Do Households Substitute Intertemporally? 10 Structural Shocks That Suggest Not

Preliminary Draft[*]

Edmund Crawley

February 2024

Abstract

I combine microdata on the intertemporal marginal propensity to consume (iMPC) with 10 different structural macro shocks to identify the role of intertemporal substitution in consumption behavior. Although some of the shocks lead to large and persistent changes in real interest rates—which in many models would induce a large intertemporal substitution effect—I find no evidence that households shift the timing of their consumption in response to these interest rate changes. Indeed, the iMPC explains almost all the aggregate consumption response, leaving no role for intertemporal substitution.

JEL: E21, E32, E52
Keywords: HANK, Monetary Policy, Consumption

---

[*] Viewpoints and conclusions stated in this paper are the responsibility of the author alone and do not necessarily reflect the viewpoints of the Federal Reserve Board.

Crawley: Federal Reserve Board, edmund.s.crawley@frb.gov.

1 Introduction

In most macroeconomic models, the real interest rate is a key input into households' consumption decisions. However, there is little empirical evidence—micro or macro—that households behave as the models suggest with respect to real interest rates. This lack of evidence partly stems from the fact that real interest rates co-move with other features of the economy that go into the household consumption decision; therefore, identifying the role played by real interest rates is challenging.

The key insight of this paper is recognizing that the effect of real interest rates on consumption can be calculated as a residual after accounting for the other inputs to the consumption—saving decision. Furthermore, in light of recent advances, economists now have good empirical evidence on how households respond to changes in the other main drivers of consumption, namely the expected paths for labor income, stock prices, and real estate prices. As a result, it is possible to infer the role of interest rates following ten different structural shocks chosen from Ramey (2016). I find that 8 of the 10 structural shocks suggest no role for intertemporal substitution in household decision making, while the other 2 point to a limited role, but these estimates are not statistically significant.

This paper is closely related to Auclert et al. (2020), who similarly find a small role for intertemporal substitution. Relative to Auclert et al. (2020), this paper contributes to the literature by narrowing the scope of interest to focus only on the consumption "block" of the model rather than attempting to estimate a full general equilibrium model. This narrower scope offers several advantages. First, it allows me to be agnostic about the rest of the model. My findings are consistent with models that include an array of financial frictions, different mechanisms for sticky prices, international trade, and time-to-build and other investment frictions, as well as many possibly yet-to-be-discovered modeling bells and whistles. Second, the focus on the consumption block of a larger model allows me to use a much wider range of structural shocks to estimate household behavior than in Auclert et al. (2020) because I do not need to model each structural shock explicitly. For a structural shock to be used for estimation in my framework, I only need to impose the condition that the way in which the shock affects household consumption decisions is mediated through the expected paths of aggregate labor income, the real fed funds rate, and asset prices. As a simple example, a productivity or monetary policy shock in the standard three-equation New Keynesian model fits into this paradigm. But more complex structural shocks from the empirical literature, such as investment-specific technology news shocks, can be used without making specific modeling choices about the exact nature of these shocks.

I impose some structure on the consumption block but allow for enough flexibility such that the estimated parameters span most empirically plausible household behavior. First, the input-output structure of the consumption block is fixed. I allow for inputs to be the expected paths for aggregate labor income, the real fed funds rate, the real return on the stock market, and the real return on real estate. The output must include aggregate consumption but can also include aggregate savings in each asset. With this input-output structure, the dynamics of the linearized version of the consumption block in sequence space are fully described by four Jacobians: the marginal change in aggregate consumption at time s to a change in each of the four inputs at time t. Once these four Jacobians are pinned down, the consumption block can, if desired, be added to a larger linear model that shares the same input-output structure for the consumption block using the toolkit described in Auclert et al. (2021).

Each of the four Jacobians that describe the dynamics of the consumption block is of infinite dimension. In order to reduce the dimensionality of the estimation problem, I use micro-empirical evidence along with theory. For the labor income Jacobian, I use evidence from Crawley and Kuchler (2023) to fit a one-asset heterogenous agent model to the marginal propensities to consume (MPC) for each quintile of the liquid wealth distribution. This resulting Jacobian is consistent with the intertemporal MPC (iMPC) implied by the Norwegian lottery data in Fagereng et al. (2021). Because the current micro-empirical literature lacks strong evidence on households' consumption responses to news shocks about future income, I allow for sticky expectations, following Carroll et al. (2020), and estimate this labor income sticky expectations parameter from the macro structural shock evidence. For the Jacobians for the stock market and real estate, I assume an annual MPC of 3 percent—the current consensus in the literature—and show that my results are robust to significant variation around that assumption.

The real fed funds rate Jacobian is of most interest in this paper as this determines how aggregate consumption reacts to changes in the fed funds rate independent of other changes to the economy. I start with the real interest rate Jacobian that comes from my one-asset toy model. Following the decomposition in Farhi et al. (2022), I then separate the intertemporal substitution effect Jacobian from the income effect Jacobian. This income effect Jacobian in the toy model does not do a good job of capturing the true income effects because the main assets households hold—namely stocks and real estate—are not well approximated by the short-term bonds of the model. As a result, I keep only the intertemporal substitution effect Jacobian and assume that the income effect is captured by the stock market and real estate Jacobians. The resulting intertemporal-substitution Jacobian shows that, in response to a future increase in the real interest rate, households reduce consumption up until the change, collect the higher interest rate, and then spend down their savings following the change. The magnitude of this effect may depend on households' elasticity of intertemporal substitution and the degree to which they pay attention to changes in the interest rate among many other possibly rational or behavioral factors. I again introduce sticky expectations, and this single parameter allows the intertemporal substitution Jacobian to span most of the empirically plausible intertemporal substitution behaviors—from nothing at all to full rational expectations with high elasticity of intertemporal substitution.

Overall, I impose enough structure from the micro-empirical evidence combined with my toy model that I am left with just two parameters—sticky expectations for labor income and for the fed funds rate—to estimate using the macro structural shocks. In order to make use of the information contained any particular structural shock series, I run Jorda projections on each of the four inputs to the consumption decision—the federal funds rate, labor income, stock and real estate returns. These Jorda projections give me what I will call the *empirical impulse response* function to the shock for every input to the consumption block. These impulse response functions tell me how much each of the four inputs to the consumption block is expected to deviate at every quarter s after the shock hits from its expected value before the arrival of the shock. If I know the consumption Jacobians with respect to each of these inputs, I can then calculate how much I expect consumption to deviate s periods after the shock hits—I will call this the *Jacobian-implied impulse response function* for consumption. This Jacobian-implied impulse response function for consumption is simply the sum of each of the four Jacobians multiplied by the empirical impulse response for their respective input. Finally, I can compare the Jacobian-implied impulse response for consumption with the empirical impulse response function for consumption calculated using Jorda projections using a distance metric. I choose the two parameters for the Jacobians in order to minimize the sum over all ten shocks of the distance between the Jacobian-implied impulse response for consumption and the empirical impulse response function for consumption. A robust finding is that the sticky expectations parameter implies the intertemporal substitution Jacobian is close to zero everywhere. That is, households do not substitute intertemporally.

Should we be surprised that households do not appear to substitute intertemporally? I would argue not. First, it is well known from the near rationality literature, starting with Akerlof and Yellen (1985), that the consumption—saving behavior that comes from standard models leans on weak incentives, and, as a result, is not robust to either small deviations in rationality or other frictions. For example, Cochrane (1989) finds that a consumer who bases their consumption—saving decision on the 10-year moving average of the real interest rate in place of the forward short-term real rate will suffer a utility loss equivalent to between \$0.08 and \$1.45 per quarter—small enough that we might expect a rational consumer to not pay attention at all. He suggests the result "implies that the theory as it stands provides few predictions about the relationship between aggregate consumption and asset price or aggregate quantity fluctuations that are robust to \$1 "mistakes" or misspecifications." Furthermore, there is a wealth of empirical evidence on consumption behavior that suggests households do not behave according to standard models in other respects. More anecdotally, Choi (2022) compares the advice of popular financial advice books with the advice that comes from models in the household finance literature. He finds many ways in which popular advice on the consumption—saving decision differs from that of academic models. He also finds a complete lack of advice on how households should change consumption—saving behavior in response to interest rates, strongly suggesting that this is not a question many households even ask. To the extent that there is any link between interest rates and consumption choices in Choi's review of the popular finance literature, it is through the advice to save enough to continue to spend a fixed fraction of your income in retirement, advice that would suggest a negative intertemporal elasticity of substitution.

The second reason it may not be surprising that households do not intertemporally substitute is that little evidence of such behavior has been found in previous studies. Hall (1988) provides a seminal contribution showing no relationship between interest rates and consumption growth in aggregate data. He states, "A detailed study of data for the twentieth-century United States shows no strong evidence that the elasticity of intertemporal substitution is positive," and, furthermore, he states of his estimates "most of them are also quite precise, supporting the strong conclusion that the elasticity is unlikely to be much above 0.1, and may well be zero." My paper can be thought of as overcoming some of the identification problems inherent in Hall (1988). A small cottage industry of papers followed in an attempt to use the Euler equation to find micro evidence on the size of intertemporal substitution—many such papers are reviewed in Browning and Lusardi (1996). These papers contained a wide array of estimates, with little in the way of consensus. Ultimately, Carroll (1997) showed that the methods used in these papers were flawed. More recently, Best et al. (2019) makes use of the unusual notching characteristics of the U.K. mortgage market to gain identification of the elasticity of intertemporal substitution under some reasonable assumptions and finds, like this paper, that it is close to zero. By contrast, Crump et al. (2022) use idiosyncratic inflation expectations from the New York Survey of expectations and find a substantially positive elasticity of intertemporal substitution.

This paper is also closely related to McKay and Wolf (Forthcoming), who show how to use the sequence space Jacobians, along with impulse response functions (IRFs) identified from several different types of monetary policy shocks, to estimate economic dynamics that are robust to the Lucas critique. Hebden and Winkler (2021), Beraja (2023), and Barnichon and Mesters (2023) all have recent theoretical papers with similar ideas.

The paper is structured as follows. The pedagogical material in section 2 explains the key ideas of the paper in the context of a simple endowment economy with two agents. Section 3 details the exact methodology that I use to estimate the intertemporal substitution Jacobian. Section 4 presents the results, section 5 examines the robustness of the results to some underlying assumptions, and section 6 concludes.

