Do Households Substitute Intertemporally? 10 Structural Shocks That Suggest Not*

Edmund Crawley

December 2024

Abstract

I combine microdata on the intertemporal marginal propensity to consume (iMPC) with 10 structural macro shocks to identify the role of intertemporal substitution in consumption behavior. Although some of the structural shocks that I examine 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 that I select from Ramey (2016). I find that all 10 structural shocks I examine suggest no role for intertemporal substitution in household decision making. The estimates are statistically significant and together provide a tight estimate close to zero.

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), who estimate a fully-specified general equilibrium model, this paper contributes to the literature by narrowing the scope of interest to focus only on the consumption "block" of a 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 both fit 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 Sequence Space Jacobian toolkit described in Auclert et al. (2021).

Each of the four Jacobians that describe the dynamics of the consumption block is of infinite dimension. The second structure I impose on the consumption block is to reduce the dimensionality of these Jacobians—to do this I use micro-empirical evidence along with theory. For the income Jacobian, I use evidence from Fagereng et al. (2021) to fit a one-asset heterogenous agent (HA) model to the marginal propensities to consume (MPC) for each of the first five years following a lottery win. 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 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 one-asset 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 increase, 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 a HA 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 from its expected value before the arrival of the shock at every quarter s after the shock hits. 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 function* 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 the reader 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 Hall (1988) 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 there has been mized evidence on intertemporal substitution. 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. The decomposition of consumption behavior in response to monetary policy shocks is studied by Holm et al. (2021) who use of Norwegian administrative data to uncover heterogenous behavior.

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:

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

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 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:

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

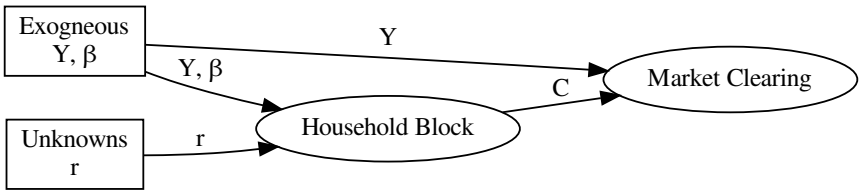


Figure 1: Directed acyclical graph for a simple endowment economy

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

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

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

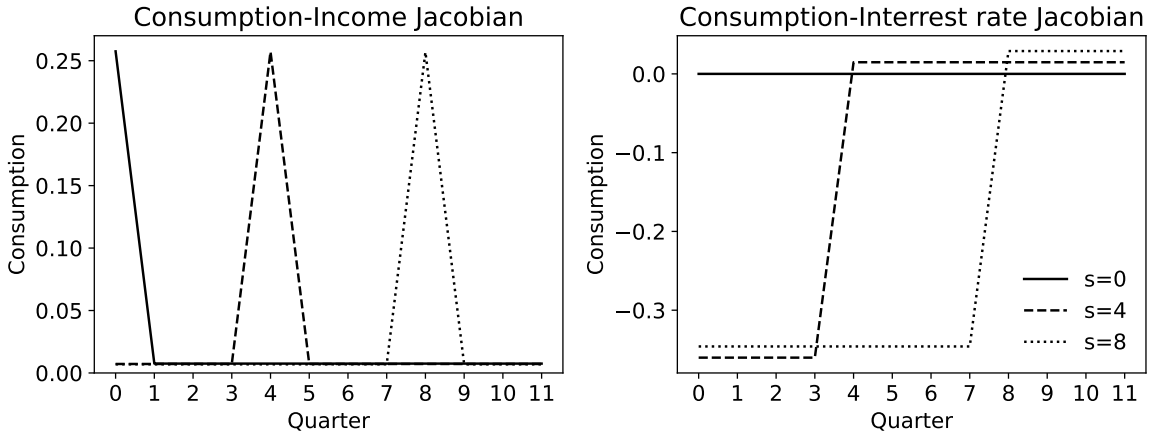


Figure 2: Consumption Jacobians for the simple two-agent model

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 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

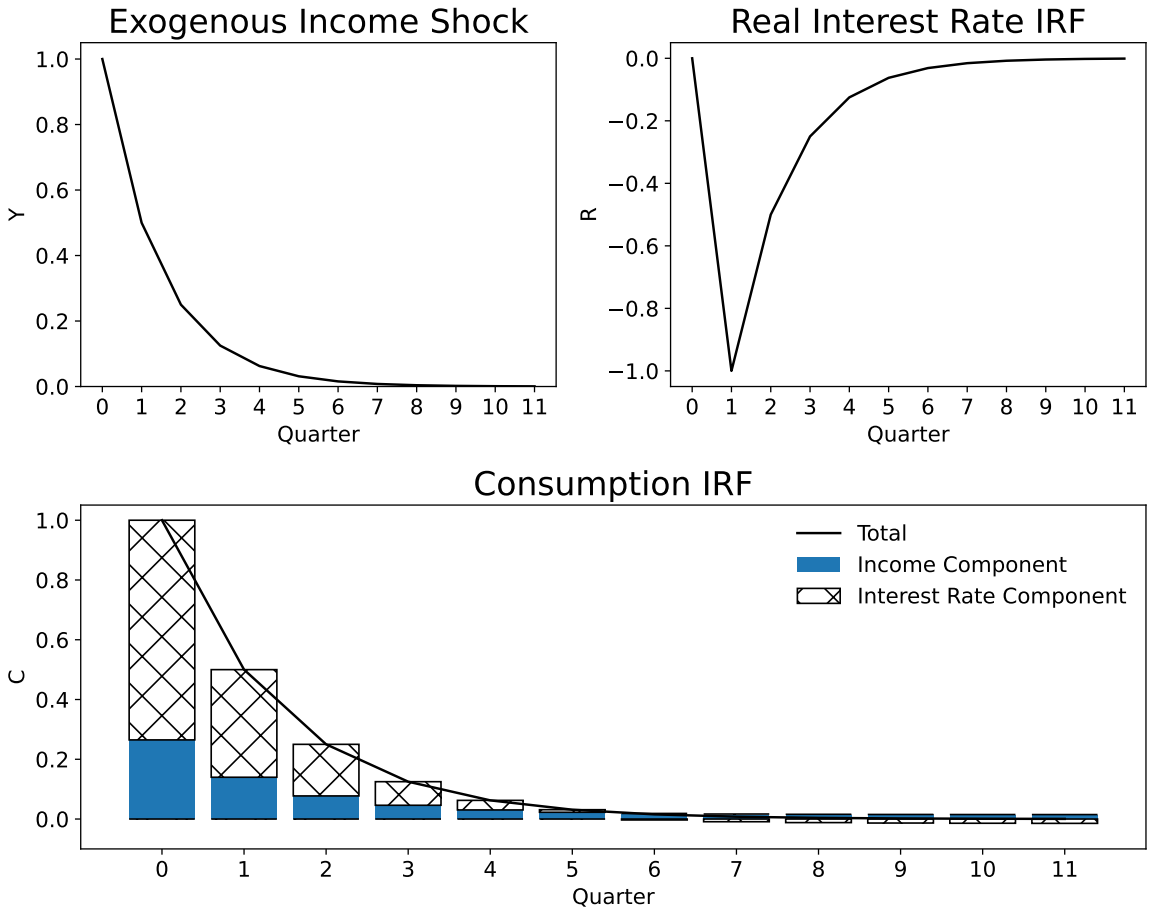


Figure 3: Impulse response to an exogneous income shock

rates. In this simple model with only two exogenous variables, these identified shocks are exogenous shocks 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.

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

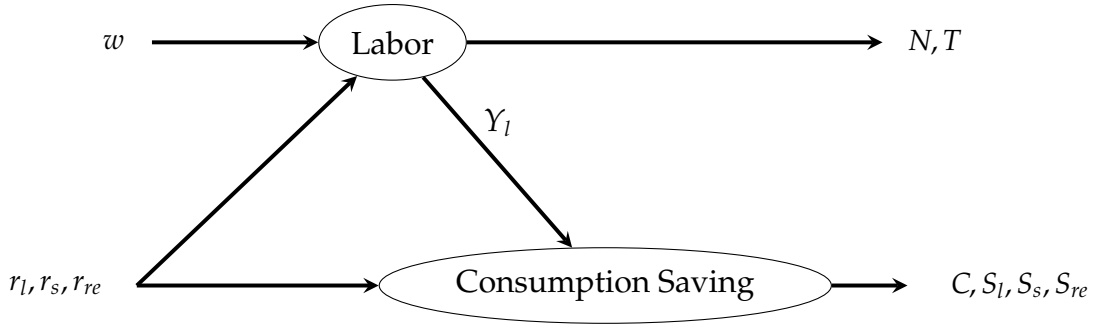


Figure 4: The household block separated into labor and consumption blocks

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 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 decisions, different aggregate shocks may be associated with different labor income distributions. As a consequence, the consumption block is not separable from the labor block through aggregate labor income alone. 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 a 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 heterogeneous 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 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

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

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, $\mathcal{J}^{C,Y_l}$, to the federal funds rate $\mathcal{J}^{C,R}$, as well as to stock and real estate returns, $\mathcal{J}^{C,\text{stocks}}$, and $\mathcal{J}^{C,\text{realestate}}$. 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. There is a fair amount of empirical evidence that I will use to fit this first column. There is 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:

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

The income process comes from Kaplan et al. (2018) along with 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 evidence from Fagereng et al. (2021) about the spending pattern over the 5 years following a shock to income.

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.

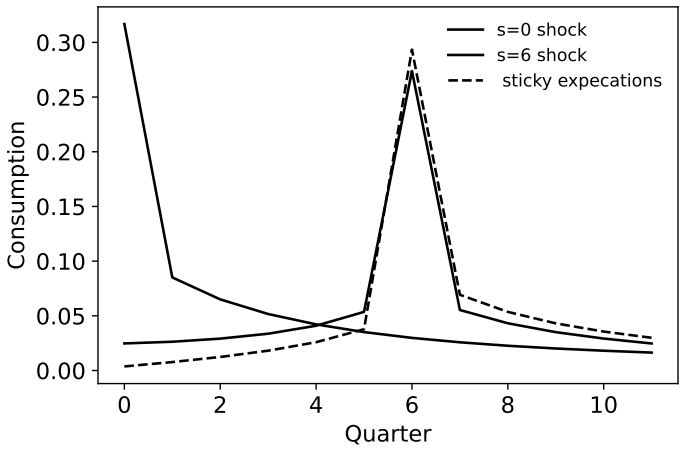


Figure 5: Columns of the consumption-to-income Jacobian, with and without sticky expectations

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.

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—and I use the one-asset model to determine the impulse response to news shocks. Encouragingly, Auclert et al. (2018) experiment with a variety of models and find that "models with very different primitives, once calibrated to the existing evidence on iMPCs out of unexpected income shocks, predict similar tent-shaped iMPCs out of expected income shocks. Given the lack of good empirical evidence on these iMPCs, this is reassuring."

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 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 the possibility of less anticipation, 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 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.[3]

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 uncer-

[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.

[3] An alternative to sticky expectations is finite horizon planning—I explore such a possibility in section 5.2. There is some evidence that the iMPC has a fat tail, as result of many two-asset models, but this feature of the iMPC is not material to my analysis—see section 5.3.

