Belief Disagreement and Business Cycles

Joao Guerreiro*

February 6, 2023

Most Recent Version [Dropbox][1]

Abstract

This paper studies how belief disagreement across households affects aggregate demand. I develop a model in which households are heterogeneously exposed to business cycles and show that the impact of disagreement can be summarized by a simple statistic–*correlated disagreement*–which captures the correlation between beliefs and individual business-cycle exposure. I endogenize disagreement via heterogeneous attention, which implies that attention increases with exposure. So, correlated disagreement is positive. Then, I show that disagreement amplifies general-equilibrium effects and acts as a propagation mechanism amplifying business cycles. I also provide evidence of this positive correlation using survey data on expectations. To quantify the implications of disagreement, I extend the analysis to a Heterogeneous-Agent New Keynesian model featuring multiple sources of heterogeneity. I show that belief disagreement can substantially amplify business-cycle fluctuations. Finally, I show that targeting spending to the most cyclical workers can significantly increase the spending multiplier.

Keywords: Belief disagreement, business cycles, fiscal policy, New Keynesian, heterogeneous agents, bounded rationality.

*Northwestern University. Email: jguerreiro@u.northwestern.edu.
I am indebted to Martin Eichenbaum, Sergio Rebelo, Matthew Rognlie, George-Marios Angeletos, and Guido Lorenzoni for their support and guidance. I am deeply grateful to Pedro Teles for his support and discussions. I would also like to thank Cassiano Alves, Adrien Auclert, Bence Bardóczy, Gadi Barlevy, Riccardo Bianchi-Vimercati, Lawrence Christiano, Carlos da Costa, Christian Hellwig, Kyle Herkenhoff, Rohan Kekre, Nicholas Kozeniauskas, Per Krusell, Simon Mongey, Joao Monteiro, Pooya Molavi, Laura Murphy, Juan Pablo Nicolini, Alessandro Pavan, Giorgio Primiceri, Miguel Santana, Laszlo Tetenyi, and Michael Weber as well as seminar participants at several institutions for useful comments. Financial support from the Robert Eisner Memorial Fellowship is gratefully acknowledged.

[1] Or download from website: https://sites.northwestern.edu/jpg0126/.

1 Introduction

Expectations about the future are central to macroeconomics. A variety of decisions made by households and firms fundamentally depend on what they expect for future income levels, interest rates, or inflation, e.g., consumption and savings, investment, or price setting. So, to understand how individuals make the decisions that shape aggregate outcomes, it is essential to understand how agents form beliefs. In this paper, I explore a particular facet of expectations: belief heterogeneity (or disagreement). Understanding the sources behind this form of heterogeneity is crucial to understand its aggregate implications.

Belief heterogeneity can be observed in surveys of expectations. In Figure 1, I plot the cross-sectional standard deviation of forecasts for one-year ahead income growth in two popular surveys of expectations: the *Survey of Professional Forecasters* (SPF) and the *Survey of Consumer Expectations* (SCE). Figure 1 shows two key findings: (1) the magnitude of disagreement is substantially high and (2) disagreement rises sharply during large events. The cross-sectional standard deviation is close to 1 percentage point in the SPF and 4.5 percentage points for the SCE.[2] Furthermore, there are also large fluctuations in disagreement over time. For example, SPF disagreement rises to 3 percentage points durring the Covid crisis.

Figure 1: Belief disagreement in survey data

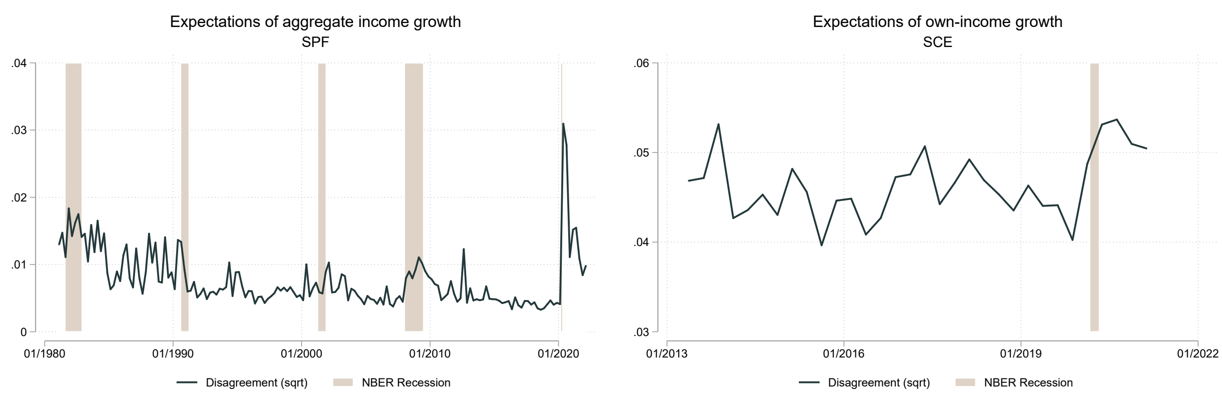

Notes: This figure displays the cross-sectional standard deviation of point forecasts for one-year ahead income growth both in the Survey of Professional Forecasters (SPF) in the left panel and the Survey of Consumer Expectations (SCE) in the right panel. See appendix A.3. for more details.

How does belief heterogeneity affect aggregate outcomes? Despite having introduced multiple forms of heterogeneity into macroeconomics models, the literature has given relatively less attention to the impact of belief heterogeneity. In part, this is a consequence of the fact that the bedrock of modern macroeconomics is the full-information and rational expectations (FIRE) assumption, which implies that everyone shares the same beliefs and so eliminates any chance for

[2] Mankiw, Reis, and Wolfers (2003) emphasize that disagreement in inflation forecasts among households is much higher than for professional forecasters.

disagreement. However, given the central importance of beliefs, understanding the sources and aggregate implications of this form of heterogeneity is crucial. In this paper, I study the impact of disagreement on aggregate demand, the transmission of macroeconomic shocks, and the efficacy of fiscal policy. I focus on aggregate demand because it has a central role in the macroeconomic transmission of shocks (see Remark 2 for a discussion).

The main findings in this paper are that (1) belief disagreement serves as a propagation mechanism which can substantially amplify business-cycle shocks and (2) the presence of belief disagreement implies that the sectoral composition of government spending affects the spending multiplier. I establish these results by considering a stylized New Keynesian model and empirical evidence. For quantification purposes, I then develop a Heterogeneous Agent New Keynesian (HANK) framework that embeds various forms of heterogeneity.

Theory I begin by considering a simple model with household heterogeneity and nominal rigidities. Households differ both in their beliefs and in their exposure to the business cycle due to heterogeneous income cyclicality. As is standard in the literature, I focus on the first-order response of this economy to shocks starting from a steady state. This model delivers two main insights.

First, I show that a single statistic summarizes the impact of belief heterogeneity on aggregate demand. This statistic, which I refer to as *correlated disagreement*, captures the extent to which belief heterogeneity is correlated with individual income cyclicality. Suppose that aggregate income (y) rises at time $t+h$. I show that the change in aggregate demand at time t, c_t, is given by:

$$\frac{\partial c_t}{\partial y_{t+h}} = \text{MPC}_h \cdot (1 + \text{CD}) \cdot \overline{\lambda}. \tag{1}$$

MPC_h denotes the marginal propensity to consume out of this increase in future income, and $\overline{\lambda}$ captures the response of average beliefs of income to the actual change in income. This term is equal to one with full information and rational expectations, but generally differs from one away from that benchmark. CD is the correlated disagreement term, which is measured as the covariance between income cyclicality and the response of individual beliefs. Intuitively, this term arises from the fact that the beliefs of individuals more exposed to aggregate-income changes (higher cyclicality) receive a larger weight in determining the response of aggregate consumption. Belief disagreement is relevant to the extent that it correlates with other individual characteristics which determine individual consumption response to macroeconomic shocks.

Second, I show that the sign of correlated disagreement determines the extent to which shocks propagate through the economy. Equation (1) shows that correlated disagreement affects the strength of the impact of future income on current aggregate demand. When correlated disagreement is positive, this channel is stronger relative to an economy where all individuals have the same average belief (and so $\text{CD} = 0$). In this case, disagreement is a propagation mechanism that amplifies the initial shock. In other words, the aggregate demand response is higher than predicted by the simple average level of attention, so the shock's impact is larger. Conversely, when correlated disagreement is negative, the general-equilibrium (GE) channel is muted relative to the homogeneous-attention economy. In this case, disagreement dampens the consequences of the shock.

Whether business cycles are amplified or dampened crucially depends on the sign of correlated disagreement. It is ex ante unclear whether we should expect the responsiveness of beliefs to be positively or negatively related to income exposure. In this paper, I provide both theoretical and empirical evidence that the correlation is positive.

Endogenous beliefs First, I endogenize beliefs via a model of behavioral inattention in the tradition of Gabaix (2014, 2016). I show that this model unambiguously predicts positively correlated disagreement. Households with higher income cyclicality are more exposed to changes in aggregate conditions, i.e., a given change in aggregate income implies a larger individual-income response for high-cyclicality than for low-cyclicality workers. So, there is an incentive for workers with high income cyclicality to track shocks more closely, i.e., pay more attention. The result is a positive correlation between attention and income cyclicality.

Evidence Second, I provide empirical support for this positive correlation. I analyze the size of forecast errors as a function of income cyclicality. Forecast errors are the difference between realized income growth and expected income growth. Through the lens of the model, the magnitude of forecast errors is informative of how attentive people are. Using survey data from the *Survey of Consumer Expectations* (SCE) and data on actual outcomes from the *Current Population Survey* (CPS), I construct average forecast errors at the state-quarterly level from 2013 to 2021. I show that the magnitude of forecast errors is decreasing in the state's average income cyclicality. A 0.1 increase in average income elasticity at the state level is associated with a 16.3 percent decrease in the magnitude of forecast errors. This result supports the implication that attention is increasing in income cyclicality.

Quantifying business-cycle amplification Since correlated disagreement is positive; aggregate demand is more responsive to changes in macroeconomic conditions, and business cycles are amplified relative to a homogeneous-attention benchmark.

To assess the quantitative relevance of the mechanisms described in the simple model, I develop a Heterogeneous Agent New Keynesian model. This model aims to develop a more realistic description of aggregate demand, the crucial object of the analysis. The quantitative model extends the previous analysis along the following dimensions: (1) it introduces incomplete markets in the form of uninsurable idiosyncratic income risk and borrowing constraints; (2) income risk increases in recessions and decreases in expansions, as emphasized by the empirical literature,[3] (3) it allows for government spending, debt, and proportional taxation; (4) it assumes that monetary

[3] For evidence on this fact, see Guvenen, Schulhofer-Wohl, Song, and Yogo (2017) and Coibion, Gorodnichenko, Kueng, and Silvia (2017).

policy is conducted according to a standard Taylor rule; and (5) it allows for time-varying prices and nominal wages, subject to standard nominal frictions. I maintain the assumption that individuals are heterogeneously exposed to aggregate conditions and that they optimize their level of attention. In this more general model, individuals must form beliefs not only about their own income, but also about the real interest rate and the tax rate.

Models with this considerable degree heterogeneity are known for their computational complexity. In this paper, non-rational heterogeneous beliefs create an additional dimension of complexity. I make the simplifying assumption that individuals choose their optimal attention in an ex-ante stage and not based on the current temporary state. Under this assumption, I can leverage recent advances in the literature to write the problem in a computationally-tractable way. The computational method is described in Section 4.6.

I calibrate the model to standard targets in the literature in addition to evidence on income cyclicality and forecast errors. I consider the response of aggregate output to a variety of standard business-cycle shocks: discount-factor, government-spending, monetary-policy, and productivity shocks. I show that amplification is more substantial when the shock is more persistent, or when the response of monetary policy to inflation is weaker. For example, calibrating productivity shocks to capture the effects of an oil shock implies that the response of aggregate income on impact is 16 percent larger because of correlated disagreement. The amount of amplification more than doubles, to 33 percent, if the response of monetary policy is weaker.

The intuition for these results is as follows. First, the more persistent the shock, the more relevant beliefs about future output are in shaping current actions. So, the correlated-disagreement mechanism becomes stronger if the shock is more long-lasting. Second, the weaker the monetary policy response, the higher the relative importance of general-equilibrium effects working from aggregate income to aggregate demand. If monetary policy is less responsive, the amplification mechanism becomes relatively more important, leading to further amplification.

Government-spending multipliers I then consider the impact of correlated disagreement on the effects of government-spending policy. In particular, I ask how the composition of government spending across worker cyclicality groups affects the government-spending multiplier, given the heterogeneity in attention.

When all individuals have the same attention or under full-information and rational expectations, the government-spending multiplier is independent of the group-composition of spending in this model.[4] In contrast, when attention is heterogeneous, the multiplier is higher if government spending targets the most cyclical groups of workers. The reason for this result is as follows. The first-round effect of spending increases workers' incomes. In response, workers that see their incomes increasing also choose to increase their consumption. The relevant statistics to determine the aggregate-demand response to this change in income are the *effective marginal propensities to*

[4]To focus on the implications of heterogeneous attention, I assume that the distribution of MPCs is orthogonal to the workers' income cyclicality. So, under FIRE, targeting becomes irrelevant. Baqaee and Farhi (2018) and Flynn, Patterson, and Sturm (2021) study the implications of MPC heterogeneity for the design of fiscal policy.

consume, i.e., the MPC weighted by the attention to income changes. It follows that, when spending targets the most cyclical groups, the government increases incomes for people with higher effective MPCs, because they are also the most attentive. The result is a larger first-round effect of government spending, which increases the government-spending multiplier.

Using the simple model, I show that if the government targets the most-cyclical workers, the multiplier can be larger than the FIRE multiplier, even if people are fully attentive to taxes.[5] Using the calibrated model, I quantify the consequences of targeting for the size of the government-spending multiplier. I show that the multiplier can depend substantially on the composition of spending. The multiplier is less than one when the government targets the least cyclical workers but rises above 1.2 if the government targets the most cyclical workers.

Literature. This paper belongs to a large literature analyzing the transmission of shocks and policies without the FIRE assumption. A large section of the literature focuses on informational frictions and shows how this deviation from FIRE affects the response of the economy to shocks, see, e.g, Woodford (2001), Mankiw and Reis (2002), Lorenzoni (2009), Angeletos and La'O (2010, 2013), Nimark (2014), or Angeletos and Lian (2018). Another strand of the literature focuses instead on bounded rationality. Ilut and Schneider (2014) analyzes the implications of ambiguity aversion, and Bianchi, Ilut, and Saijo (2021) analyzes the impact of diagnostic expectations for the transmission of business cycles. Woodford (2018) and Woodford and Xie (2019, 2022) analyze fiscal and monetary policy when people have finite planning horizons. García-Schmidt and Woodford (2019), Iovino and Sergeyev (2018) and Bianchi-Vimercati et al. (2021) evaluate the effectiveness of monetary and fiscal policies in models in which people have level-k thinking. Gabaix (2020) shows how to modify the standard New Keynesian model to account for expectational frictions in the form of cognitive discounting, while Angeletos and Sastry (2021) develop a model of shallow reasoning to analyze the relative performance of different forms of policy communication. However, these papers considered models where agents are ex-ante identical, implying that belief dispersion does not have first-order consequences to macroeconomic aggregates.

A more recent literature analyzes that question in models which allow for agent heterogeneity. Angeletos and Huo (2018) and Auclert, Rognlie, and Straub (2020) study environments with heterogeneous agents and incomplete and dispersed information, while Farhi and Werning (2019) and Farhi et al. (2020) analyze monetary and fiscal policies, respectively, with household heterogeneity and level-k thinking. Pappa, Ravn, and Sterk (2023) study a HANK model with search and match frictions and incomplete information. Bardoczy, Bianchi-Vimercati, and Guerreiro (2022) studies the stabilization effects of unemployment insurance in a general model with non-FIRE beliefs. However, all of these papers assume that the belief response is orthogonal to other sources of heterogeneity, which implies that correlated disagreement is always zero.

This paper shares the focus on the correlation between expectations and other individual char-

[5]The literature has generally found that government-spending multipliers decrease when deviating from FIRE, see Angeletos and Lian (2018), Woodford and Xie (2019), Farhi, Petri, and Werning (2020), and Bianchi-Vimercati, Eichenbaum, and Guerreiro (2021).

acteristics with Broer, Kohlhas, Mitman, and Schlafmann (2021). They document systematic heterogeneity in expectations across the income distribution about the macroeconomy and rationalize these findings with a theory of information in which individuals choose their optimal level of attention at every point in time. They argue that this leads to higher macroeconomic volatility. Instead, I focus on systematic heterogeneity in expectations across individuals with different income cyclicalities and assume that attention is a permanent characteristic of households. I derive in closed form the implications of heterogeneous beliefs in a simple model and calibrate a HANK model to study the impact of this form of heterogeneity for business cycles and fiscal stabilization policy.

There is a large literature on the empirical determinants of expectations using survey data, see, e.g., Coibion and Gorodnichenko (2012, 2015), Bordalo, Gennaioli, Ma, and Shleifer (2020), or Angeletos, Huo, and Sastry (2021). Relative to this literature, this paper provides empirical evidence on a determinant of heterogeneous attention to the macroeconomy and its implications for correlated disagreement. For further evidence on belief disagreement, see, e.g., Zarnowitz and Lambros (1987) and Mankiw et al. (2003), or more recently, Bordalo et al. (2020), Angeletos et al. (2021), and D'Acunto, Hoang, Paloviita, and Weber (2019).

This paper shares the interest in targeted spending multipliers and the effects of the composition of government spending with a growing literature. Ramey and Shapiro (1998) focus on the implications of costly capital reallocation across sectors. Cox, Müller, Pastén, Schoenle, and Weber (2020) analyze how the composition of spending affects the spending multiplier due to heterogeneous degrees of nominal rigidities in a multi-sector model. Baqaee (2015) and Bouakez, Rachedi, and Santoro (2020) study how the production network affects the impact of government spending. Baqaee and Farhi (2018) and Flynn et al. (2021) focus instead on the role of MPC heterogeneity. In this paper, I contribute another perspective regarding heterogeneous attention and show how this implies that the composition of spending affects the associated multiplier.

Finally, there has been a long tradition of studying the implications of belief heterogeneity about asset returns in asset pricing and macrofinance following the seminal contributions of Harrison and Kreps (1978) and Scheinkman and Xiong (2003). Most closely related, Caballero and Simsek (2020) analyze the implications of disagreement about financial market returns for aggregate demand. Hassan and Mertens (2017) study how small correlated forecast errors across traders can be amplified by financial trading and distort real investment. They also show that their mechanism can deliver time-varying belief disagreement in a way that is consistent with empirical evidence. Instead, I focus on beliefs about future income and on how heterogeneity in those beliefs affects aggregate demand. In this paper, I abstract from the interaction between belief heterogeneity and the valuation of financial assets.

Outline. The paper is organized as follows. In Section 2, I develop the simple model. In Section 3, I present the empirical evidence. In Section 4, I introduce the quantitative model. In Section 5, I develop the implications for the spending multiplier. Finally, Section 6 concludes. The proofs for all propositions are contained in the appendix.

2 Simple model

In this section, I describe the simple model. This model allows me to characterize the main mechanisms transparently and in closed form. In Section 4, I generalize this framework to a more complete HANK framework and use that model for quantification purposes. Throughout the paper, I restrict attention to the first-order response of the economy to shocks, starting from the flexible price steady state.

Consider a simple New Keynesian economy in discrete time $t = 0, 1, ...$ For tractability, I assume that nominal wages are fully rigid.

This economy is populated by a continuum of households $i \in [0, 1]$. Households have preferences over consumption, $C_{i,t}$, and labor $N_{i,t}$, are given by

$$\mathcal{U}_i = \sum_{t=0}^{\infty} \left(\prod_{s=0}^{t-1} \beta_{i,s} \right) \left[u(C_{i,t}) - v(N_{i,t}) \right], \tag{2}$$

where $u(C) = C^{1-\sigma^{-1}}/(1-\sigma^{-1})$ and $v(N) = N^{1+\psi^{-1}}/(1+\psi^{-1})$, $C_{i,t}$ and $N_{i,t}$ denote individual consumption and labor hours, respectively, and σ and ψ are the elasticity of intertemporal substitution and the Frisch elasticity, respectively. I assume that the steady-state discount factor $\beta \in (0, 1)$ is perturbed by discount-factor shocks and $\beta_{i,t}$ captures the effective subjective discount factor between periods t and $t + 1$. These demand shocks are the only disturbance present in this simple model. Furthermore, for simplicity, I also assume that the path of $\{\beta_{i,t}\}$ is realized at time zero, so there is no aggregate uncertainty in this economy from $t = 0$ on.

For simplicity, I assume that the production function is linear in labor $Y_t = N_t$, where Y_t and N_t denote aggregate income and labor, respectively. The goods market clearing condition is given by:

$$Y_t = \int_0^1 C_{i,t} di, \tag{3}$$

and labor market clearing requires $N_t = \int_0^1 N_{i,t} di$.

I assume that the economy is initially at a steady state and normalize the steady state output to one, $Y = N = 1$.

2.1 Firms

Firms are perfectly competitive and maximize profits, given by: $P_t Y_t - W_t N_t$. An interior solution to a firm's problem requires $W_t = P_t$, where P_t denotes the final good's price. Note that these assumptions also imply that, in equilibrium, there are no profits and that the real wage is given by:

$$w_t = \frac{W_t}{P_t} = 1. \tag{4}$$

Because wages are fully rigid, there is no price inflation:

$$\frac{P_{t+1}}{P_t} = \frac{W_{t+1}}{W_t} = 1. \tag{5}$$

2.2 Monetary policy

I assume that monetary policy directly controls the nominal interest rate i_t. Furthermore, to keep the analysis simple, I assume that the monetary authority keeps the nominal interest rate constant and equal to the steady-state interest rate:

$$i_t = r = \beta^{-1} - 1. \tag{6}$$

Remark 1. *It is well known that there may be multiple equilibria when monetary policy pegs the nominal interest rate. In this simple model, I maintain this interest-rate peg assumption and sidestep discussions regarding indeterminacy by assuming that the economy converges back to its flexible price steady state. In the model of section 4, I assume that monetary policy is given by a standard Taylor rule.*

2.3 Households

Labor income Each household belongs to a group $g = 1, ..., n$. The total mass of group g is given by π_g, where $\sum_g \pi_g = 1$. If household i belongs to group g, then their labor income is given by:

$$Y_{g,t} = w_t N_{g,t}, \tag{7}$$

where group labor supply, $N_{g,t}$, satisfies $\sum_g \pi_g N_{g,t} = N_t$.

Because nominal wages are fully rigid, labor supply is determined by firm demand, i.e., the amount of labor supplied needs to meet firms' labor demand at this nominal wage. To model heterogeneous income cyclicality, I assume that changes in aggregate demand have a different incidence across groups. Formally, I assume that:

$$N_{g,t} = \Gamma_g(N_t). \tag{8}$$

I assume that in a steady state, all people work the same number of hours $\Gamma_g(N) = N$. Furthermore, the income elasticity of group g is defined as:

$$\gamma_g \equiv \Gamma'_g(N) \geq 0, \tag{9}$$

which must satisfy $\sum_g \pi_g \gamma_g = 1$. If we think of g as different sectors, then this assumption generates sectoral heterogeneity in the incidence of business cycles, sidestepping the exact microfoundations necessary to achieve this result.[6]

[6] A way of achieving this result would require developing a multi-sector model in which workers had sector-specific skills and non-homothetic preferences.

The household's time-t flow of funds constraint is given by

$$C_{i,t} + A_{i,t+1} = Y_{g,t} + (1+r)A_{i,t}, \tag{10}$$

where $A_{i,t}$ denotes savings from period $t-1$ brought to period t. They are also subject to a standard no-Ponzi games condition $\lim_{T\to\infty} \frac{A_{i,T}}{(1+r)^T} \geq 0$.

Beliefs At every point in time, the household must anticipate the behavior of future variables which are relevant to their problem. In particular, the household must anticipate the discount-factor shock, labor income, and the real interest rate. I maintain the assumption that people have the correct expectations regarding their preference shock $\beta_{i,t}$. Furthermore, since real interest rates remain constant at their steady-state value, I assume that people correctly anticipate that $r_t = r$. So, the household's forecasting problem is simply one of forecasting their own future income.

For now, I let beliefs about future income be arbitrary and write $E_{i,t}[\cdot]$ as the expectation operator given household i's beliefs. In line with the literature, I assume that people know their current level of income. This assumption implies that markets clear in the current period given households' present consumption and savings decisions, and that basic macroeconomic identities hold. In what follows, I begin by describing household behavior and equilibrium properties for these beliefs. In subsection 2.5, I endogenize beliefs and show what it implies for disagreement and aggregate dynamics.

Consumption and savings decisions The household's utility maximization is standard; they choose consumption and savings to maximize expected utility $E_{i,t}[\mathcal{U}_i]$ subject to (10).

In Appendix A.1, I show that to first order the individual demand can be written as:

$$c_{i,t} = (1-\beta)\left\{\sum_{h=0}^{\infty}\beta^h E_{i,t}[y_{g,t+h}] + \beta^{-1}a_{i,t}\right\} + \sigma\sum_{h=0}^{\infty}\beta^{h+1}r^n_{i,t+h}, \tag{11}$$

where lower-case letters c_t, $y_{g,t}$, and $a_{i,t}$ denote the deviation of consumption, income, and assets from their steady-state values, respectively, and $r^n_{i,t} \equiv -d\log(\beta_{i,t})$ is the discount-factor shock. The interpretation of this equation follows from standard Permanent Income Hypothesis logic: $(1-\beta)$ denotes the household's marginal propensity to consume and $\sum_{h=0}^{\infty}\beta^h E_{i,t}[y_{g,t+h}]$ denotes the household's expected permanent income. The term $\beta^{-1}a_{i,t}$ denotes the household's financial wealth, while the final term $\sum_{h=0}^{\infty}\beta^{h+1}r^n_{i,t+h}$ captures the demand-shift induced by the discount-factor shock, which is multiplied by the intertemporal elasticity of substitution σ. Finally, note that $y_{g,t} = \gamma_g y_t$ where y_t denotes the deviation of aggregate output from steady state. I assume that workers understand that $y_{g,t} = \gamma_g y_t$.

2.4 Aggregation and equilibrium

Because this economy features constant wages and prices, equilibrium output is fully demand determined, i.e., output and employment adjust so as to clear the goods market $n_t = y_t = c_t$ where c_t denotes aggregate demand:

$$c_t \equiv \int_0^1 c_{i,t} di,$$

and where $c_{i,t}$ is given by (11). So, the crucial object is aggregate demand.

In characterizing aggregate demand, it also becomes evident how belief heterogeneity matters for aggregates. In what follows, I emphasize two results: (1) within-group belief heterogeneity is irrelevant, and (2) correlated disagreement determines the strength of general equilibrium forces working through aggregate income. To show this, I proceed in steps: first, I compute average consumption in group g:

$$\bar{c}_{g,t} \equiv \int_{i \in g} \frac{1}{\pi_g} c_{i,t} di,$$

and then aggregate across groups $c_t = \sum_g \pi_g \bar{c}_{g,t}$.

Aggregating individual demand (11) for all members of group g, implies that:

$$\bar{c}_{g,t} = (1-\beta) \left\{ \sum_{h=0}^{\infty} \beta^h \overline{E}_{g,t}[y_{g,t+h}] + \beta^{-1} \bar{a}_{g,t} \right\} + \sigma \sum_{h=0}^{\infty} \beta^{h+1} \bar{r}^n_{g,t+h}. \tag{12}$$

This is exactly the same expression as (11), except that individual beliefs, assets, and discount-factor shock have been replaced by group average belief, $\overline{E}_{g,t}[y_{g,t+h}]$, assets, $\bar{a}_{g,t}$, and demand shock, $\bar{r}^n_{g,t+h}$, respectively. Only the average belief of group g affects the group's average consumption. In other words, belief heterogeneity within the groups is irrelevant in the sense that the dispersion of beliefs around the average belief does not affect aggregate demand. The intuition for this result is as follows. More optimistic people perceive a higher permanent income than pessimistic people do and so choose to consume more today. However, because all households have the same MPC, the higher demand of relatively optimistic people exactly offsets the lower demand of relatively pessimistic ones. So, this type of belief dispersion does not affect average group consumption.

Next, we want to use average group demand to find aggregate demand. First, note that by asset market clearing:

$$\int_0^1 a_{i,t} di = \sum_g \pi_g \bar{a}_{g,t} = 0.$$

Furthermore, aggregating expectations of future income implies that:

$$\sum_g \pi_g \cdot \overline{E}_{g,t}[\underbrace{y_{g,t+h}}_{=\gamma_g y_{t+h}}] = \sum_g \pi_g \cdot \gamma_g \cdot \overline{E}_{g,t}[y_{t+h}] = (1 + \text{CD}_{t,h}) \cdot \overline{E}_t[y_{t+h}], \tag{13}$$

where $\overline{E}_t[y_{t+h}]$ denotes the economy-wide average belief, and $\text{CD}_{t,h} = \mathbb{C}\text{ov}\left(\gamma_g, \overline{E}_{g,t}[y_{t+h}] / \overline{E}_t[y_{t+h}]\right)$, denotes the covariance between income cyclicality γ_g and "normalized" average beliefs. I call this term *correlated disagreement*. It captures the extent to which belief disagreement is correlated with individual income cyclicality. It follows that aggregate demand is given by:

$$c_t = (1-\beta)\sum_{h=0}^{\infty}\beta^h \cdot (1+\text{CD}_{t,h}) \cdot \overline{E}_t[y_{t+h}] + \sigma\sum_{h=0}^{\infty}\beta^{h+1} r^n_{t+h}, \tag{14}$$

and $r^n_t \equiv \int_0^1 r^n_{i,t} di$ denotes the average discount-factor shock. Aggregate demand is not solely a function of average beliefs and the demand shock. Correlated disagreement determines the aggregate response of consumption to changes in income, i.e., it affects the strength of the general-equilibrium channel. The reason for this result is intuitive. Because agents are heterogeneously exposed to the cycle, not all beliefs matter equally. The beliefs of more cyclical individuals are more relevant for aggregate consumption than those of less cyclical individuals, which follows from the fact that, given a change in aggregate income, more cyclical workers have higher changes to individual income and adjust their individual consumption by more than less cyclical workers. Correlated disagreement captures exactly the term that corrects for the fact that beliefs of more cyclical workers receive a higher weight in determining aggregate demand. For instance, the weight on the beliefs of a worker with zero income cyclicality would be zero because this worker does not adjust their individual consumption in response to changes in their beliefs about aggregate income.

Equating aggregate demand to aggregate output, we find

$$y_t = (1-\beta)\sum_{h=1}^{\infty}\beta^{h-1} \cdot (1+\text{CD}_{t,h}) \cdot \overline{E}_t[y_{t+h}] + \sigma\sum_{h=0}^{\infty}\beta^{h} r^n_{t+h}, \tag{15}$$

which solves for output at time t, given beliefs future output and the discount-factor shock.

Note that under full information and rational expectations, then $\overline{E}_{i,t}[y_{t+h}] = \mathbb{E}_t[y_{t+h}] = y_{t+h}$. This fact also implies that $\text{CD}_{t,h} = 0$. So, under FIRE, this model collapses to an as-if representative agent model where heterogeneous income cyclicality is irrelevant. More generally, as long as beliefs are homogeneous, we find that $\text{CD}_{t,h} = 0$. So the economy behaves as an as-if representative agent model with average beliefs, and heterogeneous income cyclicality would be irrelevant. Instead, suppose that beliefs are heterogeneous, but all workers have the same income cyclicality. It follows that $\text{CD}_{t,h} = 0$. So, the model collapses to an as-if representative agent model with non-FIRE beliefs, where belief heterogeneity is irrelevant conditional on average beliefs. This logic shows that only the combination (and correlation) of heterogeneous income cyclicality and heterogeneous beliefs affects aggregate output. One form of heterogeneity without the other does not affect aggregate quantities.

To further analyze the effects of correlated disagreement, it is useful to make a further structural assumption on beliefs. I assume that average beliefs of group g move proportionally with the rational expectations belief:

$$\overline{E}_{g,t}[y_{t+h}] = \lambda_{g,t}\mathbb{E}_t[y_{t+h}]. \tag{16}$$

This assumption can be satisfied by several widely used models expectations which deviate from FIRE. For example, this assumption nests: (1) incomplete and dispersed information following the tradition of Lucas (1972), (2) rational inattention in the tradition of Sims (2003) in which individuals obtain signals about their permanent income, (3) sticky information as in Mankiw and Reis (2002), (4) behavioral inattention or sparsity as in Gabaix (2014, 2016), or (5) shallow reasoning as in Angeletos and Sastry (2021).[7] I refer to $\lambda_{g,t}$ as the "attention" of individuals in group g, with the understanding that in different models of beliefs it may be a consequence of different microfoundations.

For simplicity, I further assume that attention is constant over time $\lambda_{g,t} = \lambda_g$. The average level of attention is given by $\overline{\lambda} = \sum_g \pi_g \lambda_g$. This object has been the focus of study in the empirical and theoretical literature, where the consensus is that $\overline{\lambda} < 1$. In this paper, I focus instead on the correlated disagreement term. Under these assumptions, the correlated disagreement term is also constant over time:

$$\text{CD}_{t,h} = \mathbb{C}\text{ov}\left(\gamma_g, \lambda_g/\overline{\lambda}\right) \equiv \text{CD}. \tag{17}$$

Proposition 1. *Suppose that beliefs satisfy* (16) *with* $\lambda_{g,t} = \lambda_g$, *then*

$$y_t = (1-\beta)\sum_{h=1}^{\infty}\beta^{h-1}\cdot(1+CD)\cdot\overline{\lambda}y_{t+h} + \sigma\sum_{h=0}^{\infty}\beta^h r^n_{t+h}. \tag{18}$$

1. *If $CD = 0$, then the economy behaves as if it was populated by a representative agent with the average level of attention $\overline{\lambda}$. Correlated disagreement is zero if one of the following conditions hold: (1) attention is constant $\lambda_g = \overline{\lambda}$ for all g, (2) income cyclicality is constant $\gamma_g = 1$ for all g, or (3) attention and cyclicality are orthogonal.*

2. *If $CD > 0$, then the effects of changes in future output are amplified with respect to an economy with homogeneous attention, i.e., it is as if the* MPC *was higher. Correlated disagreement is positive if λ_g is increasing in γ_g.*

3. *If $CD < 0$, then the effects of changes in future output are dampened with respect to an economy with homogeneous attention, i.e., it is as if the* MPC *was lower. Correlated disagreement is negative if λ_g is decreasing in γ_g.*

Proposition 1 summarizes the effects of correlated disagreement. If correlated disagreement is zero, the economy is equivalent to a representative-agent (RA) economy with the average level of attention $\overline{\lambda}$.[8] Interestingly, this situation occurs whenever there is no heterogeneity in income

[7]The assumption outlined in equation 16 implicitly assumes that attention to output changes at different dates does not depend on the forecasting horizon, h. This assumption is always satisfied if people obtain information that helps them determine their permanent income. As it turns out, this assumption is not consequential for the analysis, and it is possible to allow the attention parameter to also vary with the horizon. I elaborate on this more complex model in Appendix A.5.

[8]I call this case the RA benchmark, with the understanding that under some models of beliefs, it would require belief heterogeneity (e.g., incomplete and dispersed information or sticky expectations). However, this type of belief heterogeneity would not be consequential for aggregate outcomes.

cyclicality or there is no heterogeneity in attention. This fact also demonstrates that each type of heterogeneity in isolation would be irrelevant to first-order output dynamics.

Instead, when these two forms of heterogeneity are present, there can be departures from the RA benchmark. If attention increases with income cyclicality, then the individuals most affected by business cycles have higher levels of attention, which implies that their beliefs move by more. This correlation then results in an amplification of general-equilibrium forces when compared to the RA benchmark. If attention decreases with income cyclicality, then the individuals most affected by business cycles have lower levels of attention, implying that their beliefs move by less. Compared to the RA benchmark, this correlation results in a dampening of general-equilibrium forces.

