Household Consumption Does Not Respond Directly to Interest Rates: Evidence From 10 Macroeconomic Shocks

Edmund Crawley and William L Gamber*

Federal Reserve Board

May 27, 2026

Abstract

We estimate how much household spending responds directly to changes in interest rates. We develop a Bayesian procedure that uses the empirical impulse responses to macroeconomic shocks to discipline the consumer block of a HANK model. The procedure can be applied shock-by-shock or pooled jointly. We apply this method in two ways using 10 macroeconomic shocks: a structural model with sticky expectations over both income and interest rates, and a non-parametric estimation of the consumption-to-interest-rate Jacobian. We find no evidence that households respond directly to interest rates at any horizon, leaving essentially no role for a direct interest rate channel.

*Contact information: edmund.s.crawley@frb.gov and will.gamber@frb.gov. Note that this paper supersedes a previously circulated paper by Edmund Crawley titled "Do Households Substitute Intertemporally? 10 Structural Shocks That Suggest Not." The views expressed in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or any other person associated with the Federal Reserve System.

1. Introduction

How does monetary policy affect household consumption? Economists' understanding of the answer has evolved considerably over the past two decades. In the representative household New Keynesian framework that dominated macroeconomics in the 1990s and 2000s, the consumption Euler equation provides the primary transmission mechanism: when interest rates rise, forward-looking households postpone consumption to take advantage of higher returns on savings. Under this framework, monetary policy operates primarily through the *direct channel*, in which interest rate changes directly influence households' intertemporal consumption decisions. More recent heterogeneous-agent New Keynesian (HANK) models have revised this view (Kaplan, Moll and Violante, 2018). In these models, because many households face borrowing constraints, they cannot easily adjust consumption in response to interest rate changes. So, in HANK models, the *indirect channel* becomes more important: monetary policy affects consumption by changing labor income, as interest rate movements influence firms' production and hiring decisions, which in turn affect household spending. The relative importance of these direct and indirect channels determines how monetary policy affects the economy and bears on how we design macroeconomic models.

In this paper, we measure the strength of the direct interest rate channel. We develop an estimation approach that incorporates both microdata evidence on the response of spending to income receipt, as well as macroeconomic dynamics following a wide range of identified macroeconomic shocks. We find that household consumption does not respond directly to interest rates at all. We thus find that there is no role for the direct channel of monetary policy; instead, our estimates imply that monetary policy affects household spending entirely through its effects on labor income. When the Fed raises interest rates, consumption falls because higher rates reduce employment and wages. We find that even the standard full-information rational expectations one-asset HANK model that we study overstates the size of the direct channel.

This finding contradicts standard macroeconomic models, in which the direct interest rate channel plays a central role, and it has direct implications for the forward guidance puzzle, the observation that announcements about future interest rates generate implausibly large immediate effects in standard models (Del Negro, Giannoni and Patterson, 2023; Carlstrom, Fuerst and Paustian, 2015). If households do not respond to future interest rates, this puzzle largely disappears.

We make two contributions. First, we develop a new methodology for estimating the direct household response to interest rates and income that uses information from both micro- and macro-data. We use the sequence-space representation of heterogeneous-agent models (Auclert et al., 2021) to express aggregate consumption as a function of two inputs: the path of interest rates and the path of labor income. The mapping from these inputs to consumption is summarized by two Jacobian matrices, which capture how consumption at each date responds to shocks to interest rates or income at any other date. We then estimate these Jacobians by matching impulse responses across many different types of externally identified shocks. Our second contribution is the substantive finding that the interest-rate-to-consumption Jacobian is close to 0.

We take two Bayesian approaches to estimating these Jacobians. In both approaches, we begin by estimating the response of consumption, interest rates, and labor income to externally identified shocks using linear projections. In the first approach, we extend a heterogeneous-agent consumption-savings model to allow for sticky expectations over both income and interest rates (Carroll et al., 2020; Auclert, Rognlie and Straub, 2020). After calibrating the model to match microdata evidence on intertemporal marginal propensities to consume from Fagereng, Holm and Natvik (2021), we then estimate how much expectation stickiness is needed to match the macroeconomic impulse responses. More formally, we choose estimates of the sticky expectations parameters that are most likely given the observed impulse responses. Our second approach is non-parametric: We estimate the consumption-to-interest-rate Jacobian flexibly using basis functions, imposing only the minimal smoothness properties shared by many structural models.

Our second contribution is the substantive finding that household consumption does not respond directly to interest rates. We establish this result by implementing our procedure using ten structural macroeconomic shocks spanning monetary, fiscal, and technology shocks, identified using high-frequency methods, narrative approaches, and VAR techniques drawn from Ramey (2016), a review of shock identification in the Handbook of Macroeconomics. Both of our estimation approaches deliver the same answer. In our structural estimation, the posterior distribution for the sticky expectations parameter on interest rates collapses tightly onto the value implying households essentially never update their expectations about interest rates. This finding holds whether we estimate shock-by-shock or jointly. As we show, the standard full-information rational expectations HANK model cannot rationalize the empirical impulse responses that we estimate. In our non-parametric estimation, the posterior mean Jacobian is near zero at all horizons, confirming that household consumption does not respond directly to interest rates regardless of the functional form we impose.

Our results build on three literatures. First, a long line of research has sought to measure the intertemporal elasticity of substitution—an important component of the sensitivity of consumption to interest rates—using micro and macro data. Early work using aggregate Euler equations (Hall, 1988) found very low elasticities, but subsequent studies produced conflicting estimates and Carroll (1997) raised serious econometric concerns about the Euler equation approach. A meta-analysis of these and more recent studies finds that, after correcting for publication bias, estimates of the elasticity of intertemporal substitution are low (Havránek, 2015). Our approach circumvents the problems of the Euler equation approach by using multiple identified shocks to discipline the entire dynamic response of consumption to interest rate changes, rather than estimating a single elasticity parameter from Euler equation regressions. Holm, Paul and Tischbirek (2021) also circumvents these issues by controlling for household-level income in regressions of household consumption on monetary policy shocks using Norwegian microdata.

Second, the heterogeneous-agent New Keynesian (HANK) literature has shown that borrowing constraints and heterogeneity in marginal propensities to consume substantially dampen the direct interest rate channel (Kaplan, Moll and Violante, 2018). In these models, monetary policy affects consumption primarily through the indirect channel we emphasize. Our findings suggest that even HANK models calibrated to match microdata on intertemporal MPCs overstate the direct channel. Rationalizing the macro evidence requires a friction—extreme information stickiness, finite planning horizons, or something observationally equivalent—to weaken the direct interest rate channel.

Methodologically, our paper contributes to a growing literature that uses the sequence-space representation of macroeconomic models to connect structural models with empirical evidence (Auclert et al., 2021). This approach recognizes that in heterogeneous-agent models, the response of aggregate variables like consumption can be summarized by Jacobian matrices that map the paths of driving forces, such as interest rates and labor income, into outcomes. These Jacobians serve as sufficient statistics for how households respond to economic shocks, independent of what generated those shocks in the first place. Recent work has exploited this insight to perform structural inference that is robust to the Lucas critique (McKay and Wolf, 2023) and to characterize fiscal multipliers using intertemporal marginal propensities to consume as sufficient statistics (Auclert, Rognlie and Straub, 2024). This approach has also been used to develop semi-structural estimation frameworks (Hebden and Winkler, 2021; Beraja, 2023) and to evaluate optimal policies (Barnichon and Mesters, 2023). We extend this literature by using sequence-space Jacobians to estimate structural parameters of the household consumption block by matching impulse responses across many different types of identified shocks. Our non-parametric approach further demonstrates how this framework can flexibly estimate the consumption-to-interest-rate mapping without imposing strong functional form assumptions.

Our paper is related to Auclert, Rognlie and Straub (2020) and Bayer, Born and Luetticke (2024) which both estimate a full general equilibrium HANK model. Like Auclert, Rognlie and Straub (2020), we also estimate a sticky-expectations model to match both micro-level evidence, as well as monetary policy shock impulse responses. That said, our paper innovates in important ways. First, in our approach, we estimate only the household block of the model. This decision has several advantages, including that it does not require that we take a stand on the remaining components of the model, and our results are broadly applicable across a wide class of models. Second, this flexibility allows us to estimate our model on impulse responses estimated from 10 shocks, rather than just a single monetary policy shock. The fact that our findings are consistent across all 10 of these shocks should obviate concerns that there is something peculiar about any single shock that spuriously generates our results. Rather, our findings appear to be consistent with macroeconomic dynamics following a broad array of shocks.

The paper proceeds as follows. Section 2 develops the theoretical framework and describes the structural model. Section 3 presents the econometric approach and data. Section 4 reports results from structural estimation. Section 5 presents the non-parametric approach and results. Section 6 concludes.

2. Theoretical framework

2.1. Household problem. The economy is populated by a unit mass of households. These households face idiosyncratic uncertainty over labor efficiency z, which evolves according to a Markov process with transition matrix P. Each household chooses its consumption c and savings a each period to maximize the present discounted value of future utility, subject to budget and borrowing constraints, taking as given aggregate post-tax labor income Y_t^L and the interest rate r_t. Following the standard practice in the sequence-space literature, households assign no uncertainty to the future paths of aggregate variables, though they do face idiosyncratic uncertainty over z. The individual states faced by each household are its beginning-of-period asset holdings a_- and productivity z. Idiosyncratic income is a deterministic function of the household's productivity and aggregate post-tax labor income $y(z, Y_t^L)$.[1] End-of-period assets are denoted by a.

Each household has a permanent type g, and there are μ_g households of type g. Households of each group have a different discount factor β_g but are otherwise ex-ante identical. This permanent heterogeneity in patience allows us to match the dynamic response of average household spending to income shocks.

Each household solves the following dynamic programming problem:

$$
\begin{aligned}
V_{g,t}(a_-, z) &= \max_{c,a} u(c) + \beta_g \mathbb{E}[V_{g,t+1}(a, z')] \\
\text{such that} \quad a + c &= y(z, Y_t^L) + (1 + r_t)a_- \\
a &\geq \underline{a}.
\end{aligned}
\tag{1}
$$

[1] For example, consider the standard New Keynesian union setup as in Auclert, Rognlie and Straub (2020), in which a union supplies aggregate hours N_t and allocates them equally across workers, so each worker supplies N_t hours regardless of $z_{i,t}$ at a uniform wage w_t per efficiency unit. If post-tax income is a time-invariant, strictly increasing function ϕ of pre-tax income, then $y_{i,t} = \phi(w_t z_{i,t} N_t)$. Integrating over the stationary cross-section of z gives $Y_t^L = G(w_t N_t)$, where $G(x) \equiv \int \phi(xz)\, dD(z)$ inherits strict monotonicity from ϕ and is therefore invertible. Hence $w_t N_t = G^{-1}(Y_t^L)$ and $y_{i,t} = \phi(z_{i,t}\, G^{-1}(Y_t^L))$, a function of $z_{i,t}$ and Y_t^L alone, satisfying the required condition. If instead the allocation of total hours across individuals depends on anything beyond $z_{i,t}$—for example, idiosyncratic wealth—the representation fails, because individual hours (and hence individual income) bring in additional state variables.

Let $c^*_{g,t}(z, a_-)$ and $a^*_{g,t}(z, a_-)$ denote the policy functions that solve the Bellman equation. Also denote by $D_{g,t}(z, A_-) \equiv Pr(z_t = z, a_{t-1} \in A_- \mid g)$ the distribution of households in group g with productivity state z that own assets in a set A_- at the start of date t. The distribution $D_{g,t}$ has law of motion

$$D_{g,t+1}(z', A) = \sum_z D_{g,t}(z, {a^*_{g,t}}^{-1}(z, A))P(z, z'), \tag{2}$$

where ${a^*_{g,t}}^{-1}(z, \cdot)$ is the inverse of $a^*_{g,t}(z, \cdot)$. Since the policy function $a^*_{g,t}$ depends only on β_g and $\{r_s, Y^L_s\}_{s\geq t}$, equation (2) implies that $D_{g,t}$ is determined recursively by the initial distribution $D_{g,0}$ and the paths $\{r_s, Y^L_s\}_{s\geq t}$.