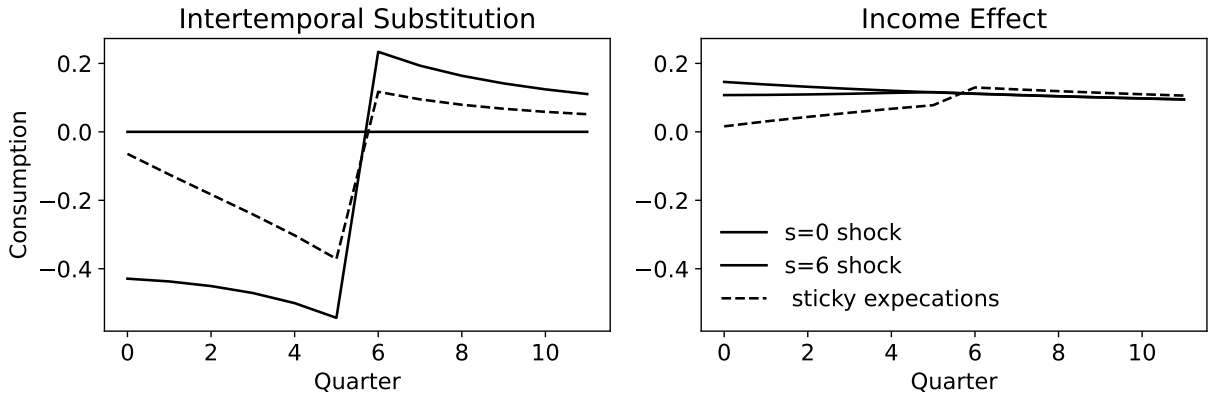


Figure 6: Columns of the consumption-to-interest-rate Jacobian, divided into the substitution and income effect, with and without sticky expectations

tainty. 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 in the budget constraint 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 by 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 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 shape of 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. However, Crawley and Kuchler (2023) show that the income effect from assets and liabilities is small in the US. Accordingly, 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.

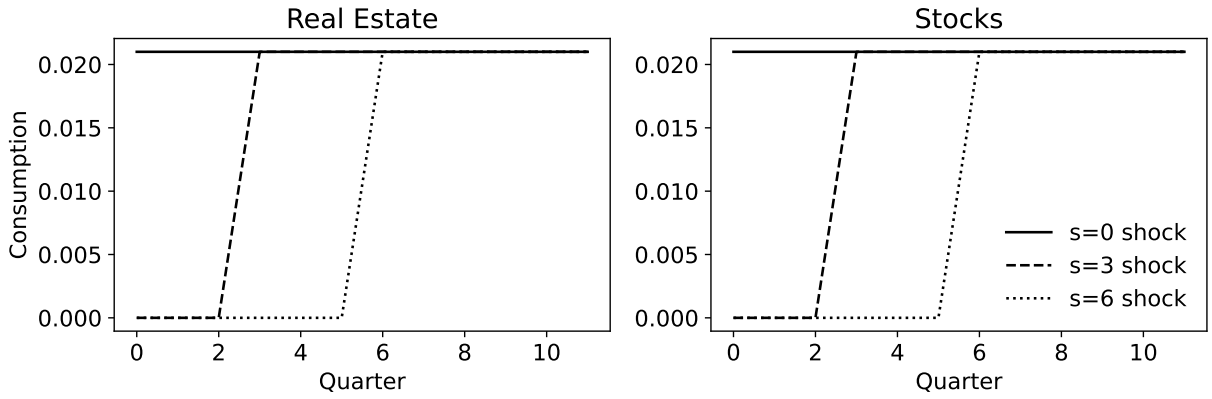


Figure 7: Columns of the consumption-to-stock-market and consumption-to-real-estate Jacobians

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 investments.[4] This measure of income is

[4]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.

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 $J_h^{O,q}$, 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_t^q at time t. Following Ramey (2016), I include in the set of controls X_t^q two lags each of log industrial production, the unemployment rate, log of the consumer price index, log of a commodity price index, and the federal funds rate. I also include 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.[5] Following the advice in Montiel Olea and Plagborg-Møller (2021), I include six lags of the shock series and calculate standard errors for $\widehat{J_h^{O,q}}$ without Newey and West (1987) adjustment.

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

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

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 choose all 10 of the shock series for which Jorda projections are plotted in the figures contained 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 such as high-frequency identification, Cholesky decomposition, maximum forecast error variance, and narrative methods. The aim of choosing a broad range of shock types and identification methods is to cover my bases: if the reader has a dislike for some particular shock-identification methodologies, I hope she is able to find others to her taste. Ultimately, all the shock series used here point in the same direction of little intertemporal substitution.

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 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

The fiscal shocks series I examine include both shocks to government spending and shocks to anticipated taxes.

3. Ramey and Zubairy (2018) military news. This method identifies changes in the expected present value of government purchases, caused by military events. Specifically, the method involves reading *BusinessWeek* for such spending events in an effort to capture the news of the event, rather than relying on spending data that may have already been anticipated.

4. Ben Zeev and Pappa (2017) defense spending shocks. This paper identifies news about future defense spending following the methodology of Barsky and Sims (2011).

These shocks are those that best explain future movements in defence spending over a five years time horizon, and that are orthogonal to current defence spending.

5. Blanchard and Perotti (2002) government spending. Government spending shocks are identified via a Cholesky decomposition of a VAR in which government spending is ordered first.

6. Mertens and Ravn (2011) tax news shocks. The shocks in the paper build upon the tax shock series from Romer and Romer (2010) by dividing that shock series into anticipated and unanticipated shocks based on the delay between the passing of the legislation and the implementation of the legislation. In turn, Romer and Romer (2010) use a narrative approach to identify tax shocks, looking at presidential speeches and congressional reports.

7. Leeper et al. (2012) expected taxes from one to five years forward. This measure of expected tax changes is based on the spread between federal and municipal bonds. The insight here is that, if asset markets are efficient, the spread between tax-exempt municipal bonds and treasury bonds will reflect anticipated changes to taxes.

Technology Shocks

8. Ben Zeev and Khan (2015) investment-specific (IST) news shocks. Following Greenwood et al. (2000), this paper identifies investment-specific technology as the inverse of the real price of investment. Then, using an adapted version of the maximum forecast error variance identification approach of Barsky and Sims (2011), IST news shock are identified by finding the linear combination of reduced-form innovations that are orthogonal to both current TFP and current IST that maximizes the sum of contributions to IST forecast error variance over a finite horizon.

9. Fernald (2014) utilization-adjusted TFP. This paper begins by measuring the Solow residual using inputs, including capital, to create a TFP series. This series is then

adjusted to account for variations in factor utilization following the method in Basu et al. (2006). The shock series is calculated as the quarterly changes to this utilization-adjusted TFP series.

10. Francis et al. (2014) unanticipated TFP shocks. This paper identifies technology shocks by maximizing the contribution to the forecast-error variance of labor productivity at a long horizon.

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:

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

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

Set of Structural Shocks (Q)	$\hat{\theta}_{sub}$	$\hat{\theta}_{inc}$
All	-0.02	0.22
	(0.04)	(0.17)
Romer and Romer (2004)	-0.80	0.76
	(0.79)	(1.19)
Gertler and Karadi (2015)	-0.13	0.00
	(1.17)	(0.25)
Ramey and Zubairy (2018) military news	-0.03	0.13
	(0.17)	(0.37)
Ben Zeev and Pappa (2017) defence spending shocks	-0.01	0.11
	(0.10)	(0.43)
Blanchard and Perotti (2002) government spending	-0.07	0.81
	(0.31)	(2.33)
Mertens and Ravn (2011) tax news shocks	0.05	0.01
	(0.13)	(0.96)
Leeper et al. (2012) expected taxes from one to five years forward	-0.00	1.00
	(0.10)	(5.77)
Ben Zeev and Khan (2015) investment specific news shocks	-0.14	0.09
	(0.18)	(0.13)
Fernald (2014) utilization-adjusted TFP	-0.36	0.00
	(1.19)	(0.48)
Francis et al. (2014) unanticipated TFP shocks	-0.24	0.78
	(3.92)	(1.00)

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

minimizes the sum of losses over all structural shocks in Q.

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

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 1 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.

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 (2011), Leeper et al. (2012), and Ben Zeev and Khan (2015) are all below 0.15, while others such as those from Fernald (2014) 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 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 0.2. That is, each quarter, about 20 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 and relatively large standard errors 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 on the loss value.

In the rest of this results section, I will examine the results in the context of the two shock series in detail: those from Romer and Romer (2004) and Ben Zeev and Pappa (2017). I will then summarize the decomposition of 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. 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.[6] This estimated impulse response is close to the empirical impulse response, although it doesn't capture the full extents of 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.

[6]This estimation restricts $\hat{\theta}_{sub}$ to be positive, which in practice means it is equal to zero. $\hat{\theta}_{inc}$ is estimated to be 0.22.

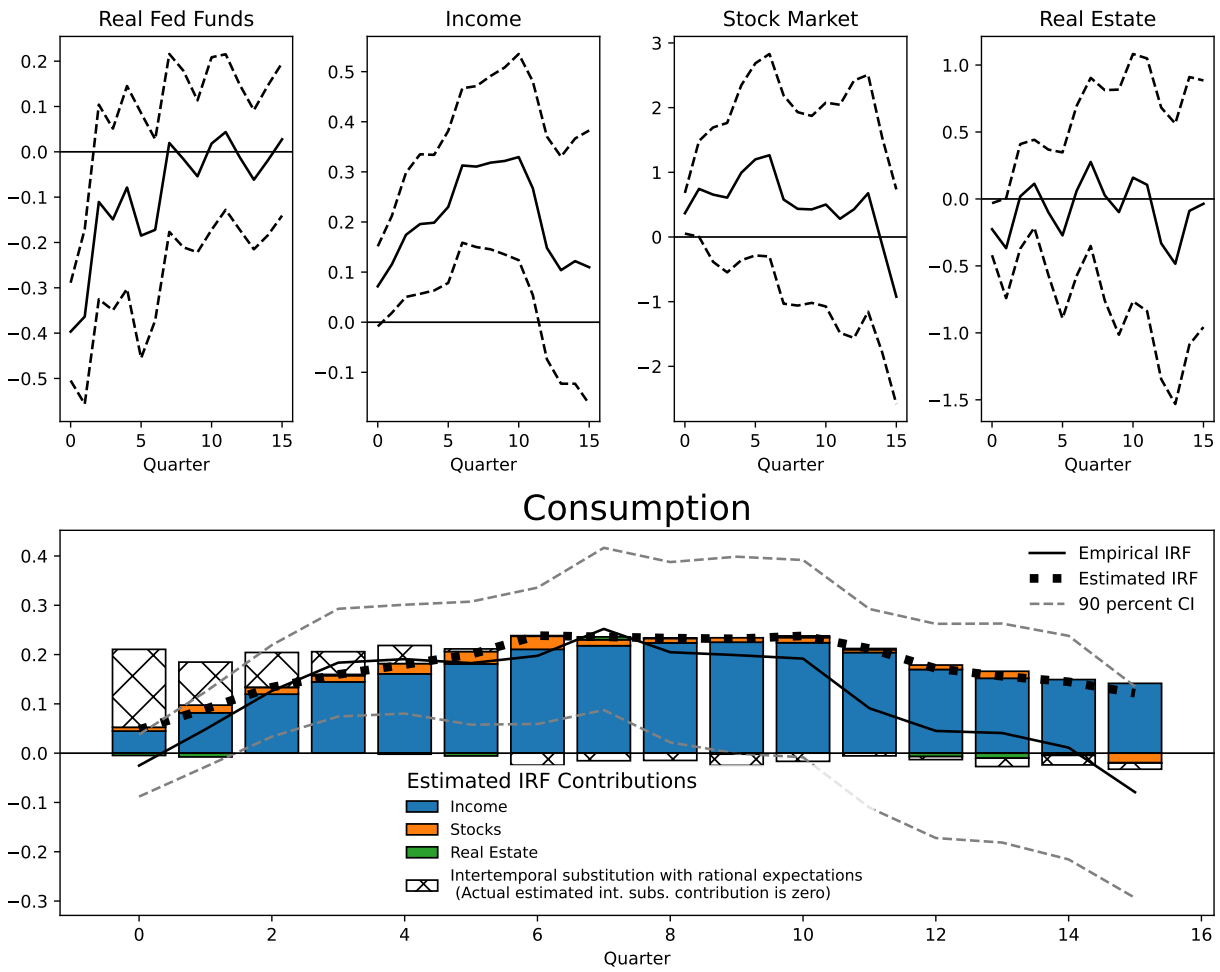


Figure 8: Inputs and output IRFs to the consumption block following a Romer and Romer (2004) shock

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 not significantly different from zero, equivalent to households paying no attention to the real fed funds rate when making their consumption—saving decisions.