We can also think about correlated disagreement as affecting the effective marginal propensity to consume out of income changes at the aggregate level. In the RA economy, the MPC out of changes to aggregate income would be $(1-\beta)\overline{\lambda}_t$, which is the micro-level MPC multiplied by the average attention to aggregate income changes. Instead, when correlated disagreement is present, it is as if the MPC was $(1-\beta)(1+\mathrm{CD})\overline{\lambda}$. This term is higher if more cyclical households are more attentive and it is lower if more cyclical households are less attentive.[9]

Remark 2. *In this paper, I focus on the implications of correlated disagreement for aggregate demand. However, belief disagreement naturally affects multiple dimensions of individual decision-making and market interaction, shaping aggregate outcomes. I focus solely on the consequences for aggregate demand to cleanly characterize this particular channel and because aggregate demand is known to play a crucial role in the macroeconomic transmission of business-cycle shocks. Furthermore, stimulating aggregate demand is also the central focus of various monetary and fiscal policies.*

Propagation mechanism We can use the previous results to think about how correlated disagreement affects the propagation of shocks. The main finding is that correlated disagreement can amplify shocks if it is positive or dampen shocks if it is negative.

Suppose that the aggregate shock is such that $r_t^n = \rho^t r_0^n$, where $\rho \in [0,1]$ denotes the persistence of the shock and $r_0^n < 0$ denotes the initial impulse. I am assuming the shock is negative, which implies that the economy is entering a recession. The logic for an expansion would be symmetrical. Using equation (18), we can show that the equilibrium is

$$y_t = \frac{\rho^t \sigma r_0^n}{1-\rho\left\{\beta + (1-\beta)\cdot(1+\mathrm{CD})\cdot\overline{\lambda}\right\}}, \tag{19}$$

Note that replacing $\overline{\lambda}=1$ and $\mathrm{CD}=0$ this expression obtains $y_t^{\mathrm{FIRE}} = \rho^t \sigma r_0^n/(1-\rho)$, i.e., the full-information and rational expectations equilibrium. If, instead, we replace $\mathrm{CD}=0$ but maintain

[9] The formal logic behind these results is similar to the relationship between static MPC heterogeneity and the incidence of a change in output developed in Patterson (2019). However, note that the mechanisms presented in this paper result from heterogeneous attention rather than MPC heterogeneity. Furthermore, these mechanisms are related to attention to the economy's future performance, i.e., it is an inherently dynamic problem.

$\overline{\lambda} < 1$ we obtain the homogeneous-attention equilibrium:

$$y_t^{\text{RA}} = \frac{\rho^t \sigma r_0^n}{1 - \rho\left\{\beta + (1-\beta)\cdot\overline{\lambda}\right\}}.$$

As it is well known, inattention $\left(\overline{\lambda} < 1\right)$ dampens general-equilibrium forces compared to the FIRE benchmark. This fact also implies that the response of the economy is muted versus that benchmark, i.e., the recession is less severe $y_t^{\text{RA}} > y_t^{\text{FIRE}}$.

However, in this paper, I am interested in comparing the response of the economy with disagreement versus the economy without disagreement keeping the average level of attention constant, i.e., I am interested in comparing y_t and y_t^{RA}. If CD > 0, the economy's response is larger than in the homogeneous attention economy. Instead, if CD < 0, the response is smaller than in the homogeneous attention economy. The intuition for this result is as follows. When correlated disagreement is positive, then attention and incidence are positively correlated, i.e., the individuals who are more cyclical and thus more responsive are also more attentive. In this case, as shown in the previous section, the general-equilibrium effects are amplified, leading to larger cuts in aggregate consumption and, thus, a larger recession. When correlated disagreement is negative, attention and incidence are negatively correlated, and the logic is exactly reversed.

To evaluate the amount of amplification generated by correlated disagreement, I define amplification as the proportional response relative to the RA benchmark:

$$\mathcal{A}_t \equiv \frac{y_t - y_t^{\text{RA}}}{y_t^{\text{RA}}}. \tag{20}$$

Proposition 2 summarizes the main results regarding the amount of amplification generated by correlated disagreement.

Proposition 2. *Suppose that $r_t^n = \rho^t r_0^n$, then amplification, as defined in equation (20), is constant over time and given by*

$$\mathcal{A}_t = \frac{(1-\beta)\,\rho\cdot CD\cdot\overline{\lambda}}{1-\rho\left\{\beta + (1-\beta)\,(1+CD)\cdot\overline{\lambda}\right\}}. \tag{21}$$

Furthermore,

1. *Amplification is increasing in correlated disagreement, $d\mathcal{A}_t/dCD > 0$.*

2. *Amplification is increasing with persistence, ρ, if and only if correlated disagreement is positive, $sign\left(d\mathcal{A}_t/d\rho\right) = sign\left(CD\right)$.*

3. *Amplification is decreasing with the discount factor if and only if correlated disagreement is positive, $sign\left(d\mathcal{A}_t/d\beta\right) = -sign\left(CD\right)$.*

Proposition 2 shows that when correlated disagreement is positive, the response of output in the economy is amplified by the presence of disagreement. Instead, when correlated disagreement is negative, the response of output in the economy is dampened by correlated disagreement.

Furthermore, I also show that the higher is correlated disagreement, the more amplification is generated, which follows from the logic described above. Finally, I provide two additional results.

First, I show that the effect of shock persistence, ρ, depends on the sign of correlated disagreement. If correlated disagreement is positive, higher shock persistence leads to further amplification. Instead, if correlated disagreement is negative, higher persistence leads to less amplification. This result follows from the fact that the higher the shock's persistence, the stronger the effect of expectations about the future in determining present consumption. As a result, the correlated disagreement propagation mechanism is stronger the more persistent shocks are.

Second, I show that the effect of a higher marginal propensity to consume, i.e., a lower discount factor, also depends on the sign of correlated disagreement. A higher marginal propensity to consume, if correlated disagreement is positive, leads to larger amplification. Instead, if correlated disagreement is negative, higher MPCs lead to less amplification. A higher marginal propensity to consume makes the general-equilibrium effects through which correlated disagreement operates more relevant. It follows that, with a higher MPC, the impact of correlated disagreement is more substantial. Thus, if CD is positive, it leads to more significant amplification; if CD is negative, it leads to less amplification.

This discussion emphasizes the central role of the sign of correlated disagreement in determining the amount of amplification in this economy. If correlated disagreement is positive, then it amplifies business cycles. Instead, if correlated disagreement is negative, it dampens business cycles. But what sign should we expect for correlated disagreement? It is ex ante unclear whether the correlation between beliefs and business-cycle exposure should be positive or negative. In what follows, I provide both theoretical and empirical evidence on the sign of this correlation. First, I show that, with endogenous attention, correlated disagreement is unambiguously positive. Second, I provide empirical evidence in favor of this theoretical prediction in Section 3.

2.5 Endogenous attention

I endogeneize beliefs by modelling attention following the sparsity-based bounded rationality model introduced by Gabaix (2014) and further extended to dynamic programming problems by Gabaix (2016).[10] I assume that

$$E_{i,t}\left[y_{g,t+h}\right] = \lambda_i \mathbb{E}_t[y_{g,t+h}], \tag{22}$$

where $\lambda_i \in [0,1]$ is the attention to permanent income. This equation is similar to equation (16) but cast in terms of individual beliefs. When $\lambda_i = 0$, this person "does not pay attention", and so their belief is just equal to their default belief. When $\lambda_i = 1$, this person "pays full attention" and has FIRE beliefs. When $\lambda_i \in (0,1)$, they have a partial perception of the true value of income.

Attention and consumption-savings decisions are made as follows. In the first ex-ante stage, individuals choose their optimal level of attention λ_i to minimize expected inattention costs due to misoptimized consumption-savings choices, plus a cognitive cost of attention $\kappa\lambda_i$, where $\kappa > 0$.

[10] A useful review of models of behavioral inattention can be found in Gabaix (2019).

Assuming that the cost of attention is linear simplifies the exposition but is not central and can be easily generalized. In the second stage, shocks realize, and beliefs are determined by equation (22). Given these beliefs, individuals make their consumption and savings choices at each date, which implies that consumption is given by equation (11) replacing beliefs.

Remark 3. *A strict interpretation of this behavioral inattention model requires the assumption that individuals are boundedly rational in the second stage. Otherwise, they would be able to infer the correct beliefs using their knowledge of λ_i and $E_{i,t}[y_{t+h}]$. In the language of Gabaix (2019), this model features deterministic attention and action. Instead, a large literature considers models where agents optimally choose to receive noisy signals, as in models of rational inattention in the tradition of Sims (2003). In this tradition, individuals need not be boundedly rational in the second stage. However, this type of orthogonal noise in forecasts is not consequential for the aggregate implications of the model in this paper. So, for simplicity of exposition, I eliminate the noise. A further consequence of this assumption is that agents do not have any uncertainty regarding their point forecast. Because I restrict attention to the first-order effects of shocks, this implication is also not consequential for the aggregate dynamics in this economy.*

As in Gabaix (2016), I assume that people choose their optimal level of attention to minimize the second-order losses from inattention relative to the full-information level of utility. In Appendix A.4, I show that, to second order, the costs of inattention are given by:

$$C_g(\lambda_i) \equiv -\frac{1}{2}\frac{\partial^2 v}{\partial c^2} \cdot \sum_{h=1}^{\infty}\sum_{\tilde{h}=1}^{\infty} \frac{\partial c}{\partial Y_h}\frac{\partial c}{\partial Y_{\tilde{h}}} \cdot (1-\lambda_i)^2 \cdot \gamma_g^2 \sigma_{h,\tilde{h}}, \tag{23}$$

where $\partial^2 v/\partial c^2 = \beta^{-1}u''(1)$ is the utility cost of consumption misoptimization, $\partial c/\partial Y_h = (1-\beta)\beta^h$ denotes the response of consumption today to an increase in income in h periods, i.e., the individual intertemporal marginal propensity to consume, and $\sigma_{h,\tilde{h}}$ denotes the perceived covariance between aggregate output at horizons h and $\tilde{h}$.

Intuitively, this formula can be read as follows: $\gamma_k^2\sigma_{h,\tilde{h}}$ denotes the expected magnitude of the changes to income, $(1-\lambda_i)$ captures the extent of inattention and the MPC captures how these expectational mistakes translate into consumption decisions. Finally, $\frac{1}{2}\frac{\partial^2 \tilde{v}}{\partial c^2}$ captures how these forecast errors matter for individual utility losses.

I also assume that attention creates a psychological cost. For simplicity, I assume that this cost is linear $\kappa\lambda_i$, for $\kappa > 0$. The optimal attention then solves

$$\min_{\lambda_i \in [0,1]} C_g(\lambda_i) + \kappa\lambda_i. \tag{24}$$

Proposition 3. *Optimal attention is given by*

$$\lambda_i = \lambda_g \equiv \max\left\{0, 1 - \frac{\kappa}{\Lambda\gamma_g^2}\right\}, \tag{25}$$

where $\Lambda \equiv -\frac{1}{2}\frac{\partial^2 v}{\partial c^2}\sum_{h=1}^{\infty}\sum_{\tilde{h}=1}^{\infty}\frac{\partial c}{\partial Y_h}\frac{\partial c}{\partial Y_{\tilde{h}}}\sigma_{h,\tilde{h}}$. *This expression shows that* λ_g *is increasing in* γ_g.

Proposition 3 shows that optimal attention is increasing in income cyclicality. This result is intuitive: since more cyclical people have more volatile incomes, they would suffer more from not paying as much attention to changes in economic conditions. As a result, they optimally choose to pay a higher psychological cost to pay more attention. With linear costs, it may be true that individuals with very low income cyclicality decide not to pay any attention.

Note that, under the assumptions in Proposition 3, we find that the correlated disagreement term is constant over time and always positive:

$$\text{CD} \equiv \mathbb{C}\text{ov}\left(\gamma_g, \lambda_g/\overline{\lambda}\right) > 0.$$

Corollary 1. *If attention is endogenous and given by* (25)*, then correlated disagreement is positive, CD >* 0*. This fact implies that the effects of changes in future output are amplified with respect to an economy with homogeneous attention.*

Endogeneizing attention implies that general-equilibrium effects are amplified relative to the RA benchmark with homogeneous attention. As a result, considering the determinants of heterogeneous attention implies that discount-factor shocks are amplified. In other words, disagreement behaves as a propagation mechanism amplifying the effects of shocks on the economy.

The lesson behind this model is more general than the simple framework considered here. First, the lesson that attention increases in individual exposure to shocks is general and also holds if one considers different models for endogenous attention, such as rational inattention. Second, the implication that this type of correlation should also predict a larger response than would be obtained in a simple homogeneous attention model also has broader implications than just for the framework considered here.

3 Forecast errors and income cyclicality

In this section, I present empirical evidence in favor of the main implication that correlated disagreement is positive. For this purpose, I use microdata on household expectations from the Survey of Consumer Expectations (SCE),[11] and to measure realized outcomes, I use the Current Population Survey (CPS).[12] The SCE is a household panel surveying expectations that has been running since 2013 and covers around 1,300 households each month. For comparability with the model, I use mean forecasts for one-year ahead own income growth, as measured by the New York Fed using the response to question 24.[13] I consider the magnitude of forecast errors as a function of income cyclicality. To show why forecast errors are informative of the level of attention of these individuals, I need to introduce some more notation.

[11] This data is publicly available from the New York Federal Reserve Bank (NYFed) website: https://www. newyorkfed.org/microeconomics/sce#/.

[12] This data can be obtained from IPUMS, see Flood, King, Rodgers, Ruggles, Warren, and Westberry (2021).

[13] This question asks individuals to suppose that, in 12 months, they are still working the same job and then report the probability that their income growth falls within various ranges. The NYFed estimates a probability density function using the method described in Engelberg, Manski, and Williams (2009).

Let $\Delta y_{i,t+h} \equiv y_{i,t+h} - y_{i,t}$ denote individual income growth, then the individual forecast error is defined as follows:

$$\text{FE}_{i,t} \equiv \Delta y_{i,t+h} - E_{i,t}[\Delta y_{i,t+h}]. \tag{26}$$

Note that, under the assumptions on beliefs in Section 2.5, we can write

$$\text{FE}_{i,t} = (1-\lambda_i)\mathbb{E}_t[\Delta y_{i,t+h}] + \varepsilon_{i,t+h} = (1-\lambda_i)\gamma_i\mathbb{E}_t[\Delta y_{t+h}] + \varepsilon_{i,t+h}, \tag{27}$$

where $\varepsilon_{i,t+h} \equiv \Delta y_{i,t+h} - \mathbb{E}_t[\Delta y_{i,t+h}]$ denotes the unforecastable component of income growth.

I derive the relationship between the magnitude of forecast errors, $\text{FE}^2_{i,t}$, and income cycliality as follows. Fixing an individual, the expected magnitude of forecast error is given by:

$$\mathbb{E}\left[\text{FE}^2_{i,t}\right] = (1-\lambda_i)^2\gamma_i^2 C + \sigma^2_{i,\varepsilon},$$

where $\sigma^2_{i,\varepsilon}$ denotes the variance of $\varepsilon_{i,t+h}$ and C is a strictly positive constant. So, the expected magnitude of forecast errors can decrease in cyclicality only if λ_i increases in γ_i. The intuition for this result is as follows. First, note that keeping attention constant, higher cyclicality means that the magnitude of forecast errors rises via a mechanical effect on the variance of income growth. If attention is not increasing in γ_i, then the magnitude of forecast errors unambiguously increases. If attention is increasing in income cyclicality, then it provides a force in the opposite direction, pushing the magnitude of forecast errors closer to zero. Only if this opposing force is sufficiently strong is it possible to observe the magnitude of forecast errors declining with γ_i.

In evaluating this implication in the data, the ideal experiment would be to be able to know individual income cyclicality and match that to belief data. However, due to data constraints, this is not possible. Instead, I exploit the state-level dimension to merge the two separate surveys.

First, I construct the state's average income elasticity $\overline{\gamma}_{S,t}$ using individual-level regressions to back out elasticities at the industry level using detailed micro-data on individual income growth from the Current Population Survey (CPS). I then aggregate these elasticities using state industry shares also measured in CPS data. More details about this procedure can be found in appendix B.2. Then, using the SCE, I construct the average income growth forecast at the state level by aggregating all forecasts within a quarter, $\overline{E}_{S,t}[\Delta y_{i,t+h}]$ where S stands for state and t denotes the respective quarter. In the SCE, we observe one-year ahead forecasts, which implies that the horizon is four quarters, $h = 4$. Then, I use direct observation of individuals within a state to construct the average yearly income growth using $\Delta y_{S,t+h}$ using CPS data. I define the state-level corrected forecast error analogously to the individual one

$$\overline{\text{FE}}_{S,t} \equiv \Delta y_{S,t+h} - \overline{E}_{S,t}[\Delta y_{i,t+h}]. \tag{28}$$

The main result can be found in Figure 2 and is presented as a binscatter plot. This figure's y-axis is normalized so that 1 represents the average squared corrected-forecast-error.

Figure 2: Income cyclicality and the magnitude of corrected-forecast errors

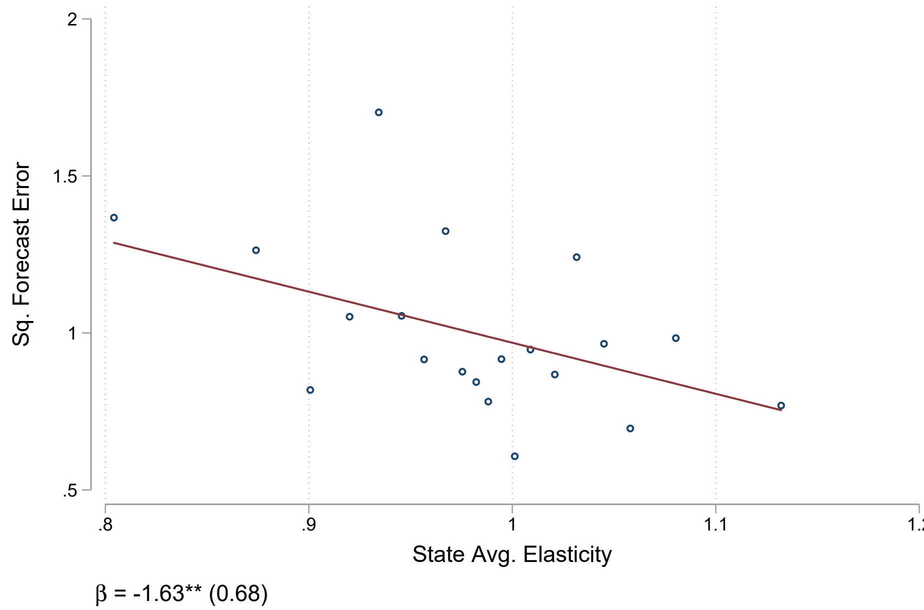

Notes: This figure shows the relationship between states' average income cyclicality and the magnitude of forecast errors. See text for more details.

Figure 2 shows that the magnitude of forecast errors falls with the state's average income cyclicality. This means that the average forecast error is lower in states with higher income cyclicality.The coefficient is statistically significant and economically significant: a 0.1 increase in the state's cyclicality is associated with a decrease in the magnitude of forecast errors by 16.3 percent on average. This result suggests that the average level of attention is higher in states with higher levels of income cyclicality. In Appendix B.3, I show that this result is robust to including a variety of controls including quarter fixed effects, the share of high-skill workers, among others.

4 Quantitative model

In Section 2, I emphasize the role of correlated disagreement in determining the strength of GE effects and the amplification of business-cycle shocks. To deliver clean results, the model in that section is intentionally stylized. In this section, I generalize the simple model in various dimensions to evaluate the quantitative relevance of the main mechanisms described in this paper. The extensions included in this section are done to achieve a more realistic description of aggregate demand, the crucial object in the analysis.

The modeling approach is based on the rapidly expanding literature on Heterogeneous-Agent New-Keynesian (HANK) models.[14] I augment a standard HANK model with the two ingredients

[14]See, e.g., Gornemann, Kuester, and Nakajima (2016), McKay and Reis (2016), Guerrieri and Lorenzoni (2017), Auclert (2019), McKay, Nakamura, and Steinsson (2016), Kaplan, Moll, and Violante (2018), Werning (2015) or Auclert,

present in the simple model. First, I assume that each household belongs to a group g which determines their income cyclicality. Second, I assume that households may be inattentive to future economic variables relevant to their decision-making. As in the simple model, I choose parameters to normalize the steady-state level of output to one, $Y = 1$.

4.1 Households

I extend the household description in Section 2.3 to allow for incomplete markets and countercyclical uninsurable idiosyncratic income risk. The economy is inhabited by a continuum of infinitely-lived households indexed by $i \in [0, 1]$.

Labor income Each household belongs to a group $g = 1, ..., n$. The average labor income of group g is given by:

$$Y_{g,t} = w_t N_{g,t}, \tag{29}$$

where $N_{g,t} = \Gamma_g(N_t)$ and $\gamma_g \equiv \Gamma'_g(1)$. To model income risk, I assume that individual households draw idiosyncratic productivity states $z_{i,t}$ from a finite support and with Markov transition matrix Π^z. Note that this transition matrix is constant across groups. I let $\pi^z(z)$ denote the steady-state mass of households in state z and normalize productivity levels so that $\sum_z \pi^z(z) z = 1$. If a household in group g has productivity $z_{i,t}$, then their pre-tax labor income is given by

$$y(z_{i,t}, Y_{g,t}) = \chi(z_{i,t}, Y_{g,t}) \cdot z_{i,t} \cdot Y_{g,t}, \tag{30}$$

where the function $\chi(z, Y)$ satisfies:

$$\chi(z, 1) = 1. \tag{31}$$

This function allows us to parameterize the cyclicality of income risk conveniently. The standard assumption in models of idiosyncratic income risk is $\chi(z, Y) = 1$. In this case, the cross-sectional variance of log-income is constant $\mathbb{V}\text{ar}(\log(y)) = \mathbb{V}\text{ar}(\log(z))$. To allow the variance of income to vary over the business cycle, I use the simple parameterization in Auclert and Rognlie (2018):

$$\chi(z, Y) = \frac{z^{\zeta \log Y}}{\sum_z \pi_z z^{1+\zeta \log Y}}. \tag{32}$$

In this case, the cross-sectional variance of the log income of individuals in group g is given by

$$\mathbb{V}\text{ar}_g\left(\log(y)\right) = \mathbb{V}\text{ar}\left(\log(z)\right) \cdot \left[1 + \zeta \log(Y_g)\right].$$

If ζ is negative, recessions lead to an endogenous widening of the distribution of income, i.e., income risk rises during a recession and falls during an expansion. Instead, if ζ is positive, income risk rises during an expansion and falls during a recession. Recent evidence suggests that $\zeta < 0$ is

Rognlie, and Straub (2018, 2020).

the empirically relevant case, see, e.g., Coibion et al. (2017) and Guvenen et al. (2017).

Assets and budget and borrowing constraints I assume that the household can trade one-period non-contingent risk-free government bonds. The household enters period t with $a_{i,t}$ assets, on which they earn the real interest rate r_t. The household's time-t budget constraint is given by

$$c_{i,t} + a_{i,t+1} = (1 - \tau_t) y(z_{i,t}, Y_{g,t}) + (1 + r_t) a_{i,t}, \tag{33}$$

where $c_{i,t}$ and $a_{i,t+1}$ denote the choices of consumption and savings, respectively, τ_t is the labor-income tax rate, and r_t is the real interest rate on savings.[15] I further assume that the household is subject to a standard no-borrowing constraint

$$a_{i,t+1} \geq 0. \tag{34}$$

Beliefs The household must forecast the discount-factor shock, labor income, tax rates, and the real interest rate. I maintain the assumption that people have the correct expectations regarding their preference shock $\beta_{i,t}$. Furthermore, the household's labor income depends on individual productivity $z_{i,t}$ and group-average income $Y_{g,t}$. I assume that people know their current productivity state and the distribution from which these states are drawn. This assumption also means that individuals are fully attentive to idiosyncratic shocks but may be inattentive to the aggregate shocks, which determine the aggregate component of their labor income.[16]

As in the simple model, I begin by allowing for general beliefs and write $E_{i,t}[\cdot]$ as the expectation operator given household i's beliefs. As before, I also assume that people know their current level of income, taxes, and real interest rates. In this model with borrowing constraints, this assumption is also useful in guaranteeing that these constraints are not violated because of misperception of current income. I discuss how to endogenize attention in Section 4.7.

Consumption and savings decisions The household's consumption and savings decisions are characterized by the problem:

$$V_{i,t}(a,z) = \max_{c,a'} u(c) - v(n) + \beta_t E_{i,t}[V_{i,t+1}(a',z')] \tag{35}$$

$$c + a' = (1 - \tau_t) y(z, Y_{g,t}) + (1 + r_t) a_{i,t} \tag{36}$$

$$a' \geq 0. \tag{37}$$

[15]Note that lower-case letters now represent the level of consumption of individual i at time t and not deviations from steady state.

[16]The assumption of full attention to idiosyncratic income shocks is extreme. I maintain this assumption to focus on how inattention to aggregates affects the transmission of business-cycle shocks. This assumption also makes the computational task easier because it implies that bounded rationality does not affect the economy's steady state. See the discussion in Section 4.6.

This problem defines a consumption and assets policy function: $c_{i,t}(a,z)$ and $a_{i,t}(a,z)$, respectively. The term β_t captures discount-factor shocks. This dynamic problem implicitly assumes that the law of iterated expectations holds at the individual level, i.e., $E_{i,t}[E_{i,t+1}[\cdot]] = E_{i,t}[\cdot]$. This implies that individuals expect not to make forecast revisions in the future.[17]

Remark 4. *This model assumes that households in different groups are symmetric across all dimensions except for their income cyclicality and expectations. In particular, I have assumed that all households draw from the same idiosyncratic productivity distribution and have the same discount factors. Furthermore, I assume that, in a steady state, all groups have the same average income. It follows that groups will be perfectly symmetric in a steady state. These assumptions also imply that the marginal propensities to consume out of own income are the same across groups. Patterson (2019) provides evidence of a positive correlation between income cyclicality and the marginal propensity to consume out of contemporaneous income. In this model, I have decided to abstract from this correlation for two reasons: (1) to emphasize the role of heterogeneous attention, and (2) because the mechanism at play in this paper relates more to forward-looking* MPC *(the* MPC *out of income in the future) for which we have fewer data. However, note that if these forward-looking* MPC *are also positively related to income cyclicality, this will work to reinforce the mechanism which is the focus of this paper.*

4.2 Firms

Firms are perfectly competitive and maximize profits. They operate a linear technology:

$$Y_t = \Theta_t N_t, \tag{38}$$

where Y_t denotes aggregate output, N_t is aggregate labor, and Θ_t denotes productivity. I maintain the assumption that firms have flexible prices, which implies that the real wage is:

$$w_t = \frac{W_t}{P_t} = \Theta_t. \tag{39}$$

This equilibrium condition implies that inflation $\pi_t = \log(P_t/P_{t-1})$ is given by:

$$\pi_t = \pi_t^w - \log(\Theta_t/\Theta_{t-1}), \tag{40}$$

where $\pi_t^w = \log(W_t/W_{t-1})$ denotes wage inflation.

4.3 Unions and sticky wages

I follow the New Keynesian sticky-wage literature and model sticky wages via a model of labor unions, as in Erceg, Henderson, and Levin (2000), Schmitt-Grohé and Uribe (2005) and Auclert et al. (2018).

[17]This assumption is also present in the anticipated-utility approach pioneered by Kreps (1998), in which agents behave as if they believe their expectations wouldn't change. See also Cogley and Sargent (2008) for a further discussion.

There is a continuum of labor unions that determine wages and labor supply. At each date, person i supplies $n_{u,i,t}$ hours of work to union u, where $u \in [0,1]$ and $n_{i,t} = \int n_{u,i,t} du$ denotes total hours of work by person i. Each union aggregates the efficiency hours of work into total hours for their specific task $N_{u,t} = \int \chi(z_{i,t}, Y_{g(i),t}) \cdot z_{i,t} \cdot n_{u,i,t} di$. The union-specific labor supply is aggregated with those of other unions by a competitive labor-market packer via a CES technology

$$N_t = \left(\int_0^1 N_{u,t}^{\frac{\mu_w - 1}{\mu_w}} du \right)^{\frac{\mu_w}{\mu_w - 1}}, \tag{41}$$

which then sells these labor services to the final goods producer at W_t.

In order to generate nominal rigidities, I assume that unions face a quadratic cost of wage adjustment $\frac{1}{2\tilde{\kappa}_w} \left(\frac{W_{u,t}}{W_{u,t-1}} - 1 \right)^2$, which is measured in terms of household utility. At every date, the union chooses a new wage $W_{u,t}$ and is required to elicit labor by each of its members according to a uniform rule $n_{u,i,t} = N_{u,t}$. I assume that the union sets the wage $W_{u,t}$ and labor supply $N_{u,t}$ to maximize the aggregate welfare valuation of its income and labor supply. Furthermore, I assume that all unions are symmetric. So, in equilibrium, all unions set the same wage. In Appendix C.1, I outline these details and show that the linearized NK wage Phillips curve is given by:

$$\pi_t^w = \kappa_w \left[\sigma^{-1} \hat{c}_t + \psi^{-1} \hat{n}_t - (\hat{y}_t - \hat{\tau}_t - \hat{n}_t) \right] + \beta \mathbb{E}_t [\pi_{t+1}^w], \tag{42}$$

where κ_w denotes the wage-stickiness parameter, and σ and ψ denote the intertemporal elasticity of substitution and the Frisch elasticity, respectively. Also, $\hat{c}_t$, $\hat{n}_t$ and $\hat{y}_t$ denote the log-deviation of aggregate consumption, labor, and output from steady state, respectively, and $\hat{\tau}_t \equiv d\tau_t/(1-\tau)$ denotes the deviation of taxes from steady state.

4.4 Fiscal and monetary policy

I assume that the government spends $\{G_t\}$, issues debt $\{B_t\}$, and taxes labor income at the rate $\{\tau_t\}$. The government budget constraint is given by:

$$G_t + (1 + r_t) B_t = \tau_t Y_t + B_{t+1}. \tag{43}$$

I assume that the government sets a path for spending exogenously and that government debt is constant at its steady-state level $B_t = B$. So, the tax τ_t adjusts every period to make this budget constraint hold. In Appendix E.3, I consider a more general case in which government debt is allowed to vary over time.

I assume that the monetary authority controls the nominal interest rate. Monetary policy is given by a Taylor interest-rate rule:

$$1 + i_t = (1 + r_t^*) \cdot e^{\phi_\pi \pi_t}, \tag{44}$$

where $\{r_t^*\}$ denotes a monetary policy shock and $\phi_\pi > 1$ is the Taylor coefficient. The real interest rate is given by:

$$(1 + r_t) = (1 + i_t)/e^{\pi_t}. \tag{45}$$

4.5 Aggregation and equilibrium

As before, aggregate demand is given by:

$$C_t = \int_0^1 c_{i,t} di,$$

and aggregate asset demand is given by:

$$A_{t+1} = \int_0^a a_{i,t+1} di.$$

Given beliefs, initial conditions on wages, government debt, and the distribution of individuals over assets and productivity states, and the path of government spending, an equilibrium is a sequence for prices $\{W_t, P_t, w_t, r_t, \pi_t\}$, aggregate quantities $\{C_t, N_t, Y_{g,t}, N_{g,t}\}$, policies $\{\tau_t, i_t\}$, and individual allocations $\{c_{i,t}, n_{i,t}, a_{i,t+1}\}$, such that households optimize, firms optimality conditions are satisfied, unions optimize, the government budget constraint is satisfied, interest rates satisfy the Taylor rule, and markets clear:

$$C_t + G_t = Y_t = \Theta_t N_t,$$
$$A_t = B.$$

4.6 Computational method

The goal is to use a calibrated version of this model to quantify the response of aggregate variables to four different shocks: discount-factor, $\{\beta_t\}$, government spending, $\{G_t\}$, monetary policy, $\{r_t^*\}$, and productivity, $\{\Theta_t\}$, shocks. There are two features of this model that make the computation difficult. The first is the presence of household heterogeneity with incomplete markets and uninsurable idiosyncratic income risk. The second is the presence of heterogeneous non-FIRE beliefs. I leverage recent contributions in the Sequence-Space representation HANK literature by Auclert, Bardóczy, Rognlie, and Straub (2021) and Auclert et al. (2020).

The computational strategy is to first solve for the stationary equilibrium of this economy and then solve for the first-order impulse-response functions for a given shock in sequence space.

Steady state In a steady state, aggregate quantities and prices are constant over time. I assume that monetary policy sets the nominal interest rate to support the zero inflation equilibrium, $\pi = \pi^w = 0$. In steady state, there are no shocks so $\beta_t = \beta$, $G_t = G$, $r_t^* = r$, and $\Theta_t = 1$.

I assume that the steady state is common knowledge. So, if the economy stays in a steady state, everyone understands that the relevant variables will remain at their steady-state levels. In other words, in a steady-state allocation, people make no errors in forecasting aggregate income, interest rates, or taxes.[18] This implies that the problem of a consumer is given by:

$$\begin{aligned} V(a,z) = &\max_{c,a'} u(c) - v(N) + \beta \mathbb{E}[V(a',z')|z], \\ &c + a' = (1-\tau) \cdot z \cdot Y + (1+r)a \\ &a' \geq 0, \end{aligned}$$

where the expectation $\mathbb{E}[\cdot|z]$ is taken with respect to the distribution of z' given z. Note that group heterogeneity is irrelevant in a steady state because all groups are symmetric. The solution to this problem determines policy rules $c(a,z)$ and $a'(a,z)$, which determine optimal consumption and asset holdings given the individual state variables.

Aggregate consumption and asset demand are given by:

$$C = \sum_z \int c(a,z) D(da,z), \quad \text{and} \quad A = \sum_z \int a'(a,z) D(da,z), \tag{46}$$

where $D(\cdot,\cdot)$ denotes the endogenous distribution of asset holdings and productivities. Note also that because all groups are symmetrical, the distribution D is constant across groups in a steady state. The market clearing conditions are:

$$C + G = Y \quad \text{and} \quad A = B.$$

Computing the equilibrium requires solving for quantities and prices that satisfy all agents' private optimality, the above steady state restrictions, and market clearing.

Transition dynamics with FIRE First, I discuss how to compute the full-information rational expectations (FIRE) equilibrium in this economy. Then, I discuss how this approach can be generalized to models with more general beliefs.

With FIRE computing the transition dynamics would require solving for the path of each variable satisfying all the conditions. I use the method described in Auclert et al. (2021) to solve the equilibrium. This requires splitting the model into "blocks" that take in specific inputs and produce other aggregate sequences as outputs. For instance, a group-g household block can be constructed taking in the sequences of discount factor shocks, group average income, taxes, and real interest rates while outputting sequences for average consumption and savings for this group.

Following Auclert et al. (2021), I construct the Jacobians, $\mathcal{J}$, of each block. These Jacobians summarize the partial derivative of a given output of the block with respect to that block's inputs. For example, one Jacobian of the household block is $\mathcal{J}_g^{C,r} = [\partial \overline{C}_{g,t} / \partial r_h]_{t=0,1,\ldots;h=0,1,\ldots}$, which summarizes how average demand of group g at each date t responds to an increase in the real

[18] This assumption can be justified by the fact that the households have spent a long time in a steady state and have thus learned the stationary equilibrium.

interest rate at time h. These Jacobians summarize the relevant responses of the different blocks of the economy to every variable. They can thus be used to compute the first-order response in this economy to various exogenous impulses/shocks. In this model, the computationally complex components are the household blocks. I elaborate on these blocks below.

We are interested in the response of group g's average consumption and asset demand to changes in the objects that are relevant to their decisions: group-average income, $Y_{g,t}$, tax rates, τ_t, real interest rates, r_t, and discount factors, β_t. I describe how the approach solves for changes in average consumption demand, with the understanding that similar expressions can be written for asset demand. The response of average consumption is given by:

$$d\overline{\mathbf{C}}_g = \mathcal{J}_g^{C,Y} \cdot d\mathbf{Y}_g + \mathcal{J}_g^{C,\tau} \cdot d\boldsymbol{\tau} + \mathcal{J}_g^{C,r} \cdot d\mathbf{r} + \mathcal{J}_g^{C,\beta} \cdot d\boldsymbol{\beta}, \tag{47}$$

where bold letters denote the column-vector of realizations of that variable for each date, e.g., $\mathbf{Y}_g = \begin{bmatrix} Y_{g,0} & Y_{g,1} & \dots \end{bmatrix}'$. Auclert et al. (2021) provide efficient ways of computing the relevant Jacobians that matter to solve for the response of this average consumption. We can then aggregate these blocks to find aggregate demand:

$$d\mathbf{C} = \sum_g \pi_g d\overline{\mathbf{C}}_g. \tag{48}$$

It follows from the fact that all groups are symmetrical in a steady state that the partial equilibrium responses summarized by the Jacobians are the same across all groups, i.e., $\mathcal{J}_g^{C,X} = \mathcal{J}^{C,X}$ for each variable X. Furthermore, as before, we know that $d\mathbf{Y}_g = \gamma_g \cdot d\mathbf{Y}$. As a result, the change in aggregate consumption is given by:

$$\begin{aligned} d\mathbf{C} &= \sum_g \pi_g \left[\mathcal{J}^{C,Y} \cdot \gamma_g \cdot d\mathbf{Y} + \mathcal{J}^{C,\tau} \cdot d\boldsymbol{\tau} + \mathcal{J}^{C,r} \cdot d\mathbf{r} + \mathcal{J}^{C,\beta} \cdot d\boldsymbol{\beta} \right] \\ &= \mathcal{J}^{C,Y} \cdot d\mathbf{Y} + \mathcal{J}^{C,\tau} \cdot d\boldsymbol{\tau} + \mathcal{J}^{C,r} \cdot d\mathbf{r} + \mathcal{J}^{C,\beta} \cdot d\boldsymbol{\beta}, \end{aligned}$$

where the equality follows from the fact that $\sum_g \pi_g \gamma_g = 1$. This result shows that, with FIRE, the aggregate demand response in this economy is exactly the same as in an economy without group income cyclicality heterogeneity. Intuitively, as in the simple model, this is a consequence of the fact that the average marginal propensities to consume are the same across groups.