2 Example: A Two-Agent Endowment Economy with $C = Y$

It is useful to see how the methodology works in a simple two-agent endowment economy model in which consumption is equal to income. In this economy, all households receive an exogenous stream of income, Y_t. A fraction of households, λ, are hand to mouth and spend all their income each period: $C_{htm,t} = Y_t$. The remaining fraction, $1-\lambda$, optimize their consumption according to the consumption problem:

$$\underset{\{C_{opt,t}\}}{\text{maximize}} \quad \mathbb{E}_0 \sum_{t=0}^{\infty} \left(\Pi_{i=0}^{t} \beta_i\right) \frac{C_{opt,t}^{1-\sigma}}{1-\sigma} \tag{1}$$

$$\text{subject to:}$$

$$C_{opt,t} + A_t \leq (1+r_t)A_{t-1} + Y_t. \tag{2}$$

Here, β_i is the exogenously given discount factor for each period i, while the real interest rate, r_t, is taken as given by the optimizing household but will be endogenously determined in equilibrium to clear the market such that $C_t = \lambda C_{htm,t} + (1-\lambda)C_{opt,t} = Y_t$. A directed acyclical graph representation of the model is shown in figure 1, in which the two

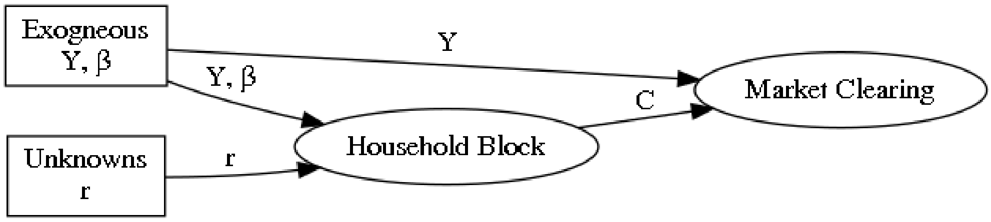

Figure 1: Directed acyclical graph for a simple endowment economy

agents have been aggregated to a household block with input paths for Y, β, and r and an output path for C. Assuming $A_{-1} = 0$, the household block is described by the following function:

$$\{C\}_{t=0}^{\infty} = \mathcal{C}(\{Y\}_{t=0}^{\infty}, \{r\}_{t=0}^{\infty}, \{\beta\}_{t=0}^{\infty}). \tag{3}$$

I will linearize this model around steady-state values. Consequently, the dynamics of the household block are fully described by three Jacobians:

$$\mathcal{J}_{t,s}^{C,Y} = \frac{dC_t}{dY_s}, \qquad \mathcal{J}_{t,s}^{C,r} = \frac{dC_t}{dr_s}, \qquad \mathcal{J}_{t,s}^{C,\beta} = \frac{dC_t}{d\beta_s}.$$

Economists are particularly interested in the first two Jacobians: how consumption changes with income and with interest rates. In the case of our two-agent model, these Jacobians are graphically represented in figure 2. The left-hand panel of 2 shows the impulse response of consumption over time to an expected one unit increase in income in zero, four, and eight quarters' time, holding interest rates and discount factors constant. This Jacobian is called the intertemporal marginal propensity to consume (iMPC) in the literature. In this example, I have set the hand-to-mouth share to be 0.25 and this shows through in the large spikes in the consumption response at the time the income is received. The optimizing agents increase their consumption in every period by $r_{ss}/(1+r_{ss})$ multiplied by the present value of the expected increase in income—this increase can only just be made out in the figure. The right-hand panel of figure 2 shows the impulse response of consumption over time to an expected one unit increase in the realized interest rate in 0, 4, and 8 quarters' time, holding income and discount rates constant. While hand-to-mouth households do not change their consumption in response to a change in interest rates, optimizing agents choose to save up to the quarter of the higher interest rate and then—after they have received the higher interest rate on their savings—consume thereafter at a slightly higher level than before.

2.1 Two-steps to find the consumption—interest rate Jacobian

The thought experiment I want to entertain is one in which an economist is able to directly measure the consumption—income Jacobian, say by running experiments in which households are given income s periods in the future but cannot run a similar randomized experiment for interest rates. Nevertheless, the economist has aggregate data on income and interest rates (and consumption, though this is equal to income by construction) and wishes to estimate the consumption—interest rate Jacobian. My

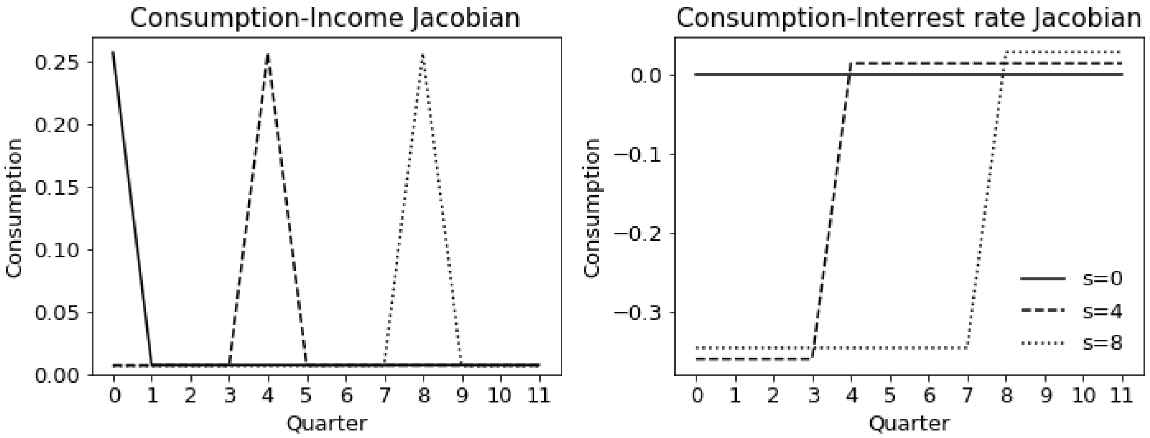

Figure 2: Consumption Jacobians for the simple two-agent model

approach is to follow two steps.

Step 1: Identify a shock in which β does not move. In many macroeconomic models, including the simple one presented here, households may change their consumption decision for reasons unrelated to income or interest rates. In this case, I have modeled this possibility as a shock to the discount factor. The existence of such shocks makes it difficult to measure the Jacobian of the interest rate because in equilibrium the interest rate will move to offset changes to the discount factor. In the absence of shocks to income, these discount factor shocks would result in a constant level of consumption—since $C = Y$—despite a changing interest rate. A naive economist might be led to the conclusion that consumption is insensitive to changes in the interest rate.

My strategy to overcome this problem, in the context of this example, is to identify shocks that are orthogonal to changes in household preferences and therefore have an effect on consumption only through changes to the path for expected income and interest rates. In this simple model with only two exogenous variables, these identified shocks are exogenous to the path of income. Such shocks to income can, however, vary in timing and persistence.

In figure 3, I show the impulse responses for income, real interest rates, and consumption to an exogenous shock to income that takes the form of an AR(1) decaying at a rate of 0.5. Because these IRFs are a result of general equilibrium, the IRF for consumption is equal to that for income. However, from the point of view of individual households, their consumption choice is a function of the IRFs for both income and real interest rates, which they take as given.

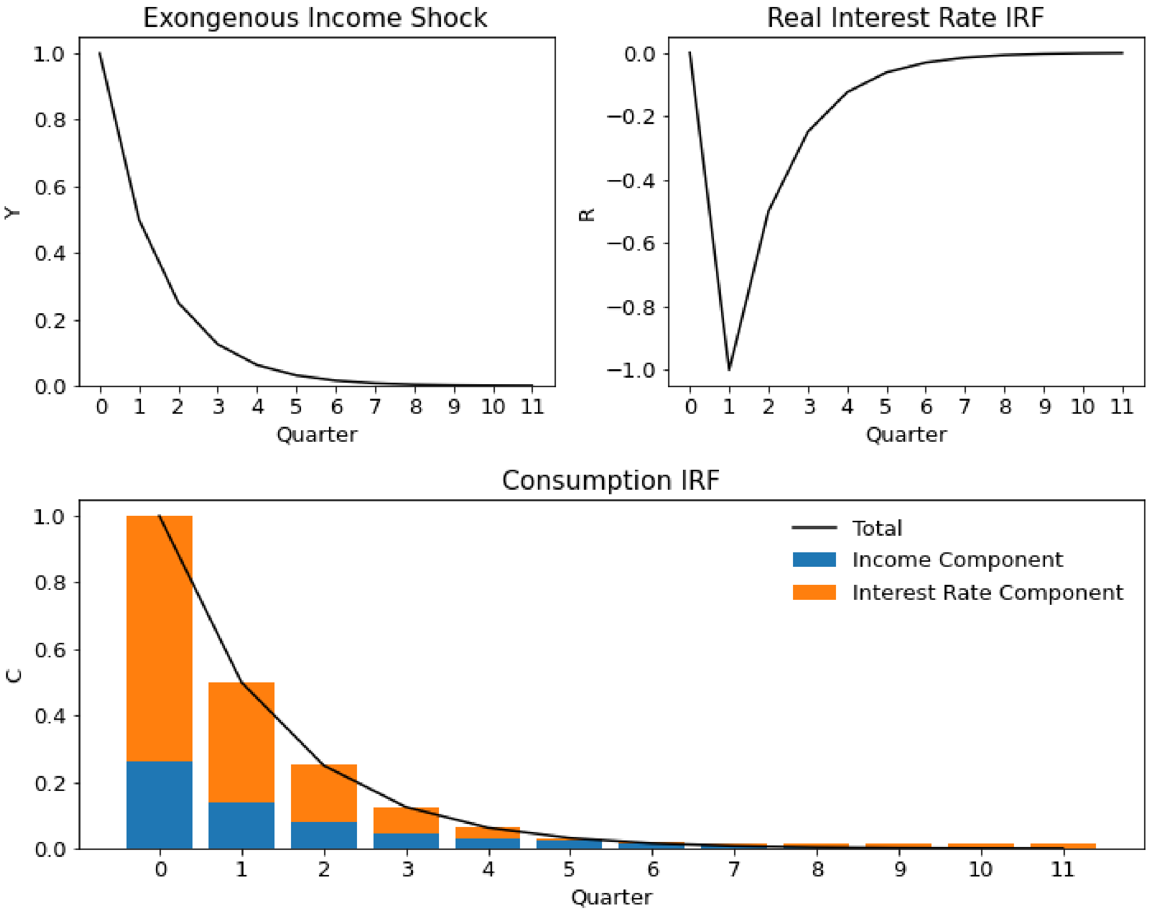

Figure 3: Impulse response to an exogneous income shock

Step 2: Find the consumption—interest rate Jacobian that explains the residual consumption IRF. The consumption IRF in figure 3 can be further decomposed into two components: the household response to the expected income path and the household response to the expected real interest rate path.