Overall, the Romer and Romer (2004) shocks provide some evidence that households do not substitute intertemporally. However, the point estimate of $\hat{\theta}_{sub}$ using this data is large and negative (-0.8) with very large standard errors (0.76). These large standard errors result from the fact the shock is relatively short lived—households that are slow to react will end up not reacting at all to these shocks even if they would eventually change their consumption behavior substantially in reaction to a more persistent shock. Furthermore, identification of monetary policy shocks (as with all macro shocks) is challenging and Romer and Romer (2004) is only one way of doing so.

One of the contributions of the method in this paper over others that require full specification of a general equilibrium model, for example Auclert et al. (2020), is that it allows many other structural shocks to be used for estimation. This can help reduce standard errors around estimates and increase the robustness of the results to one particular shock identification methodology. Next, I describe the results using the defence spending shocks from Ben Zeev and Pappa (2017) that overcome some of the limitations of the Romer and Romer (2004) shock series.

Results using Ben Zeev and Pappa (2017) Defense Spending Shocks. In contrast to the short-lived change in the real fed funds rate following a Romer and Romer (2004) shock, the change in the real fed funds rate following a defense spending shock as identified by Ben Zeev and Pappa (2017) is persistent. The top-left panel of 9 shows that the real fed funds rate impulse response function remains negative for 4 years following the shock. The remaining three panels in the top row of figure 9 show the impulse response functions for income, the stock market, and real estate prices. Income increases gradually, reaching

a maximum 2 to 3 years after the shock, while the stock market shows no significant change and real estate prices increase.

The lower panel of figure 9 shows the decomposition of the impulse response of consumption to a defense spending shock. The solid black line shows the empirical response as estimated by a Jorda projection. The dotted black line shows the consumption IRF calculated from each of the four inputs multiplied by their respective Jacobians, where the Jacobian parameters have been estimated using all 10 shocks and intertemporal substitution restricted to be positive (and hence estimated to be zero). The blue, orange, and green bars show the decomposition of this IRF into contributions from income, stocks, and real estate.

The hashed white bars in the lower panel of 9 show the contribution to the consumption IRF that would be implied by the path of the real fed funds rate in the absence of sticky expectations. Under this shock, these intertemporal substitution effects are large and positive, and would push the Jacobian-implied consumption IRF well outside of the confidence intervals of the empirical consumption IRF. As such, it is clear that the Jacobian-implied IRF fits that data far better when there is no intertemporal substitution. This shock—assuming it is well identified—provides strong evidence that households do little in the way of intertemporal substitution.

Results for All 10 Structural Shocks. Figure 10 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.[7] The top-left panel of figure 10 is a repeat of the lower panel in figure 8, except that the contributions from stocks and real estate have been combined. Similarly, the second-row right panel is a repeat of figure 9. The other eight 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 10.

[7] Appendix A shows the input IRFs (income, real federal funds, stocks, and real estate) as well as the consumption decomposition shown in figure 10 for all 10 structural shocks.

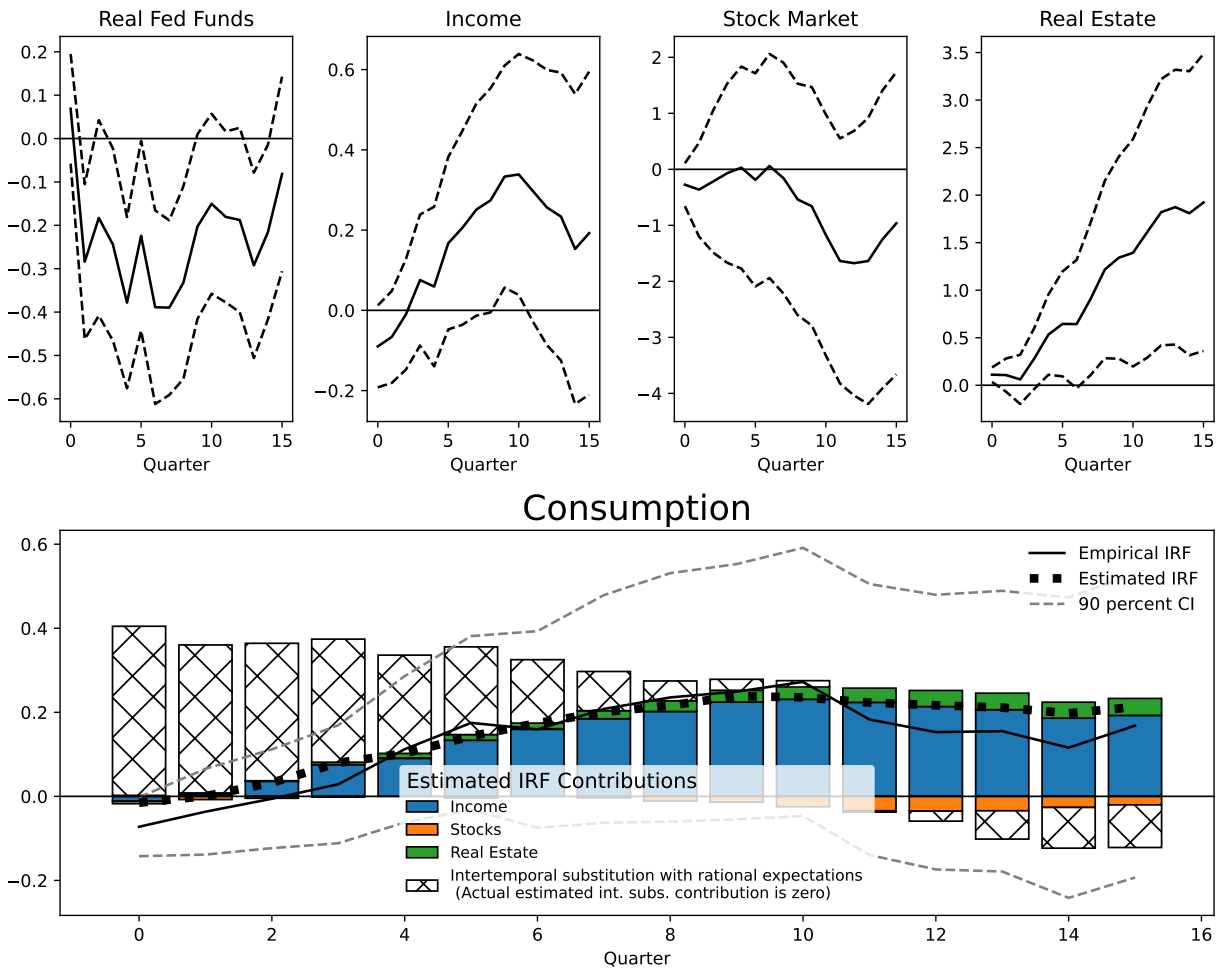


Figure 9: Inputs and output IRFs to the consumption block following a Ben Zeev and Pappa (2017) shock