Transition dynamics without FIRE However, the goal of this paper is to be able to compute the equilibrium in this economy for more general beliefs, which need not coincide with the realization (people may make forecast errors), which may be heterogeneous (to capture disagreement), and may change over time (to capture learning). In this section, I discuss how, under some restrictions, the Jacobians of the household block can be computed at almost no additional computational cost from the FIRE Jacobians. The central insight and computational method used here were originally developed in Auclert et al. (2020).

The computational complexity arises from the fact that, without FIRE, the average consumption function of group g is a function not only of the realized path for each input but also of the entire distribution of beliefs that individuals hold about this path at every point in time. Auclert et al. (2020) show that, assuming that the distribution of beliefs is orthogonal to the individual states (a, z), the Jacobians without FIRE can be computed directly from the FIRE Jacobians at almost no extra computational cost.

In this paper, I use a different representation of the Jacobians without FIRE from the one in Auclert et al. (2020). This representation is also used in Bardoczy et al. (2022), where the equivalence between the alternative forms is discussed in detail.[19] The response of group g's consumption can be written as:

$$d\mathbf{C}_g = \sum_{\mathbf{X}\in\{Y,\tau,r,\beta\}} \left[\mathcal{J}^{C,X} \cdot \underbrace{\overline{E}_{g,0}[d\mathbf{X}]}_{\text{Initial belief}} + \sum_{t\geq 1} \mathcal{R}_t^{C,X} \cdot \underbrace{\left(\overline{E}_{g,t}[d\mathbf{X}] - \overline{E}_{g,t-1}[d\mathbf{X}]\right)}_{\text{Forecast Revision at time } t} \right], \tag{49}$$

where $\overline{E}_{g,t}[\cdot]$ denotes the average expectation of group g at time t and

$$\mathcal{R}_t^{C,X} \equiv \begin{bmatrix} 0 & \mathbf{0}_t' \\ \mathbf{0}_t' & \mathcal{J}^{C,X} \end{bmatrix}.$$

Note that if the cross-sectional distribution of beliefs is orthogonal to the distribution of idiosyncratic states (a, z), then all that matters to first order is the response of average beliefs in group g. This assumption is essential in allowing us to compute the relevant Jacobians, $\mathcal{R}$, from their FIRE counterparts, $\mathcal{J}$. The FIRE Jacobian multiplies the initial beliefs $\overline{E}_{g,0}[\cdot]$ and transformations of this Jacobian then multiply the successive forecast revisions that people make at each date t, $\overline{E}_{g,t}[d\mathbf{X}] - \overline{E}_{g,t-1}[d\mathbf{X}]$. Note that, by construction, the t element of the forecast revision is the forecast error, i.e., the t-th element of the vector $\overline{E}_{g,t}[d\mathbf{X}] - \overline{E}_{g,t-1}[d\mathbf{X}]$ is $dX_t - \overline{E}_{g,t-1}[dX_t]$. So, the forecast revision term also captures the impact of forecast errors.

This expression has a natural interpretation. Note that the FIRE Jacobian $\mathcal{J}^{C,X}$ multiplies the initial beliefs. Because the shocks are unanticipated, the initial response in beliefs is unanticipated both under FIRE and for any other model of beliefs. It follows that the slopes that determine the response to the initial change in beliefs are always the same. However, as time advances, people can learn more about the shocks and revise their beliefs. Forecast revisions would never happen under FIRE because people have perfect foresight. Without FIRE, people change their views over time as they suffer forecast errors and learn more. These successive forecast revisions lead people to adjust their consumption behavior relative to their original plan. The slopes which determine the revision in consumption decisions are captured by the matrix $\mathcal{R}_t^{C,X}$, where t denotes the time in which the forecast revision occurs. This matrix implies that there is no consumption response prior to date t. This result follows from the fact that people could not have anticipated

[19] In Appendix C.2, I briefly present the details behind these results.

the forecast revision before date t when those decisions were taken. Furthermore, the way in which consumption at current and subsequent dates is revised is captured exactly by the FIRE Jacobian $\mathcal{J}^{C,X}$. Intuitively, this result is also a consequence of the fact that the forecast revision was not anticipated. The response in current and future consumptions to the forecast revision is the same as if these had been time-0 belief updates, approppriately shifted. To further develop these intuitions, I now discuss two particular cases.

First, suppose that beliefs are never updated. Then, $\overline{E}_{g,t}[d\mathbf{X}] = \overline{E}_{g,0}[d\mathbf{X}]$. Since $\overline{E}_{g,t}[dX_t] = dX_t$, it follows that $\overline{E}_{g,t}[d\mathbf{X}] = d\mathbf{X}$, i.e., the initial beliefs were correct. It follows that in this case households consume exactly the same at every date as if they had full information and rational expectations. More generally, households may make forecast revisions and errors. However, at time 0 they do not anticipate any forecast error and so they make their consumption decisions as if the initial beliefs were fully accurate. This logic shows why the FIRE Jacobian multiplies the initial beliefs.

Now, suppose that beliefs are revised at some date t:

$$\overline{E}_{g,t}[d\mathbf{X}] - \overline{E}_{g,t-1}[d\mathbf{X}] = \begin{bmatrix} 0 & \dots & X_t - \overline{E}_{g,t-1}[X_t] & \overline{E}_{g,t}[X_{t+1}] - \overline{E}_{g,t-1}[X_{t+1}] & \cdots \end{bmatrix}'.$$

The consumption response to this forecast revision is given by $\mathcal{R}_t^{C,X}$. For instance, its implications for consumption at time t are given by

$$\mathcal{R}^{C,X}_{t,(t,:)} \cdot \left(\overline{E}_{g,t}[d\mathbf{X}] - \overline{E}_{g,t-1}[d\mathbf{X}]\right) = \sum_{h\geq 0} \frac{\partial C_0}{\partial X_h} \cdot \left(\overline{E}_{g,t}[X_{t+h}] - \overline{E}_{g,t-1}[X_{t+h}]\right).$$

Note that the way in which time-t consumption responds to a forecast revision is exactly the same as the way in which time-0 consumption would react to an unanticipated perfect-foresight shock under FIRE. Intuitively, this result follows from the fact that the forecast error could not have been anticipated, and so, to first order, it leads to the same consumption response as if it was a time-0 unanticipated shock.

To impose more structure, suppose that beliefs respond proportionally to the full-information and rational expectation: $\overline{E}_{g,t}[dX_{t+h}] = \lambda_{g,t,h} dX_{t+h}$. Under this assumption, we can write:

$$\overline{E}_{g,t}[d\mathbf{X}] = \Lambda_{g,t} d\mathbf{X}, \tag{50}$$

where $\Lambda_{g,t} = \text{diag}\left(\{1, ..., 1, \lambda_{t,1}, \lambda_{t,2}, ...\}\right)$ is a diagonal matrix. It follows that:

$$d\mathbf{C}_g = \sum_{\mathbf{X}\in\{Y,\tau,r,\beta\}} \tilde{\mathcal{J}}_g^{C,X} d\mathbf{X}, \tag{51}$$

where $\tilde{\mathcal{J}}_g^{C,X} \equiv \mathcal{J}^{C,X} \cdot \Lambda_{g,t} + \sum_{t\geq 1} \mathcal{R}_t^{C,X} \cdot \left(\Lambda_{g,t} - \Lambda_{g,t-1}\right)$ are simple manipulations of the FIRE Jacobians.

4.7 Optimal attention

As in the simple model, I endogenize beliefs following Gabaix (2016). I extend that model by assuming that:

$$E_{i,t}[dX_{t+h}] = \lambda^X_{i,h} \cdot \mathbb{E}_t[dX_{t+h}] + (1 - \lambda^X_{i,h}) \cdot E_{i,t-1}[dX_{t+h}], \tag{52}$$

with initial condition $E_{i,-1}[dX_{t+h}] = 0$. The additional term implies that individuals learn over time, so new information accumulates to past knowledge. I also allow the attention variables to depend on the forecast horizon, allowing individuals to have more accurate forecasts regarding variables that are closer in time than those that are farther away. However, I maintain the assumption that attention for all states and periods is chosen once and for all in the pre-period, and households cannot re-optimize this plan in future times or states.

In Appendix C.3, I show that the utility cost of inattention takes a similar form to that found in the simple model:

$$\mathcal{C}_g(\boldsymbol{\lambda}_i, a, z) = -\frac{1}{2}\frac{\partial^2 v(c(a,z);a,z)}{\partial c^2} \sum_{X,\tilde{X},h,\tilde{h}} \frac{\partial c(a,z)}{\partial X_h}\frac{\partial c(a,z)}{\partial \tilde{X}_{\tilde{h}}} \left(1 - \lambda^X_{i,h}\right)\left(1 - \lambda^{\tilde{X}}_{i,\tilde{h}}\right) \sigma_{X_h,\tilde{X}_{\tilde{h}}}, \tag{53}$$

where $c(a,z)$ denotes the steady-state policy function.

The utility costs of inattention are now a function of the idiosyncratic asset and productivity states. I assume that the attention costs are linear $\kappa^X \lambda^X_{i,h}$. The cost of attention does not depend on the horizon but may depend on the variable being forecasted. This assumption will allow us to calibrate the model to match survey-data facts, see Section 4.8.

I assume that attention is chosen once and for all, to minimize ex-ante expected costs of inattention weighted by the ergodic distribution, i.e.,

$$\boldsymbol{\lambda}_i = \arg\min_{\boldsymbol{\lambda}} \sum_z \int \mathcal{C}_g(\boldsymbol{\lambda}, a, z) D(da, z) + \sum_{X,h} \kappa^X \lambda^X_{i,h}. \tag{54}$$

This assumption implies that attention is constant for all members of group g, $\boldsymbol{\lambda}_i = \boldsymbol{\lambda}_g$ and so beliefs are orthogonal to asset and productivity states. Under this assumption, we can use the computational method discussed in the previous section.

Following Gabaix (2014), I make the simplifying assumption that people believe the correlation across variables to be zero.[20] This assumption implies that the optimal attention can be easily solved and it is given by:

$$\lambda^Y_{g,h} = \max\left\{0, 1 - \frac{\kappa^Y}{\sum_z \int \frac{\partial^2 v(a,z)}{\partial c^2}\left(\frac{\partial c(a,z)}{\partial Y_{g,h}}\right)^2 D(da,z) \cdot \gamma_g^2 \sigma_Y^2}\right\} \tag{55}$$

[20] In the appendix, I generalize these results to allow people to perceive correlations across variables. The results are consistent with the ones in the baseline model.

and

$$\lambda_{g,h}^{X} = \max\left\{0, 1 - \frac{\kappa^{X}}{\sum_z \int \frac{\partial^2 v(a,z)}{\partial c^2}\left(\frac{\partial c(a,z)}{\partial X_h}\right)^2 D(da,z)\cdot\sigma_X^2}\right\} \tag{56}$$

for $X = \tau, r$.

Forecast errors Note that an individual's forecast error in predicting a variable h periods ahead is given by:

$$\mathrm{FE}_{i,t,t+h}^{X} = X_{t+h} - E_{i,t}[X_{t+h}] = \frac{1-\lambda_{i,h}^{X}}{\lambda_{i,h}^{X}}\cdot \mathrm{FR}_{i,t,t+h}^{X} + \varepsilon_{t,t+h}^{X},$$

where $\varepsilon_{t,t+h}^{X} \equiv X_{t+h} - \mathbb{E}_t[X_{t+h}]$ denotes the unpredictable component of forecast errors and

$$\mathrm{FR}_{i,t,t+h}^{X} \equiv E_{i,t}[X_{t+h}] - E_{i,t-1}[X_{t+h}]$$

denotes the individuals forecast revision at time t. This result means that it would be possible to obtain individual attention, $\lambda_{i,h}^{X}$, from the regressions of forecast errors on forecast revisions in Coibion and Gorodnichenko (2015), Bordalo et al. (2020), and Angeletos et al. (2021).

Remark 5. *In endogenizing attention, I assume that beliefs are chosen in an ex-ante stage, so they are not conditional on individual productivity and asset states. As discussed in the previous section, this assumption greatly facilitates the computational task. It is unclear how allowing for a correlation between* (a, z) *would affect the results in this paper. In this remark, I briefly comment on the consequences of allowing for this correlation.*

On the one hand, individuals with fewer assets or lower productivity are more likely to be borrowing constrained. It follows that they have lower MPC *out of future income. In the limit, a borrowing-constrained individual has a zero* MPC *out of future income. All else equal, this fact would imply that individuals with fewer assets or lower productivity have a lower incentive to pay attention to changes in future income than individuals with more assets or higher productivity. Similar logic to heterogeneous income cyclicality would generate an even stronger correlation between attention and responsiveness, reinforcing the results in this paper.*

On the other hand, individuals with lower assets and productivity also have higher marginal utility of consumption. It follows that consumption misoptimization is more costly for these individuals than for individuals with more assets or high productivity. This force would then work in the opposite direction mitigating the correlation.

In general, it is not clear which force dominates. Therefore, I think of the assumptions here as a useful and conservative benchmark to study the implications of heterogeneous income cyclicality.

Finally, note that we can write beliefs of an individual in group g at time t as $E_{i,t}[dX_{t+h}] =$

dX_{t+h} if $h \leq 0$ and

$$E_{i,t}[dX_{t+h}] = \sum_{s=0}^{t} \lambda^X_{g,h+s} \prod_{m=0}^{s-1} \left(1 - \lambda^X_{g,h+m}\right) \cdot dX_{t+h}, \tag{57}$$

if $h \geq 1$. It follows that this framework fits into the framework in equation (50).

4.8 Calibration

I first discuss the calibration of the economy's steady state and then elaborate on the calibration of the remaining parameters. The model is calibrated to a quarterly frequency. In steady state, I shut down shocks. So productivity is normalized to one $\Theta = 1$, the discount factor is equal to its steady-state value β, government spending is constant G, nominal interest rates are equal to real interest rates $i = r$, and inflation is zero $\pi = \pi^w = 0$.

Table 1 shows the calibrated parameters relevant to compute the steady state. I assume the household's utility function has constant elasticity over consumption and labor. This means that $u(c) = c^{1-\sigma^{-1}}/(1-\sigma^{-1})$ where σ is the intertemporal elasticity of substitution. Following Auclert et al. (2018), I set this elasticity to 0.5. The disutility of labor is given by $v(n) = \xi n^{1+\psi^{-1}}/(1+\psi^{-1})$, where ψ is the Frisch elasticity. Following Chetty, Guren, Manoli, and Weber (2011), I set the Frisch elasticity to 0.5 and calibrate the disutility parameter ξ so that the steady state features zero inflation with $Y = N = 1$. This calibration yields $\xi = 0.64$.

The productivity shocks are drawn from a discretized AR(1) process with persistence $\rho_z = 0.95$ and standard deviation $\sigma_z = 0.5$, which is in line with the parameters traditionally used in the literature. I set the interest rate to an annual rate of 2 percent or 0.5 percent quarterly. The government spending-to-GDP ratio is calibrated to 16 percent. I choose the level of assets-to-GDP and the discount factor to match an average marginal propensity to consume of 0.25. This yields $B/Y = 1.92$ and $\beta = 0.97$, which is in line with the values found in the literature.

Table 1: Calibration

Parameter	Description	Value	Param.	Description	Value
σ	IES	0.5	r	Real int. rate	0.5%
ψ	Frisch	0.5	G/Y	Spending-to-GDP	16%
ρ_z	Persistence z	0.95	B/Y	Assets-to-GDP	1.92
σ_z	St. Dev. z	0.5	β	Discount factor	0.97
ζ	Cyclicality of income risk	-0.5	ϕ_π	Taylor Coefficient	1.5
κ_w	Wage rigidity	0.0062	ρ_β	β shock – Persistence	0.9
ρ_G	Spending shock – Persistence	0.91	ρ_r	r^* shock – Persistence	0.89
ρ_Θ	Productivity shock – Persistence	0.98	ξ	Labor disutility	0.64

I assume there are 14 household groups, $n = 14$, one for each census industry group. The estimation of the elasticities γ_g follows the procedure described in Appendix B.2. I assume that the group shares π_g are equal to the shares of each industry in the US economy in 2018. I also use CPS data to estimate these shares. The results can be found in Table 2.

Table 2: Group shares and income cyclicality

	Industry	π_g	γ_g		Industry	π_g	γ_g
1	Agriculture, Forestry, Fishery	1.76%	0.05	8	Non-durable Man.	4.35%	0.80
2	Public Administration	5.84%	0.12	9	Durable Man.	7.60%	1.44
3	Bus. and Repair Services	7.10%	0.14	10	Retail Trade	14.69%	1.77
4	Prof. and Related Serv.	30.90%	0.43	11	Wholesale Trade	2.63%	2.26
5	Mining	0.58%	0.48	12	Personal Services	2.39%	2.41
6	Transp., Commun., Public Util.	7.34%	0.59	13	Finance, Insur., Real Est.	7.02%	2.45
7	Construction	6.09%	0.62	14	Ent. and Recr. Serv.	1.70%	4.24

Following Auclert and Rognlie (2018), I assume that $\zeta = -0.5$, which implies that income risk is countercyclical and provides a good fit to the empirical findings with a single parameter. Furthermore, I assume that the Taylor coefficient is $\phi_\pi = 1.5$, which is standard in the literature. Consistent with the findings in Hazell, Herreno, Nakamura, and Steinsson (2022), I set the wage flexibility parameter to $\kappa_w = 0.0062$.

I calibrate the attention cost parameters so that the average level of attention matches the regression results in Bordalo et al. (2020). To match these results, I calibrate the cost parameters κ^X, so that $\overline{\lambda}_3^Y = 0.69$ and $\overline{\lambda}_3^r = 0.45$. Because there are no forecasts for the tax rate, I cannot obtain $\overline{\lambda}_3^\tau$ in this way. Instead, I assume that attention to taxes is the same as attention to the aggregate component of income.

I consider four different shocks: to discount factors, β, to government spending, G, to the monetary policy rate, r^*, and productivity, Θ. For each shock, I assume that the initial impulse evolves geometrically over time $X_{t+1} = \rho_X X_t$, where ρ_X captures the persistence of the shock. I set the persistence for each shock in line with standard parameters in the literature. The persistence of discount factor shocks is set to 0.9, see Justiniano, Primiceri, and Tambalotti (2010). The persistence of government-spending shocks is set to 0.91, see Auclert et al. (2018). The persistence of monetary policy shocks is set to 0.89, as estimated by Auclert et al. (2020). Finally, I set the persistence of TFP shocks to 0.98, which captures an oil shock in reduced form, see Blanchard and Gali (2007).

Remark 6. *Note that the model calibration does not directly use the empirical findings in section 3. Instead, I calibrate the model to a standard target in the literature and let the forces at play in the model determine attention heterogeneity. I do not pursue a more data-driven approach to recovering heterogeneous attention for three reasons. First, due to data constraints, I cannot conduct the analysis in Section 3 directly at the sector level, which could allow us to recover attention at the sector level. It is not easy to back out these attentions from state-aggregated data due to time-varying state characteristics, such as industrial composition. Second, because we do not have multiple forecast horizons in the SCE, we cannot compute the necessary forecast revision statistics which allow us to run the regressions in Coibion and Gorodnichenko (2015), Bordalo et al. (2020), and Angeletos et al. (2021). Finally, calibrating to the SPF's empirical findings has become standard in the literature. So, this choice also maximizes the comparability of my results to those in the literature.*

4.9 Quantitative results

In this section, I present the main quantitative findings. I begin by discussing the optimal level of attention generated by the model. I then discuss how heterogeneous attention impacts the response of group demand to changes in aggregate income. Finally, I present the main quantitative results on business cycle amplification.

Optimal attention Figure 3 displays the optimal level of attention to income, taxes, and interest rates on the left, middle, and right panels, respectively, for four different groups $g = 3, 5, 8$, and 13.

Figure 3: Optimal attention

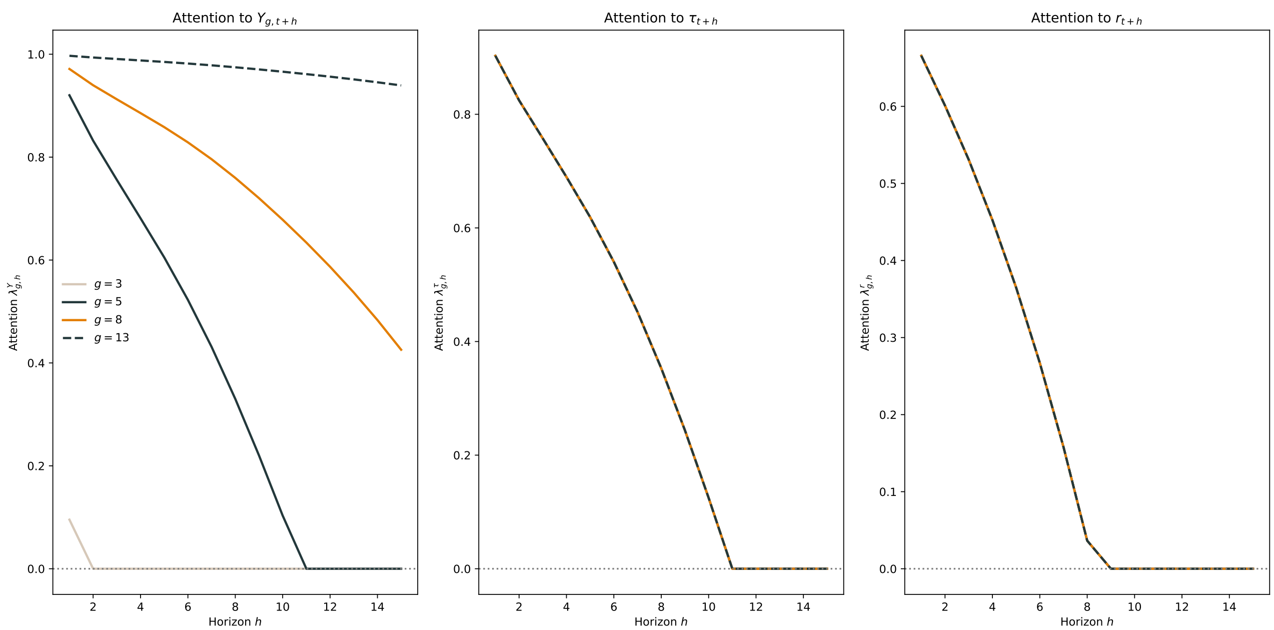

Notes: This figure displays optimal attention in the quantitative model for four different household groups $g = 3, 5, 8$ and 13, where $\gamma_3 = 0.14$, $\gamma_5 = 0.48$, $\gamma_8 = 0.8$, and $\gamma_{13} = 2.45$. The left panel displays attention to the aggregate component of income at various horizons, the middle panel displays attention to tax rates, and the right panel displays attention to real interest rates. See text for more details.

As expected from equation (56), the optimal attention to tax and interest rates is not affected by income cyclicality. So, all groups have the same levels of attention for every horizon. Instead, attention to changes in income depends on income cyclicality. Workers in more cyclical occupations choose a higher level of attention than workers who are less exposed to changes in aggregate conditions. We see that people in $g = 13$ are very close to full attention. Instead, people in $g = 3$ have such a low cyclicality that they optimally devote no attention to the aggregate component of their income from $h = 2$ on. In this model, people only disagree about income changes, not tax or interest rates. It would be easy to modify the assumptions to allow for disagreement about these other variables. However, since all agents respond equally to tax and interest rate changes, this form of disagreement would not affect aggregate outcomes in this economy.

Overall, we see that attention decreases with the horizon h of the forecast. The reason for this result is that shocks to variables that are far in the future have a lower impact on present decisions than shocks to variables that are closer in time, e.g., the marginal propensity to consume out of income two quarters ahead is higher than the marginal propensity to consume of income ten quarters ahead. As a result, people devote less cognitive effort to forecasting far-off variables. At far enough horizons, the value of predicting a variable is so tiny that individuals choose not to pay any attention, so $\lambda^X_{g,h} \to 0$ and $h \to \infty$. It is interesting to note that an additional contribution of this framework is to provide a microfoundation by which people behave as if they had finite planning horizons as in Woodford (2018) and Woodford and Xie (2019, 2022).

Response to aggregate income changes How does inattention affect the response of consumption to increases in aggregate income? To shed light on this question, in Figure 4, I compare the full-information and rational expectations response to an increase in income at time 5, with the response obtained in the economy with heterogeneous attention, for two groups $g = 3$ and $g = 13$. I also compute these responses in the counterfactual economy in which all individuals have the same level of attention (homogeneous attention), which coincides with the average level of attention in the baseline economy.

Figure 4: Consumption response

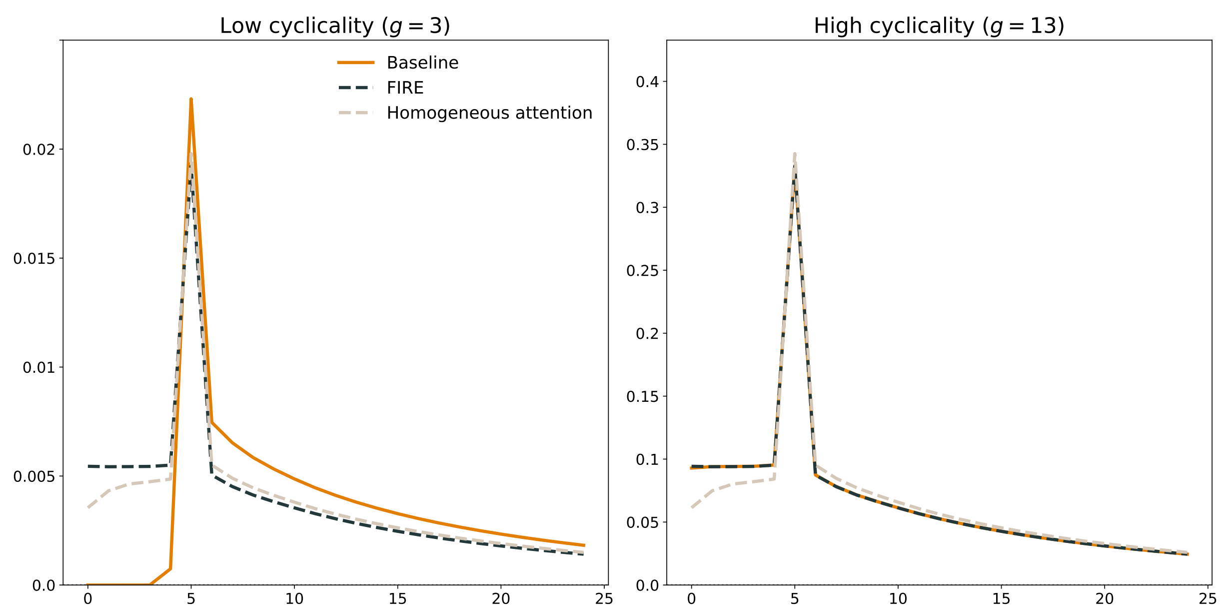

Notes: This figure displays the fifth column of the Jacobian multiplied by γ_g, or the partial-equilibrium response of consumption to an increase in aggregate income at time 5, of two groups: $g = 3$ in the left panel and $g = 13$ in the right panel, to an increase in aggregate income at time 5. For each group, the figure plots the response of consumption under full information and rational expectations (FIRE), heterogeneous attention (Baseline), and the counterfactual homogeneous attention (Homogeneous attention), which assumes that all agents have the same level of attentiveness.

Note that the shape of the FIRE responses in the two panels is exactly the same but they have different magnitudes. This result follows from the fact that, under FIRE, the Jacobians are the same for all groups. The different magnitudes result from the fact that an increase in aggregate income has a higher impact on the incomes of households with higher income cyclicality than for households with lower income cyclicality.

The same facts regarding the shape and size of the response are true for the economy with homogeneous attention. However, relative to the FIRE response, we see that the initial impact is dampened, i.e., agents consume less in anticipation of higher income in the future. Since individuals are inattentive, they do not fully incorporate how much their future income is rising into their present decicions. So, they do not consume as much early on. Since they have dissaved less relative FIRE, the increase in consumption at time 5 and subsequent dates is higher.

Figure 5: Consumption response for all groups

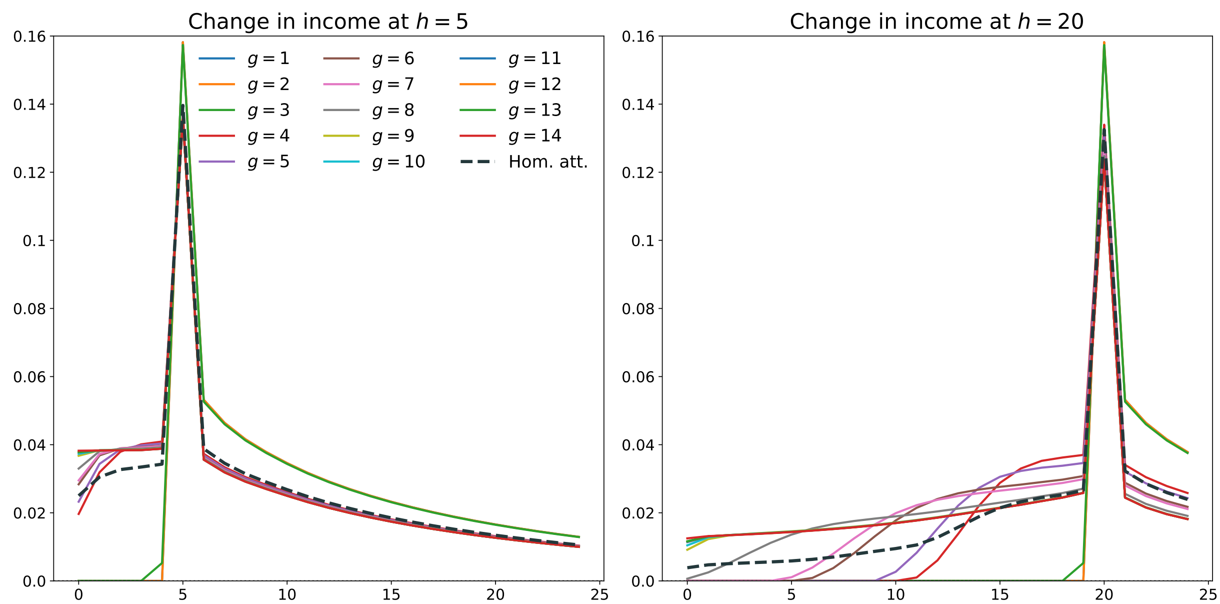

Notes: This figure displays the partial-equilibrium response of consumption for all groups to an increase in aggregate income at horizons 5 and 20 in the left and right panels, respectively. The figure also displays the response that would be obtained under homogeneous attention. The responses are divided by γ_g for comparability.

Instead, the consumption responses to an increase in aggregate income in the baseline economy are quite different. The differences are not just in magnitude but also in their shape. The low cyclicality group chooses to pay almost no attention to this income component. So, they do not increase consumption before time 4. At time 4, they become aware that there will be some increase in their future income, and so we see a mild response in average consumption. At the moment of the income increase, they finally become fully aware and increase their consumption. Because they have not consumed as much in earlier periods, they have not dissaved, so their consumption increase and time 5 and subsequent dates is higher relative to FIRE or homogeneous beliefs. Instead, the high cyclicality group displayed on the right panel has a very high level of attention to

Y_5. As a result, their consumption in the baseline economy is essentially the same as under FIRE.

Figure 5 displays the response of consumption for all groups to an increase in aggregate income at two horizons $h = 5$ and $h = 20$. For comparison, I also plot the same response in the counterfactual economy with homogeneous attention. I divide the responses by γ_g to fit the same scale. We can see that the message from the analysis above extends to all groups. Generally, the higher the cyclicality of a group, the higher their level of attention. This fact implies a higher consumption response before the income realization.

Amplification of Business Cycles How does disagreement affect the transmission and propagation of business cycles? In Section 2, I argue that heterogeneous attention can amplify discount factor shocks in a simple model with rigid wages. I now discuss how disagreement affects the amplification of business cycles in the quantitative model. For example, I show that the impact of an oil shock interpreted in the model as a productivity shock with high persistence can be amplified on impact by over 17 percent.

Figure 6: Business-cycle amplification

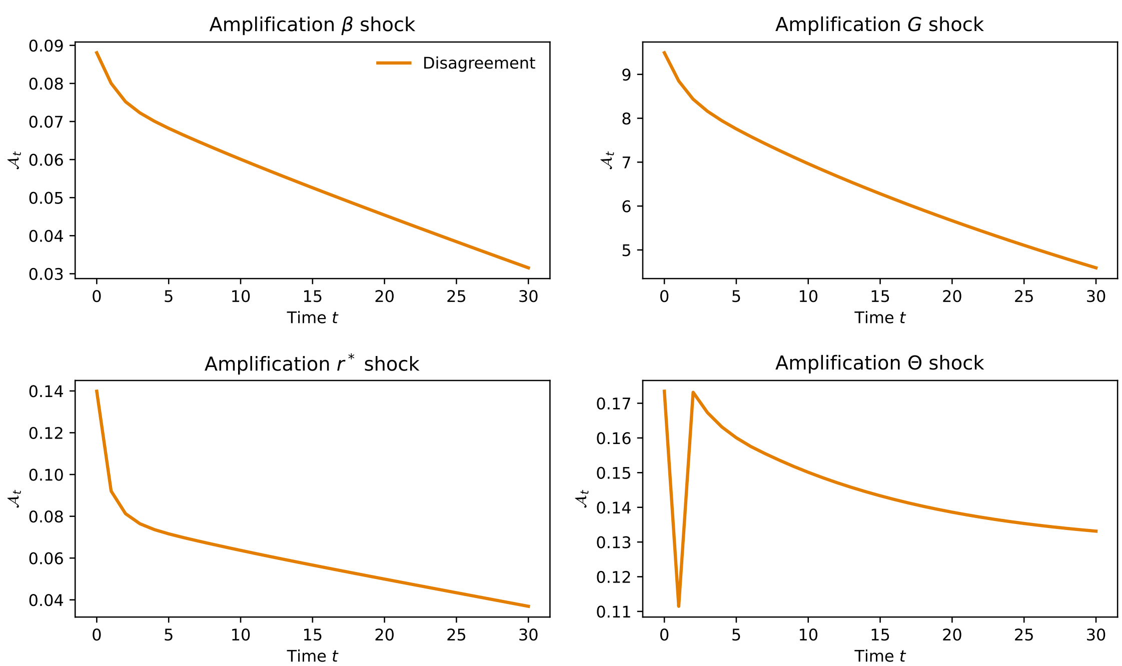

Notes: This figure displays amplification in the response of output, as defined in equation (20). I consider the response of the economy to four different shocks: a discount-factor shock in the top left panel, a government spending shock in the top right panel, an interest rate shock in the bottom left panel, and a productivity shock in the bottom right panel. See text for more details.

I compute the response of the economy in response to four different shocks: discount factor, β, government spending, G, interest rate, r^*, and productivity, Θ, shocks. For each of these, I compute the impulse response function to an innovation at time 0 dissipating with persistence ρ_x, where $\rho_\beta = 0.9$, $\rho_G = 0.91$, $\rho_r = 0.89$, and $\rho_\Theta = 0.98$. I compute the impulse response function under heterogeneous attention and in the counterfactual economy with homogeneous attention and use them to compute amplification as in equation (20).