$$dC_t = \underbrace{\sum_{s=0}^{\infty} \mathcal{J}^{C,Y}_{t,s} dY_s}_{\text{Income component}} + \underbrace{\sum_{s=0}^{\infty} \mathcal{J}^{C,r}_{t,s} dr_s}_{\text{Interest rate component}} \tag{4}$$

In my setup, I can observe the following IRFs: dC_t, dY_t, and dr_t. Furthermore, I have assumed that the consumption—income Jacobian, $\mathcal{J}^{C,Y}_{t,s}$, is observable, say from natural experiments or randomized control trials. Under these assumptions, I can calculate the interest rate component of the consumption IRF:

$$\text{Interest rate component} = \sum_{s=0}^{\infty} \mathcal{J}^{C,r}_{t,s} dr_s = dC_t - \sum_{s=0}^{\infty} \mathcal{J}^{C,Y}_{t,s} dY_s. \tag{5}$$

This calculation allows me to identify the partial equilibrium consumption response to a particular path for the real interest rate. However, there are many consumption—interest rate Jacobians that are consistent with any one particular path. In order to identify the entire consumption—interest rate Jacobian, it will be necessary to parameterize the Jacobian with a finite set of parameters. I can then find the set of parameters that best fits the IRF, or IFRs, that I have observed. This set of parameters allows me to calculate the entire consumption—interest rate Jacobian.

3 Methodology

The core identification ideas are outlined in the specific example of a simple two-agent model in section 2. In this section, I build on those ideas and show how the methodology can be applied to much richer models that include rich household heterogeneity, firms that make investment decisions, sticky prices and wages, multiple assets, international trade, and financial frictions. I will also show some of the limitations of the methodology, both in theory and in practice.

After I have set up the theory, I show how I choose to parameterize the consumption Jacobians of interest. I then explain my methodology for estimating the impulse responses to structural shocks. Finally, I detail the method I use to estimate the Jacobian parameters and, hence, the full consumption Jacobians.

3.1 The Consumption Choice Embedded in a Rich Model

The household block structure—inputs Y, r, and β and output C—shown in figure 1 appears as a sub-block of many standard and not-so-standard DSGE models of the economy. For example, this block is embedded in the textbook three-equation model from chapter 3 of Galí (2015). Often in such models, the household block is described as both a consumption and labor decision, with inputs, including the wage in place of income, and outputs, including hours worked. Figure 4 shows how, in many models, this household block can be further separated into a labor decision and a consumption decision (the figure also allows for three asset types—liquid assets, stocks, and real estate—and includes savings in these three assets as outputs of the consumption—saving block). The ability to separate

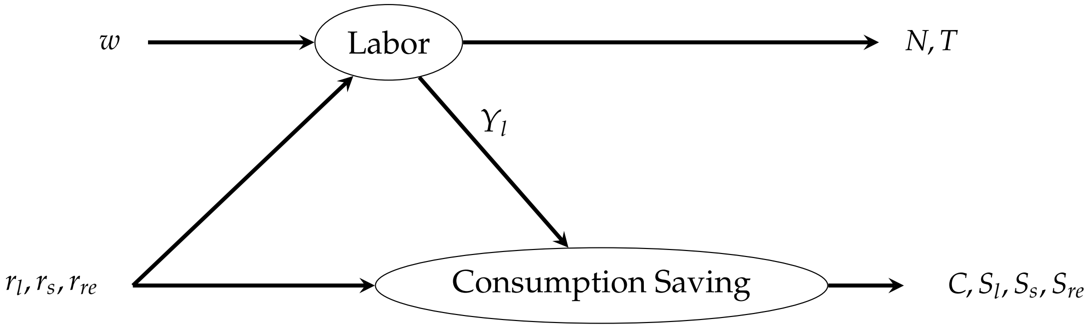

Figure 4: The household block separated into labor and consumption blocks

the labor choice from the consumption—saving decision through aggregate labor income is my key assumption:

Key assumption: The impulse response function for aggregate consumption is a function of the expected paths for aggregate labor income, the returns on available assets, and other inputs that are independent of the aggregate shocks analyzed.

Figure 4 is useful in demonstrating when this key assumption will not be valid. For example, if leisure and consumption are not separable in the utility function of the household, then it is not possible to separate the labor and consumption decisions in this way. Furthermore, in a two-agent or heterogenous agent model in which households make their own labor decision, liquidity-constrained households may change their labor decision in response to wage changes but not to changes in the interest rate on liquid assets. Therefore, the distribution of changes in income can depend on its origin—aggregate wage changes or interest rate changes—and, as a consequence, the consumption block is not separable from the labor block through aggregate labor income alone and, instead, the whole distribution of income is needed to separate the labor and consumption blocks. My methodology cannot strictly apply to such models. However, the method can be applied to model in which household labor is allocated by a union and the distribution of this allocation is independent of the origin of the change in aggregate hours worked, such as is the case in many heterogenous agent models with sticky wages. Furthermore, empirical evidence does not point to clear differences in how aggregate income fluctuations are distributed except in the case of progressive or regressive changes in taxes.[1]

As long as the key assumption above is satisfied, the consumption—saving block can be envisaged as part of a far more complex model and the methodology for identifying households' responses to interest rate shocks will remain valid. For example, such a block

---

[1] For a detailed analysis of how labor income varies heterogenously over the business cycle, seePatterson (2023).

can form part of a model with many of the bells and whistles that researchers have added to both the real business cycle and New Keynesian models: investment, firm financial frictions, heterogenous firms, international trade, sticky prices or menu costs, and labor market frictions. A key advantage of the methodology presented here is that it is not necessary to specify the rest of the model—even though a shock may affect every other block in the model, it is enough to know the inputs and outputs of the consumption block to estimate household consumption behavior.

3.2 Parameterizing the Consumption Jacobians

Under the key assumption from section 3.1 and assuming a linearized model, the dynamics of aggregate consumption are determined by the Jacobian of aggregate consumption to aggregate labor income, $\mathcal{J}^{C,Y_l}$, as well as the Jacobians of aggregate consumption to the real asset returns: $\mathcal{J}^{C,r_l}$, $\mathcal{J}^{C,r_s}$, and $\mathcal{J}^{C,r_{re}}$. Because each of these objects has an infinite dimension, it is necessary to discipline these Jacobians to reduce the number of parameters required for estimation. In this section, I will show how I use a mixture of theory and microdata to leave just two parameters to be estimated with macrodata.

The Consumption-to-Labor-Income Jacobian The consumption-to-labor-income Jacobian is also known as the interterporal marginal propensity to consume, or iMPC. It measures the amount by which aggregate consumption t periods from now increases when a news shock arrives that indicates that aggregate labor income will increase by one dollar s periods from now. The first column—$s = 0$—is the impulse response of aggregate consumption to an instantaneous increase in labor income and there is a fair amount of empirical evidence that I will use to derive this first column. There is much less empirical evidence on consumption responses to news shocks, and I will use a mixture of theory and empirics to reduce the number of parameters for this Jacobian to just one sticky expectations parameter.

In order to parameterize the consumption-to-labor-income Jacobian, I will make use of a one-asset heterogeneous agent model with sticky expectations. The resulting one-parameter Jacobian fits the empirical evidence available for the first column of the Jacobian while allowing for a variety of possibilities for news shocks. The model itself it not meant to be taken too seriously but instead should be thought of as a way to cover the space of feasible Jacobians given the empirical and theoretical evidence available.

In the one-asset model, infinitely lived households maximize their utility in the face of idiosyncratic income uncertainty, along with an inability to borrow. The value function is a function of the current wage state, e, and stock of assets, a, of the household. In order to match the first column of the Jacobian, I allow for six different types of household that differ only in their discount factor β_i. For each type, the household problem can be written in Bellman form as follows:

$$\begin{aligned} V_i(e,a) &= \max_{c,a'} u(c) + \beta_i \mathbb{E}[V_i(e',a')|e] \\ \text{s.t. } a' + c &= (1+R)a + y(e) \\ a' &\geq 0 \end{aligned}$$

The income process comes from Kaplan et al. (2018) along a progressive tax structure and is identical across each of the six household types. I choose power utility with a coefficient of relative risk aversion of 1 and a real interest rate of 3 percent. Note that with a different income process and parameters, the space spanned by the one-parameter Jacobian will be similar—that is, all my results are robust to the specific specifications here.

With this model, I then estimate the values of β_i's that best fit the one-year MPC for each of the five quintiles of the wealth distribution in the Danish data by using results from Crawley and Kuchler (2023). Specifically, I fit the beta values to match the MPC of each wealth quintile, assuming everyone in the wealth quintile is given an equal *dollar* amount extra to spend. The solid lines in figure 5 show the resulting iMPC columns for $s = 0$ and $s = 6$, assuming all household's in the economy are hit with an equal *percentage* increase in their before-tax income—note that the progressive tax means that after-tax income will increase relatively more for low-income households.

The first column of the Jacobian that comes out of this calibrated model, shown as the $s = 0$ solid line in figure 5, shows an MPC of 0.3 in the first quarter. Marginal spending drops rapidly in the second year to below 0.1 and then declines more gradually in the years following. The MPC of 0.3 in the first quarter is in line with the now-large literature on MPCs, and, furthermore, the years that follow match the more limited data available—in particular, that from evidence on spending out of Norwegian lotteries in Fagereng et al. (2021).