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

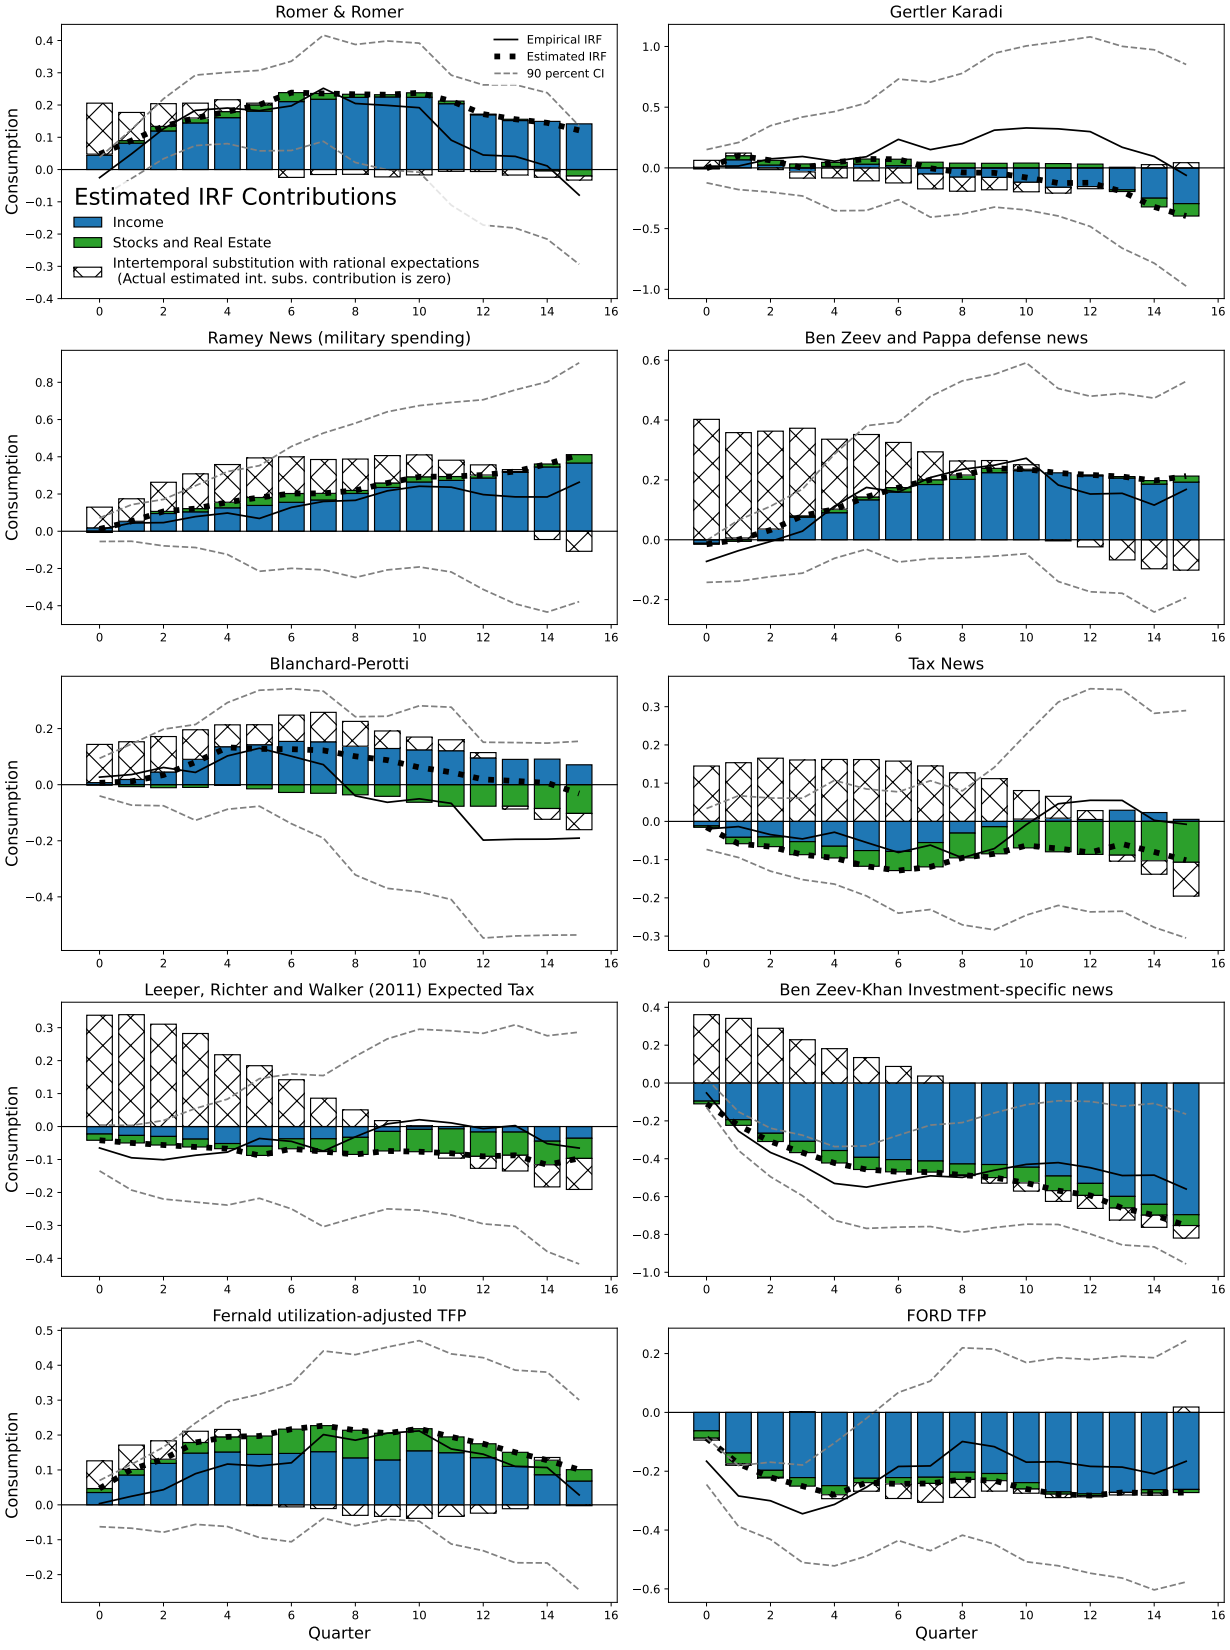


Figure 10: The consumption IRFs for all 10 shocks

MPC—the fact that these price changes play such a small role in the decomposition shows that the main results of this paper are 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 with an intertemporal elasticity equal to 1, 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" shock series does a little bit of intertemporal substitution—sticky expectations parameter of 0.05—even help move the estimated IRF toward the empirical IRF. For the other nine 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 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

For the baseline estimation, I took the consumption—income Jacobian from a one-asset buffer-stock model with an elasticity of substitution equal to 1 as the starting point. I

Set of Structural Shocks (Q)	EIS=1		EIS=0.5		EIS=0.33		EIS=0.25	
All	**-0.04**	(0.05)	**-0.10**	(0.09)	**-0.20**	(0.27)	**-0.33**	(0.46)
Romer and Romer (2004)	**-0.80**	(0.79)	**-1.00**	(1.61)	**-1.00**	(2.44)	**-1.00**	(3.16)
Gertler and Karadi (2015)	**-0.13**	(1.17)	**-0.38**	(7.73)	**-0.42**	(12.68)	**-0.42**	(17.18)
Ramey and Zubairy (2018)	**0.05**	(0.19)	**-0.68**	(2.03)	**-0.55**	(2.63)	**-0.52**	(3.16)
Ben Zeev and Pappa (2017)	**-0.01**	(0.10)	**-0.02**	(0.20)	**-0.04**	(0.34)	**-0.07**	(0.47)
Mertens and Ravn (2011)	**0.05**	(0.13)	**0.14**	(0.41)	**0.43**	(1.48)	**0.52**	(2.17)
Leeper et al. (2012)	**-0.00**	(0.10)	**-0.01**	(0.20)	**-0.02**	(0.32)	**-0.04**	(0.45)
Ben Zeev and Khan (2015)	**-0.14**	(0.18)	**-0.17**	(0.40)	**-0.17**	(0.60)	**-0.19**	(0.81)
Fernald (2014)	**-0.36**	(1.19)	**-0.75**	(3.24)	**-0.96**	(5.10)	**-1.00**	(6.57)
Fernald (2014)	**-0.68**	(0.95)	**-1.00**	(2.09)	**-1.00**	(3.14)	**-1.00**	(4.03)
Francis et al. (2014)	**-0.24**	(3.92)	**-0.30**	(13.38)	**-0.28**	(17.09)	**-0.24**	(15.89)

Table 2: Estimates of $\hat{\theta}_{sub}$ starting with different values for the elasticity of intertemporal substitution using all structural shocks together and each individually.

then estimated the sticky-expectations parameter which best fit the impulse responses. I claim that this methodology spans the space of most reasonable-looking iMPCs and intertemporal substitution Jacobians. In this robustness section, I push that claim by considering three alternative parameterizations for these Jacobians.

5.1 Robustness to the model's elasticity of substitution

The one-asset model that I begin with in the main analysis uses an elasticity of substitution (EIS) equal to 1. I then estimate a sticky expectations parameter of zero, implying that households do not in practice substitute intertemporally. Here, I test whether that conclusion is robust to starting with a lower EIS in the initial model.

Table 2 shows estimates for $\hat{\theta}_{sub}$ starting with different values for the EIS using all structural shocks together and each individually. The EIS=1 column replicates the estimates in the baseline analysis. The estimates using EIS values less than one are also statistically insignificant with nine out of ten estimates negative in each column. However, relative to the baseline, the magnitude of the point estimates tend to be larger—a consequence of the fact that the size of the consumption response is approximately proportional to the EIS multiplied by the sticky expectations parameter when the sticky expectations parameter is

small. Overall, these results are confirm a negligible response of consumption to interest rates.

5.2 Finite horizon planning in place of sticky expectations

In the main analysis, I have chosen to examine a sticky expectations version of a standard buffer-stock model in order to allow for the hump-shaped response to shocks that shocks that are often observed in the data. A popular alternative deviation from rational expectations is a model of finite horizon planning (FHP). Under finite horizon planning, agents have fully rational expectations about the state of the economy in the current period, but their expectations about deviations from the steady state become increasingly biased towards zero the further into the future these expectations relate to. For example, with an FHP parameter of 0.5, agents' expectations about deviations from the steady state two quarters from now are one quarter (0.5^2) of what a rational agent would expect. Finite horizon planning is one way for a model to avoid the forward guidance puzzle. However, unlike a sticky expectations model, finite horizon planning does not lead to hump shaped impulse responses.

Table 3 shows the estimates for the FHP parameters associated with the intertemporal substitution Jacobian ($\hat{FHP}_{sub}$) and the income Jacobian ($\hat{FHP}_{inc}$). As with the sticky expectations estimates, all but one of the FHP estimates suggest negative intertemporal substitution, and none of the 10 individual-shock estimates, nor the estimate that uses all 10 shocks together, are statistically significant. The point estimate using all 10 structural shocks is -0.45, which translates into only a small (and opposite-signed) intertemporal substitution response.

5.3 Fatter tails and lower intial MPCs in the intertemporal MPC

One concern is that the one-asset buffer-stock model that I start with has some known problems, and, in particular, the way I have calibrated it to match the one-year MPC for

Set of Structural Shocks (Q)	$\hat{FHP}_{sub}$	$\hat{FHP}_{inc}$
All	-0.45	0.87
	(0.25)	(0.12)
Romer and Romer (2004)	-0.56	0.80
	(0.67)	(0.49)
Gertler and Karadi (2015)	-0.62	0.14
	(1.23)	(8.61)
Ramey and Zubairy (2018) military news	-0.26	0.83
	(2.03)	(0.51)
Ben Zeev and Pappa (2017) defence spending shocks	-0.31	0.83
	(0.80)	(0.58)
Blanchard and Perotti (2002) government spending	-0.65	0.99
	(0.76)	(0.28)
Mertens and Ravn (2011) tax news shocks	0.64	-0.08
	(0.38)	(11.35)
Leeper et al. (2012) expected taxes from one to five years forward	-0.33	0.79
	(0.83)	(3.46)
Ben Zeev and Khan (2015) investment specific news shocks	-0.46	0.89
	(0.49)	(0.13)
Fernald (2014) utilization-adjusted TFP	-0.70	-0.16
	(0.81)	(11.20)
Francis et al. (2014) unanticipated TFP shocks	-0.75	0.99
	(1.53)	(0.10)

Table 3: Finite horizon planning parameter estimates using all structural shocks together and each individually.

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

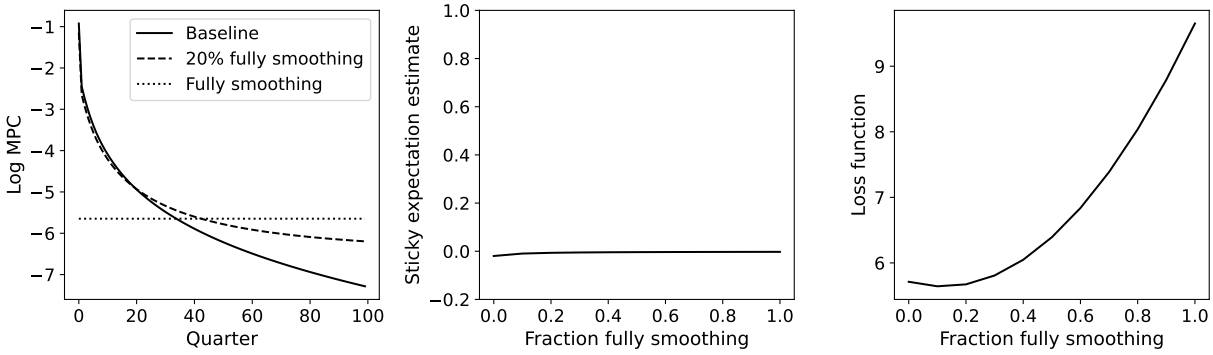


Figure 11: Log MPC out up to 40 years (left-hand panel), sticky-expectations parameter estimate (middle panel), and loss function (right-hand panel) for different fractions of fully-smoothing agents in the model.

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 11 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 iMPC of a model in which all agents fully smooth consumption. 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 11 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 11, increasing the fraction of fully-smoothing agents beyond around 0.2 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.