Figure 6 displays the amount of amplification for each date t. We find that the correlated disagreement mechanism can significantly amplify the output response on impact. Discount-factor and government spending shocks can be amplified almost 10 percent, while interest rate shocks are amplified by over 14 percent, and productivity shocks are amplified by over 17 percent.

As highlighted in the simple model, the amount of amplification generated by correlated disagreement increases in the shock's persistence. To see this, Figure 7 displays impact and cumulative amplification for each of these shocks. Impact amplification is defined as $\mathcal{A}_0$ as in equation (20). Cumulative amplification summarizes the extent of amplification for the whole impulse response function, and I define it as:

$$\mathcal{CA} \equiv \frac{\sum_{t\geq 0}(1+r)^{-t}(y_t - y_t^{RA})}{\sum_{t\geq 0}(1+r)^{-t}y_t^{RA}}. \tag{58}$$

Figure 7: Business-cycle amplification: The role of persistence

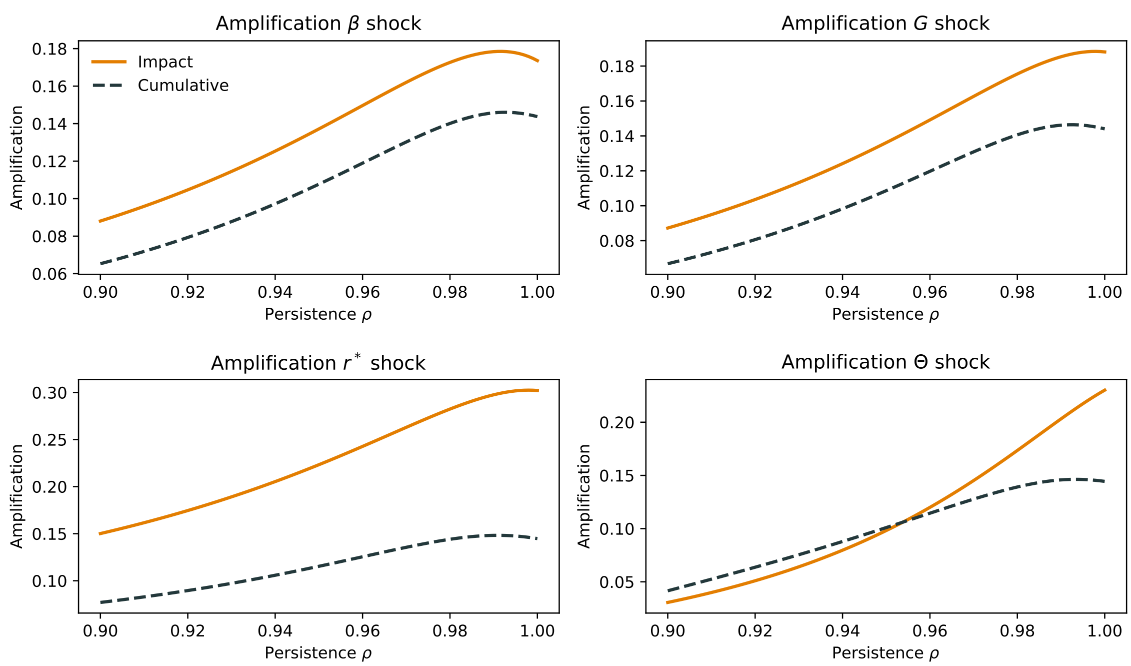

Notes: This figure displays impact and cumulative amplification in the response of output as a function of the persistence of the shock. Impact amplification is defined by $\mathcal{A}_0$, and cumulative amplification is defined as (58). I consider the response of the economy to four different shocks: a discount-factor shock in the top left panel, a government spending shock in the top right panel, an interest rate shock in the bottom left panel, and a productivity shock in the bottom right panel. See text for more details.

Figure 7 shows that impact amplification in the quantitative model can increase substantially with shock persistence. The amplification of discount-factor and spending shocks almost increases to over 18 percent, while it is as high as 25 percent for productivity shocks. Interest rate shocks are amplified by over 30 percent.

The role of monetary policy In this model, disagreement amplifies the response of output because it affects the response of aggregate demand to the general-equilibrium income channel. The strength of this channel crucially depends on how strongly monetary policy reacts to inflation. So, a natural question is how monetary policy affects the amount of amplification resulting from correlated disagreement. To answer this question, Figure 8 reconsiders the exercise of Figure 7 but with a lower Taylor coefficient of $\phi_\pi = 1.1$ instead of 1.5.

Figure 8: Business-cycle amplification: The role of monetary policy

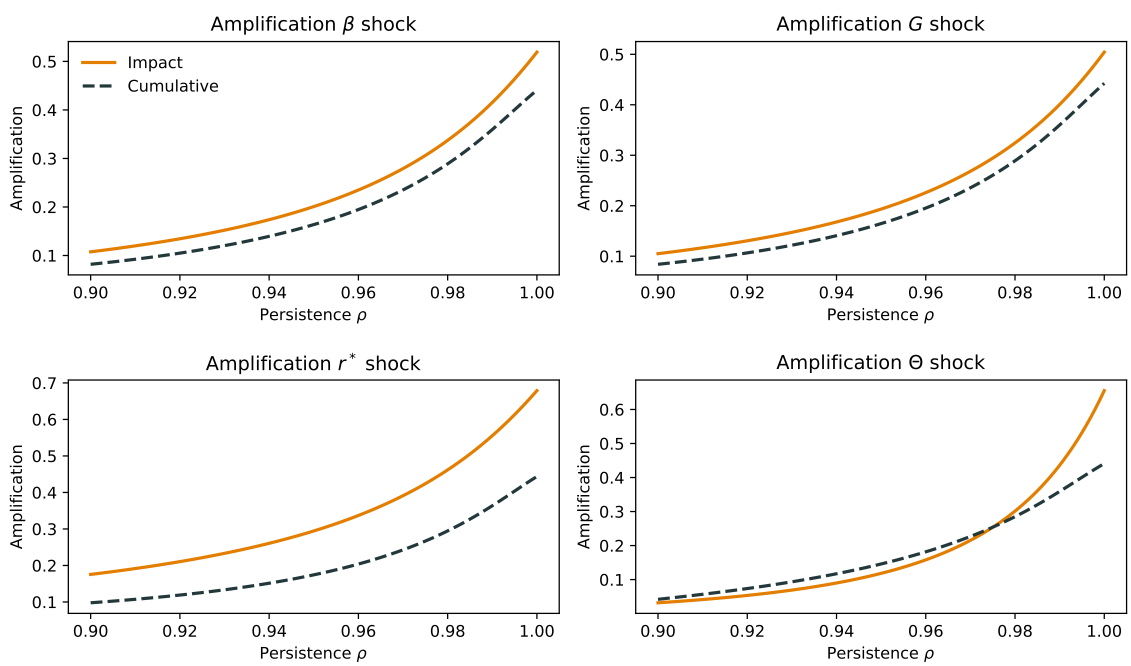

Notes: This figure displays impact and cumulative amplification in the response of output as a function of the persistence of the shock when the Taylor coefficient is reduced from $\phi_\pi = 1.5$ to $\phi_\pi = 1.1$. Impact amplification is defined by $\mathcal{A}_0$, and cumulative amplification is defined as (58). I consider the response of the economy to four different shocks: a discount-factor shock in the top left panel, a government spending shock in the top right panel, an interest rate shock in the bottom left panel, and a productivity shock in the bottom right panel. See text for more details.

Comparing Figure 8 to 7 shows that there is much stronger amplification if the monetary policy response to inflation is weaker. The response to an oil shock is now amplified by almost 30 percent on impact and similarly if evaluated in terms of cumulative effects. When monetary policy response is weaker, a more significant share of general-equilibrium forces operates via the income channel. This fact means that the correlated disagreement mechanism has a larger role when monetary policy is relatively unresponsive, leading to more considerable amplification.

5 Fiscal policy

In this section, I analyze the impact of correlated disagreement on the transmission of fiscal policy. In particular, I analyze how the composition of government spending can affect the size of the spending multiplier. I conduct this analysis from a purely positive perspective and do not consider the welfare and distributional consequences of this policy. It is well known that the desirability of stabilizing spending policy in stimulating aggregate demand depends crucially on which other constraints are imposed on monetary policy and other fiscal instruments. In this section, I do not try to assess the desirability of government spending policy.

To highlight the central intuition, I first analyze fiscal policy in the simple model of Section 2. I then evaluate these results quantitatively in Section 5.2.

5.1 Fiscal policy in the simple model

I extend the model in Section 2 to include government spending $\{G_t\}$ and proportional labor taxation $\{\tau_t\}$. For simplicity, I assume that there is no government debt, so the government runs a balanced budget. The government budget constraint is given by:

$$G_t = \tau_t Y_t, \tag{59}$$

and the modified household budget constraint is given by:

$$C_{i,t} + A_{i,t+1} = (1 - \tau_t) \cdot Y_{g,t} + (1 + r)A_{i,t}. \tag{60}$$

The market clearing condition is now given by:

$$C_t + G_t = Y_t. \tag{61}$$

For simplicity, I assume that steady-state spending and taxes equal zero.

Untargeted spending I first assume that the government buys units of the final good. In Appendix D.1, I show that to first order the individual demand can be written as:

$$c_{i,t} = (1 - \beta) \sum_{h=0}^{\infty} \beta^h E_{i,t}[y_{g,t+h} - \tau_{t+h}] + (1 - \beta)\beta^{-1} a_{i,t} + \sigma \sum_{h=0}^{\infty} \beta^{h+1} r^n_{i,t+h}. \tag{62}$$

Suppose furthermore that beliefs for income and taxes are proportional to their realized counterpart, $E_{i,t}[y_{g,t+h}] = \lambda^Y_g y_{g,t+h}$ and $E_{i,t}[\tau_{t+h}] = \lambda^\tau_g \tau_{t+h}$. Aggregate demand is given by:

$$c_t = (1 - \beta) \sum_{h=0}^{\infty} \beta^h [(1 + \mathrm{CD}) \cdot \overline{\lambda}^Y y_{t+h} - \overline{\lambda}^\tau \tau_{t+h}] + \sigma \sum_{h=0}^{\infty} \beta^{h+1} r^n_{t+h}, \tag{63}$$

where $\text{CD} \equiv \mathbb{C}\text{ov}\left(\gamma_g, \lambda_g^Y / \overline{\lambda}^Y\right)$ denotes correlated disagreement and $\overline{\lambda}^Y$ and $\overline{\lambda}^\tau$ denote the average attention to income and tax rates, respectively.

Equation (63) is the modified aggregate demand taking government into account. In equilibrium, it must be that taxes equal government spending, $\tau_t = G_t$, and that output equals private and public demand, $y_t = c_t + G_t$. So, equilibrium output is given by:

$$y_t = G_t + (1-\beta)\sum_{h=1}^{\infty}\beta^{h-1}[(1+\text{CD})\cdot\overline{\lambda}^Y y_{t+h} - \overline{\lambda}^\tau G_{t+h}] + \sigma\sum_{h=0}^{\infty}\beta^h r_{t+h}^n. \tag{64}$$

The government-spending multiplier, dy_t/dG_t, can be computed recursively using:

$$\frac{dy_t}{dG_t} = 1 + (1-\beta)\sum_{h=1}^{\infty}\beta^{h-1}\left[(1+\text{CD})\cdot\overline{\lambda}^Y\frac{dy_{t+h}}{dG_{t+h}} - \overline{\lambda}^\tau\right]\frac{dG_{t+h}}{dG_t}. \tag{65}$$

This equation relates the time t government-spending multiplier to people's beliefs about future spending multipliers. First, suppose that individuals are fully attentive. It follows that the spending multiplier is equal to one, $dy_t/dG_t = 1$, as in the FIRE analysis conducted in Woodford (2011) and Bilbiie (2011). When agents are inattentive to income and taxes, this multiplier is modified to take into account how expectations of future disposable income affect consumption choices today. The expectations of future disposable income are a function of expectations for future taxes and the effect that higher future spending has on future incomes, i.e., the future government-spending multipliers. All else equal, higher future spending multipliers increase the spending multiplier at time t, and higher correlated disagreement or greater attention to income also increases the spending multiplier if future spending multipliers are positive. The intuition for these results is that if agents expect future income to be higher, they start consuming more today, resulting in a larger spending multiplier. Instead, all else equal, a higher level of attention to taxes implies that the spending multiplier is lower because it reduces people's perceived disposable income and leads them to curtail private spending.

Proposition 4. *Suppose that $dG_t = \rho_G^t dG_0$, then the government-spending multiplier is given by:*

$$\frac{dy_t}{dG_t} = \frac{1-\varrho_G\overline{\lambda}^\tau}{1-\varrho_G\cdot(1+CD)\cdot\overline{\lambda}^Y}, \tag{66}$$

where $\varrho_G \equiv (1-\beta)\rho_G/(1-\beta\rho_G) \in (0,1)$. It follows that:

1. *The government-spending multiplier is increasing in correlated disagreement.*

2. *The government-spending multiplier is larger than under FIRE if and only if*

$$(1+CD)\cdot\overline{\lambda}^Y \geq \overline{\lambda}^\tau.$$

Proposition 4 shows that the spending multiplier is constant over time and depends on the average level of attention to taxes, correlated disagreement, and the average level of attention to income changes. Other things equal, a higher level of attention to taxes decreases the multiplier, while a higher level of attention to income or higher correlated disagreement increases the spending multiplier.

Suppose that people are fully attentive to taxes, $\overline{\lambda}^{\tau} = 1$. The spending multiplier is always lower than obtained under full information and rational expectations. Because people are fully attentive to taxes, they immediately react by decreasing consumption in expectation of higher future taxes. Because they are inattentive to income changes, they do not fully incorporate how, in general equilibrium, future higher spending translates into higher future income. In other words, the positive general-equilibrium effect of future government spending on consumption is dampened. The net effect is a lower government-spending multiplier than under FIRE. This result has been previously emphasized by Farhi et al. (2020) and Bianchi-Vimercati et al. (2021).

Instead, suppose the average attentions to taxes and income are the same. Then, the spending multiplier is larger than the one obtained under FIRE if and only if correlated disagreement is positive. While people are heterogeneously exposed to changes in income, they are equally exposed to an increase in the tax rate. It follows that the response of aggregate demand to higher taxes is captured by the economy-wide average level of attention to taxes, while the response to higher income must take into account correlated disagreement. If the average attention to taxes and income are equal, but correlated disagreement is positive, then the relevant attention to income is higher, generating a larger spending multiplier.

Targeted spending But what if the government can affect the composition of spending by directly eliciting labor from different groups? How does the composition of spending affect the spending multiplier?

Note that the analysis above implicitly assumes that government purchases in goods produced from each member of group g are given by:

$$G_{g,t} = \frac{\Gamma_g(Y_t)}{Y_t} G_t.$$

Suppose, instead, that the government can affect the composition of spending so that

$$G_{g,t} = \left(\frac{\Gamma_g(Y_t)}{Y_t} + \omega_g \right) G_t,$$

where ω_g are the targeting parameters which satisfy $\sum_g \pi_g \omega_g = 0$. This section investigates the impact of targeting via ω_g for the size of the spending multiplier. In doing so, I assume that this is a one-time unanticipated policy, so I keep the people's level of attention unchanged. See Remark 7 for a discussion.

In this case, the individual demand function (62) continues to hold, but now the individual's change in income is given by:

$$y_{g,t} = \gamma_g y_t + \omega_g G_t.$$

It follows that equilibrium output can be written as:

$$y_t = G_t + (1-\beta)\sum_{h=1}^{\infty}\beta^{h-1}\left[(1+\text{CD})\cdot\overline{\lambda}^Y y_{t+h} + \left(\text{TC}\cdot\overline{\lambda}^Y - \overline{\lambda}^\tau\right)G_{t+h}\right] + \sigma\sum_{h=0}^{\infty}\beta^h r^n_{t+h}, \tag{67}$$

where $\text{TC} \equiv \mathbb{C}\text{ov}\left(\omega_g, \lambda^Y_g/\overline{\lambda}^Y\right)$ captures the covariance between the targeting parameters ω_g and the attention parameters λ^Y_g. Using this equation, we can compute the spending multiplier recursively using the following relationship:

$$\frac{dy_t}{dG_t} = 1 + (1-\beta)\sum_{h=1}^{\infty}\beta^{h-1}\left[(1+\text{CD})\cdot\overline{\lambda}^Y\frac{dy_{t+h}}{dG_{t+h}} + \text{TC}\cdot\overline{\lambda}^Y - \overline{\lambda}^\tau\right]\frac{dG_{t+h}}{dG_t}. \tag{68}$$

Compared to equation (65), equation (68) displays a new term, $\text{TC} \equiv \mathbb{C}\text{ov}\left(\omega_g, \lambda^Y_g/\overline{\lambda}^Y\right)$, which captures the covariance between the targeting parameters ω_g and the level of attention of the group λ^Y_g. This term captures the fact that by changing the composition of spending, the government increases the incomes of certain groups more than others. To the extent that these groups have different levels of attention, they will also react heterogeneously to this income increase. People with higher income cyclicality are more attentive and will react more to the increase in income which results from higher spending. Instead, people with lower income cyclicality are less attentive and will react less to the increase in income which results from higher spending.

It follows that if the government targets the most cyclical/most attentive workers, i.e., if $\text{TC} > 0$, then spending increases the income of people with high cyclicality and attention. Because these workers are more attentive, it follows that they will increase consumption by more in response to higher spending, leading to a larger spending multiplier than without targeting. Instead, if the government targets the most cyclical workers, i.e., if $\text{TC} < 0$, then the opposite happens, and the resulting spending multiplier is lower.

Proposition 5. *Suppose that $dG_t = \rho^t_G dG_0$, then the government-spending multiplier is given by:*

$$\frac{dy_t}{dG_t} = \frac{dy^u_t}{dG_t} + \frac{\varrho_G \cdot TC \cdot \overline{\lambda}^Y}{1-\varrho_G\cdot(1+CD)\cdot\overline{\lambda}^Y}, \tag{69}$$

where $\frac{dy^u_t}{dG_t}$ denotes the untargeted government-spending multiplier, defined in equation (66)*. It follows that:*

1. *With homogeneous attention, $\lambda^Y_g = \overline{\lambda}^Y$, then targeting is irrelevant since $TC = 0$.*
2. *With heterogeneous attention, the government-spending multiplier increases if the government targets the most cyclical workers.*

Proposition 5 computes the government-spending multiplier with targeting. It shows that the government spending is equal to the untargeted spending multiplier plus an additional term which accounts for how targeting correlates with heterogeneous attention.

With homogeneous attention or with FIRE, targeted spending does not affect the spending multiplier. In this model, since all workers share the same marginal propensity to consume out of income, then targeting would be irrelevant.[21] Instead, with disagreement, targeting the most attentive workers increases the government-spending multiplier. This result follows from the fact that targeting highly attentive workers magnifies the positive effect that government spending has on aggregate demand.

The spending multiplier exceeds the one obtained under FIRE if and only if:

$$\text{TC} \geq \frac{\overline{\lambda}^{\tau} - (1 + \text{CD}) \cdot \overline{\lambda}^{Y}}{\overline{\lambda}^{Y}}. \tag{70}$$

A large multiplier is possible if the government targets the most attentive workers. Note that, even if there is full attention to taxes $\overline{\lambda}^{\tau} = 1$, it is possible to obtain a multiplier exceeding the FIRE multiplier by appropriately targeting spending.

Remark 7. *The analysis in this section assumes that attention is not reoptimized after the government policy change. So, this section can best be considered analyzing a one-time unanticipated policy. If the government were to adopt a policy that would systematically alter people's income processes, they would eventually reoptimize their levels of attention in a way that may affect the conclusions derived in this section.*

5.2 Targeted spending in the quantitative model

Figure 9: Targeted spending multipliers

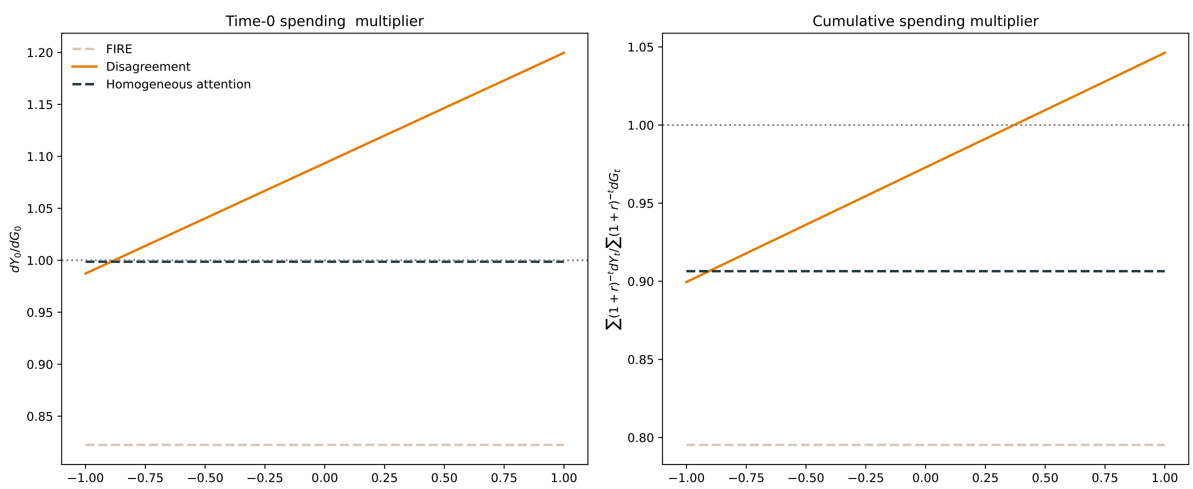

Notes: This figure shows the impact and cumulative spending multiplier as a function of the targeting parameter α for the baseline economy, FIRE, and for the economy with homogeneous attention. See text for more details.

[21] Instead, if marginal propensities to consume are heterogeneous across groups, targeting would affect the spending multiplier in a way that is similar to how heterogeneous attention affects the spending multiplier. See Baqaee and Farhi (2018) and Flynn et al. (2021).

In this section, I evaluate the quantitative implications of targeted spending for the government-spending multiplier. I extend the model in Section 4 to allow for targeted spending in the same way as in the simple model above. This assumption implies that:

$$dY_{g,t} = \gamma_g dY_t + \omega_g dG_t, \tag{71}$$

where ω_g captures the targeting parameters which satisfy $\sum_g \pi_g \omega_g = 0$. Furthermore, I assume that

$$\omega_g = \alpha \cdot (\gamma_g - 1). \tag{72}$$

This expression implies that the targeting of a particular group g is proportional to the difference between their income cyclicality, γ_g, to the average income cyclicality 1. The proportionality parameter, α, captures the strength of targeting in this policy. If $\alpha > 0$, spending targets the most cyclical groups, while if $\alpha < 0$, spending targets the least cyclical groups. The higher the targeting parameter, the higher the level of targeting to high-cyclicality workers.

Figure 9 displays the government-spending multiplier as a function of the targeting parameter α. The persistence of government spending is calibrated in the same way as in Section 4. I plot the impact and cumulative multipliers on the left and right panels, respectively, for the baseline economy, FIRE, and the counterfactual economy with homogeneous attention. The impact multiplier is defined as dY_0/dG_0 while the cumulative multiplier is defined as $\sum_t (1+r)^{-t} dY_t / \sum_t (1+r)^{-t} dG_t$.

Figure 10: Targeted spending multipliers: The role of monetary policy

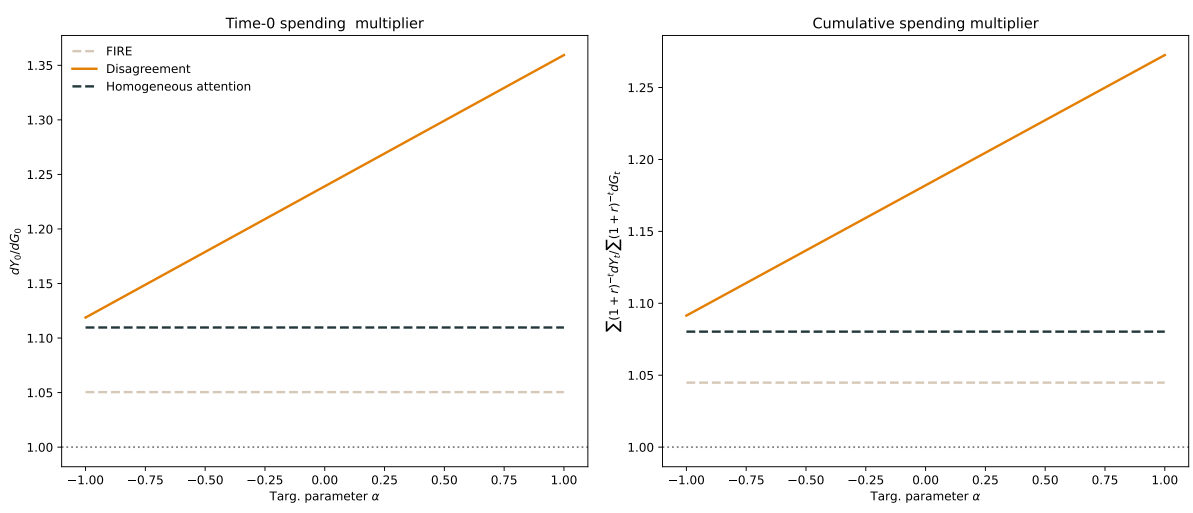

Notes: This figure shows the impact and cumulative spending multiplier as a function of the targeting parameter α assuming that the Taylor parameter is $\phi_\pi = 1.1$ instead of $\phi_\pi = 1.5$. See text for more details.

Figure 9 shows that under homogeneous attention or under FIRE, the spending multiplier is not affected by targeting. This result follows directly from the fact that group heterogeneity does not affect aggregate outcomes in this economy without belief heterogeneity. Instead, the spending multiplier is increasing in the targeting parameter in the economy with disagreement. The more spending targets highly cyclical workers, the larger the spending multiplier. In this model, we see that moving from $\alpha = -1$ to $\alpha = 1$ increases the impact spending multiplier by more than 0.20 and similarly for the cumulative multiplier.

Figure 10 redoes the same exercise but assumes a weaker response of monetary policy to inflation, $\phi_\pi = 1.1$. As in the analysis of business-cycle amplification, Figure 10 shows that the results for the spending multiplier are magnified if the monetary policy response is weaker.

These results show that the government-spending multiplier can depend substantially on which groups see their incomes rising. The power of targeted spending is higher the more accommodative monetary policy is.

6 Conclusion

This paper studies the aggregate implications of belief disagreement for the transmission of business cycles and fiscal spending policy. In particular, I study the impact of belief disagreement in shaping how aggregate demand responds to macroeconomic shocks and policies. I conduct this analysis through the lens of standard New Keynesian models with two sources of heterogeneity: heterogeneous beliefs about future income and heterogeneous income cyclicality.

The results establish the determinant role of *correlated disagreement* (CD) in shaping aggregate demand. This statistic summarizes the covariance between individual income cyclicality and heterogeneity in the response of beliefs about future income. In other words, CD summarizes the covariance between exposure and attention to shocks. I show that CD affects the general-equilibrium channel from higher future income, feeding into an expansion of aggregate demand contemporaneously, i.e., the effective marginal propensity to consume (MPC) out of future aggregate income. I show that when CD is positive, this channel is magnified relative to a counterfactual economy without heterogeneous beliefs but the same average level of attention, i.e., the effective MPC out of future income is higher. Instead, when CD is negative, the channel is dampened relative to that counterfactual economy, i.e., the effective MPC out of future income is lower.

When CD is positive, business-cycle shocks can be amplified relative to the homogeneous attention economy. I show that amplification is more significant the more persistent shocks are. This result follows from the fact that more persistence shocks attribute higher quantitative importance to expectations about the future in determining consumption and savings decisions. Instead, when CD is negative, these results are reversed.

I then endogenize beliefs via behavioral inattention as in Gabaix (2014). This model allows us to establish theoretical predictions for the sign of correlated disagreement. I show that endogeneizing beliefs implies that the sign of correlated disagreement is positive because people who are more exposed to the shock choose to pay more attention. Because they are more exposed, people with higher income cyclicality see their incomes varying more following changes in macroeconomic conditions. So, the benefit of paying attention to these shocks is higher for these individuals. It follows that attention is positively related to income exposure, implying a positive sign for correlated disagreement. I show that this implication has empirical support. Using survey data on beliefs, I compute average forecast errors in predicting income growth at the state level. I show that the magnitude of these forecast errors decreases with state average income cyclicality. Through the lens of the model, this result must be a consequence of rising levels of attention as income exposure increases.

I quantify the relevance of this propagation mechanism in a quantitative model in the Heterogeneous Agent New Keynesian tradition with countercyclical income risk, incomplete markets, and borrowing constraints. I leverage the recent computational advances to show how the model can be written in a computationally tractable way despite the large extent of income, wealth, and belief heterogeneity. I show that the correlated disagreement mechanism can lead to substantial business cycle amplification. For example, an oil shock may be propagated by as much as 16 percent. The amplification of oil shocks rises above 30 percent if monetary policy is less reactive to inflation.

Finally, I turn to a fiscal policy application. I analyze how the composition of government spending affects the fiscal spending multiplier. I show that this multiplier is higher the more the government targets workers with more cyclical incomes. If government spending targets the services of high-income cyclicality workers, it will increase the income of the most attentive people. Because they are more attentive, these workers respond more to the increase in incomes generated by government spending, leading to a more significant increase in aggregate demand than would occur if spending was targeted towards people with low income cyclicality. It follows that the spending multiplier becomes larger by targeting highly-attentive individuals. I show that, quantitatively, the differences in spending multipliers can be quite substantial as a function of the level of targeting.

References

ACEMOGLU, D. AND D. AUTOR (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of labor economics*, Elsevier, vol. 4, 1043–1171.

ANGELETOS, G.-M. AND Z. HUO (2018): "Myopia and Anchoring," Working Paper 24545, National Bureau of Economic Research.

ANGELETOS, G.-M., Z. HUO, AND K. A. SASTRY (2021): "Imperfect Macroeconomic Expectations: Evidence and Theory," *NBER Macroeconomics Annual*, 35, 1–86.

ANGELETOS, G.-M. AND J. LA'O (2010): "Noisy Business Cycles," *NBER Macroeconomics Annual*, 24, 319–378.

——— (2013): "Sentiments," *Econometrica*, 81, 739–779.

ANGELETOS, G.-M. AND C. LIAN (2018): "Forward Guidance Without Common Knowledge," *American Economic Review*, 108, 2477–2512.

ANGELETOS, G.-M. AND K. A. SASTRY (2021): "Managing Expectations: Instruments Versus Targets," *The Quarterly Journal of Economics*, 136, 2467–2532.

AUCLERT, A. (2019): "Monetary Policy and the Redistribution Channel," *American Economic Review*, 109, 2333–67.

AUCLERT, A., B. BARDÓCZY, M. ROGNLIE, AND L. STRAUB (2021): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 89, 2375–2408.

AUCLERT, A. AND M. ROGNLIE (2018): "Inequality and Aggregate Demand," Working Paper 24280, National Bureau of Economic Research.

AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2018): "The Intertemporal Keynesian Cross," Working Paper 25020, National Bureau of Economic Research.

——— (2020): "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," Working Paper 26647, National Bureau of Economic Research.

BAQAEE, D. R. (2015): "Targeted Fiscal Policy," Tech. rep., Working Paper.

BAQAEE, D. R. AND E. FARHI (2018): "Macroeconomics with Heterogeneous Agents and Input-Output Networks," Working Paper 24684, National Bureau of Economic Research.

BARDOCZY, B., R. BIANCHI-VIMERCATI, AND J. GUERREIRO (2022): "Unemployment Insurance in Macroeconomic Stabilization with Bounded Rationality," mimeo.

BIANCHI, F., C. ILUT, AND H. SAIJO (2021): "Diagnostic Business Cycles," .

BIANCHI-VIMERCATI, R., M. S. EICHENBAUM, AND J. GUERREIRO (2021): "Fiscal Policy at the Zero Lower Bound without Rational Expectations," Working Paper 29134, National Bureau of Economic Research.

BILBIIE, F. O. (2011): "Nonseparable Preferences, Frisch Labor Supply, and the Consumption Multiplier of Government Spending: One Solution to a Fiscal Policy Puzzle," *Journal of Money, Credit and Banking*, 43, 221–251.

BLANCHARD, O. J. AND J. GALI (2007): "The Macroeconomic Effects of Oil Shocks: Why are the 2000s So Different from the 1970s?" Working Paper 13368, National Bureau of Economic Research.

BORDALO, P., N. GENNAIOLI, Y. MA, AND A. SHLEIFER (2020): "Overreaction in Macroeconomic Expectations," *American Economic Review*, 110, 2748–82.

BOUAKEZ, H., O. RACHEDI, AND E. SANTORO (2020): "The Sectoral Origins of the Spending Multiplier," Tech. rep., Working paper.

BROER, T., A. KOHLHAS, K. MITMAN, AND K. SCHLAFMANN (2021): "Information and Wealth Heterogeneity in the Macroeconomy," .

CABALLERO, R. J. AND A. SIMSEK (2020): "A Risk-Centric Model of Demand Recessions and Speculation," *The Quarterly Journal of Economics*, 135, 1493–1566.

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, 471–75.

COGLEY, T. AND T. J. SARGENT (2008): "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making," *International Economic Review*, 49, 185–221.

COIBION, O. AND Y. GORODNICHENKO (2012): "What Can Survey Forecasts Tell us About Information Rigidities?" *Journal of Political Economy*, 120, 116–159.

——— (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105, 2644–78.

COIBION, O., Y. GORODNICHENKO, L. KUENG, AND J. SILVIA (2017): "Innocent Bystanders? Monetary Policy and Inequality," *Journal of Monetary Economics*, 88, 70–89.

COX, L., G. MÜLLER, E. PASTÉN, R. SCHOENLE, AND M. WEBER (2020): "Big G," Working Paper 27034, National Bureau of Economic Research.

D'ACUNTO, F., D. HOANG, M. PALOVIITA, AND M. WEBER (2019): "IQ, Expectations, and Choice," Working Paper 25496, National Bureau of Economic Research.

ENGELBERG, J., C. F. MANSKI, AND J. WILLIAMS (2009): "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters," *Journal of Business & Economic Statistics*, 27, 30–41.

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, 281–313.

FARHI, E., M. PETRI, AND I. WERNING (2020): "The Fiscal Multiplier Puzzle: Liquidity Traps, Bounded Rationality, and Incomplete Markets," mimeo.

FARHI, E. AND I. WERNING (2019): "Monetary Policy, Bounded Rationality, and Incomplete Markets," *American Economic Review*, 109, 3887–3928.

FLOOD, S., M. KING, R. RODGERS, S. RUGGLES, J. R. WARREN, AND M. WESTBERRY (2021): "Integrated Public Use Microdata Series, Current Population Survey: Version 9.0," .

FLYNN, J. P., C. PATTERSON, AND J. STURM (2021): "Fiscal Policy in a Networked Economy," Working Paper 29619, National Bureau of Economic Research.

GABAIX, X. (2014): "A Sparsity-Based Model of Bounded Rationality," *The Quarterly Journal of Economics*, 129, 1661–1710.

——— (2016): "Behavioral Macroeconomics Via Sparse Dynamic Programming," Working Paper 21848, National Bureau of Economic Research.

——— (2019): "Behavioral Inattention," in *Handbook of Behavioral Economics: Applications and Foundations 1*, Elsevier, vol. 2, 261–343.

——— (2020): "A Behavioral New Keynesian Model," *American Economic Review*, 110, 2271–2327.

GARCÍA-SCHMIDT, M. AND M. WOODFORD (2019): "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis," *American Economic Review*, 109, 86–120.

GORNEMANN, N., K. KUESTER, AND M. NAKAJIMA (2016): "Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy," Manuscript.

GUERRIERI, V. AND G. LORENZONI (2017): "Credit Crises, Precautionary Savings, and the Liquidity Trap," *The Quarterly Journal of Economics*, 132, 1427–1467.

GUVENEN, F., S. SCHULHOFER-WOHL, J. SONG, AND M. YOGO (2017): "Worker Betas: Five Facts about Systematic Earnings Risk," *American Economic Review*, 107, 398–403.

HARRISON, J. M. AND D. M. KREPS (1978): "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *The Quarterly Journal of Economics*, 92, 323–336.

HASSAN, T. A. AND T. M. MERTENS (2017): "The Social Cost of Near-Rational Investment," *American Economic Review*, 107, 1059–1103.

HAZELL, J., J. HERRENO, E. NAKAMURA, AND J. STEINSSON (2022): "The Slope of the Phillips Curve: Evidence from US States," *The Quarterly Journal of Economics*, 137, 1299–1344.

HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): "Optimal Tax Progressivity: An Analytical Framework," *The Quarterly Journal of Economics*, 132, 1693–1754.

ILUT, C. L. AND M. SCHNEIDER (2014): "Ambiguous Business Cycles," *American Economic Review*, 104, 2368–99.

IOVINO, L. AND D. SERGEYEV (2018): "Central Bank Balance Sheet Policies Without Rational Expectations," .

JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2010): "Investment Shocks and Business Cycles," *Journal of Monetary Economics*, 57, 132–145.

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

KREPS, D. M. (1998): "Anticipated Utility and Dynamic Choice," *Econometric Society Monographs*, 29, 242–274.

LORENZONI, G. (2009): "A Theory of Demand Shocks," *American Economic Review*, 99, 2050–84.

LUCAS, R. E. J. (1972): "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4, 103–124.

MANKIW, N. G. AND R. REIS (2002): "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 117, 1295–1328.

MANKIW, N. G., R. REIS, AND J. WOLFERS (2003): "Disagreement About Inflation Expectations," *NBER Macroeconomics Annual*, 18, 209–248.

MCKAY, A., E. NAKAMURA, AND J. STEINSSON (2016): "The Power of Forward Guidance Revisited," *American Economic Review*, 106, 3133–58.

MCKAY, A. AND R. REIS (2016): "The Role of Automatic Stabilizers in the US Business Cycle," *Econometrica*, 84, 141–194.

NIMARK, K. P. (2014): "Man-bites-dog Business Cycles," *American Economic Review*, 104, 2320–67.

PAPPA, E., M. O. RAVN, AND V. STERK (2023): "Expectations and Incomplete Markets," in *Handbook of Economic Expectations*, Elsevier, 569–611.

PATTERSON, C. (2019): "The Matching Multiplier and the Amplification of Recessions," Working paper.

RAMEY, V. A. AND M. D. SHAPIRO (1998): "Costly Capital Reallocation and the Effects of Government Spending," in *Carnegie-Rochester conference series on public policy*, Elsevier, vol. 48, 145–194.

SCHEINKMAN, J. A. AND W. XIONG (2003): "Overconfidence and Speculative Bubbles," *Journal of political Economy*, 111, 1183–1220.

SCHMITT-GROHÉ, S. AND M. URIBE (2005): "Optimal Fiscal and Monetary Policy in a Medium-Scale Macroeconomic Model," *NBER Macroeconomics Annual*, 20, 383–425.

SIMS, C. A. (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665–690.

WERNING, I. (2015): "Incomplete Markets and Aggregate Demand," Working Paper 21448, National Bureau of Economic Research.

WOODFORD, M. (2001): "Imperfect Common Knowledge and the Effects of Monetary Policy," Working Paper 8673, National Bureau of Economic Research.

——— (2011): "Simple Analytics of the Government Expenditure Multiplier," *American Economic Journal: Macroeconomics*, 3, 1–35.

——— (2018): "Monetary Policy Analysis When Planning Horizons are Finite," *NBER Macroeconomics Annual*, 33, 1–50.

WOODFORD, M. AND Y. XIE (2019): "Policy Options at the Zero Lower Bound When Foresight is Limited," *AEA Papers and Proceedings*, 109, 433–37.

——— (2022): "Fiscal and Monetary Stabilization Policy at the Zero Lower Bound: Consequences of Limited Foresight," *Journal of Monetary Economics*, 125, 18–35.

ZARNOWITZ, V. AND L. A. LAMBROS (1987): "Consensus and Uncertainty in Economic Prediction," *Journal of Political economy*, 95, 591–621.

A Appendix to section 2

A.1 Individual demand

The household enters time t with assets $A_{i,t}$ and chooses consumption and savings to solve

$$\max E_{i,t} \sum_{h=0}^{\infty} \prod_{s=0}^{h-1} (\beta_{i,t+s}) \left[u(C_{i,t}) - v\left(N_{i,t}\right) \right], \quad \text{subject to}$$
$$C_{i,t+h} + A_{i,t+h+1} = Y_{g,t+h} + (1+r)A_{i,t+h}.$$

The Euler equation is given by:

$$u'(C_{i,t}) = \prod_{s=0}^{h-1} \beta_{i,t+s}(1+r)E_{i,t}[u'(C_{i,t+h})]. \tag{73}$$

I log-linearize the solution to the household problem around a steady-state equilibrium in which $(1+r)\beta = 1$ and all households are symmetrical: $A_i = 0$, $Y_g = Y = 1$, and $C_i = C = Y = 1$. Log-linearizing the Euler equation (73) obtains

$$E_{i,t}[c_{i,t+h}] = c_{i,t} - \sigma \sum_{s=0}^{h-1} r^n_{i,t+s}, \tag{74}$$

where $c_{i,t} \equiv d\log(C_i, t)$ and $r^n_{i,t} \equiv -d\log(\beta_{i,t})$.

Linearizing the budget constraint, we obtain

$$c_{i,t+h} + a_{i,t+h+1} = y_{g,t+h} + \beta^{-1} a_{i,t+h}, \tag{75}$$

where $a_{i,t} = dA_{i,t}$ and $(1+r) = \beta^{-1}$. Multiplying this equation by β^h and iterating forward for each we obtain

$$\sum_{h=0}^{\infty} \beta^h c_{i,t+h} = \sum_{h=0}^{\infty} \beta^h y_{g,t+h} + \beta^{-1} a_{i,t}. \tag{76}$$

Taking expectations and replacing equation (74) obtains

$$\sum_{h=0}^{\infty} \beta^h \left[c_{i,t} - \sigma \sum_{s=0}^{h-1} r^n_{i,t+s} \right] = \sum_{h=0}^{\infty} \beta^h E_{i,t}[y_{g,t+h}] + \beta^{-1} a_{i,t}$$
$$\Leftrightarrow \frac{1}{1-\beta} c_{i,t} = \sum_{h=0}^{\infty} \beta^h E_{i,t}[y_{g,t+h}] + \beta^{-1} a_{i,t} + \sum_{s=0}^{\infty} \sum_{h=s+1}^{\infty} \beta^h r^n_{i,t+s}$$
$$\Leftrightarrow c_{i,t} = (1-\beta) \sum_{h=0}^{\infty} \beta^h E_{i,t}[y_{g,t+h}] + (1-\beta)\beta^{-1} a_{i,t} + \sum_{h=0}^{\infty} \beta^{h+1} r^n_{i,t+h}.$$

A.2 Proof of proposition 1

Computing group-average demand we obtain

$$\bar{c}_{g,t} = (1-\beta)\sum_{h=0}^{\infty}\beta^h \gamma_g \overline{E}_{g,t}[y_{t+h}] + (1-\beta)\beta^{-1}\bar{a}_{g,t} + \sum_{h=0}^{\infty}\beta^{h+1}\bar{r}^n_{g,t+h}.$$

Aggregating across groups, we obtain

$$\begin{aligned}
c_t &= \sum_g \pi_g \bar{c}_{g,t} \\
\Leftrightarrow c_t &= \sum_g \pi_g \left[(1-\beta)\sum_{h=0}^{\infty}\beta^h \gamma_g \overline{E}_{g,t}[y_{t+h}] + (1-\beta)\beta^{-1}\bar{a}_{g,t} + \sum_{h=0}^{\infty}\beta^{h+1}\bar{r}^n_{g,t+h}\right] \\
\Leftrightarrow c_t &= (1-\beta)\sum_{h=0}^{\infty}\beta^h \sum_g \pi_g \gamma_g \overline{E}_{g,t}[y_{t+h}] + \sum_{h=0}^{\infty}\beta^{h+1} r^n_{t+h},
\end{aligned}$$

where $\sum_g \pi_g \bar{a}_{g,t} = 0$ by asset market clearing and $\sum_g \pi_g \bar{r}^n_{g,t+h} \equiv r^n_{t+h}$. Finally, note that

$$\begin{aligned}
\sum_g \pi_g \gamma_g \overline{E}_{g,t}[y_{t+h}] &= \underbrace{\left[\sum_g \pi_g \gamma_g\right]}_{=1} \cdot \underbrace{\left[\sum_g \pi_g \overline{E}_{g,t}[y_{t+h}]\right]}_{=\overline{E}_t[y_{t+h}]} + \mathbb{C}\text{ov}\left(\gamma_g, \overline{E}_{g,t}[y_{t+h}]\right) \\
&= \overline{E}_t[y_{t+h}] + \mathbb{C}\text{ov}\left(\gamma_g, \frac{\overline{E}_{g,t}[y_{t+h}]}{\overline{E}_t[y_{t+h}]}\right)\overline{E}_t[y_{t+h}] = (1+\text{CD}_{t+h})\cdot \overline{E}_t[y_{t+h}].
\end{aligned}$$

Aggregate demand is thus given by

$$c_t = (1-\beta)\sum_{h=0}^{\infty}\beta^h (1+\text{CD}_{t+h})\cdot \overline{E}_t[y_{t+h}] + \sum_{h=0}^{\infty}\beta^{h+1} r^n_{t+h}. \tag{77}$$

Finally, market clearing for goods market requires $c_t = y_t$, and so equilibrium output solves

$$y_t = (1-\beta)y_t + (1-\beta)\sum_{h=1}^{\infty}\beta^h (1+\text{CD}_{t+h})\cdot \overline{E}_t[y_{t+h}] + \sum_{h=0}^{\infty}\beta^{h+1} r^n_{t+h} \tag{78}$$

$$\Leftrightarrow y_t = (1-\beta)\sum_{h=1}^{\infty}\beta^{h-1} (1+\text{CD}_{t+h})\cdot \overline{E}_t[y_{t+h}] + \sum_{h=0}^{\infty}\beta^{h} r^n_{t+h}. \tag{79}$$

Assuming $\overline{E}_{g,t}[y_{t+h}] = \lambda_g y_{t+h}$ and defining $\bar{\lambda} \equiv \sum_g \pi_g \lambda_g$, we can write $\overline{E}_t[y_{t+h}] = \bar{\lambda} y_{t+h}$ and

$$\text{CD}_{t,h} = \mathbb{C}\text{ov}\left(\gamma_g, \frac{\overline{E}_{g,t}[y_{t+h}]}{\overline{E}_t[y_{t+h}]}\right) = \mathbb{C}\text{ov}\left(\gamma_g, \frac{\lambda_g}{\bar{\lambda}}\right) \equiv \text{CD}.$$

Replacing these expressions in the equation above, we finally obtain the following:

$$y_t = (1-\beta)\sum_{h=1}^{\infty}\beta^{h-1}\cdot(1+\text{CD})\cdot\overline{\lambda}y_{t+h} + \sigma\sum_{h=0}^{\infty}\beta^h r^n_{t+h}. \tag{80}$$

A.3 Proof of proposition 2

Note that equation (18) can be equivalently written as

$$y_t = (1-\beta)(1+\text{CD})\cdot\overline{\lambda}y_{t+1} + \sigma r^n_t + \beta\underbrace{\left[(1-\beta)\sum_{h=2}^{\infty}\beta^{h-2}\cdot(1+\text{CD})\cdot\overline{\lambda}y_{t+h} + \sigma\sum_{h=1}^{\infty}\beta^{h-1}r^n_{t+h}\right]}_{=y_{t+1}}$$

$$y_t = \left[\beta + (1-\beta)(1+\text{CD})\cdot\overline{\lambda}\right]y_{t+1} + \sigma r^n_t.$$

If $r^n_t = \rho^t r_0$, then the unique solution to this difference with $\lim_{t\to\infty} y_t = 0$ satisfies $y_t = \rho^t y_0$ and

$$\rho^t y_0 = \left[\beta + (1-\beta)(1+\text{CD})\cdot\overline{\lambda}\right]\rho^{t+1}y_0 + \sigma\rho^t r^n_0 \Leftrightarrow y_0 = \frac{\sigma r^n_0}{1-\rho\left[\beta + (1-\beta)(1+\text{CD})\cdot\overline{\lambda}\right]}. \tag{81}$$

The solution with disagreement is

$$y_t = \frac{\sigma\rho^t r^n_0}{1-\rho\left[\beta + (1-\beta)(1+\text{CD})\cdot\overline{\lambda}\right]} \tag{82}$$

and with constant attention

$$y_t = \frac{\sigma\rho^t r^n_0}{1-\rho\left[\beta + (1-\beta)\cdot\overline{\lambda}\right]}. \tag{83}$$

Computing amplification we obtain

$$\mathcal{A}_t = \frac{(1-\beta)\rho\cdot\text{CD}\cdot\overline{\lambda}}{1-\rho\left[\beta + (1-\beta)(1+\text{CD})\cdot\overline{\lambda}\right]}.$$

Note that

$$\frac{d\mathcal{A}_t}{d\text{CD}} = \frac{\left(1-\rho\left(\overline{\lambda}+\beta-\beta\overline{\lambda}\right)\right)(1-\beta)\rho\cdot\overline{\lambda}}{\left(1-\rho\left[\beta + (1-\beta)(1+\text{CD})\cdot\overline{\lambda}\right]\right)^2} > 0 \tag{84}$$

and so amplification is increasing in correlated disagreement.

Furthermore, the comparative static concerning persistence is given by

$$\frac{d\mathcal{A}_t}{d\rho} = \frac{(1-\beta)\cdot\text{CD}\cdot\overline{\lambda}}{\left(1-\rho\left[\beta + (1-\beta)(1+\text{CD})\cdot\overline{\lambda}\right]\right)^2} \tag{85}$$

This derivative has the same sign as CD.

Finally, the comparative static concerning β is given by

$$\frac{d\mathcal{A}_t}{d\beta} = -\frac{\beta\rho\,(1-\rho)\cdot \text{CD}\cdot\overline{\lambda}}{\left(1-\rho\left[\beta+(1-\beta)(1+\text{CD})\cdot\overline{\lambda}\right]\right)^2}, \tag{86}$$

which has the opposite sign as CD.

A.4 Utility cost of inattention

In this appendix, I derive the expression for the utility costs of attention, equation (23). The exposition here translates the discussion in Gabaix (2016) to the setting in this paper. As the main text discusses, individuals face no uncertainty around their forecasts. I define the value function of an individual with full attention as follows:

$$V_{i,t}(A;\{Y_{g,t+h}\}_{h\geq 0}) = \max_C\{u(C)+\beta_{i,t}V_{i,t+1}(Y_{g,t}+(1+r)A-C)\} \tag{87}$$

and the objective function in this problem is

$$v_{i,t}(C) \equiv u(C)+\beta_{i,t}V_{i,t+1}(Y_{g,t}+(1+r)A-C). \tag{88}$$

Instead, the problem of an inattentive individual is given by

$$u(C)+\beta_{i,t}E_{i,t}[V_{i,t+1}(Y_{g,t}+(1+r)A-C)]. \tag{89}$$

Note that the individual acts assuming they will not update their beliefs. First, we want to make a second-order approximation of the objective function around the point in which $C=1$, $A=0$, $Y_{g,t}=1$, i.e., around the unshocked steady-state equilibrium. This approximation yields

$$v_{i,t}(c) = v(0)+\frac{1}{2}\frac{\partial^2 v}{\partial C^2}c^2+\sum_{h=0}^{\infty}\frac{\partial^2 v}{\partial C\partial Y_h}\cdot c\cdot y_{g,t+h}+\frac{\partial^2 v}{\partial C\partial A}\cdot c\cdot a+\sum_{h=0}^{\infty}\frac{\partial^2 v}{\partial C\partial \beta_{i,h}}\beta\cdot c\cdot r^n_{i,t+h}$$
$$+\text{ terms independent of } C, \tag{90}$$

where $v(0)=\frac{u(1)}{1-\beta}$, and note that because the Euler equation $\frac{\partial v_{i,t}(C)}{\partial C}=0$ holds, then we can write the following second-order derivatives. First, the curvature in C is given by

$$\frac{\partial^2 v}{\partial C^2} = u''(1)+\beta\frac{\partial V}{\partial A^2} = \beta^{-1}u''(1).$$

because

$$\frac{\partial V}{\partial A^2} = \beta^{-1}\frac{\partial u'(C^*(A))}{\partial A} = u''(1)(1-\beta)\beta^{-2}.$$

By similar logic, we can write

$$\frac{\partial^2 v}{\partial C \partial A} = -\frac{\partial^2 v}{\partial C^2} \underbrace{(1-\beta)\beta^{-1}}_{=\partial C/\partial A},$$

$$\frac{\partial^2 v}{\partial C \partial Y_h} = -\frac{\partial^2 v}{\partial C^2} \underbrace{(1-\beta)\beta^{h}}_{=\partial C/\partial Y_h},$$

and

$$\frac{\partial^2 v}{\partial C \partial \beta_{i,h}} = -\frac{\partial^2 v}{\partial C^2} \underbrace{\sigma\beta^{h}}_{=\partial C/\partial \beta_h},$$

Note that the solution to the problem of maximizing expected utility in (90) yields the same solution we have derived before:

$$c^*_{g,t}(a, \lambda_i) = (1-\beta)\sum_{h=0}^{\infty} \beta^h E_{i,t}[y_{g,t+h}] + (1-\beta)\beta^{-1} a + \sigma \sum_{h=0}^{\infty} \beta^{h+1} r^n_{i,t+h}.$$

Note that, for any C, we can write the utility value as

$$v_{i,t}(c) = v(0) + \frac{1}{2}\frac{\partial^2 v}{\partial C^2} c^2 - \frac{\partial^2 v}{\partial C^2} \cdot c \cdot c^*_{g,t}(a,1) + \text{terms independent of } C,$$

where $c^*_{g,t}(a,1)$ denotes the rational expectations demand.

The realized utility cost of inattention is given by

$$v_{i,t}(c^*_{g,t}(a,1)) - v_{i,t}(c^*_{g,t}(a,\lambda_i)) = -\frac{1}{2}\frac{\partial^2 v}{\partial C^2}\left(c^*_{g,t}(a,1) - c^*_{g,t}(a,\lambda_i)\right)^2, \quad (91)$$

and note that

$$\left(c^*_{g,t}(a,1) - c^*_{g,t}(a,\lambda_i)\right)^2 = \left(\sum_{h=1}^{\infty} \frac{\partial C}{\partial Y_h}(1-\lambda_i) y_{g,t+h}\right)^2 = \sum_{h=1}^{\infty}\sum_{\tilde{h}=1}^{\infty} \frac{\partial C}{\partial Y_h}\frac{\partial C}{\partial Y_{\tilde{h}}}(1-\lambda_i)^2 y_{g,t+h} y_{g,t+\tilde{h}}. \quad (92)$$

It follows that the ex-ante utility cost of inattention is given by

$$\mathcal{C}_g(\lambda_i) = -\frac{1}{2}\frac{\partial^2 v}{\partial C^2}\sum_{h=1}^{\infty}\sum_{\tilde{h}=1}^{\infty} \frac{\partial C}{\partial Y_h}\frac{\partial C}{\partial Y_{\tilde{h}}}(1-\lambda_i)^2 \cdot \gamma_g^2 \sigma_{h,\tilde{h}}, \quad (93)$$

where $\sigma_{h,\tilde{h}}$ denotes the ex-ante perceived covariance between y_{t+h} and $y_{t+\tilde{h}}$ which is assumed to depend only on the horizons and not date t.

A.5 Horizon-varying attention

We may think about situations where individuals have different attentions based on the forecast horizon. So, suppose that attention varies with the horizon, then we replace the structural relation

(16) with

$$\overline{E}_{g,t}[y_{t+h}] = \lambda_{g,h} y_{t+h}. \tag{94}$$

In this case,

$$\overline{E}_t y_{t+h} = \overline{\lambda}_h y_{t+h} \quad \text{and} \quad \text{CD}_h \equiv \mathbb{C}\text{ov}\left(\gamma_g, \lambda_{g,h}/\overline{\lambda}_h\right),$$

where $\overline{\lambda}_h \equiv \sum_g \pi_g \lambda_{g,h}$.

Equilibrium output now satisfies

$$y_t = (1-\beta) \sum_{h=1}^{\infty} \beta^{h-1} \cdot (1 + \text{CD}_h) \cdot \overline{\lambda}_h y_{t+h} + \sigma \sum_{h=0}^{\infty} \beta^h r_{t+h}^n. \tag{95}$$

It follows that the results in Proposition 1 still hold under the caveat that we must require the properties to be met for all horizons.

Defining amplification in the same way, we find that

$$\mathcal{A}_t = \sum_{h=1}^{\infty} \{\text{CD}_h + (1 + \text{CD}_h)\, \mathcal{A}_{t+h}\} \cdot \frac{(1-\beta) \cdot \beta^{h-1} \cdot \overline{\lambda}_h y_{t+h}}{y_t^{\text{RA}}} \tag{96}$$

It follows that amplification is positive if $\text{CD}_h > 0$ for all h and negative if $\text{CD}_h < 0$ for all h. However, exactly how amplification depends on persistence, and the discount factor becomes less clear. This expression, however, suggests two facts: (1) the longer the effects on output, the more significant is the impact of correlated disagreement, and (2) the more important the general equilibrium channel, the more significant the impact of correlated disagreement. The first can be achieved via a persistent shock, and the second with a higher marginal propensity to consume.

To endogenize attention, I allow individuals to optimize their level of attention for each horizon $\lambda_{i,h}$. It turns out that the utility costs of inattention can be written analogously to what we have found before

$$\mathcal{C}_g(\boldsymbol{\lambda}_i) = -\frac{1}{2} \frac{\partial^2 v}{\partial C^2} \sum_{h=1}^{\infty} \sum_{\tilde{h}=1}^{\infty} \frac{\partial C}{\partial Y_h} \frac{\partial C}{\partial Y_{\tilde{h}}} (1 - \lambda_{i,h})(1 - \lambda_{i,\tilde{h}}) \cdot \gamma_g^2 \sigma_{h,\tilde{h}}. \tag{97}$$

Albeit not essential, I assume, as in Gabaix (2014), that people perceive no correlation across the variables. Optimal attention thus solves

$$\min_{\boldsymbol{\lambda}_I} -\frac{1}{2} \frac{\partial^2 v}{\partial C^2} \sum_{h=1}^{\infty} \left(\frac{\partial C}{\partial Y_h}\right)^2 (1 - \lambda_{i,h})^2 \cdot \gamma_g^2 \sigma^2 + \kappa \sum_h \lambda_{i,h}. \tag{98}$$

Proposition 6. *Optimal attention to horizon h is given by*

$$\lambda_{i,h} = \lambda_{g,h} \equiv \max\left\{0, 1 - \frac{\kappa}{\Lambda_h \gamma_g^2}\right\},$$

where $\Lambda_h \equiv -\frac{1}{2} \frac{\partial^2 v}{\partial C^2} \left(\frac{\partial C}{\partial Y_h}\right)^2 \sigma^2$. *It follows that:*

1. *Attention $\lambda_{g,h}$ is increasing in γ_g.*

2. *Attention is decreasing in horizon h. There exists $H > 0$ such that $\lambda_{g,h} = 0$ for all $h > H$.*

This extended model still holds the central result that attention is increasing in the income cyclicality. Furthermore, we also find that attention is decreasing in the forecast horizon. The reason for this implication is as follows. The present the value of a change in income at date h is given by β^h, which means that the longer the horizon, the lower the impact that those changes in income have on contemporaneous consumption. So, people choose to devote more attention to incomes that are relatively close in time than to incomes that are further away in the future.

Note also that for sufficiently far-off events, individuals become fully inattentive. Intuitively, for enough distant events, their impact on current consumption would be so small that it does not pay off to exert the cognitive effort of trying to forecast them. Interestingly, this model generates an endogenous "finite planning horizon", a behavioral feature analyzed in Woodford (2018) and Woodford and Xie (2019, 2022).

B Appendix to section 3

In this appendix, I describe the data used in section 3. The data comes from the Survey of Consumer Expectations (SCE) and the Current Population Survey (CPS).

Survey of Consumer Expectations The SCE is a monthly internet rotating panel survey of one thousand and three hundred (1,300) households that started in June 2013. New respondents are drawn to match demographic factors from the American Community Survey, ensuring population representativeness, and stay on the panel for up to twelve months. To increase data availability, I aggregate individual responses to the quarterly level by averaging within that time frame.

In this paper, I use the responses to the following question: "*Suppose again that, 12 months from now, you are working in the exact same/main job at the same place you currently work, and working the exact same number of hours. In your view, what would you say is the percent chance that 12 months from now your earnings on this job, before taxes and deductions, will have...*" Respondents are asked to assign probabilities to ten different bins: higher than 12%, between 8% and 12%, between 4% and 8%, between 2% and 4%, between 0% and 2%, between -2% and 0%, between -4% and -2%, between -8% and -4%, between -12% and -8%, and lower than -12%.

The SCE estimates a density distribution for household forecasts using the approach in Engelberg et al. (2009). I assume that this estimated mean captures $E_{i,t}[\Delta y_{i,t+h}]$ where the horizon $h = 4$ quarters or 1 year.

Current Population Survey The CPS is a monthly survey of around sixty thousand U.S. households (60,000) conducted by the BLS, starting from 1940. This survey contains detailed microdata on employment and income characteristics of members within a household.

To estimate the income cyclicality parameters, I use yearly data from the ASEC March Supplement of the CPS from 2000 to 2019. I remove the Covid-19 recession from this estimation due to its unusual features in terms of labor market incidence. Using the monthly responses for 2012 to 2021, I also compute state-level average income growth $\Delta\overline{y}_{S,t}$ for 2013-21.

B.1 Forecast error

In the CPS data, I consider individuals who are in the labor force aged 20 to 64, who are active in the labor force, and who are not in the military. As in Acemoglu and Autor (2011), I multiply top-coded weekly earnings and hourly wages by 1.5. When not available, I compute weekly earnings using the information on hourly wages and weekly hours of work. I deflate these weekly earnings by the CPI to measure real earnings. I then use weekly earnings to compute average income growth at the state level using the sample earnings weights. I also aggregate the SCE responses to obtain a state-level average forecast using the sample weights.

Using this data, the state-level average forecast error is defined as

$$\overline{\text{FE}}_{S,t} \equiv \Delta y_{i,t+h} - \overline{E}_{S,t}[\Delta y_{i,t+h}]. \tag{99}$$

B.2 Income cyclicality

To estimate γ_g, I use March Supplement CPS data. I restrict the analysis to households aged 20 to 64 active in the labor force and not in the military. I focus on the set of 14 census industries by matching the 1990 industry information to their corresponding industry. The precise matching can be found in table 3.

Table 3: Census industry and 1990 Industrial Class. System

	Industry	Start	End		Industry	Start	End
1	Agriculture, Forestry, Fishery	10	32	8	Non-durable Man.	100	229
2	Public Administration	900	932	9	Durable Man.	230	392
3	Bus. and Repair Services	721	760	10	Retail Trade	580	691
4	Prof. and Related Serv.	812	893	11	Wholesale Trade	500	571
5	Mining	40	50	12	Personal Services	761	791
6	Transp., Commun., Public Util.	400	472	13	Finance, Insur., Real Est.	700	712
7	Construction	60	60	14	Ent. and Recr. Serv.	800	810

As in Acemoglu and Autor (2011), I multiply top-coded weekly earnings and hourly wages by 1.5. When not available, I compute weekly earnings using the information on hourly wages and weekly hours of work. I deflate these weakly earnings by the CPI to measure real earnings. I use the unique individual identifier to match individuals across consecutive years.[22] I compute income growth for each individual and calculate nationwide aggregate income growth using individual earnings weights.

[22] I use only matches for which race and sex coincide and for which age is consistent across the two observations.

At the industry level, I regress individual income growth on aggregate income growth and recover the estimated parameter $\tilde{\gamma}_g$. I include a vector of controls for a cubic polynomial of age, sex, race, state, and level of education. In practice, I find that the conclusions do not change if we exclude the vector of controls. I renormalize $\gamma_g = \tilde{\gamma}_g / (\sum_g \pi_g \tilde{\gamma}_g)$ using the industry shares in 2018.

B.3 Robustness

Table 4: Robustness exercises 2

	Magnitude of Forecast Errors							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\overline{\gamma}_S$	-1.02	-1.73^{**}	-1.62^{**}	-1.56^{**}	-1.76^{**}	-1.64^{**}	-1.63^{**}	-1.19^{*}
Quarter F.E.	✓							✓
High-skill share		✓						✓
Numeracy share			✓					✓
Avg. Tenure				✓				✓
Avg. Age					✓			✓
Med. Income Share						✓		✓
High Income Share							✓	✓

Corrected forecast errors

In the baseline empirical exercise, I use the state's average forecast error as obtained in the data. However, note that in that baseline analysis, the unpredictable component of forecast errors is also increasing in income cyclicality. So, the expected magnitude of forecast errors can decrease in cyclicality only if attention increases sufficiently fast with income cyclicality.

In this appendix, in order to remove the dependence on the state's business cycle exposure, I divide the state's forecast error by its average exposure, $\overline{\gamma}_S$ and looking at the following regression:

$$(\overline{\text{FE}}_{S,t} / \overline{\gamma}_{S,t})^2 = \alpha + \beta \overline{\gamma}_{S,t} + \varepsilon_{S,t}. \tag{100}$$

The main result can be found in Figure 11 and the robustness results taking into account a variety of relevant control variables in Table 5.

Figure 11: Income cyclicality and the magnitude of forecast errors

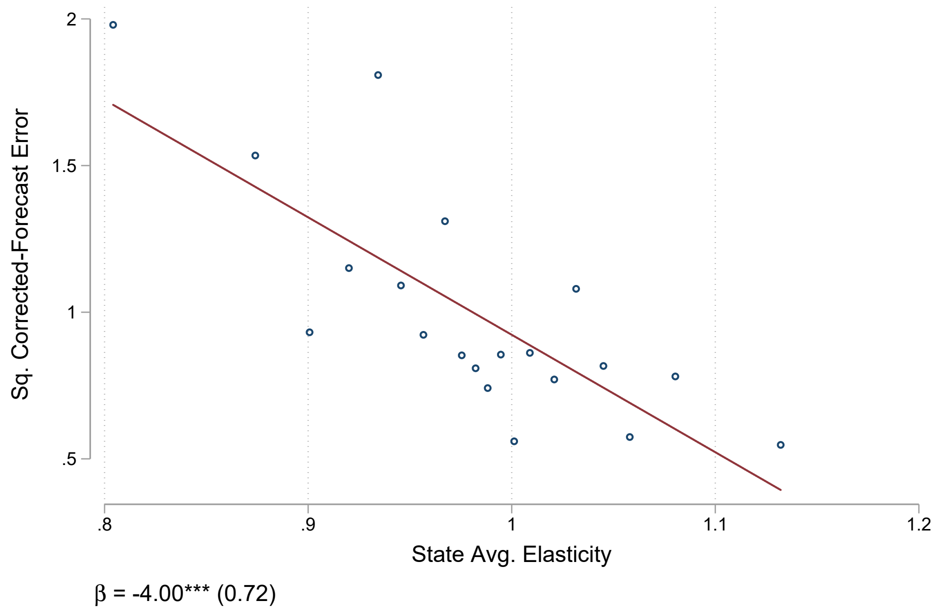

Table 5: Robustness exercises

	Magnitude of Forecast Errors							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\overline{\gamma}_S$	-3.37^{***}	-4.10^{***}	-4.00^{***}	-3.94^{***}	-4.15^{***}	-4.02^{***}	-4.02^{***}	-3.55^{***}
Time F.E.	✓							✓
High-skill share		✓						✓
Numeracy share			✓					✓
Avg. Tenure				✓				✓
Avg. Age					✓			✓
Med. Income Share						✓		✓
High Income Share							✓	✓

C Appendix to section 4

C.1 Unions and labor supply

In this section, I derive the wage Phillips curve. At time t, union u sets the wage to maximize

$$\sum_{h\geq 0}\beta^h\left[u'\left(C_{t+h}\right)\left(1-\tau_{t+h}\right)\frac{W_{u,t+h}N_{u,t+h}}{P_{t+h}}-v'\left(N_{t+h}\right)N_{u,t+h}-\frac{1}{2\tilde{\kappa}_w}\left(\frac{W_{u,t+h}}{W_{u,t+h-1}}-1\right)^2\right].$$

I assume that the union works to maximize a utility valuation of the income derived from the union labor supply and the utility cost of labor supply, subject to quadratic wage adjustment costs. Note that, to measure the utility valuation, I use the aggregates for consumption and labor supply, which implies that the union ignores the distributional consequences of its decisions. Alternatively, we could assume that the union maximizes an average utility valuation considering these distributional consequences, as in Auclert et al. (2018). In practice, this would make little quantitative difference but would have the computational cost of computing the average marginal utility of consumption and labor at each point. For these reasons, I focus on this more straightforward representation of the Phillips curve.

When setting the wage $W_{u,t}$, the union behaves monopolistically, taking into account the response of demand which is given by

$$N_{u,t} = \left(\frac{W_{u,t}}{W_t}\right)^{-\mu_w} N_t,$$

Taking first-order conditions, we obtain the following non-linear Phillips curve

$$\left(e^{\pi_t^w} - 1\right) e^{\pi_t^w} = \tilde{\kappa}_w \left(\mu_w - 1\right) \left[-u'\left(C_t\right)\left(1-\tau_t\right)Y_t + \frac{\mu_w}{\mu_w - 1} v'\left(N_t\right) N_t\right] + \beta_t \left(e^{\pi_{t+1}^w} - 1\right) e^{\pi_{t+1}^w}.$$

Linearizing this equation, we obtain

$$\pi_t^w = \kappa_w \left[\sigma^{-1} c_t + \psi^{-1} n_t - \left(y_t - \hat{\tau}_t - n_t\right)\right] + \beta \pi_{t+1}^w, \tag{101}$$

where $\kappa_w \equiv \tilde{\kappa}_w \mu_w v'\left(N\right) N$.

C.2 Jacobians without FIRE in Section 4.6

In this appendix, I provide the central sketch for the result in Section 4.6. Following Auclert et al. (2021) and Auclert et al. (2020), I consider a generic representation of a heterogeneous-agent problem as a mapping from some inputs $\mathbf{X}_t$ to a time-path of aggregates $\mathbf{C}_t$. In the model of this paper, the heterogeneous-agent blocks are each group of households, the aggregates are the group's average consumption and savings, and the inputs are their incomes, taxes, and interest rates. To simplify, I will work with a representation with a single input and output, but the analysis can be easily extended to multiple inputs and outputs, see Auclert et al. (2021). Furthermore, I also assume that all individuals in a single group share the same beliefs. This can be easily extended.

Let v_t denote the marginal utility of consumption The generic problem is

$$\mathbf{v}_t = u\left(\mathbf{v}_{t+1}^{e,t}, X_t\right), \quad \text{for } t \geq 0 \tag{102}$$

$$\mathbf{v}_s^{e,t} = u\left(\mathbf{v}_{s+1}^{e,t}, X_s^{e,t}\right), \quad \text{for } t \geq 0, s \geq t+1 \tag{103}$$

$$\mathbf{D}_{t+1} = \Lambda\left(\mathbf{v}_{t+1}^{e,t}, X_t\right)' \mathbf{D}_t, \quad \text{for } t \geq 0 \tag{104}$$

$$C_t = c\left(\mathbf{v}_{t+1}^{e,t}, X_t\right)' \mathbf{D}_t. \tag{105}$$

Here $\mathbf{v}_t$ is the marginal utility of consumption which is related to the future expected future marginal utility of consumption $\mathbf{v}_{t+1}^{e,t}$ and the input today X_t. The problem is discretized to n_g grid points, so $\mathbf{v}_t$ is $n_g \times 1$. The distribution over these grid points is given by D_{t+1} and the individual consumption choices are given by $c\left(\mathbf{v}_{t+1}^{e,t}, X_t\right)$. Equation (102) thus represents the Euler equation, and (103) defines the predicted future Euler equations. Equation (104) determines how the distribution is updated given some transition matrix Λ. Finally, equation (105) determines how individual choices are aggregated.