Aggregate consumption is obtained by summing over types and integrating over individual states:

$$C_t = \sum_g \mu_g \sum_z \int_{a_-} c^*_{g,t}(z, a_-)\, D_{g,t}(z,\, da_-). \tag{3}$$

The key advantage of this sequence-space approach is that household consumption depends only on the expected paths of interest rates and labor income, not on the underlying economic structure that generates these paths. This "sufficiency result" allows us to estimate household behavior by matching consumption responses to these two inputs across multiple shocks, without requiring a fully specified general equilibrium model. We formalize this insight in Proposition 1.

Proposition 1 (Jacobian representation of household consumption)**.** *Suppose the economy is in steady state prior to date 0, with constant interest rate and labor income (r_{ss}, Y^L_{ss}) and corresponding stationary distributions $D_{g,ss}$ for each type g. At date 0, the future paths $\{r_s, Y^L_s\}_{s\geq 0}$ are revealed (an MIT shock—a perfect-foresight transition from steady state). Assume these paths are such that the household dynamic program admits a unique solution and the implied sequence of distributions $\{D_{g,t}\}_{t\geq 0}$ converges.Then:*

(i) Aggregate consumption at each date t is determined by the paths $\{r_s, Y^L_s\}_{s\geq 0}$ alone. That is, there exist functions $C_t(\{r_s, Y^L_s\}_{s\geq 0})$ such that equation (3) holds.

(ii) To a first-order approximation around steady state, the response of aggregate consumption to deviations of the interest rate and aggregate labor income from steady state is

$$\mathbf{dC} \approx \mathcal{J}^{ss}_{C,R}\, \mathbf{dr} + \mathcal{J}^{ss}_{C,Y^L}\, \mathbf{dY^L}, \tag{4}$$

where $\mathbf{dC} = (C_0 - C_{ss}, C_1 - C_{ss}, \ldots)'$, and $\mathcal{J}^{ss}_{C,R}$ and $\mathcal{J}^{ss}_{C,Y^L}$ are the Jacobians of the sequence of aggregate consumption with respect to the interest rate and aggregate labor income paths, evaluated at steady state.

Proof. Part (i): Since the economy begins in steady state, $D_{g,0} = D_{g,ss}$ for each g. As noted above, the policy function $c^*_{g,t}$ depends only on β_g and $\{r_s, Y^L_s\}_{s\geq t}$. By equation (2), $D_{g,t}$ for $t \geq 1$ is determined recursively by

$D_{g,ss}$ and $\{a^*_{g,s}\}_{s<t}$, which in turn depends only on the future paths for interest rates and labor income. Hence the integrand and the measures in (3) are determined by $\{r_s, Y^L_s\}_{s\geq 0}$, establishing the claim.

Part (ii): Part (i) establishes that C_t is a function of the infinite-dimensional price paths. Differentiating C_t with respect to these paths and evaluating at steady state yields the Jacobian matrices $\mathcal{J}^{ss}_{C,R}$ and $\mathcal{J}^{ss}_{C,Y^L}$. The first-order Taylor expansion around steady state then gives (4). □

Remark 1. *Part (ii) assumes that C_t is differentiable with respect to the interest and labor income paths at steady state. While individual policy functions have kinks at the borrowing constraint, aggregation across the distribution of households smooths these out, making this assumption reasonable for the aggregate economy.*

Remark 2. *The Jacobian matrices $\mathcal{J}^{ss}_{C,R}$ and $\mathcal{J}^{ss}_{C,Y^L}$ encode the complete dynamic response of aggregate consumption to interest rate and labor income shocks. Entry (t,s) of $\mathcal{J}^{ss}_{C,R}$, for example, gives the effect of date-0 news about a date-s change in the interest rate on date-t consumption, inclusive of the redistribution across households induced by changes in the distribution $D_{g,t}$.*

Remark 3. *Proposition 1 places no restriction on the long-run behavior of the input paths. In particular, it accommodates permanent changes in labor income (as arise, for example, from a permanent productivity shock) as well as persistent deviations of the interest rate from steady state. The Jacobian columns $\mathcal{J}^{ss}_{C,Y^L}[\cdot,s]$ and $\mathcal{J}^{ss}_{C,R}[\cdot,s]$ encode the response of date-t consumption to a unit perturbation of the corresponding input at date s, however far in the future, and the linear approximation in (4) is valid for any small bounded perturbation $\mathbf{dY^L}$, $\mathbf{dr}$.*

2.2. Closing the model. The implication of Proposition 1 is that the remaining details of the economy—i.e., how Y^L_t and r_t are determined—do not matter for computing the response of aggregate consumption to any given sequence of shocks, conditional on knowing $\{Y^L_t, r_t\}$. As such, we can embed the household problem specified previously into any general equilibrium framework, as long as it generates paths for $\{Y^L_t, r_t\}$.

2.3. Sticky expectations. We assume that households have sticky expectations as in Carroll et al. (2020). That is, each period, for each aggregate variable $X \in \{R, Y^L\}$, a fraction $1-\theta^X$ of households update to the true expected path of the aggregate state. We assume that the probability of updating, θ^X, is the same across household groups g and that the probability of updating the aggregate path X is independent over time, across households, and across aggregate variables X. Households form expectations about the future paths of income Y^L_t and the interest rate r_t separately, treating these variables as exogenous to their own decisions. We denote the sticky expectations parameters for these states by θ^Y and θ^R, respectively. Full details on how we implement sticky expectations are described in Appendix A.

We assume that, in the case of labor income, all households know the actual income they have already received. We make a different assumption for the interest rate—we instead assume that households update beliefs about their past asset returns only at the point that they update their expectations for the future path of interest rates. We make this choice to reflect the fact that the arrival of labor income is usually highly salient, while households may be much slower to look at the value of their pensions or other financial assets.

In a single-asset model with rational expectations, households react to changes in wealth exactly as they do to changes in income. In practice, this symmetry generates responses to changes in wealth that are too front-loaded and too large when compared with empirical evidence.[2] Our choice to limit households' ability to observe changes in their wealth brings the model responses in line with the data. One could think of this choice as a stand-in for adjustment costs or attention frictions that prevent households from accessing or observing changes in stock portfolio, pension, or real estate wealth.

To check that this assumption is not driving our key results, we also estimated the model under the alternative assumption that households immediately observe changes in their asset values (i.e., treating the period-0 portfolio return like labor income with respect to information). The headline result remains unchanged: households do not respond directly to interest rates. However, this alternative specification fits the empirical consumption impulse responses less well, producing counterfactually large date-0 consumption responses that are not present in the data. Lastly, the non-parametric exercise that we describe in section 5 does not impose any assumption about when households observe change in their wealth driven by interest rate changes.

2.4. Parameterization. We assume that households have log utility over consumption. The pre-tax income process follows a discretization of the jump-drift continuous process from Kaplan, Moll and Violante (2018).[3] The process is estimated to match eight moments of annual earnings changes from Social Security Administration data, estimated in Guvenen et al. (2021). Households face a progressive tax system in which a one percent increase in pre-tax income leads to a $(1-\tau)$ percent increase in after-tax income, where $\tau = 0.181$ is the tax progressivity parameter. This parameterization captures the compressive effect of progressive taxation on the earnings distribution.[4]

We assume that there are two groups of households, $g \in \{0, 1\}$. We calibrate β_0, β_1, and population share μ_0 to match (1) an aggregate wealth-to-income ratio of 6.6 and (2) the population-mean response of household spending to a one-time lottery win over 5 years from Fagereng, Holm and Natvik (2021). Figure

[2] Many papers estimate a wealth effect on consumption of around 3 cents on the dollar. (See, for example Chodorow-Reich, Nenov and Simsek, 2021) As Di Maggio, Kermani and Majlesi (2020) find, households respond to unrealized capital gains by much less than dividend flows, consistent with our assumption that households respond to changes in the flow of income by more than changes in asset values.

[3] See Appendix B for details.

[4] The value $\tau = 0.181$ is from Heathcote, Storesletten and Violante (2017).

Figure 1: Intertemporal MPC in model and data

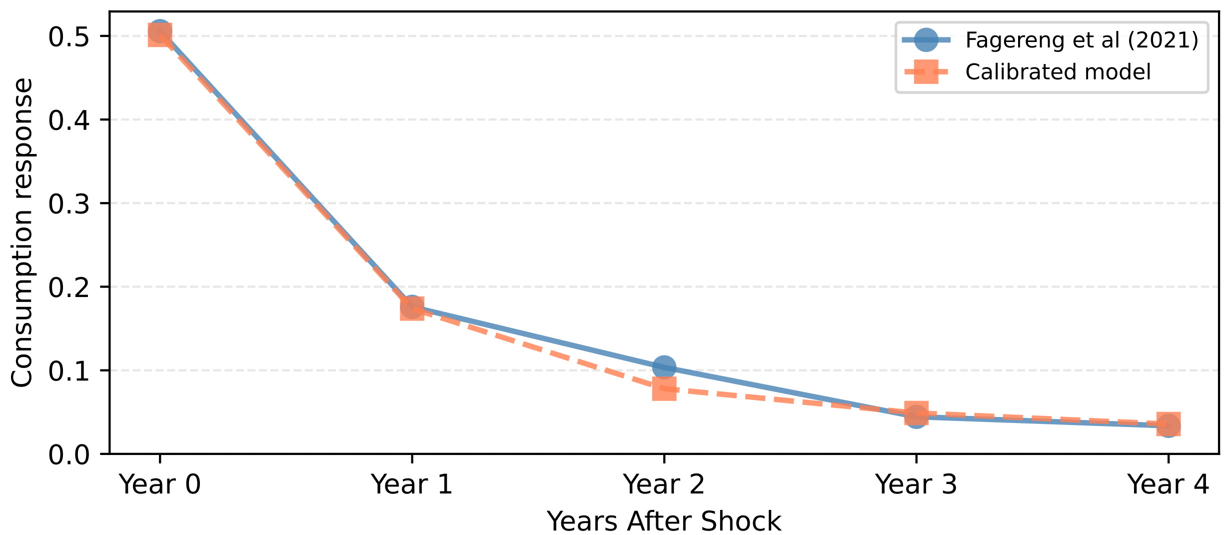

Notes: Figure depicts the population-mean consumption response to a one-time lump-sum transfer received at date 0, as a share of the transfer amount. Blue round dots and solid lines show the estimates from Fagereng, Holm and Natvik (2021), and orange squares and dashed lines show the estimates in the calibrated model.

1 shows the intertemporal MPC in the model and in the data. With just these three parameters, we are able to very closely match the target empirical intertemporal MPC path.

2.5. Jacobians. As shown in equation (4), the Jacobians of consumption with respect to labor income and the interest rate determine the response of consumption to aggregate shocks. These Jacobians are determined by the structural parameters of the model, including θ^Y and θ^R. Given the important role of these two parameters in shaping the Jacobians, we estimate these to match the dynamics of aggregate consumption following macroeconomic shocks.

Consumption-to-labor-income Jacobian We calibrated the model's discount factors and population shares to match the consumption response to lump-sum transfers from Fagereng, Holm and Natvik (2021) (as shown in Figure 1). Holding fixed the values of θ^Y and θ^R, these same structural parameters determine the consumption-to-labor-income Jacobian, also known as the intertemporal MPC to an aggregate labor income shock, which captures the response to a proportional income shock subject to progressive taxation.[5] The sticky expectations parameter θ^Y then modifies this Jacobian to account for gradual learning about future income paths.

Figure 2 shows several columns of this Jacobian. The solid line shows the first column which displays

[5]Note that the shock implies an equal percent change in gross income for all households. The progressive tax means that after-tax income will increase relatively more for low-income households.