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

Akerlof, George A. and Janet L. Yellen, "A Near-Rational Model of the Business Cycle, With Wage and Price Inertia," *The Quarterly Journal of Economics*, 1985, *100*, 823–838.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 2021, *89* (5), 2375–2408.

_ , Matthew Rognlie, and Ludwig Straub, "The Intertemporal Keynesian Cross," Working Paper 25020, National Bureau of Economic Research September 2018.

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

Barnichon, Régis and Geert Mesters, "A Sufficient Statistics Approach for Macro Policy," *American Economic Review*, November 2023, *113* (11), 2809–45.

Barsky, Robert B. and Eric R. Sims, "News shocks and business cycles," *Journal of Monetary Economics*, 2011, *58* (3), 273–289.

Basu, Susanto, John G. Fernald, and Miles S. Kimball, "Are Technology Improvements Contractionary?," *American Economic Review*, December 2006, *96* (5), 1418–1448.

Ben Zeev, Nadav and Evi Pappa, "Chronicle of a War Foretold: The Macroeconomic Effects of Anticipated Defence Spending Shocks," *Economic Journal*, August 2017, *127* (603), 1568–1597.

_ and Hashmat Khan, "Investment-Specific News Shocks and U.S. Business Cycles," *Journal of Money, Credit and Banking*, 2015, *47* (7), 1443–1464.

Beraja, Martin, "A Semistructural Methodology for Policy Counterfactuals," *Journal of Political Economy*, 2023, *131* (1), 190–201.

Best, Michael Carlos, James S Cloyne, Ethan Ilzetzki, and Henrik J Kleven, "Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches," *The Review of Economic Studies*, 05 2019, *87* (2), 656–690.

Blanchard, Olivier and Roberto Perotti, "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output," *The Quarterly Journal of Economics*, 2002, *117* (4), 1329–1368.

Browning, Martin and Annamaria Lusardi, "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature*, 1996, *34* (4), 1797–1855.

Carroll, Christopher D, "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.

Carroll, Christopher D., Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White, "Sticky Expectations and Consumption Dynamics," *American Economic Journal: Macroeconomics*, July 2020, *12* (3), 40–76.

Choi, James J., "Popular Personal Financial Advice versus the Professors," *Journal of Economic Perspectives*, November 2022, *36* (4), 167–92.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Chapter 2 Monetary policy shocks: What have we learned and to what end?," in "Handbook of Macroeconomics," Vol. 1, Elsevier, 1999, pp. 65–148.

Cochrane, John H., "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *The American Economic Review*, 1989, *79* (3), 319–337.

Crawley, Edmund and Andreas Kuchler, "Consumption Heterogeneity: Micro Drivers and Macro Implications," *American Economic Journal: Macroeconomics*, January 2023, *15* (1), 314–41.

Crump, Richard K., Stefano Eusepi, Andrea Tambalotti, and Giorgio Topa, "Subjective Intertemporal Substitution," *Journal of Monetary Economics*, 2022, *126*, 118–133.

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik, "MPC Heterogeneity and Household Balance Sheets," *American Economic Journal: Macroeconomics*, October 2021, *13* (4), 1–54.

Farhi, Emmanuel, Alan Olivi, and Iván Werning, "Price Theory for Incomplete Markets," Working Paper 30037, National Bureau of Economic Research May 2022.

Fernald, John G., "A quarterly, utilization-adjusted series on total factor productivity," Working Paper Series 2012-19, Federal Reserve Bank of San Francisco April 2014.

Francis, Neville, 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*, 2014, *96* (4), 638–647.

Galí, Jordi, *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.

Ganong, Peter and Pascal Noel, "Consumer Spending during Unemployment: Positive and Normative Implications," *American Economic Review*, July 2019, *109* (7), 2383–2424.

Gertler, Mark and Peter Karadi, "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, January 2015, *7* (1), 44–76.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, "The role of investment-specific technological change in the business cycle," *European Economic Review*, 2000, *44* (1), 91–115.

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

Havranek, Tomas and Anna Sokolova, "Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say "probably not"," *Review of Economic Dynamics*, 2020, *35*, 97–122.

Hebden, James 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.

Holm, Martin Blomhoff, Pascal Paul, and Andreas Tischbirek, "The Transmission of Monetary Policy under the Microscope," *Journal of Political Economy*, 2021, *129* (10), 2861–2904.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante, "Monetary Policy According to HANK," *American Economic Review*, March 2018, *108* (3), 697–743.

Kueng, Lorenz, "Excess Sensitivity of High-Income Consumers*," *The Quarterly Journal of Economics*, 06 2018, *133* (4), 1693–1751.

Leeper, Eric M., Alexander W. Richter, and Todd B. Walker, "Quantitative Effects of Fiscal Foresight," *American Economic Journal: Economic Policy*, May 2012, *4* (2), 115–44.

McKay, Alisdair and Christian Wolf, "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?," *Econometrica*, Forthcoming.

Mertens, Karel and Morten O. Ravn, "Understanding the aggregate effects of anticipated and unanticipated tax policy shocks," *Review of Economic Dynamics*, 2011, *14* (1), 27–54. Special issue: Sources of Business Cycles.

_ **and** _ , "Empirical Evidence on the Aggregate Effects of Anticipated and Unanticipated US Tax Policy Shocks," *American Economic Journal: Economic Policy*, May 2012, *4* (2), 145–81.

Newey, Whitney K. and Kenneth D. West, "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 1987, *55* (3), 703–708.

Olea, José Luis Montiel and Mikkel Plagborg-Møller, "Local Projection Inference Is Simpler and More Robust Than You Think," *Econometrica*, 2021, *89* (4), 1789–1823.

Patterson, Christina, "The Matching Multiplier and the Amplification of Recessions," *American Economic Review*, April 2023, *113* (4), 982–1012.

Ramey, V.A., "Chapter 2 - Macroeconomic Shocks and Their Propagation," in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics*, Vol. 2, Elsevier, 2016, pp. 71–162.

Ramey, Valerie A. and Sarah Zubairy, "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 2018, *126* (2), 850–901.

Romer, Christina D. and David H. Romer, "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, September 2004, *94* (4), 1055–1084.

_ **and** _ , "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, June 2010, *100* (3), 763–801.

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

A Input and output IRFs for the 10 structural shocks.

Figures A.1 through A.10 show the four input IRFs along with the consumption IRF deconmposed into its various components for all 10 shocks used in the paper. The Jacobians used in each decomposition are estimated using all 10 shocks with intertemporal substitution restricted to be positive (and in practice equal to zero).

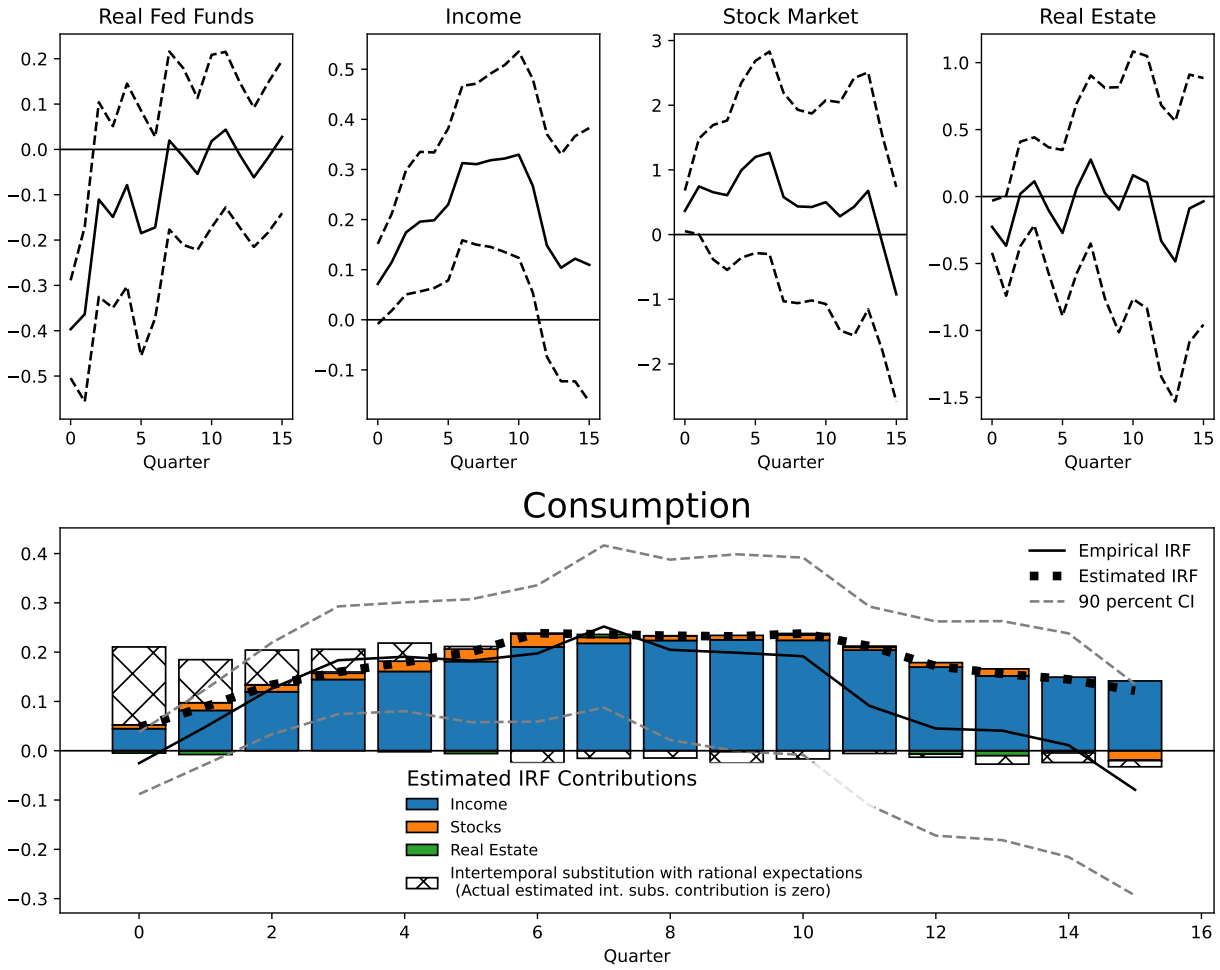


Figure A.1: Inputs and output IRFs to the consumption block following Romer and Romer (2004) monetary policy shock

