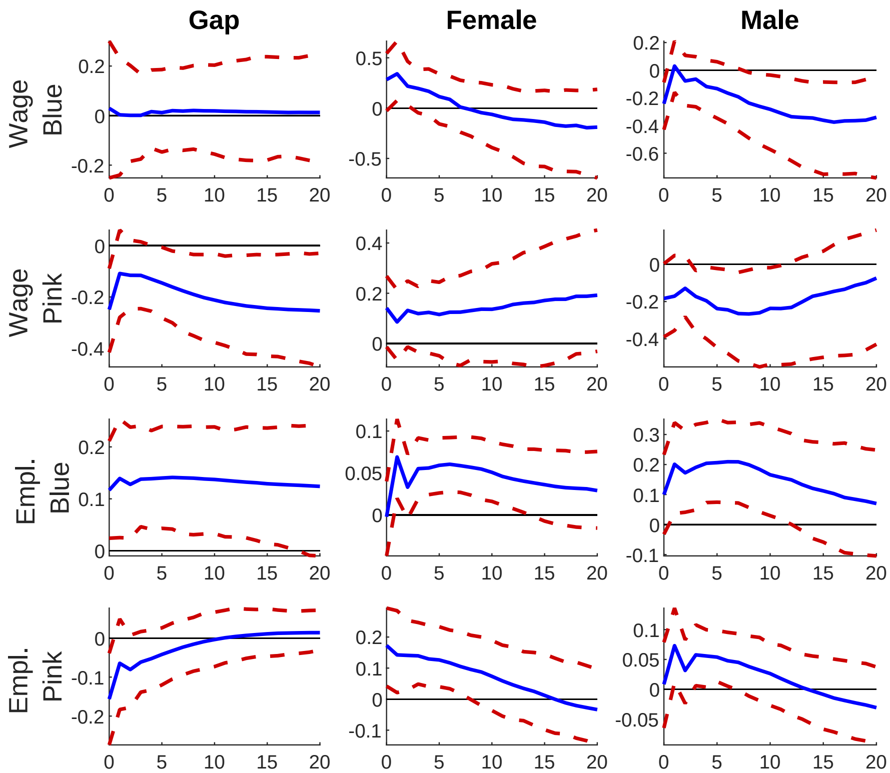

Figure 2.B.6: IRFs of gender gaps to a shock in total government spending - pink- vs. blue-collar occupations

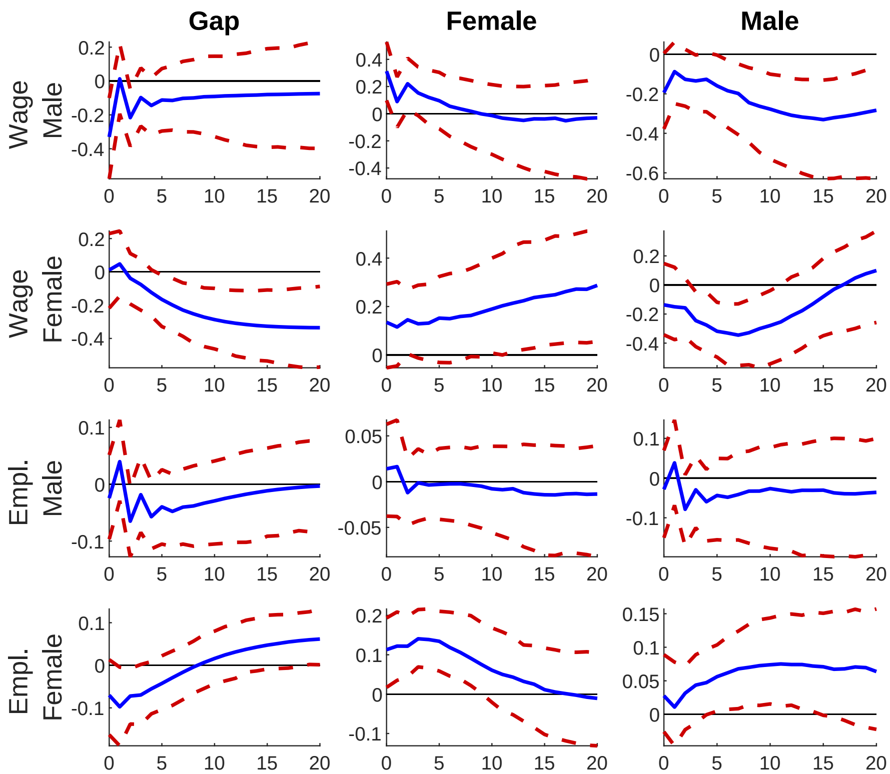

Figure 2.B.7: IRFs of gender gaps to a shock in total government spending - male vs. female dominant industry

Government Consumption Expenditure

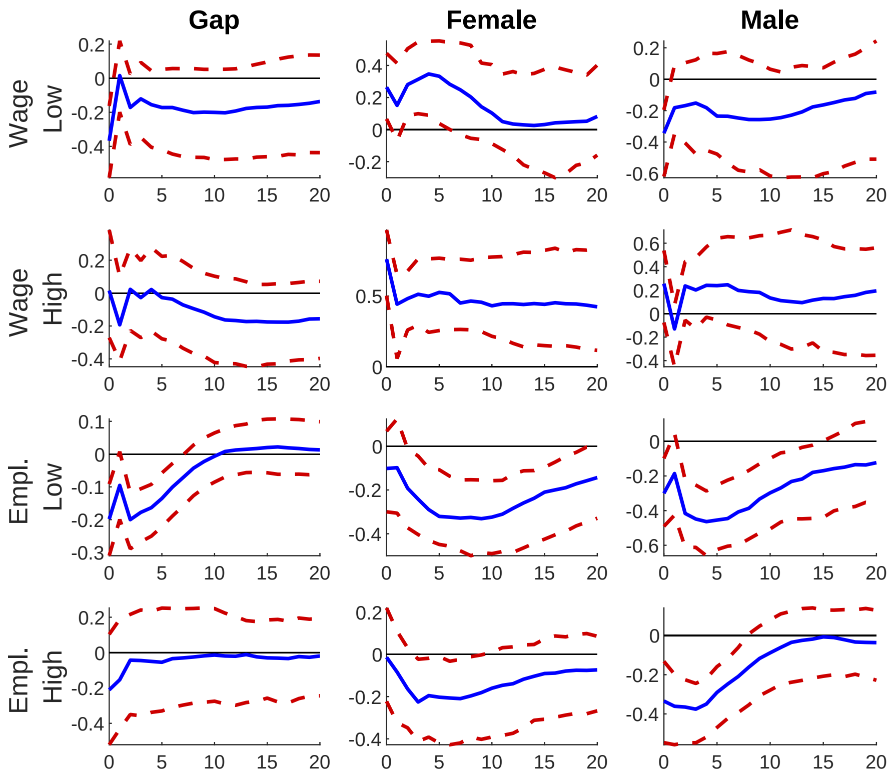

Figure 2.B.8: IRFs of gender gaps to a shock in government consumption spending - no college degree (low) vs. college degree (high)

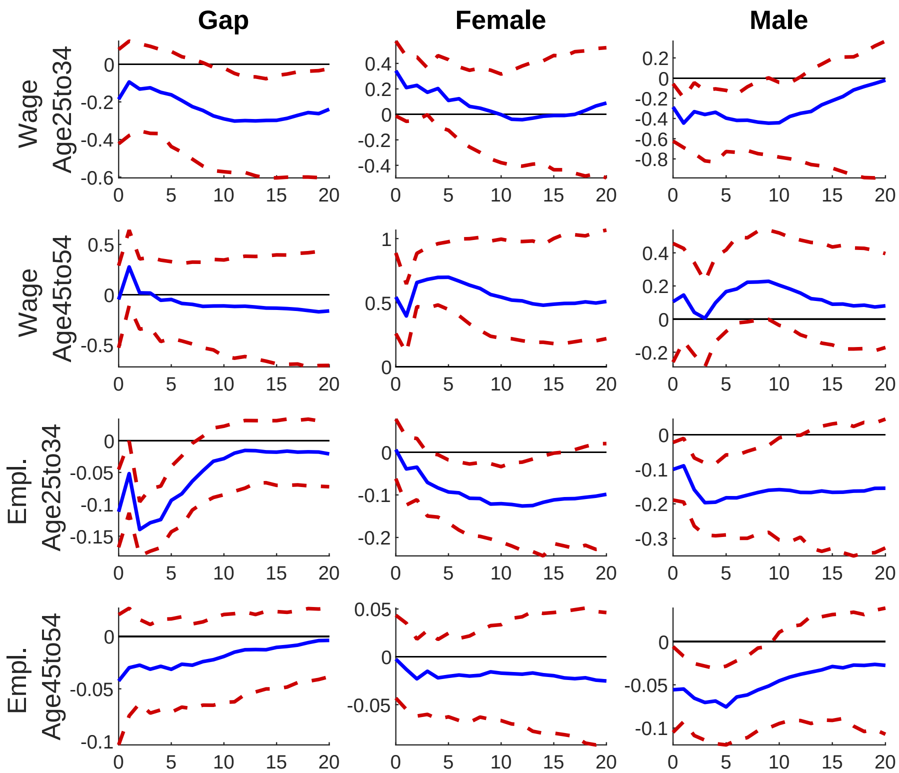

Figure 2.B.9: IRFs of gender gaps to a shock in government consumption spending - workers aged 25-34 vs. 45-54

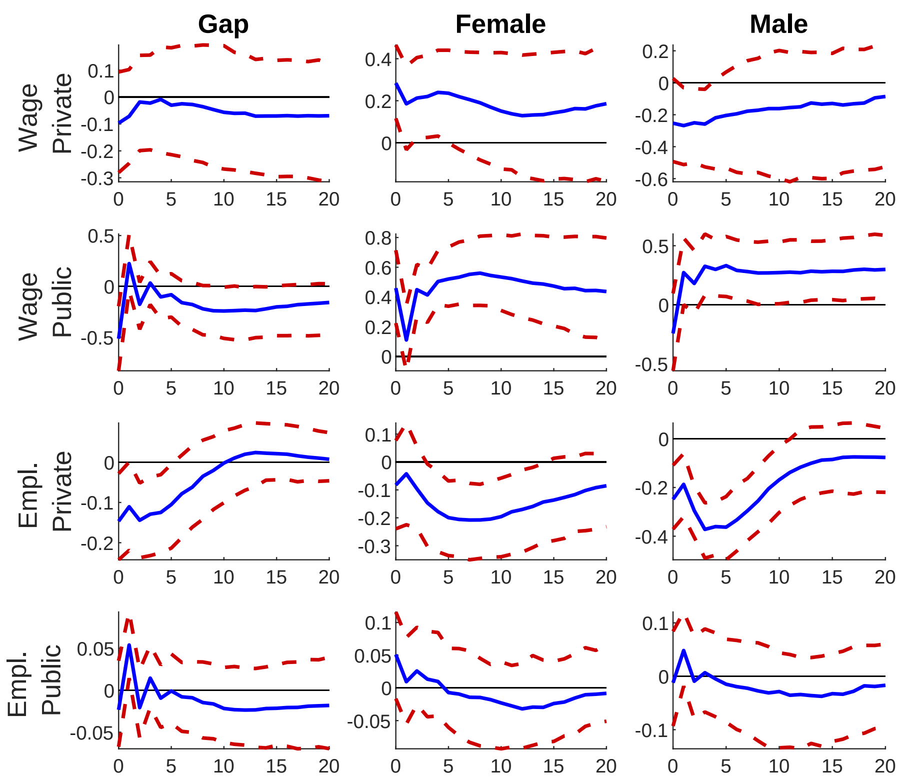

Figure 2.B.10: IRFs of gender gaps to a shock in government consumption spending - public vs. private sector

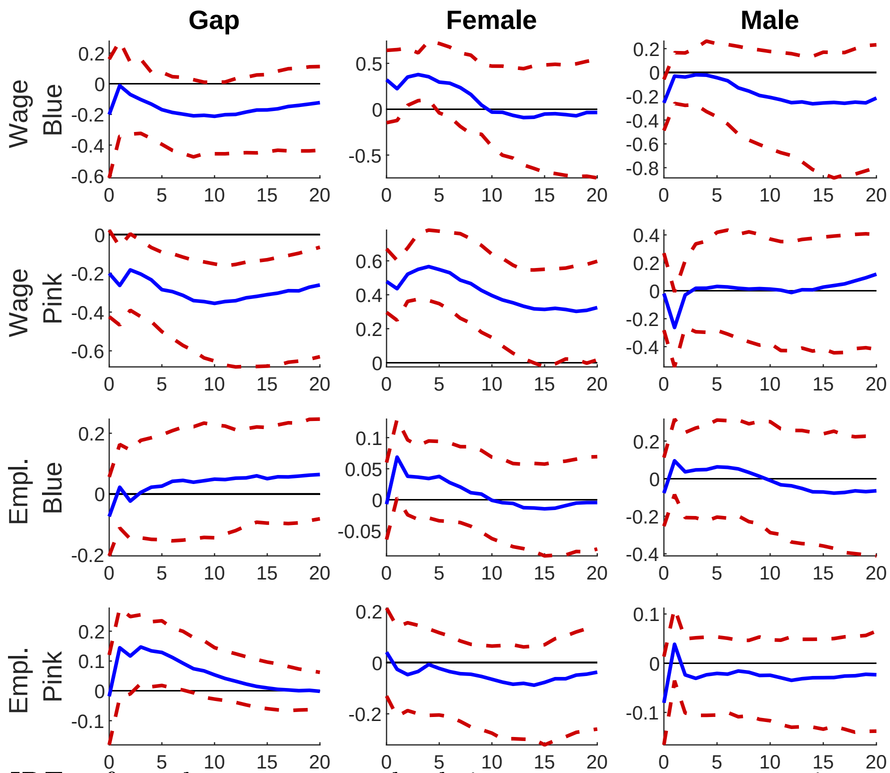

Figure 2.B.11: IRFs of gender gaps to a shock in government consumption spending - pink- vs. blue-collar occupations

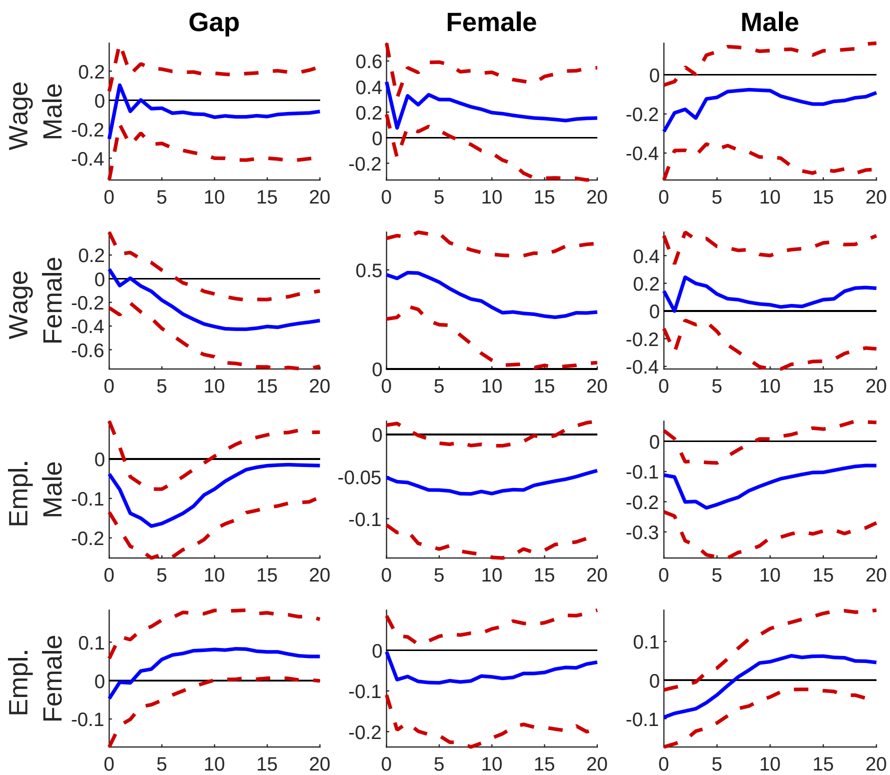

Figure 2.B.12: IRFs of gender gaps to a shock in government consumption spending - male vs. female dominant industry

Government Wage Bill Expenditure

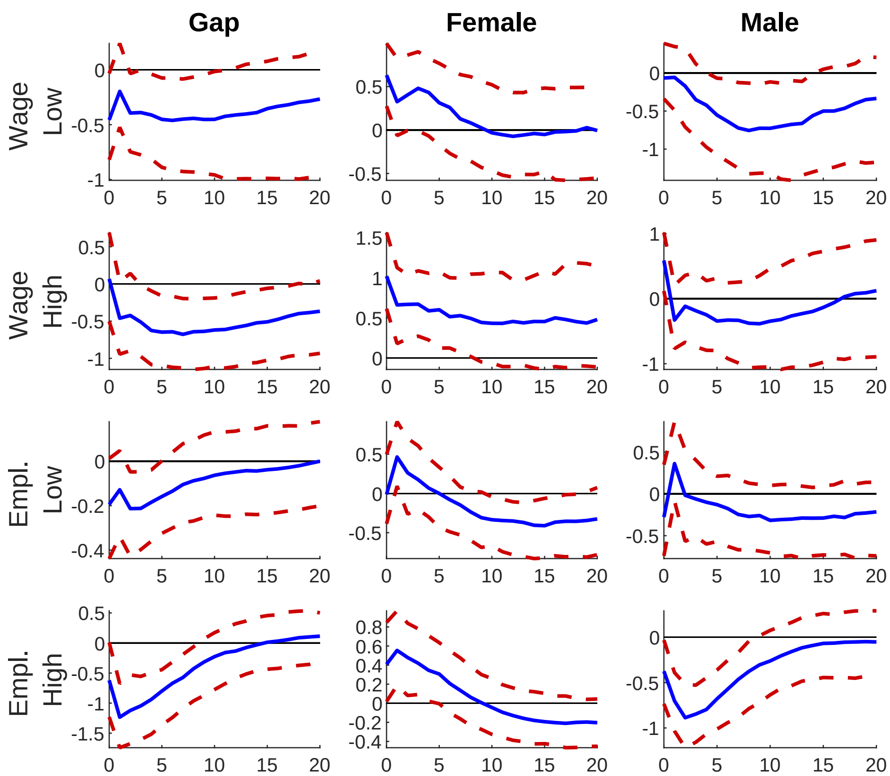

Figure 2.B.13: IRFs of gender gaps to a shock in government wage bill spending - no college degree (low) vs. college degree (high)

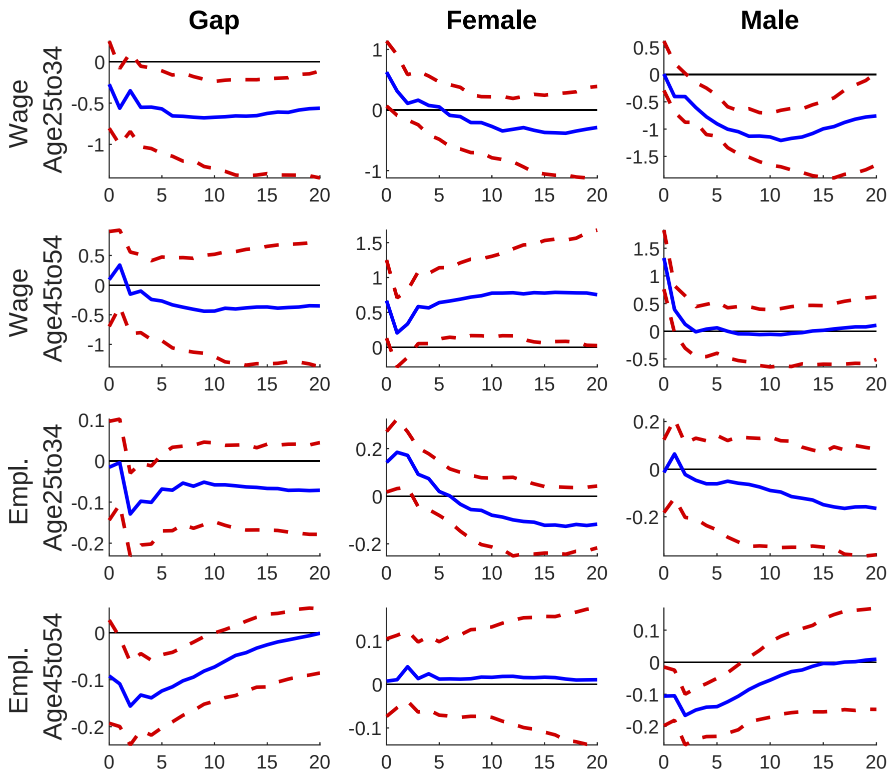

Figure 2.B.14: IRFs of gender gaps to a shock in government wage bill spending - workers aged 25-34 vs. 45-54

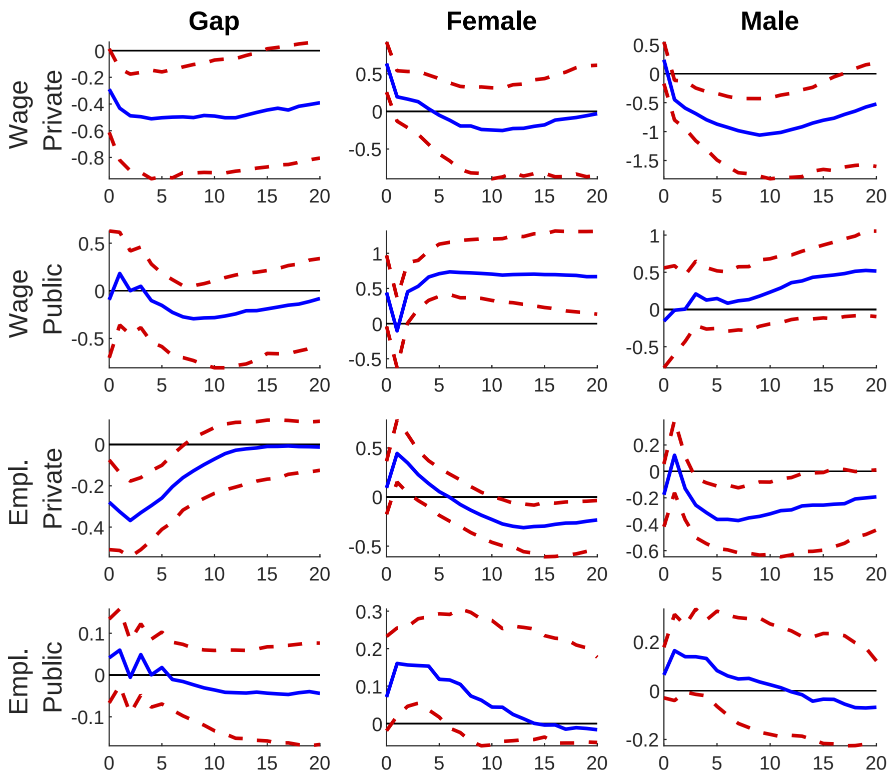

Figure 2.B.15: IRFs of gender gaps to a shock in government wage bill spending - public vs. private sector

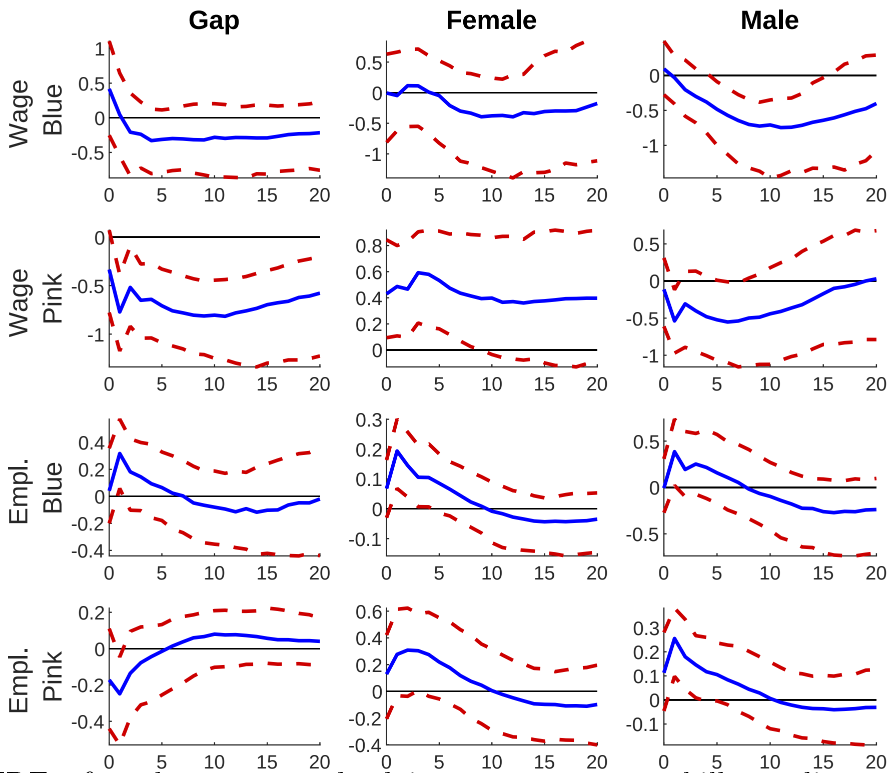

Figure 2.B.16: IRFs of gender gaps to a shock in government wage bill spending - pink- vs. blue-collar occupations

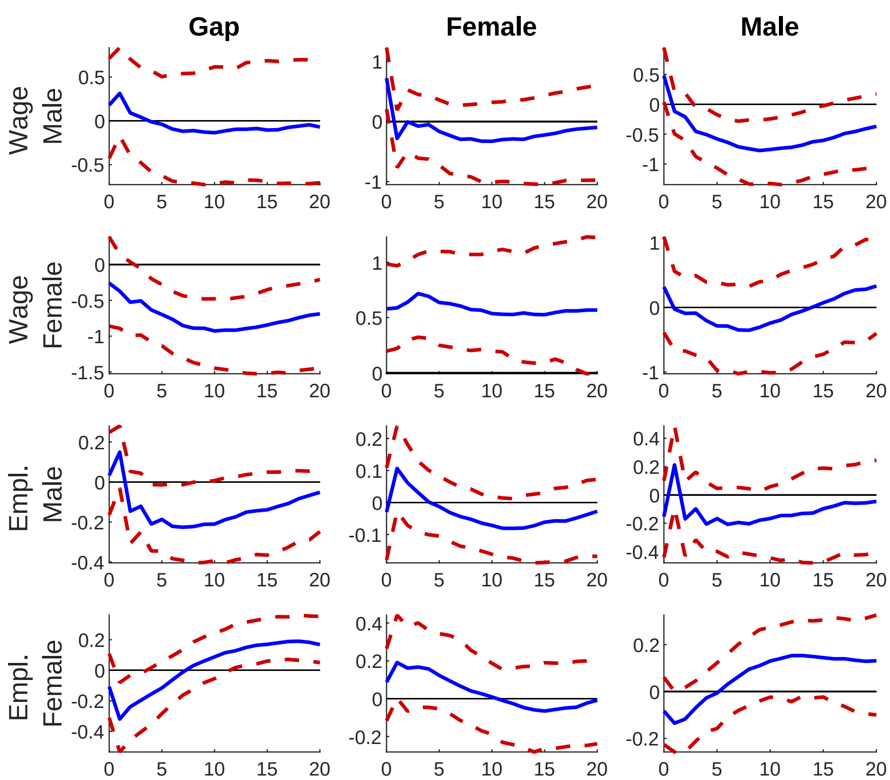

Figure 2.B.17: IRFs of gender gaps to a shock in government wage bill spending - male vs. female dominant industry

Government Investment Expenditure

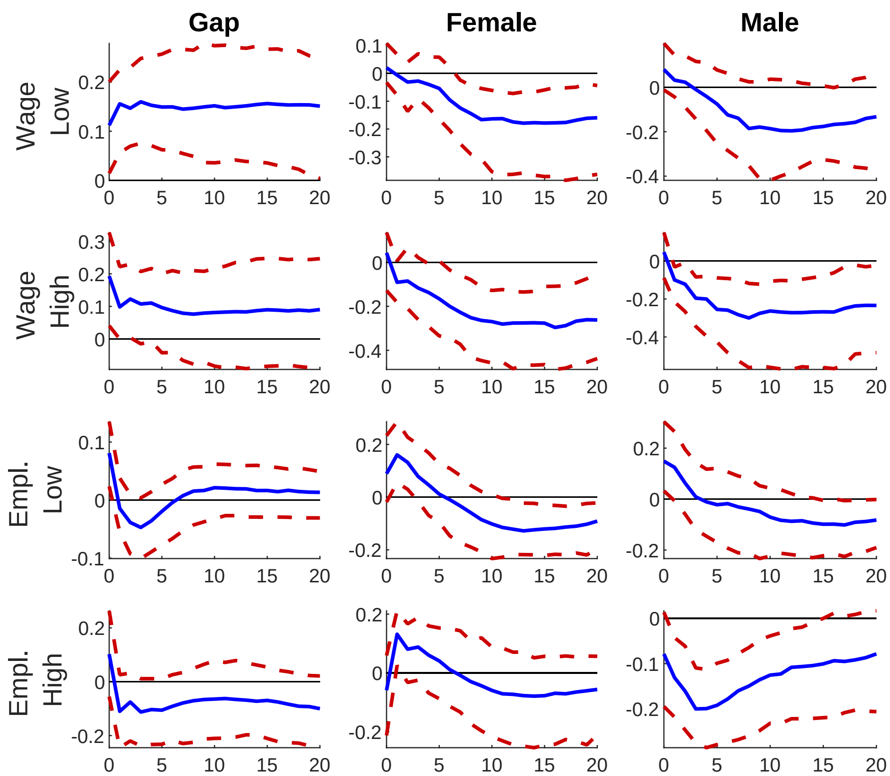

Figure 2.B.18: IRFs of gender gaps to a shock in government investment spending - no college degree (low) vs. college degree (high)

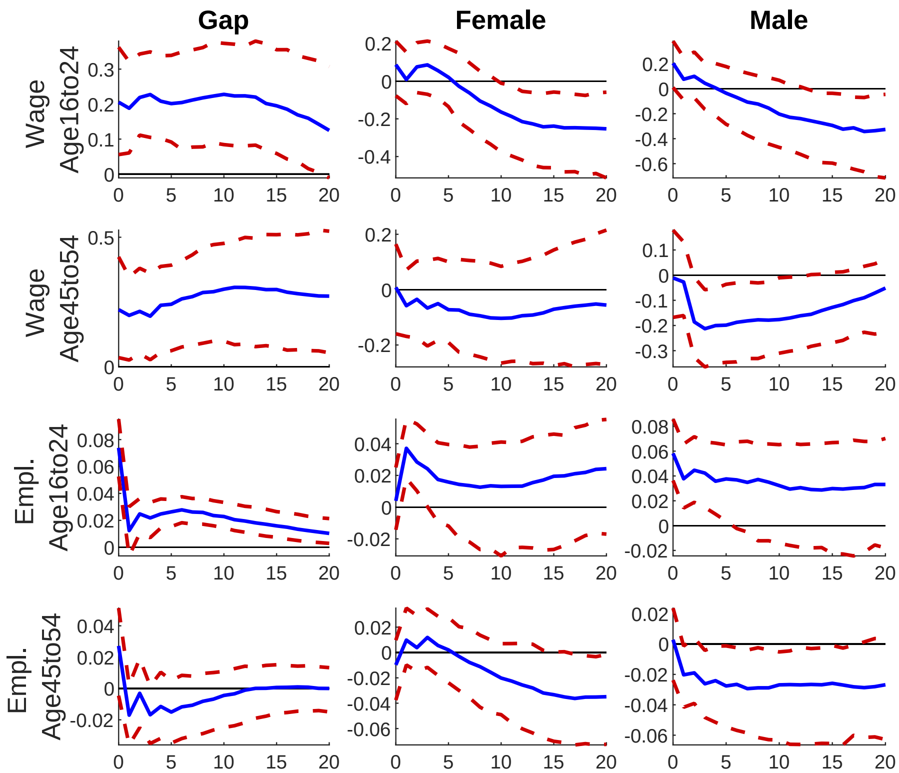

Figure 2.B.19: IRFs of gender gaps to a shock in government investment spending - workers aged 16-24 vs. 45-54