Figure 2: Columns of the consumption-to-income Jacobian, with and without sticky expectations

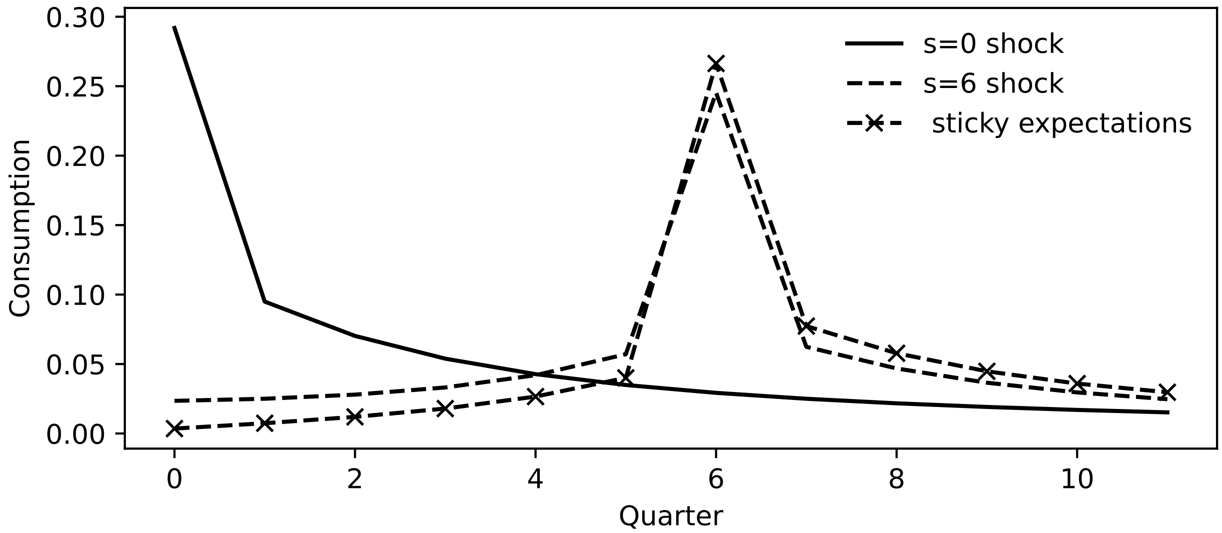

Notes: The figure shows the consumption response at different horizons (x-axis, in quarters) to aggregate labor income shocks arriving at different dates, expressed as a fraction of the income shock. The solid line shows the response to a shock arriving at date 0. The dashed lines show the response to a shock arriving in six quarters' time: the plain dashed line assumes full information rational expectations (FIRE), while the dashed line with crosses assumes sticky expectations with $\theta^Y = 0.85$ (15% of households update their expectations each quarter).

an MPC of around 0.3 in the first quarter, dropping rapidly to below 0.1 in the second quarter and declining more gradually thereafter. This first column is the same under full information rational expectations (FIRE) and sticky expectations. The dashed line shows the response to an income shock arriving in six quarters under FIRE. The dashed line with crosses shows the same column under sticky expectations, where only 15 percent of households update their expectations each quarter. As shown, sticky expectations reduce anticipatory consumption and increase the response once income arrives, consistent with evidence from Kueng (2018) and Ganong and Noel (2019) that households exhibit less anticipatory behavior than standard models predict.

Consumption-to-real-interest-rate Jacobian The response of consumption to a real interest rate shock can be decomposed into substitution and income effects.[6] The substitution effect captures changes in the timing of consumption driven by changes in the relative price of consumption across periods, holding the budget set fixed. The income effect captures changes in consumption driven by the wealth gains or losses households experience when interest rates change, holding intertemporal prices fixed. In our model, we isolate these effects by varying the interest rate that enters the Euler equation (substitution) separately from the interest rate that enters the budget constraint (income).

[6] Farhi, Olivi and Werning (2022) lays out the theory in the context of an incomplete market with uncertainty.

Figure 3: Columns of the consumption-to-interest-rate Jacobian, divided into the substitution and income effect, with and without sticky expectations

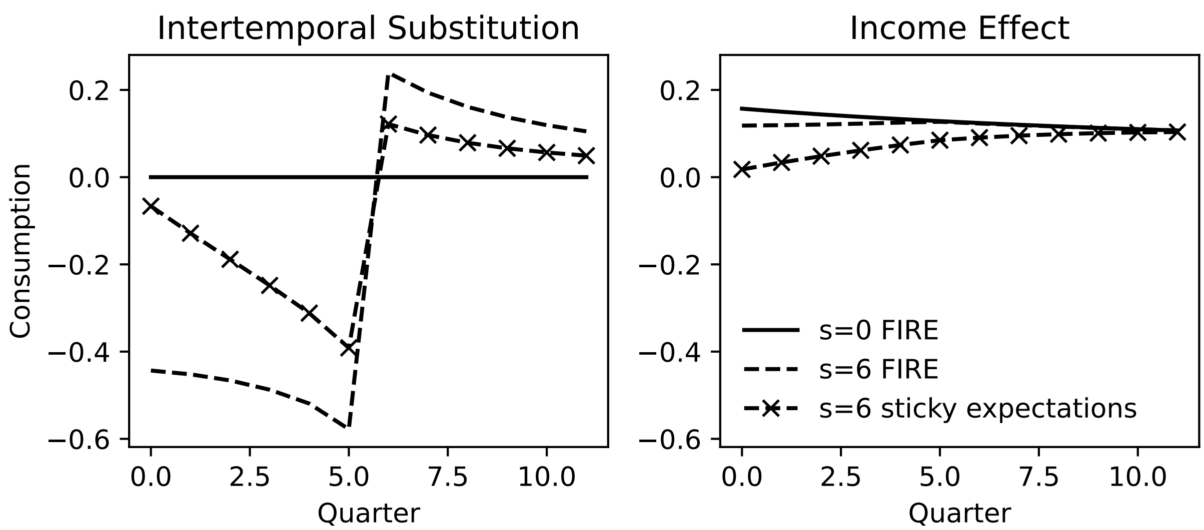

Notes: The figure shows the consumption response at different horizons to real interest rate shocks, decomposed into substitution effects (left panel) and income effects (right panel). Solid lines show the response to a shock arriving at date 0. Dashed lines show the response to a shock arriving in six quarters: the plain dashed line assumes FIRE, while the dashed line with crosses assumes sticky expectations with θ^R governing the fraction of households that do not update beliefs about future interest rates each period.

Figure 3 shows these two components separately; their sum is the full consumption-to-real-interest-rate Jacobian. The left panel displays the substitution-effect Jacobian, and the right panel displays the income-effect Jacobian. In each panel, solid lines show the response to a shock arriving at date 0, while dashed lines show the response to a shock arriving in six quarters. The dashed lines with crosses show the sticky expectations case where only $(1 - \theta^R)$ of households update their beliefs about future interest rates each period.

The substitution effect (left panel) exhibits the standard pattern: news of a future interest rate increase causes households to save more before the rate increase in order to consume more afterward. This response is budget-neutral by construction. Note that under our timing convention, there is no substitution effect for a shock at date 0, since households use the interest rate r_{t+1} to discount utility from period $t + 1$ to t.

The income effect (right panel) shows a positive consumption response to higher interest rates. News of a future interest rate increase (dashed line under FIRE) leads households to increase current consumption in anticipation of future wealth gains. Under sticky expectations (dashed line with crosses), this anticipatory response is dampened, with the consumption response building more gradually as households learn about the interest rate change.

A note on r_0: The term r_0 captures the immediate revaluation of household asset holdings when news about future interest rates arrives. In our model with no aggregate risk and a single asset, the asset can be valued using a discounted cash flow model. Following an MIT shock, the asset will revalue according to the change in its expected cash flows and the change in future interest rates. The return on the asset from this revaluation is equal to r_0, the return received on assets held going into the quarter that the news shock arrives. A shock to r_0 has an income effect but no substitution effect on consumption. The exact nature of the asset (short-term or long-term bonds, equities, or something else) will be determined by the rest of the model and does not need to be specified for our methodology, as long as the return can be measured. For a financial portfolio consisting of long-term assets, a news shock that raises the expected future path of interest rates without changing cash flows will lead to a negative r_0 and an expectation of higher future returns.

Comparison The response of consumption to future changes in income is quite different from its response to future changes in interest rates. Even with FIRE, current consumption moves very little in response to future changes in income, whereas current consumption responds strongly to anticipated changes in interest rates. This difference arises because different types of households drive each response. The income response is determined primarily by high-MPC households close to their borrowing constraint. Because these households cannot or do not wish to borrow against future income, they adjust their consumption little in anticipation of future income changes. The interest rate response, by contrast, is driven primarily by unconstrained households with significant assets. For these households, the Euler equation governs consumption decisions, creating a tight link between current consumption and expectations about future interest rates. As we show later in the paper, we can tightly identify θ^R as this parameter changes the shape of the interest-rate Jacobian significantly. By contrast, θ^Y will be less tightly identified in the data because variation in this parameter has less influence over the aggregate consumption response.

3. Structural estimation methodology

3.1. Methodology. We seek to estimate the parameters θ^Y and θ^R, which govern how much households respond to future changes in labor income and interest rates. Our approach is a Bayesian impulse-response matching estimator. It follows in the tradition of Christiano, Eichenbaum and Evans (2005) and Christiano, Trabandt and Walentin (2011), among others.[7]

Let $\theta = (\theta^Y, \theta^R)$ denote this parameter vector. The "data" we observe are estimated empirical impulse response functions (IRFs) for consumption ($\hat{\zeta}_h^{C,q}$), post-tax labor income ($\hat{\zeta}_h^{Y,q}$), and returns ($\hat{\zeta}_h^{R,q}$), for shock

[7] For more on limited-information and quasi-Bayesian inference from moment conditions, see Kim (2002); Chernozhukov and Hong (2003); Chib, Shin and Simoni (2018).

q, and horizon h, along with the covariance of these estimates $\hat{\Omega}$. We split the returns IRF into the period-0 asset return $\hat{\zeta}_0^{R,q}$ which we set to the return on the average financial portfolio in the U.S. in the period that the shock arrives, followed by $\hat{\zeta}_+^{R,q} = (\hat{\zeta}_1^{R,q}, \hat{\zeta}_2^{R,q}, \ldots, \hat{\zeta}_{H-1}^{R,q})'$, the impulse response of the real Federal Funds rate, consistent with the expected return on all assets being equal following an MIT shock. Stacking these impulse responses into a vector ζ, we assume that these are normally distributed around the truth:

$$\hat{\zeta} \sim \mathcal{N}(\zeta, \Omega). \tag{5}$$

Here ζ is the true vector of impulse responses and Ω is the covariance matrix for these estimates. We seek to use these data to characterize the distribution of θ. So, we evaluate the posterior distribution:

$$P(\theta|data) \propto P(data|\theta)P(\theta), \tag{6}$$

where $P(data|\theta)$ is the *quasilikelihood function* and $P(\theta)$ is the prior distribution of θ.

Priors We assume θ^Y and θ^R are independent with uniform priors on $[0, 1]$. These priors allow the data to drive the posterior estimates.

Quasilikelihood function We evaluate the quasilikelihood function on a grid over θ. Define the vector of input impulse responses for shock q as

$$\hat{\zeta}^{I,q} = \begin{bmatrix} \hat{\zeta}^{Y,q} \\ \hat{\zeta}_0^{R,q} \\ \hat{\zeta}_+^{R,q} \end{bmatrix}. \tag{7}$$

Note that the composite interest rate input R in the Jacobian $\mathcal{J}^{C,R}$ corresponds to this stacked vector: the initial portfolio return $r_0^{\text{portfolio}}$ followed by the Federal Funds rate path $(r_1, \ldots, r_{H-1})$.