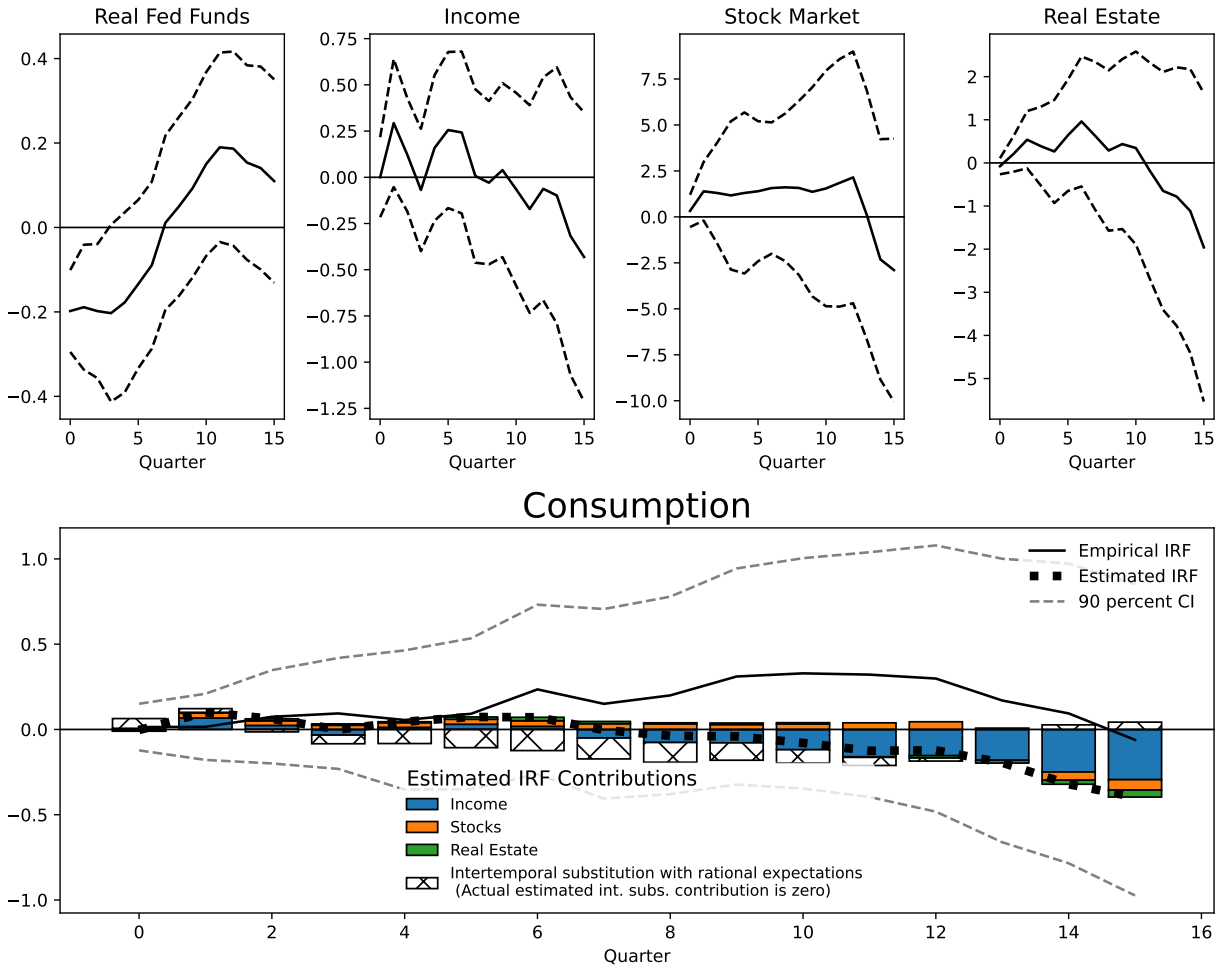


Figure A.2: Inputs and output IRFs to the consumption block following Gertler and Karadi (2015) monetary policy shock

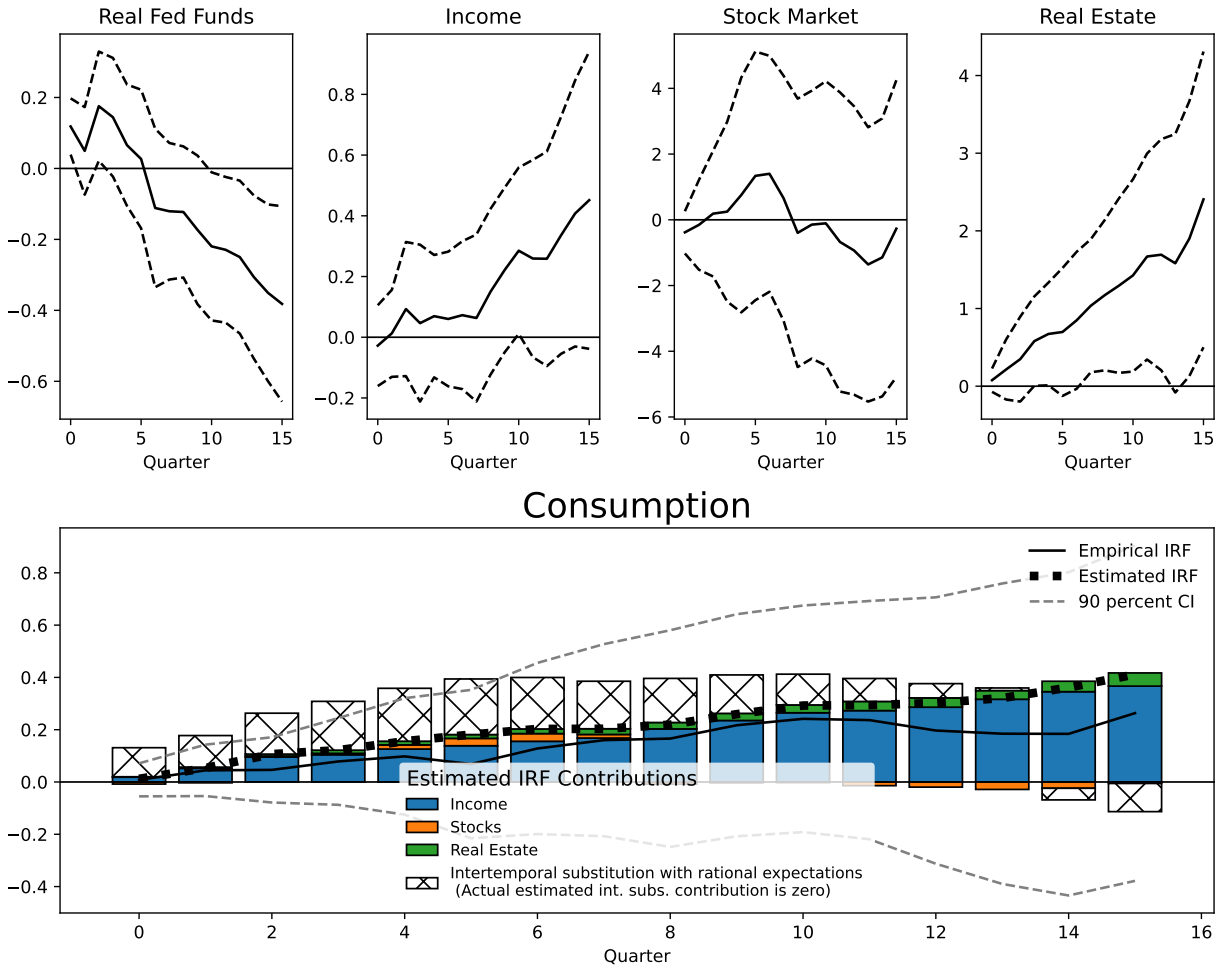


Figure A.3: Inputs and output IRFs to the consumption block following Ramey and Zubairy (2018) military news shock

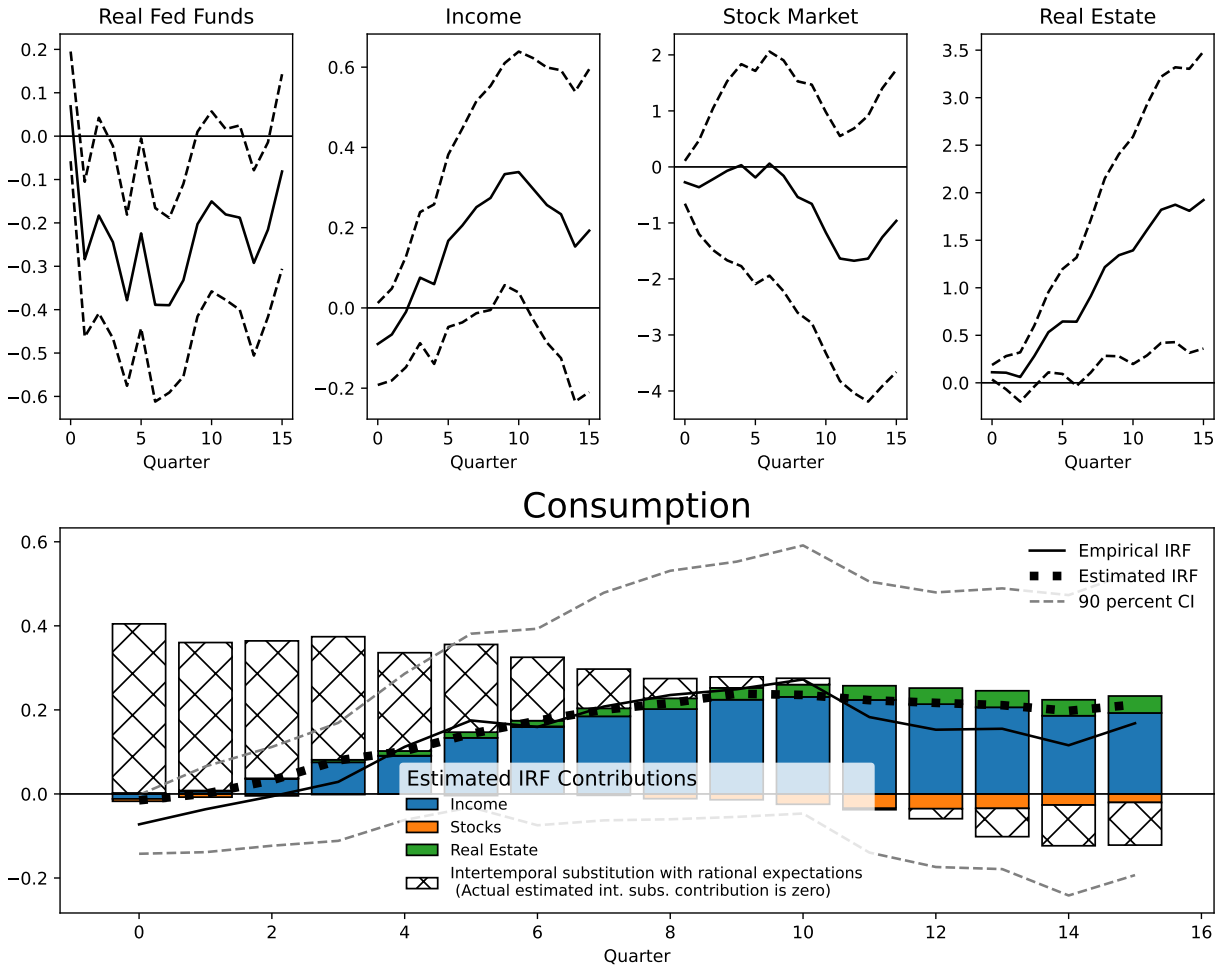


Figure A.4: Inputs and output IRFs to the consumption block following Ben Zeev and Pappa (2017) defence spending shock