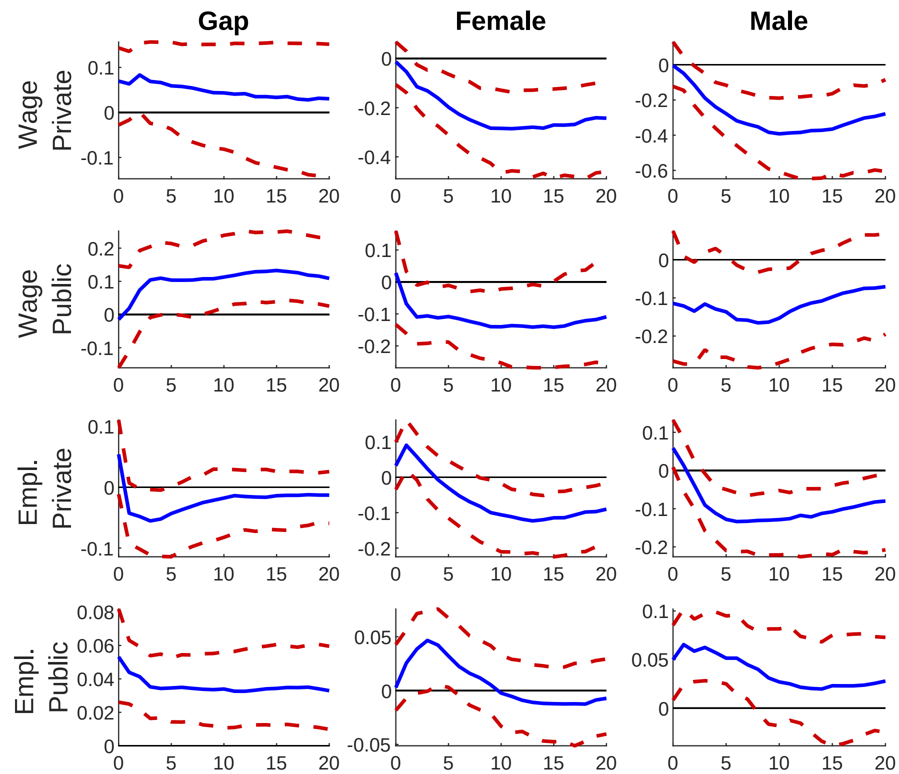

Figure 2.B.20: IRFs of gender gaps to a shock in government investment spending - public vs. private sector

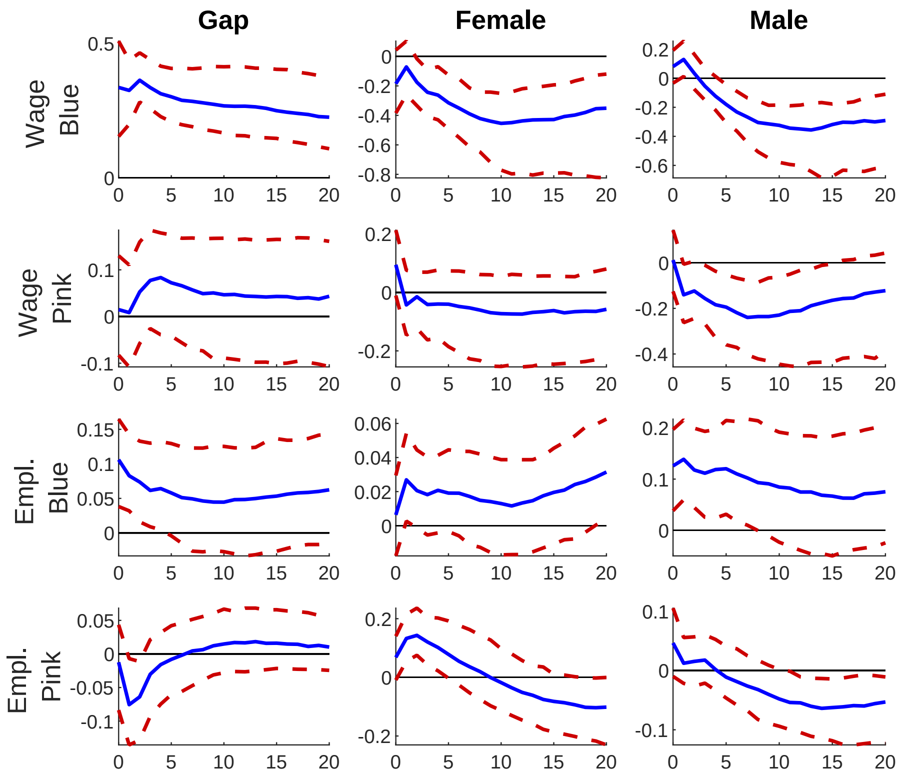

Figure 2.B.21: IRFs of gender gaps to a shock in government investment spending - pink- vs. blue-collar occupations

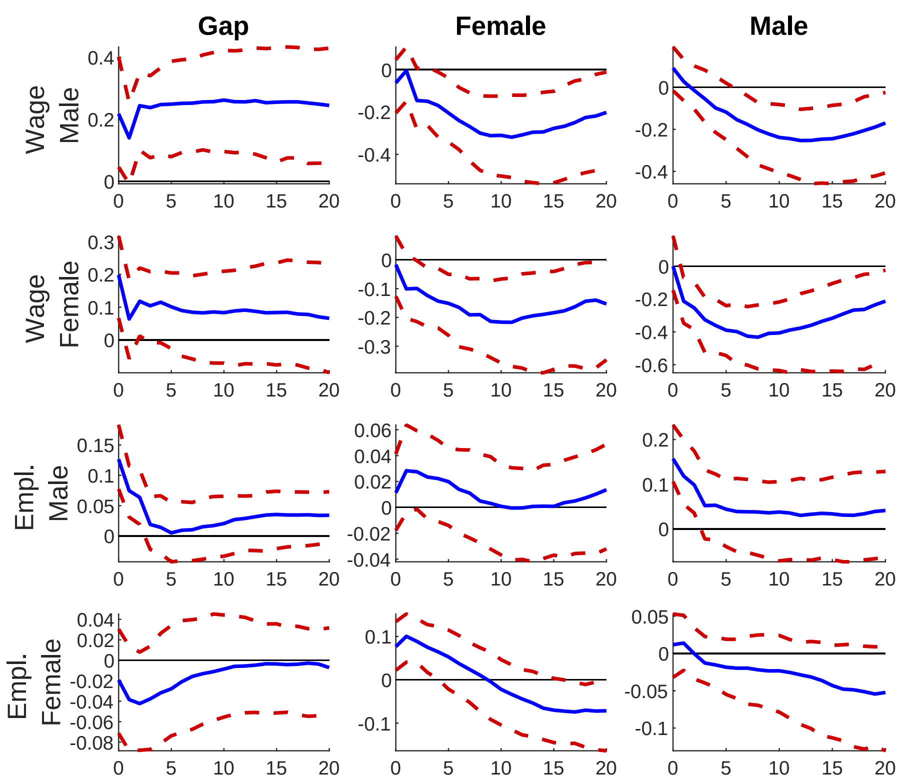

Figure 2.B.22: IRFs of gender gaps to a shock in government investment spending - male vs. female dominant industry

Government Benefits Expenditure

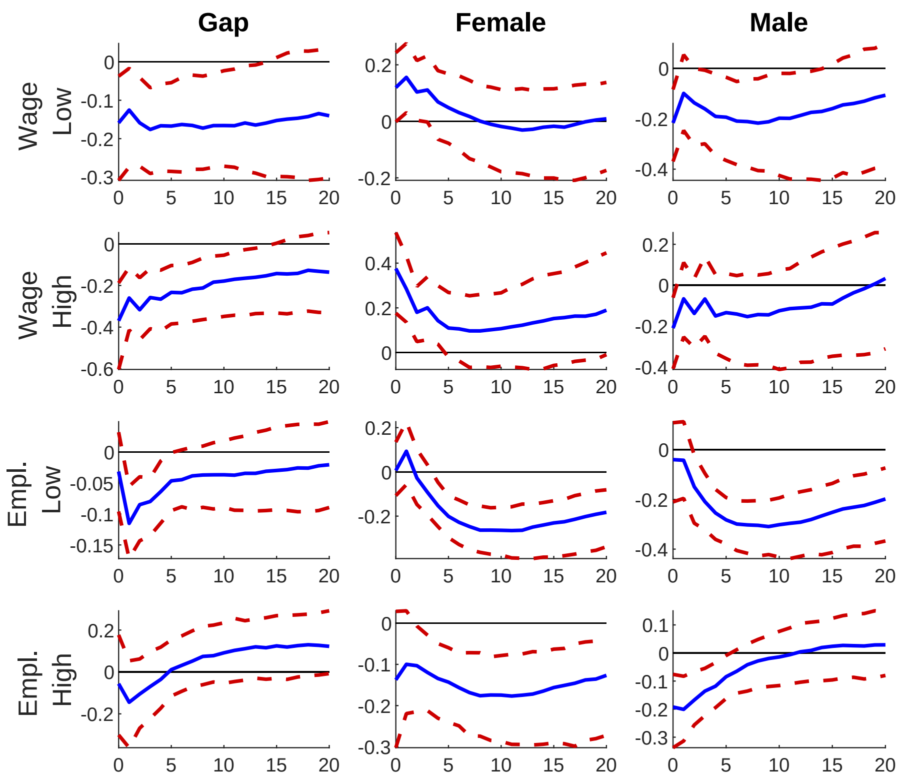

Figure 2.B.23: IRFs of gender gaps to a shock in government benefits spending - no college degree (low) vs. college degree (high)

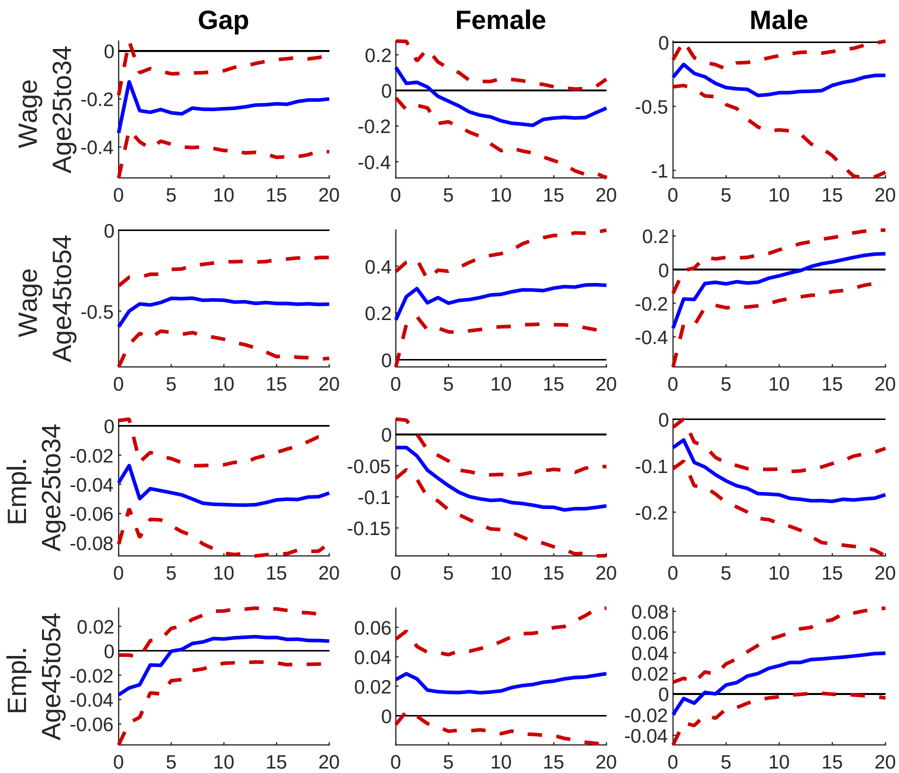

Figure 2.B.24: IRFs of gender gaps to a shock in government benefits spending - workers aged 25-34 vs. 45-54

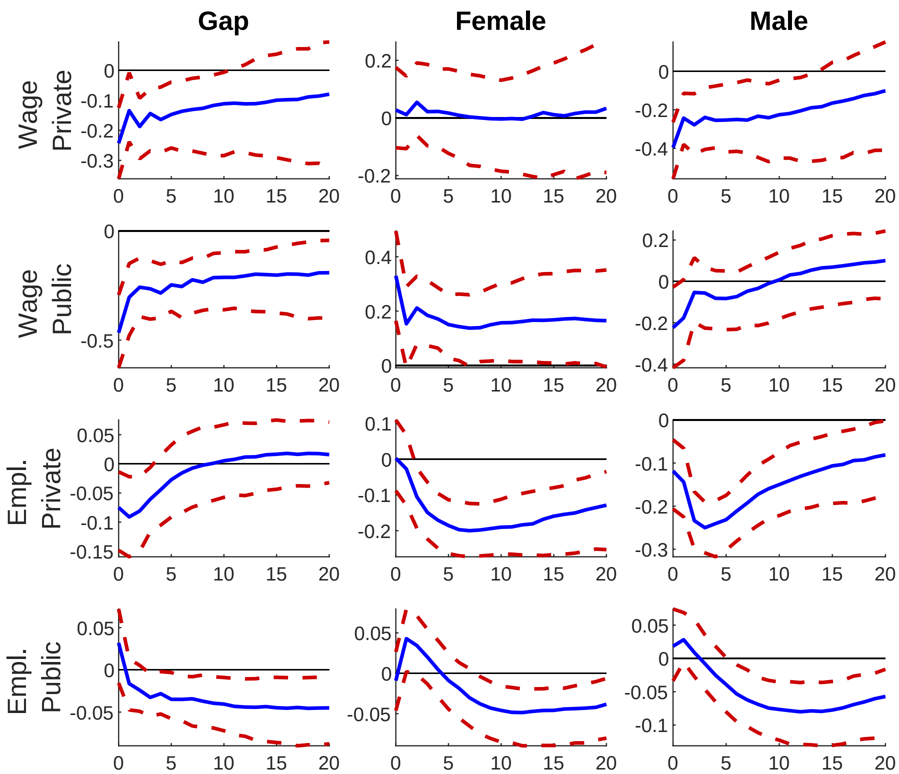

Figure 2.B.25: IRFs of gender gaps to a shock in government benefits spending - public vs. private sector

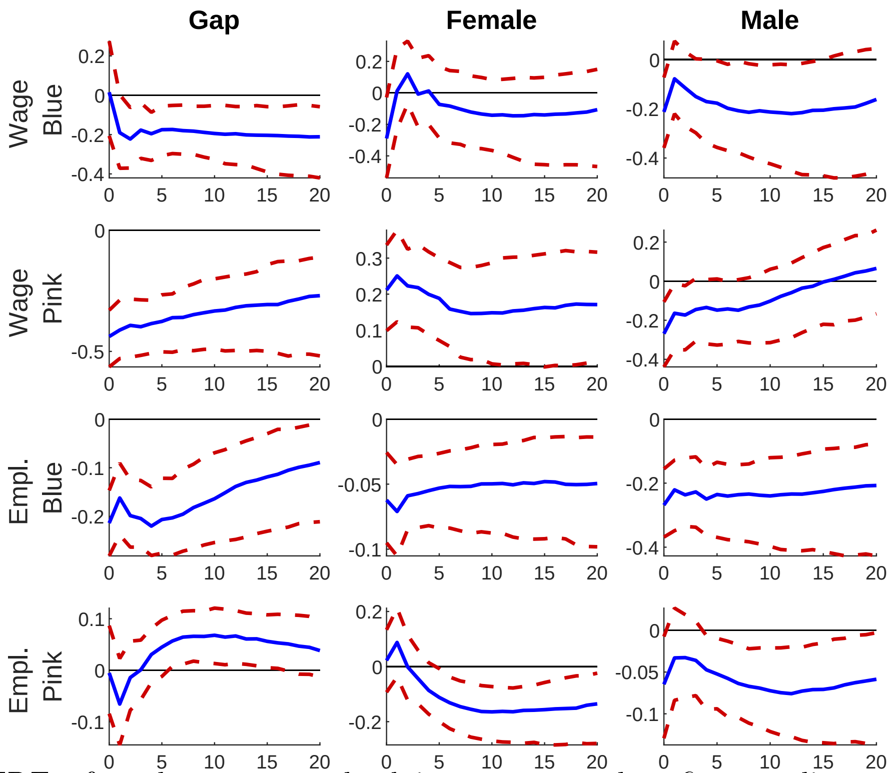

Figure 2.B.26: IRFs of gender gaps to a shock in government benefits spending - pink- vs. blue-collar occupations

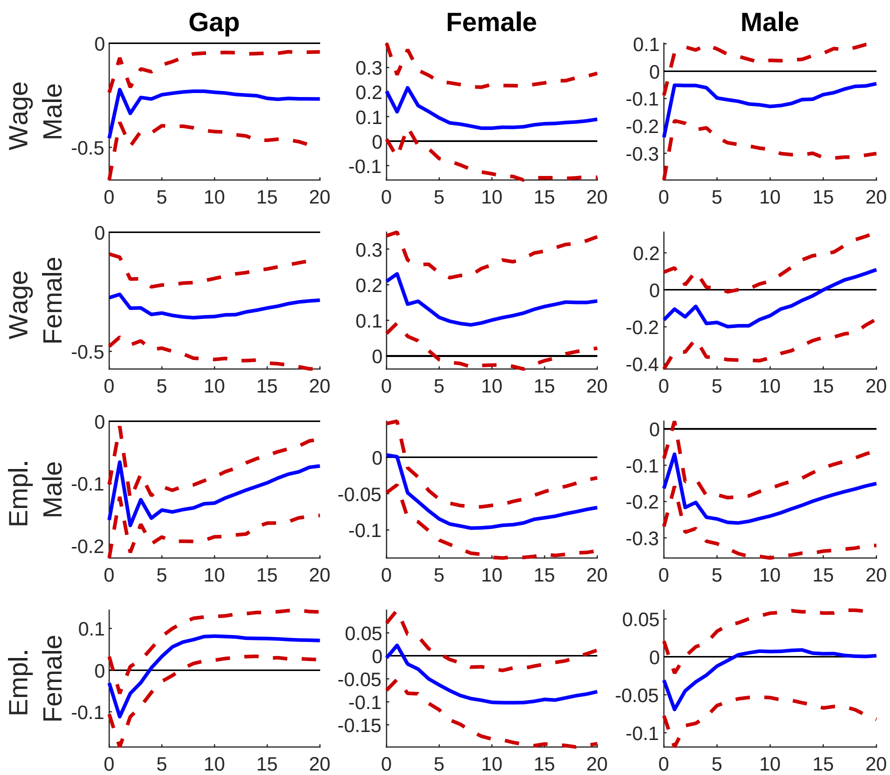

Figure 2.B.27: IRFs of gender gaps to a shock in government benefits spending - male vs. female dominant industry

Government Defense Expenditure

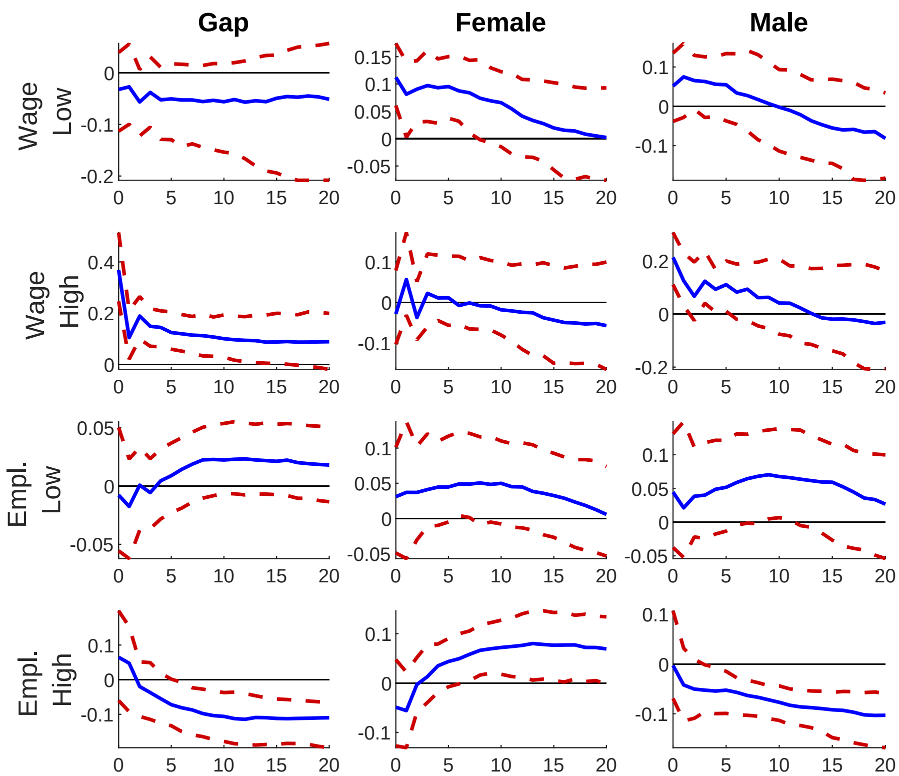

Figure 2.B.28: IRFs of gender gaps to a shock in government defense spending - no college degree (low) vs. college degree (high)

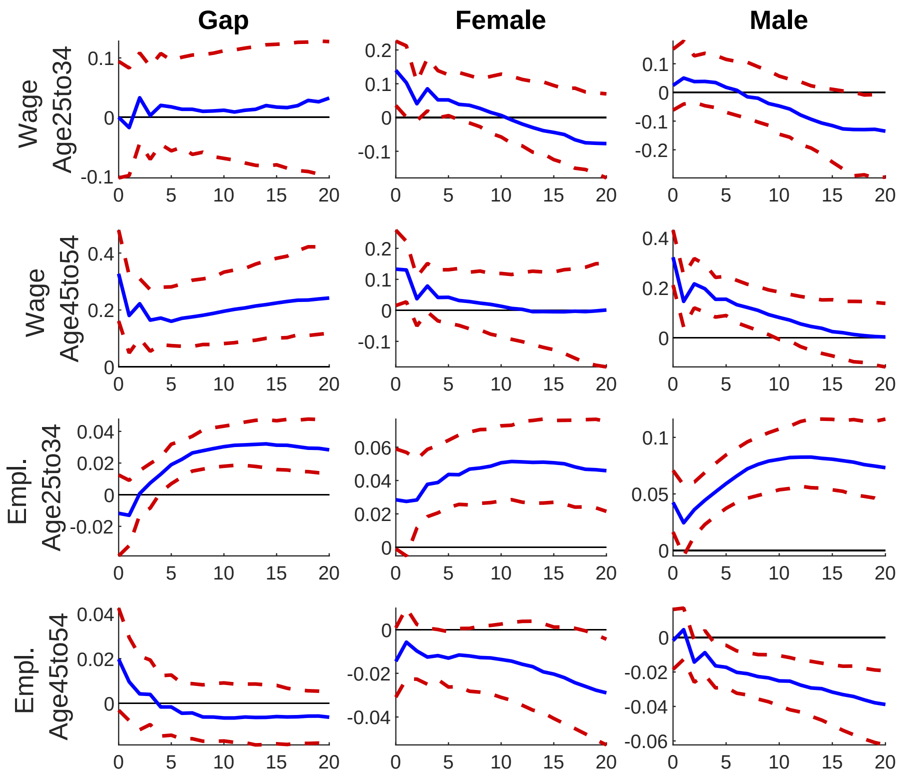

Figure 2.B.29: IRFs of gender gaps to a shock in government defense spending - workers aged 25-34 vs. 45-54

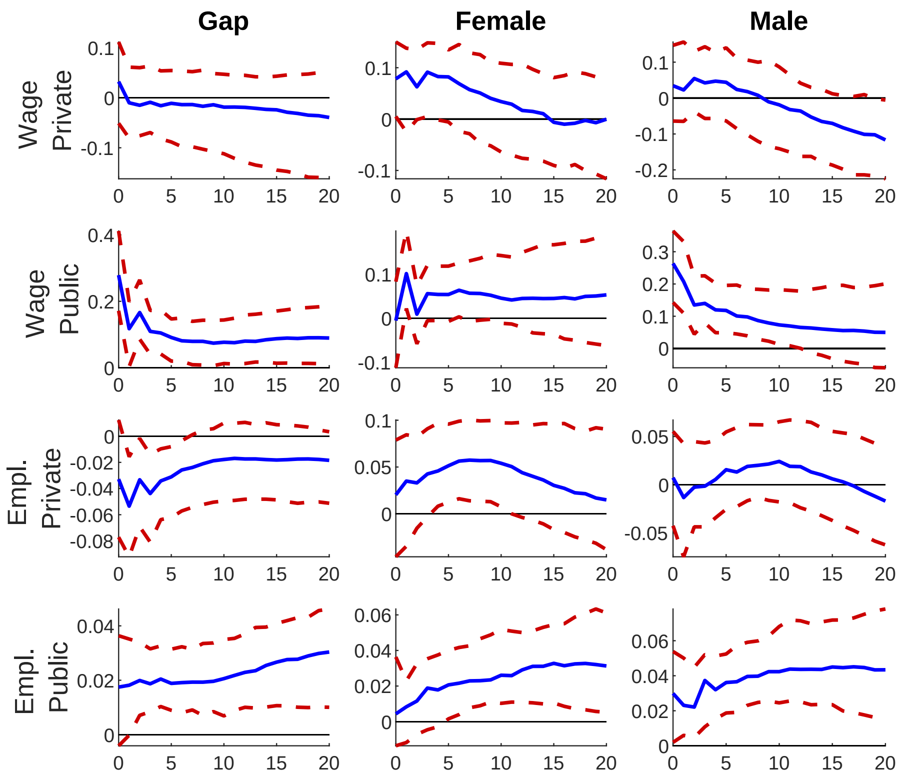

Figure 2.B.30: IRFs of gender gaps to a shock in government defense spending - public vs. private sector

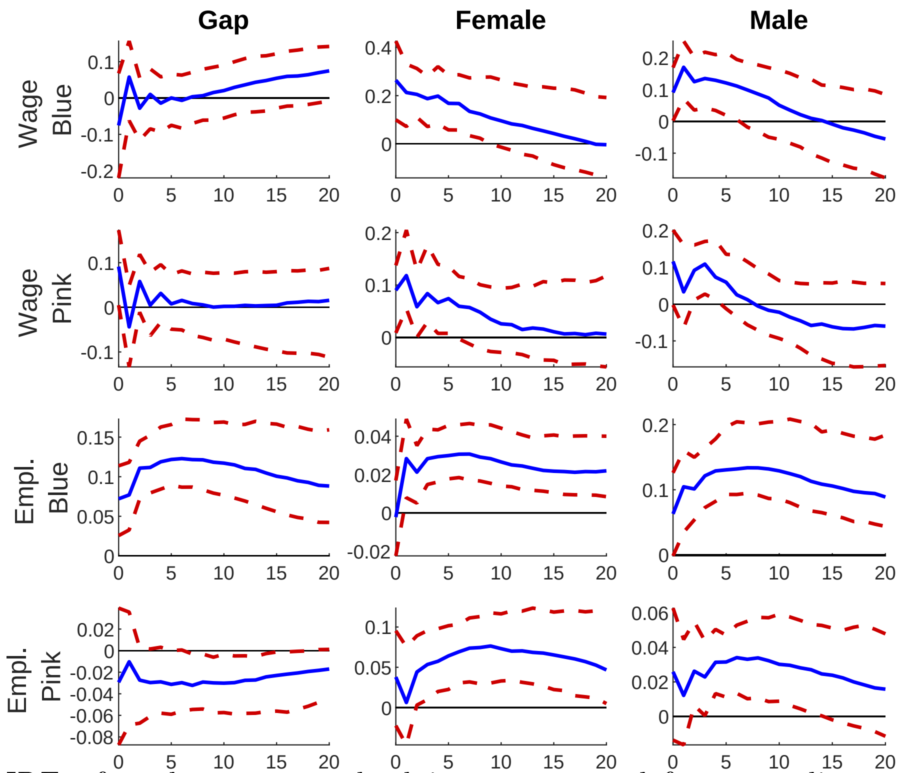

Figure 2.B.31: IRFs of gender gaps to a shock in government defense spending - pink- vs. blue-collar occupations

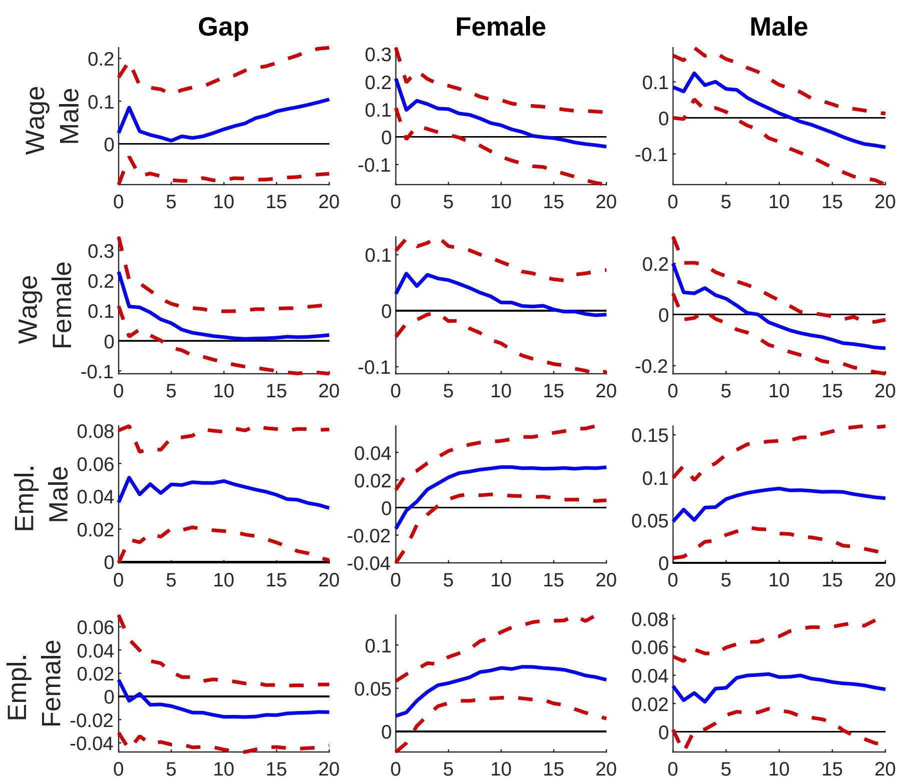

Figure 2.B.32: IRFs of gender gaps to a shock in government defense spending - male vs. female dominant industry