For each q, we can partition Ω into blocks of the following form:

$$\Omega^{q,q'} = \begin{bmatrix} \Omega_{CC}^{q,q'} & \Omega_{CI}^{q,q'} \\ \Omega_{IC}^{q,q'} & \Omega_{II}^{q,q'} \end{bmatrix}, \tag{8}$$

where $\Omega_{XY}^{q,q'}$ corresponds to the covariance between $\zeta^{X,q}$ and $\zeta^{Y,q'}$. Note that $\Omega_{CC}^{q,q'} \in \mathbb{R}^{H \times H}$, $\Omega_{CI}^{q,q'} \in \mathbb{R}^{H \times 2H}$, $\Omega_{IC}^{q,q'} \in \mathbb{R}^{2H \times H}$, and $\Omega_{II}^{q,q'} \in \mathbb{R}^{2H \times 2H}$.

To evaluate the quasilikelihood function for a given θ and given input impulse responses $\hat{\zeta}^{I,q}$, first observe that each value of θ implies a vector of "true" impulse responses for consumption ζ^C. This true impulse response function is a linear function of the input impulse responses:

$$\zeta^C(\theta) = A(\theta) \cdot \zeta^I, \tag{9}$$

where $A(\theta)$ is a matrix that is made up of the model-implied Jacobians of consumption with respect to these inputs:

$$A(\theta) = \begin{bmatrix} \mathcal{J}^{C,Y}(\theta^Y) & \mathcal{J}^{C,R}(\theta^R) \end{bmatrix}. \tag{10}$$

The matrix $A(\theta)$ has dimension $H \times 2H$, so it operates on input paths truncated at horizon H. By Proposition 1, however, consumption at any date depends on the *entire* future path of labor income and interest rates. Closing the gap requires an assumption on how each input evolves beyond horizon H.

Assumption A1 (Continuation of input paths beyond horizon H)**.**

(i) Labor income remains at its horizon-$(H-1)$ *value:* $dY^L_s = dY^L_{H-1}$ *for all* $s \geq H$.

(ii) The interest rate returns to steady state: $dr_s = 0$ *for all* $s \geq H$.

Assumption A1(i) reflects the substantive nature of the shocks we study. In particular, productivity shocks may contain a unit root and therefore imply a permanent change in the level of labor income. For these shocks, allowing the income IRF to persist at its terminal value is the natural restriction. For shocks whose income response is transitory, such as monetary policy shocks, the IRF at horizon $H-1$ is close to zero, and the assumption degrades gracefully. Operationally, (i) implies that column $H-1$ of the truncated income Jacobian used in $A(\theta)$ is constructed as $\mathcal{J}^{C,Y^L}_{\cdot,H-1} + \sum_{s \geq H} \mathcal{J}^{C,Y^L}_{\cdot,s}$, capturing the cumulative effect on consumption of an income change that begins at $H-1$ and persists indefinitely.

Assumption A1(ii) reflects the fact that none of the shocks we study imply a permanent change in the real interest rate, which is a stationary object in any standard structural model. The contribution of columns $s \geq H$ of $\mathcal{J}^{C,R}$ is therefore negligible and is set to zero.

Equations (5) and (9) imply

$$\hat{\zeta}^{C,q} = A(\theta)\zeta^{I,q} + \epsilon^q_C, \tag{11}$$

$$\hat{\zeta}^{I,q} = \zeta^{I,q} + \epsilon^q_I, \tag{12}$$

where the errors are normally distributed $(\epsilon^q_C, \epsilon^q_I) \sim \mathcal{N}(0, \Omega)$. Define the difference between the direct empirically estimated consumption impulse response and the Jacobian-implied consumption impulse response as $\hat{m}^q \equiv \hat{\zeta}^{C,q} - A(\theta)\hat{\zeta}^{I,q}$. Then, stacking across shocks:

$$\hat{m} = \begin{bmatrix} m^1 \\ m^2 \\ \vdots \\ m^Q \end{bmatrix}. \tag{13}$$

This vector is normally distributed with mean 0 and variance Σ, where $\Sigma \in \mathbb{R}^{HQ \times HQ}$ is a block matrix with Q blocks of dimension $H \times H$:

$$\Sigma^{q,q'}(\theta) = \Omega_{CC}^{q,q'} + A(\theta)\Omega_{II}^{q,q'}A(\theta)' - \Omega_{CI}^{q,q'}A(\theta)' - A(\theta)\Omega_{IC}^{q,q'}. \tag{14}$$

Given our observed impulse responses $\hat{\zeta}$ and estimated covariance matrix $\hat{\Omega}$, we evaluate the quasilikelihood as:

$$P(data|\theta) = |\hat{\Sigma}|^{-1/2}\phi\Big(\hat{\Sigma}^{-1/2}\hat{m}\Big), \tag{15}$$

where ϕ is the probability density function for a multivariate random normal variable with mean 0 and variance I.[8] Note that the posterior distribution conditions on the first-stage covariance estimate $\hat{\Omega}$, and so the credible sets should be interpreted as conditional on this covariance estimator.

3.2. Estimating empirical impulse response functions. The first step in our estimation procedure is to recover the "data": our point estimate $\hat{\zeta}$ and the covariance matrix $\hat{\Omega}$. We combine multiple externally identified shock series into a vector Q_t and run the following local projection regression for each outcome variable O, horizon h, and time period t:

$$O_{t-1,t+h} = \zeta_h^O Q_t + \beta_h^O X_t + \epsilon_{t,h}^O, \tag{16}$$

where $O_{t-1,t+h}$ denotes the cumulative change in outcome O from period $t-1$ to $t+h$. We estimate this regression for 100 times log changes in consumption and income at horizons $h = 0, 1, \ldots, H-1$, for the portfolio return at horizon $h = 0$, and for changes in the federal funds rate at horizons $h = 1, 2, \ldots, H-1$, where $H = 16$ (four years of quarterly horizons). The portfolio return at $h = 0$ captures the immediate revaluation effect on household assets when the shock arrives, while the federal funds rate path starting at

[8] In practice, we apply a ridge normalization to Σ to address numerical instability issues.

$h = 1$ determines expected returns.

We then stack these estimates into $\hat{\zeta}$ and compute the covariance matrix $\hat{\Omega}$ across all horizons and outcome variables.[9] In light of Montiel Olea and Plagborg-Møller (2021), we include four lags of the controls and the shock series to calculate standard errors without the Newey and West (1987) adjustment.

Our sample period is 1970Q2–2024Q4. The data used in the estimation are defined as follows:

1. Consumption is measured as total real PCE, deflated using the PCE deflator.
2. Post-tax labor income consists of the components of real disposable income that are not derived from investments, in order to accord with the conception of labor income in our model.[10]
3. The portfolio return is constructed as a value-weighted average of equity and bond returns.[11]
4. The real federal funds rate is measured as the nominal federal funds rate minus median one-year-ahead expected CPI inflation from the Survey of Professional Forecasters.[12]
5. The control variables X_t are log industrial production, the unemployment rate, log CPI, log PPI for commodities, and the federal funds rate, following Ramey (2016).

Choice of structural shocks The methodology in this paper allows us to use any structural shock series, as long as the shock only affects consumption through its effect on labor income and asset 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 effects on labor 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. For example, monetary policy shocks have been identified in a number of ways but each of these suffers from some limitations for our purposes. First, monetary policy shocks are thought to be responsible

[9] We use HC1 heteroskedasticity-consistent standard errors.

[10] Specifically, the definition of income we 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.

[11] The market portfolio return is constructed as $R_t^{mkt} = 400 \times [\log(1 + r_t^{mkt,ex} + r_t^{rf}) - \pi_t]$, where $r_t^{mkt,ex}$ is the value-weighted excess return, r_t^{rf} is the risk-free rate, and π_t is quarterly PCE inflation. The excess return combines equity (Fama-French market factor Mkt-RF) and bond returns (5-year Gürkaynak, Sack and Wright (2007) zero-coupon holding-period return minus the risk-free rate), weighted by lagged market values of corporate equities and debt securities from the Federal Reserve's Financial Accounts (Z.1). This gives an annualized real return in percentage points.

[12] Pre-1981Q3, the expected GDP deflator is used in place of expected CPI inflation, which is not available for that period.

for only a small fraction of the total forecast variance. Second, the identification methods used 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 (Nakamura and Steinsson, 2018) draws into question whether these shocks are truly shocks to monetary policy or if they are in reaction to other macroeconomic events.[13] We therefore take the approach of using a wide variety of structural shocks, including but not only monetary policy 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.

The Ramey (2016) handbook chapter provides an overview of well-established structural shocks spanning multiple shock types and identification strategies, making it a natural source for our analysis. In order to limit the number of shock series to a manageable number while not cherry-picking series, we choose all ten of the shock series for which local projections are plotted in the figures contained in Ramey (2016). This handbook chapter 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 both to show the robustness of the results and bring a wide variety of evidence to answer the question of how households' consumption responds to interest rate changes. Table 1 summarizes these shocks.

[13] While Romer and Romer (2004) is an attempt to reduce the information effect, it only does so to the extent that the information is contained in the Greenbook forecasts.

Table 1: Summary of Macroeconomic Shocks

Panel A: Monetary Policy Shocks	
1. Romer and Romer (2004)	Monetary policy shocks identified by regressing the federal funds target rate on Greenbook forecasts at each FOMC meeting. The residual is taken as the shock.
2. Gertler and Karadi (2015)	High frequency identification using changes to three-month-ahead fed funds futures around a window of FOMC announcements to capture surprise changes to the policy rate.
Panel B: Fiscal Shocks	
3. Ramey (2011)	Military news shocks capturing changes in the expected present value of government purchases caused by military events. Identified by reading newspaper accounts (e.g., BusinessWeek) to capture news of spending events.
4. Ben Zeev and Pappa (2017)	Defense spending news shocks following Barsky and Sims (2011). Shocks that best explain future movements in defense spending over a five-year horizon and are orthogonal to current defense spending.
5. Blanchard and Perotti (2002)	Government spending shocks identified via Cholesky decomposition of a VAR with government spending ordered first.
6. Mertens and Ravn (2011)	Tax news shocks building on Romer and Romer (2010), dividing tax shocks into anticipated and unanticipated based on delay between legislation passage and implementation.
7. Leeper, Richter and Walker (2012)	Expected tax changes one to five years forward based on the spread between federal and municipal bonds, reflecting anticipated tax changes under market efficiency.
Panel C: Technology Shocks	
8. Ben Zeev and Khan (2015)	Investment-specific technology (IST) news shocks. IST measured as inverse of real price of investment. Shocks identified using maximum forecast error variance approach, orthogonal to current TFP and IST.
9. Fernald (2014)	Utilization-adjusted TFP shocks. Solow residual adjusted for variations in factor utilization following Basu, Fernald and Kimball (2006). Shock series calculated as quarterly changes to adjusted TFP.
10. Francis et al. (2014)	Unanticipated TFP shocks identified by maximizing the contribution to forecast-error variance of labor productivity at a long horizon.

4. Structural estimation results

We begin by presenting results from individual shock series to illustrate how our methodology identifies the sticky expectations parameters. We focus on two detailed examples—a monetary policy shock (Romer and Romer, 2004) and a fiscal policy shock (Ben Zeev and Pappa, 2017)—to show how different types of shocks lead to similar conclusions about household behavior. We then present parameter estimates from all ten shock series individually, showing that the findings are robust across different shock types and identification strategies. Finally, we present results from joint estimation across all ten shocks, which pools the evidence to provide our tightest parameter estimates.

4.1. Results with one shock: Romer and Romer (2004). We use the results from the Romer and Romer (2004) shocks, a widely-used set of monetary policy shocks, to illustrate the estimation methodology. Figure 4 shows the response of real PCE, our real labor income measure, the real interest rate, and the real market portfolio return. In response to a contractionary shock, real PCE falls, real income falls, the real interest rate rises, and the excess return on the market portfolio falls for one period.

Figure 5 shows the resulting marginal posterior distributions for the sticky expectations parameters. Both parameters range from 0 to 1; at a value of 0, households have full information over that aggregate variable, and at a value of 1, they never update their expectations. As the left panel shows, the posterior distribution is maximized with θ^R at 1, indicating completely sticky expectations over the interest rate. At this value, households do not respond to interest rates at all. The sticky expectations parameter for income is imprecisely measured, though the maximum a posteriori (MAP) estimate is about 0.9, implying that about 10 percent of households update their expectations about future income each quarter.