There is less empirical evidence on the response of consumption to a news shock about future income—the columns after the first column in the income Jacobian. The solid line with a peak in the sixth quarter in figure 5 shows the model-implied consumption response to an increase in income in six quarters' time. It more or less fits the evidence

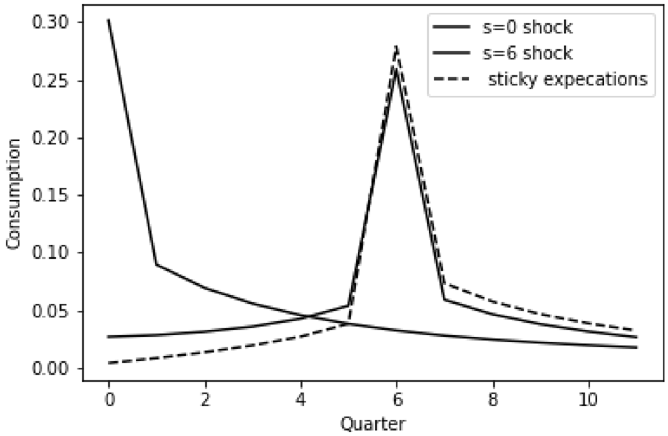

Figure 5: Columns of the consumption-to-income Jacobian, with and without sticky expectations

available: There seems to be little difference in the consumption response to anticipated and unanticipated income shocks. However, there is some evidence that the consumption response to a news shock may not be as symmetric around the arrival date of income as implied by a standard heterogenous agent model. For example, Kueng (2018), using the announcement and arrival of Alaska oil payments, and Ganong and Noel (2019), using the known expiration of unemployment benefits, both find much less anticipatory behavior than a standard model would suggest. To account for this possibility, I allow for households to have sticky expectations with respect to income. Under sticky expectations and following an income news shock, a fixed fraction of households who have yet to update their expectations, θ_{inc}, learn about the news each quarter.[2] Once the income arrives on households' balance sheets, all households learn about the income change. The mechanism is similar to that described in Carroll et al. (2020) and Auclert et al. (2020) but is applied only to the income Jacobian. The effect of sticky expectations is to reduce the size of the consumption response in anticipation of future income and increase the consumption response once the income has arrived. The dotted line in figure 5 shows the consumption response to a news shock six quarters from now with sticky expectations—in this example 15 percent of households update their expectations each quarter.

The resulting income Jacobian is described by one parameter, θ_{inc}, which adjusts the degree to which households anticipate income news shocks. The first column of the Jacobian is pinned down by microdata, while the columns that follow have some flexibility

---

[2] In this paper, θ_{inc} and θ_{sub} denote the fraction of households that update each quarter. Some other papers, such as Carroll et al. (2020), label the fraction that do *not* update.

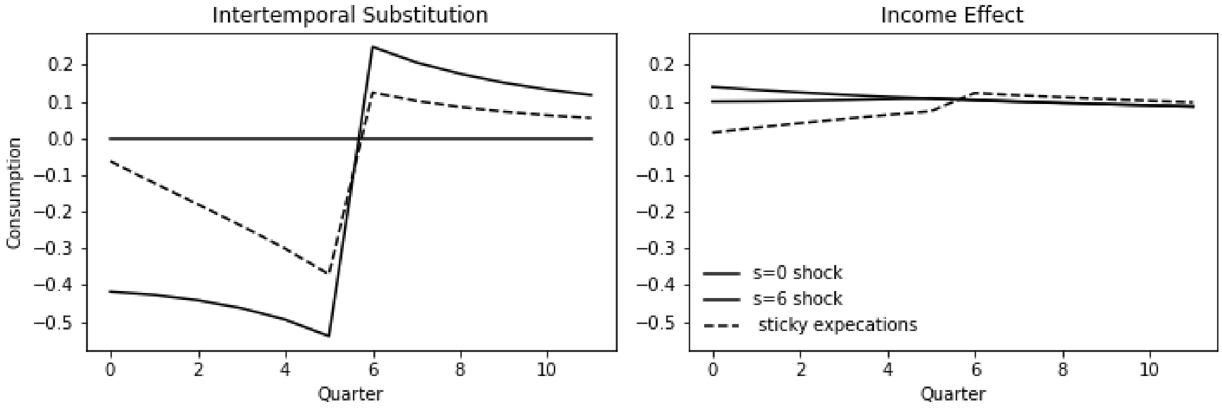

Figure 6: Columns of the consumption-to-interest-rate Jacobian, divided into the substitution and income effect, with and without sticky expectations

remaining. I will estimate θ_{inc} using macrodata. My reading of the empirical literature is that the possible Jacobians spanned by $\theta_{inc} \in [0, 1]$ cover most empirically plausible estimates of the income Jacobian.

The Consumption-to-Real-Interest-Rate Jacobian The response of consumption to a real interest rate news shock can be broken down into a substitution and income effect. Farhi et al. (2022) lays out the theory in the context of an incomplete market with uncertainty. In the heterogeneous agent model used as a starting point for the Jacobians in this paper, the substitution effect can be calculated by changing the interest rate in the Euler equation—affecting the marginal utility of spending in each period—while keeping the interest rate used in the budget constraint fixed. Keeping the rate fixed ensures that the feasible set of consumption plans does not change and therefore that a change in consumption is solely the result of intertemporal substitution. Similarly, the income effect can be calculated bu fixing the interest rate in the Euler equation and changing the interest rate that appears in the budget constraint, ensuring that any change in behavior is derived from a change in the feasible set of consumption plans and not from intertemporal substitution.

Figure 6 plots the first and sixth columns of both the substitution Jacobian (left-hand panel) and the income-effect Jacobian (right-hand panel) in solid black. The dotted lines show the respective sixth columns that assume households have sticky expectations to news shocks about future interest rates.

The basic shape of the substitution Jacobian is similar across many models: News of a future interest rate increase causes households to save more before the rate increase in

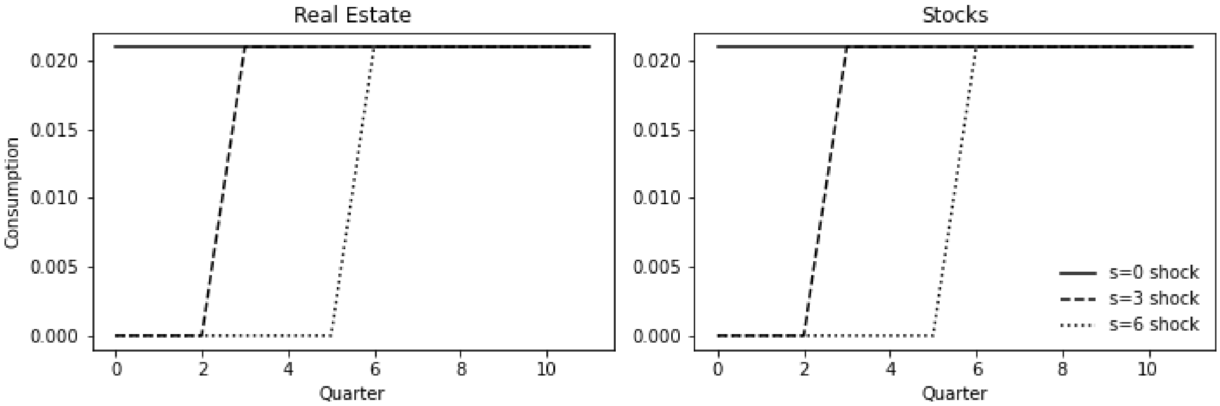

Figure 7: Columns of the consumption-to-stock-market and consumption-to-real-estate Jacobians

order to spend more after the rate increase. Overall, this substitution is budget neutral by construction. In theory, the shape of this Jacobian will depend on the elasticity of intertemporal substitution, the degree of precautionary saving, and other features of the model. In practice, I find that a single parameter for sticky expectations suffices to approximately span a wide range of possible model Jacobians. In order to allow for the possibility of negative intertemporal substitution, and so as to avoid corner solutions, I set $\mathcal{J}(\theta_{sub}) = -\mathcal{J}(-\theta_{sub})$ when $\theta_{sub} < 0$ for the intertemporal substitution Jacobian. This formulation allows the parameters to smoothly move from the standard sign for intertemporal substitution, $\theta_{sub} < 0$, through no intertemporal substitution, $\theta_{sub} = 0$, to intertemporal substitution in the opposite direction of standard theory, $\theta_{sub} > 0$.

The income Jacobian, shown in the right-hand panel of figure 6, shows a positive income effect consumption response to an increase in interest rates. Although the substitution Jacobian—the left-hand panel—is relatively robust to model misspecification, the shape of the income Jacobian from the model is derived from the unrealistic model assumption that households hold only short-term liquid real bonds. In practice, the income effect is governed by the types of assets and liabilities households hold: stocks, bonds, real estate, and fixed and floating mortgages, to name just a few. A model that fully captures all of these, along with their use as borrowing collateral, is beyond the scope of this paper. Instead, I discard the income-effect Jacobian from the model and replace it with empirical estimates of the consumption response to changes in stocks and real estate, the two largest asset classes held by households.

The Consumption-to-Stocks and Consumption-to-Real-Estate Jacobians I set the Jacobians for both the stock market and real estate such that households increase their consumption permanently, but without anticipating when the value of each asset goes up. I calibrate the MPC to both the stock market and real estate to be 0.03, in line with the existing literature, and set the stock of each asset owned by households to be about half of GDP to match the historical average since 1970. Selected columns of the real estate (left-hand panel) and stock market (right-hand panel) Jacobians are shown in figure 7.