Alternative Identification Scheme

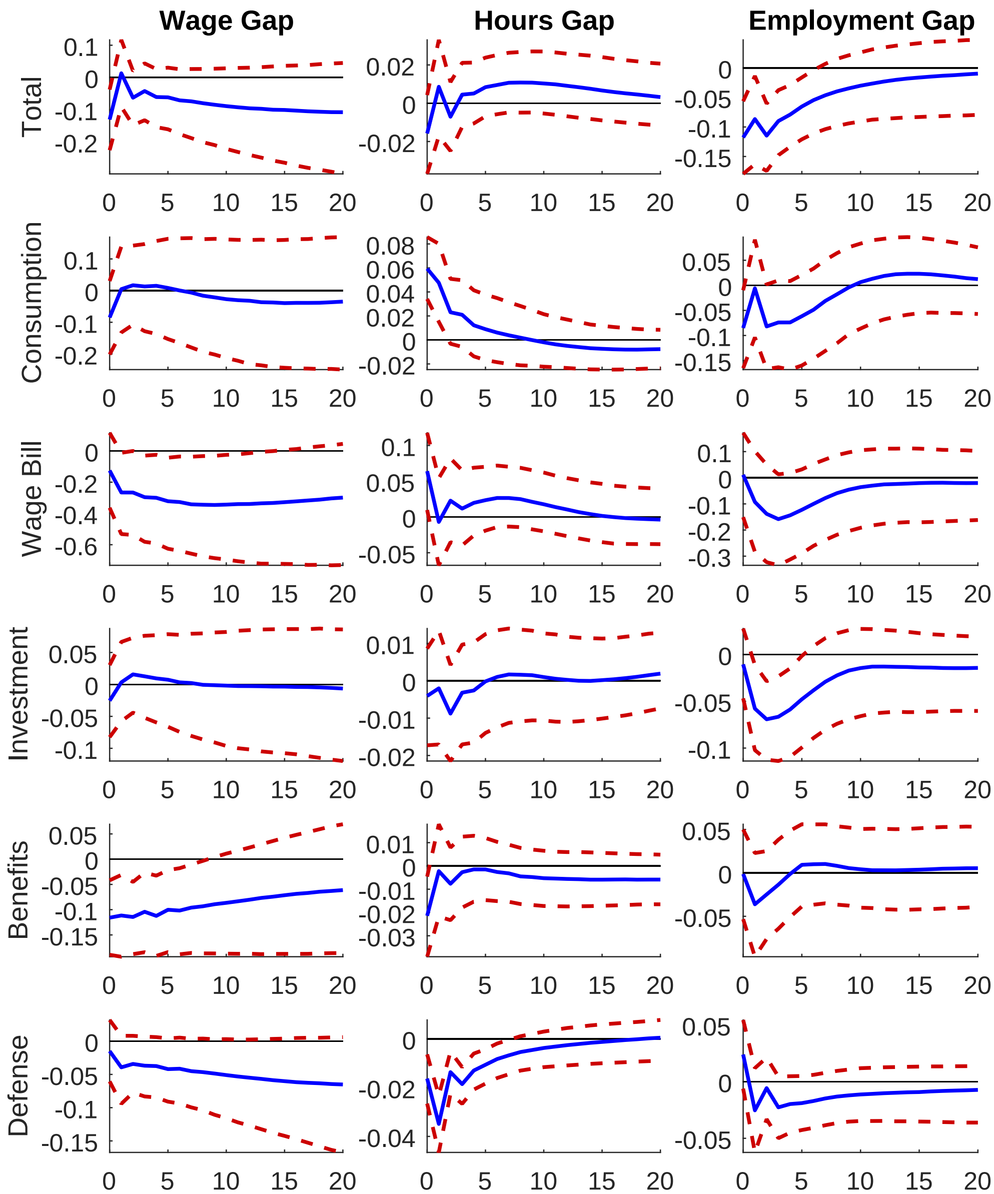

Figure 2.B.33: IRFs of gender gaps to government expenditures shocks - Cholesky identification

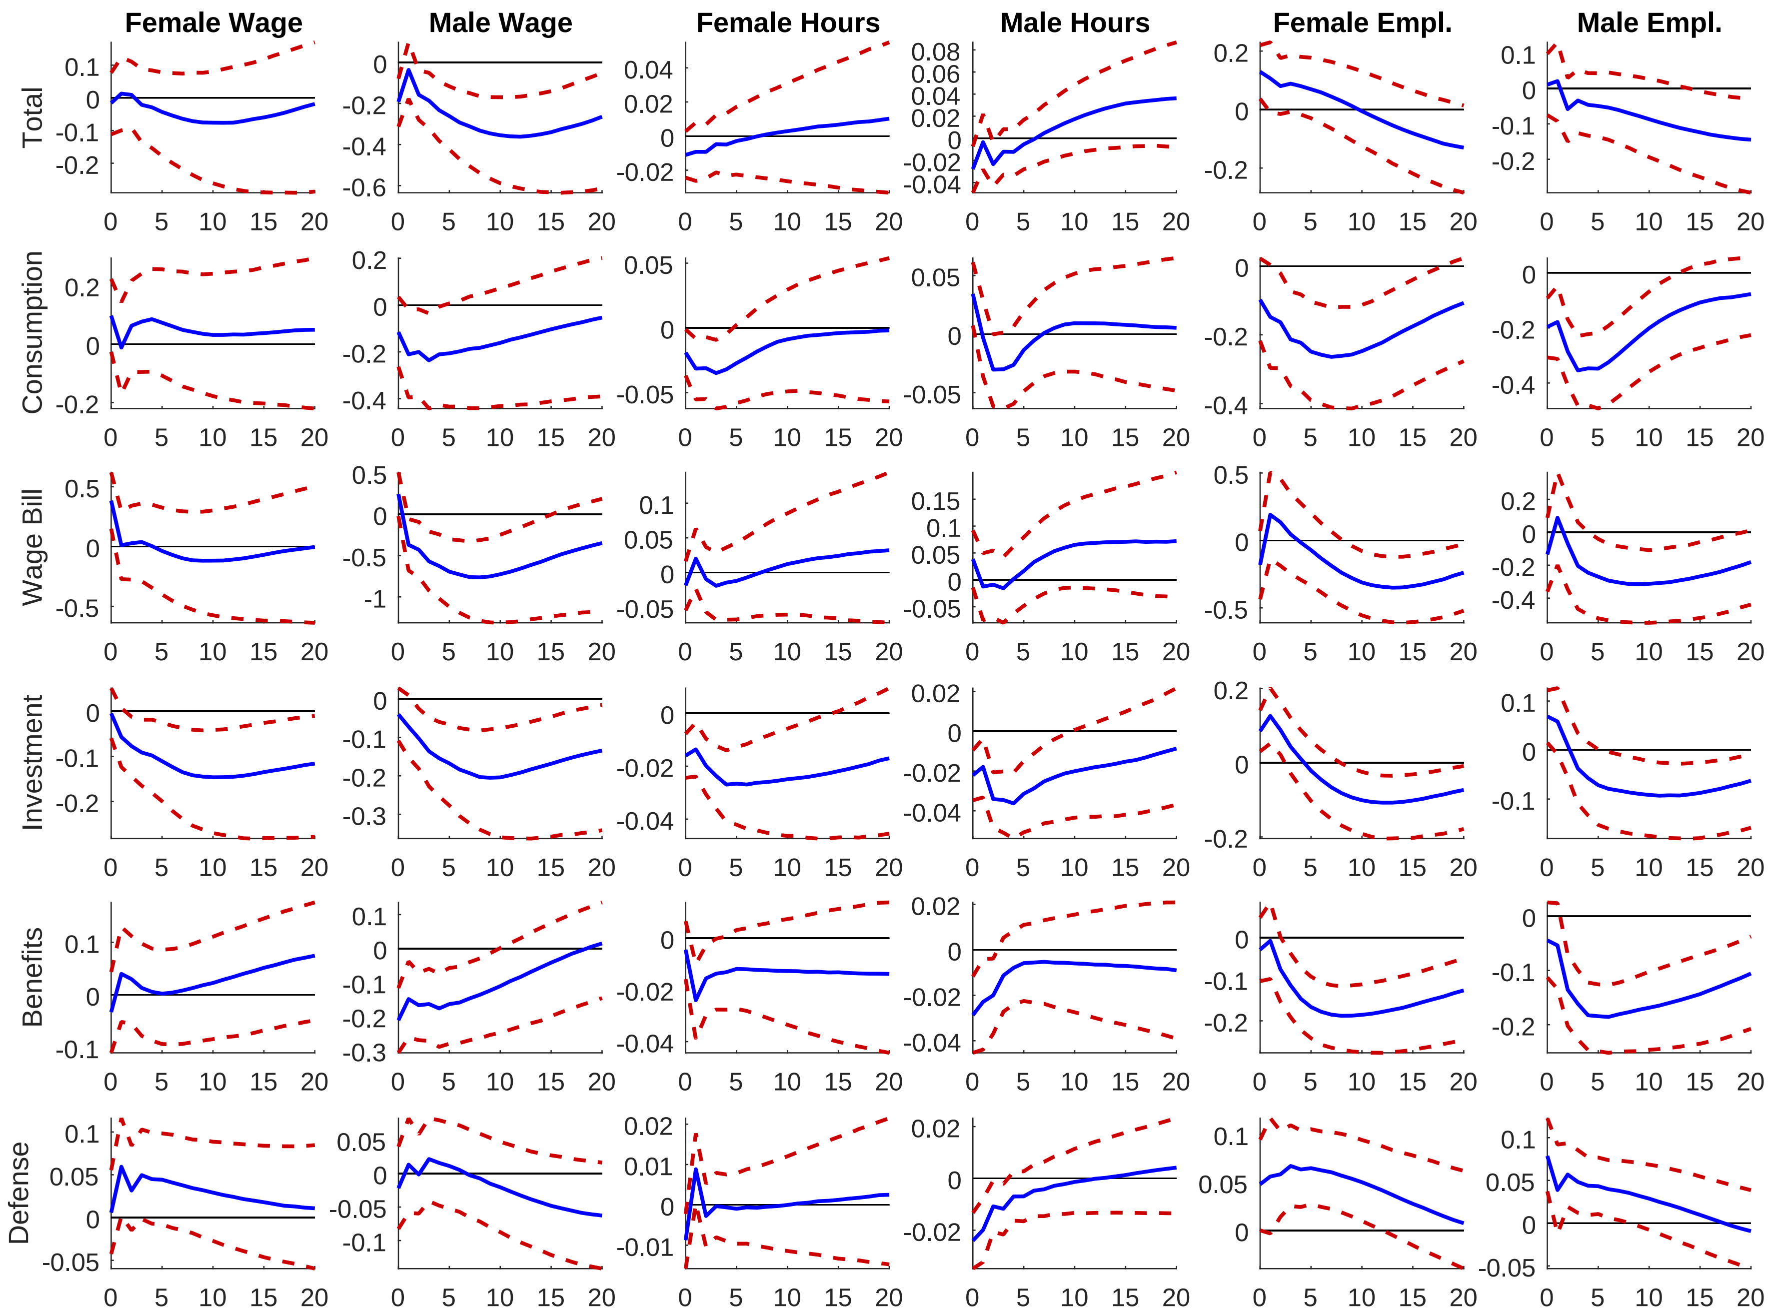

Figure 2.B.34: IRFs of male and female labor market variables to government expenditures shocks - Cholesky identification

Part-time workers

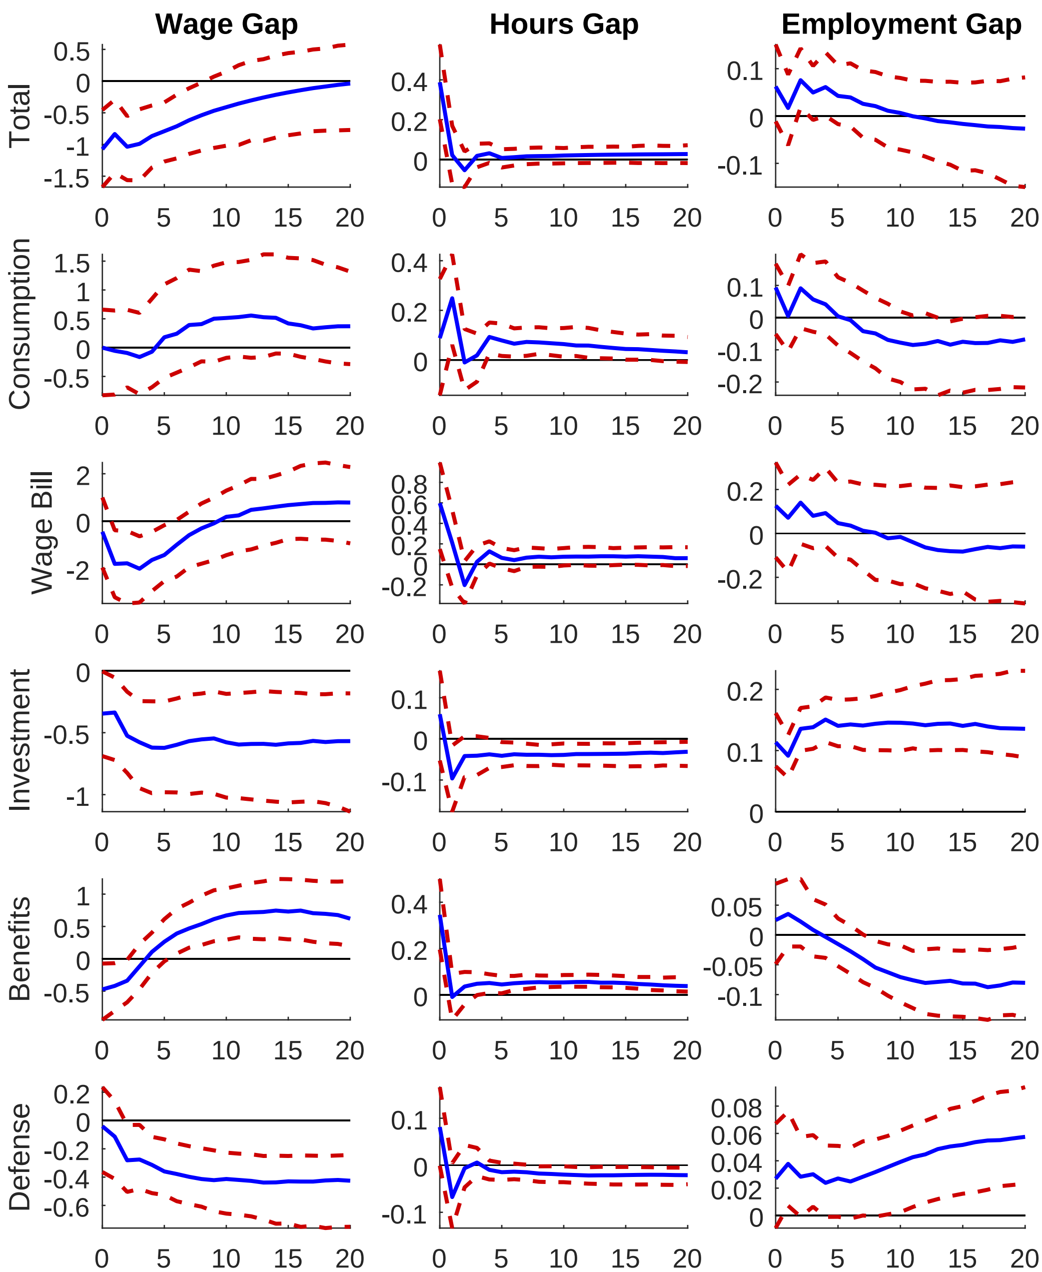

Figure 2.B.35: IRFs of gender gaps to shocks in different spending components - Part-time workers

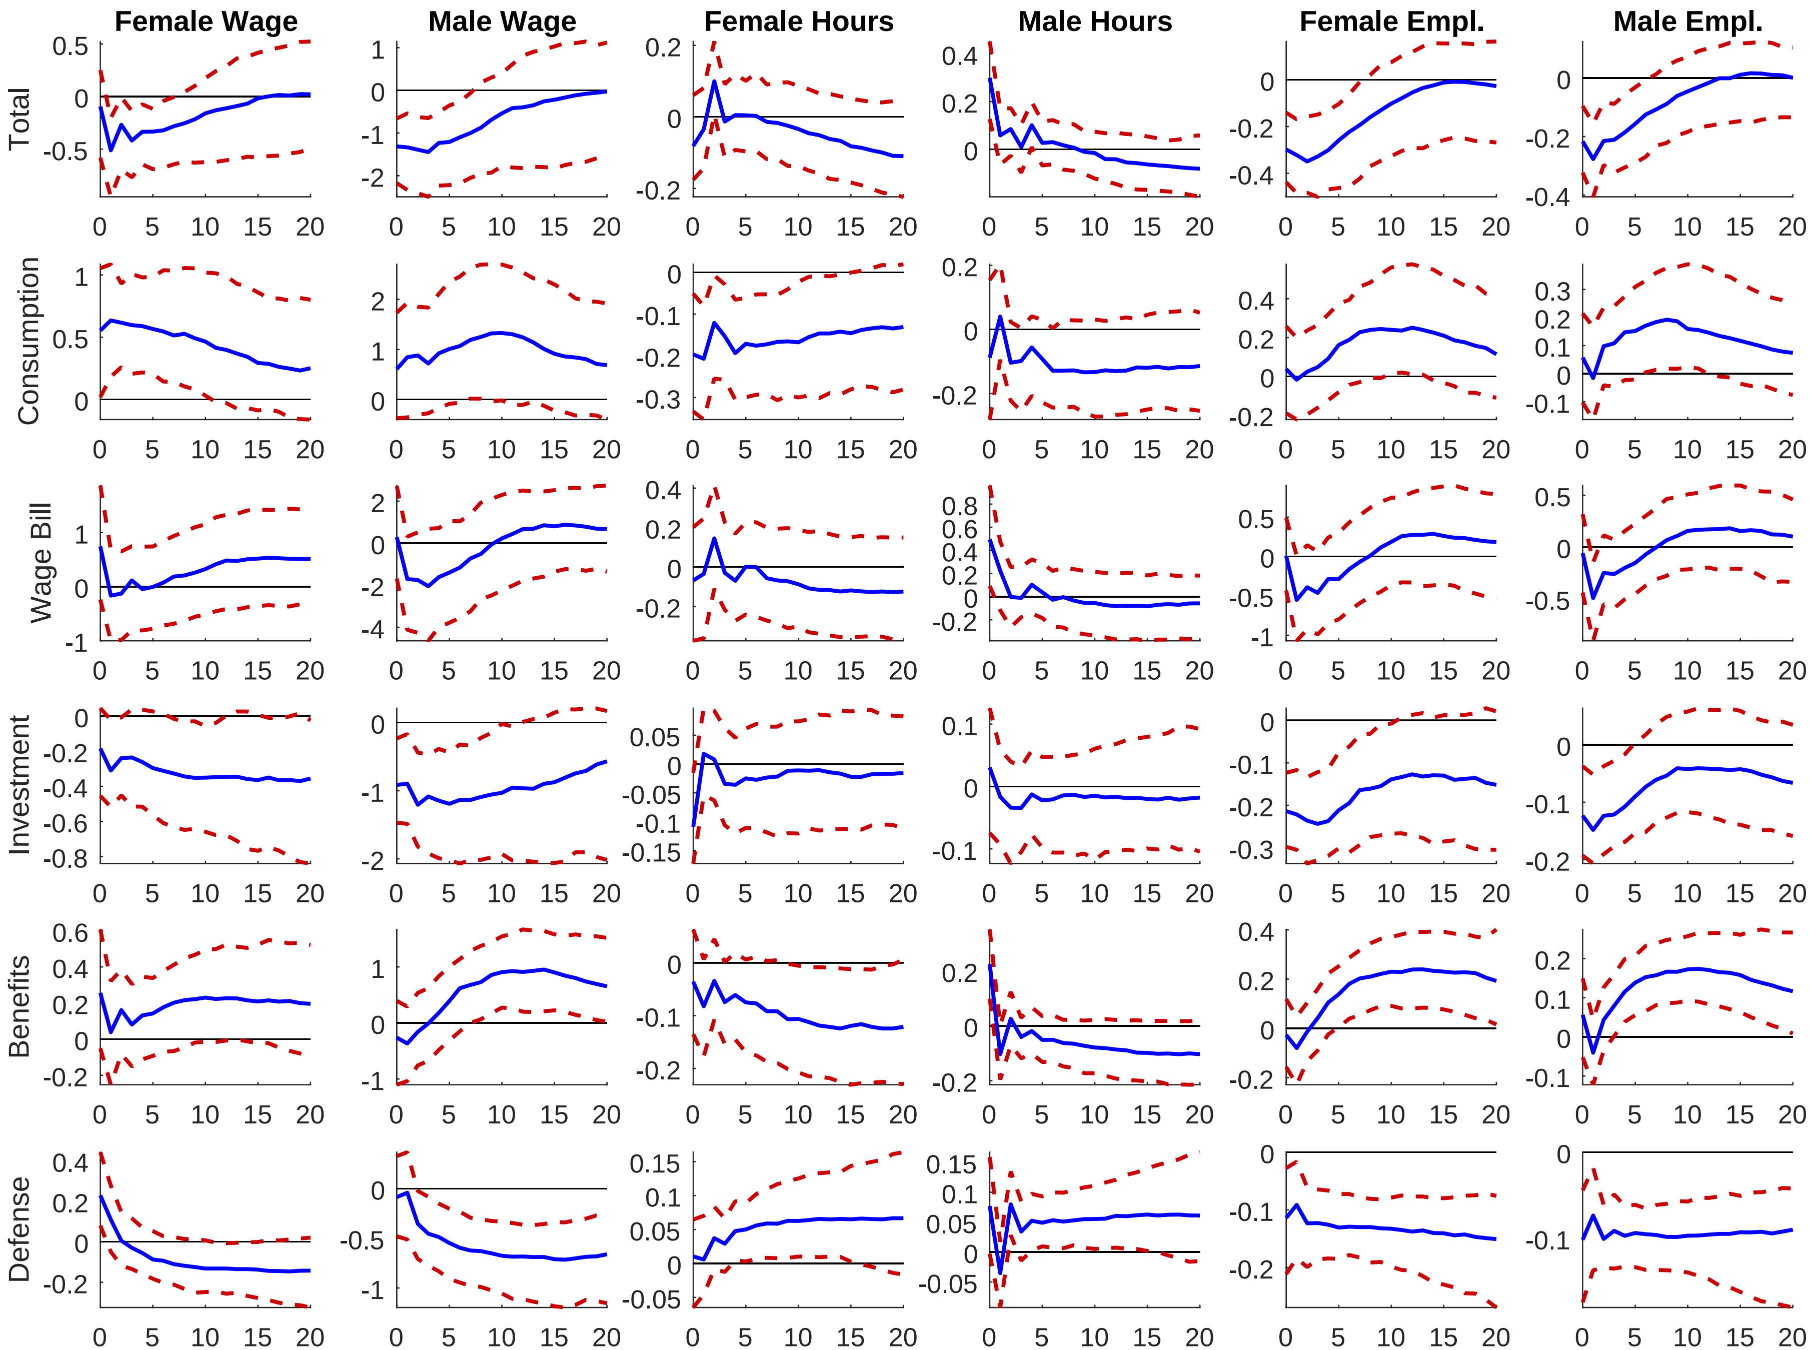

Figure 2.B.36: IRFs of labour market outcomes by gender to shocks in different spending components - Part-time workers

Unmarried workers

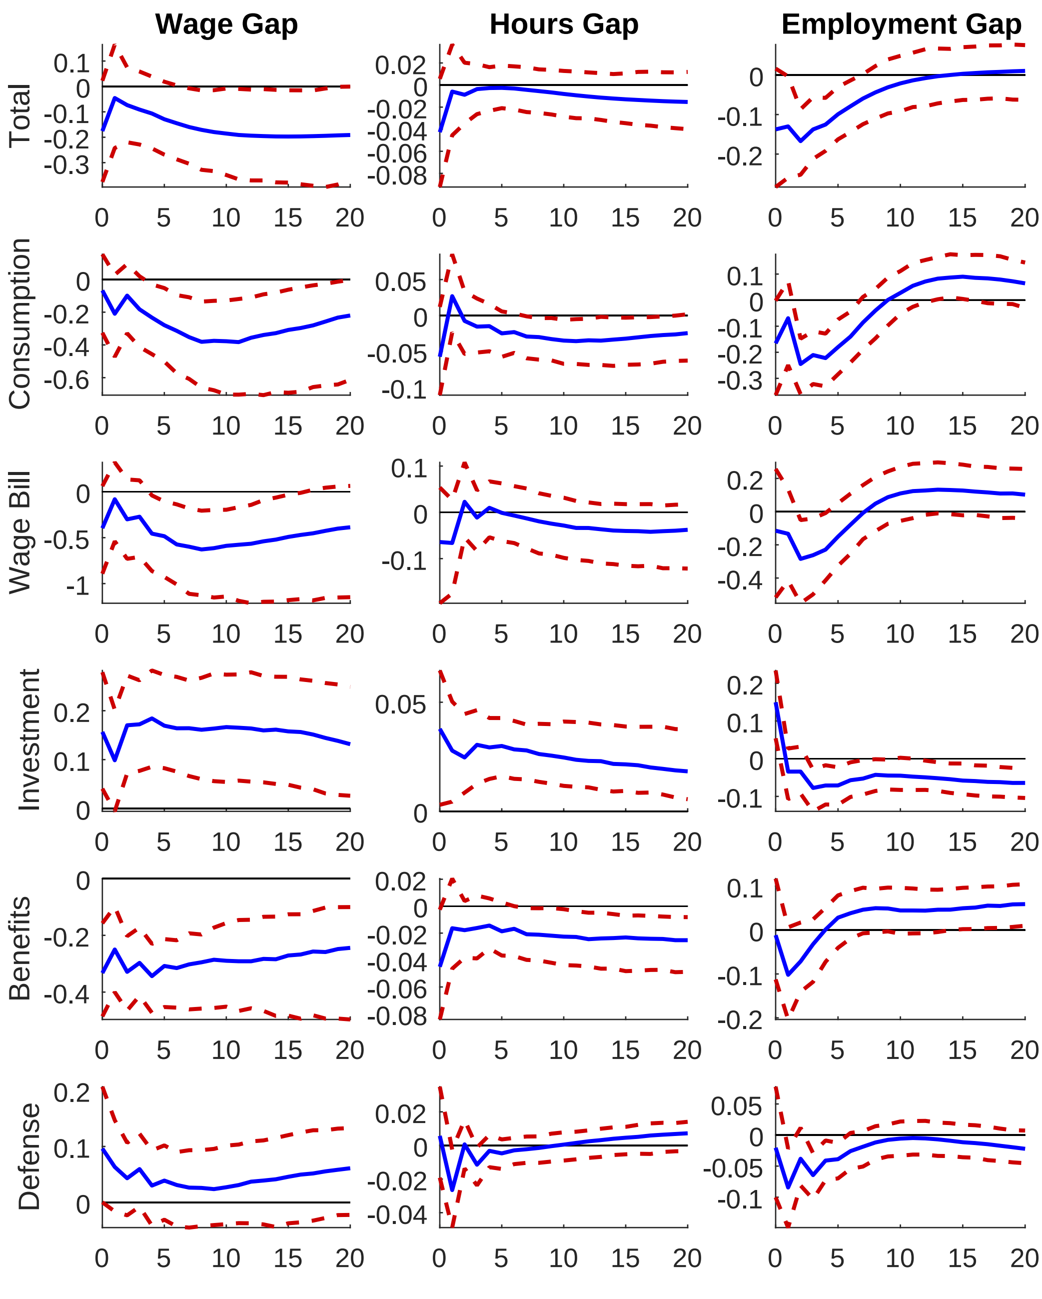

Figure 2.B.37: IRFs of gender gaps to shocks in different spending components - Unmarried workers