Figure 6 depicts the estimated empirical impulse response for consumption from the Romer and Romer (2004) shocks, along with several model-based simulated paths for consumption. To construct each of these paths for consumption, we feed in the empirical impulse response for labor income, the period-0 portfolio return, and the period-1 onward impulse response for the real federal funds rate to the Jacobian model described in equation (10). The red dashed line shows the model-based simulation, using the MAP estimates for θ^Y and θ^R. As it shows, the model-based consumption path matches the empirical path closely, capturing the hump-shaped response of consumption to the shock.

The purple dotted line shows the simulated path of consumption in a FIRE model, where all households have perfect foresight over the paths of the interest rate and labor income following the realization of the shock. The FIRE model misses the hump-shaped response of consumption to the shock, implying a sizeable, immediate decline in household spending following the shock. This response incorporates

Figure 4: Romer and Romer (2004) shock impulse responses

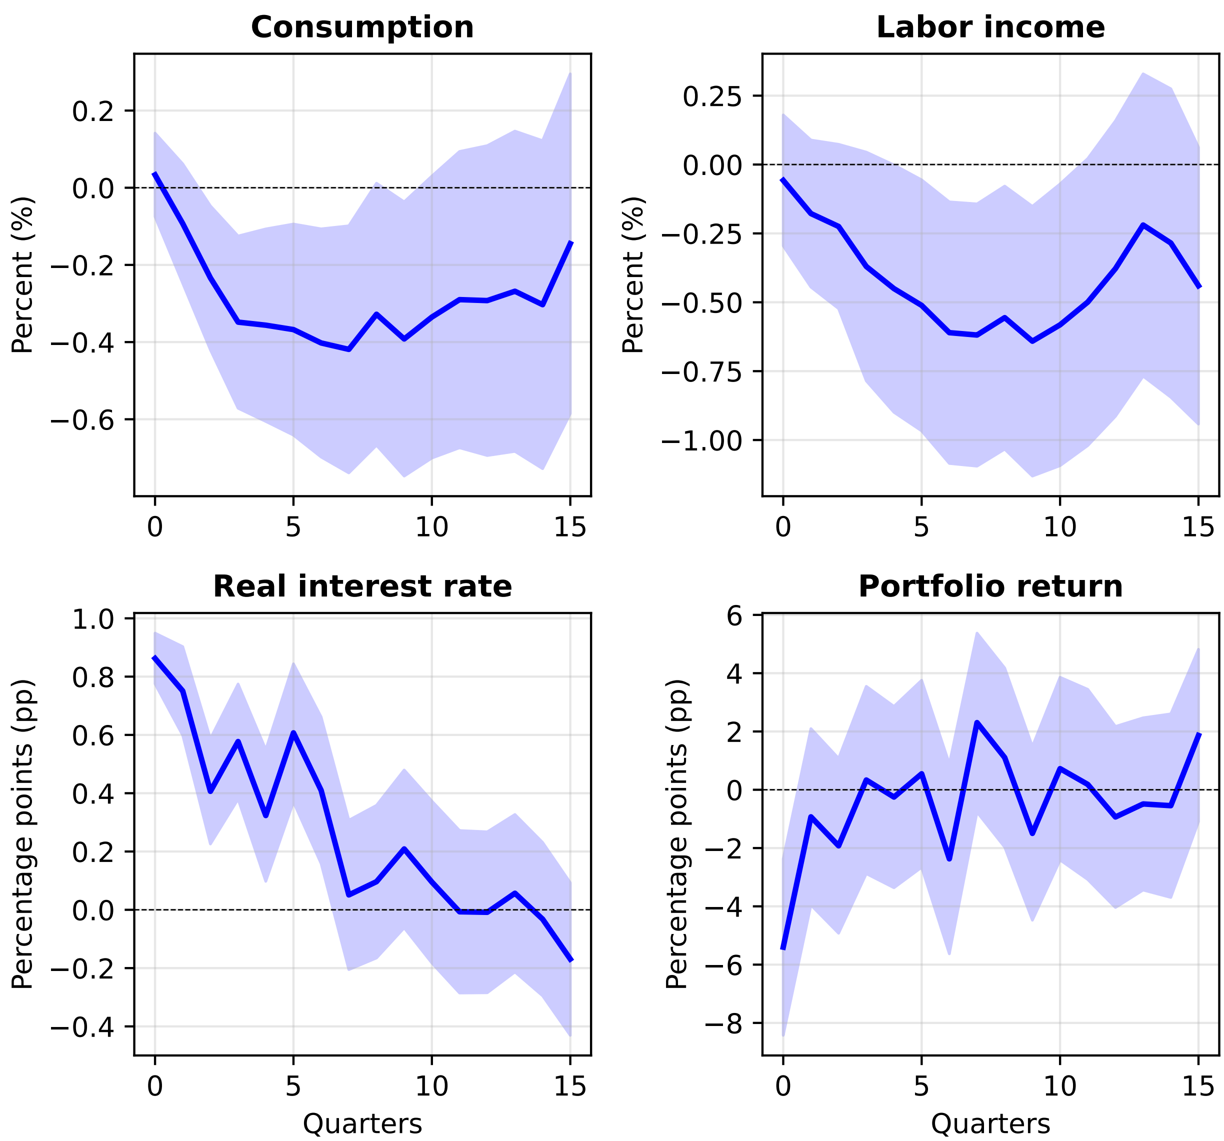

Notes: Shaded regions show 90 percent confidence intervals.

Figure 5: Marginal posterior distributions for sticky expectations parameters: Romer and Romer (2004) shocks

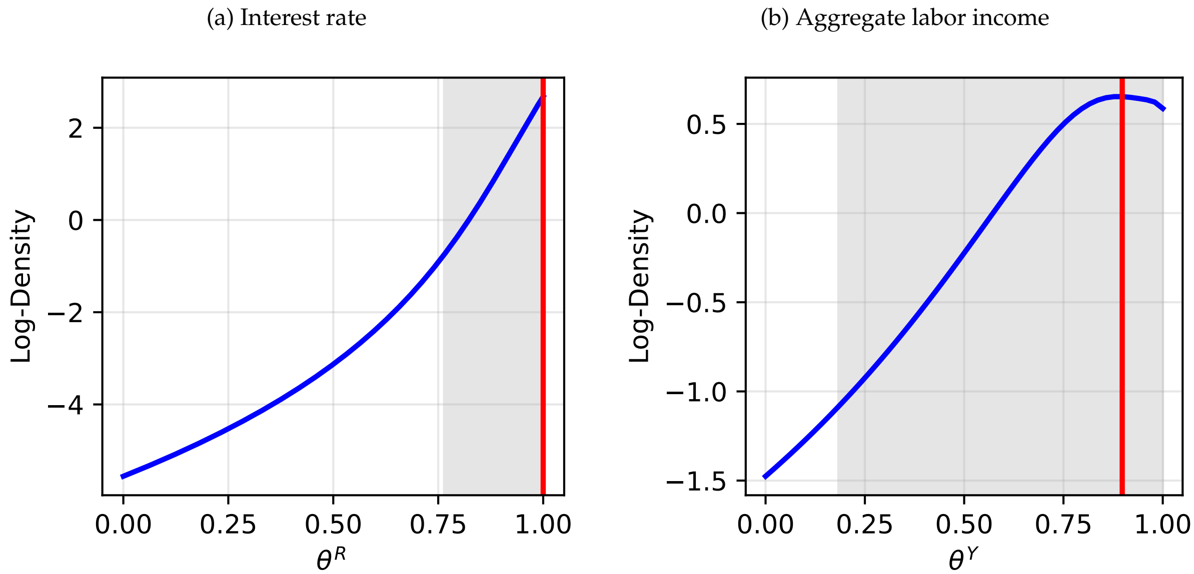

Notes: Blue solid lines depict the log marginal posterior density over the support of the estimated parameter. Red vertical solid lines indicate the posterior mode, and grey shaded areas indicate the 95 percent highest posterior density regions.

households' reaction to the expected decline in income, as well as the immediate rise in interest rates and decline in their portfolio value.

The remaining two lines show simulated paths of consumption with sticky expectations on only one of the variables. Comparing the green dot-dash line, which shows the consumption path in a model with perfect foresight on interest rates but sticky expectations for labor income, to the purple dotted line (FIRE) shows that sticky labor income expectations have a relatively small effect on the consumption path. This is, in part, due to the fact that the estimated parameter for sticky expectations over labor income implies that most households update their expectations after just two quarters. Comparing the yellow dashed line to the purple dotted line shows that sticky interest rate expectations matter much more, generating the hump-shaped response that we see in the data.

How are these parameters identified? Examining the impulse responses in Figure 4 and the Jacobians in Figure 2 and 3 helps illuminate how our method identifies these parameters. Looking first at interest rates: The interest rate Jacobian implies that households with perfect foresight respond close-to-immediately to both future and current changes in the interest rate. So, with perfect foresight over interest rates, any change in interest rates, either delayed or immediate, implies an immediate decline in consumption that is not present in the data. This finding connects to the literature on the strength of the forward guidance puzzle;

Figure 6: Comparison of the consumption path to Romer and Romer (2004) shock across model assumptions

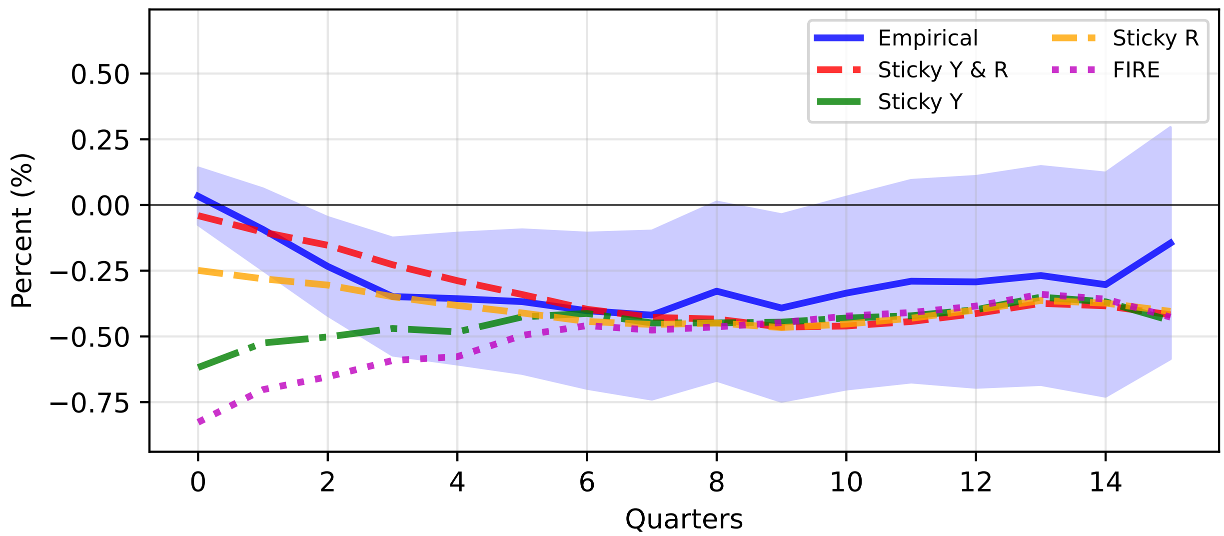

Notes: Blue solid line shows the empirical impulse response (shaded region shows 90 percent confidence intervals). Red dashed line shows the model with MAP estimates of both θ^Y and θ^R (sticky expectations for both labor income and interest rates). Yellow dashed line shows the model with sticky interest-rate expectations only (θ^R at MAP, $\theta^Y = 0$). Green dot-dash line shows the model with sticky labor-income expectations only (θ^Y at MAP, $\theta^R = 0$). Purple dotted line shows the full-information rational expectations (FIRE) model ($\theta^Y = \theta^R = 0$).

see, for example, Hagedorn et al. (2019) and McKay, Nakamura and Steinsson (2016). Sticky expectations are needed to tamp down on this immediate response to current and future changes in the interest rate.