3.3 Estimating Empirical Impulse Response Functions

For each structural shock series, I need to calculate the empirical impulse functions for the following outcomes: consumption, income, the real federal funds rate, the real return to the stock market, and the real return to real estate assets. To obtain these impulse response functions, I use Jordà (2005) projections. That is, for each of these outcomes O, I run the following regressions for shock series q:

$$O^q_{t,t+h} = J^{O,q}_h \epsilon^q_t + \beta^{O,q}_h X^q_t + \varepsilon^{O,q}_{t,h}. \tag{6}$$

Here $O^q_{t-1,t+h}$ is the outcome variable. For consumption, income, the stock market, and real estate, this outcome is the percent change from period $t-1$. The real federal funds rate outcome is measured as the federal funds rate minus the median expected one-year-ahead inflation from the Survey of Professional Forecasters. Consumption is total real personal consumption expenditure (PCE), and my measure of income captures the portion of real disposable income that is not derived from capital.[3] This measure of income is chosen to align with the income Jacobian that is calibrated to match the MPC literature. By contrast to shocks to labor income and government transfers, it is well known that shocks to capital income, which are highly skewed to the very wealthy, have a much lower MPC. The stock market measure uses prices from Fama French and is adjusted for PCE inflation, and the real estate measure comes from the Case-Shiller house price index, also adjusted for PCE inflation. The object of interest that I wish to estimate is $J^{O,q}_h$, the impulse response for outcome variable O after h periods following a structural shock of type q and magnitude one. The structural shock series of type q takes the value ϵ^q_t at time t. Following Ramey (2016), I include in the set of controls X^q_t two lags each of log industrial production, the unemployment rate, log of the consumer price index, log of a commodity price index, and

---

[3] Specifically, the definition of income I use is the sum of compensation of employees, proprietors' income with inventory valuation and capital consumption adjustments, and transfers minus contributions to Social Security and 80 percent of personal taxes. The taxes are chosen to align with the proportion of personal taxes paid on non-capital income.

the federal funds rate. I also include two lags of the shock series itself, which is why the controls X_t^q are indexed by q as well as t.

For each outcome, I run Jordà projections at a monthly frequency for a horizon of up to 48 months.[4] I compute the standard errors for $\widehat{J_h^{O,q}}$ using the correction from Newey and West (1987).

Choice of Structural Shocks The methodology in this paper allows me to use any structural shock series, so long as the shock only affects consumption through its effect on income and asset market returns. The recent pandemic is an example of a shock that clearly violates these assumptions—households chose, or were obliged, to cut back on their normal consumption activities in order to socially distance and limit the spread of disease. However, many of the shocks studied in the literature are thought to affect consumption only indirectly. For example, a typical New Keynesian model will feature monetary policy shocks, total factor productivity (TFP) shocks, and government spending shocks that only change consumption behavior through their effect on income and asset returns.

A large number of structural shocks have been proposed in the literature, along with different methods to identify them. However, for any one particular structural shock there is no consensus on whether it is well identified. I therefore take the approach of using a wide variety of structural shocks in the hope that—while no single shock will be sufficient to convince the reader of households' consumption behavior—the aggregate evidence will be overwhelming.

In order to limit the number or shock series to a manageable number while not cherry-picking series, I start with the universe of shocks that is described in Ramey (2016). This handbook chapter is an overview of the shock literature and includes monetary policy shocks, government spending shocks, tax shocks, and technology shocks. Furthermore, the chapter covers a variety of different identification methodologies. Starting from the shock series in this handbook, I throw out those that are identified explicitly using a DSGE model that conflicts with the income Jacobian that I am assuming in this paper and also those that are identified using a structural vector auto-regression (SVAR) model.

Monetary Policy Shocks

Monetary policy shocks are some of the most studied structural shocks because economists and central bankers are naturally interested in the effects of monetary policy on economic outcomes. Ramey (2016) examines representative shock series from three

---

[4] Where monthly data is not available, I interpolate quarterly data to a monthly frequency.

different methods for identifying monetary policy shocks. The first, from Christiano et al. (1999), uses a recursive assumption in a SVAR model. My analysis in this paper is on Jordà projections, and I will not examine this shock series. The next two monetary policy shock series are included in my analysis.

1. Romer and Romer (2004). In this method, monetary policy shocks are identified by regressing the federal funds target rate on the Greenbook forecasts at each FOMC meeting and taking the residual to be the shock.
2. Gertler and Karadi (2015). This high frequency identification method uses changes to the three-month-ahead fed funds futures around a window of FOMC announcements to find surprise changes to the policy rate.

In most macroeconomic models, monetary policy shocks satisfy the requirement that the effect of these shocks on consumption comes only through income and asset price changes. However, these shocks suffer from some limitations. First, these shocks are thought to be responsible for only a small fraction of the total forecast variance. Second, the identification methods used here only pick up a further small fraction of total monetary policy shocks. Consequently, the estimated IRFs have large standard errors and are sensitive to the exact time period over which they are estimated. Furthermore, the so-called Fed information effect draws into question whether these shocks are truly shocks to monetary policy or if they are in reaction to other macroeconomic events.

Fiscal Shocks

3. Ramey and Zubairy (2018) military news.
4. Ben Zeev and Pappa (2017) defense spending shocks.
5. Mertens and Ravn (2012) tax news shocks.
6. Leeper et al. (2012) expected taxes from one to five years forward.

Technology Shocks

7. Ben Zeev and Khan (2015) investment-specific news shocks.
8. Fernald (2012) utilization-adjusted TFP.
9. Fernald (2012) utilization-adjusted investment TFP.
10. Francis et al. (2014) unanticipated TFP shocks.

3.4 Estimating the Consumption Jacobians

After parameterizing the consumption Jacobians in section 3.2 and estimating the empirical impulse responses to structural shocks in section 3.3, the last step is to estimate the parameters of the Jacobians to best fit the empirical impulse response functions. I do this with a standard minimum-distance estimation method.

Given parameters for $\boldsymbol{\theta} = (\theta_{inc}, \theta_{sub})$ and input-estimated IRFs $\hat{\boldsymbol{J}}^q = (\hat{J}^{Y,q}, \hat{J}^{R,q}, \hat{J}^{\text{stocks},q}, \hat{J}^{\text{realestate},q})$, I can calculate the implied consumption IRF for a structural shock of type q:

$$C(\hat{\boldsymbol{J}}^q, \boldsymbol{\theta}) = \sum_{O \in \{Y,R,\text{stocksrealestate}\}} \mathcal{J}^{C,O}(\boldsymbol{\theta})\hat{J}^{O,q}.$$

This consumption impulse response implied by the Jacobians and the empirical input IRFs can then be compared to the empirical IRF for consumption, $\hat{C}^q = \hat{J}^{C,q}$. For each structural shock, I create a loss function as a function of the parameters, $\mathcal{L}^q(\theta)$.

$$\mathcal{L}^q(\boldsymbol{\theta}) = (C(\hat{\boldsymbol{J}}^q, \boldsymbol{\theta}) - \hat{C}^q)' \Sigma_q^{-1} (C(\hat{\boldsymbol{J}}^q, \boldsymbol{\theta}) - \hat{C}^q) \tag{7}$$

Here, Σ_q is a diagonal matrix of estimated consumption IRF variances. Using a set of structural shocks Q—which may consist of just one or up to all of the shocks described in section 3.3—my estimator for the Jacobian parameters is the parameter vector that minimizes the sum of losses over all structural shocks in Q.

$$\hat{\boldsymbol{\theta}} = \underset{\theta}{\operatorname{argmin}} \sum_{q \in Q} \mathcal{L}^q(\boldsymbol{\theta})$$

4 Results

Table 1 shows the estimation results for the sticky expectations parameters for intertemporal substitution and income along with standard errors in parentheses below. The top row shows the result of estimation using all 10 structural shocks, and the following rows show the result of estimation using each shock series individually.

The main result of the paper is that, in the estimation using all the structural shocks together, I find $\hat{\theta}_{sub} \approx 0$. This finding suggests that households do not intertemporally substitute—their consumption response to these structural shocks is best explained only through their response to income and wealth effects. When estimated individually for each selected structural shock, only 2 out of the 10 shocks result in positive estimates for $\hat{\theta}_{sub}$ intertemporal substitution. For all the shocks, $\hat{\theta}_{sub}$ is estimated to be statistically indifferent from zero, implying a negligible role for intertemporal substitution.

Set of Structural Shocks (Q)	$\hat{\theta}_{sub}$	$\hat{\theta}_{inc}$
All	-0.06	0.31
	(0.04)	(0.20)
Romer and Romer (2004)	-0.08	0.29
	(0.25)	(0.89)
Gertler and Karadi (2015)	-0.27	1.00
	(0.65)	(13.04)
Ramey and Zubairy (2018) military news	0.09	0.00
	(0.39)	(0.16)
Ben Zeev and Pappa (2017) defence spending shocks	0.00	0.02
	(0.11)	(0.17)
Mertens and Ravn (2012) tax news shocks	0.05	0.00
	(0.15)	(0.39)
Leeper et al. (2012) expected taxes from one to five years forward	-0.06	0.31
	(0.11)	(2.91)
Ben Zeev and Khan (2015) investment specific news shocks	-0.10	0.13
	(0.13)	(0.16)
Fernald (2012) utilization-adjusted TFP	-0.38	0.97
	(1.02)	(3.15)
Fernald (2012) utilization-adjusted investment TFP	-0.29	0.00
	(0.53)	(0.86)
Francis et al. (2014) unanticipated TFP shocks	-0.83	0.79
	(2.64)	(0.72)

Table 1: Parameter estimates using all structural shocks together and each individually

While all 10 estimates of θ_{sub} are statistically indistinguishable from zero, some of the structural shocks are far more informative about the possible size of the intertemporal substitution effect than others. For example, the standard errors from the estimates using Ben Zeev and Pappa (2017), Mertens and Ravn (2012), Leeper et al. (2012), and Ben Zeev and Khan (2015) are all below 0.15, while others such as those from Fernald (2012) and Francis et al. (2014) are above 1.0. These differences come from the fact that some of the structural shocks feature real interest rate IRFs that are both large in magnitude and persistent, while in others the real interest rate barely moves. The standard errors from typical monetary policy shocks—Romer and Romer (2004) and Gertler and Karadi (2015)—are somewhere in between, highlighting the advantage being able to use a wider range of structural shocks to better identify intertemporal substitution behavior.