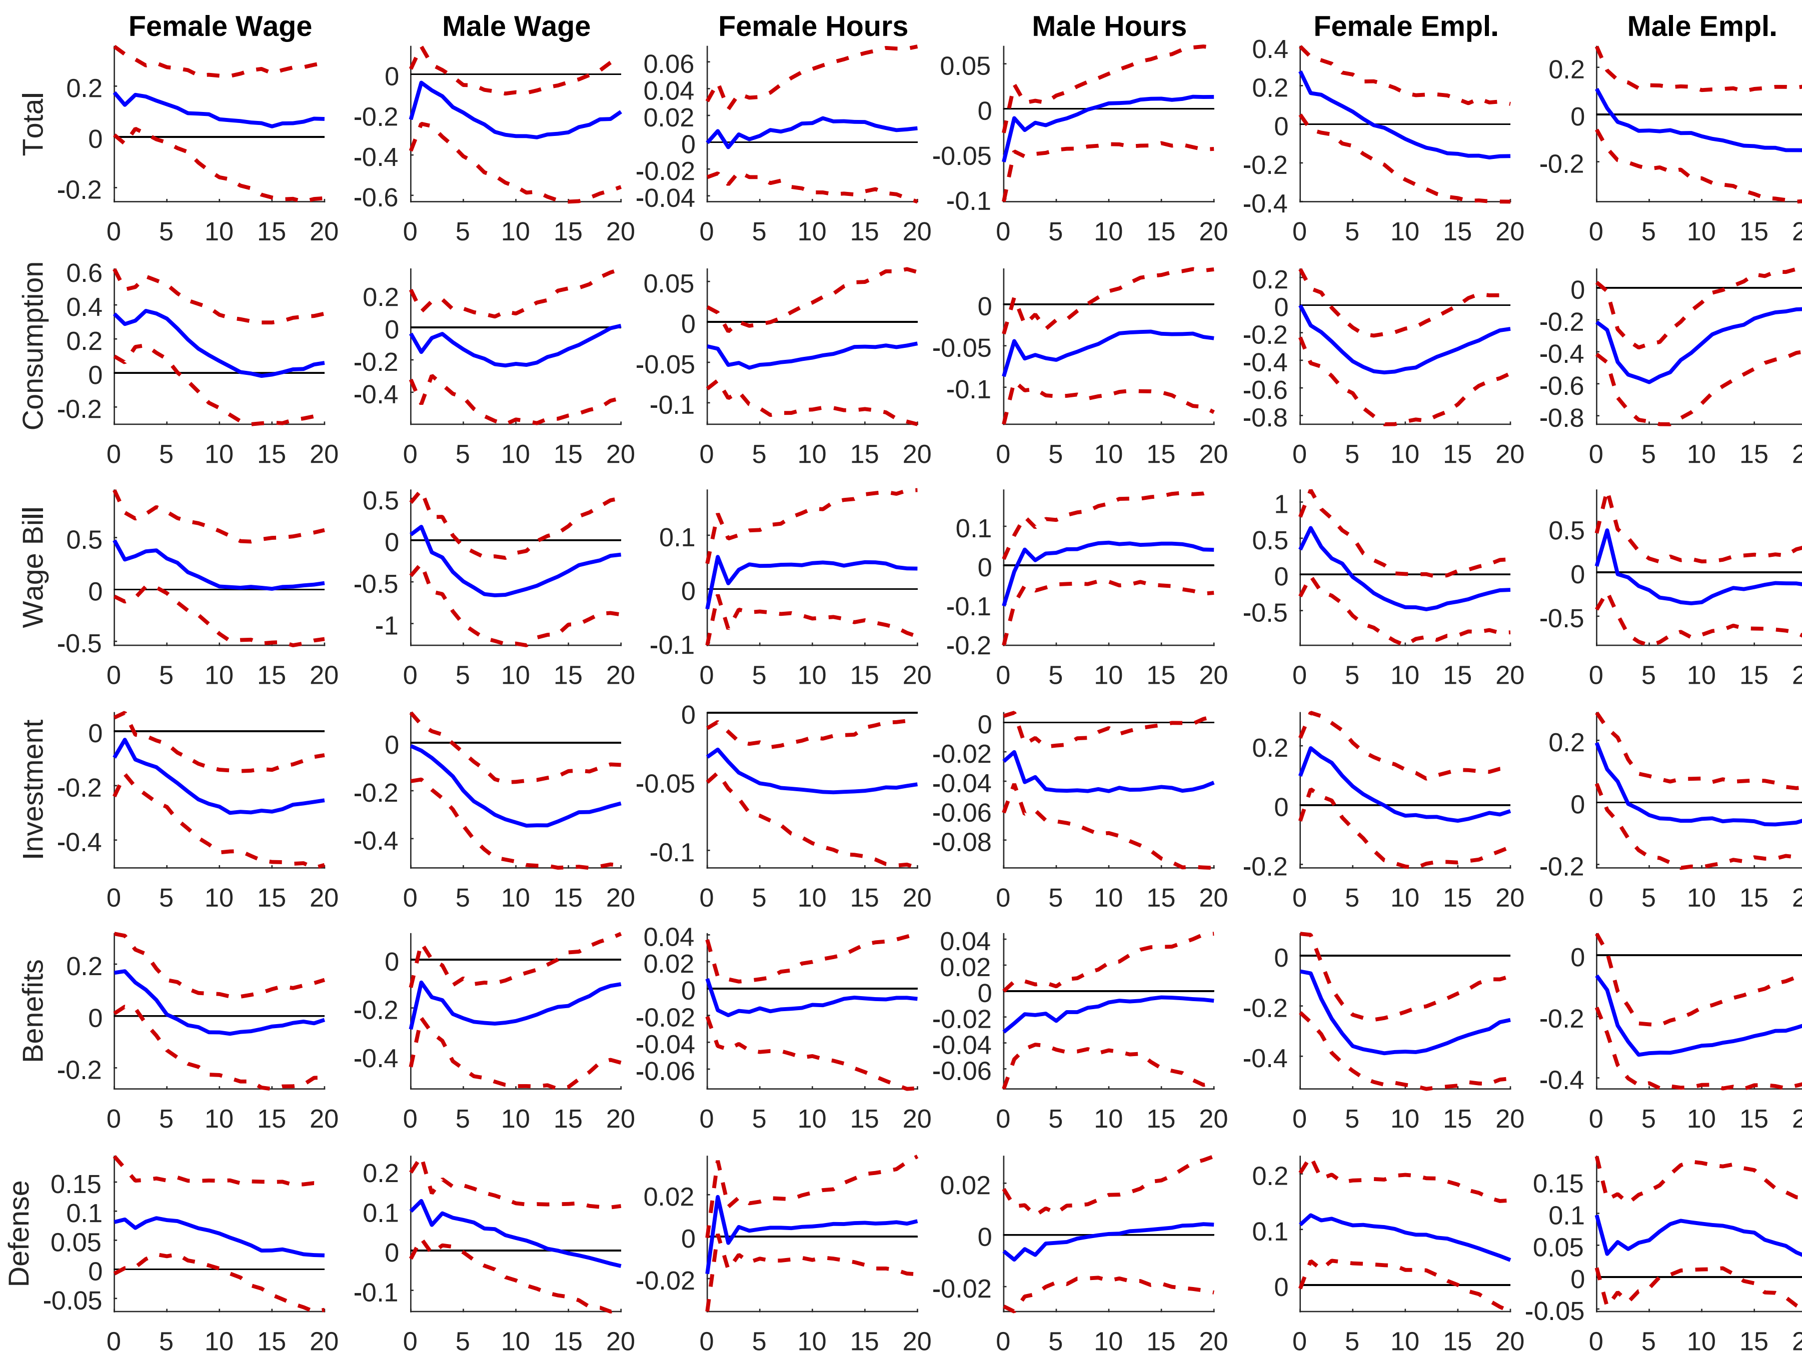

Figure 2.B.38: IRFs of labour market outcomes by gender to shocks in different spending components - Unmaried workers

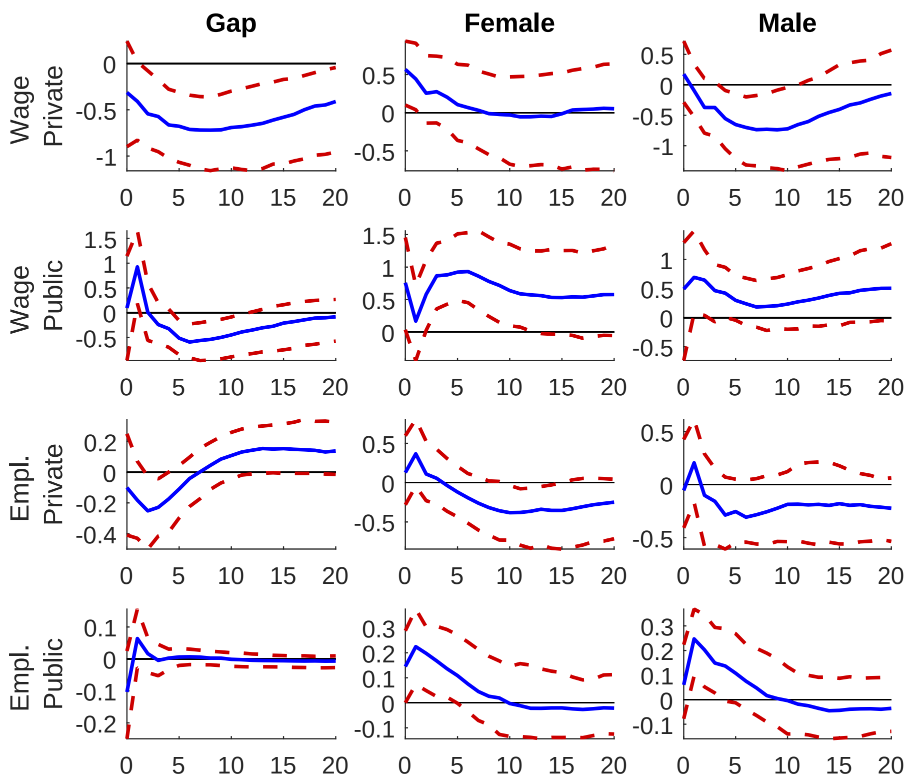

Figure 2.B.39: IRFs of gender gaps to a shock in government wage bill - public vs. private sector, unmarried workers

2.A Stylized Facts

2.A.1 Government spending components

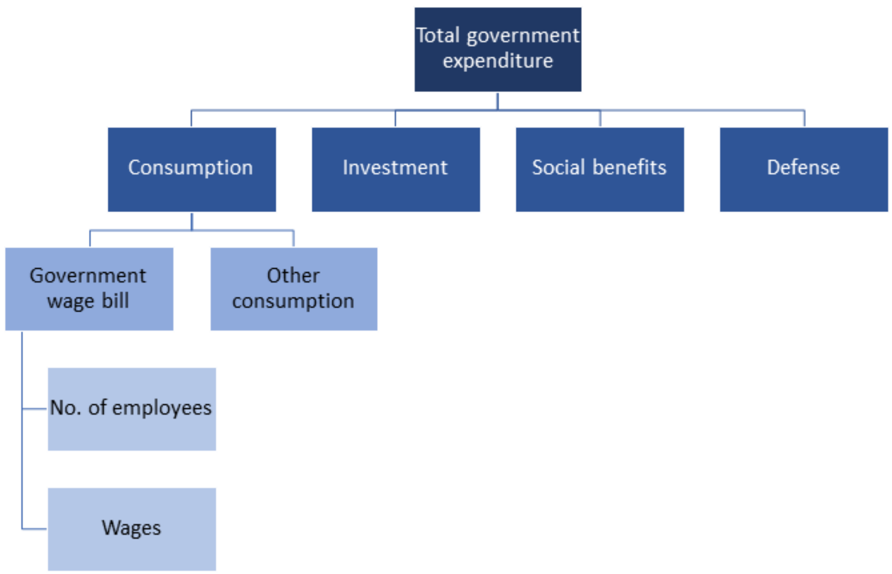

Figure 2.A.1: Hierarchical decomposition of government spending components

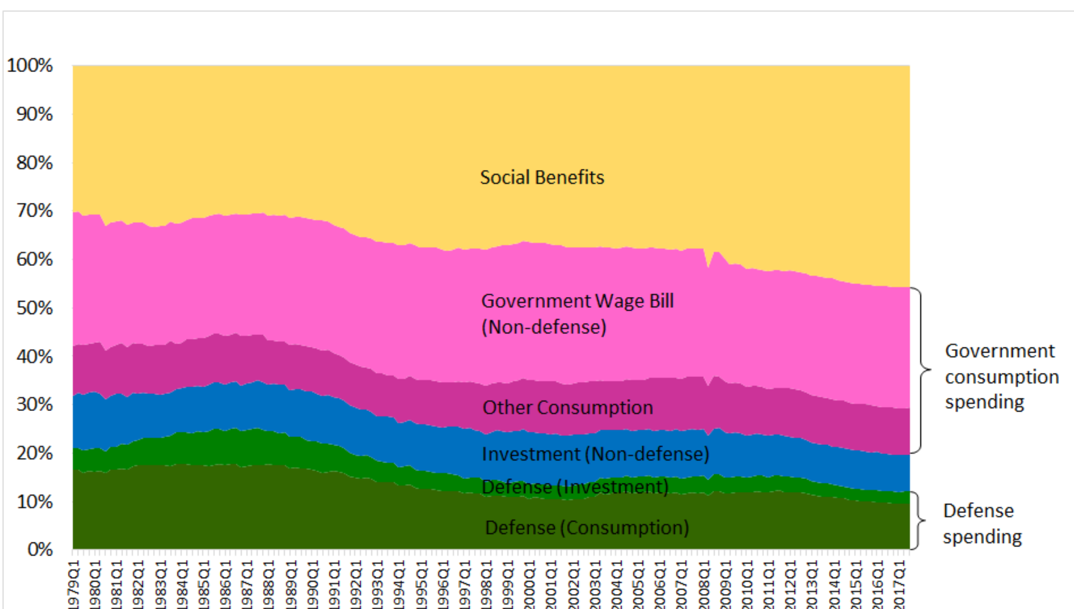

Figure 2.A.2: Historical evolution of government spending components

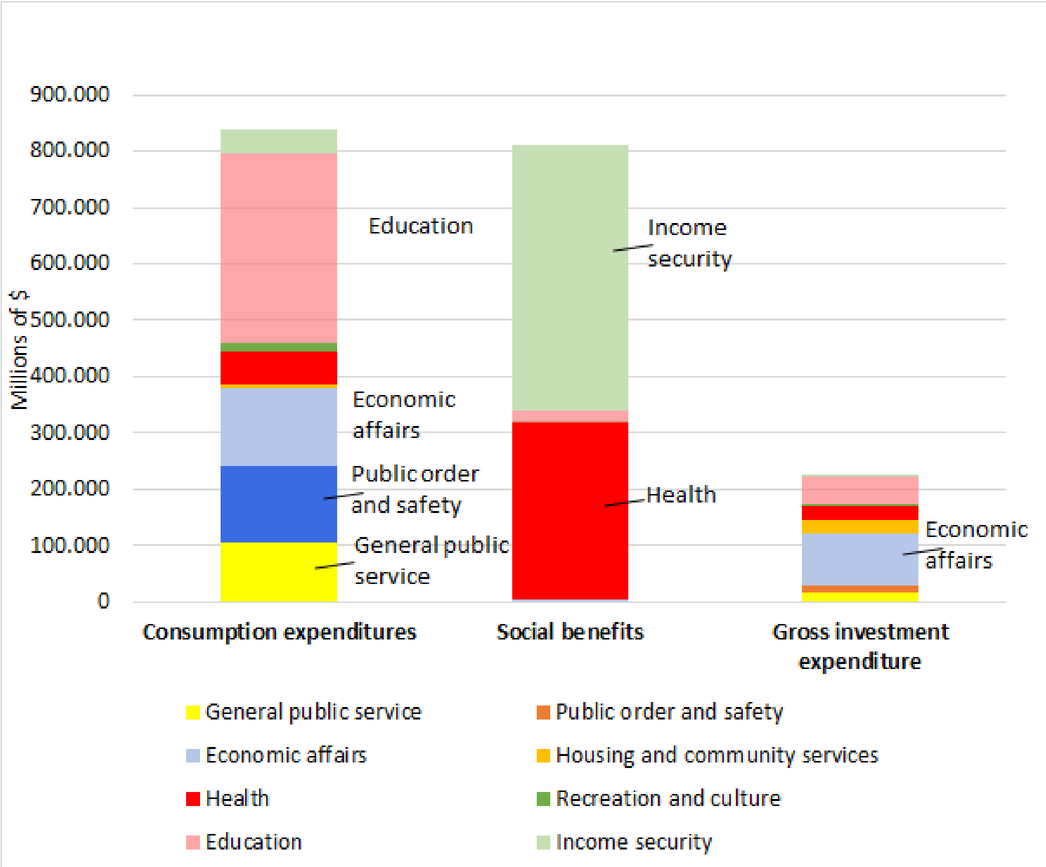

Figure 2.A.3: Selected spending components by function. Averages over 1979-2007. Source: BEA, NIPA Table 3.17. Note: General public services include tax collection and financial management, interest payments, executive and legislative spending; Public order and safety includes police, fire, law courts, prisons; Economic affairs includes mainly transportation and space; Education includes mostly elementary and secondary; Income security mostly consists of social insurance funds, including old age, survivors, disability insurance (social security), and railroad retirement and excludes government employee retirement plans.

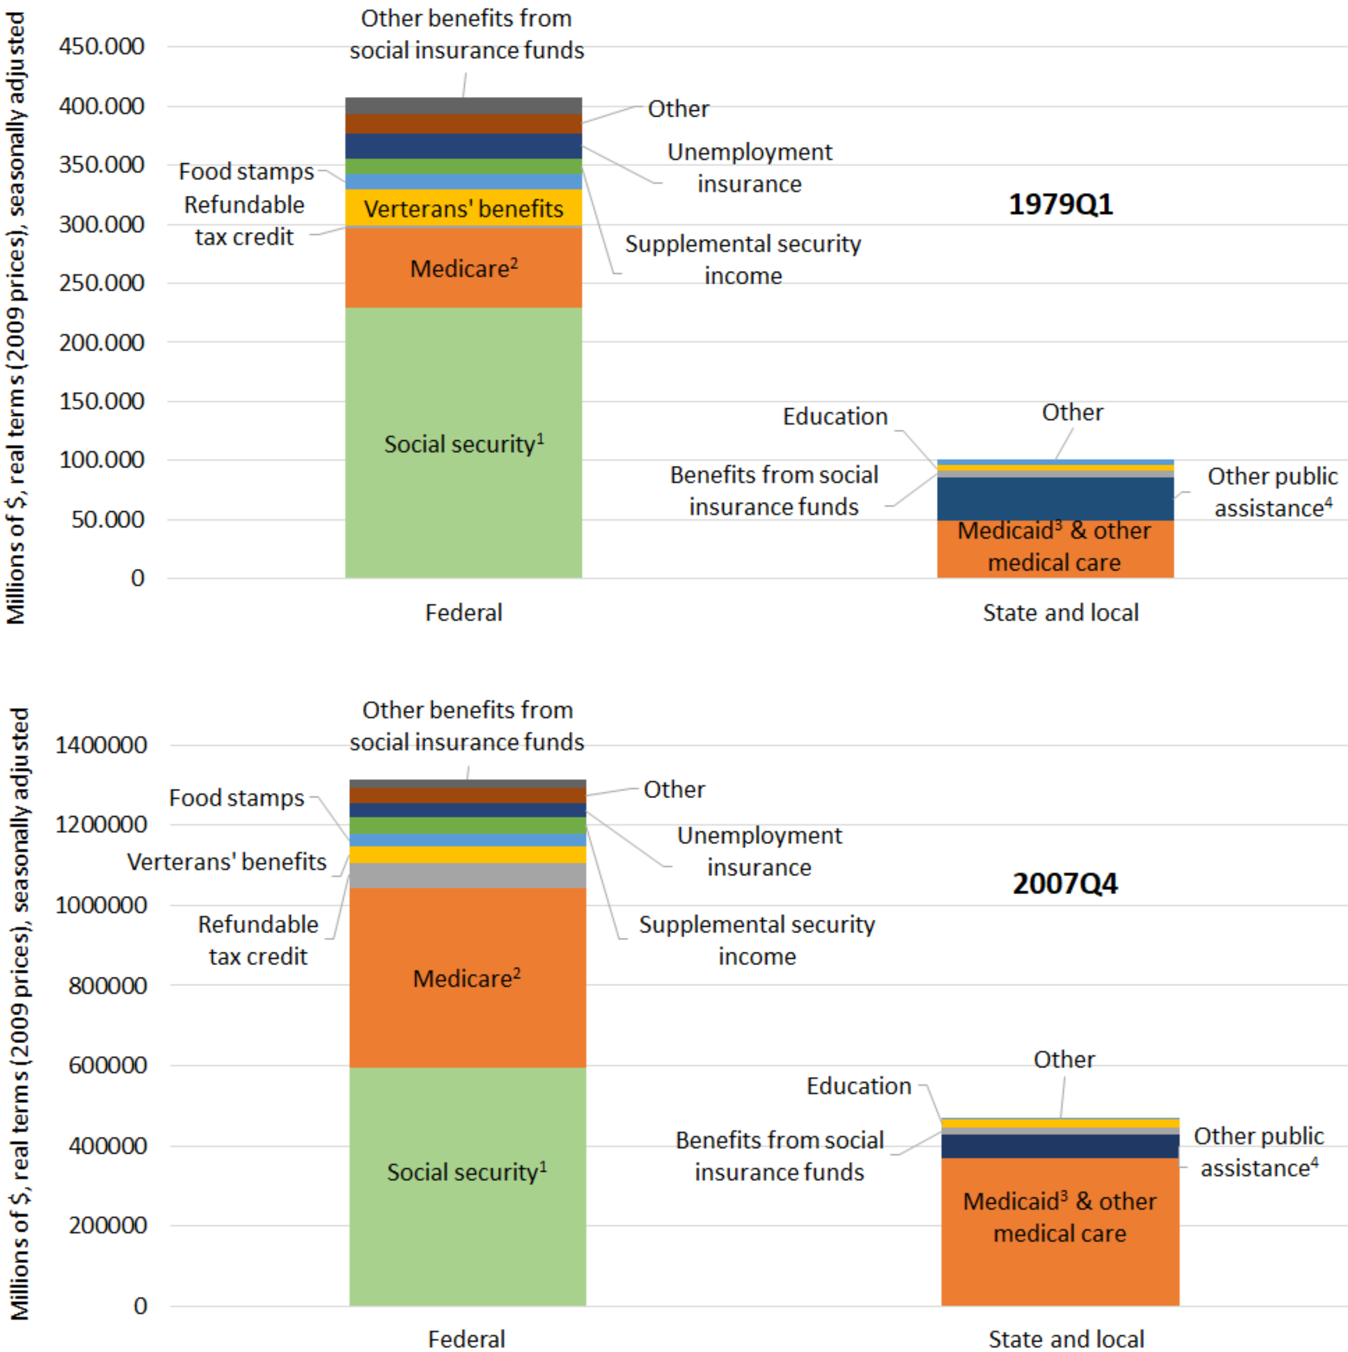

[1] includes old-age, survivors, and disability insurance benefits that are distributed from the federal old-age and survivors insurance trust fund and the disability insurance trust fund.
[2] includes hospital and supplementary medical insurance benefits that are distributed from the federal hospital insurance trust fund and the supplementary medical insurance trust fund (to those over 64 or those with severe disabilities).
[3] provides health coverage for those with a low income.
[4] includes supplemental security income, energy assistance, family and general assistance.

Figure 2.A.4: Decomposition of **social benefits** by government level for **1979Q1** and **2007Q4**

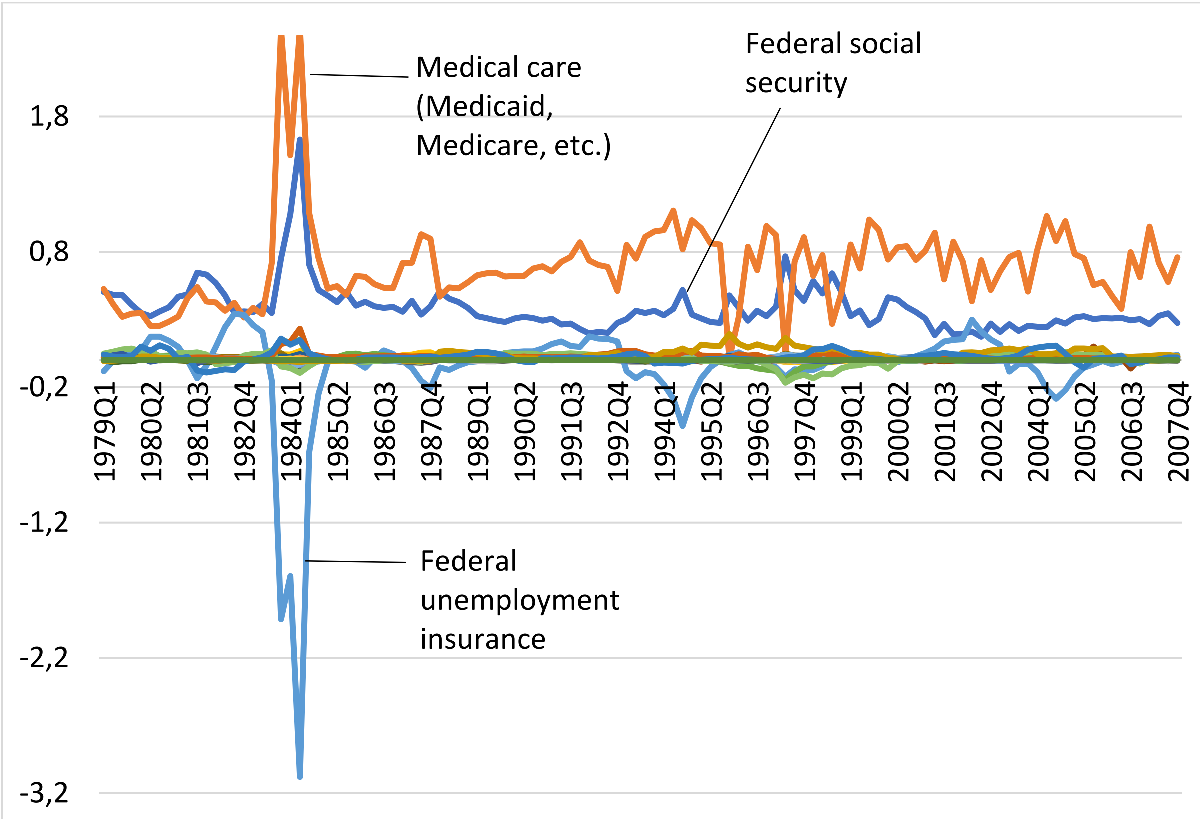

Figure 2.A.5: Quarter-to-quarter growth in benefit components relative to growth in total social benefits. Source: BEA, NIPA Table 3.12, authors' calculations.

2.A.2 Female shares across occupations

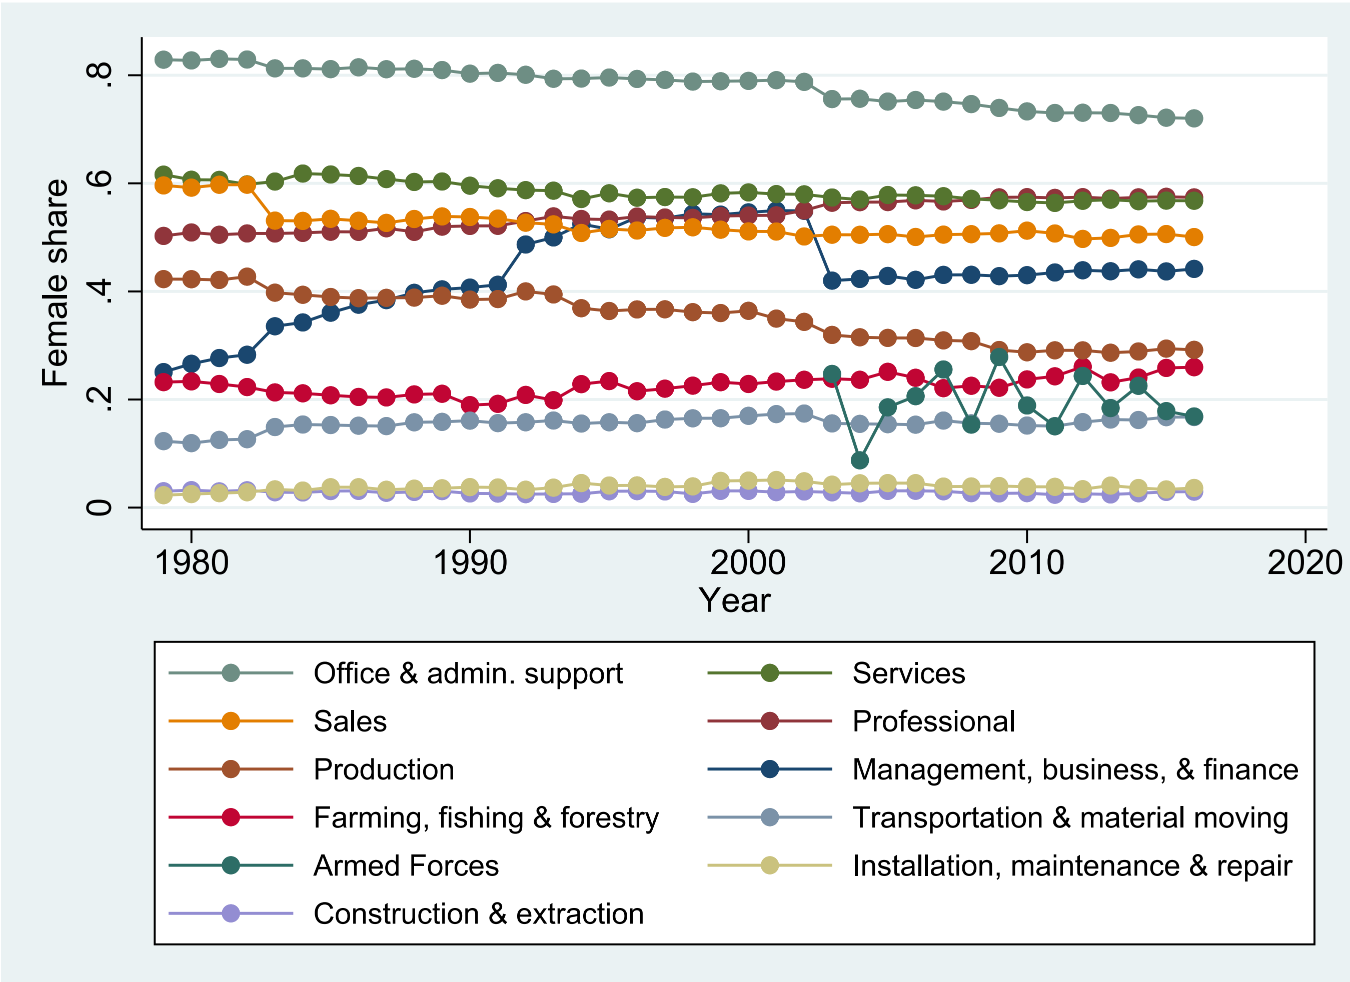

Figure 2.A.6: Female shares across occupations. Source: Own calculations based on CEPR-CPS.