Turning now to income: The Jacobian for income has a very different shape, with a much more pronounced peak in the period when the shock arrives. Following the Romer and Romer (2004) shock, income responds only gradually, meaning that even with relatively flexible expectations, consumption moves gradually.

4.2. Results with one shock: Ben Zeev and Pappa (2017). Monetary policy shocks have often been used to understand how consumption responds to interest rates. Here we show how our methodology can also leverage a fiscal shock, the defense spending news shock from Ben Zeev and Pappa (2017), to provide similar insights. Following this shock, real interest rates increase and remain elevated for a period of 3 years or more. Consumption and labor income both move down over time, before recovering.[14] This shock provides a useful comparison: although it features a more persistent increase in the interest rate than the Romer and Romer (2004) example, the consumption and income responses are quite similar, and the resulting estimate of θ^R is essentially the same.

Bayesian estimation using this shock alone leads to similar parameter posteriors as using Romer and

[14]The full set of impulse responses and posterior distributions for all shocks are shown in online Appendix D.

Romer (2004) shocks alone. The posterior mode estimates are $\theta^Y = 0.94$ and $\theta^R = 0.99$, both very close to those from the Romer and Romer shocks ($\theta^Y = 0.90$ and $\theta^R = 1.00$). Once again, the stickiness parameter for the interest rate is near the upper bound, indicating households do not update their expectations about future interest rate changes.

Figure 7 depicts the estimated empirical impulse response for consumption from the Ben Zeev and Pappa (2017) shocks, along with several model-based simulated paths. The red dashed line shows the model-based simulation using the MAP estimates for θ^Y and θ^R. As with the Romer and Romer shocks, the model-based consumption path closely tracks the empirical path, capturing the delayed decline in consumption. The purple dotted line shows the simulated path in a FIRE model with perfect foresight. The FIRE model again predicts an immediate, sharp decline in household spending that is absent in the data, driven by households' immediate response to the anticipated path of higher interest rates and lower income. The remaining two lines show simulated paths with sticky expectations on only one variable. Comparing the green dot-dash line (sticky income expectations only) to the purple dotted line shows that sticky labor income expectations contribute modestly to the gradual consumption response. In contrast, comparing the yellow dashed line (sticky interest rate expectations only) to the purple line shows that sticky interest rate expectations are again the dominant factor, preventing the immediate consumption decline predicted by the FIRE model.

4.3. Individual estimation results for each of the 10 shocks. Before pooling information across all shocks in a joint estimation, we first examine parameter estimates from each shock individually to demonstrate that the main finding is robust across different shock types and identification strategies.

Table 2 shows the posterior modes and 95 percent highest posterior density intervals for the two sticky expectations parameters when estimation is conducted for each shock individually.

The parameter that governs how household spending responds to interest rates, θ^R, is estimated to be close to 1.0. For all 10 shocks, 1.0 is included in the 95 percent highest posterior density interval, and 1.0 is the mode of the posterior distribution for eight of ten shocks. While some of the shocks used provide relatively tight identification on this parameter, others, such as the Ben Zeev and Khan (2015) IST news shock, have wider confidence intervals. Nevertheless, all the evidence is consistent with household spending not responding directly to interest rates.

The parameter that governs sticky expectations for income, θ^Y, is of secondary interest to our analysis. This parameter is estimated with low precision for all 10 shocks individually. As discussed in Section 2.5, the income Jacobian has a pronounced peak in the period when the shock arrives, meaning that consumption responds primarily to contemporaneous income rather than expected future income. Because the empirical

Figure 7: Comparison of the consumption path to Ben Zeev and Pappa (2017) shock across model assumptions

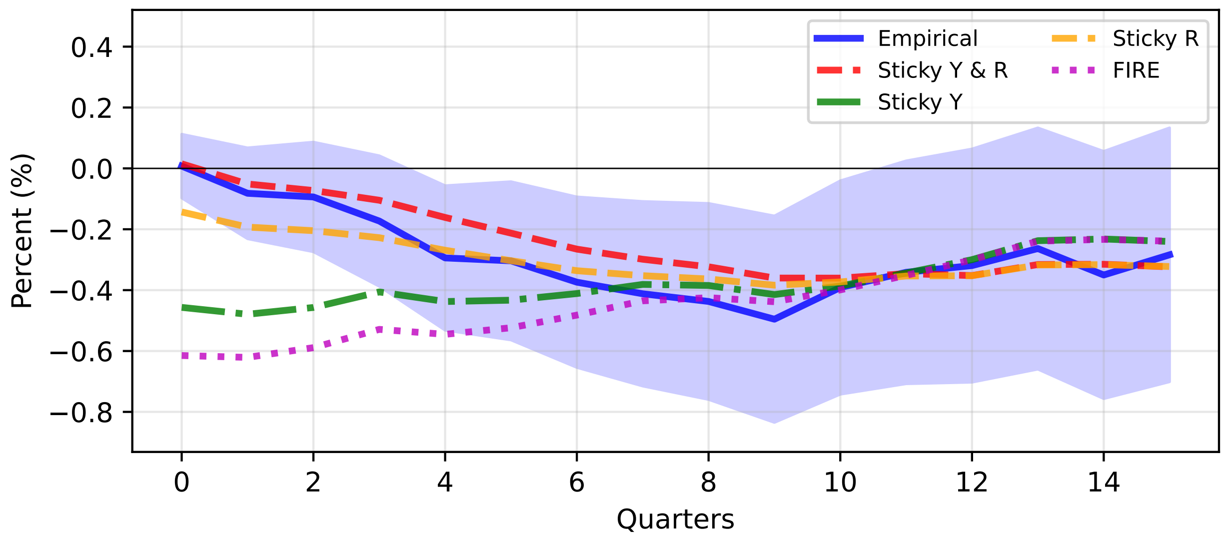

Notes: Blue solid line shows the empirical impulse response (shaded region shows 90 percent confidence intervals). Red dashed line shows the model with MAP estimates of both θ^Y and θ^R (sticky expectations for both labor income and interest rates). Yellow dashed line shows the model with sticky interest-rate expectations only (θ^R at MAP, $\theta^Y = 0$). Green dot-dash line shows the model with sticky labor-income expectations only (θ^Y at MAP, $\theta^R = 0$). Purple dotted line shows the full-information rational expectations (FIRE) model ($\theta^Y = \theta^R = 0$).

Table 2: Estimated Sticky Expectations Parameters (Individual Shock Estimations)

Shock	θ^Y	θ^R
	0.90	1.00
Romer & Romer	(0.18, 1.00)	(0.76, 1.00)
	0.45	1.00
Gertler-Karadi	(0.00, 0.92)	(0.23, 1.00)
	0.94	1.00
Ramey News	(0.06, 1.00)	(0.10, 1.00)
	0.94	0.99
Ben Zeev & Pappa	(0.08, 1.00)	(0.45, 1.00)
	0.73	0.93
Blanchard-Perotti	(0.00, 0.92)	(0.06, 1.00)
	0.86	1.00
Tax News	(0.06, 1.00)	(0.09, 1.00)
	0.78	1.00
Leeper-Richter-Walker	(0.06, 1.00)	(0.70, 1.00)
	0.43	1.00
Ben Zeev & Khan IST News	(0.00, 0.73)	(0.00, 1.00)
	0.80	1.00
Fernald TFP Growth	(0.04, 0.98)	(0.06, 1.00)
	0.31	1.00
FORD TFP	(0.00, 0.76)	(0.12, 1.00)

Notes: θ^Y is the stickiness parameter for income shocks, θ^R is the stickiness parameter for interest rate. Point estimates are posterior modes (MAP). 95% highest posterior density (HPD) intervals shown in parentheses. HPD intervals are the narrowest intervals containing 95% of the posterior probability and are guaranteed to include the MAP estimate.

income impulse responses evolve gradually for most shocks, the consumption path is also gradual regardless of whether income expectations are sticky or flexible. This makes θ^Y harder to identify than θ^R, where the contrast between the Jacobian's immediate response to future interest rates and the empirical IRFs' delayed response provides sharp identification.

Despite the considerable uncertainty for θ^Y across shocks, all individual estimates for θ^R cluster tightly near 1.0, with eight of ten shocks yielding a posterior mode exactly at the upper bound.

4.4. Joint estimation results. Our methodology allows us to pool the evidence from all ten shocks, even when they are correlated as might be expected with different methods of identifying monetary or fiscal shocks. This pooling has the advantage that we can identify the parameters with more precision, but it has the disadvantage that it requires the identification methods for all ten shocks to be valid. By contrast, the individual-shock parameter estimates require only that the shock used has a valid identification.

Figure 8 shows the results when we estimate θ^Y and θ^R on all ten shocks jointly. As the figure shows, the sticky expectations parameter on the interest rate is tightly estimated at 1, implying households never update their expectations about interest rates.[15] Compared to the individual estimates, the joint estimation substantially narrows the posterior distribution on θ^R. The sticky expectations parameter on labor income is estimated at $\theta^Y = 0.71$ (95% HPD: [0.39, 0.88]), implying that about 30 percent of households update their expectations of future income each period.

We now ask whether this finding survives without the parametric structure of the sticky-expectations model.

5. Non-parametric estimation

So far, we have shown that a heterogeneous-agent model with sticky expectations best rationalizes the macroeconomic data when expectations over the interest rate are completely sticky and households never update their expectations about the interest rate. To check that this result is not an artifact of the sticky-expectations parameterization, we next estimate the response of households to interest rates non-parametrically.

Rather than parameterizing the interest rate Jacobian with a single parameter, θ^R, we approximate it using a flexible basis function expansion. The basis imposes smoothness within economically motivated regions of the Jacobian—a restriction that is satisfied by essentially all structural models—while remaining agnostic about the specific structural form of the model. The approach serves two purposes. First, it acts as

[15] A natural concern is that the corner estimate suggests that our model is misspecified. In an alternative procedure, we allow θ^R to rise above 1. Although $\theta^R > 1$ has no structural interpretation, we implement this parameterization as $\mathcal{J}(\theta^R) = -\mathcal{J}(2 - \theta^R)$. Under this alternative, we estimate $\theta^R = 1.02$, very close to our estimate of 1.00.

Figure 8: Posterior distribution for sticky expectations parameters: all shocks

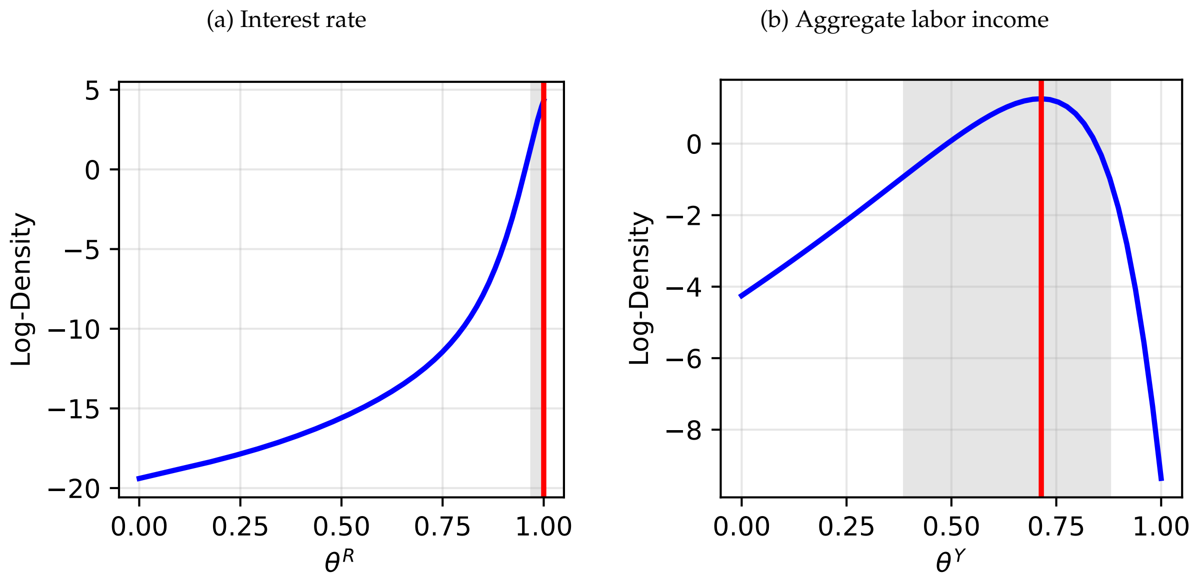

Notes: Blue solid lines depict the log marginal posterior density over the support of the estimated parameter. Red vertical solid lines indicate the posterior mode, and grey shaded areas indicate the 95 percent highest posterior density regions.

a robustness check: if the data truly imply that households do not respond to interest rates, we should find this result whether or not we impose the parametric structure of the sticky expectations model. Second, by encoding only the minimal structure shared across a broad class of structural models, it allows us to assess whether the data prefer a Jacobian that departs from the predictions of any standard model.