We are interested in the consumption response to an increase in the actual variable X_t and the expectations expectation $X_s^{e,t}$. As it turns out, we can write the response to an unanticipated ∂X_t as follows

$$\frac{\partial C_t}{\partial X_s} = \begin{cases} 0 & \text{if } t < s, \\ \mathcal{J}_{t-s,0} & \text{if } t \geq s, \end{cases}$$

where $\mathcal{J}$ denotes the FIRE Jacobian. The effects of a shock to beliefs can be written as

$$\frac{\partial C_t}{\partial X_s^{e,m}} = \begin{cases} 0 & \text{if } t < m \text{ or } s \leq m, \\ \frac{\partial C_{t-m}}{\partial X_{s-m}^{e,0}} = \mathcal{J}_{t-m,s-m} - \mathcal{J}_{t-m-1,s-m-1} & \text{if } t > m \text{ and } s > m, \\ \mathcal{J}_{0,s-m} & \text{if } t = m \text{ and } s > m. \end{cases}$$

It follows that

$$dC_t = \sum_{s=0}^{t} \mathcal{J}_{t-s,0} dX_s + \sum_{m=0}^{t-1} \sum_{s=m+1}^{\infty} \left(\mathcal{J}_{t-m,s-m} - \mathcal{J}_{t-m-1,s-m-1}\right) dX_s^{e,m} + \sum_{s=t+1}^{\infty} \mathcal{J}_{0,s-t} dX_s^{e,t}$$

$$\Leftrightarrow dC_t = \sum_{s=0}^{t} \mathcal{J}_{t-s,0}\left(dX_s - dX_s^{e,s-1}\right) + \sum_{m=1}^{t} \sum_{s=m+1}^{\infty} \mathcal{J}_{t-m,s-m}\left(dX_s^{e,m} - dX_s^{e,m-1}\right) + \sum_{s=0}^{\infty} \mathcal{J}_{t,s} dX_s^{e,0}.$$

Using the defition that $X_s^{e,t} = X_s$ if $s \leq t$, we can write the above expression in vector form:

$$d\mathbf{C} = \mathcal{J} \cdot \underbrace{d\mathbf{X}^{e,0}}_{\text{Initial belief}} + \sum_{t \geq 1} \mathcal{R}_t \cdot \underbrace{\left(d\mathbf{X}^{e,t} - d\mathbf{X}^{e,t-1}\right)}_{\text{Forecast Revision at time } t}, \tag{106}$$

where $\mathcal{J} \equiv [\mathcal{J}_{t,s}]$ and

$$\mathcal{R}_t \equiv \begin{bmatrix} 0 & \mathbf{0}_t' \\ \mathbf{0}_t' & \mathcal{J} \end{bmatrix}.$$

C.3 Utility cost of inattention

For every (a,z) such that the policy function implies $c^*_{g,t}(a,z) < (1-\tau_t)y(z_{i,t}, Y_{g,t}) + (1-r_t)a$, we can proceed in the same way as before. I define the value function of an individual with full attention as follows:

$$V_{i,t}(a,z) = \max_c \{u(c) + \beta_{i,t}\mathbb{E}_t[V_{i,t+1}((1-\tau_t)y(z_{i,t}, Y_{g,t}) + (1-r_t)a - c, z')]\} \tag{107}$$

and the objective function in this problem is

$$v_{i,t}(c;a,z) \equiv u(c) + \beta_{i,t}\mathbb{E}_t[V_{i,t+1}((1-\tau_t)y(z_{i,t}, Y_{g,t}) + (1-r_t)a - c, z'). \tag{108}$$

Instead, the problem of an inattentive individual is given by

$$u(c) + \beta_{i,t}E_{i,t}[V_{i,t+1}((1-\tau_t)y(z_{i,t}, Y_{g,t}) + (1-r_t)a - c, z')]. \tag{109}$$

Following the same steps, we find that the realized utility cost of inattention is given by

$$v_{i,t}(c^*_{g,t}(a,1);a,z) - v_{i,t}(c^*_{g,t}(a,\lambda_i);a,z) = -\frac{1}{2}\frac{\partial^2 v(c(a,z);a,z)}{\partial c^2}\left(c^*_{g,t}(a,z,1) - c^*_{g,t}(a,z,\boldsymbol{\lambda}_i)\right)^2$$
$$= -\frac{1}{2}\frac{\partial^2 v(c(a,z);a,z)}{\partial c^2}\sum_{X,\tilde{X},h,\tilde{h}}\frac{\partial c(a,z)}{\partial X_h}\frac{\partial c(a,z)}{\partial \tilde{X}_{\tilde{h}}}(1-\lambda^X_{i,h})(1-\lambda^{\tilde{X}}_{i,\tilde{h}})X_{t+h}X_{t+\tilde{h}}.$$

It follows that the ex-ante utility cost of inattention is given by

$$\mathcal{C}_g(a,z)(\boldsymbol{\lambda}_i) = -\frac{1}{2}\frac{\partial^2 v(c(a,z);a,z)}{\partial c^2}\sum_{X,\tilde{X},h,\tilde{h}}\frac{\partial(a,z)}{\partial X_h}\frac{\partial c(a,z)}{\partial \tilde{X}_{\tilde{h}}}\left(1-\lambda^X_{i,h}\right)\left(1-\lambda^{\tilde{X}}_{i,\tilde{h}}\right)\sigma_{X_h,\tilde{X}_{\tilde{h}}}. \tag{110}$$

where $\sigma_{X_h,\tilde{X}_{\tilde{h}}}$ denotes the ex-ante perceived covariance between X_{t+h} and $\tilde{X}_{t+\tilde{h}}$.

If $c^*_{g,t}(a,z) = (1-\tau_t)y(z_{i,t}, Y_{g,t}) + (1-r_t)a$ then the individual is at the borrowing constraint. Note that if an individual is borrowing constrained, then their consumption is not changing given changes in future variables, i.e.,

$$\frac{\partial c(a,z)}{\partial X_h} = 0, \tag{111}$$

for $h \geq 1$. This implies that the misoptimization costs of inattention to future variables are exactly zero, i.e., $\mathcal{C}_g(a,z)(\boldsymbol{\lambda}_i) = 0$. These two facts put together allow us to write

$$\mathcal{C}_g(a,z,\boldsymbol{\lambda}_i) = -\frac{1}{2}\frac{\partial^2 v(c(a,z);a,z)}{\partial c^2}\sum_{X,\tilde{X},h,\tilde{h}}\frac{\partial c(a,z)}{\partial X_h}\frac{\partial c(a,z)}{\partial \tilde{X}_{\tilde{h}}}\left(1-\lambda^X_{i,h}\right)\left(1-\lambda^{\tilde{X}}_{i,\tilde{h}}\right)\sigma_{X_h,\tilde{X}_{\tilde{h}}}, \tag{112}$$

which equals zero since the partial derivatives of the consumption function are equal to zero.

D Appendix to section 5

D.1 Individual demand

The household enters time t with assets $A_{i,t}$ and chooses consumption and savings to solve

$$\max E_{i,t} \sum_{h=0}^{\infty} \prod_{s=0}^{h-1} (\beta_{i,t+s}) \left[u(C_{i,t}) - v\left(N_{i,t}\right)\right], \quad \text{subject to}$$
$$C_{i,t+h} + A_{i,t+h+1} = (1-\tau_{t+h})Y_{g,t+h} + (1+r)A_{i,t+h}.$$

The Euler equation is still given by (73) and its linearized form (74).

Linearizing the budget constraint, we obtain

$$\begin{aligned} & c_{i,t+h} + a_{i,t+h+1} = (1-\tau)\cdot y_{g,t+h} - d\tau_{t+h}\cdot Y + \beta^{-1}a_{i,t+h} \\ \Leftrightarrow & c_{i,t+h} + a_{i,t+h+1} = y_{g,t+h} - \tau_{t+h} + \beta^{-1}a_{i,t+h} \end{aligned}$$

since $\tau = 0$ and $Y = 1$ in steady state. We can again aggregate flow-of-funds constraints and obtain

$$\sum_{h=0}^{\infty} \beta^h c_{i,t+h} = \sum_{h=0}^{\infty} \beta^h [y_{g,t+h} - \tau_{t+h}] + \beta^{-1}a_{i,t}. \tag{113}$$

Proceeding as before finally shows that

$$c_{i,t} = (1-\beta)\sum_{h=0}^{\infty} \beta^h E_{i,t}[y_{g,t+h} - \tau_{t+h}] + (1-\beta)\beta^{-1}a_{i,t} + \sum_{h=0}^{\infty} \beta^{h+1} r^n_{i,t+h}.$$

D.2 Proof of proposition 4

The government-spending multiplier satisfies

$$\frac{dy_t}{dG_t} = 1 + (1-\beta)\sum_{h=1}^{\infty} \beta^{h-1}\left[(1+\text{CD})\cdot\overline{\lambda}^Y \frac{dy_{t+h}}{dG_{t+h}} - \overline{\lambda}^\tau\right]\rho_G^h. \tag{114}$$

I guess and verify that the multiplier is constant over time $dy_t/dG_t = \Omega$:

$$\begin{aligned} \Omega &= 1 + \varrho_G\left[(1+\text{CD})\cdot\overline{\lambda}^Y\Omega - \overline{\lambda}^\tau\right] \\ \Leftrightarrow \Omega &= \frac{1-\varrho_G\cdot\overline{\lambda}^\tau}{1-\varrho_G\cdot(1+\text{CD})\cdot\overline{\lambda}^Y} > 0, \end{aligned}$$

where $\varrho_G \equiv (1-\beta)\rho_G/(1-\beta\rho_G) \in (0,1)$.

Note that

$$\frac{d\Omega}{d\text{CD}} = \Omega \cdot \frac{\varrho_G \cdot \overline{\lambda}^Y}{1 - \varrho_G \cdot (1 + \text{CD}) \cdot \overline{\lambda}^Y} > 0 \tag{115}$$

and

$$\Omega \geq 1 \Leftrightarrow \frac{1 - \varrho_G \cdot \overline{\lambda}^\tau}{1 - \varrho_G \cdot (1 + \text{CD}) \cdot \overline{\lambda}^Y} \geq 1 \Leftrightarrow (1 + \text{CD}) \cdot \overline{\lambda}^Y \geq \overline{\lambda}^\tau. \tag{116}$$

D.3 Proof of proposition 5

The government-spending multiplier satisfies

$$\frac{dy_t}{dG_t} = 1 + (1 - \beta) \sum_{h=1}^{\infty} \beta^{h-1} \left[(1 + \text{CD}) \cdot \overline{\lambda}^Y \frac{dy_{t+h}}{dG_{t+h}} + \text{TC} \cdot \overline{\lambda}^Y - \overline{\lambda}^\tau \right] \rho_G^h. \tag{117}$$

I guess and verify that the multiplier is constant over time $dy_t/dG_t = \Omega$:

$$\begin{aligned} \Omega &= 1 + \varrho_G \left[(1 + \text{CD}) \cdot \overline{\lambda}^Y \Omega - \overline{\lambda}^\tau \right] \\ \Leftrightarrow \Omega &= \frac{1 - \varrho_G \cdot (\overline{\lambda}^\tau + \text{TC} \cdot \overline{\lambda}^Y)}{1 - \varrho_G \cdot (1 + \text{CD}) \cdot \overline{\lambda}^Y} \\ \Leftrightarrow \Omega &= \frac{dy_t^0}{dG_t} + \frac{\varrho_G \cdot \text{TC} \cdot \overline{\lambda}^Y}{1 - \varrho_G \cdot (1 + \text{CD}) \cdot \overline{\lambda}^Y} \end{aligned}$$

where $\varrho_G \equiv (1 - \beta)\rho_G / (1 - \beta\rho_G) \in (0, 1)$, and $dy_t^0 / dG_t \equiv \frac{1 - \varrho_G \cdot \overline{\lambda}^\tau}{1 - \varrho_G \cdot (1 + \text{CD}) \cdot \overline{\lambda}^Y}$.

With homogeneous beliefs $\lambda_g^Y = \overline{\lambda}^Y$ (of which FIRE is a special case with $\lambda_g = 1$), we find that

$$TC = \mathbb{C}\text{ov}\left(\omega_g, \overline{\lambda}^Y / \overline{\lambda}^Y \right) = 0.$$

This means that

$$dy_t / dG_t = \frac{1 - \varrho_G \cdot \overline{\lambda}^\tau}{1 - \varrho_G \cdot (1 + \text{CD}) \cdot \overline{\lambda}^Y},$$

and so targeting does not affect the spending multiplier.

Instead, suppose that attention is heterogeneous. Then,

$$\frac{d\Omega}{d\text{TC}} = \frac{\varrho_G \cdot \overline{\lambda}^Y}{1 - \varrho_G \cdot (1 + \text{CD}) \cdot \overline{\lambda}^Y} > 0, \tag{118}$$

which implies that the spending multiplier increases if the covariance between ω_g and λ_g is higher.

E Extensions

E.1 Horizon-independent attention

The baseline quantitative model allows attention to vary with horizon. In this appendix, I assess the robustness of the quantitative results to assuming that horizon must be constant over time. The interest for this analysis stems from the fact that this independence on horizon is closest to what would be obtained in a model with rational expectations in which individuals receive noisy signals of the fundamental shocks, as in Angeletos and Huo (2018), or in models of sticky information, as in Mankiw and Reis (2002).

Formally, I now assume that beliefs are given by

$$E_{i,t}[dX_{t+h}] = \lambda_i^X \cdot \mathbb{E}_t[dX_{t+h}] + (1-\lambda_i^X) \cdot E_{i,t-1}[dX_{t+h}], \tag{119}$$

and proceed as before to optimize for λ_i^X for each variable X.

$$\lambda_g^Y = \max\left\{0, 1 - \frac{\kappa^Y}{\sum_z \int \frac{\partial^2 v(a,z)}{\partial c^2} \sum_{h=1}^{\infty} \left(\frac{\partial c(a,z)}{\partial Y_{g,h}}\right)^2 D(da,z) \cdot \gamma_g^2 \sigma_Y^2}\right\} \tag{120}$$

and

$$\lambda_g^X = \max\left\{0, 1 - \frac{\kappa^X}{\sum_z \int \frac{\partial^2 v(a,z)}{\partial c^2} \sum_{h=1}^{\infty} \left(\frac{\partial c(a,z)}{\partial X_h}\right)^2 D(da,z) \cdot \sigma_X^2}\right\} \tag{121}$$

for $X = \tau, r$.

The quantitative results can be found below. In sum, the results emphasized in the main text also hold in this case.

Figure 12: Optimal attention

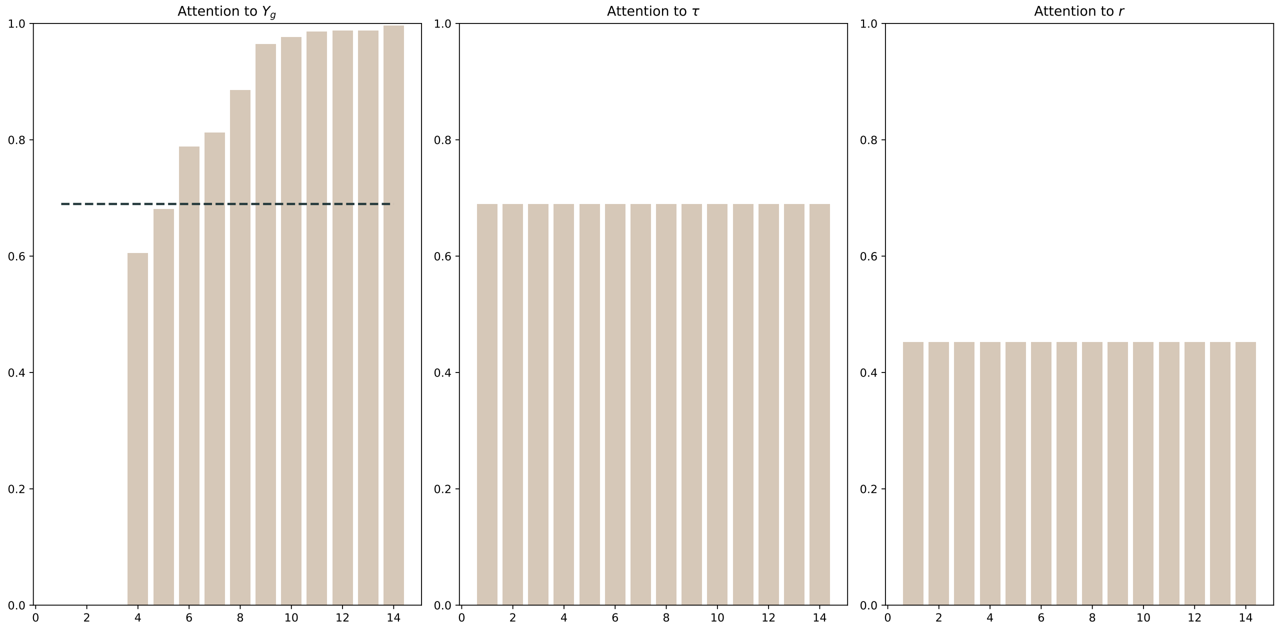

Figure 13: Consumption response

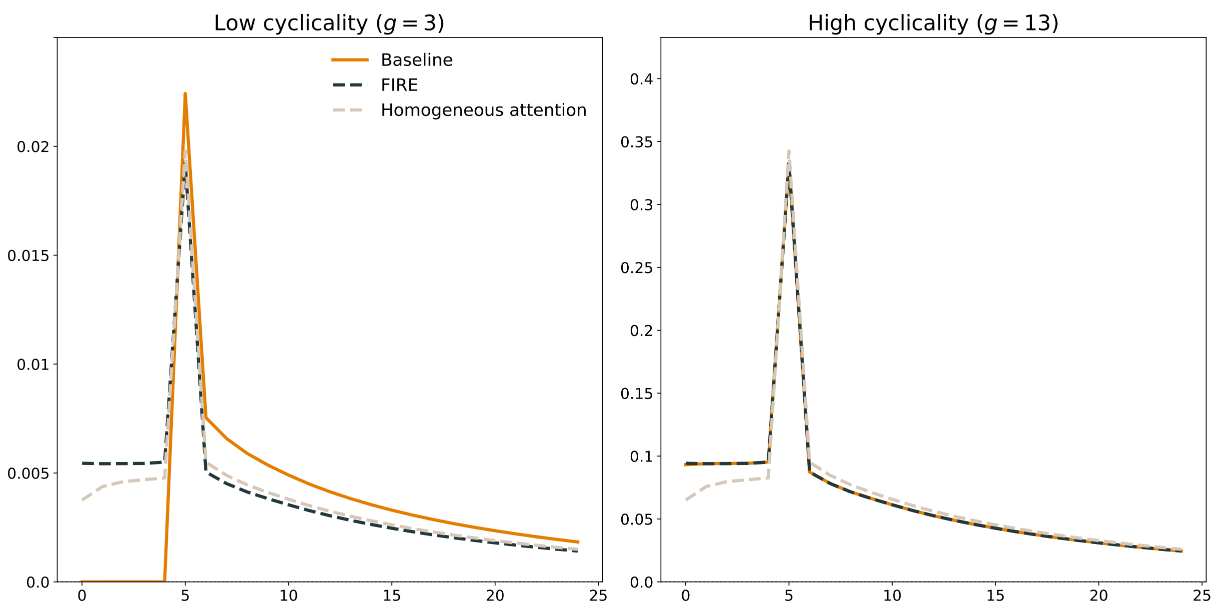

Figure 14: Consumption response for all groups

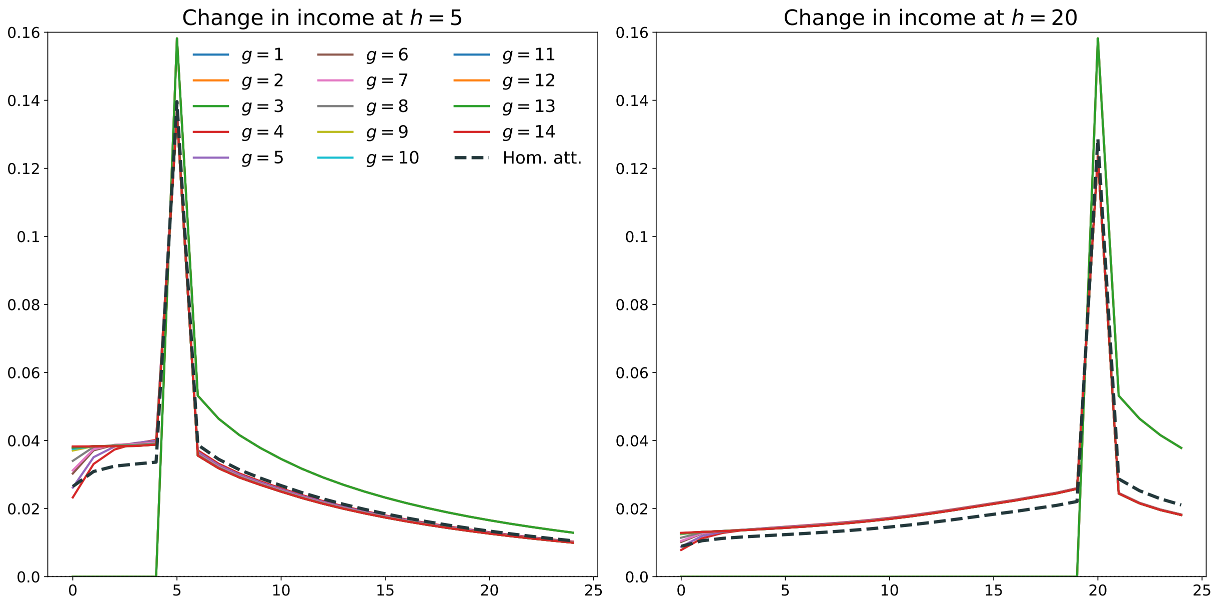

Figure 15: Business-cycle amplification

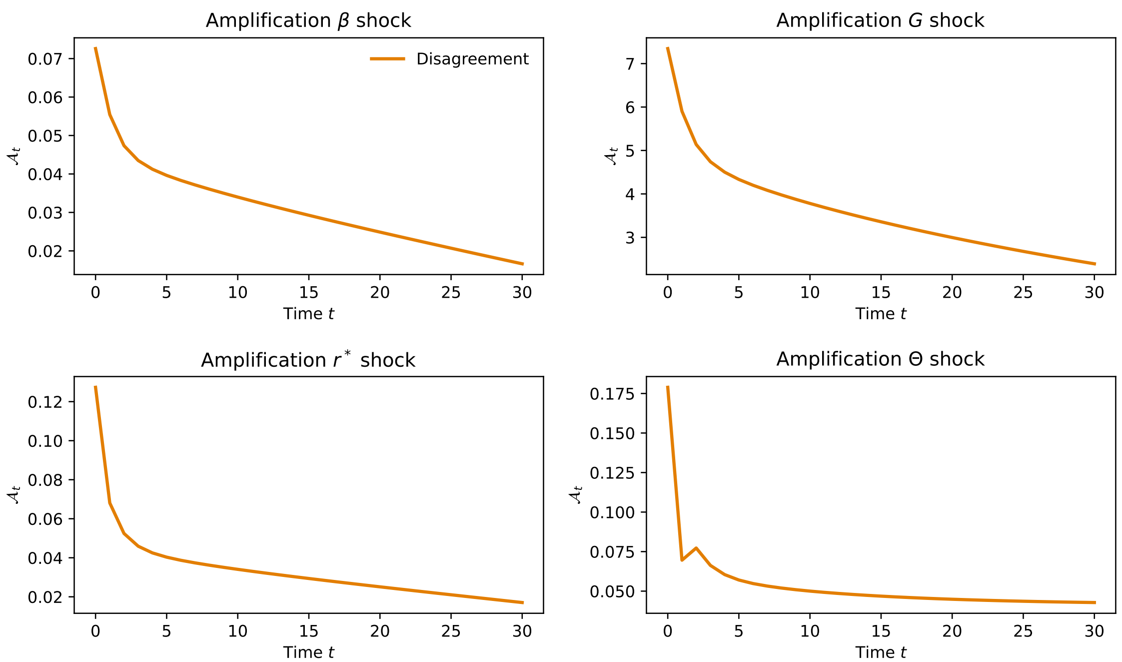

Figure 16: Business-cycle amplification: The role of persistence

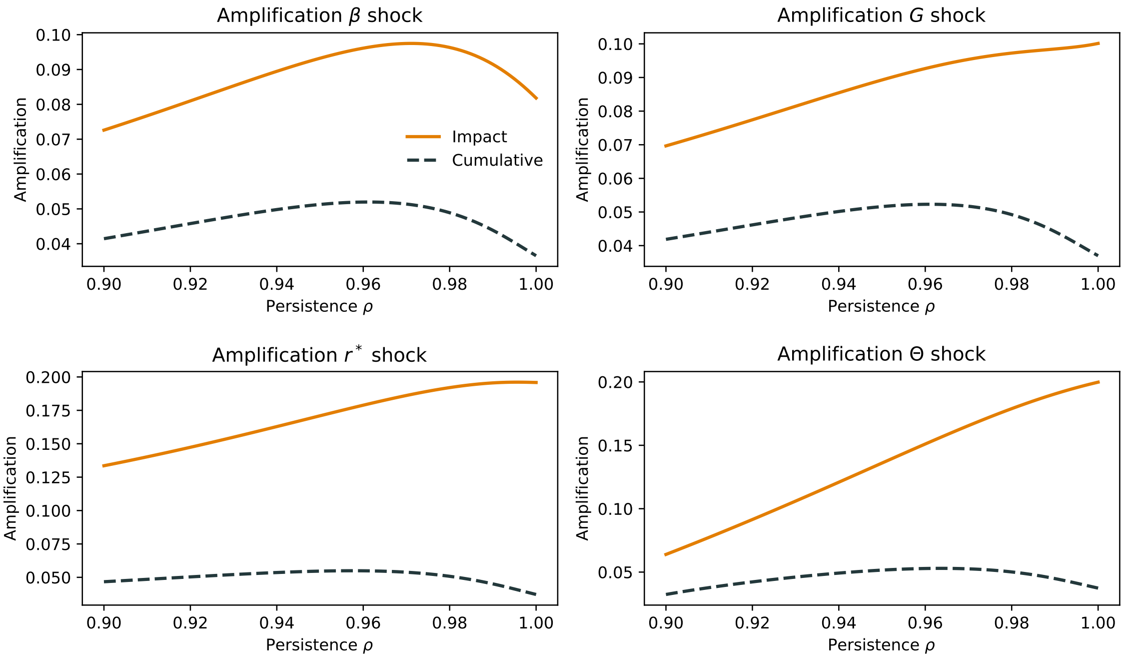

Figure 17: Business-cycle amplification: The role of monetary policy

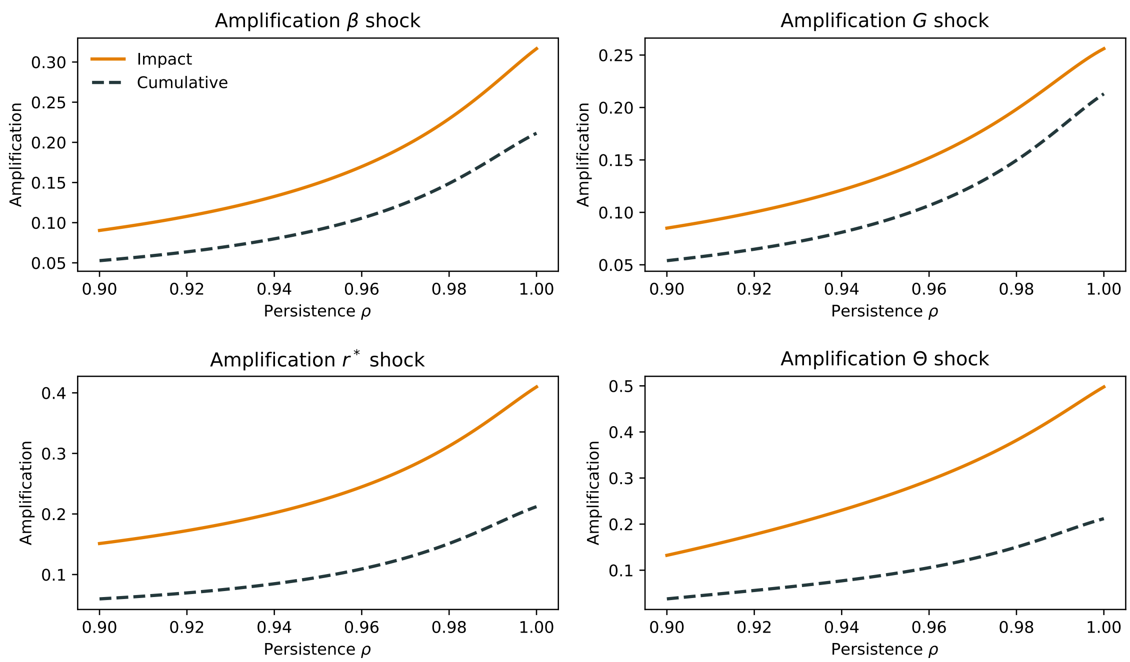

Figure 18: Targeted spending multipliers

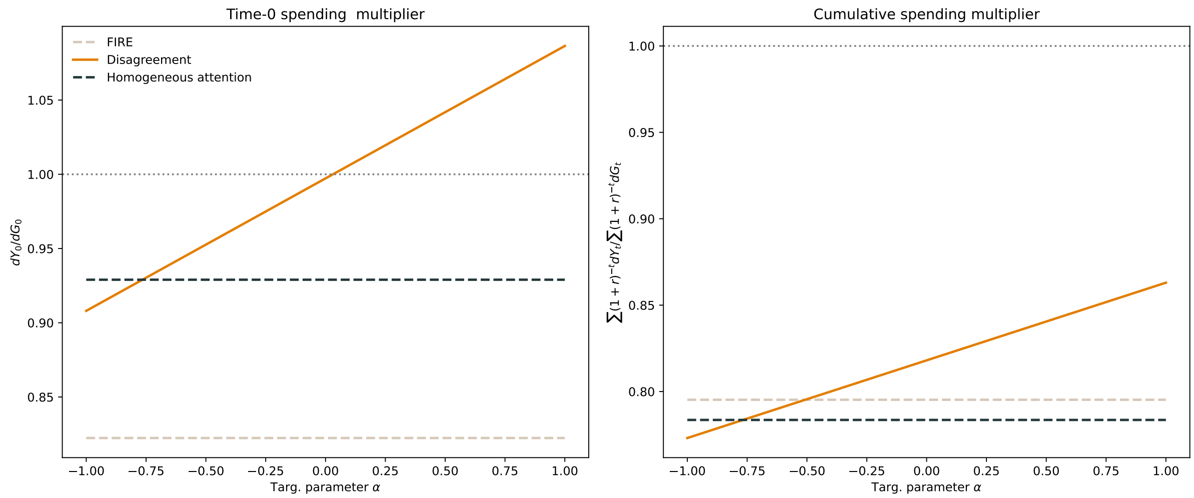

Figure 19: Targeted spending multipliers: The role of monetary policy

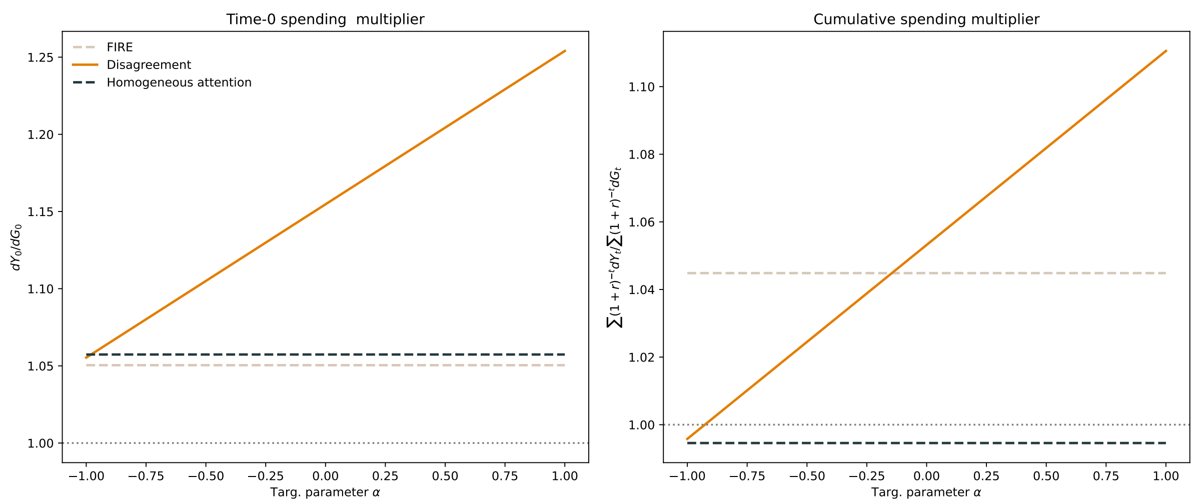

E.2 Allowing for perceived correlations

Following Gabaix (2014), in the baseline quantitative model, I assume that individuals perceive all variables to be uncorrelated. In this appenix, I show that the quantitative results are robust to allowing for perceptions in correlations. Specifically, I assume that individuals perceive each variable to evolve as an AR(1):

$$dX_{t+1} = \rho_X dX_t + \eta^X_{t+1}, \tag{122}$$

where ρ_X denotes the persistence of variable X and $\mathbb{V}\text{ar}(dX_t) = \sigma^2_X$ and $\mathbb{C}\text{orr}(dX_t, d\tilde{X}_t) = \rho_{X,\tilde{X}}$. I assume that peoples perceived correlations are equal to their empirical counterparts and use U.S. data to estimate these correlations. The empirical standard deviations and correlations can be found in table 6.

Table 6: Empirical covariances

Param.	Description	Value	Param.	Description	Value
σ_y	St. Dev. GDP	0.03	ρ_r	Persistence real rate	0.59
σ_r	St. Dev. real rate	0.13	$\rho_{y,\text{tax}}$	Correl. GDP-taxes	0.51
σ_{tax}	St. Dev. taxes	0.01	$\rho_{y,r}$	Correl. GDP-real rate	0.12
ρ_y	Persistence GDP	0.87	$\rho_{\text{tax},r}$	Correl. taxes-real rate	-0.09
ρ_{tax}	Persistence Taxes	0.89			

Solving the optimal attention problem becomes more difficult, since the problem is no longer separable across forecasting variables. However, it is still possible to find a closed form solution to this problem which takes the form:

$$\boldsymbol{\lambda}_g = 1 - \boldsymbol{\Psi}^{-1}\boldsymbol{\kappa}$$

where

$$\boldsymbol{\lambda}_g = \begin{bmatrix} \lambda_1^Y \\ \lambda_2^Y \\ \dots \\ \lambda_1^\tau \\ \dots \\ \lambda_1^r \\ \dots \end{bmatrix}, \quad \boldsymbol{\Psi} \equiv \begin{bmatrix} \Lambda_{Y_0,Y_0}\sigma_{Y_0,Y_0} & \Lambda_{Y_0,Y_1}\sigma_{Y_0,Y_2} & \dots & \Lambda_{Y_0,\tau_0}\sigma_{Y_0,\tau_0} & \dots \\ \Lambda_{Y_1,Y_0}\sigma_{Y_1,Y_0} & \Lambda_{Y_1,Y_1}\sigma_{Y_1,Y_2} & \dots & \Lambda_{Y_1,\tau_0}\sigma_{Y_1,\tau_0} & \dots \\ \dots & \dots & \dots & \dots & \dots \\ \Lambda_{\tau_0,Y_0}\sigma_{\tau_0,Y_0} & \Lambda_{\tau_0,Y_1}\sigma_{\tau_0,Y_2} & \dots & \Lambda_{\tau_0,\tau_0}\sigma_{\tau_0,\tau_0} & \dots \\ \dots & \dots & \dots & \dots & \dots \\ \Lambda_{r_0,Y_0}\sigma_{r,Y_0} & \Lambda_{\tau_0,Y_1}\sigma_{r_0,Y_2} & \dots & \Lambda_{r_0,\tau_0}\sigma_{r_0,\tau_0} & \dots \\ \dots & \dots & \dots & \dots & \dots \end{bmatrix}, \quad \boldsymbol{\kappa} \equiv \begin{bmatrix} \kappa_Y \\ \kappa_Y \\ \dots \\ \kappa_\tau \\ \dots \\ \kappa_r \\ \dots \end{bmatrix}.$$

The quantitative results can be found below. In sum, the results emphasized in the main text also hold in this case.