2.A.3 Comparing men and women in the public versus the private sector

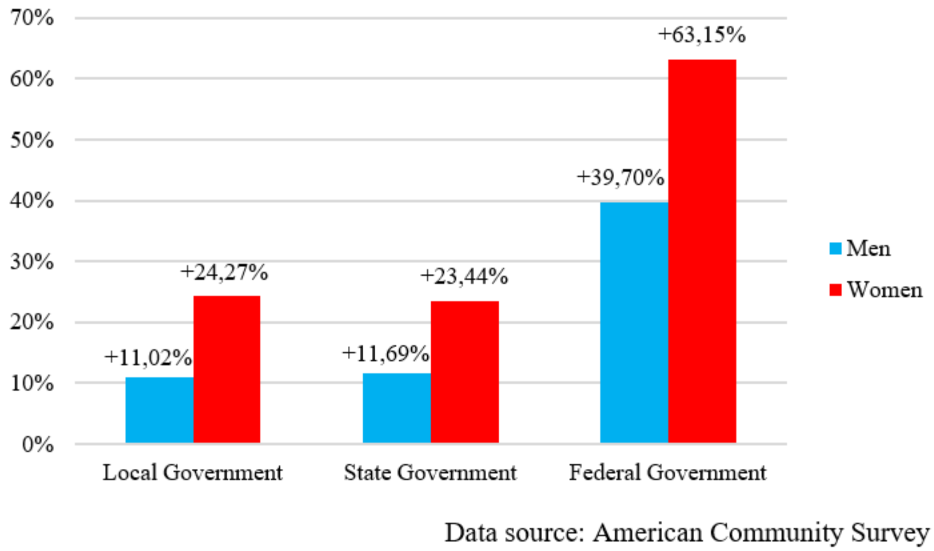

Figure 2.A.7: Public sector wage premium compared to for-profit private sector (median full-time earnings) in 2015

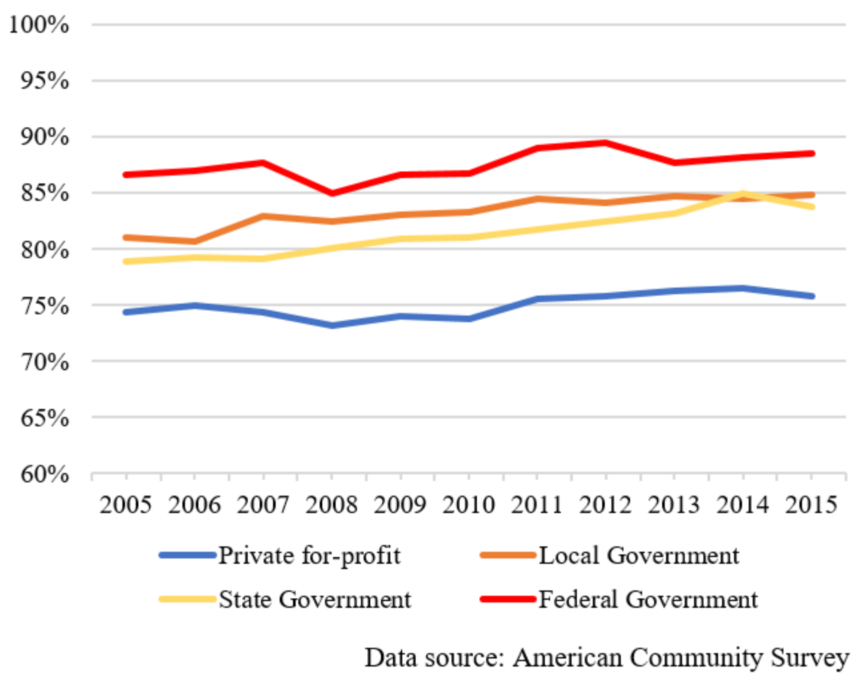

Figure 2.A.8: Female-to-male full-time median earnings ratio by sector

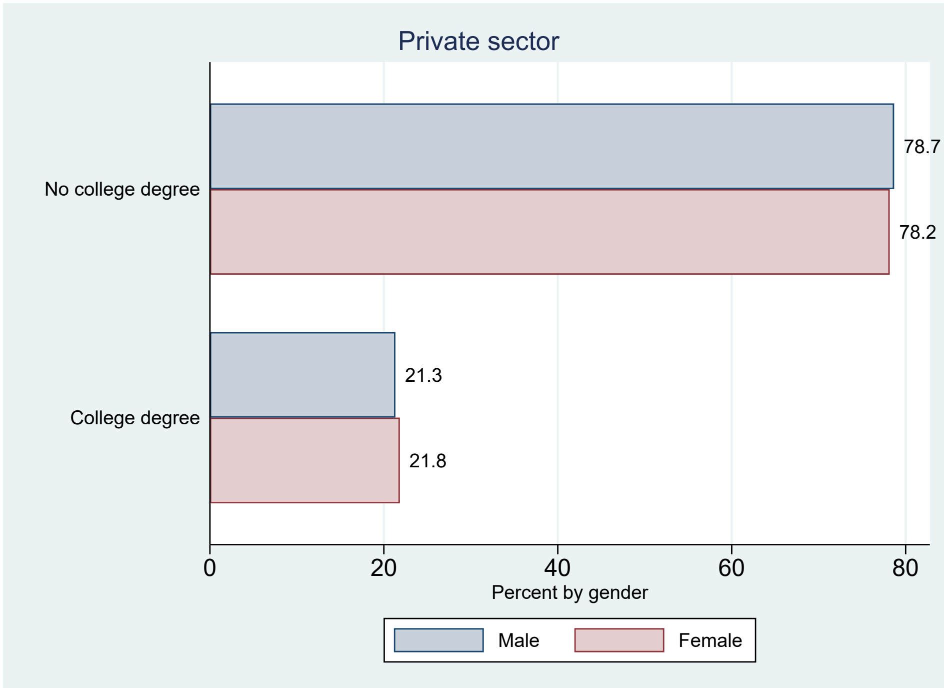

Figure 2.A.9: Educational attainment of private sector employees by gender (averages for the period 2003-2007). Source: CEPR-CPS, own calculations

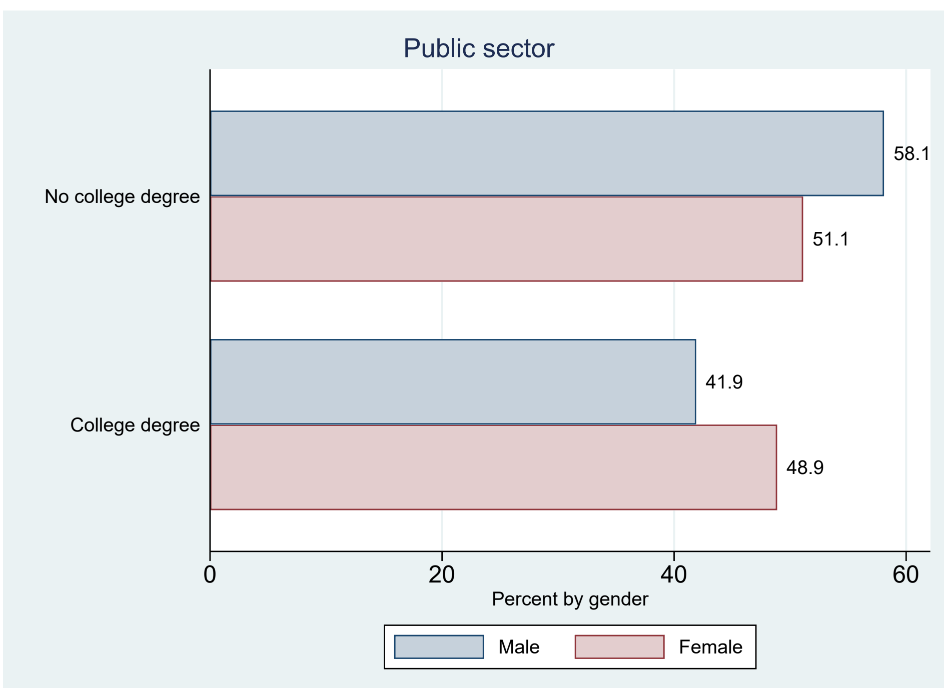

Figure 2.A.10: Educational attainment of public sector employees by gender (averages for the period 2003-2007). Source: CEPR-CPS, own calculations

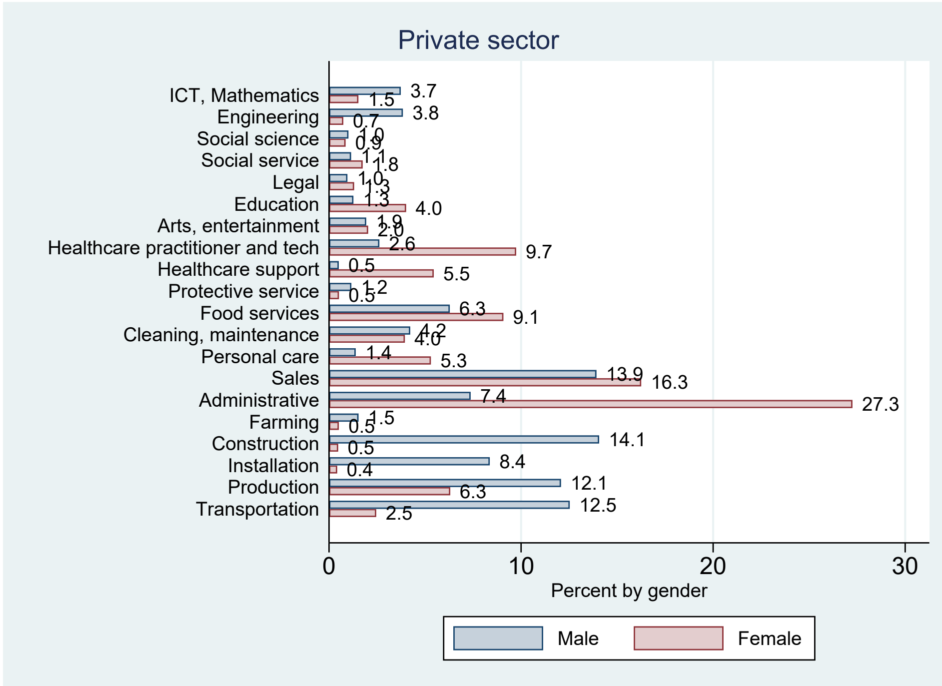

Figure 2.A.11: Distribution of men and women across private sector occupations (averages for the period 2003-2007). Source: CEPR-CPS, own calculations

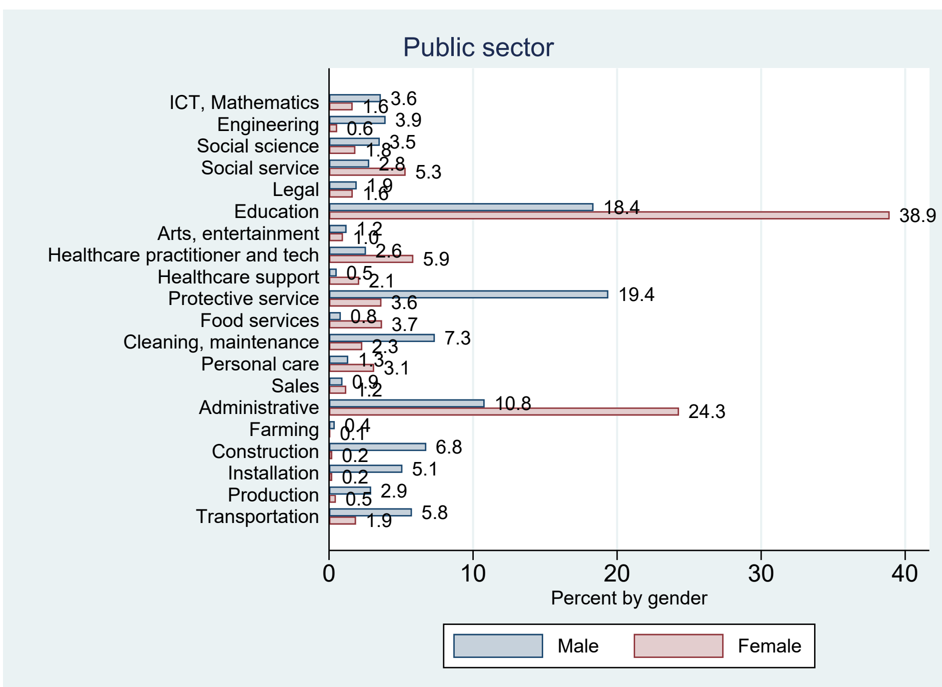

Figure 2.A.12: Distribution of men and women across public sector occupations (averages for the period 2003-2007). Source: CEPR-CPS, own calculations

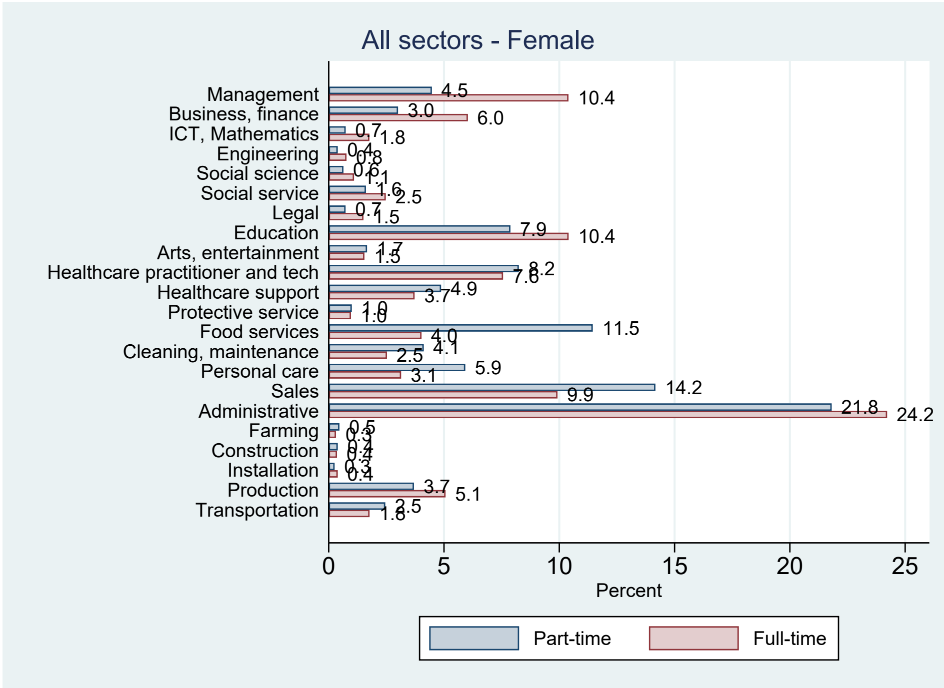

Figure 2.A.13: Distribution of full- and part-time women across occupations (averages for the period 2003-2007). Source: CEPR-CPS, own calculations

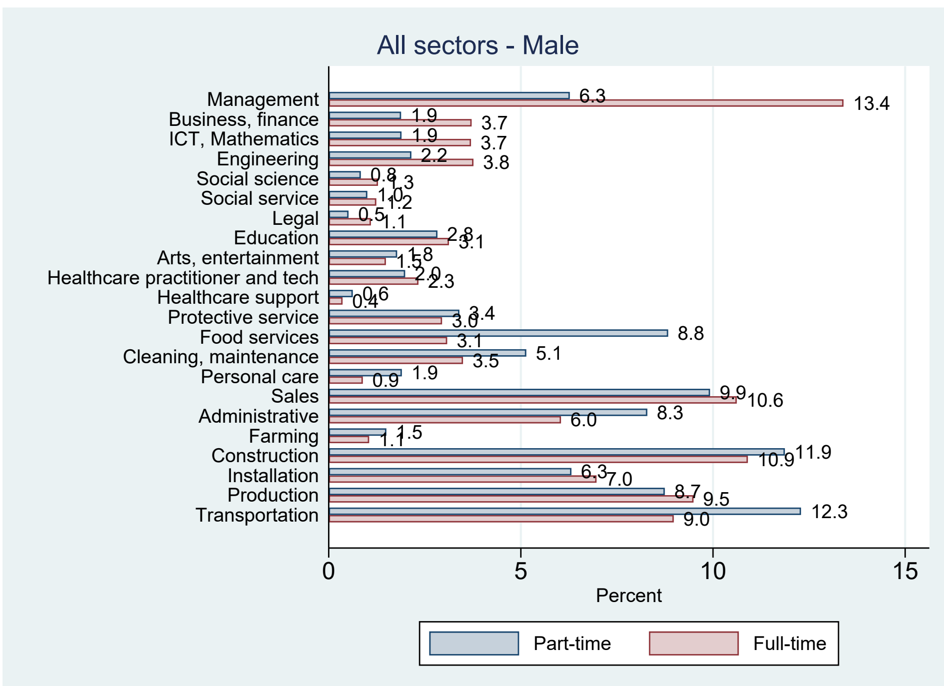

Figure 2.A.14: Distribution of full- and part-time men across occupations (averages for the period 2003-2007). Source: CEPR-CPS, own calculations

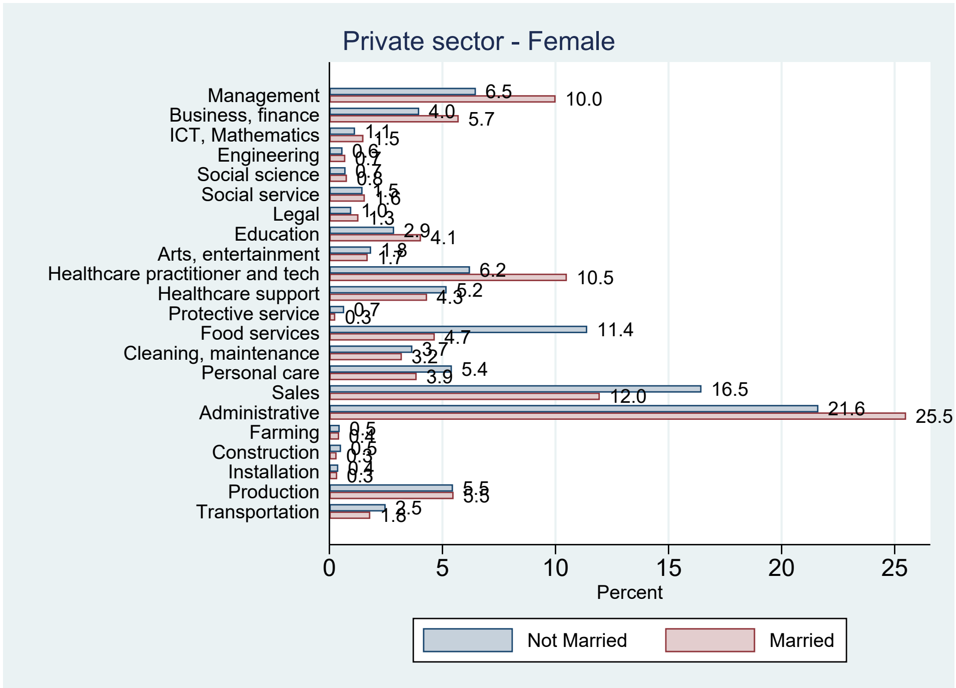

Figure 2.A.15: Distribution of women across private sector occupations depending on their marital status (averages for the period 2003-2007). Source: CEPR-CPS, own calculations

2.B Data Definitions and Sources

Variable	Source	Definition
Government expenditures		
Government total expenditures	Bureau of Economic Analysis	Government total expenditures (Item 37, Table 3.1)
Government investment	Bureau of Economic Analysis	Gross government investment (Item 39, Table 3.1) minus defense investment expenditure (Item 19, Table 3.9.5)
Government defense	Bureau of Economic Analysis	National defense expenditure (Item 17, Table 3.9.5)
Government consumption	Bureau of Economic Analysis	Government consumption expenditure (Item 21, Table 3.1) minus defense consumption expenditure (Item 18, Table 3.9.5)
Government wage bill	Bureau of Economic Analysis	Compensation of general government employees (Item 4) minus compensation of defense government employees (Item 26), Table 3.10.5
Government benefits	Bureau of Economic Analysis	Government social benefits (Item 23, Table 3.1)
Net tax revenue	Bureau of Economic Analysis	Current tax receipts (Item 2) plus Contributions for government social insurance (Item 7) plus Current transfer receipts (Item 13) minus Current transfer payments (Item 19) minus Subsidies (Item 27), Table 3.1.
Other macroeconomic variables		
Total output	Bureau of Economic Analysis	Gross Domestic Product, Item 1, Table 1.1.5 Gross Domestic Product
Interest rate	FRED	Federal Funds Rate (Item FEDFUNDS)
GDP deflator	Bureau of Economic Analysis	Item 1, Table 1.1.4 Prices Indexes for Gross Domestic Product
Inflation rate		Quarterly growth rate of GDP deflator
Population	U.S. Bureau of Labor Statistics	Civilian noninstitutional population (item LNU00000000)
Labor market variables		
Hourly wage	CEPR extracts of CPS MORG - Authors' calculations	Average real hourly wage of male (female) workers
Weekly hours	CEPR extracts of CPS MORG - Authors' calculations	Average weekly hours of male (female) workers
Employment rate	CEPR extracts of CPS MORG - Authors' calculations	Male (female) employment rate is constructed as the ratio of men (women) employed to total male (female) working-age population

Notes: Table 3.1: Government Current Receipts and Expenditures; Table 3.9.5: Government Consumption Expenditures and Gross Investment; Table 3.10.5: Government Consumption Expenditures and General Government Gross Output.

Table 2.B.1: Data Definitions and Sources

2.C VAR Estimation Method and Algorithm for Computation of IRFs

The procedure to identify the shocks follows the approach described in Arias et al. (2014) to make independent draws from the posterior distribution of structural parameters conditional on the sign and zero restrictions.

The VAR model can be written in the following general form:

$$y_t' A_0 = \sum_{k=1}^{p} y_{t-k}' A_k + c + \epsilon_t' \tag{2.1}$$

where y_t is the vector of n endogenous variables, ϵ_t a $n \times 1$ vector of exogenous structural shocks. The reduced form representation of this model is:

$$y_t' = x_t' D + u_t' \tag{2.2}$$

where $D = BA_0^{-1}$, $u_t' = \epsilon_t' A_0^{-1}$ and $E(u_t u_t') = \Sigma = (A_0 A_0')^{-1}$, and $B' = [A_1' \ ... \ A_p' \ c']$. The matrices D and Σ are the reduced-form parameters, A_0 and B the structural parameters.

Let h be any continuously differentiable mapping from the set of symmetric positive definite $n \times n$ matrices into the set of $n \times n$ matrices such that $h(X)'h(X) = X$. In particular, $h(X)$ could be the Cholesky decomposition of X. We have $(A_0, B) = (h(\Sigma)^{-1}, Dh(\Sigma)^{-1})$. Denoting $f(h(\Sigma)^{-1}, Dh(\Sigma)^{-1})$ a function, with dimensions $nr \times n$, which stacks the impulse responses for the r horizons where sign restrictions are imposed, such that it satisfies $f(h(\Sigma)^{-1}Q, Dh(\Sigma)^{-1}Q) = f(h(\Sigma)^{-1}, Dh(\Sigma)^{-1})Q$ for any orthogonal matrix $Q \in O(n)$. Zero restrictions can be defined using matrices Z_j of dimension $z_j \times nr$, with z_j the number of zero restrictions imposed on $f(h(\Sigma)^{-1}, Dh(\Sigma)^{-1})$. The parameters (D, Σ) satisfy the zero restrictions if $Z_j f(h(\Sigma)^{-1}Q, Dh(\Sigma)^{-1}Q)e_j = 0$, for $1 \leq j \leq n$, where e_j is the jth column of the identity matrix I_n.

The main steps of the algorithm are the following:

1. Draw (D, Σ) from the posterior distribution of the reduced-form parameters.

2. Draw $X = [x_1, ..., x_n]$ from an independent standard normal distribution.

3. Let $Q = \left[\frac{N_1 N_1' x_1}{\|N_1' x_1\|} \dots \frac{N_n N_n' x_n}{\|N_n' x_n\|}\right]$ where the columns of matrix N_j form an orthonormal basis for the null space of the $(j-1+z_j) \times n$ matrix M_j:
   $M_j = \left[\frac{N_1 N_1' x_1}{\|N_1' x_1\|} \dots \frac{N_{j-1} N_{j-1}' x_{j-1}}{\|N_{j-1}' x_{j-1}\|} \; Z_j f(D, \Sigma)\right]$ for $1 \leq j \leq n_s$, where n_s is the number of structural shocks considered.

4. Keep the draw if it satisfies all the sign restrictions.

5. Repeat steps 2-4 for M draws of orthogonal matrices Q.

6. Repeat steps 1-5 for N draws from the posterior distribution of the VAR parameters.

7. For all accepted draws, compute and save the corresponding impulse response.

8. Lastly, calculate the median and the 16th and 84th percentiles of all the impulse responses.

2.D State-level Analysis

In this section, we provide further evidence on the effects of fiscal policy on gender gaps using U.S. state-level data. This exercise provides us with a greater set of experiments and thus serves as a robustness check. Another advantage of using state-level data is that monetary policy can be taken as exogenous. This reduces interactions between monetary and fiscal policy and improves the identification of shocks. Results are in line with our baseline findings, as state and local government expenditure on consumption and the wage bill reduce gaps in all labor market variables, while expenditure on investment increase them.

We use annual data for all 50 U.S. states from 1979 to 2007. Data on state and local (S&L) level spending and taxes is taken from the regional government finance series of the Bureau of the Census.[39] The gross state products and U.S. aggregate data come from the Bureau of Economic Analysis. As in our baseline analysis, micro-level data on several labor market variables is from the CEPR extracts of the Current Population Survey (CPS) Merged Outgoing Rotation Groups. For each state and gender, we build annual series of weekly hours worked, real hourly wages, and employment rates. As in the baseline, we only consider full-time workers aged 16-64.

[39] Unfortunately, quarterly macroeconomic data is not available at the state level.

The structure of the VAR model that we estimate resembles the one in the baseline. The vector of endogenous variables includes the log of real per capita total state and local government level expenditure, the log of real per capita net tax revenue, the log of real per capita gross state product, the gender wage gap and another labour market gap variable.[40] The latter alternates between the gender gap in (i) employment rates, (ii) hours. All gender gaps are measured as in the baseline. In addition, we rotate in three different fiscal subcomponents: (i) S&L government consumption expenditure, (ii) S&L government investment expenditure and (iii) expenditure on the S&L government wage bills. Furthermore, we include total S&L expenditure net of the respective fiscal variable of interest as an anchor. Exogenous variables include two lags of an exogenous war dummy to control for fiscal foresight, as well as aggregate U.S. government expenditure, GDP and the federal funds rate to control for the state of the aggregate business cycle. Since we only have annual data, we limit the lag length to one following Pappa (2009). The VAR model is estimated for each state separately using sign restrictions. We follow the same recursive identification strategy as in the main text except without identifying the monetary policy shock. Hence, we only identify a generic business cycle shock, a tax shock and the government spending shock(s).

Figure 2.D.1 presents "typical" (mean) responses of gender gaps in wages, hours, LFP and employment for U.S. states. The first column reports the responses to a shock to total government spending, the second to government consumption, the third to the wage bill and the fourth to government investment. Following Pappa (2009), solid lines are the average of individual states' median responses to a shock that raises the fiscal expenditure component by 1% on impact. The confidence intervals are one standard error bands computed as the square root of the average variance across states. Qualitatively, these estimates confirm our baseline results. Total S&L government spending, spending on S&L consumption as well as the S&L wage bill reduce gaps in all labour market variables. In contrast, S&L investment spending widens gaps in wages, LFP and employment. However, the hours gap response is insignificant in our baseline but significantly falls after an investment shock at the state level. The magnitude of these estimates is low due to heterogeneous responses across states.

This is illustrated by Figure 2.D.2 which plots the median impact responses of the wage gap and the hours gap for each state to a total S&L government spending shock. Filled dots imply that the gender wage gap changed significantly. Although there are a number of outliers, the state data

[40] We have also considered state and local level government expenditure separately and the results do not change.

confirm previous results since the majority clusters in the third quadrant, suggesting that hours and wage gaps tend to close after a shock to total S&L government spending. For 16 out of 50 states the gender wage gap closes significantly. Interestingly, for four out of the five states in which the gender wage gap widens significantly, female dominated industries account for a below-average share of state personal income.[41] In contrast, Nevada, which is the state with the highest concentration of female dominated industries and pink collar occupations (i.e. services), shows the strongest decline in the gender wage gap. This is again in line with our baseline analysis where we find that the gender wage gap declines more among pink-collar workers and those in female dominated industries after an aggregate U.S. total spending shock.

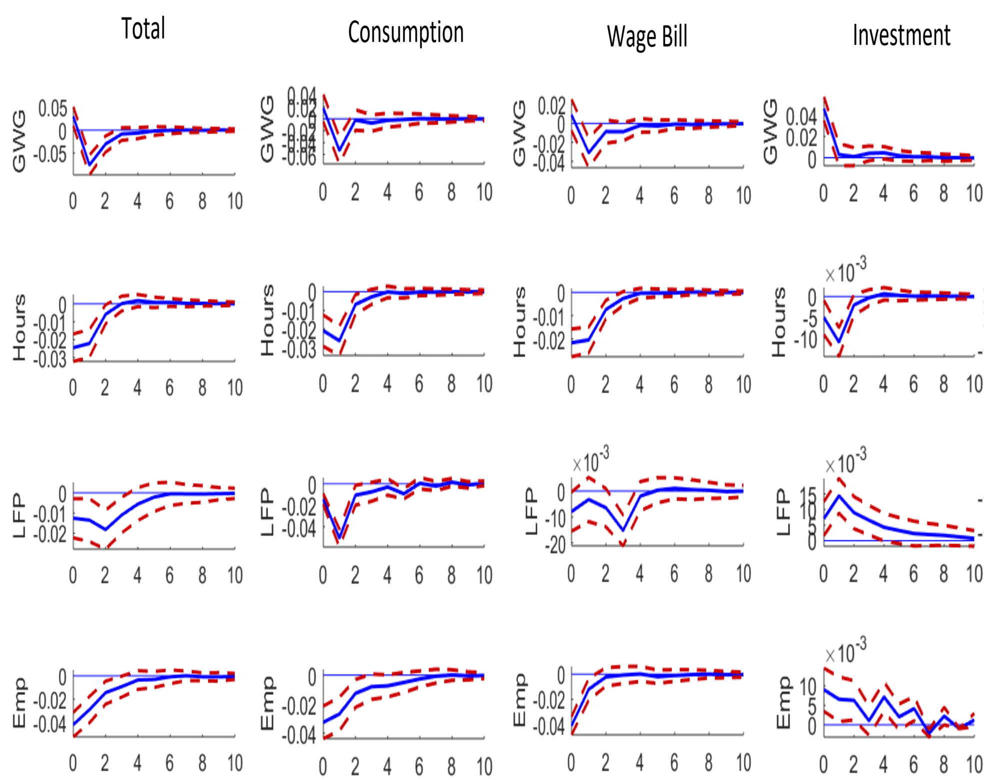

Figure 2.D.1: “Typical" (mean) responses of labour market gaps for U.S. states by spending components

[41] We make this statement on the basis of personal income data by NAICS industry obtained from the BEA.

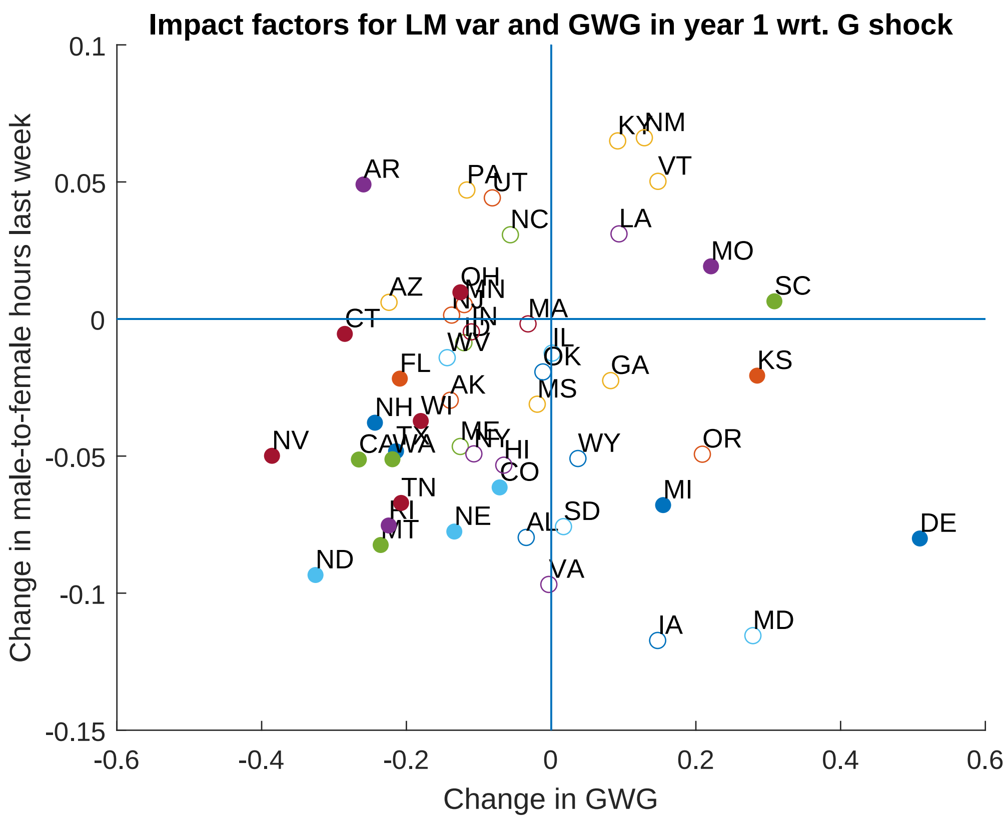

Figure 2.D.2: Impact multipliers for U.S. states to a total S&L government spending shock.

2.E Business Cycles and Gender Gaps

In this section we report the impulse responses of male and female labor market variables and gender gaps to a generic business cycle shock, identified as a shock that raises output and government net tax revenues for four quarters after the shock.[42] We first present the responses for total population. We next explore male and female labor market dynamics over the business cycle across population subgroups, along the same dimensions as in our baseline analysis for fiscal shocks: age, education, occupation, public vs. private sector.

Impulse responses for total population are displayed in Figure 2.E.1. Business cycles shocks unambiguously affect disproportionately men, in line with previous litterature.[43] Wage, hours and

[42] We acknowledge that this shock could be driven by a range of factors, including TFP.
[43] See for instance Clark and Summers (1981), Hoynes et al. (2012), Bredemeier et al. (2017).

employment gaps all increase after a positive business cycle shock. Several studies have pointed out the importance of gender differences in occupations and industries.[44] Men are more represented in highly cyclical industries such as construction and manufacturing, while women are more likely to be employed in services and public administration, which are less exposed to cycles. Our results broadly confirm the role of gender composition of occupations and industries in explaining the larger exposure of men to cyclical fluctuations.