5.1. Econometric framework. Following (4), for each shock q, the impulse response of consumption to a shock can be decomposed into the component coming from labor income and the component coming from the interest rate:

$$\zeta^{C,q} = \mathcal{J}^{C,Y}\zeta^{Y,q} + \mathcal{J}^{C,R}\zeta^{R,q}. \tag{17}$$

In our baseline estimation, both Jacobians are functions of $\theta = (\theta^Y, \theta^R)$. Here, we instead hold $\mathcal{J}^{C,Y}$ fixed at its structural value—evaluated at the posterior mode from the joint estimation in Section 4.4—and treat $\mathcal{J}^{C,R}$ as a matrix to be estimated from the data, subject to smoothness constraints encoded in a basis function expansion. Because we condition on the posterior mode from our structural estimation $\hat{\theta}^Y_{\text{MAP}}$ rather than treating θ^Y as a free parameter, the posterior we compute is formally

$$p\left(\mathcal{J}^{C,R} \,\middle|\, \hat{\zeta}, \hat{\Omega}, \hat{\theta}^Y_{\text{MAP}}\right), \tag{18}$$

not the marginal posterior for $\mathcal{J}^{C,R}$. We therefore report results below as conditional on the income Jacobian.

Define the *income-adjusted consumption response* as the consumption IRF less the component explained by income:

$$\tilde{\zeta}^{C,q} \equiv \zeta^{C,q} - \mathcal{J}^{C,Y}\zeta^{Y,q} \tag{19}$$

$$= \mathcal{J}^{C,R}\zeta^{R,q}. \tag{20}$$

We observe estimated impulse responses $\hat{\zeta}^{C,q}$, $\hat{\zeta}^{Y,q}$, and $\hat{\zeta}^{R,q}$ from local projections. As in the baseline estimation, these estimates are normally distributed around their true values with a known covariance structure $\hat{\Omega}$ (estimated from the local projection regressions). We henceforth drop the $\hat{}$ from estimated impulse responses; all matrices $\tilde{C}$, R, y in what follows are formed from estimated objects.

Arranging the estimated impulse responses for each shock as columns of the matrices R and $\tilde{C}$ for the interest rate and income-adjusted responses, respectively, equation (20) becomes a matrix equation:

$$\tilde{C} = \mathcal{J}^{C,R}R + \text{error}. \tag{21}$$

5.2. Basis function parameterization. We approximate the interest rate Jacobian $\mathcal{J}^{C,R}$ using a tensor-product B-spline basis. We write

$$\text{vec}(\mathcal{J}^{C,R}) = B\vartheta + \text{approximation error}, \tag{22}$$

where $B \in \mathbb{R}^{H^2 \times p}$ is a basis evaluation matrix and $\vartheta \in \mathbb{R}^p$ is a vector of basis coefficients, with $p \ll H^2$. The basis we use imposes smoothness on parts of the Jacobian, while allowing for a discontinuity that is consistent with many theories of household behavior.

Two-region structure The interest rate Jacobian has a natural partition along its diagonal. The element $\mathcal{J}^{C,R}_{h,s}$ gives the response of consumption at horizon h to an interest rate change at horizon s. When $h \geq s$, the interest rate change has already been realized—these elements capture the *causal* response of consumption to past rate changes. When $h < s$, the interest rate change lies in the future—these elements capture the *anticipatory* response to expected future rate changes.

In all structural models we are aware of, the Jacobian is smooth within each of these two regions but exhibits a discontinuity along the diagonal. The economic intuition is straightforward. In the lower-triangular region, the response of consumption at horizon h to a rate change at horizon $s \leq h$ reflects how households adjust their spending after the rate change has occurred. These adjustments evolve smoothly as the time since the rate change increases: the response to a rate change that occurred three quarters ago is similar to the response to a rate change two quarters ago. In the upper-triangular (anticipatory) region, the response reflects how households adjust spending *before* the rate change arrives. These anticipatory adjustments also evolve smoothly: the response to a rate change expected in six quarters is similar to one expected in seven quarters.

However, the transition from anticipation to realization—from $h < s$ to $h \geq s$—involves a discrete shift in household behavior. Before an anticipated one-period increase in the interest rate, households may adjust through intertemporal substitution—they are incentivized to save more so that they can take advantage of the future higher interest rate. After the rate increase is realized, consumption jumps as household wealth is above its steady-state level, reflecting both the increased savings from before the interest rate change as well as the increased return at the time of the shock. This creates a structural break along the diagonal that is present across a wide class of models, including models with finite planning horizons, myopia, and varying degrees of forward-looking behavior.

Figure 9 illustrates this regularity by plotting column 8 of the interest rate Jacobian—the consumption response at each horizon h to an interest rate change at horizon $s = 8$. We compare four structural models: the baseline FIRE model, the sticky expectations model ($\theta^R = 0.8$), a finite horizon planning model, and a myopia model.[16] All four Jacobians are smooth within the lower-triangular and upper-triangular regions, with a discrete break at the diagonal ($h = s$). Despite their different microfoundations, they share the same qualitative two-region structure. The figure also shows that the tensor-product B-spline basis approximation (described below) closely tracks each model's Jacobian, confirming that the basis can flexibly represent the Jacobian shapes implied by a wide range of structural assumptions.

We construct the basis separately for each region. For the lower-triangular region ($h \geq s$) and the upper-triangular region ($h < s$), we define a two-dimensional tensor-product B-spline basis:

$$\phi_{i,j}(h, s) = B_i^{(h)}(h) \cdot B_j^{(s)}(s), \tag{23}$$

where $B_i^{(h)}$ and $B_j^{(s)}$ are one-dimensional B-spline basis functions of degree k (we use cubic splines, $k = 3$) in the response-horizon and shock-horizon directions, respectively. With n_h basis functions in the h-direction and n_s in the s-direction, each region has $n_h \times n_s$ basis functions. The total number of parameters is $p = 2 \times n_h \times n_s$. In our baseline specification, $n_h = n_s = 6$, yielding $p = 72$ parameters. This specification is flexible enough to fit a wide range of structural models, as shown in Figure 9.

[16] The details of these models are described in Appendix C.

Figure 9: Eighth column of the interest rate Jacobian in four structural models

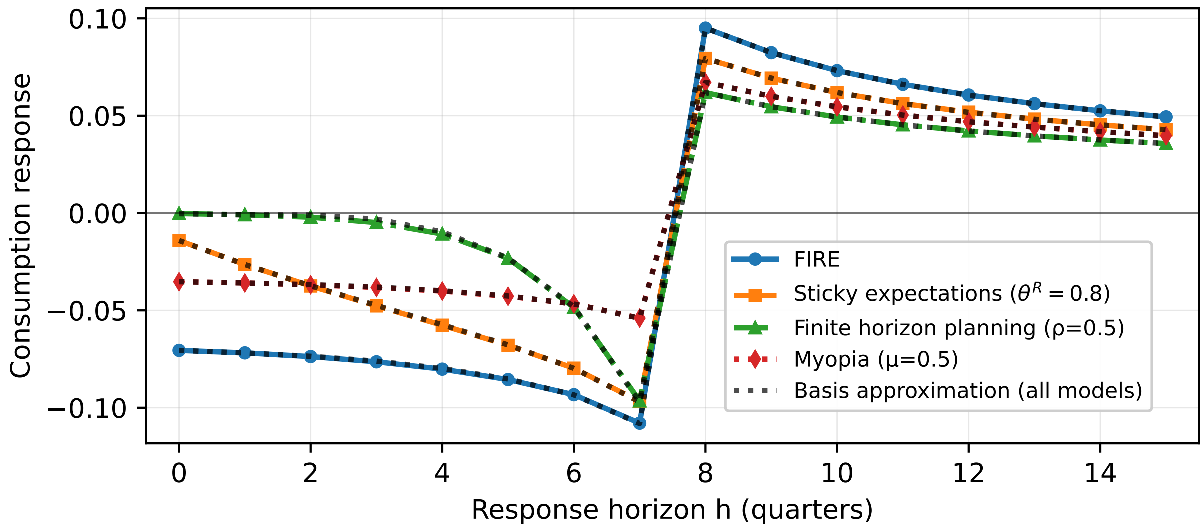

Notes: The figure plots column 8 of the interest rate Jacobian ($\mathcal{J}^{C,R}_{h,8}$) for four structural models: the baseline full-information rational expectations (FIRE) model, the sticky expectations model with $\theta^R = 0.8$, a finite horizon planning model, and a myopia model. Each curve shows the consumption response at horizon h to an interest rate change at horizon $s = 8$. The black dotted lines show the corresponding basis function approximations using the tensor-product B-spline parameterization. All models exhibit smooth behavior within the causal ($h \geq s$) and anticipatory ($h < s$) regions, with a structural break at the diagonal—the minimal regularity encoded in the basis.

5.3. Vectorization and the Bayesian linear regression. Substituting equation (22) into the vectorized model, we obtain

$$\text{vec}(\tilde{C}) = (R' \otimes I_H)\, B\vartheta + \varepsilon, \tag{24}$$

where $\otimes$ denotes the Kronecker product, I_H is the $H \times H$ identity matrix, and ε is a normally distributed error term. Define $y \equiv \text{vec}(\tilde{C}) \in \mathbb{R}^{HQ}$, $X \equiv R' \otimes I_H \in \mathbb{R}^{HQ \times H^2}$, and $X_B \equiv XB \in \mathbb{R}^{HQ \times p}$. The model is then a Bayesian linear regression in the basis coefficients:

$$y = X_B\vartheta + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \Omega_0), \tag{25}$$

$$\vartheta \sim \mathcal{N}(\mu_0, \Sigma_0),, \tag{26}$$

where $\Omega_0 \in \mathbb{R}^{HQ \times HQ}$ is the covariance matrix of the errors, $\mu_0 \in \mathbb{R}^p$ is the prior mean, and $\Sigma_0 \in \mathbb{R}^{p \times p}$ is the prior covariance.

This approach is equivalent to regressing the income-adjusted consumption response $\tilde{C}$ on the interest rate response R to recover $\mathcal{J}^{C,R}$. However, rather than estimating all H^2 elements of the Jacobian directly, we estimate $p = 72$ basis coefficients that impose economically motivated smoothness within the causal and anticipatory regions.

Error covariance The error covariance Ω_0 is constructed exactly as in the baseline estimation. It accounts for the fact that both the consumption and interest rate impulse responses are estimated with error, and depends on the full covariance structure of the local projection estimates across outcomes and shocks. We evaluate Ω_0 at a reference Jacobian—the structural model Jacobian at the posterior mode from the baseline estimation—and hold it fixed. Treating the error covariance as known conditional on a first-stage reference parameter follows the logic of two-step estimation (Newey and McFadden, 1994), analogous to feasible GLS where the covariance matrix is estimated in a preliminary stage. Holding Ω_0 fixed also preserves the conjugacy of the Bayesian linear regression and avoids the need to re-evaluate Ω_0 at each candidate $\mathcal{J}^{C,R}$.