Figure 20: Optimal attention

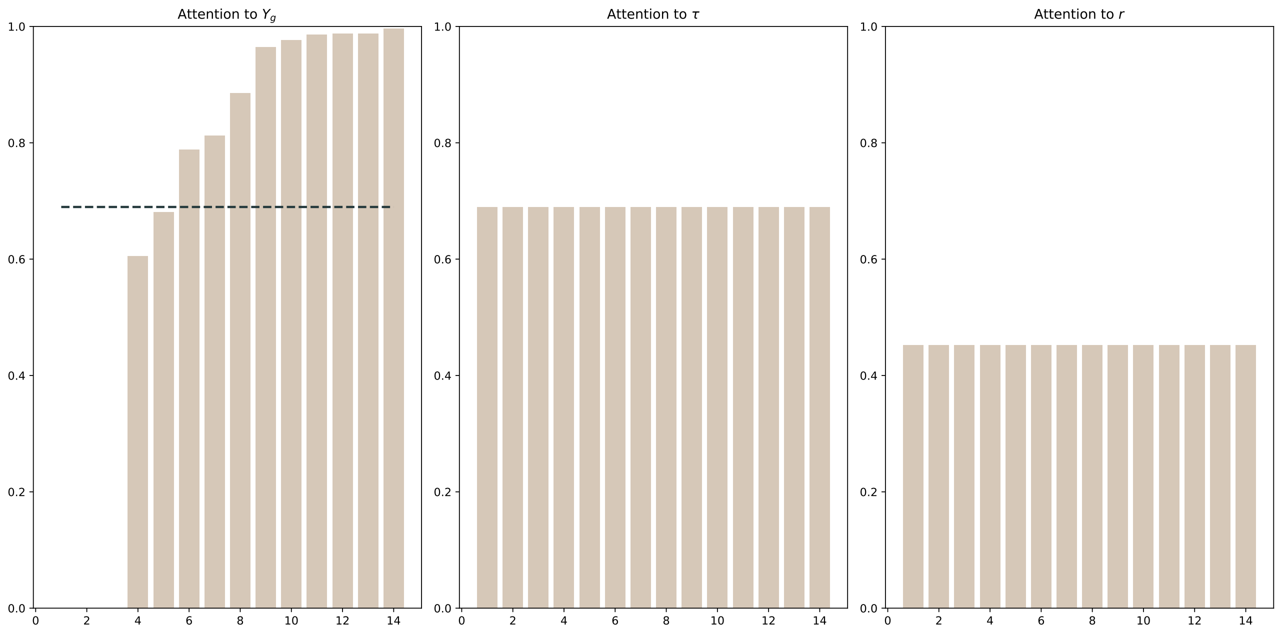

Figure 21: Consumption response

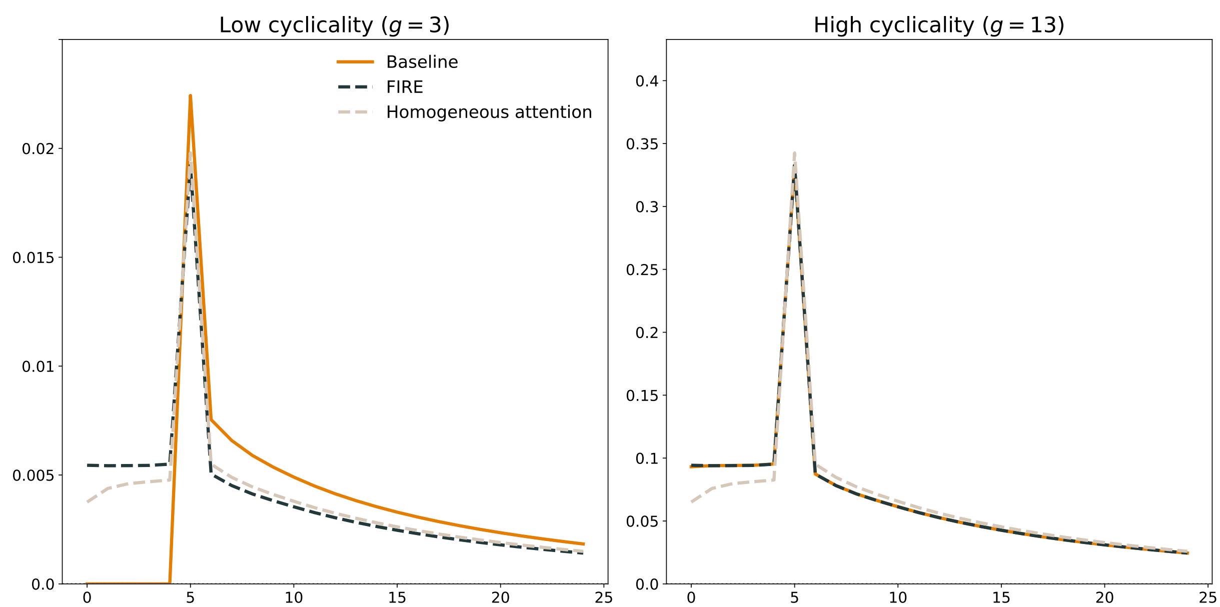

Figure 22: Consumption response for all groups

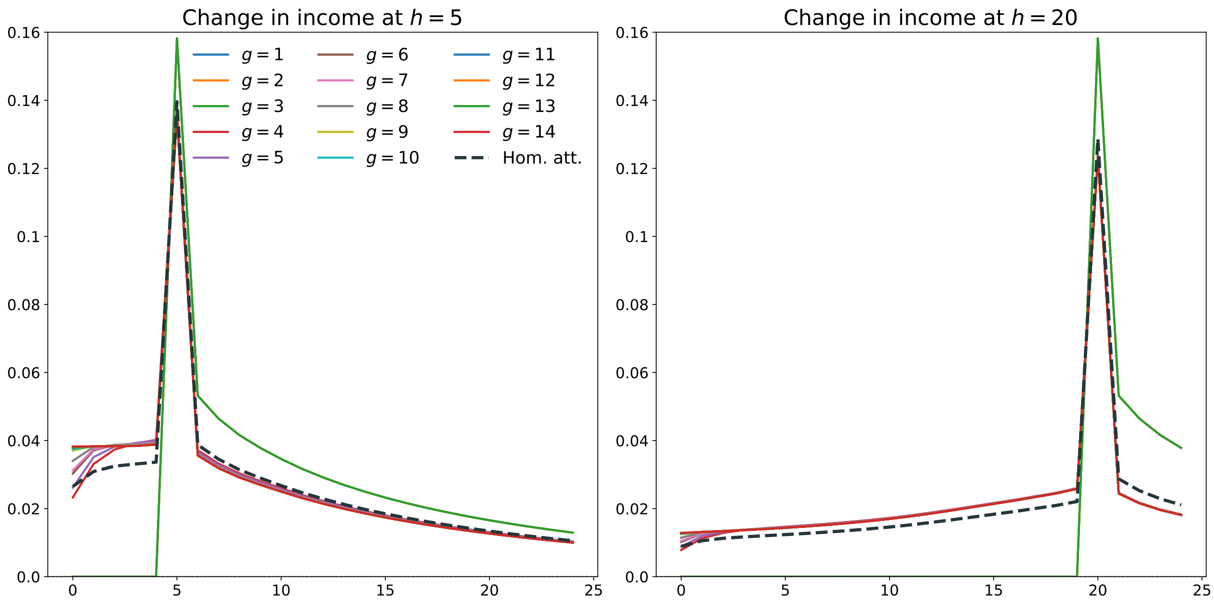

Figure 23: Business-cycle amplification

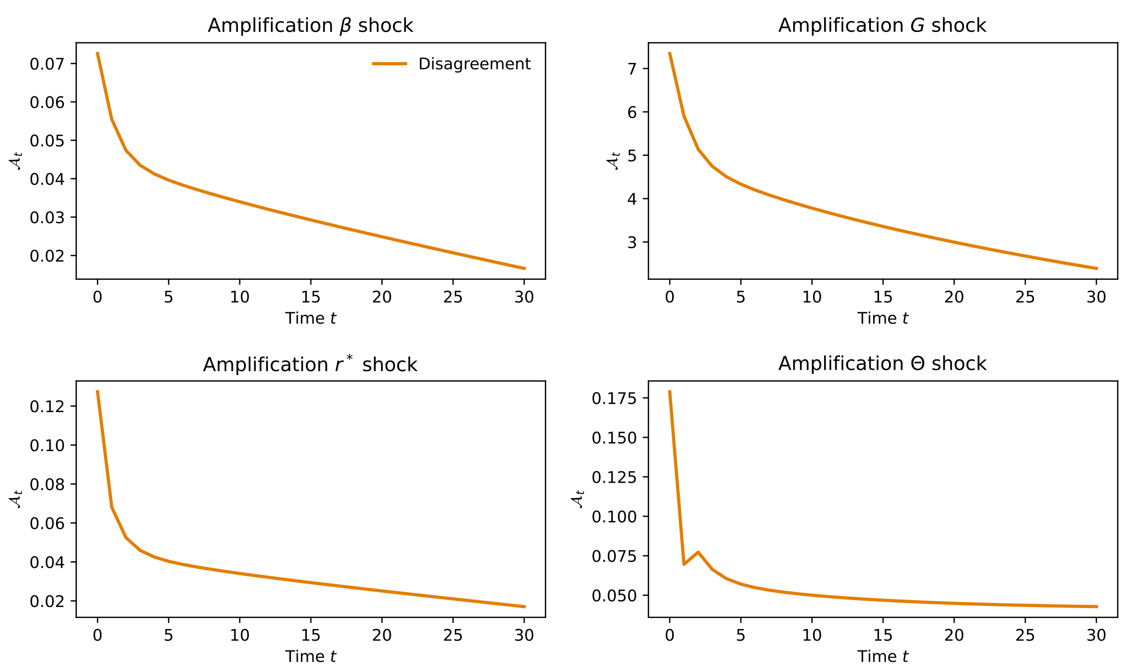

Figure 24: Business-cycle amplification: The role of persistence

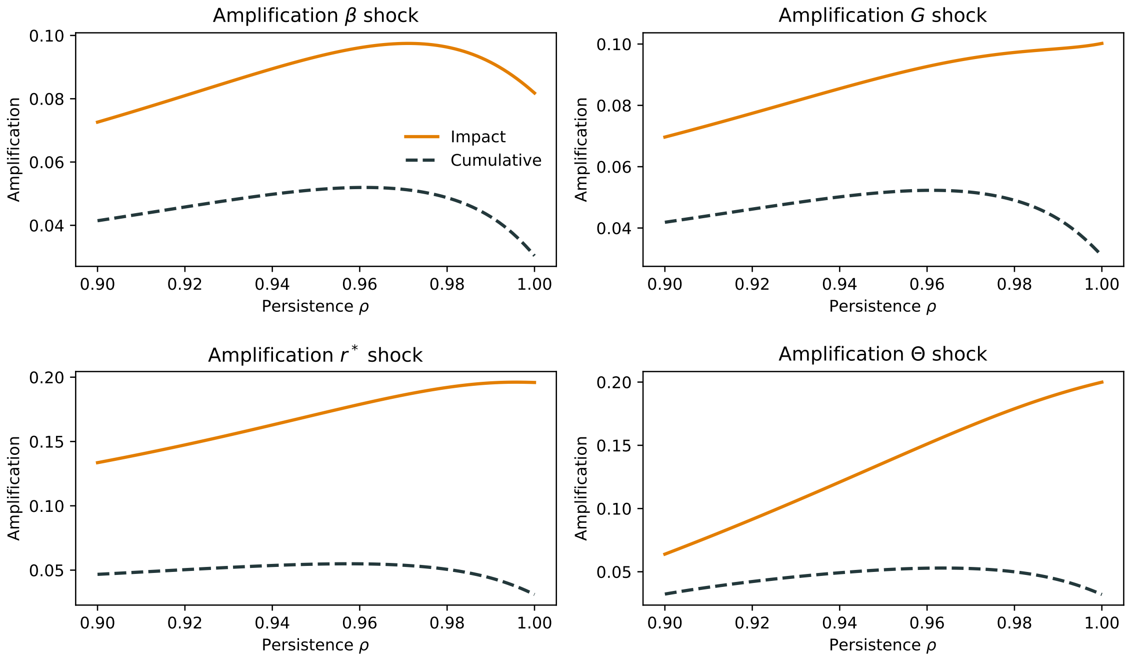

Figure 25: Business-cycle amplification: The role of monetary policy

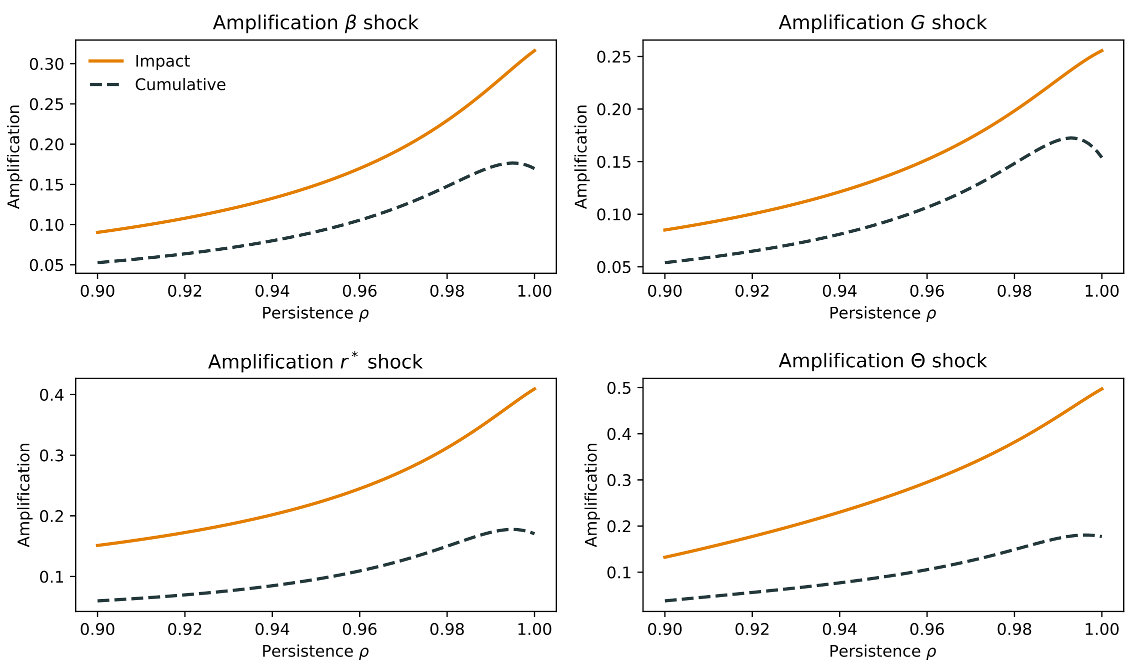

Figure 26: Targeted spending multipliers

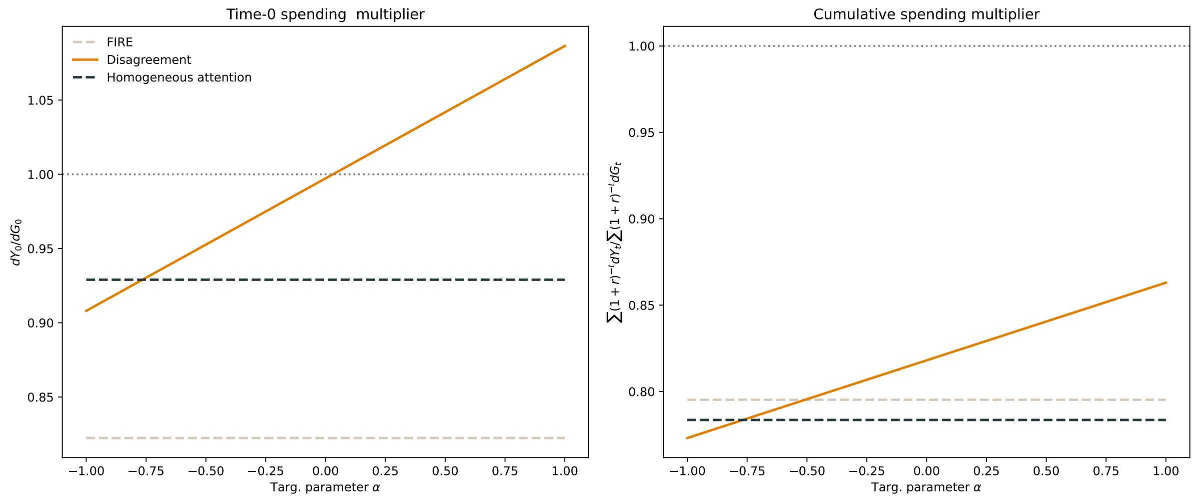

Figure 27: Targeted spending multipliers: The role of monetary policy

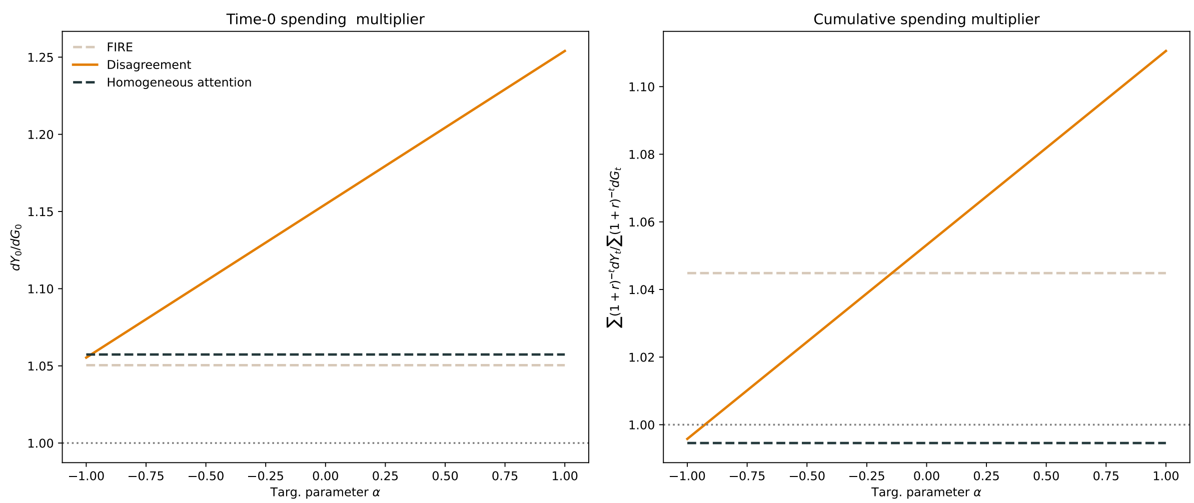

E.3 Budget deficits

The baseline quantitative model assumes that debt is kept constant at its steady-state level $B_t = B$. In this appendix, I allow debt to vary over time by assuming a fiscal rule for taxes as in Auclert et al. (2020). Formally, I assume that

$$\tau_t = \tau + \psi \frac{B_t - B}{Y}. \tag{123}$$

This expression implies that taxes are updated smoothly so that in the long-run government debt converges back to the original steady state. However, note that it also implies that upon increasing spending, the government only starts raising taxes a period later. In this sense, this extension allows for budget deficits. Following Auclert et al. (2020), I assume that the response of taxes to deviations of debt is given by $\psi = 0.1$ per annum, which is in line with the empirical results of the fiscal literature.

Figure 28: Optimal attention

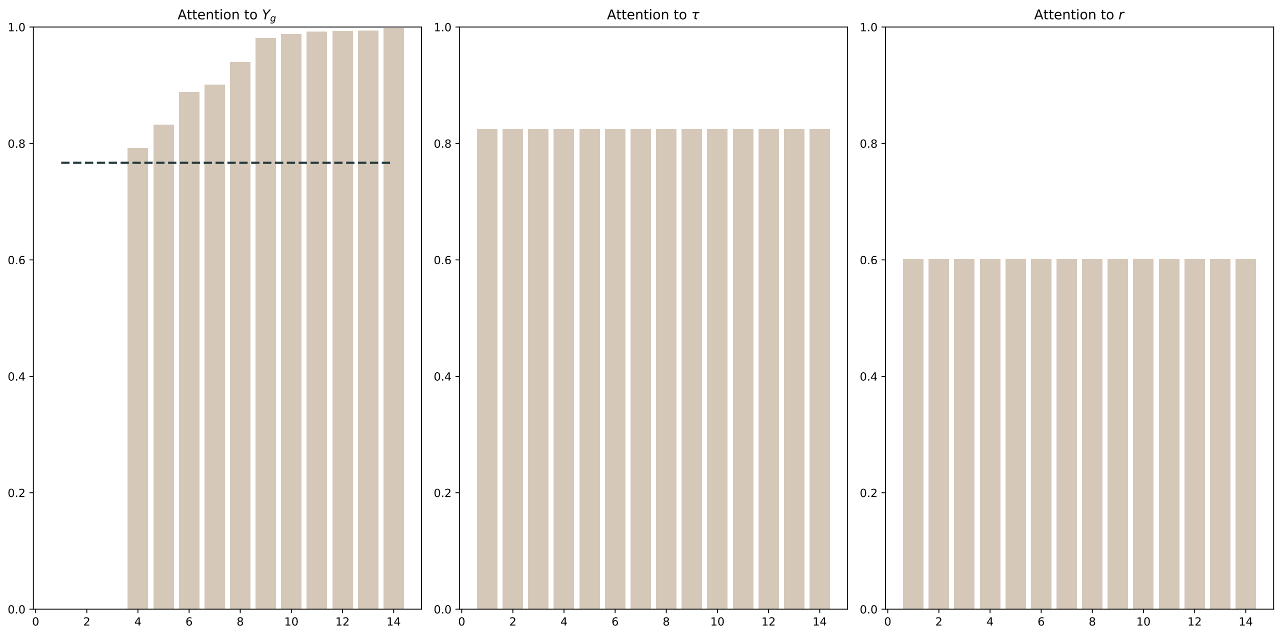

Figure 29: Consumption response

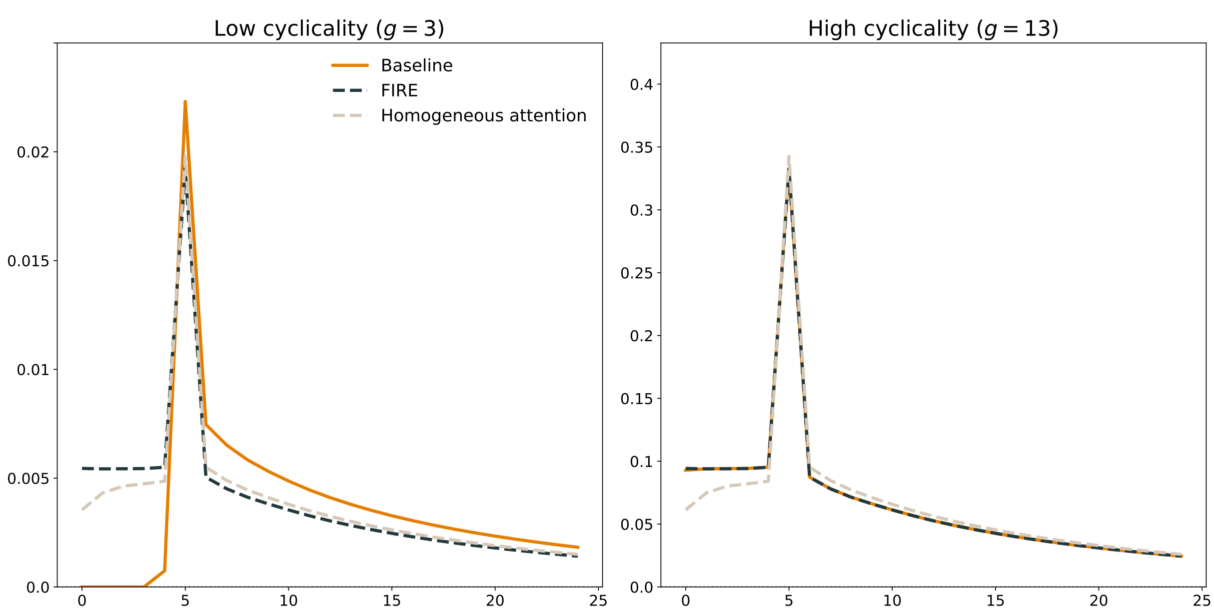

Figure 30: Consumption response for all groups

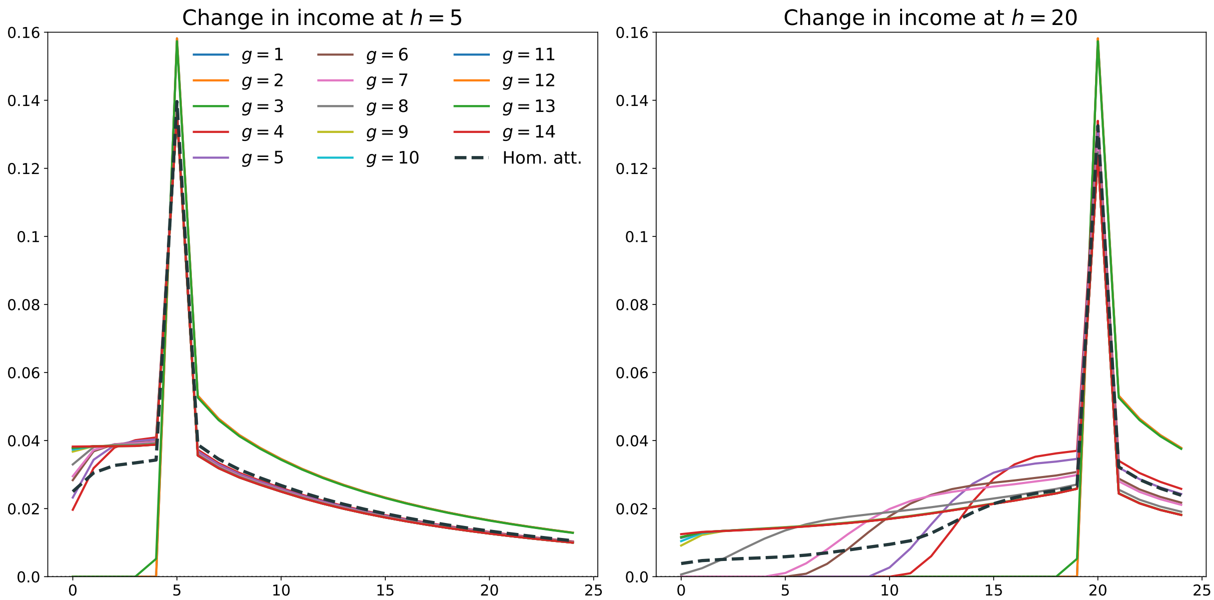

Figure 31: Business-cycle amplification

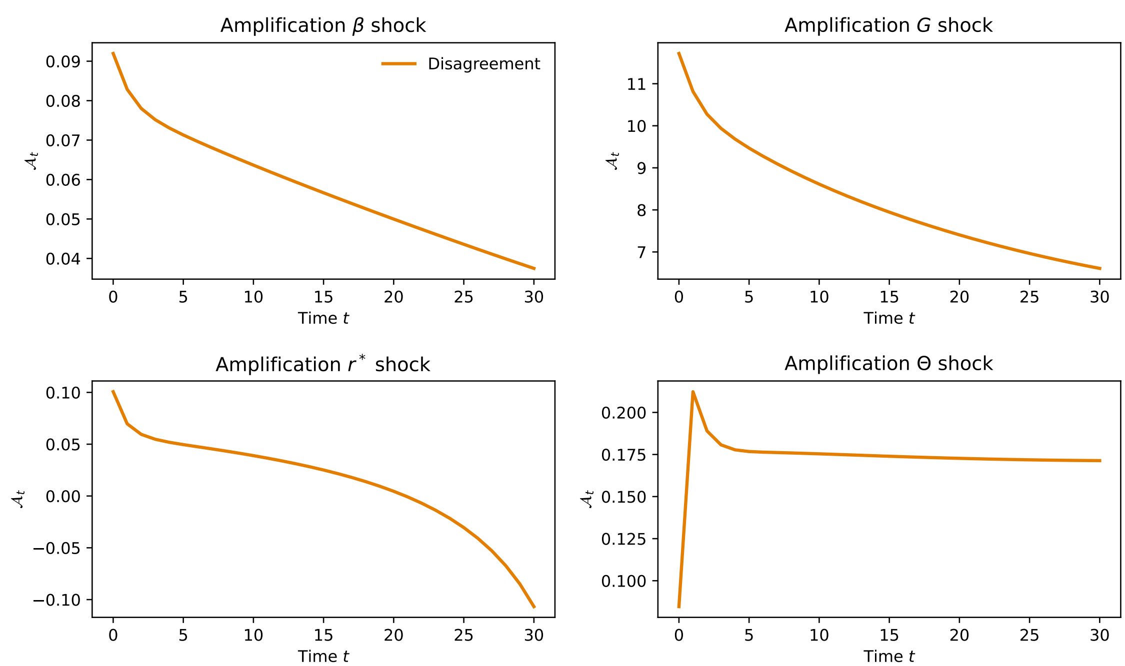

Figure 32: Business-cycle amplification: The role of persistence

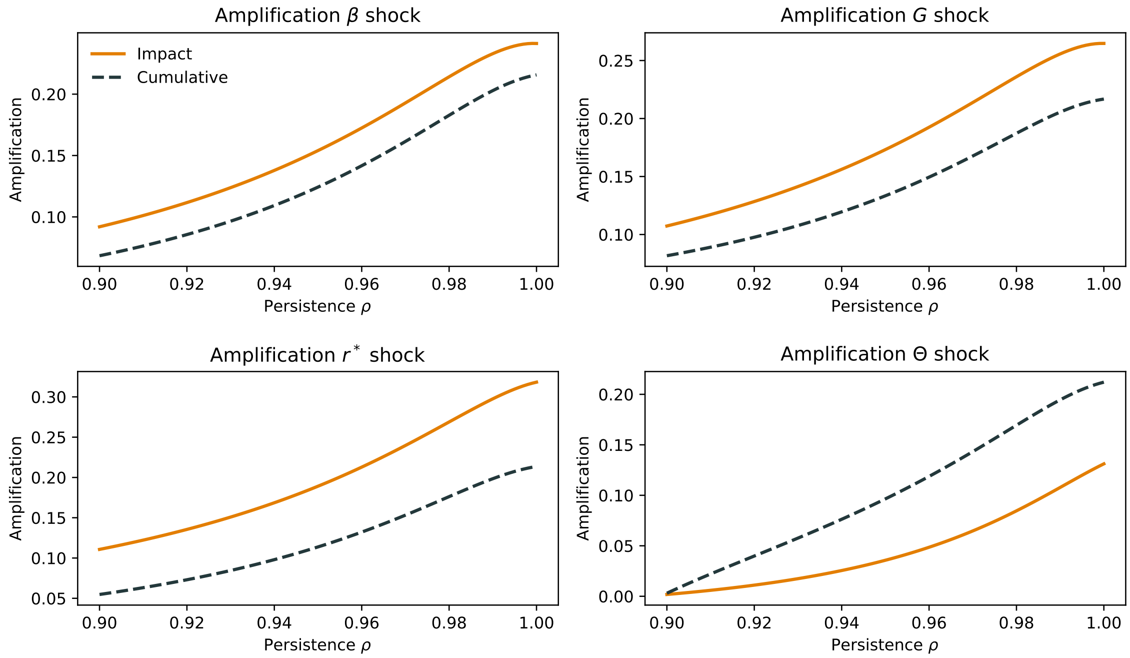

Figure 33: Business-cycle amplification: The role of monetary policy

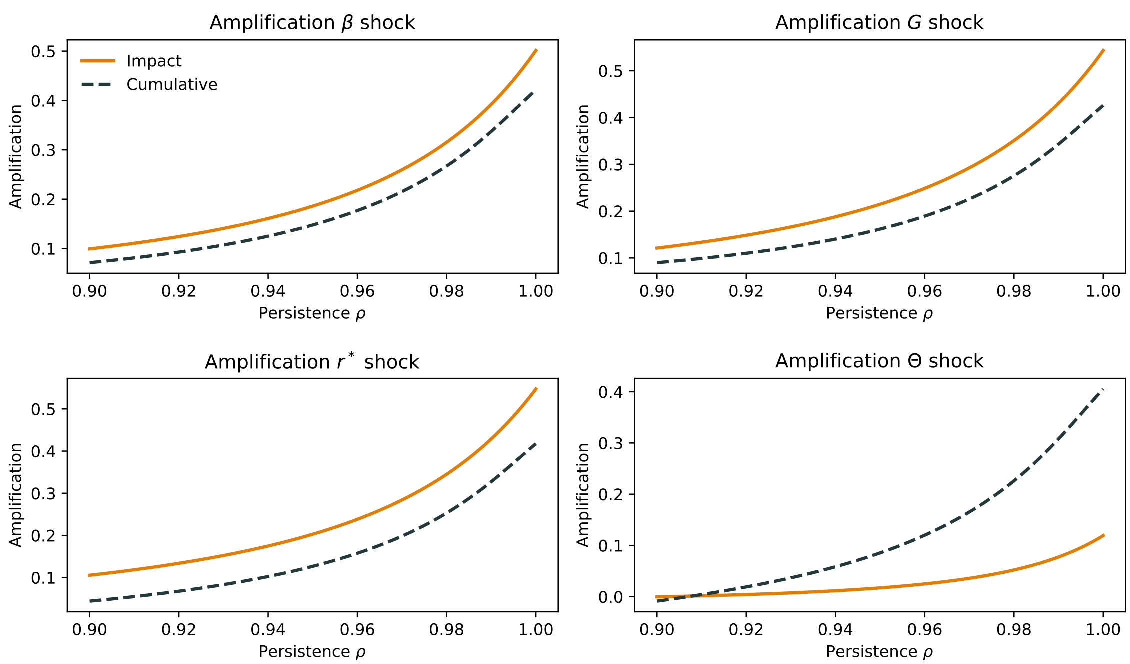

Figure 34: Targeted spending multipliers

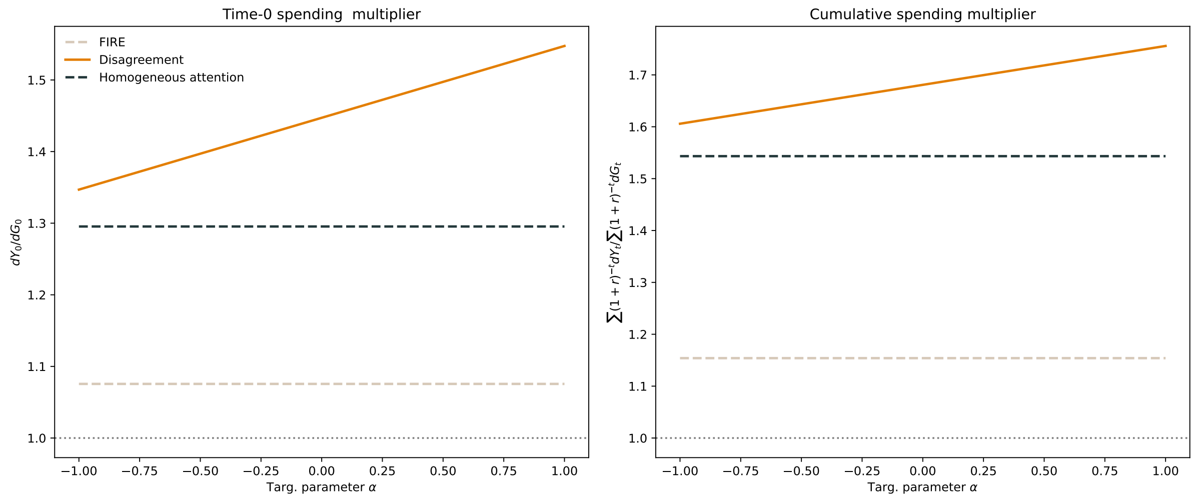

Figure 35: Targeted spending multipliers: The role of monetary policy

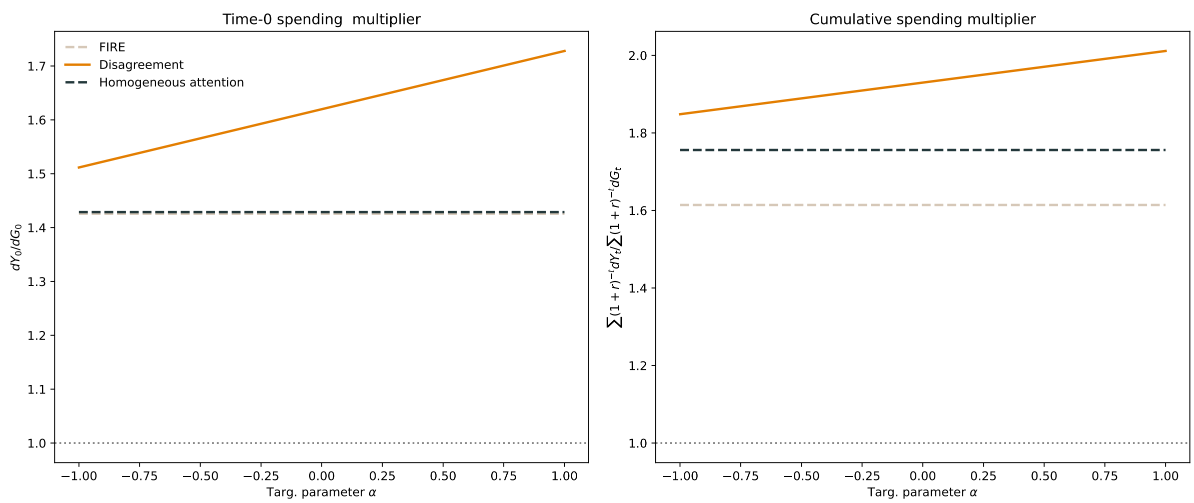

E.4 Progressive taxation

The baseline quantitative model assumes that labor income taxation is constant. In this appendix, I allow taxes to be progressive by assuming the constant-elasticity retention function in Heathcote, Storesletten, and Violante (2017). Formally, I assume that if an individual's pretax labor income is given by $y_{i,t}$ then their after tax labor income is given by $\tau_t y_{i,t}^{1-p}$, where τ_t and p control the average level of taxation and progressivity, respectively. The modified government budget constraint is given by

$$G_t + (1+r_t)B_t = \int_0^1 \left(y_{i,t} - \tau_t y_{i,t}^{1-p} \right) di + B_{t+1}. \tag{124}$$

Following Heathcote et al. (2017), I calibrate progressivity to $p = 0.181$. In my quantitative exercises, I fix progressivity and let the level of taxation τ_t vary so as to clear the government budget constraint.