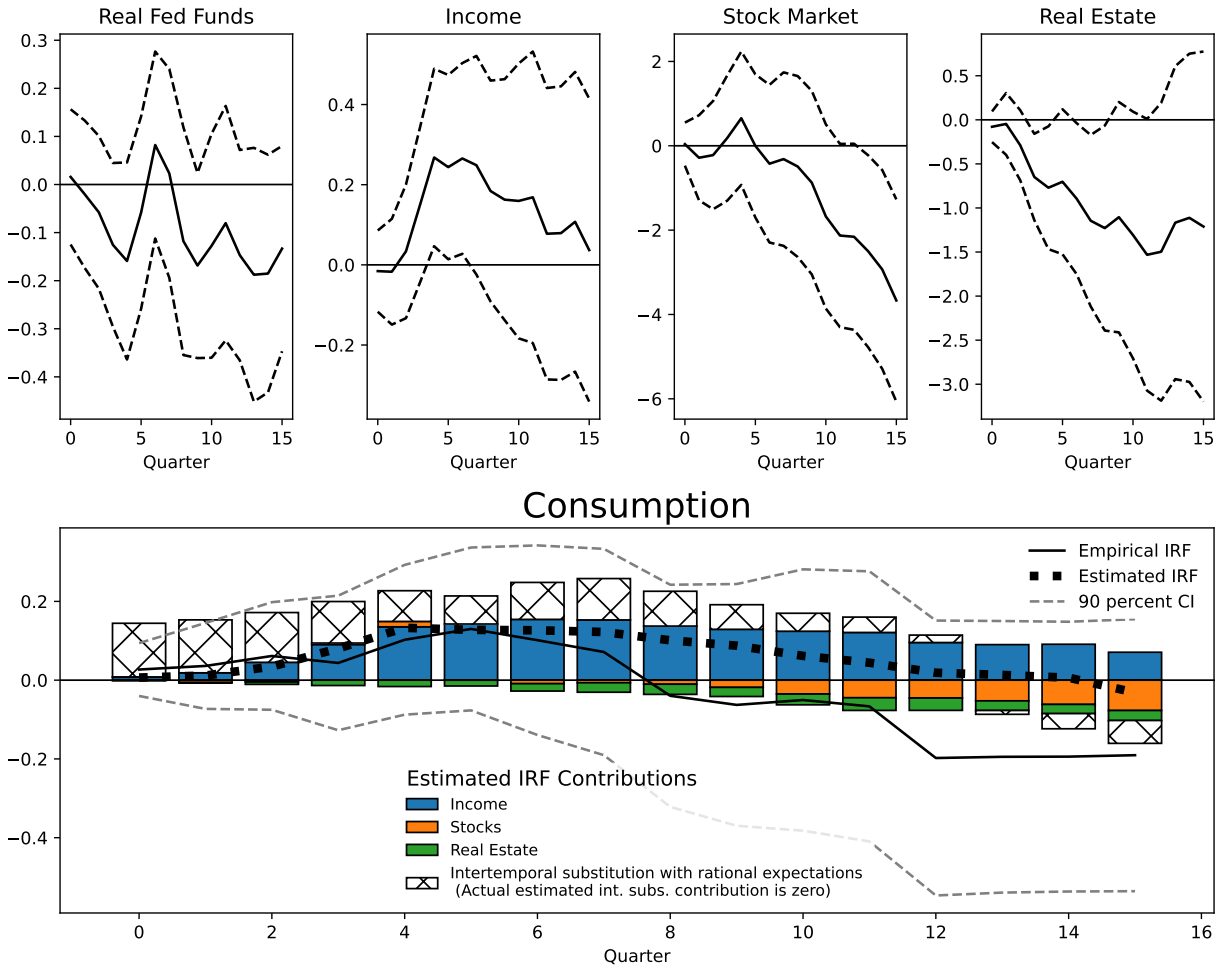


Figure A.5: Inputs and output IRFs to the consumption block following Blanchard and Perotti (2002) government spending shock

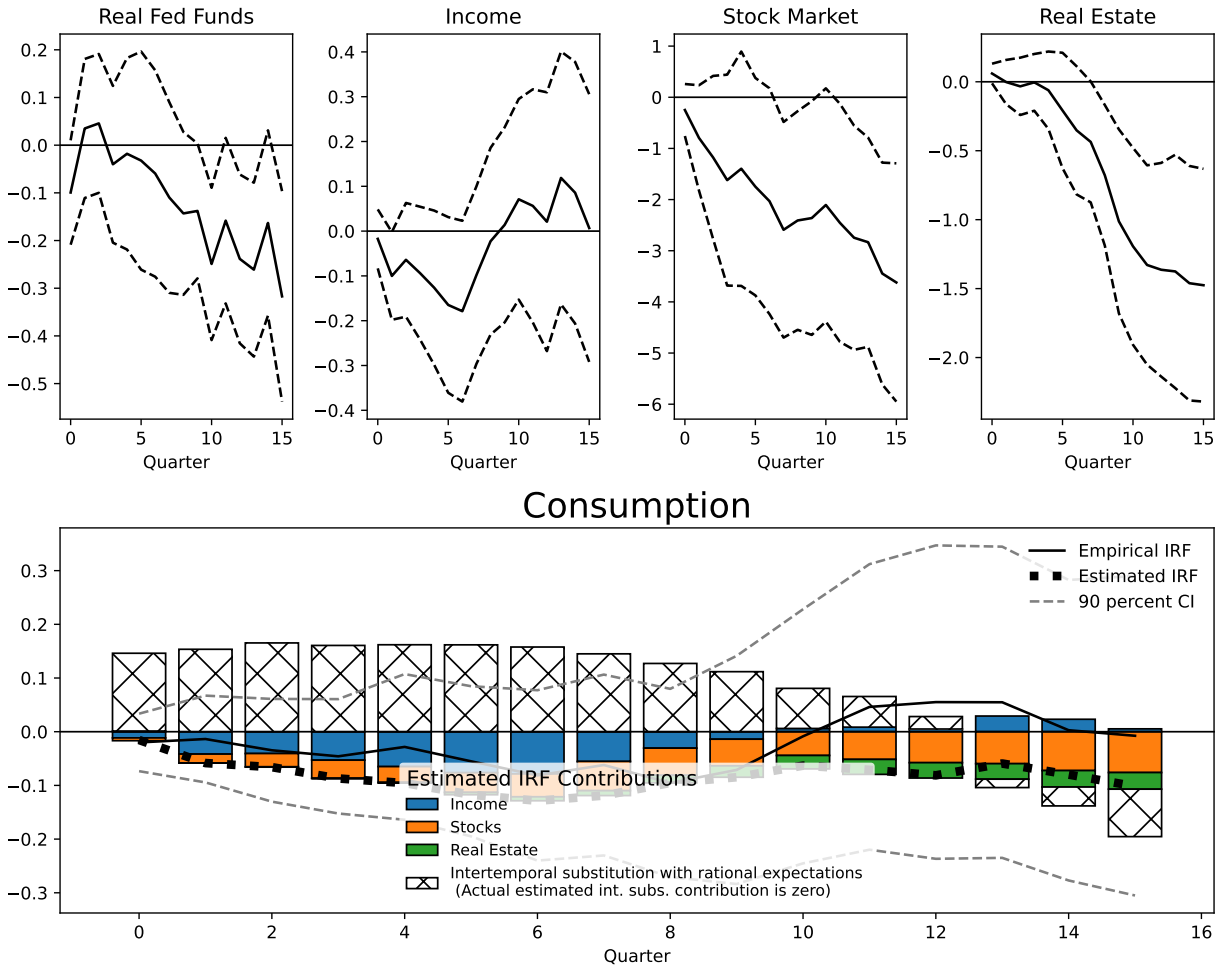


Figure A.6: Inputs and output IRFs to the consumption block following Mertens and Ravn (2012) tax news shock

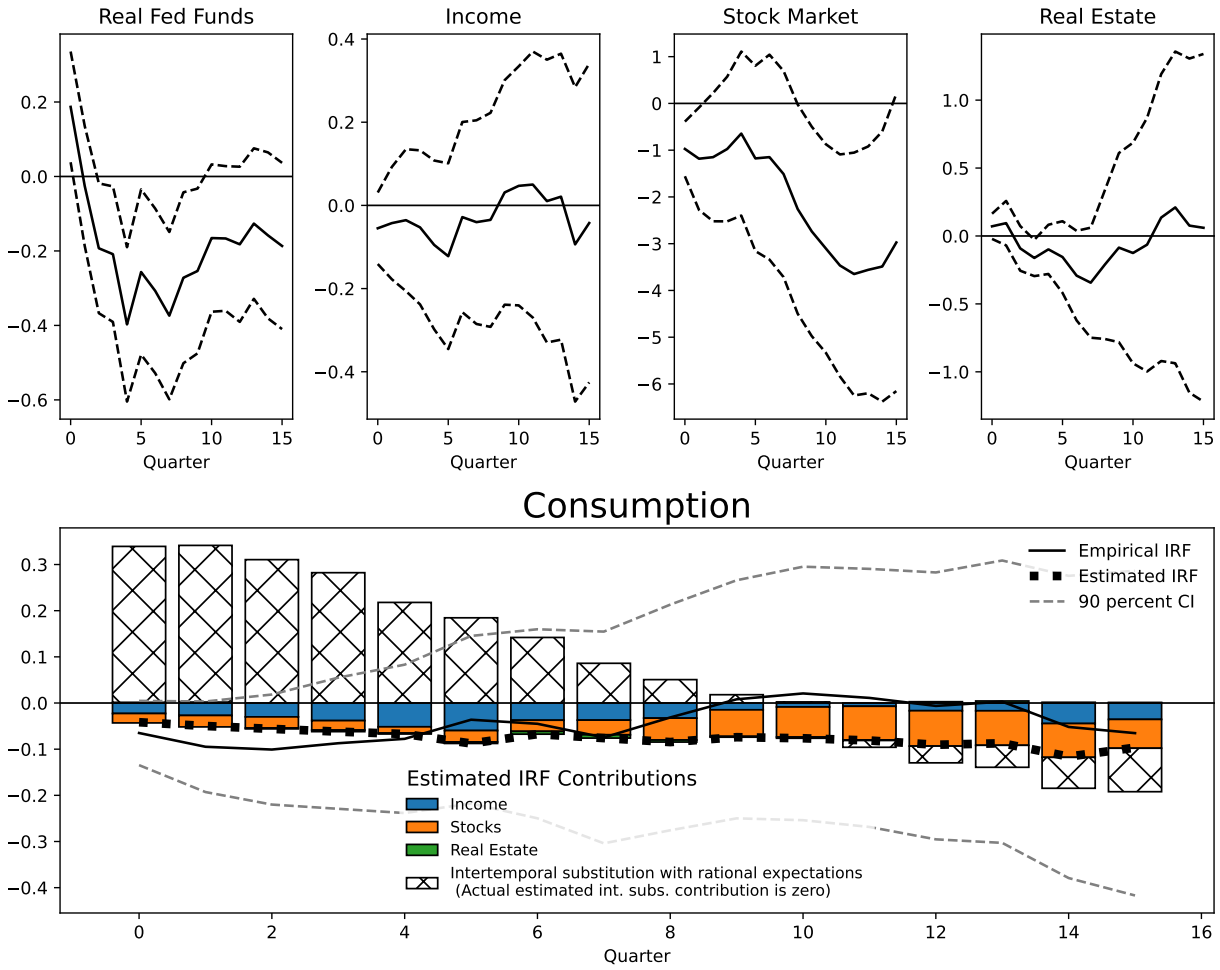


Figure A.7: Inputs and output IRFs to the consumption block following Leeper et al. (2012) expected taxes shock