The responses for pink- vs. blue-collar occupations, displayed in Figure 2.E.2, show that wages and employment strongly increase for blue-collar workers, especially for men. In contrast, the increase is less pronounced and less persistent among pink-collar occupations. Despite these differences in male and female outcomes across occupation groups, the increase in employment gap is of similar magnitude, as women working in blue-collar occupations also benefit more from the boom. However, we can note a more persistent increase of the wage gap among pink-collar occupations, as female wages are not affected by the business cycle shock, except negatively in the short-run, while male wages increase. The hours gap increases significantly on impact among pink-collar workers, as women do not adjust their hours worked, while women in blue-collar occupations increase them. Results for female vs. male dominated industries, displayed in Figure 2.E.3, yield similar conclusions.

Regarding the different dynamics between private and public sector, impulse responses shown in Figure 2.E.4 indicate that wages and employment in the private sector display high pro-cyclicality, while they are weakly related to the business cycle in the public sector which offers more job security. Furthermore, men working in the private sector benefit more than women from booms, leading to a strong increase in gender gaps in this sector.

Next, we consider groups split by educational attainment. Impulse responses are displayed in Figure 2.E.5. Wage and employment rise more persistently for the less educated after a positive business cycle shock, especially for men. Workers with lower education are more likely to work in construction and manufacturing sectors, thus are more exposed to business cycles. However, the increase in gender gaps is of similar magnitude across subgroups, as women benefit less from booms than men, irrespectively of their education level.

Along the age dimension, the differences between the responses to a business cycle shock are particularly stark between the youngest (16-24) and middle-aged (45-54) groups, as shown in

[44] See for instance Altonji and Blank (1999), Blau and Kahn (2000), Hoynes et al. (2012), Bredemeier et al. (2017).

Figure 2.E.6. In particular, the strong increase in wages and employment reveal substantial cyclical sensitivity among the young, especially for men. This could be explained by the higher concentration of youth employment in industries more exposed to cycles, such as restaurants, construction, manufacturing, wholesale and retail trade. In contrast, among the middle-aged, wages tend to decrease in the short-run, especially for women, and employment only increase for men, leading to a stronger increase of the employment gap in this age group.

Overall, we find that men experience more cyclical labor market outcomes. In particular, the population groups which are most affected by business cycles are young, low-educated, and blue-collar workers, as well as private sector employees, with a higher propensity to be employed in cyclically sensitive industries, as documented by previous studies. However, gender gaps tend to increase across most population subgroups, as women benefit less from booms, both in terms of wage and employment dynamics. These results imply that cyclical fluctuations have distributional implications, which raises not only a trade-off between stabilizing aggregate employment and stabilizing its composition, but also a trade-off between gender equality and other forms of equality.

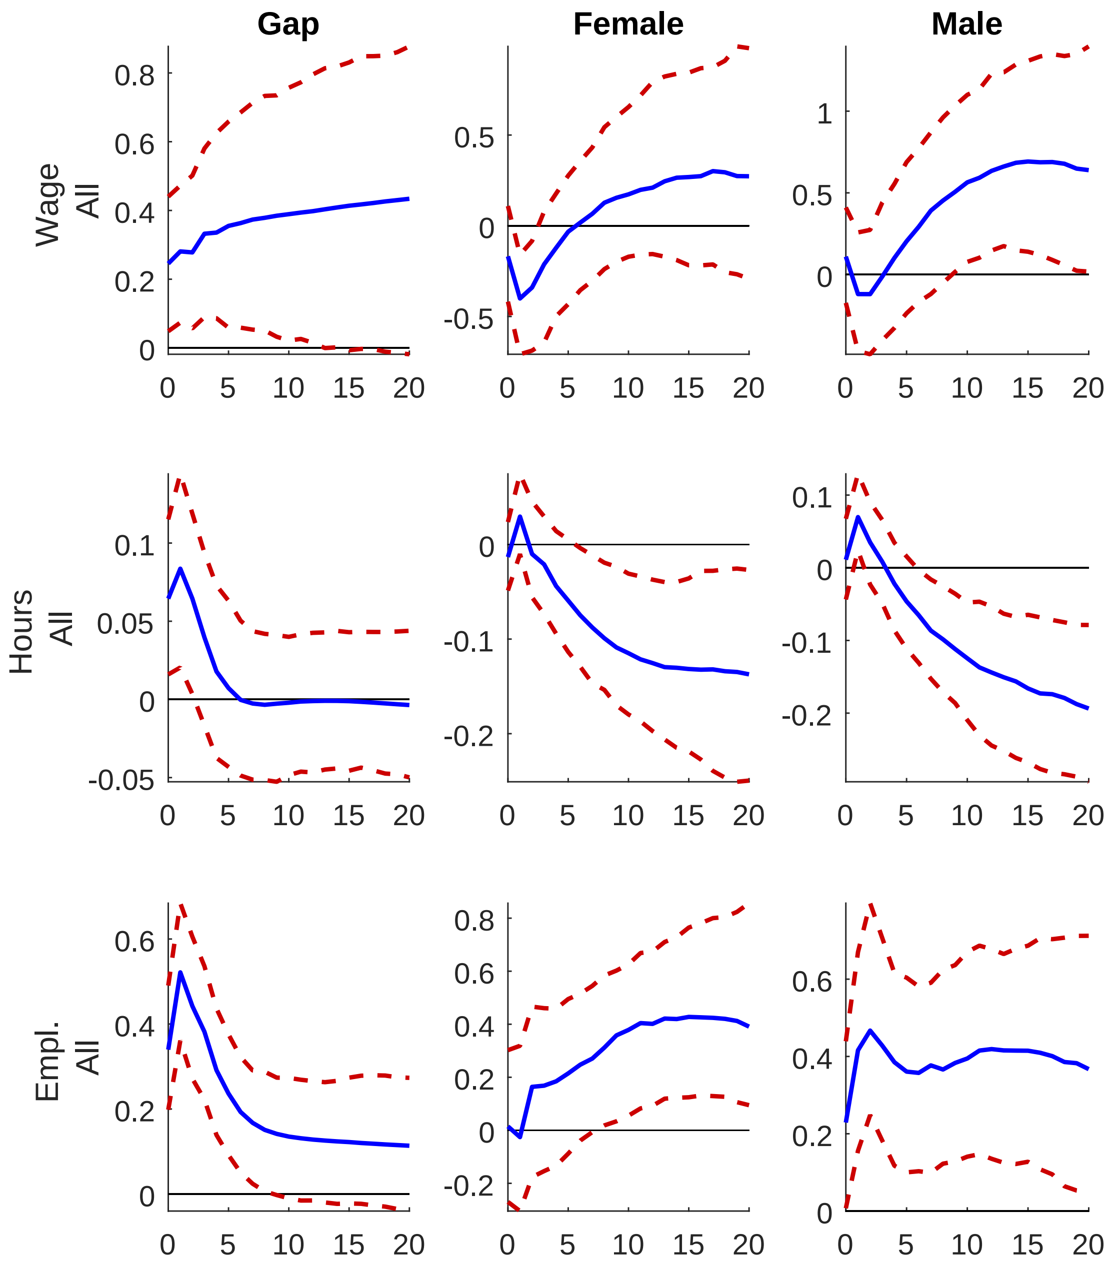

Figure 2.E.1: IRFs to business cycles shocks - total population

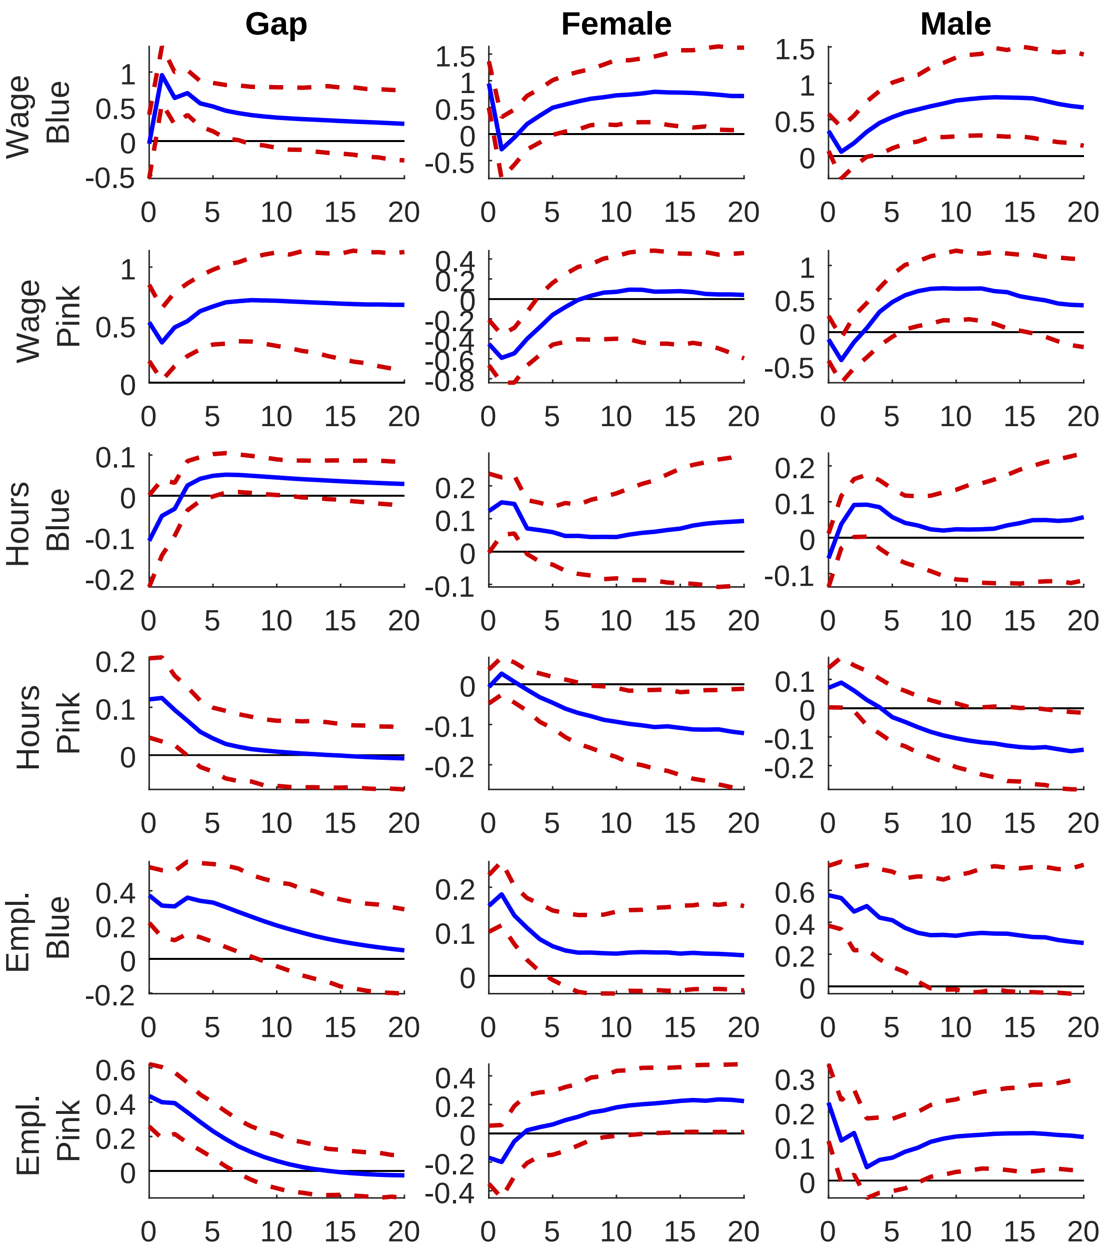

Figure 2.E.2: IRFs to business cycles shocks - pink vs. blue occupation

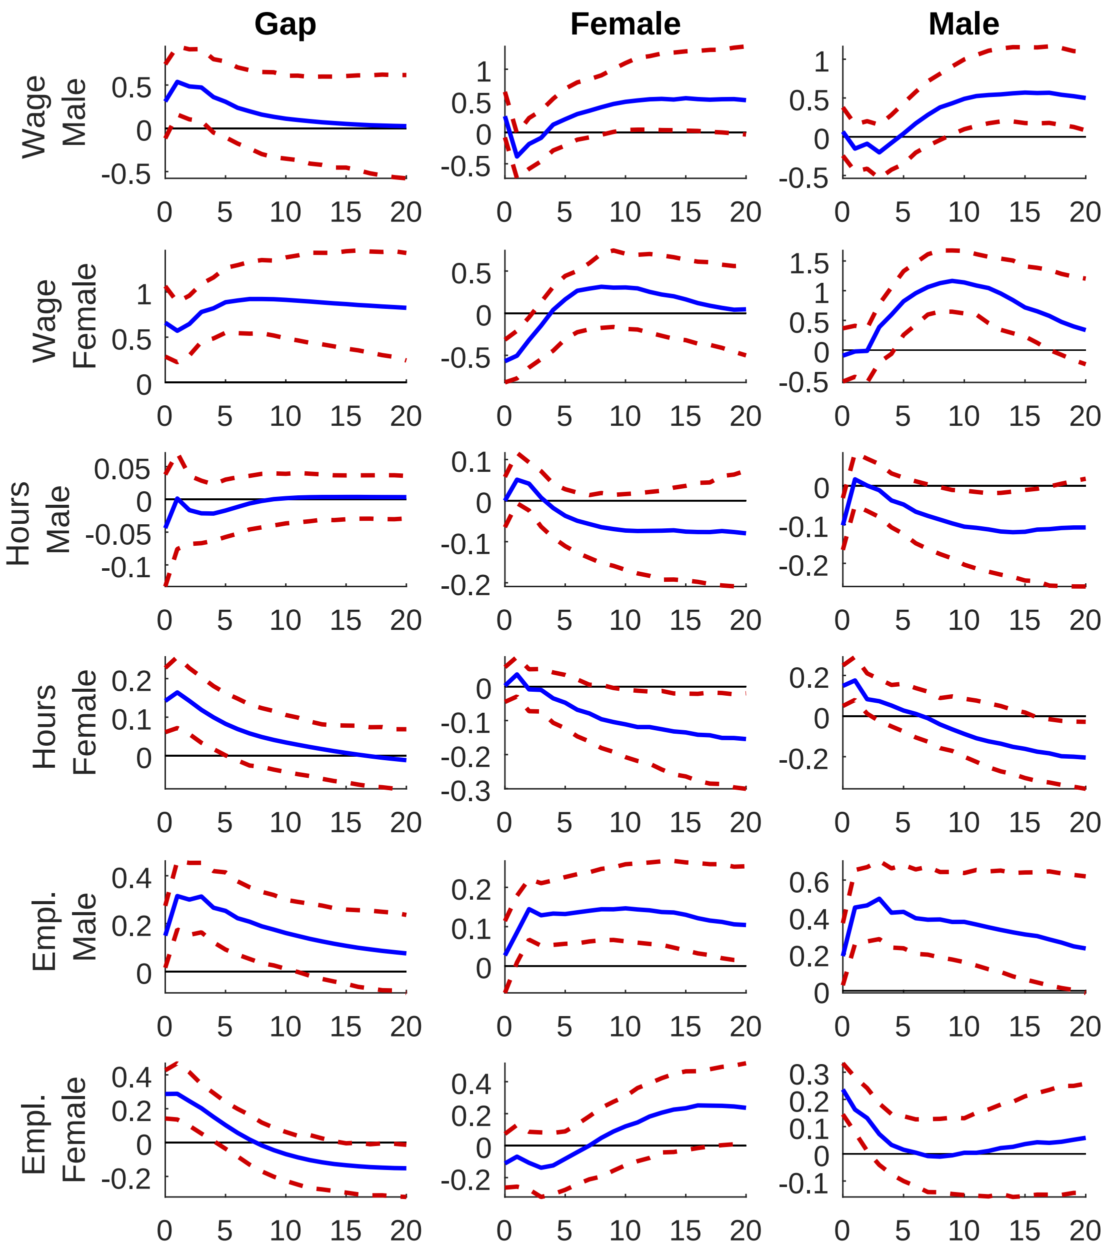

Figure 2.E.3: IRFs to business cycles shocks - female vs. male dominated industries

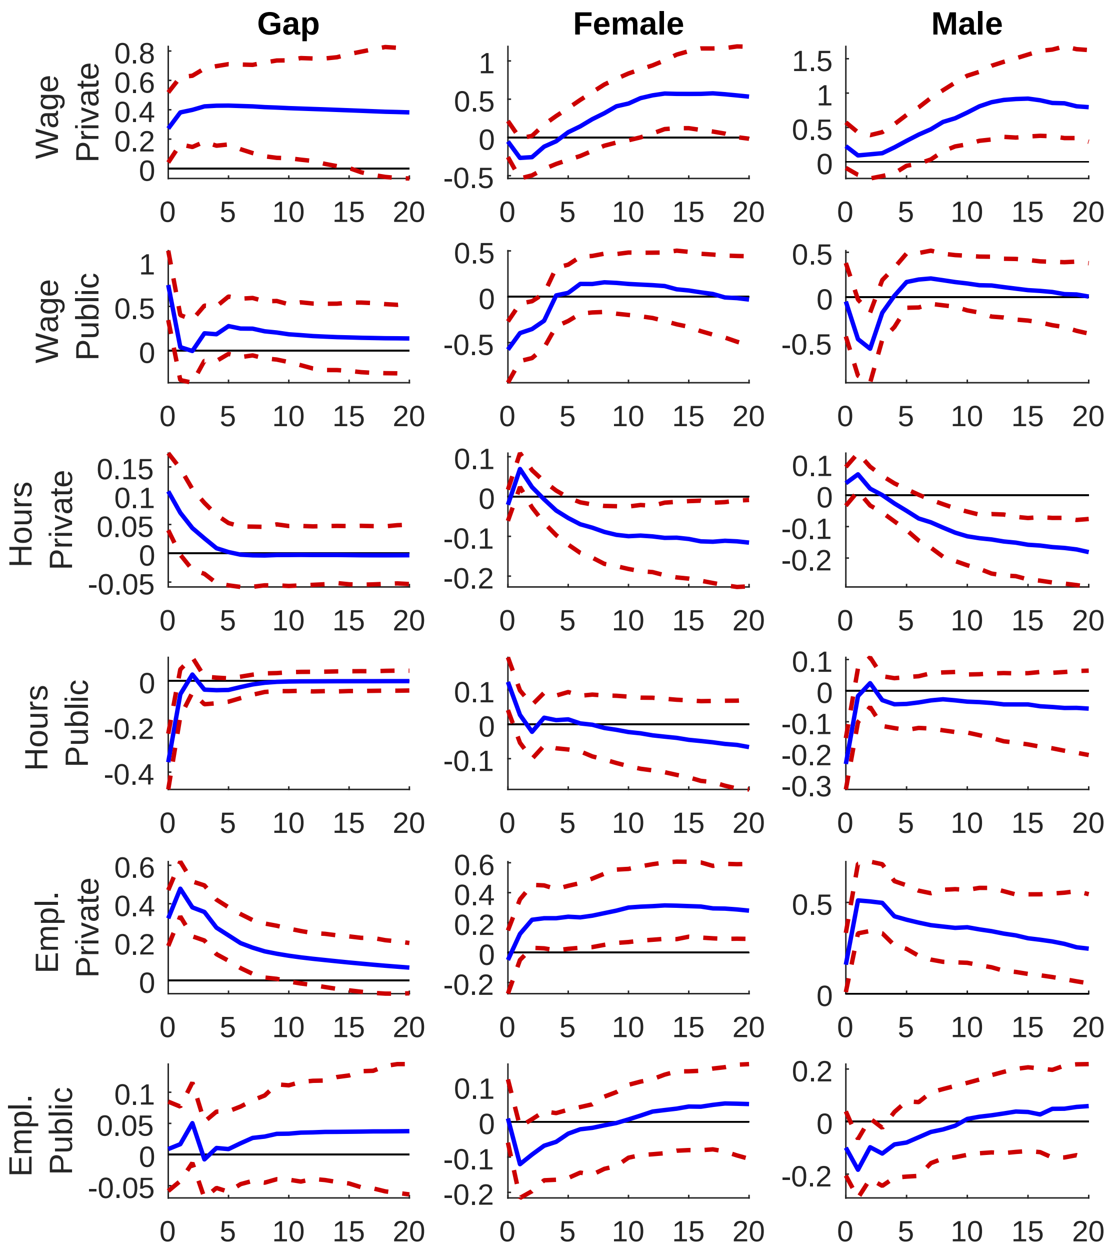

Figure 2.E.4: IRFs to business cycles shocks - public vs. private sector

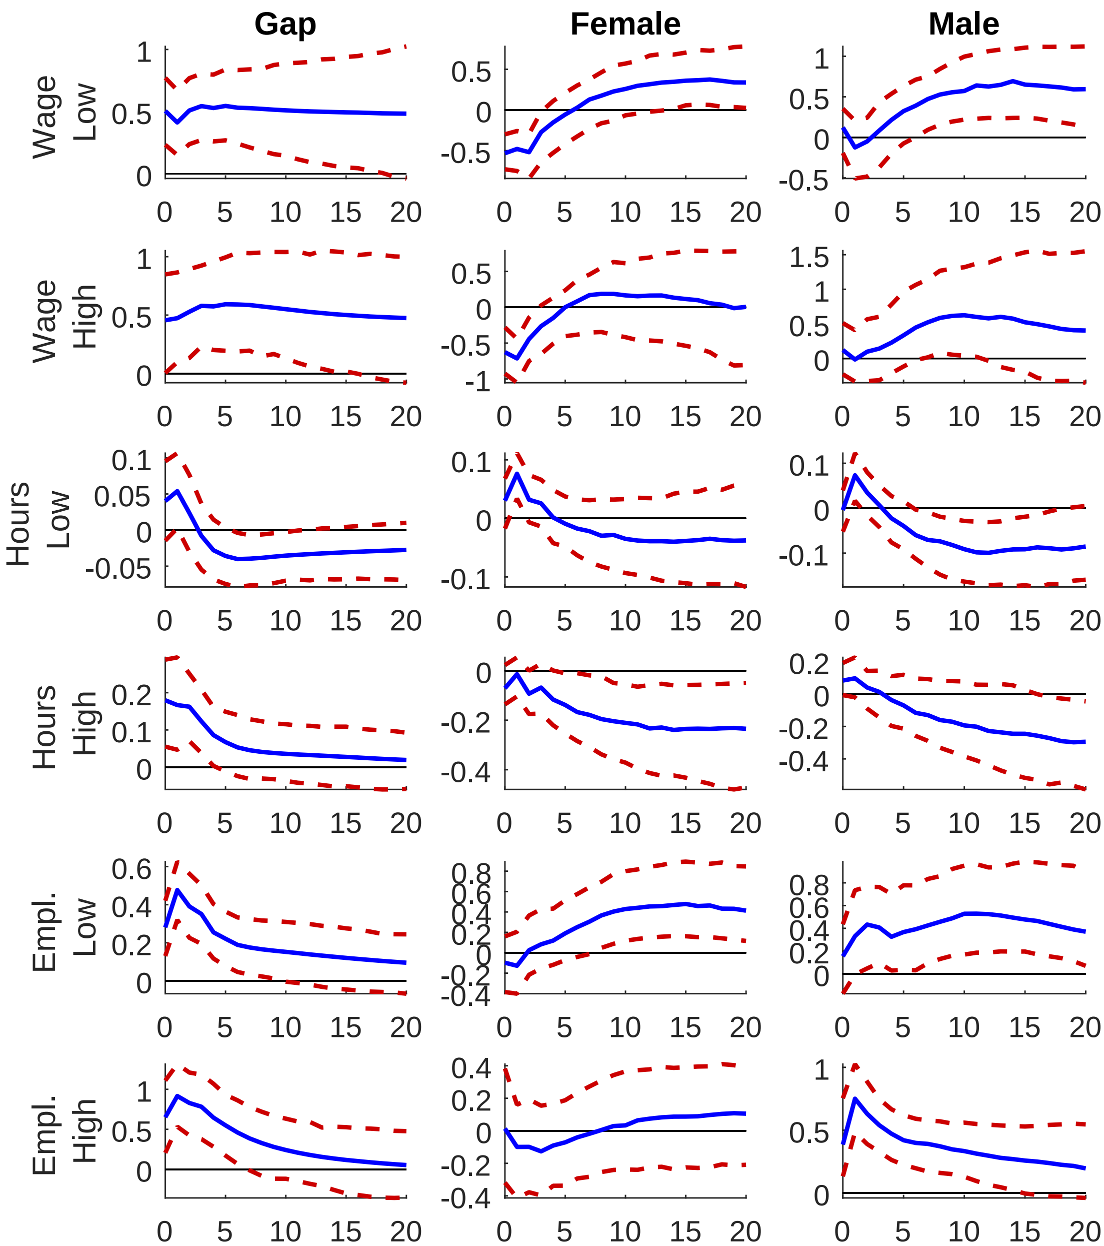

Figure 2.E.5: IRFs to business cycles shocks - high vs. low education

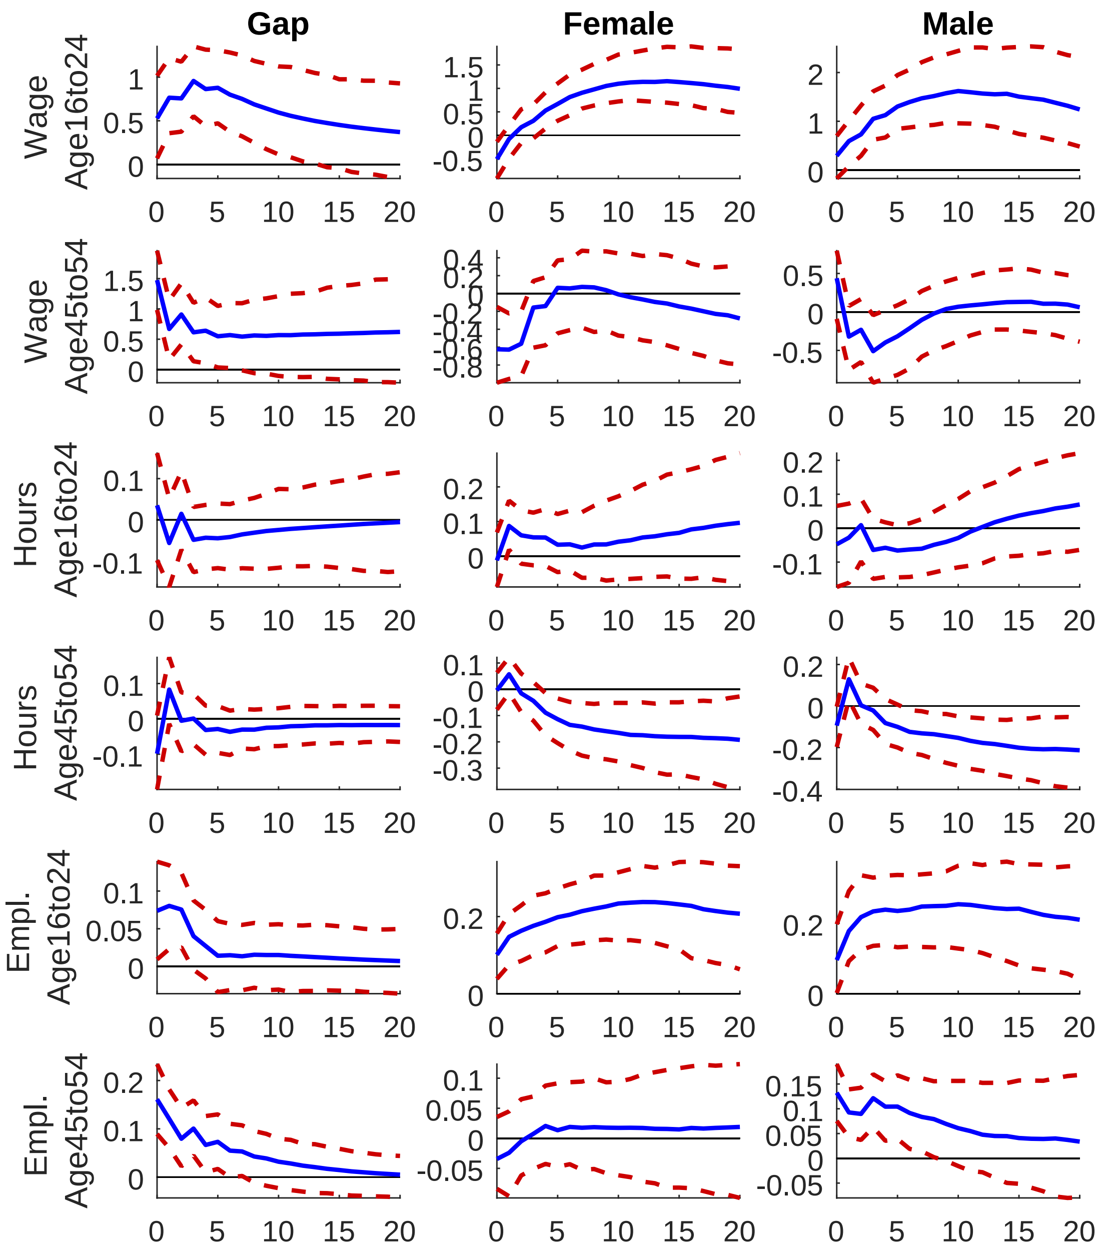

Figure 2.E.6: IRFs to business cycles shocks - workers aged 16-24 vs. 45-54

Chapter 3

Short- and Long-Run News: Evidence from Giant Mineral Discoveries

3.1 Introduction

Starting with the seminal work by Beaudry and Portier (2004, 2006), there has been a revival of the idea that aggregate changes in anticipations about the future can generate important macroeconomic fluctuations. Recently, a breakthrough has been made in the identification of these changes to anticipations–so-called 'news' shocks by the literature. Indeed, Arezki et al. (2017) (henceforth ARS) have used a plausibly exogenous giant oil and gas discoveries sample for a set of post-war country time-series data. Giant discoveries come, by definition, as a giant surprise, and therefore constitute an attractive collection of macroeconomic events for the identification of news shocks.

We contribute to the measurement of news shocks building on the same premise as ARS, but using a high-quality proprietary data set on giant mineral discoveries instead of oil and gas. Our data includes a total of 131 Tier 1 discoveries worldwide since 1960. These are discoveries of minerals such as gold, silver, diamonds, nickel, copper, iron ore, uranium, etc.

By repeating the ARS distributed-lag regression exercise, we find that there is a small or zero impact effect of the discovery on macroeconomic variables. This is not to say there is no evidence of anticipation effects. There is. Indeed, we find evidence of sizable anticipation, or news effects, but only two or three years before production starts: GDP, investment, and consumption all rise, and the current account turns negative. Similar to ARS, the current account turns negative two or three years before production, and then sharply rises a couple of years after production has started. GDP, consumption, and investment peak when production starts. Different to ARS, we do *not* find a fall in employment, instead, in our sample employment rises and comoves with GDP.

In light of our results, a question that immediately arises is: Why are the responses in the minerals sample clearly more delayed than in the oil[1] sample? An obvious candidate reason is a technological delay in the median exploitation of new mines larger than for new oil fields. Luckily, our data set contains precise information about this delay. We do indeed find that the median delay is 9. ARS report a delay of 4 to 6 years on average for oil fields, and we compute a median delay of 5 years for a similar oil sample. Thus, mines seem to require a significant longer delay than oil fields for exploitation.

We recur to the ARS model to see if a longer delay generates a delayed news response of macroeconomic variables. We simulate the ARS model using a delay equal to the median delay in our mineral discoveries sample. The ARS model correctly predicts what we observe in the data. There is a zero or small macroeconomic effect on impact, a news effect two or three years before production starts, with GDP, consumption and investment comoving. (The only feature of our data that the ARS model is not able to predict being the rise of employment.)