The estimates for sticky expectations for income cover a broad range of possibilities. Some structural shocks suggest that households do not react in anticipation of future

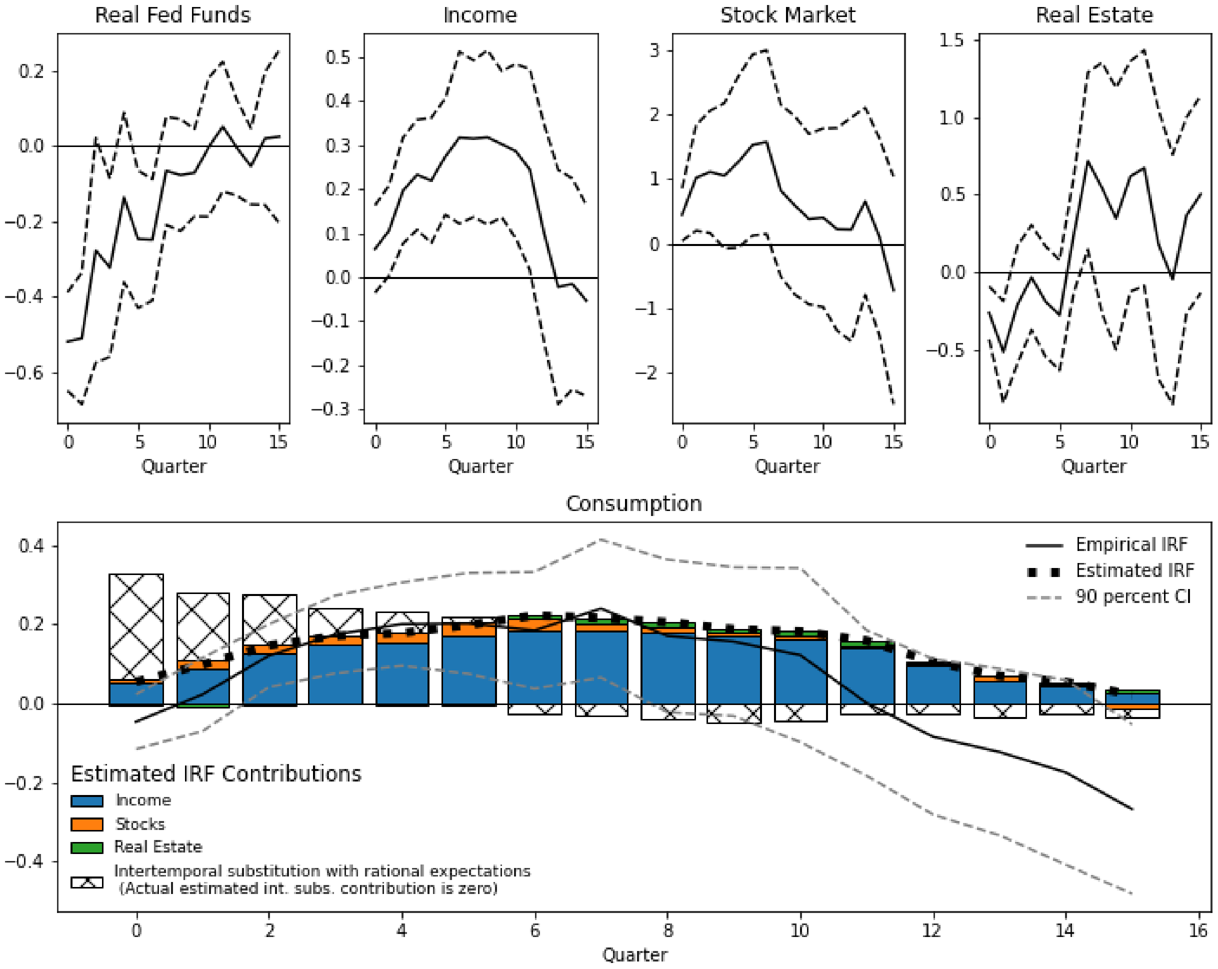

Figure 8: Inputs and output IRFs to the consumption block following a Romer and Romer (2004) shock

income changes ($\hat{\theta}_{inc} = 0$), while others suggest households fully anticipate predictable income changes ($\hat{\theta}_{inc} = 1$). Estimation based on all 10 structural shocks leads to an estimate of 0.3. That is, each quarter, about 30 percent of households update their consumption based on the true expected income at a point in time in the future. The wide array of estimates for this parameter reflects the fact that, in contrast to the value of θ_{sub}, the value of θ_{inc} has a relatively small effect on the impulse response function and therefore the loss value.

In the rest of this results section, I will examine the results in the context of the shock series from Romer and Romer (2004) before examining in more detail the consumption IRFs for all 10 shock series.

Results using Romer and Romer (2004) Monetary Policy Shocks. The top row of figure 8 shows the empirical impulse response functions, along with 90 percent confidence intervals, for a Romer and Romer shock for each of the inputs to the consumption block: the real fed funds rate, labor income, the stock market, and real estate. The plot shows results for the sticky expectations parameters estimated on all shocks with a restriction that intertemporal substitution cannot be negative (and is in practice therefore estimated to be zero). Following a shock of this type, the real fed funds rate immediately drops and then dissipates over a relatively short period of one to two years. Labor income gradually rises over this period, peaking a little less than two years following the shock, before falling back to trend by the end of the period calculated. The stock market rises a little, and real estate prices oscillate, but the change in both asset prices is only just statistically significant at the 90 percent level.

The larger lower panel shows the empirical impulse response of consumption along with detail on the impulse response estimated using the four inputs from the top row of the panel. The solid black line, the empirical impulse response, shows that, after a brief decline, consumption rises gradually over a year and a half before declining to somewhat below trend four years after the shock. The 90 percent confidence intervals show the increase to be statistically significant. The bold dotted black line shows the impulse response function calculated by applying the estimated Jacobians—using parameters from the estimation for all 10 shocks—for each of the four inputs in the top row of the panel. This estimated impulse response is close to the empirical impulse response, although it doesn't capture the decline in consumption at the end of the period.

The blue, orange, and green bars in the lower panel of figure 8 show the contribution to the estimated IRF for each of the following inputs: labor income, stocks, and real estate. These three bars add up to the bold dotted black line showing the estimated IRF. The contribution from labor income, the blue bars, accounts for the bulk of the estimated IRF, with asset price movements—which are only just statistically significant—making only a small contribution to consumption.

To illustrate why the contribution from intertemporal substitution is estimated to be zero, figure 8 also shows the size of this intertemporal substitution contribution that is implied by the model, assuming rational expectations—the hashed bars. These hashed bars push up consumption substantially over the first year following the shock, which then leads to a slight negative contribution to consumption after the shock's effect on the real fed funds rate has dissipated. These contributions from intertemporal substitution over the first year would act to pull the estimated IRF away from the empirical IRF and therefore increase the loss function. As a result, the sticky expectations parameter is estimated to be close to zero, equivalent to households paying no attention to the real fed funds rate when making their consumption—saving decisions.

Results for All 10 Structural Shocks. Figure 9 shows the empirical consumption IRFs for all 10 structural shocks used in the estimation with intertemporal substitution restricted to be positive, and therefore zero in practice. The top-left panel of figure 9 is a repeat of the lower panel in figure 8, except that the contributions from stocks and real estate have been combined. The other nine panels show the same elements of the consumption IRF and its decomposition for each of the other shocks. There are three key takeaways from figure 9.

First, the estimated IRF from applying the Jacobians to the empirical IRFs for each of the four inputs does a good job at fitting the empirical consumption IRF for all 10 shock types. Indeed, the estimated IRF rarely escapes the 90 percent confidence interval bounds and often closely tracks the central estimate. Furthermore, the labor income contribution to the consumption IRF forms the bulk of the response, with stocks and real estate playing a negligible role except in the "Tax News" shock panel. The small role played by stocks and real estate comes partly from the small movements of these asset prices following each shock relative to the size of asset price swings often observed in the market. The Jacobian I used to calculate these consumption responses to asset price changes assumed a 3 percent MPC—the fact that these price changes play such a small role in the decomposition shows that the main results of this paper would likely be robust to replacing these Jacobians with any other reasonable calibration.

Second, for all the shock series, the contribution of intertemporal substitution, if it were assumed households had rational expectations, would be to significantly increase the loss function. In almost all cases, adding the contribution from intertemporal substitution would move the estimated consumption IRF well outside of the 90 percent confidence intervals of the empirical consumption IRF. In fact, only for the"Tax News" and "Ramey News" shock series does a little bit of intertemporal substitution—sticky expectations parameters of 0.05 and 0.09, respectively—even help move the estimated IRF toward the empirical IRF. For the other eight shock series, intertemporal substitution is estimated as negative.

Third, some of the shock series feature IRFs for the real fed funds rate (not shown) that remain significantly negative for a prolonged period relative to the identified monetary policy shocks that are commonly studied. For example, the real fed funds rate remains negative throughout the period shown for the Ben Zeev and Pappa (2017) defense news shocks and the expected tax shocks from Leeper et al. (2012). The investment-specific news