Figure 36: Optimal attention

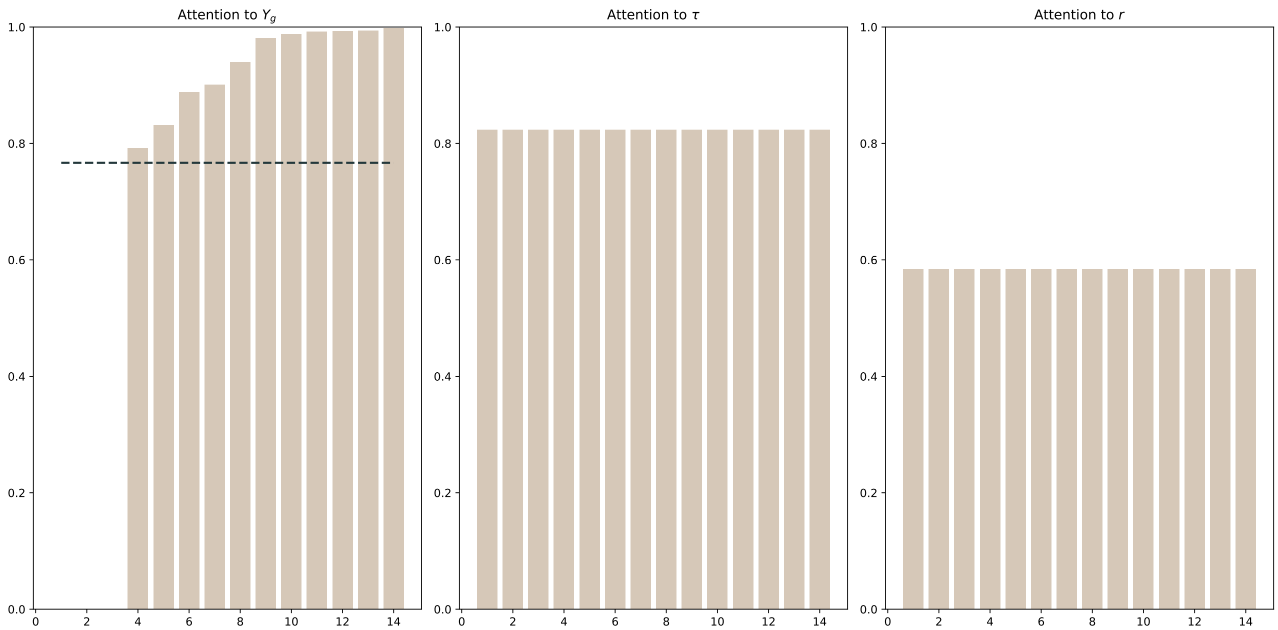

Figure 37: Consumption response

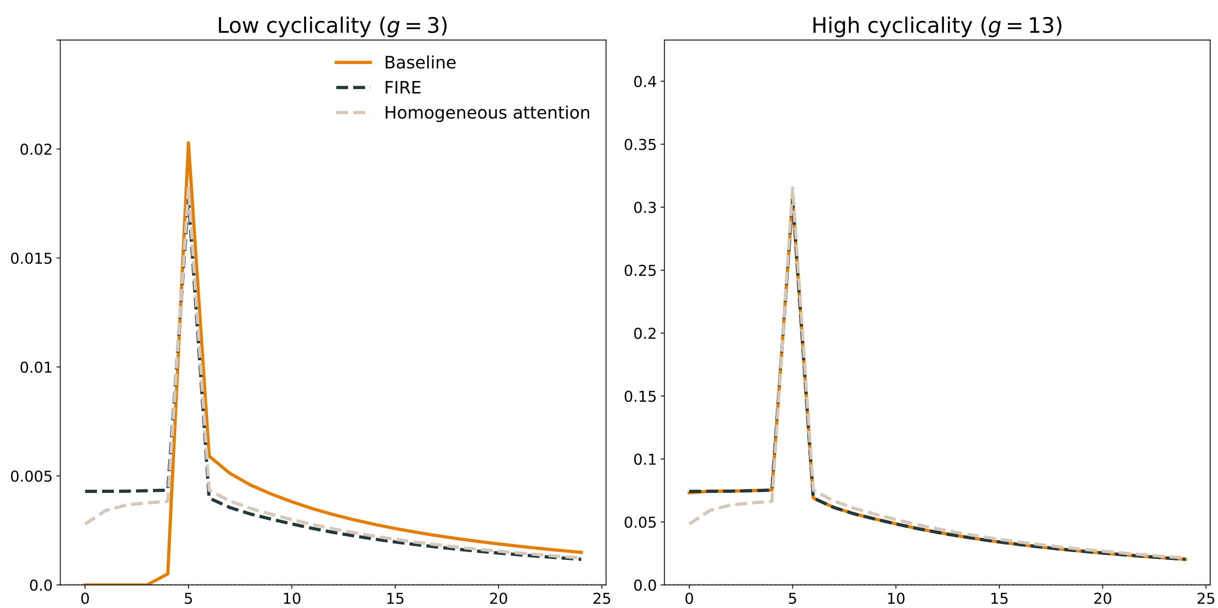

Figure 38: Consumption response for all groups

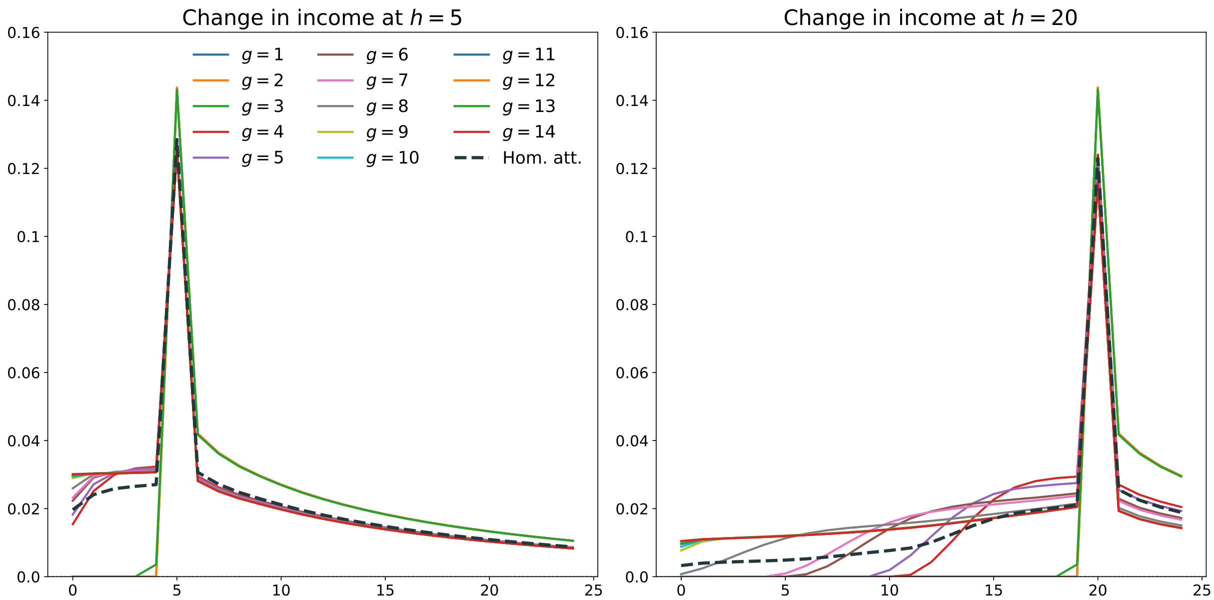

Figure 39: Business-cycle amplification

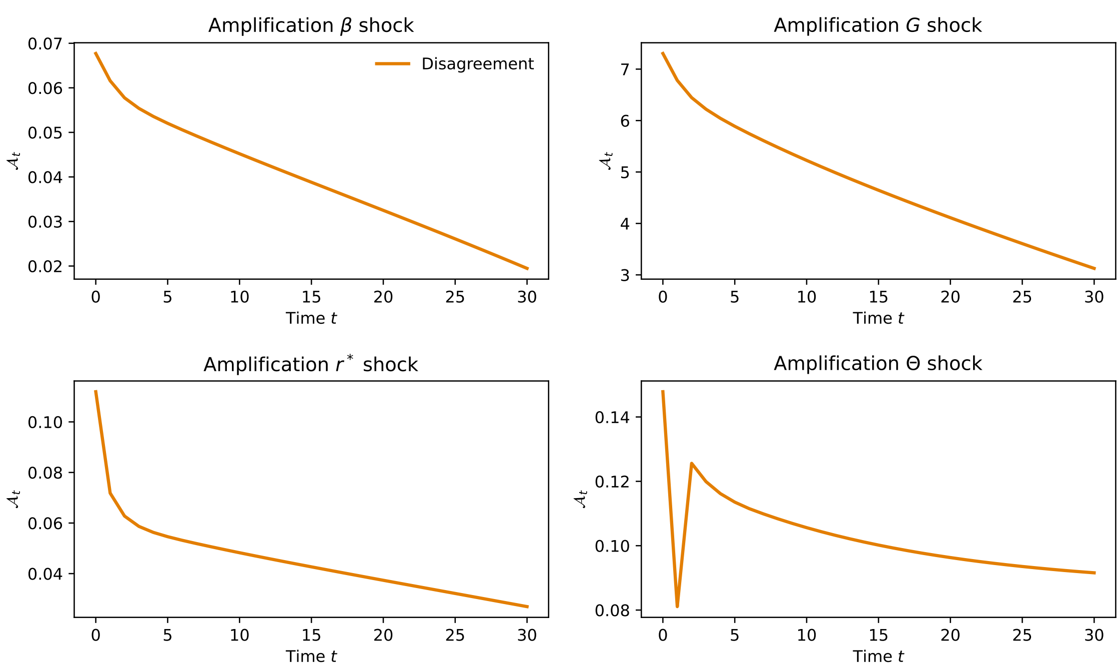

Figure 40: Business-cycle amplification: The role of persistence

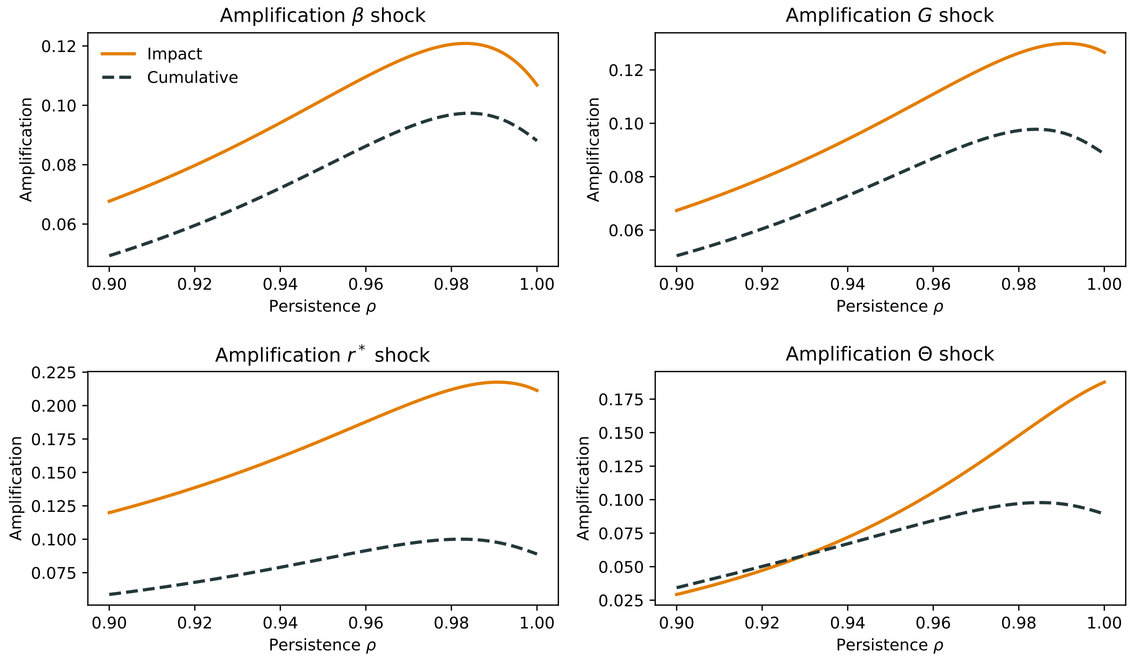

Figure 41: Business-cycle amplification: The role of monetary policy

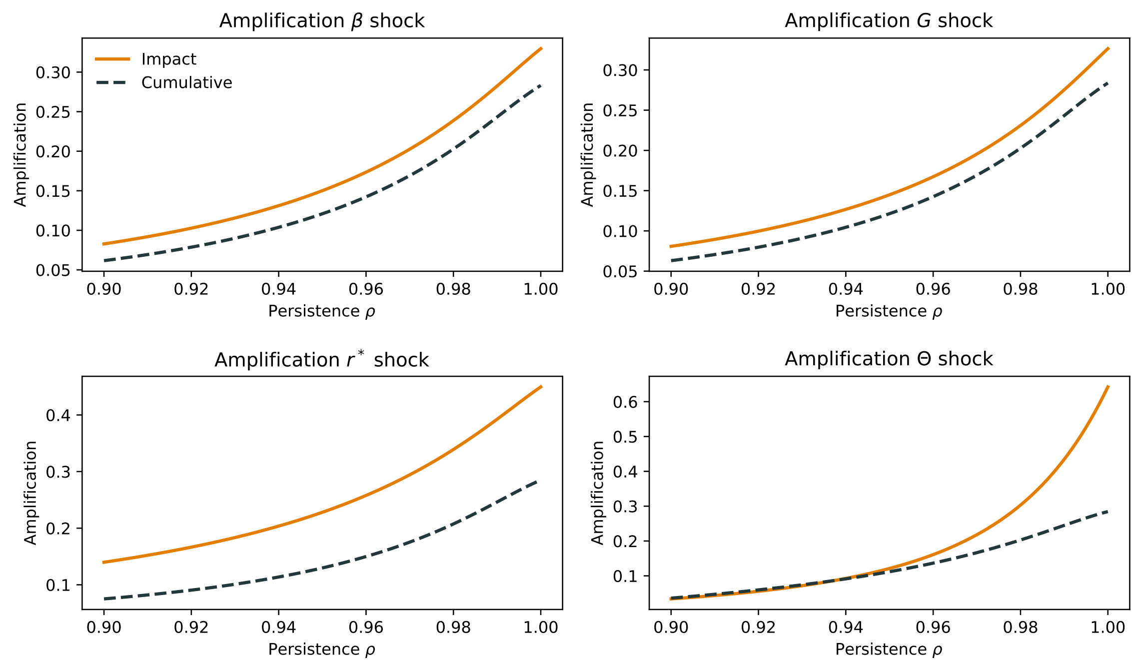

Figure 42: Targeted spending multipliers

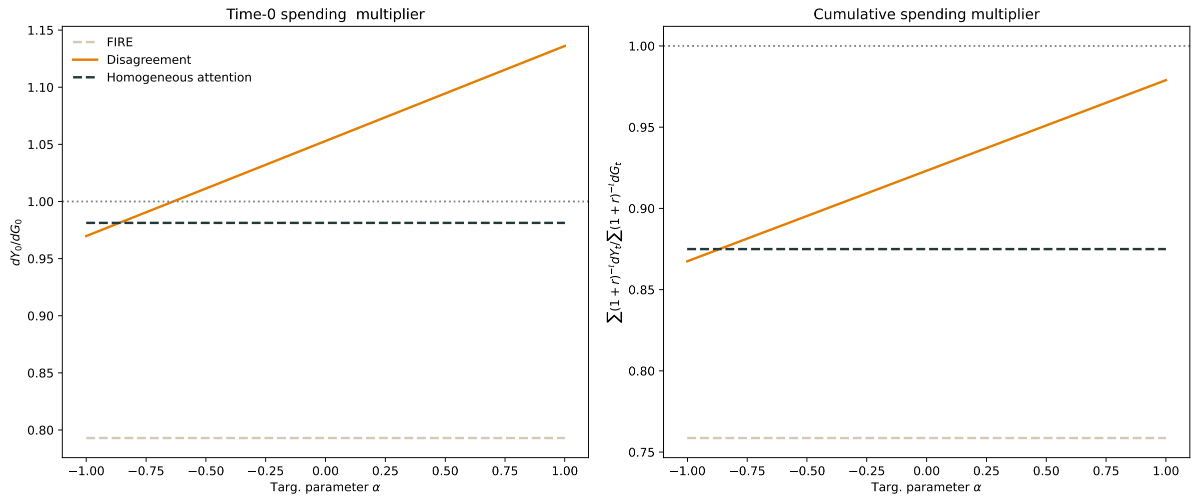

Figure 43: Targeted spending multipliers: The role of monetary policy

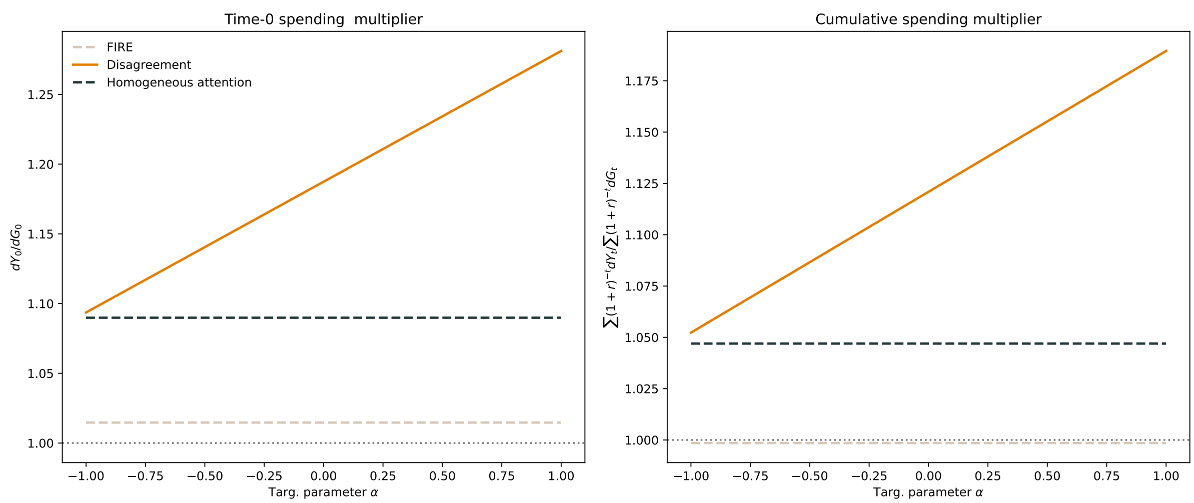

E.5 Progressive taxation

In the baseline quantitative model, I assume that income risk is countercyclical by calibrating $\zeta = -0.5$. In order to assess the robustness of the results in this paper to this assumption, this appendix assumes that $\zeta = 0$ and recomputes the main quantitative results.

Figure 44: Optimal attention

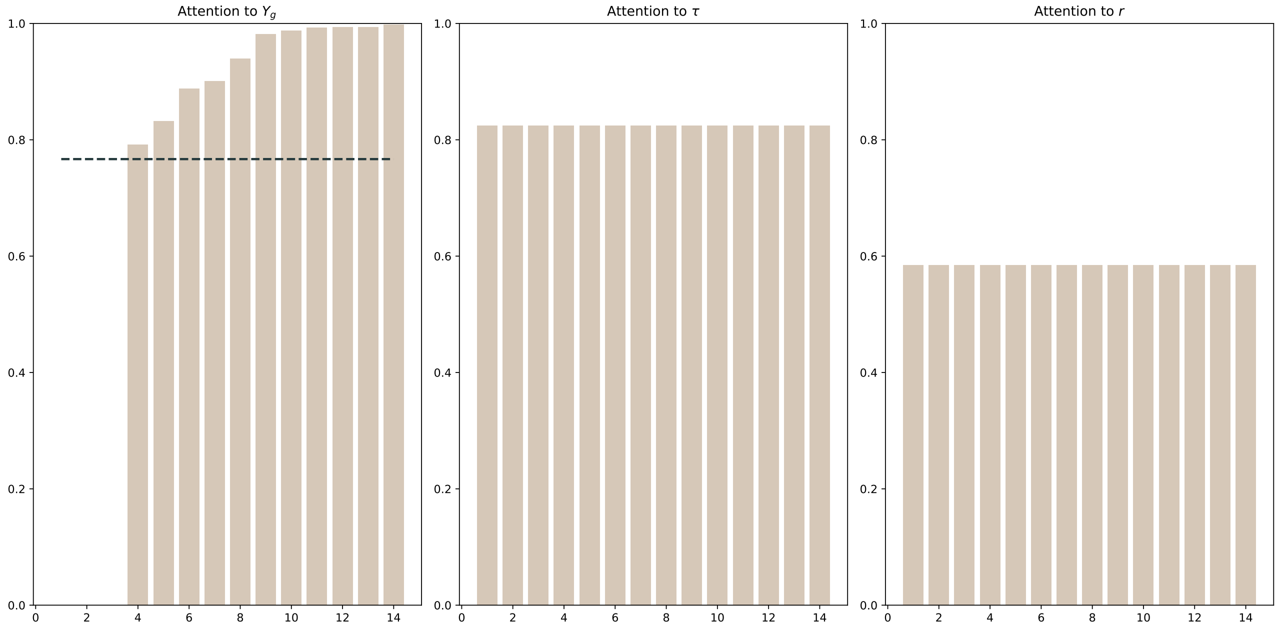

Figure 45: Consumption response

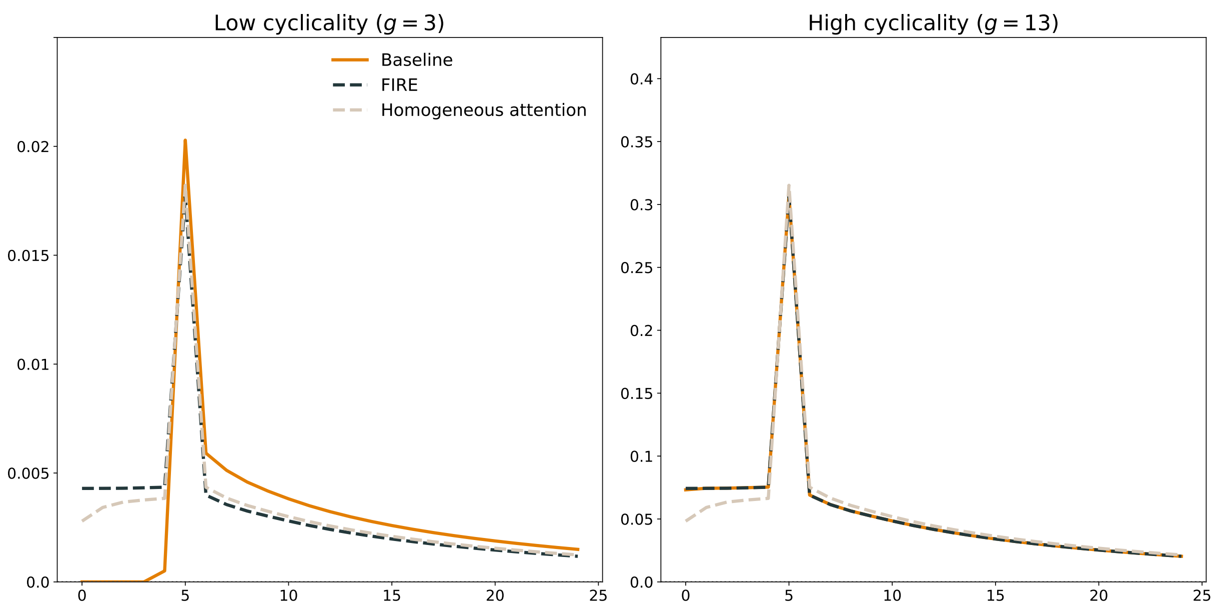

Figure 46: Consumption response for all groups

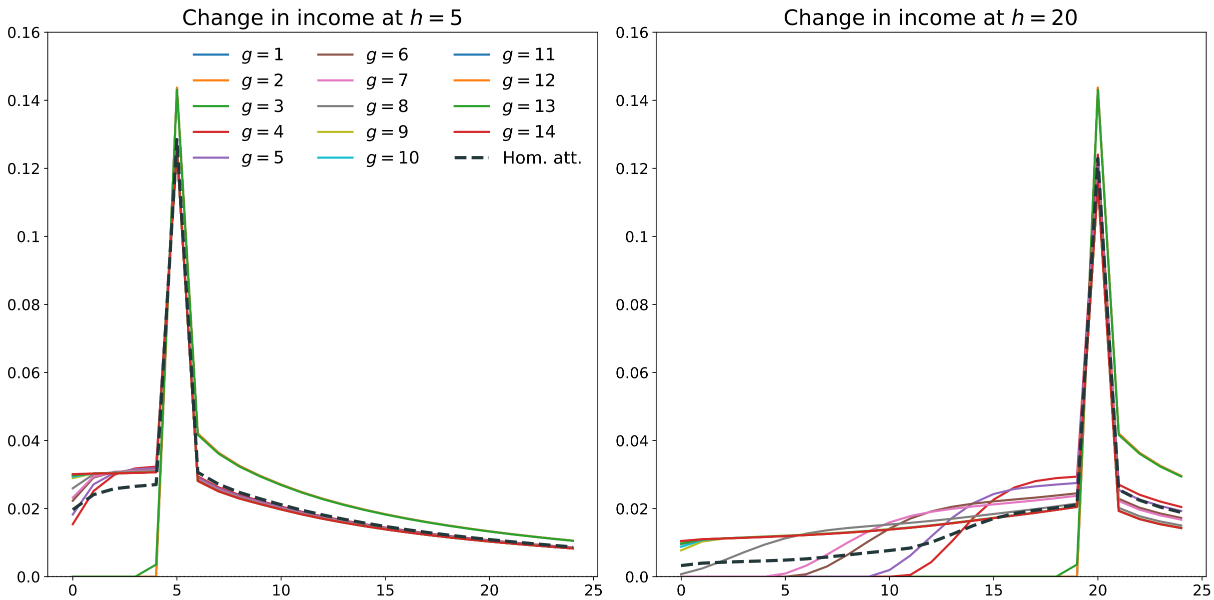

Figure 47: Business-cycle amplification

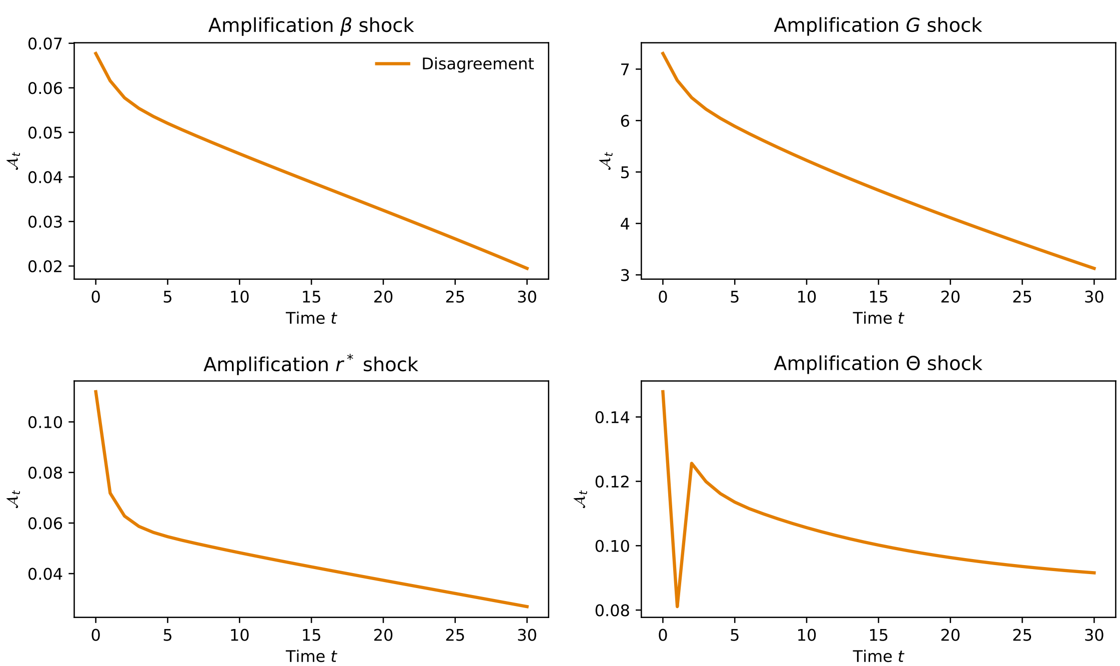

Figure 48: Business-cycle amplification: The role of persistence

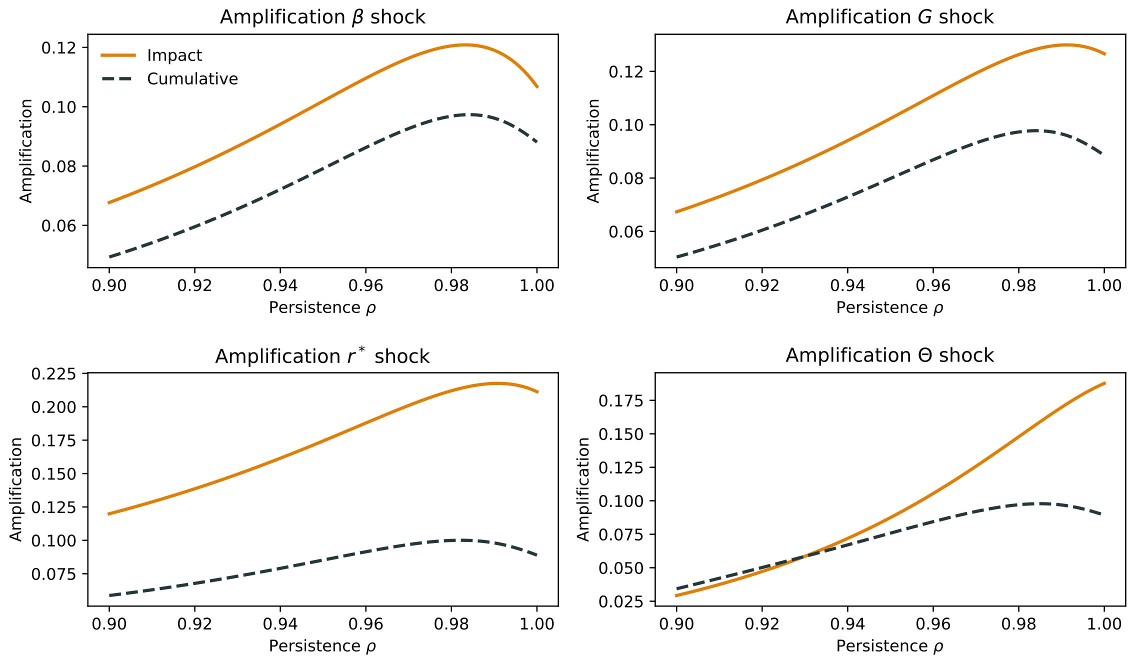

Figure 49: Business-cycle amplification: The role of monetary policy

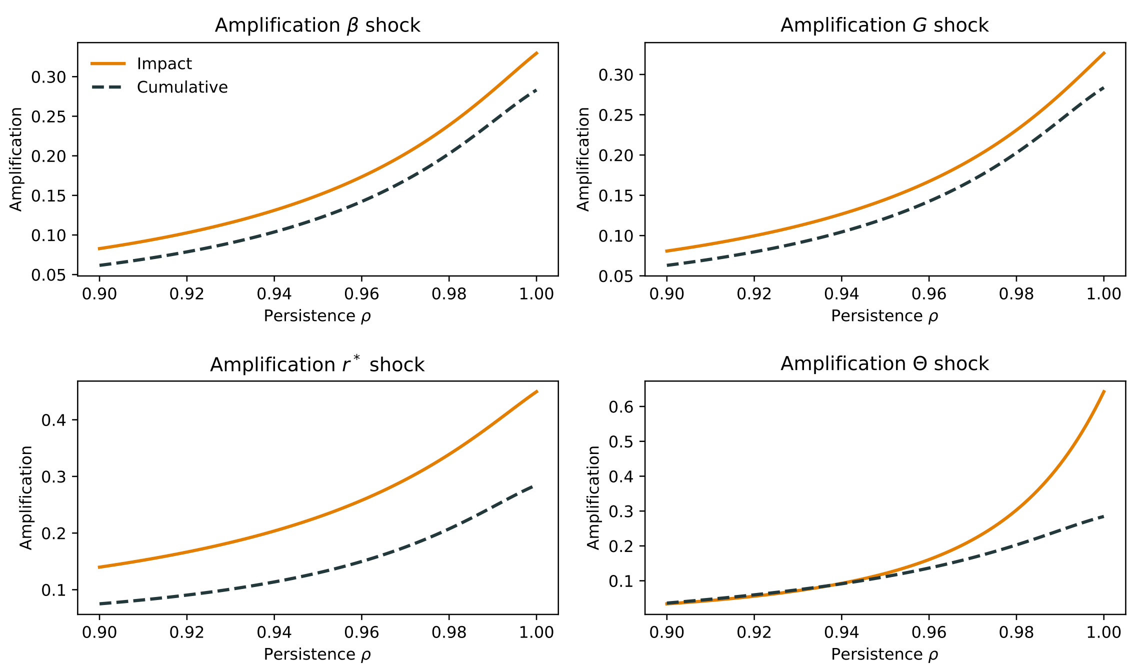

Figure 50: Targeted spending multipliers

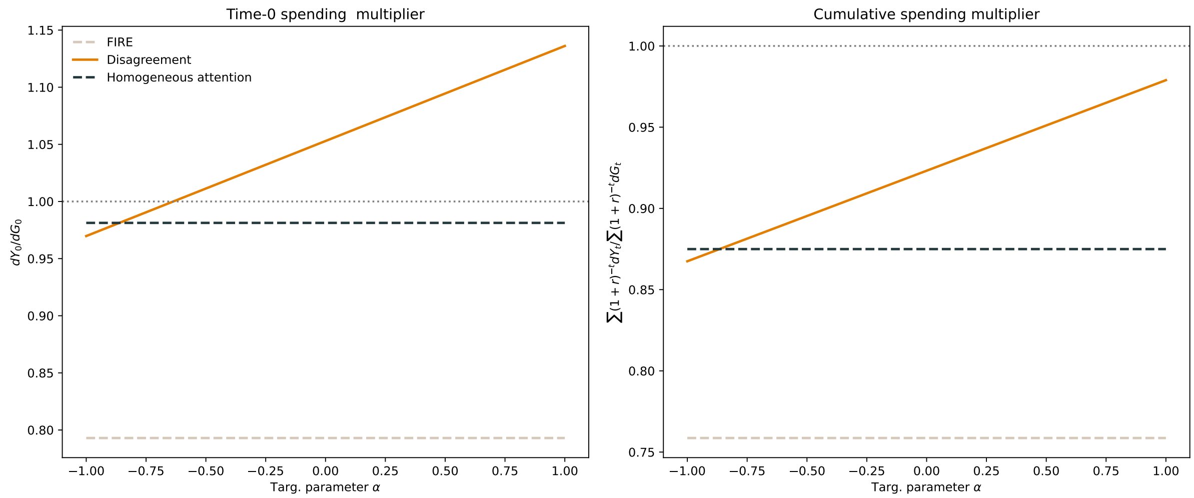

Figure 51: Targeted spending multipliers: The role of monetary policy