In order to more deeply explore the role of the production delay on the macroeconomic effects of these discoveries, we merge the ARS oil sample to our minerals sample. We then split this sample of both mineral and oil discoveries into short and long delays. In order to validate this split, we look into the correlation or the type or discovery (oil vs. minerals) with geography or GDP per capita. We do not find any correlation, and thus it is actually plausible that the delay is mostly technological and therefore exogenous. We run the ARS distributed-lag regression model on the sample split, and we indeed find that discoveries with a longer delay have similar macroeconomic effects than mineral discoveries. Discoveries with a long delay have similar macroeconomic effects than discoveries with shorter delay, but these effects arrive much later, and have roughly a nil effect on impact.

The success of the ARS model in reproducing the delay observed in the data is remarkable.

[1] Consistent with ARS, we will call oil and gas simply ‘oil’.

However, at the same time, we currently lack a direct test that the channels present in this model are indeed the ones generating the delay in the macroeconomic responses in the data. With this point in mind, we include towards the end of the paper a brief discussion of plausible channels that can generate the type of delay observed. We include in the discussion the role of limited foresight due, for example, to limited cognitive abilities or lack of common knowledge. All these channels have been extensively discussed in recent papers and thus are worth emphasizing. We do this on Section 3.5 below.

There is a large macroeconomic literature studying the effects of news shocks. Following Beaudry and Portier (2006) , Jaimovich and Rebelo (2009) and Schmitt-Grohe and Uribe (2012) showed how to capture the empirical effects of news shocks within state-of-the-art models. Barsky and Sims (2012) expand the notion of news to macroeconomic confidence. Blanchard et al. (2013) and Chahrour and Jurado (2018) deal mainly with issues of identification.

3.2 Data

In order to estimate the macroeconomic responses to giant mineral discoveries, we merge data from several sources. The first source is a data set on major mineral deposit discoveries worldwide starting in 1950. The second source consists of a collection of data sets on commodity prices. The third source is a data set of macroeconomic variables for a large number of countries. Our fourth source is a data set of giant oil discoveries, which we use for comparability with ARS. We describe each in turn and then present summary statistics.

3.2.1 Minerals

To get information on mineral discoveries, we employ a proprietary data set. This data was generously shared by one of the main mining consultants worldwide: MinEx Consulting Pty Ltd. The data contains information on the date of discovery, estimated total reserves, and date of production start (mine startup date) for 12 key minerals:

- precious metals: gold and silver;

- gems: diamonds[2];
- bulk minerals: coal, bauxite, iron ore;
- non-ferrous metals: copper, nickel, lead and zinc;
- uranium,
- and potash.

(The data set contains information on discovery and production dates for other minerals as niobium, zircon, or graphite, but it does not contain information on reserves for these minerals.)

To the best of our knowledge, this data set has not been previously used in economics.[3] A novelty in our data set—besides the presence of commodities not present in other data sets, as the Horn (2014) oil and gas giant discoveries data set—is the presence of the date of production, which allows a computation of the delay between discovery and production. As it will turn out, this delay will be important to understand the news effect of these giant discoveries.

Our data set contains a total of 131 Tier 1 discoveries of minerals mentioned above. A 'discovery' is an event in which a major deposit of mineral is discovered in a given country, in a given year (all our data is annual.) A Tier 1 discovery is a deposit defined and determined as 'major deposit' by the data set provider (MinEx Consulting).[4] These discoveries happened in 41 countries between 1960 and 2015.[5] 15 countries in the sample experience only one discovery.

The data set also contains diamond discoveries (a total of 12). Unfortunately, in the case of this mineral, the data does not contain information on estimated reserves. Given the saliency and importance of this gem, we searched for information on estimated reserved for these 12 discoveries.

[2] No production start date in the case of this commodity, see below.

[3] There exists a parallel literature in political science; this data set does not seem to have been used there either.

[4] MinEx consulting considers a major deposit (Tier 1) a mine containing an estimated, say, > 1 Mt Cu-equiv, > 100 kt Ni, > 1 Moz Au, > 10 m carats, > 25 kt U_3O_8 Schodde (2010). To give an idea of what this means; 1 Mt of copper (Cu) is equivalent to the annual production of Australia, the top-6 copper producer in the world, or two months of production of the biggest copper exporter, Chile. 1 Moz of gold (Au) is equivalent to a one-month production of China, the largest gold producer, or to the annual production of Burkina Faso, the top-20 gold producer in the world.

[5] We did not use the data before 1960.

We found this information on different sources. These include the world's leading diamond company, De Beers; the site `https://www.mining-technology.com`; and other sources available on request.[6]

A key and interesting dimension of our data set is the delay between the date of discovery and the date of production start because it determines the horizon of the news shock. Our data set contains high quality information on this delay (except, of course, in a few cases where production did not start before the end of the sample.) Moreover, the delay in the case of minerals is much larger than in the case of oil, which leads itself to an interesting analysis. We will present some descriptive statistics of this delay in this same section, below.

3.2.2 Commodity Prices

In order to determine the economic value of the discovery we need commodity prices. We collect these data from several sources. For commodity prices, our main source is the World Bank Commodity Price Data. It provides data for main commodities starting from 1960. For the earlier period (which we need for the regression of subsection 3.4.3), we rely on Jacks (2019). For diamonds we supplement the data by the U.S. Geological survey. Uranium price data comes from TradeTech (`www.uranium.info`).

An challenge in the construction of a value of discoveries dataset for economics is the lack of uniformity in conversion units across sources. For example, potash is represented by K_2O in MinEx and KCl in World Bank pricing data. A conversion ratio is typically applied with the assumption that other elements in the compound have little or no value. We followed this convention.

3.2.3 Macroeconomic Data

Our macroeconomic data sources are similar to ARS, whose main source is the IMF World Economic Outlook.[7] Our macro data set contains information on GDP, investment, consumption, the current account, the saving rate, and the employment rate for 181 countries. Our baseline

[6]For example, this includes the academic article by Pervov, V. A., Somov, S. V., Korshunov, A. V., Dulapchii, E. V., and Felix, J. T. (2011). *The Catoca kimberlite pipe, Republic of Angola: A paleovolcanological model.* Geology of Ore Deposits, 53(4), 295-308; and so on.

[7]Using data from the Penn World Tables instead delivers similar results.

estimation is based on the 1980-2012 time span at the yearly frequency.[8] All national accounts data are provided in real local currency units. For GDP, we also use a series in real USD in order to compute the value of the discovery (in USD) as a percentage of GDP. We note also that the dataset contains some extreme values, such as a drop in the current account in Kuwait from 20% of GDP in 1990 to -224% in 1991 (due to the Gulf war.) We have checked that the results are not driven by these rare instances.

3.2.4 Oil

We complement our mineral data set with the Horn data set on giant oil and gas discoveries Horn (2014). For brevity, unless explicitly overruled, throughout the paper we refer to oil and gas as simply 'oil'. This allows us to compare our results to ARS.

A crucial dimension of the minerals data set—and the main vector of comparison to the case of oil—is the delay between the date of discovery and the date when production starts. Unfortunately, the Horn data set does not have this information, a feature also discussed in ARS (p. 120). Thus, for each discovery in the Horn data set we search for the date of production start from other sources: the Uppsala Giant Oil field database Höök et al. (2014); the Petroleum Dataset compiled by the Peace Research Institute Oslo (PRIO) Lujala et al. (2007)[9]; among other. After this effort, roughly 60% of the oil discoveries remain without information about the date of production. (This is much larger compared to the minerals data set, where for only 19 discoveries out of 131 (14%) we do not have this information. As explained above, in the case of the minerals data set, we know that this is because production did not start yet at the end of the sample.)

3.2.5 Descriptive Statistics

Table 1 contains a first set of descriptive statistics for all commodities in our merged data set (minerals and oil). It lists the total number of discoveries across commodities, and their value

[8]Similar to ARS, due to data limitations, we cannot start earlier. Before 1970, the macroeconomic data is mostly available for advanced economies. Even though our estimation starts in 1980, we use the information on discoveries pre-1980 in the ADL model below. We point out also that the World Economic Outlook misses some of the series for some countries even after 1980—the most complete series being GDP— and therefore these are automatically dropped from the regression.

[9]We used the PETRODATA V1.2 update.

(reserves in physical quantities multiplied by prices at date of discovery, in 1998 USD.)

There is a total of 131 mineral discoveries in our sample. More than one third of these are gold discoveries. The second and third most frequent type of mineral deposits discovered are of copper and iron ore. The total number of oil discoveries in the sample is 428 in the Horn data, and gas discoveries amount to 392. As we will explain immediately below, many of these oil and gas discoveries happen in the same combination of country and year.[10]

Even though there are fewer mineral discoveries than oil discoveries in the data, their economic importance is similar. Indeed, the median value of a mineral deposit is USD 56 billion; the median size of an oil deposit is USD 53 billion; and the median size of a gas deposit is USD 88 billion. Aside from this, there is substantial value heterogeneity in the data.

Table 1: Giant Discoveries Merged Data Set: Type, Number, and Value (bln 1998 USD), 1960–2015

	Obs	Mean	Median	St Dev	Min	Max
Minerals						
Gold	50	13	8	12	2	57
Diamond	9	6	5	4	2	14
Iron Ore	14	80	51	60	30	232
Copper	21	153	71	201	13	886
Nickel	12	60	27	102	8	377
Uranium	10	20	17	11	11	47
Other	15	91	36	123	5	381
Total	**131**	**56**	**20**	**109**	**2**	**886**
Oil & Gas						
Oil	428	53	23	102	4	1,168
Gas	392	88	23	310	4	5,030
Total	**820**	**70**	**23**	**227**	**4**	**5,030**

It is a frequent event that more than one discovery happens in the same country-year combination. In fact, it is easy to imagine a positive correlation in the occurrence of discoveries for geological and technological reasons. If one re-defines a discovery by not allowing for such 'double-counting' in the same country-year combination (in other words, aggregating over discoveries happening in the same country in a given year,) there are 114 mineral discoveries, and 518 oil and gas discoveries over 1960–2015.

[10]The number of diamond discoveries in the 1960–2015 period is 9, whereas in the 1950–2015 period there are 12 diamond discoveries (as reported above.)

Figure 1 shows the geographical distribution of discoveries. It is interesting to note that discoveries are well spread out around the globe. There are more oil discoveries in our data, consequently a country is more likely to experience an oil discovery than a mineral discovery. Moreover, large countries have—naturally—a higher probability of multiple discoveries of both oil and minerals. On the contrary, small countries often have no discoveries.

Figure 1: Geographical Distribution of Giant Mineral and Oil Discoveries, 1960-2015

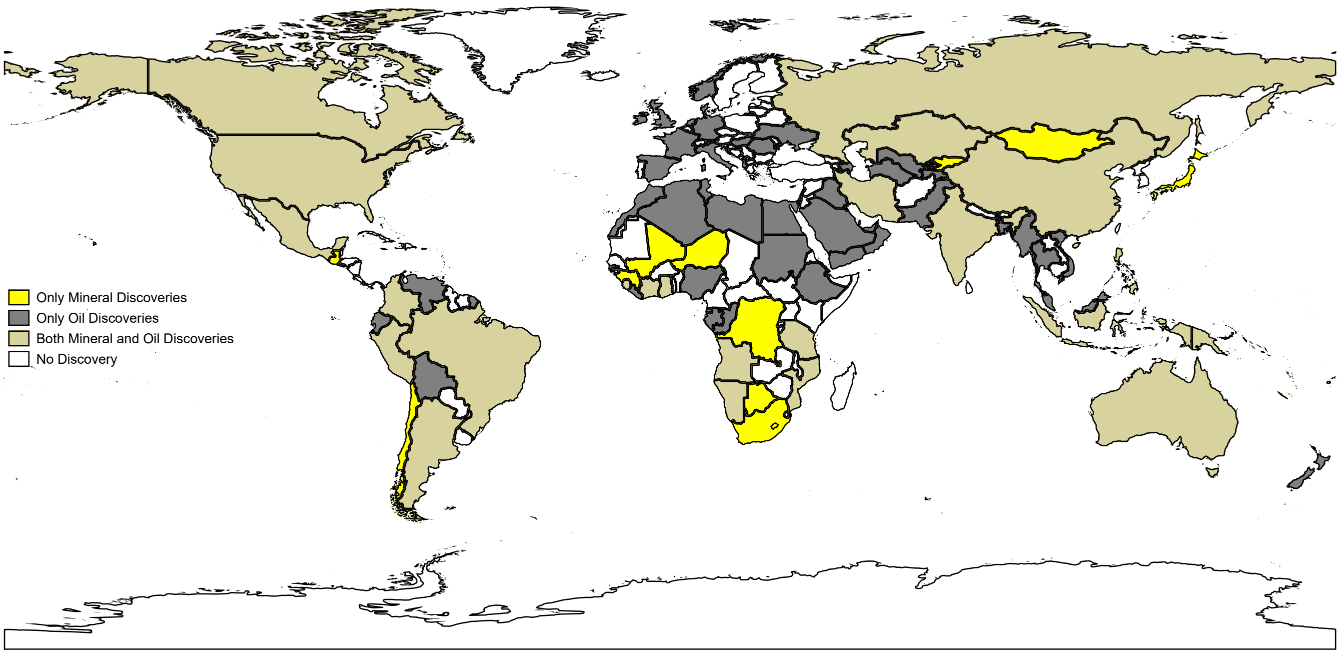

A salient difference of the mineral discovery data, when compared to the oil and gas data used by ARS, is the delay between the discovery date and the production start date. Figure 2 shows two histograms of these delays (when observed, which is the large majority in our sample): on the left, minerals, and on the right, oil. Minerals feature a larger delay than oil: the median delay for minerals is 9 years, and the median delay for oil and gas is 5 years.[11] Moreover, there is substantial delay heterogeneity across commodities ranging from 0 to 40 years, resulting in a long right tail in the empirical distributions. More importantly, the distribution of oil seems to be special in the sense

[11] The average delay for minerals is 12 years and for oil and gas it is 7.9 years. Our average measure is slightly different than in ARS because our sample is not the exactly the same: ARS computed their delay based on the Hook et al. data Höök et al. (2014), where the median delay is 3 and the average is 4.4 (5.3 for the sample of 1980-1999). However, the Hook et al. data misses all gas discoveries, many offshore discoveries as well as recent discoveries and discoveries of unconventional oil (as for instance shale oil and sand oil). All of them tend to have longer delays.

that, in the case of more than 5% of the discoveries, production started immediately.[12]

This finalizes the descriptive statistics section.

3.3 Empirical Evidence

3.3.1 Methodology

We follow ARS and use an autoregressive distributed lag (ADL) regression model to estimate the response of a number of macroeconomic variables to the discovery of a giant mineral deposit. Specifically, we estimate the linear regression

$$y_{it} = ay_{i,t-1} + B(L)Min_{it} + \alpha_i + \mu_t + \varepsilon_{it} \tag{3.1}$$

where y_{it} is the value of a dependent macroeconomic variable in country i at time t, Min_{it} is the net present value (NPV) of a mineral deposit discovered in country i at t (fully described below) normalized by GDP. α_i denotes a country fixed effect and μ_t denotes a time fixed effect. ε_{it} is a homoscedastic disturbance. $B(L)$ is a pth order lag polynomial, $p \geq 0$. In our baseline results we pick $p = 20$. This is twice the number of lags used by ARS, the reason being that the delay between the discovery and production is roughly twice for minerals than for oil.[13] Following ARS, we do not include controls (beyond country and time fixed effects) to compute our baseline estimates. In regressions using log levels of variables (rather than percent of GDP) and employment rate, we also include (country-specific) linear trends.

The mineral discovery NPV variable was constructed similar to ARS, but with a number of simplifications allowing us to accommodate very different types of minerals. Indeed, given that our dataset includes a wide range of minerals such as gold, silver, diamond, coal, bauxite, uranium, etc., production profiles and timing may vary quite a bit. Thus we benefit from this flexibility.[14]

[12] Ideally we would want to have access to the ex-ante (expected) delay once a discovery is made. This data however is not available. The implicit assumption in our statistical analysis below is that agents form a rational expectation on the delay conditional on the information of the type of discovery (essentially the type of commodity.)

[13] Using different values of p, as for instance $p = 10$ and $p = 15$, did not qualitatively change our main conclusions.

[14] ARS embed in the NPV the notion of a discounted production profile for oil fields. This is difficult to

Figure 2: Histograms: Delay from Discovery to Production Start

(a) Minerals

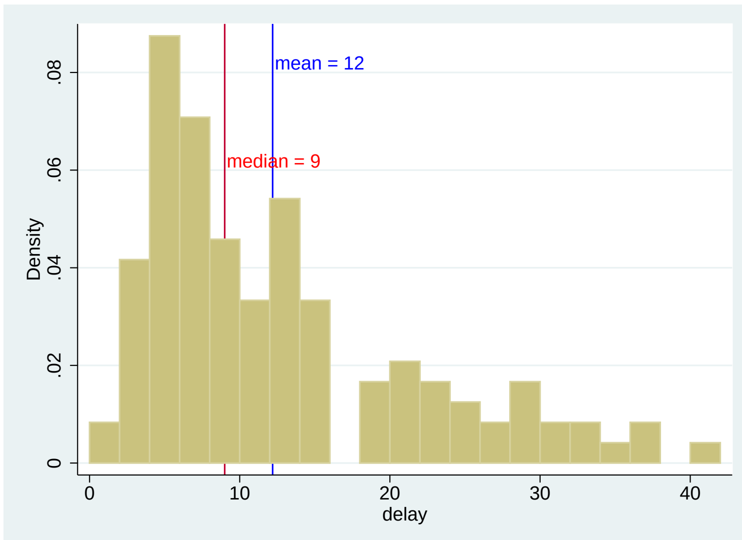

(b) Oil

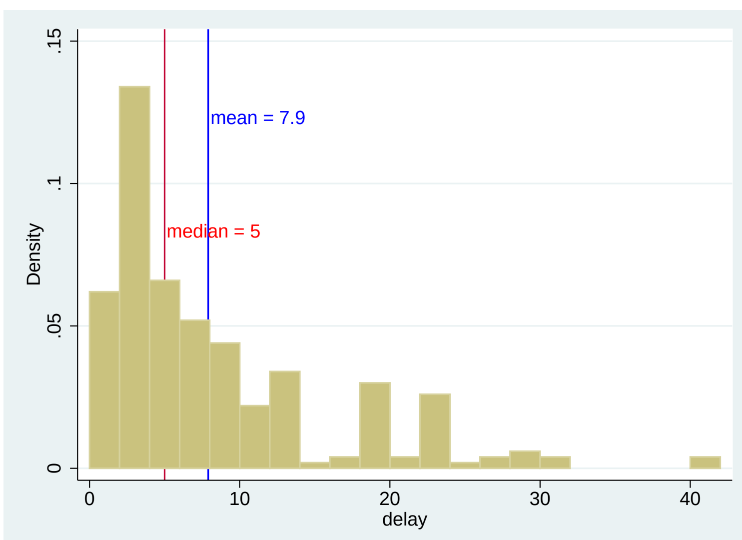

Specifically, we use the following formula for Min_{it}:

$$Min_{it} = \left\{ \frac{1}{(1+r)^{delay_{it}}} \sum_{\{\text{discovery } j\}} \left[\frac{Res_{jit} \times p_{jit}}{GDP_{it}} \right] \right\} \times 100$$

where Res_{jit} is the estimated reserves (quantity) of mineral discovery j in country i at time t, p_{jit} is the price of the mineral associated to discovery j at country i at time t, GDP_{it} is the GDP of country i at time t. The summation is over discoveries, because it is possible—and actually observed in our data—that two or more discoveries happen in a given country in the same year. r is the interest rate used for discounting, which we fix at 5%. $delay_{it}$ is the observed delay from the year of discovery to the year in which production takes place. Due to data limitations we do not observe this delay for all discoveries (see the data section for more details), in which case we use the simple average on our sample. In the case of more that one discovery per year-country, $delay_{it}$ is set to the average delay over discoveries j weighted by the dollar value of reserves (thus with weights $(Res_{jit} \times p_{jit})/(\sum_j Res_{jit} \times p_{jit})$).

We consider 6 different dependent macroeconomic variables y_{it}: the natural logarithm of investment (log investment), log private consumption, the ratio of CA over GDP, log GDP, the savings rate (the sum of investment and the CA over GDP), and the employment rate (the number of employed over civilian population).

Under the assumption—introduced by ARS—that giant discoveries are exogenous, the regression (3.1) can be estimated by OLS.

3.3.2 Results

Figures 3 and 4 present the dynamic effects of a giant mineral discovery. Figure 3 presents these effects for investment, consumption, and the current account (CA); Figure 4 presents these for GDP, the savings rate, and the employment rate.

Looking at Figure 3, investment is the variable that reacts the most. It rises and peaks 9 years after the discovery. The peak in the response of investment coincides with the median delay. Two important comments regarding the timing of the rise of investment. First, we observe relatively

do in our case given the large heterogeneity of minerals and, likely, respective extraction technologies.

little action right after the discovery. Second, investment significantly rises *before* production starts. Indeed, there is a significant rise visible already, say, 4 years after the discovery.

In terms of interpretation, let us first consider the size of the response of investment. Investment peaks at .04, meaning that, at year 9, a discovery of total size 1% of GDP raises investment by .04%. The total cumulated effect of this discovery is the integral below the impulse response, and is equal to .37. This means that the median discovery—which is of size 6% of GDP—implies a total increase of investment of $6\% \times .37 = 2.22\%$. Furthermore, the fact that investment reacts early—4 years after the discovery—suggests the presence of of a macroeconomic 'news' or anticipation effect. However, this plausible news effect does not materialize immediately after the discovery, which underlines the importance of the delay for the response of investment.[15]

We now turn to the response of consumption. Consumption also rises, but its response is of a smaller magnitude of that of investment (about half), and the estimates are less precise. To a theorist of anticipation effects, this is a bit puzzling. Indeed, studies by Barsky and Sims (2012), and especially Blanchard et al. (2013), conceptually emphasize consumption.[16] In our data, there is muted evidence of the importance of news shocks for consumption, and this is a feature also shared by the evidence in ARS. Aside from this, we do observe similar features to the reaction of investment: there is large delay in the reaction of consumption, but there is some evidence of an anticipated reaction of consumption.

Looking at the response of the CA, we observe a long initial period with a small or nil reaction, followed by a fall and a subsequent rise of the CA. The delay in the reaction of the CA is significantly larger than in the case of the investment. We do not have a good understanding for the reason—but one can think about a few options. First, we remind the reader that the distribution of delays has a fairly long right tail. Second, it is possible that production starts with relatively small quantities extracted, and therefore—in cases where the mineral is mostly exported—these exports are not large enough to immediately reverse the CA effects of investment. Similar to the insights in ARS, it is

[15]There is another interpretation of the rise in investment before the start of production. Looking at histogram (a) on Figure 2, this could merely be due to the presence of heterogeneity of the delay between discovery and production in the data. Indeed, we may see a rise of investment before the median delay simply due to the fact that several discoveries have a short delay (before production starts.) This is a point that the literature has so far not considered given that it did not have detailed information about these delays. A deeper investigation of this issue would be warranted, but it is outside the scope of this paper.

[16]Barsky and Sims (2012) emphasize consumer confidence; Blanchard et al. (2013) base their benchmark 'small scale' model solely on the responses of consumption.

Figure 3: Dynamic Effects of a Mineral Discovery: Investment, Consumption, and Current Account. Red Line: Median Delay.

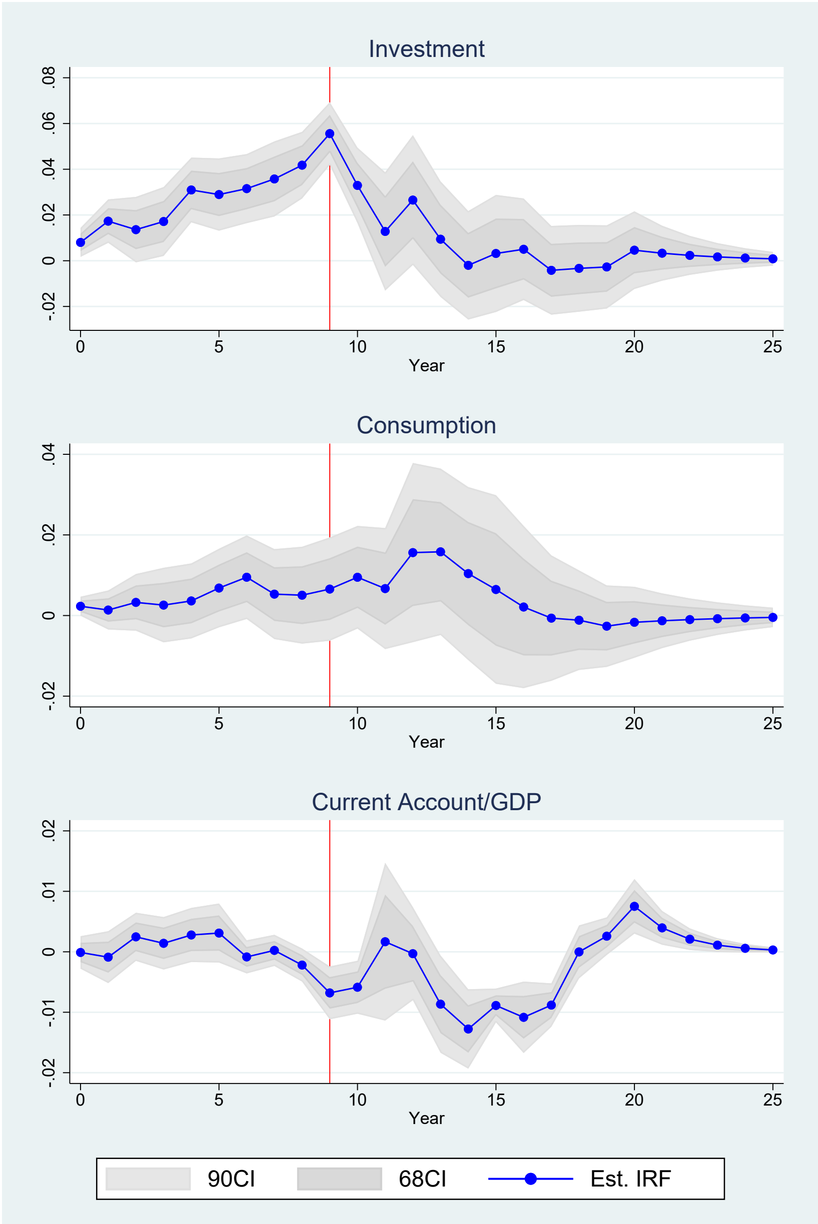

quite plausible that the fall in the CA is due to the rise in investment—which could generate a rise of imported investment goods—and the subsequent rise is due to the production start and export of minerals abroad. The rise in our case is less sharp, but it is there.

Turning to Figure 4, we now look at the responses of GDP, the savings rate, and the employment rate.

There is a strong and clearly positive reaction of GDP, which peaks between 10 to 12 years after the discovery. There is some evidence of GDP reacting early on in an anticipated way, likely because of the early rise of investment. The savings rate turns negative with a long delay, and there is little or no action before. Eventually, this completely reverses and becomes positive. In terms of employment, we do observe a rise of employment (but again with a long delay). This is a marked difference to ARS, who found a *fall* in employment.

To sum up, the most noticeable features of these responses are the long delays (between 4 to 13 years after the discovery) in the macroeconomic responses to a giant mineral discovery. There is relatively little or no action, in any of the variables, 3 years into the discovery. Moreover, in the case of the first 5 macroeconomic variables (investment, consumption, the current account, GDP, and savings) about 5 to 4 years before production starts, a standard news effect appears to kick in: this effect generates a *procyclical* response of investment, consumption, GDP and employment. At that moment, the economy borrows from abroad, generating a fall in the CA and savings. When production starts, these two variables turn positive.

3.4 Investigating the Mechanism

In this section we take a deeper look into the delay of the macroeconomic reaction to the discovery documented above. First, we simulate the ARS model with a long delay between the date of discovery and the date of production start. Second, we then estimate an ADL model with both data on mineral and oil discoveries. Because oil discoveries typically feature a smaller delay, we can use this regression to compare the macroeconomic effects of discoveries with long delays with those of short delays. Third, we merge the data on both types of discoveries and split them into long and short delays. Thus, among long (short) delays we will have both minerals *and* oil. This allows us to look more directly at the implications of the delay itself.

Figure 4: Dynamic Effects of a Mineral Discovery: GDP, Savings Rate, and Employment Ratio. Red Line: Median Delay.

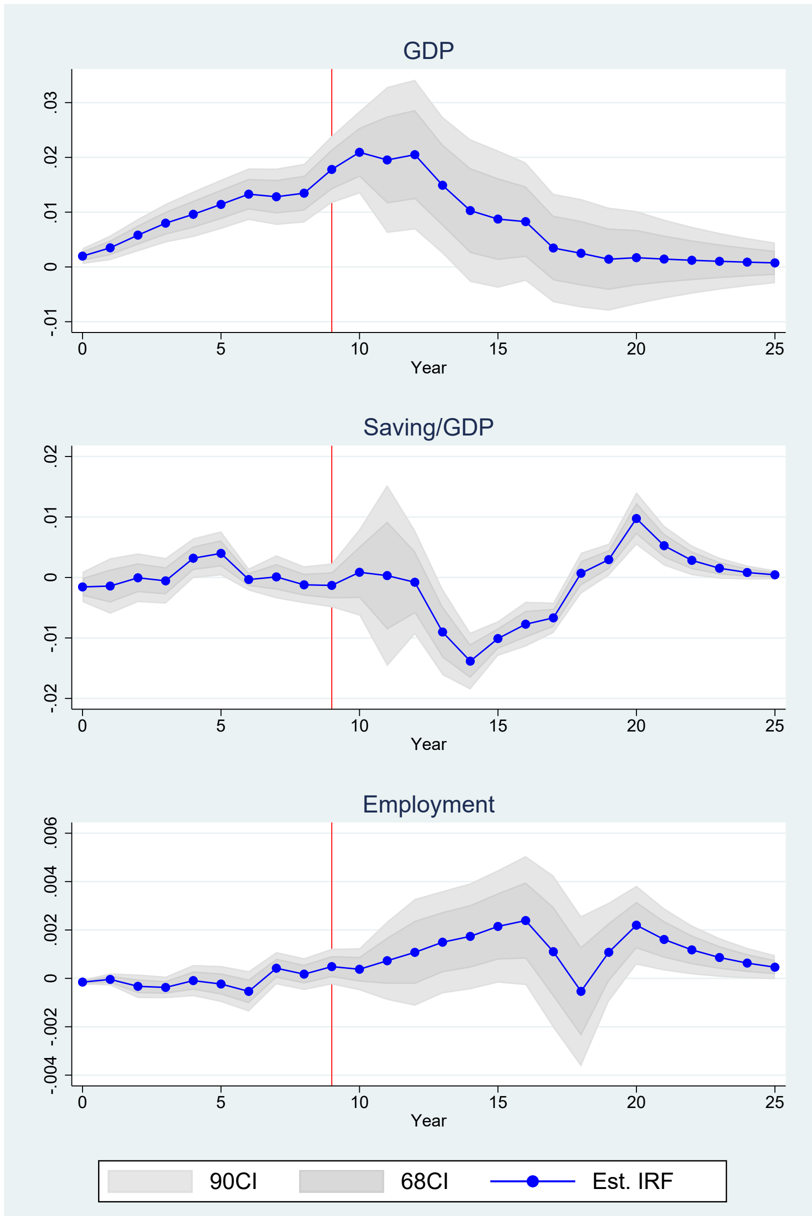

3.4.1 The Benchmark Model by Arezki, Ramey, and Sheng 2017

Arezki et al. (2017) consider a small open economy with two sectors: a commodity extraction sector and the rest of economy. The small open economy does not affect the world interest rate nor world commodity prices. We use the exact same model as ARS. However, for completeness, we reproduce the model below.