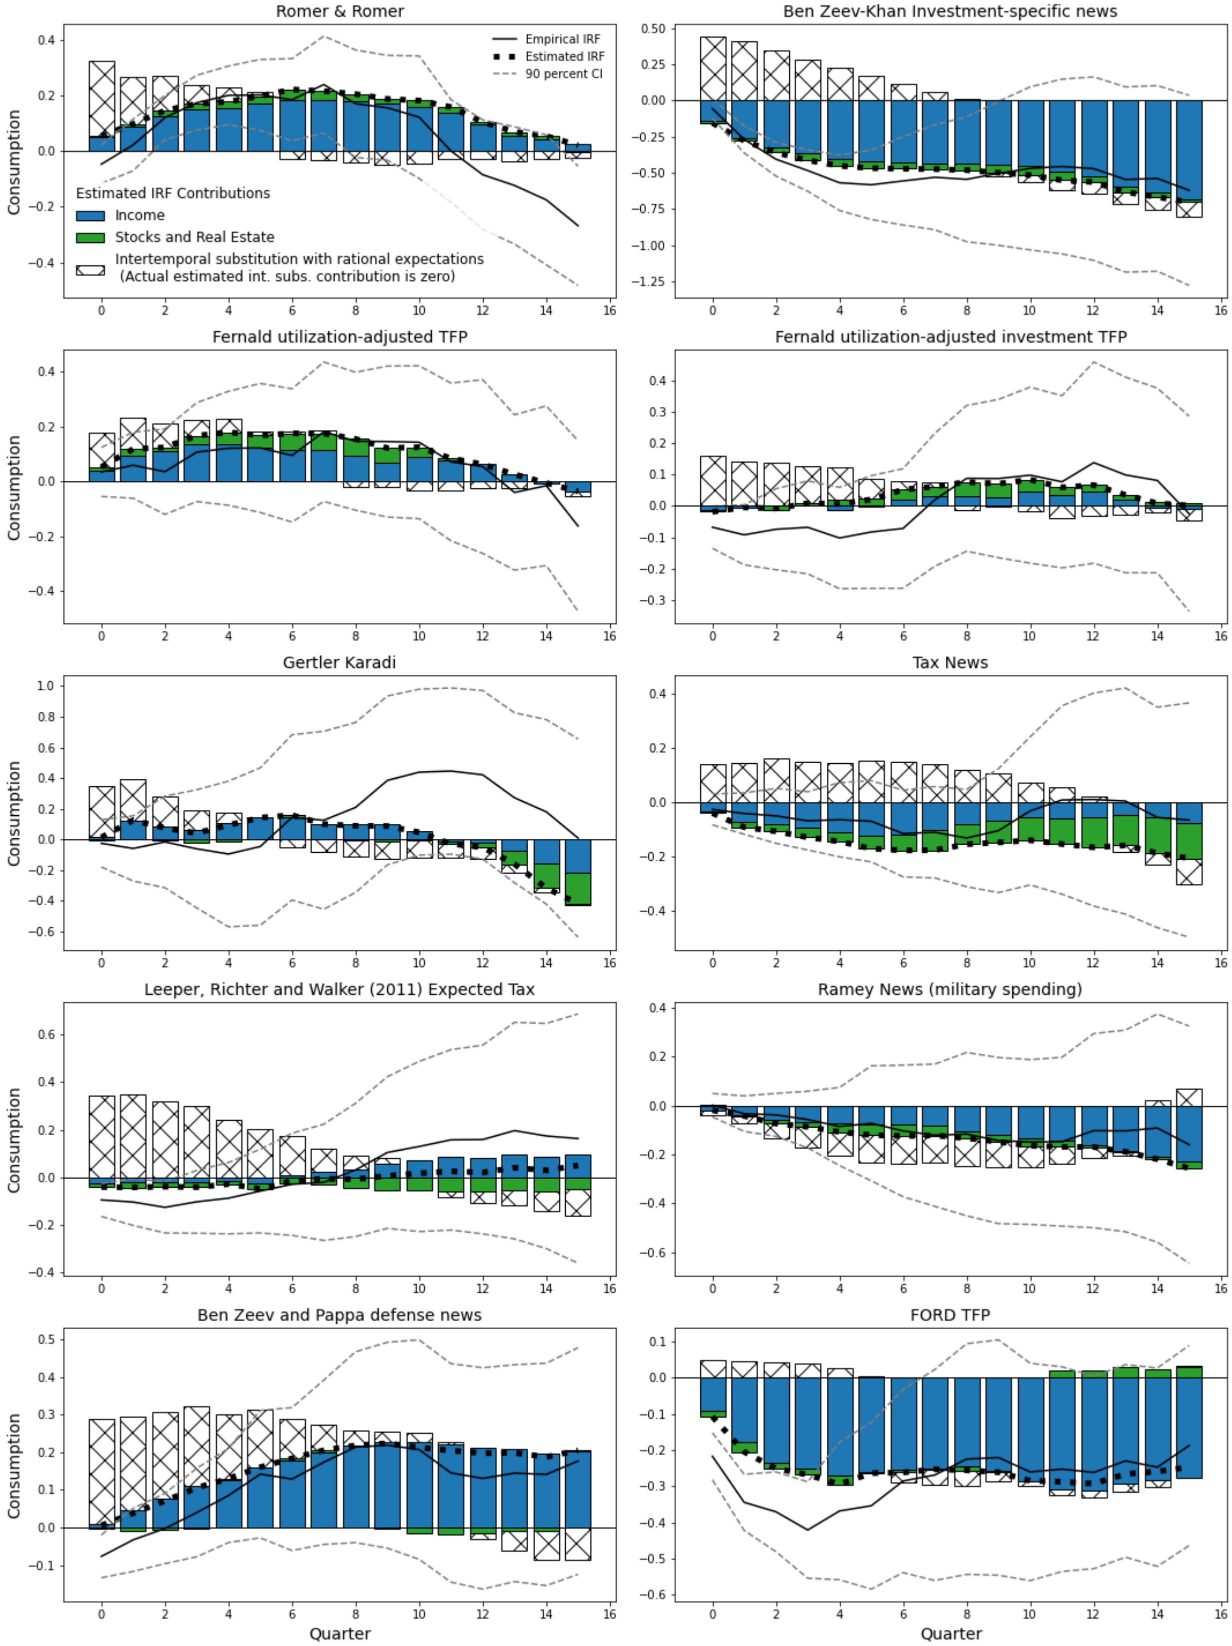

Figure 9: The consumption IRFs for all 10 shocks

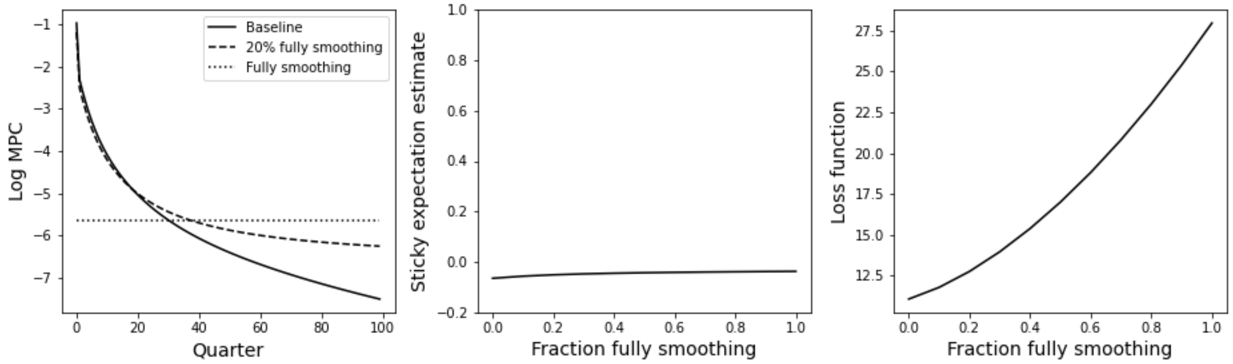

Figure 10: Log MPC out up to 40 years (left-hand panel), sticky-expectations parameter estimate (middle panel), and loss function (right-hand panel) for different fractions of fully-smoothing agents in the model.

shocks from Ben Zeev and Khan (2015) also result in a prolonged period of depressed real fed funds rates. The lack of an intertemporal substitution contribution from these shocks, despite a long period of lowered real fed funds rate, provides stronger evidence than the shorter-lived identified monetary policy shocks that households do not appear to substitute intertemporally.

5 Robustness

5.1 The shape of the intertemporal MPC

For the baseline estimation, I took the consumption—income Jacobian from a one-asset buffer-stock model as the starting point. I then estimated the sticky-expectations parameter which best fit the impulse responses. One concern of this method is that the one-asset buffer-stock model has some known problems, and, in particular, the way I have calibrated it to match the one-year MPC for different wealth quintiles results in far too little aggregate wealth. A two-asset model can address some of these concerns.

In Auclert et al. (2018), the authors examine the differences between the iMPCs in a selection of models. Somewhat reassuringly, they find few material differences in the columns of the iMPCs between the models they examine, conditional on matching the empirical evidence for the first column. The one exception is that, in their two-asset model, the marginal propensity to consume four or more years after the an income shock decays more slowly than in the other models examined. This feature of the two-asset model can lead to far larger fiscal multipliers in their general equilibrium model.

A further concern is that MPCs may be lower than suggested by the empirical evidence

I use to discipline my one-asset model. The MPC esimates I use are at the high end of those in the literature, and, furthermore, Havranek and Sokolova (2020) finds evidence of publication bias in this literature.

To address these concerns, I examine the effects of replacing the baseline iMPC I use with a linear combination of the baseline iMPC and the iMPC that comes from a fully-smoothing agent. The left-hand panel of figure 10 shows how introducing a fraction of fully-smoothing agents changes the tail behavior of MPCs. The solid line shows the log MPCs from my baseline one-asset model, the dashed line shows those from a model in which 80 percent of agents are similar to those of the baseline model and 20 percent are fully smoothing, and the dotted line shows the fixed MPC of fully-smoothing agents. Compared to the baseline model, the model with 20 percent fully-smoothing agents displays a far slower decline in the MPC in the very long tail. This behavior captures the main difference that also exists between a two-asset model and a one-asset model. Note that, relative to the one-asset model in Auclert et al. (2018), my baseline model already has a slower decline in MPCs because my baseline model includes heterogeneity in discount factors.

The middle panel of figure 10 shows that the sticky expectations parameter estimate is not sensitive to variation in the model in this way—the estimate is close to zero (and slightly negative) for all versions of the underlying model up to 100 percent fully-smoothing agents. However, as shown in the right-hand panel of figure 10, increasing the fraction of fully-smoothing agents in the model diminishes the fit of the model to the empirical impulse response functions.

This exercise shows that structural shocks identified in the literature show no evidence of intertemporal substitution, even when viewed through the lens of older representative-agent models. However, the iMPCs implied by heterogenous agent models do a better job of matching the impulse response data than representative agent models. Furthermore, the slow decline in the tail MPCs that are a feature of two-asset models, which can greatly increase fiscal multipliers in some general equilibrium settings, do not have a large effect on the estimation of sticky expectations in my set up.

5.2 Isolating the magnitude of intertemporal substitution

In the baseline estimation sticky expectations, the single parameter changes the shape of the intertemporal substitution Jacobian in two ways: 1) reducing attentiveness delays the reaction to changes in interest rates causing a hump shape, and 2) the overall magnitude of the consumption reaction is dampened. In order to test just the magnitude of the intertemporal response, in this robustness exercise I fix the value of the sticky expectations parameter at 0.2. I then introduce a scalar to allow for a bigger or smaller reaction to real interest rates. That is, $\mathcal{J}^{C,R}_{\text{magnitude only}}(\gamma) = \gamma \mathcal{J}^{C,R}(0.2)$—the profile of the intertemporal response remains fixed, but its magnitude is controlled by γ. When I estimate γ to best fit the impulse response for the 10 structural shocks, I estimate a negative coefficient with fairly wide error bands. This evidence suggests the result that households do not substitute intertemporally is robust to other likely shapes that the intertemporal substitution Jacobian might take.