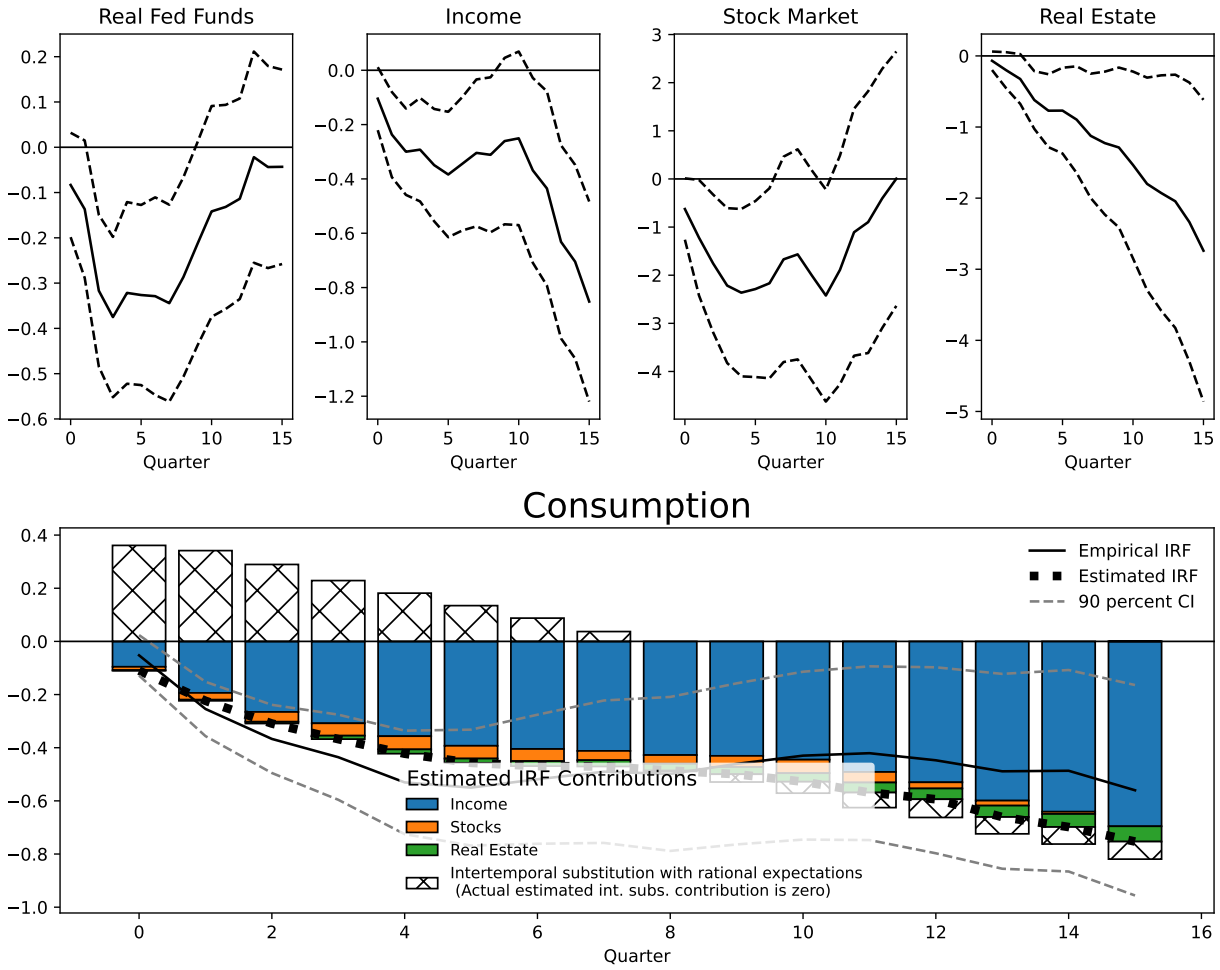


Figure A.8: Inputs and output IRFs to the consumption block following Ben Zeev and Khan (2015) investment specific news shock

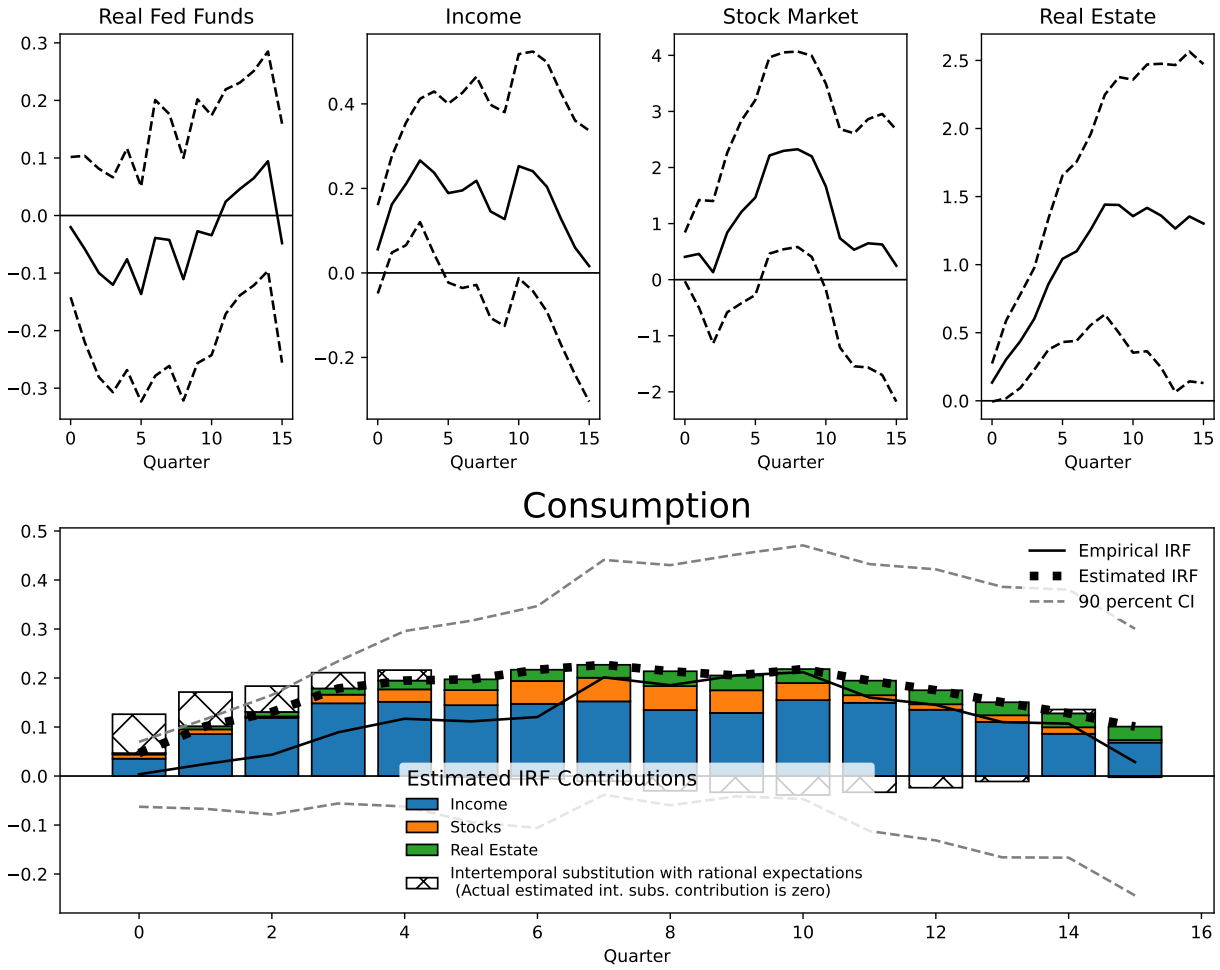


Figure A.9: Inputs and output IRFs to the consumption block following Fernald (2014) utilization-adjusted TFP shock

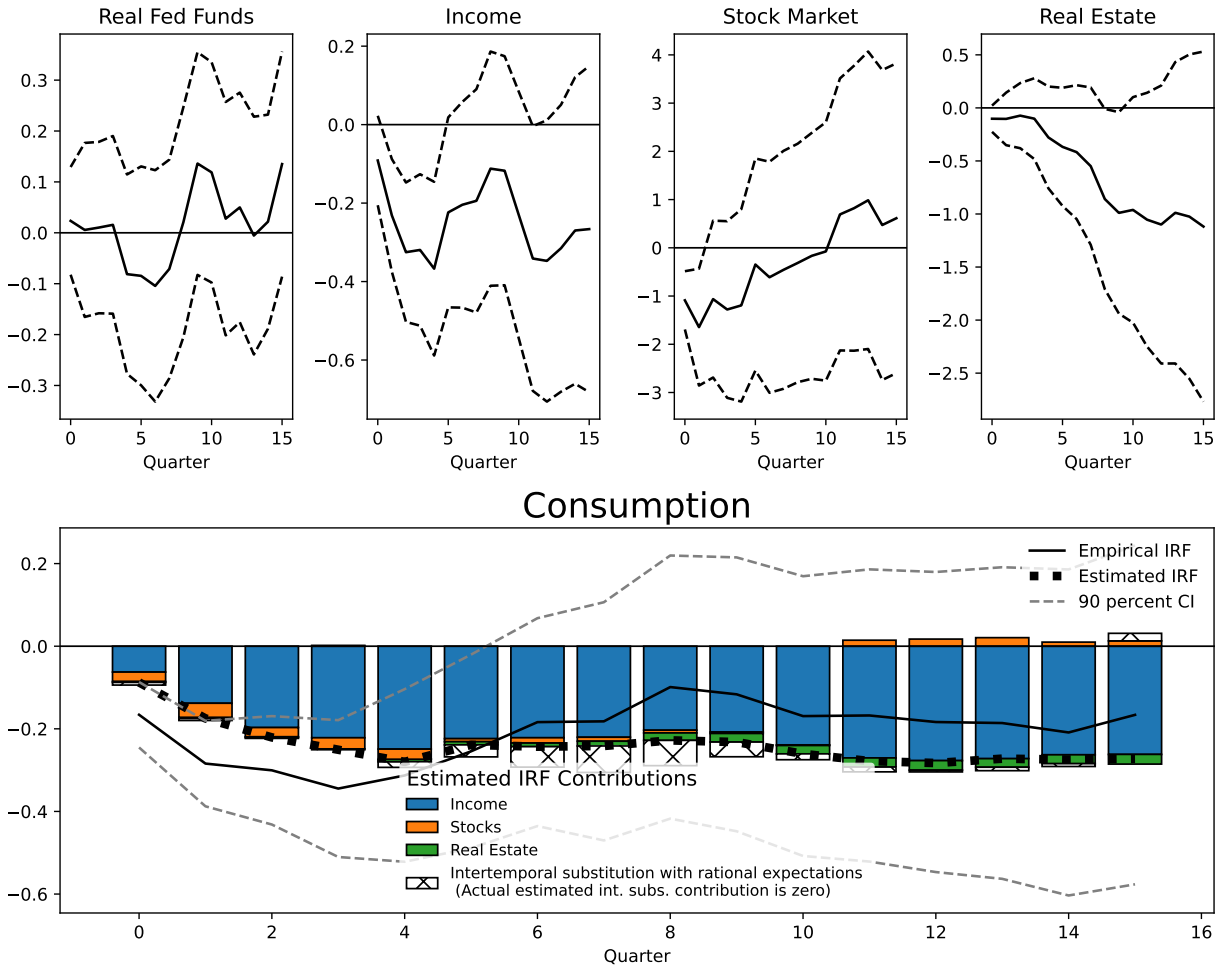


Figure A.10: Inputs and output IRFs to the consumption block following Francis et al. (2014) unanticipated TFP shock