Firms

There are two sectors in the economy: a commodity extraction sector and another sector, which we will call manufacturing. The manufacturing sector uses a Cobb-Douglas, constant returns to scale technology, which depends on capital and labor

$$Y_{1,t} = A_{1,t} N_{1,t}^{\alpha_1} K_{1,t}^{1-\alpha_1}$$

The commodity sector uses capital, labor, and the *stock of reserves* R_t also with a Cobb-Douglas, constant returns to scale, production function:

$$Y_{2,t} = A_{2,t} N_{2,t}^{\alpha_2} K_{2,t}^{\alpha_k} R_t^{1-\alpha_k-\alpha_2}$$

where $0 < \alpha_1, \alpha_2, \alpha_k < 1$.

Capital accumulation in each sector is subject to the investment adjustment cost Γr la Jaimovich and Rebelo (2009):

$$K_{s,t} = I_{s,t} \left[1 - \frac{\phi}{2} \left(\frac{I_{st}}{I_{s,t-1}} - 1 \right)^2 \right] + (1-\delta) K_{s,t-1} \quad , \quad s = 1, 2$$

where s denotes the sector, $\delta \in (0, 1)$, $\phi > 0$. Adjustment costs in steady state are equal to zero.

Households

The economy is populated by identical agents who maximize lifetime utility defined over sequences of consumption C_t and hours worked N_t. Lifetime utility is

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \frac{\left(C_t - \psi N_t^{\theta}\right)^{1-\sigma} - 1}{1-\sigma} \right]$$

It is assumed that $\theta > 1, \psi > 0$ and $\sigma > 0$. We use Greenwood et al. (1988) (GHH) preferences, which shut down the wealth effect on labor supply and are now standard in open economy models. The household supplies capital and labor in a competitive market.

Households consume only good 1, but can exchange the commodity (good 2) for good 1 on international markets. Thus, the flow budget constraint is as follows:

$$B_t = (1 + r_t)B_{t-1} + (Y_{1,t} + p_t Y_{2,t}) - (C_t + I_{1,t} + I_{2,t})$$

where p_t is the relative price of a commodity determined exogenously in the world market.

To induce stationarity of foreign bond holdings, we follow the external debt-elastic interest rate proposed by Schmitt-Grohe and Uribe (2003):

$$r_t = r^* + \chi \left[exp(\bar{B} - B_{t-1}) - 1 \right]$$

Aggregation

Aggregate output, capital, investment, and domestic labor are defined as:

$$\begin{aligned} Y_t &= Y_{1,t} + p_t Y_{2,t} \\ K_t &= K_{1,t} + K_{2,t} \\ I_t &= I_{1,t} + I_{2,t} \\ N_t &= N_{1,t} + N_{2,t} \end{aligned}$$

The current account is defined as

$$CA_t = B_t - B_{t-1} = S_t - I_t$$

where S_t is aggregate saving.

Exogenous Processes

We model the delay between discovery and production using the same 'time-to-connect' notion as ARS. Known reserves appear as soon as the commodity is discovered but become productive reserves only when the roads in the case of minerals, or pipelines in the case of oil, have been connected to capital and labor. This takes time. Thus, the stock of producing reserves evolves as follows:

$$R_t = \bar{R} + R_{t-1} - Y_{2,t} + \epsilon_{t-j}$$

This relation says that producing reserves at the end of year $t-1$, R_{t-1}, are augmented with an exogenous stream $\bar{R}$, and are endogenously depleted by the commodity production, $Y_{2,t}$. ϵ_{t-j} captures the interaction of news of a commodity discovery and the time-to-connect feature; in period $t-j$, news of a discovery arrives. Known reserves rise immediately at $t-j$; but producing reserves rise only at period t (because it takes time to connect them to the capital and labor.) Thus, the lag on ϵ_{t-j} captures the key feature that reserves are not immediately available for production when the news about the discovery is revealed. We assume $j = 9$ (i.e. the median delay) for mineral discoveries.

Results

For the parametrization, we follow ARS and assume that the sectors' TFP $A_{s,t}$, $s = 1, 2$, and the commodity price p_t stays constant over time. Hence the only shock driving the economy is the discovery news shock ϵ_{t-j}. The rest of the parametrization is reported on ARS, p. 114.

The typical delay is 9 years for mineral discoveries. Figure 5 shows the impulse responses of investment, consumption and the CA from the model. Figure 6 shows the responses of GDP, the savings rate and employment. The impulse responses of investment, the current account, GDP, and saving are delayed. The response of consumption and employment is muted. All this is in line with our empirical findings. The figures also show the responses for the typical delay in the case of oil discoveries (5 years). This shows that the macroeconomic responses for minerals essentially shift by 4 years.

Figure 5: The green dashed (blue solid) line presents the impulse responses in the model with 9 (5) years delay.

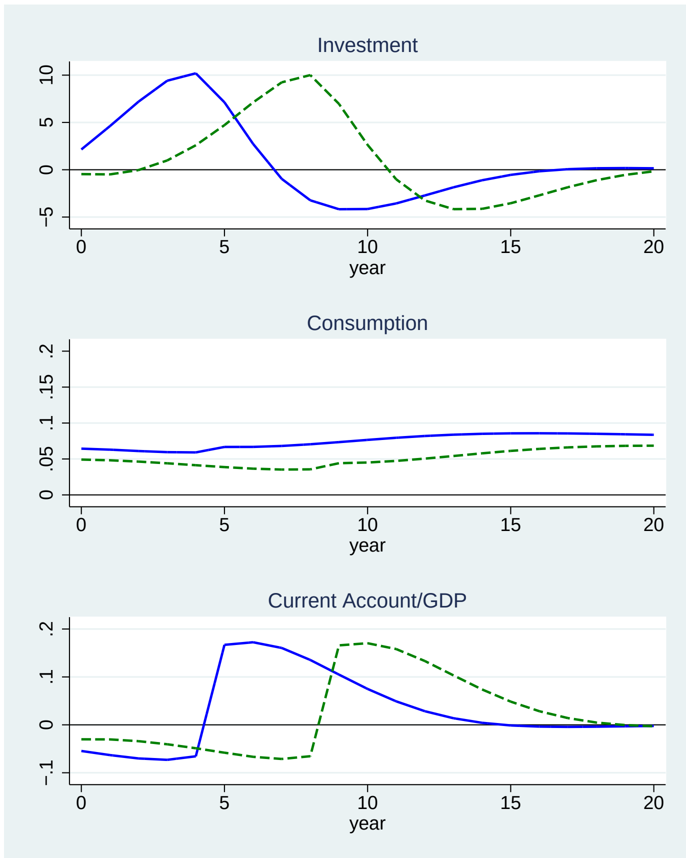

3.4.2 Comparing the Macroeconomic Effects of Giant Minerals vs. Oil Discoveries

We now compare the dynamic effects of minerals and oil. Given the results of the model in the previous section, we expect discoveries with a long delay (as minerals) to have nil or small effects on impact and for a number of years, and then a gradually appearing news-type effect several years before the date of production. Instead, with respect to discoveries with a short delay (oil), we expect a news effect to appear (almost) immediately after the discovery.

Figure 6: The green dashed (blue solid) line presents the impulse responses in the model with 9 (5) years delay.

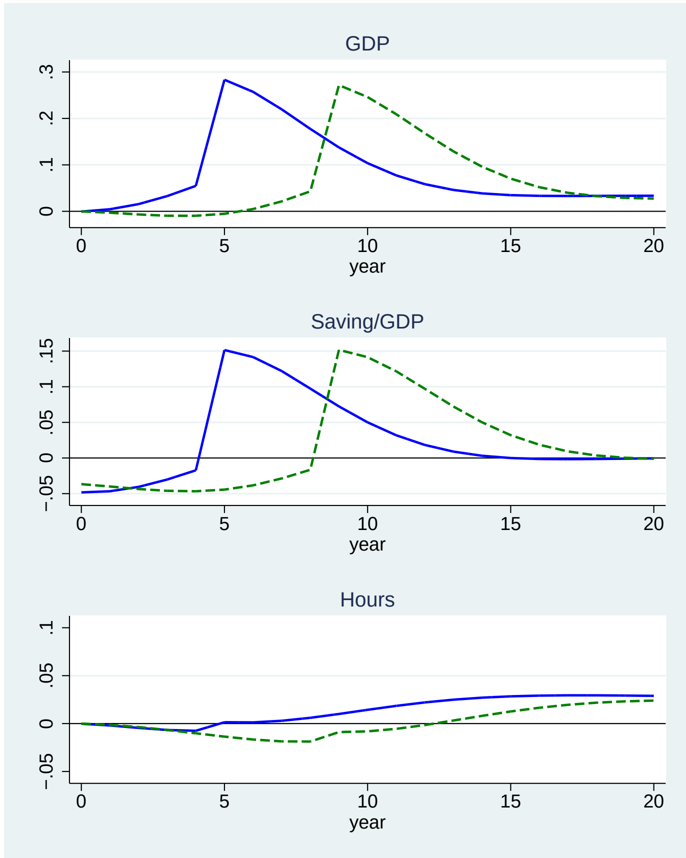

In order to do this, we merge our minerals data set with the giant oil discoveries used in ARS, and estimate the regression model

$$y_{it} = ay_{i,t-1} + B(L)Min_{it} + C(L)Oil_{it} + \alpha_i + \mu_t + \varepsilon_{it} \tag{3.2}$$

where Min_{it} is the NPV over GDP of a giant mineral discovered in country i in year t, and Oil_{it} is the NPV over GDP of a giant oil discovery in country i in year t. $B(L)$ and $C(L)$ are pth- and qth- order lag polynomials, $p, q \geq 0$. In our baseline results we pick $p = 20$ and $q = 10$. 10 is the number of lags used by ARS. (The other variables have the same definition as in our previous baseline regression (3.1) above.)

Figure 7 shows the IRFs of investment, consumption, and the CA. The results in the case of minerals (right column) look quite similar to our baseline in Section 3.3. This basically tells us that controlling for oil discoveries does not affect our baseline results. The results in the case of oil (left column) look similar to the IRFs reported by ARS (p. 128-9). Aside from the same type of remark as before—that controlling for mineral discoveries do not affect ARS' baseline results—we mainly note how delayed the response to mineral discoveries is. Essentially, the qualitative shape of the responses of the three variables is similar, but just delayed by about four years in the case of minerals. Aside from this, it is noticeable that the response of investment in the case of minerals is significantly larger that in the case of oil.

We do not report the other IRFs for the other variables for brevity. The conclusions are similar.

Given the results of the model, it is tempting to interpret the different timing in the reaction of macroeconomic variables to the type of discovery as being caused by the production delay. However, for this to be interpreted causally, we need to check that macroeconomic determinants do not affect the chances of finding one type of discovery and the other. Unreported regressions of the type of discovery on country GDP-per-capita, degree of macroeconomic openness, population density, etc. suggest this is indeed the case. For instance, Figure 8 shows that there is no visible correlation between the (log) GDP per capita and the probability of no discovery, one discovery of either type, or both types of discovery. Also, as a reminder of Figure 1, notice that the type of discovery is well-spread out around the world, and therefore it seems largely uncorrelated with geography—at least it is *not* the case that some type discoveries are concentrated per continent. Thus, although we do not formally reject reverse causality, it seems at first pass plausible to interpret the differences between the responses to a giant mineral vs. oil discovery causally. Next, we will study what happens when we split the sample by long and short delays.

3.4.3 Using the Delay as Explanatory Variable

One issue regarding the previous regression is that the type of discovery is confounded with the delay. Thus, it is not clear whether a giant mineral discovery generates a delayed macroeconomic response because of the longer delay of production compared to oil, or for other reasons. In order to investigate the effect of the delay on the macroeconomic responses more directly, here we merge the minerals and oil data and then split them according to long vs. short delays. We investigate this route in this short section.

Figure 7: Comparing Mineral and Oil Discoveries: Dynamic Effects on Investment, Consumption, and Current Account. Red Line: Median Delay.

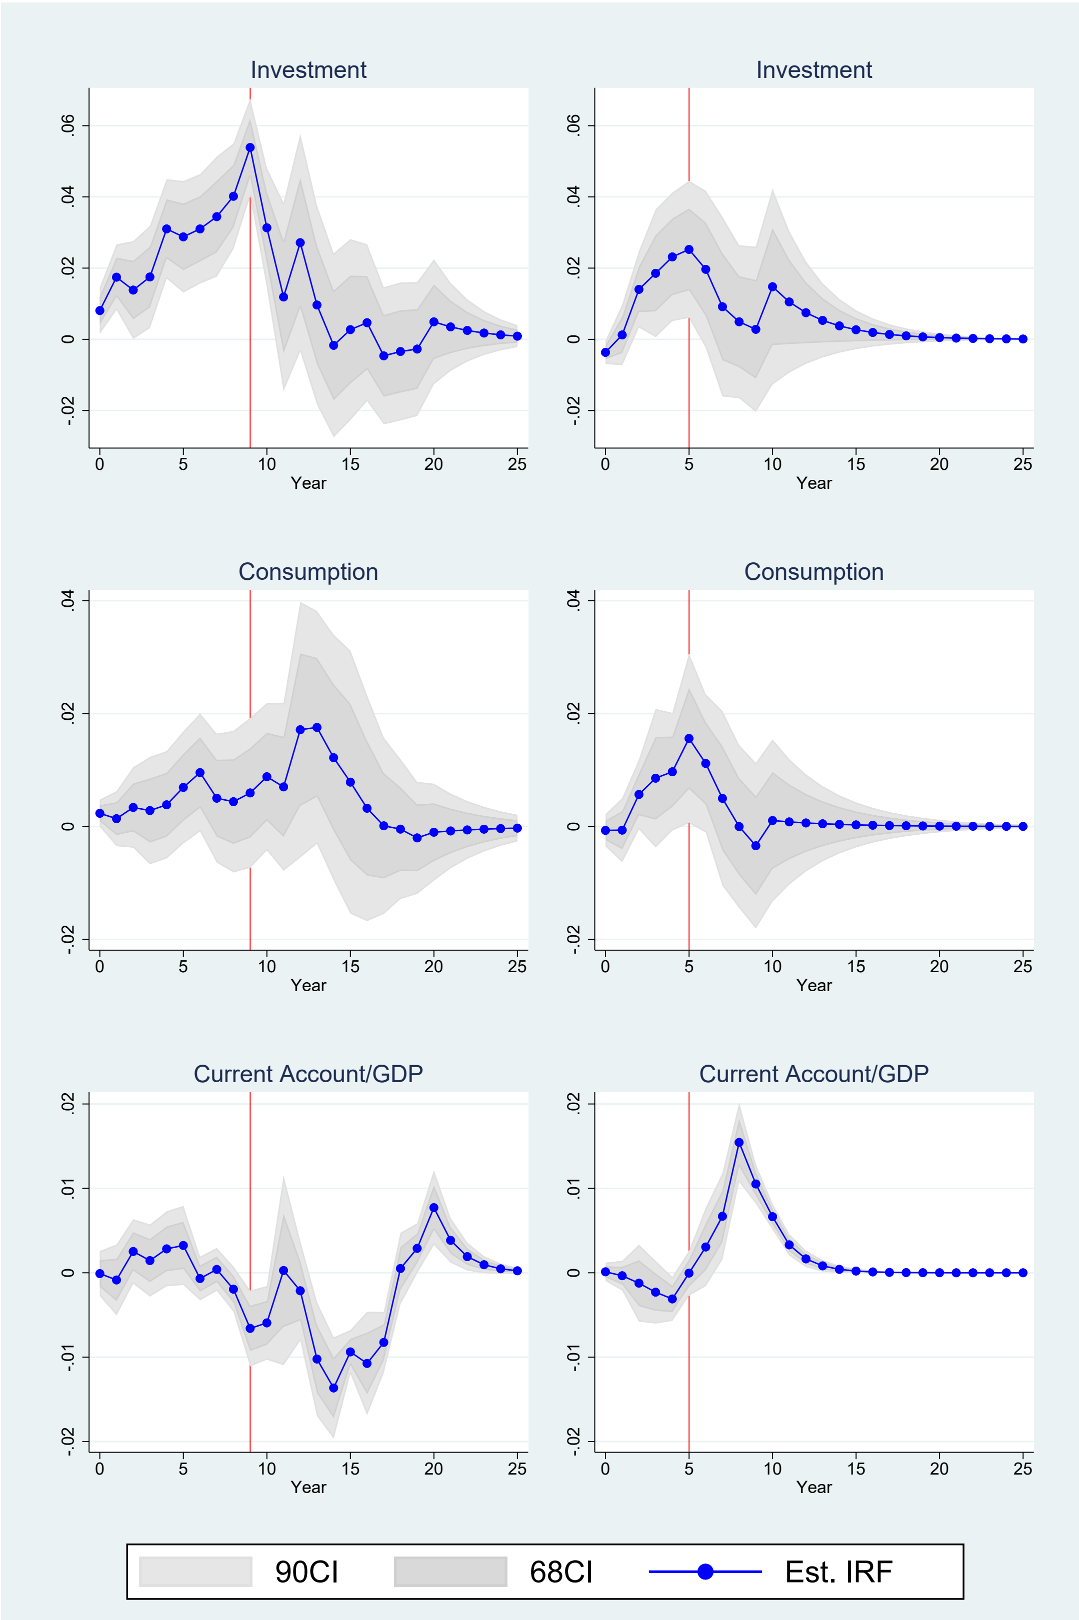

Figure 8: Log GDP-per-Capita and Type of Discovery Around the World, 1960–2015

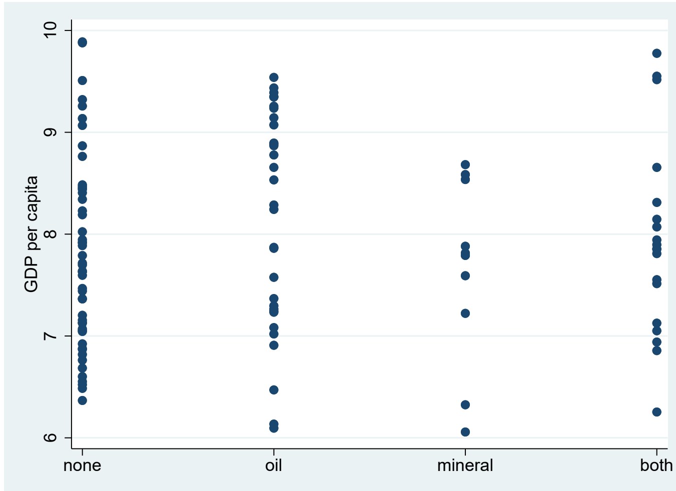

To this end, we run the regression

$$y_{it} = a y_{i,t-1} + B(L) Long_{it} + C(L) Short_{it} + \alpha_i + \mu_t + \varepsilon_{it} \tag{3.3}$$

where $Long_{it}$ is the NPV over GDP of a commodity (mineral or oil) discovered in country i at t and with a delay above the 50th percentile of the delay distribution, and $Short_{it}$ is the NPV over GDP of an oil discovery in country i at t and with a delay below the 50th percentile of the delay distribution. We remove the discoveries without information about the delay for this regression. Also, in the definition of NPV we are careful not to mix discoveries that fall in the long or short delay category. $B(L)$ and $C(L)$ are pth- and qth-order lag polynomials, $p, q \geq 0$. In our baseline results we pick $p = 25$ and $q = 10$, the reason being that the median delay by category are 13 and 3, for $Long_{it}$ and $Short_{it}$, respectively. (The other variables have the same definition as in our previous baseline regression (3.1) above.)

Ideally, we would like to have exogenous variation on the delay. Empirically, it quite possible is that the production delay is mainly technological, and therefore independent of other macroeconomic determinants. However, we have not yet made further efforts to rule out reverse causation. In any case, the IRFs obtained after estimating (3.3) by OLS are plotted on Figure 9.

Regarding investment, the response for long delays is well estimated and shows once again, little initial reaction, followed by a gradual increase *before* the median date of production start. The response for short delays is more noisy, but investment clearly rises much earlier and then also dies out faster.

Turning to consumption, the responses for long delays are again more precisely estimated. In this case, there is a clear rise of consumption several years before the median delay, and peak around the median delay, and then it goes back to steady state. This is a comforting finding given that in all previous results there was not a clear rise of consumption when compared to investment. In this case, the rise of consumption is smaller than in the case of investment, but it is precisely estimated. It is also worthwhile noticing how consumption rises several years before the median delay (more than 5). The response of consumption to discoveries with a short delay is noisier.

The CA, in the case of long delays, turns negative and then positive, but these movements are quite delayed. In the case of short delays, the response is immediately negative, and then turns positive around the median delay.

To sum up, we find qualitatively similar results in this regression using the delay as explanatory variable (regression (3.3)) to when comparing minerals (typically long delay) vs. oil (typically short delay). Longer delays imply, generally, a delayed macroeconomic response to the discovery shock.

3.5 Discussion

The previous empirical results have shown that giant mineral discoveries imply a news-type reaction of macroeconomic variables, but this reaction arrives with a significant delay compared to giant oil discoveries. Given the evidence that, typically, mineral discoveries have a longer delay than oil discoveries to pan out, this is a likely candidate explanation for the macroeconomic delay. Although we do not have an instrument for the delay on the discoveries, we have argued that it is plausible that the production delay is likely to be largely exogenous, and model simulations based on ARS also suggest a plausible causal channel.

In this section we discuss a number of issues that remain open regarding the exact economic channels determining the delay of macroeconomic adjustment. It is comforting to find that the baseline ARS model is able to generate the delayed reaction of macroeconomic variables observed in the data. There are two main reasons the model generates this delay. First, obviously the discovery raises extracted reserves only when production starts. Thus, production in the commodity sector and also GDP only benefit from the discovery after the time-to-connect. Second, investment only rises

Figure 9: Comparing Long and Short Delays: Dynamic Effects on Investment, Consumption, and Current Account. Red Line: Median Delay.

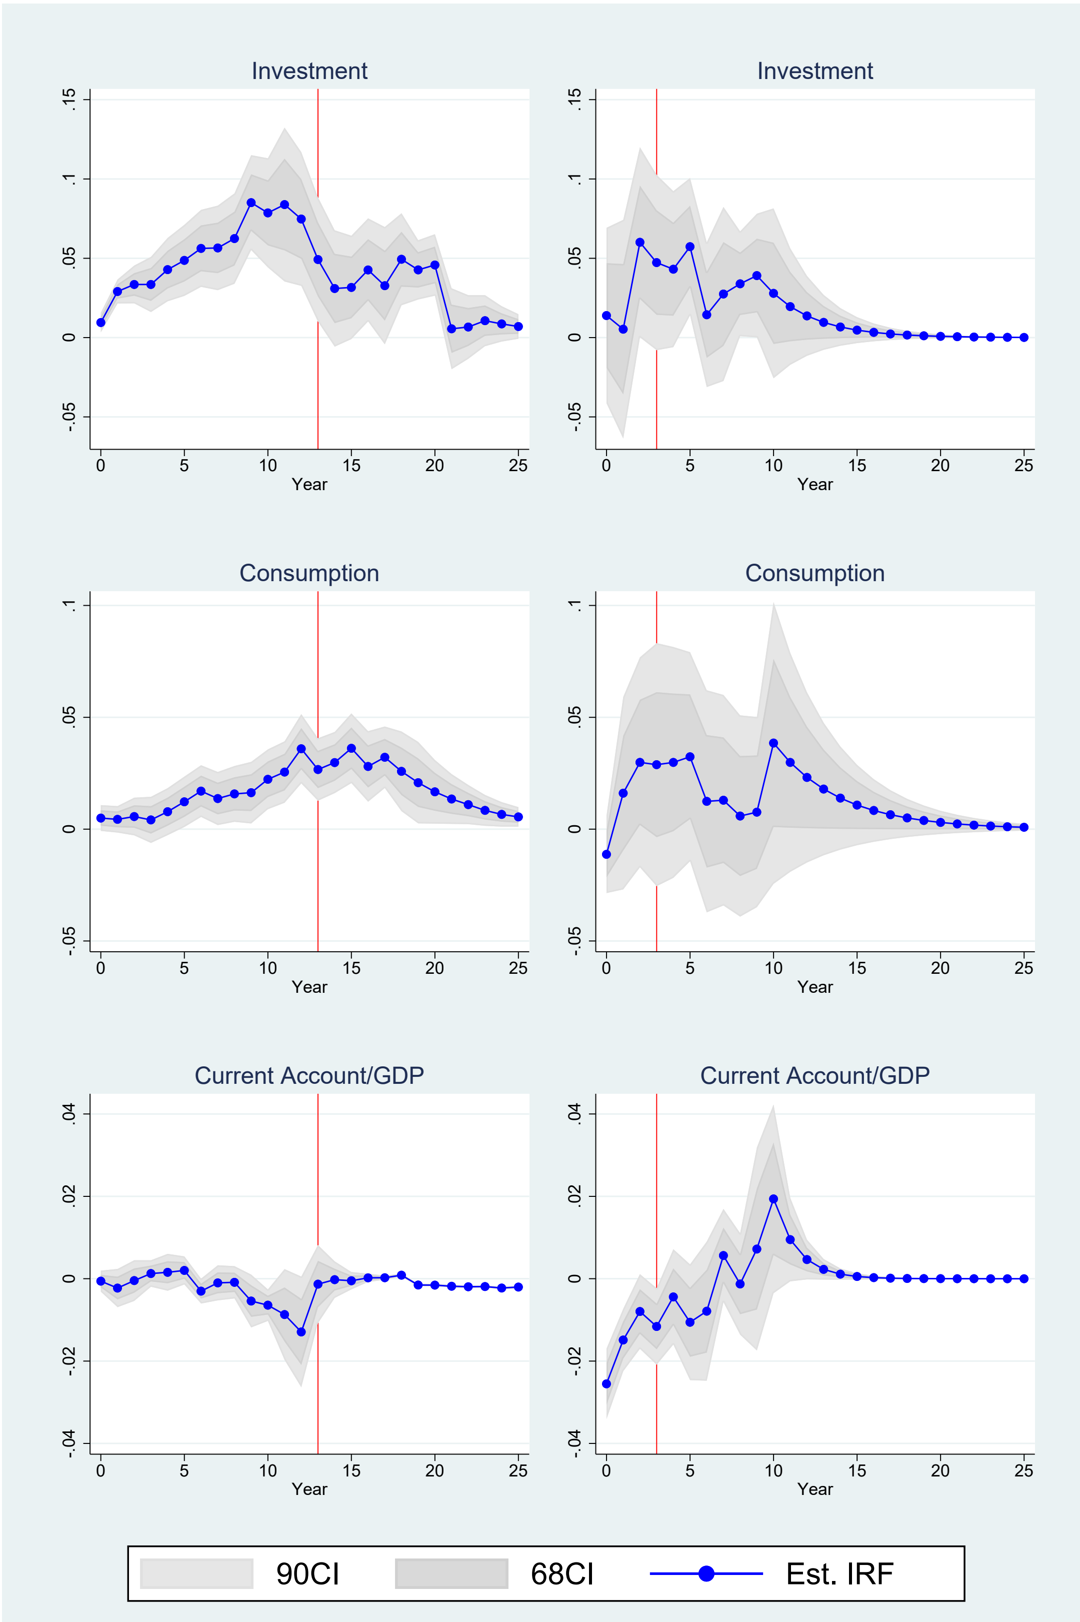

in anticipation of future returns. With adjustment costs and depreciation, it is natural to expect investment to rise only closer to the production start date. All this implies a delayed macroeconomic reaction.

The above discussion leaves open another potential channel to explore in this class of models. In principle one could expect, given that the ARS model is a small economy, that consumption would rise immediately after the discovery in anticipation of rise in future income.[17] However, ARS model the discovery as a temporary shock and thus this channel is absent. It is potentially interesting to think more about this point, and whether modeling the shock as permanent could make sense.

A last point about broader implications for macroeconomics. There is a growing interest in the literature about the degree of 'forward-lookingness' of the representative agent to policy news (Angeletos and Lian (2017); Angeletos and Huo (2018); Farhi and Werning (2019); Gabaix (2019)). In the data, we did find evidence of a very limited impact reaction to a giant commodity discovery that will take a while to materialize. Connecting this finding to forward guidance seems like a fruitful research step.

[17] See L'Huillier and Yoo (2018) for a detailed investigation of this channel.

Bibliography

Angeletos, G.-M. and Z. Huo (2018). Myopia and anchoring. *NBE Working Paper No. 24545*.

Angeletos, G.-M. and C. Lian (2017). Dampening general equilibrium: From micro to macro. *NBER Working Paper No. 23379*.

Arezki, R., V. A. Ramey, and L. Sheng (2017). News shocks in open economies: Evidence from giant oil discoveries. *The Quarterly Journal of Economics 132*(1), 103–155.

Barsky, R. B. and E. Sims (2012). Information, animal spirits, and the meaning of innovations in consumer confidence. *American Economic Review 102*(2), 1343–1377.

Beaudry, P. and F. Portier (2004). An exploration into Pigou's theory of cycles. *Journal of Monetary Economics 51*(6), 1183 – 1216.

Beaudry, P. and F. Portier (2006). Stock prices, news and economic fluctuations. *American Economic Review 96*(4), 1293–1307.

Blanchard, O. J., J.-P. L'Huillier, and G. Lorenzoni (2013). News, noise, and fluctuations: An empirical exploration. *American Economic Review 103*(7), 3045–3070.

Chahrour, R. and K. Jurado (2018). News or noise? the missing link. *American Economic Review 108*(7), 1702–36.

Farhi, E. and I. Werning (2019). Monetary policy, bounded rationality, and incomplete markets. *American Economic Review 109*(11), 3887–3928.

Gabaix, X. (2019). A behavioral New Keynesian model. *Mimeo*.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988). Investment, capacity utilization, and the real business cycle. *The American Economic Review 78*(3), 402–417.

Höök, M., S. Davidsson, S. Johansson, and X. Tang (2014). Decline and depletion rates of oil production: a comprehensive investigation. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 372*(2006).

Horn, M. K. (2014). Giant oil and gas fields of the world. http://www.datapages.com/AssociatedWebsites/GISOpenFiles/HornGiantFields.aspx.

Jacks, D. S. (2019). From boom to bust: A typology of real commodity prices in the long run. *Cliometrica 13*(2), 201–220.

Jaimovich, N. and S. Rebelo (2009). Can news about the future drive the business cycle? *American Economic Review 99*(4), 1097–1118.

L'Huillier, J.-P. and D. Yoo (2018). Where is the GE? consumption dynamics in DSGEs. *Mimeo*.

Lujala, P., J. K. Rod, and N. Thieme (2007). Fighting over oil: Introducing a new dataset. *Conflict Management and Peace Science 24*(3), 239–256.

Schmitt-Grohe, S. and M. Uribe (2012). What's news in business cycles. *Econometrica 80*, 2733–2764.

Schodde, R. (2010). *Global discovery trends 1950–2009: What, where and who found them?* Panel Discussion: Exploration expenditures are increasing, but discoveries are not. Why? MinEx Consulting.