6 Conclusion

In this paper, I have presented a new way to estimate the size of the intertemporal substitution effect on household behavior. Despite the traditional emphasis on intertemporal substitution as a pivotal factor in macroeconomic models influencing household consumption, I find no evidence of such an effect across ten different structural shocks. This evidence is inline with other recent advances in the heterogenous agent literature, but by isolating the consumption block of the model I have been able to bring to bear an much wider array of empirical evidence. The finding that intertemporal substitution has little bearing on consumption behavior invites further research on the transmission of monetary policy, with a possibility that investment behavior drives the response or that output is less responsive to monetary policy than previously thought.

References

George A. Akerlof and Janet L. Yellen. A near-rational model of the business cycle, with wage and price inertia. *The Quarterly Journal of Economics*, 100:823–838, 1985. ISSN 00335533, 15314650. URL `http://www.jstor.org/stable/1882925`.

Adrien Auclert, Matthew Rognlie, and Ludwig Straub. The intertemporal keynesian cross. Working Paper 25020, National Bureau of Economic Research, September 2018. URL `http://www.nber.org/papers/w25020`.

Adrien Auclert, Matthew Rognlie, and Ludwig Straub. Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model. Working Paper 26647, National Bureau of Economic Research, January 2020. URL `http://www.nber.org/papers/w26647`.

Adrien Auclert, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408, 2021. doi: https://doi.org/10.3982/ECTA17434. URL `https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA17434`.

Régis Barnichon and Geert Mesters. A sufficient statistics approach for macro policy. *American Economic Review*, 113(11):2809–45, November 2023. doi: 10.1257/aer.20220581. URL `https://www.aeaweb.org/articles?id=10.1257/aer.20220581`.

Nadav Ben Zeev and Hashmat Khan. Investment-specific news shocks and u.s. business cycles. *Journal of Money, Credit and Banking*, 47(7):1443–1464, 2015. ISSN 00222879, 15384616. URL `http://www.jstor.org/stable/24500920`.

Nadav Ben Zeev and Evi Pappa. Chronicle of a War Foretold: The Macroeconomic Effects of Anticipated Defence Spending Shocks. *Economic Journal*, 127(603):1568–1597, August 2017. URL `https://ideas.repec.org/a/wly/econjl/v127y2017i603p1568-1597.html`.

Martin Beraja. A semistructural methodology for policy counterfactuals. *Journal of Political Economy*, 131(1):190–201, 2023.

Michael Carlos Best, James S Cloyne, Ethan Ilzetzki, and Henrik J Kleven. Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches. *The Review of Economic Studies*, 87(2):656–690, 05 2019. ISSN 0034-6527. doi: 10.1093/restud/rdz025. URL `https://doi.org/10.1093/restud/rdz025`.

Martin Browning and Annamaria Lusardi. Household saving: Micro theories and micro facts. *Journal of Economic Literature*, 34(4):1797–1855, 1996. ISSN 00220515. URL `http://www.jstor.org/stable/2729595`.

Christopher D Carroll. Death to the log-linearized consumption euler equation! (and very poor health to the second-order approximation). Working Paper 6298, National Bureau of Economic Research, December 1997. URL `http://www.nber.org/papers/w6298`.

Christopher D. Carroll, Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. Sticky expectations and consumption dynamics. *American Economic Journal: Macroeconomics*, 12(3):40–76, July 2020. doi: 10.1257/mac.20180286. URL `https://www.aeaweb.org/articles?id=10.1257/mac.20180286`.

James J. Choi. Popular personal financial advice versus the professors. *Journal of Economic Perspectives*, 36(4):167–92, November 2022. doi: 10.1257/jep.36.4.167. URL `https://www.aeaweb.org/articles?id=10.1257/jep.36.4.167`.

Lawrence J. Christiano, Martin Eichenbaum, and Charles L. Evans. Chapter 2 monetary policy shocks: What have we learned and to what end? In *Handbook of Macroeconomics*, volume 1, pages 65–148. Elsevier, 1999. doi: https://doi.org/10.1016/S1574-0048(99)01005-8. URL `https://www.sciencedirect.com/science/article/pii/S1574004899010058`.

John H. Cochrane. The sensitivity of tests of the intertemporal allocation of consumption to near-rational alternatives. *The American Economic Review*, 79(3):319–337, 1989. ISSN 00028282. URL `http://www.jstor.org/stable/1806848`.

Edmund Crawley and Andreas Kuchler. Consumption heterogeneity: Micro drivers and macro implications. *American Economic Journal: Macroeconomics*, 15(1):314–41, January 2023. doi: 10.1257/mac.20200352. URL `https://www.aeaweb.org/articles?id=10.1257/mac.20200352`.

Richard K. Crump, Stefano Eusepi, Andrea Tambalotti, and Giorgio Topa. Subjective intertemporal substitution. *Journal of Monetary Economics*, 126:118–133, 2022. ISSN 0304-3932. doi: https://doi.org/10.1016/j.jmoneco.2021.11.008. URL `https://www.sciencedirect.com/science/article/pii/S0304393221001410`.

Andreas Fagereng, Martin B. Holm, and Gisle J. Natvik. Mpc heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics*, 13(4):1–54, October 2021. doi: 10.1257/mac.20190211. URL `https://www.aeaweb.org/articles?id=10.1257/mac.20190211`.

Emmanuel Farhi, Alan Olivi, and Iván Werning. Price theory for incomplete markets. Working Paper 30037, National Bureau of Economic Research, May 2022. URL `http://www.nber.org/papers/w30037`.

John G. Fernald. A quarterly, utilization-adjusted series on total factor productivity. Working Paper Series 2012-19, Federal Reserve Bank of San Francisco, April 2012. URL `https://ideas.repec.org/p/fip/fedfwp/2012-19.html`.

Neville Francis, Michael T. Owyang, Jennifer E. Roush, and Riccardo DiCecio. A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. *The Review of Economics and Statistics*, 96(4):638–647, 2014. ISSN 00346535, 15309142. URL `http://www.jstor.org/stable/43554945`.

Jordi Galí. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications Second edition*. Number 10495 in Economics Books. Princeton University Press, 2015. URL https://ideas.repec.org/b/pup/pbooks/10495.html.

Peter Ganong and Pascal Noel. Consumer spending during unemployment: Positive and normative implications. *American Economic Review*, 109(7):2383–2424, July 2019. doi: 10.1257/aer.20170537. URL https://www.aeaweb.org/articles?id=10.1257/aer.20170537.

Mark Gertler and Peter Karadi. Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76, January 2015. doi: 10.1257/mac.20130329. URL https://www.aeaweb.org/articles?id=10.1257/mac.20130329.

Robert E Hall. Intertemporal substitution in consumption. *Journal of political economy*, 96 (2):339–357, 1988.

Tomas Havranek and Anna Sokolova. Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say "probably not". *Review of Economic Dynamics*, 35:97–122, 2020. ISSN 1094-2025. doi: https://doi.org/10.1016/j.red.2019.05.004. URL https://www.sciencedirect.com/science/article/pii/S1094202518302849.

James Hebden and Fabian Winkler. Impulse-based computation of policy counterfactuals. *Finance and Economics Discussion Series 2021-042. Washington: Board of Governors of the Federal Reserve System*, 2021. doi: https://doi.org/10.17016/FEDS.2021.042.

Òscar Jordà. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, March 2005. doi: 10.1257/0002828053828518. URL https://www.aeaweb.org/articles?id=10.1257/0002828053828518.

Greg Kaplan, Benjamin Moll, and Giovanni L. Violante. Monetary policy according to hank. *American Economic Review*, 108(3):697–743, March 2018. doi: 10.1257/aer.20160042. URL https://www.aeaweb.org/articles?id=10.1257/aer.20160042.

Lorenz Kueng. Excess Sensitivity of High-Income Consumers*. *The Quarterly Journal of Economics*, 133(4):1693–1751, 06 2018. ISSN 0033-5533. doi: 10.1093/qje/qjy014. URL https://doi.org/10.1093/qje/qjy014.

Eric M. Leeper, Alexander W. Richter, and Todd B. Walker. Quantitative effects of fiscal foresight. *American Economic Journal: Economic Policy*, 4(2):115–44, May 2012. doi: 10.1257/pol.4.2.115. URL https://www.aeaweb.org/articles?id=10.1257/pol.4.2.115.

Alisdair McKay and Christian Wolf. What can time-series regressions tell us about policy counterfactuals? *Econometrica*, Forthcoming.

Karel Mertens and Morten O. Ravn. Empirical evidence on the aggregate effects of anticipated and unanticipated us tax policy shocks. *American Economic Journal: Economic*

Policy, 4(2):145–81, May 2012. doi: 10.1257/pol.4.2.145. URL `https://www.aeaweb.org/articles?id=10.1257/pol.4.2.145`.

Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987. ISSN 00129682, 14680262. URL `http://www.jstor.org/stable/1913610`.

Christina Patterson. The matching multiplier and the amplification of recessions. *American Economic Review*, 113(4):982–1012, April 2023. doi: 10.1257/aer.20210254. URL `https://www.aeaweb.org/articles?id=10.1257/aer.20210254`.

V.A. Ramey. Chapter 2 - macroeconomic shocks and their propagation. In John B. Taylor and Harald Uhlig, editors, *Handbook of Macroeconomics*, volume 2, pages 71–162. Elsevier, 2016. doi: https://doi.org/10.1016/bs.hesmac.2016.03.003. URL `https://www.sciencedirect.com/science/article/pii/S1574004816000045`.

Valerie A. Ramey and Sarah Zubairy. Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data. *Journal of Political Economy*, 126(2): 850–901, 2018. doi: 10.1086/696277. URL `https://ideas.repec.org/a/ucp/jpolec/doi10.1086-696277.html`.

Christina D. Romer and David H. Romer. A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084, September 2004. doi: 10.1257/0002828042002651. URL `https://www.aeaweb.org/articles?id=10.1257/0002828042002651`.