5.4. Prior specification. We center the prior on the structural model's interest rate Jacobian. Specifically, let $\mathcal{J}^{C,R}_{\text{struct}}$ denote the structural Jacobian under FIRE. We obtain the prior mean for the basis coefficients by projecting $\mathcal{J}^{C,R}_{\text{struct}}$ onto the basis:

$$\mu_0 = \arg\min_{\vartheta} \|B\vartheta - \text{vec}(\mathcal{J}^{C,R}_{\text{struct}})\|^2 = (B'B)^{-1}B'\text{vec}(\mathcal{J}^{C,R}_{\text{struct}}). \tag{27}$$

As shown in Figure 9, the basis approximation is highly accurate: the relative projection error $\|B\mu_0 - \text{vec}(\mathcal{J}^{C,R}_{\text{struct}})\|/\|\text{vec}(\mathcal{J}^{C,R}_{\text{struct}})\|$ is below 10^{-4} in our specification.

Centering the prior at the structural model under FIRE serves as a natural benchmark. For the prior covariance, we use an isotropic prior on the basis coefficients:

$$\Sigma_0 = \kappa^2 I_p, \tag{28}$$

where $\kappa > 0$ is a scalar hyperparameter. Since the B-spline basis already encodes smoothness within each region of the Jacobian, an isotropic prior on the basis coefficients is sufficient—no additional smoothness penalty is needed. A small κ pulls the posterior toward the structural model; a large κ allows the data to dominate.

5.5. Analytical posterior. The conjugacy of the Gaussian prior with the Gaussian quasilikelihood yields an analytical posterior. The posterior distribution of ϑ is:

$$\vartheta \mid y \sim \mathcal{N}(\mu_{\text{post}}, \Sigma_{\text{post}}), \tag{29}$$

where the posterior precision and mean are given by:

$$\Sigma_{\text{post}}^{-1} = X_B' \Omega_0^{-1} X_B + \Sigma_0^{-1}, \tag{30}$$

$$\mu_{\text{post}} = \Sigma_{\text{post}} \left(X_B' \Omega_0^{-1} y + \Sigma_0^{-1} \mu_0 \right). \tag{31}$$

These are the standard Bayesian linear regression updating formulas. The posterior mean is a precision-weighted average of the data-implied estimate and the prior mean (the structural model prediction).

To recover the posterior in Jacobian space, we use the linear mapping $\text{vec}(\mathcal{J}^{C,R}) = B\vartheta$:

$$\text{vec}(\mathcal{J}^{C,R}) \mid y \sim \mathcal{N}\left(B\mu_{\text{post}},\ B\Sigma_{\text{post}} B' \right). \tag{32}$$

The posterior mean, reshaped into an $H \times H$ matrix, gives the estimated interest rate Jacobian $\hat{\mathcal{J}}^{C,R}_{\text{post}}$.

5.6. Impulse response distributions. We are interested in the consumption response to a hypothetical interest rate path $r \in \mathbb{R}^H$:

$$\Delta C = \mathcal{J}^{C,R} r. \tag{33}$$

Since $\mathcal{J}^{C,R}$ is uncertain, ΔC is a random vector. We can write

$$\Delta C = P \cdot \text{vec}(\mathcal{J}^{C,R}), \quad \text{where} \quad P = r' \otimes I_H \in \mathbb{R}^{H \times H^2}. \tag{34}$$

Substituting the Jacobian posterior from equation (32), the implied consumption response is Gaussian:

$$\Delta C \mid y \sim \mathcal{N}\left(PB\mu_{\text{post}},\ PB\Sigma_{\text{post}} B' P' \right). \tag{35}$$

This gives us a full posterior distribution over the consumption response at each horizon, including pointwise credible intervals.

5.7. Results. We apply this procedure using all ten shocks jointly. Figure 10 compares the posterior mean Jacobian (panel 10b) to the prior mean Jacobian, which is that of a FIRE model (panel 10a). Under the FIRE model, the Jacobian is negative for response horizons before the shock ($h < s$) and positive for shock horizons equal to or after the shock ($h \geq s$). The estimated model Jacobian is centered close to zero at all horizons, and these figures show that the data shift the posterior sharply toward zero from the prior.

To understand what interest rate dynamics look like in this model, we compute the posterior distribution of the consumption response to an exponentially decaying interest rate shock with a 4-quarter half-life. This

Figure 10: Estimated interest rate Jacobians

(a) FIRE model

(b) Non-parametrically estimated model

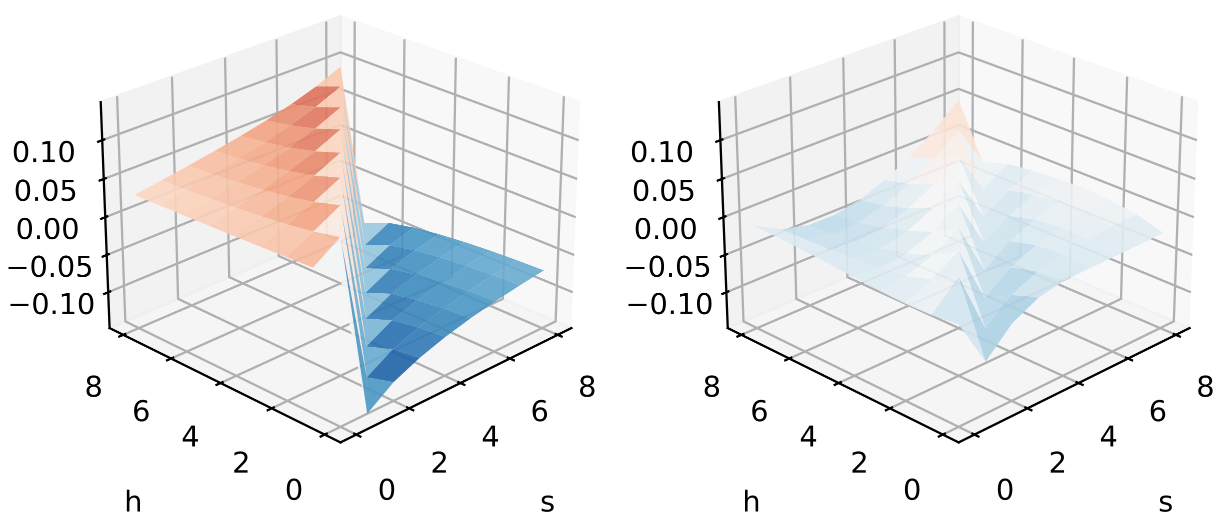

Notes: The figure plots the interest rate Jacobian $\mathcal{J}^{C,R}_{h,s}$, showing the consumption response at horizon h to an interest rate change at horizon s. Panel (a) shows the FIRE model, which serves as the prior mean. Panel (b) shows the posterior mean Jacobian estimated using all ten shocks jointly.

experiment asks: if the real interest rate rises by 1 percentage point on impact and decays exponentially, how does consumption respond? The interest rate path also implies an initial capital loss on asset holdings (for this exercise, we assume the asset is a perpetual annuity), which enters the model through the $h = 0$ component of the interest rate variable.

Figure 11 plots the prior and posterior distributions of this consumption response. The prior distribution, shown in orange, is centered on the structural model's prediction and is wide, reflecting our choice of a relatively uninformative isotropic prior over the basis coefficients. The posterior distribution, shown in blue, is notably tighter: the data are informative about the consumption response, and the posterior mean is approximately zero at every horizon beyond impact. The 95% credible interval includes zero throughout.

The posterior mean is slightly negative at horizon $h = 0$. This reflects the prior, which is centered on the structural model's prediction of a negative impact response (due to the capital loss on asset holdings). We verified that when we repeat the exercise with the prior centered at zero instead, the posterior mean is not negative at impact. This illustrates that the slight negativity at $h = 0$ is driven by the prior rather than the data.

The green dotted line shows the prediction of the structural model evaluated at FIRE ($\theta^R = 0$), which implies a meaningful consumption decline in response to the interest rate increase. The posterior clearly rejects this FIRE prediction, instead favoring a Jacobian that generates essentially no consumption response

Figure 11: Non-parametric estimate of the consumption response to a real interest rate shock

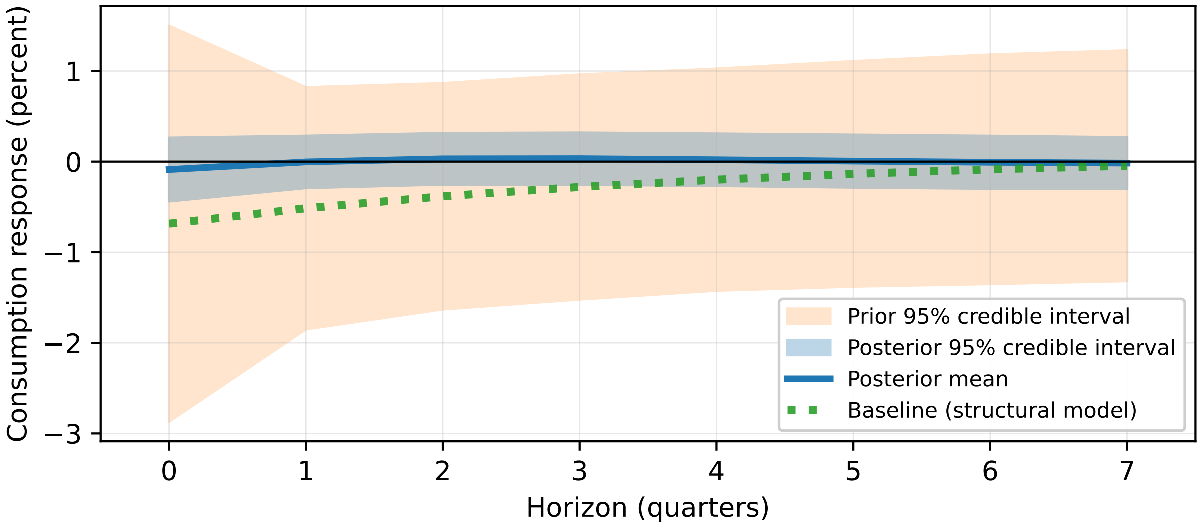

Notes: The figure plots the prior and posterior distributions of the consumption response to an exponentially decaying real interest rate shock (1 percentage point on impact, 4-quarter half-life), computed from the non-parametric interest rate Jacobian estimated using all ten shocks. The Jacobian is parameterized using a tensor-product B-spline basis with separate expansions for the lower-triangular (causal) and upper-triangular (anticipatory) regions. The orange shaded region is the 95% prior credible interval, centered on the structural model's prediction. The blue shaded region is the 95% posterior credible interval, with the posterior mean shown in the solid blue line. The green dotted line shows the structural model's prediction under full-information rational expectations ($\theta^R = 0$).

to interest rate changes.

This result corroborates the finding from the sticky expectations estimation. In that estimation, the posterior for θ^R is concentrated at 1, implying that the sticky expectations model rationalizes the data by effectively shutting down the interest rate channel. Here, without imposing the parametric structure of sticky expectations—but retaining the minimal smoothness properties shared by all structural models—we find the same conclusion: the data prefer an interest rate Jacobian that generates essentially no consumption response to interest rate changes.

5.8. Discussion. The non-parametric estimation complements the structural estimation in several ways. First, it demonstrates that the finding of no household response to interest rates is not an artifact of the functional form imposed by the sticky expectations model. Even when we allow the data to choose any smooth interest rate Jacobian with a free discontinuity along the diagonal, the estimated Jacobian implies near-zero consumption responses. The credible intervals in the non-parametric estimation are wider than those from the structural estimation, as expected: the structural estimation imposes that all $H^2 = 256$ elements of the Jacobian are determined by a single parameter θ^R, whereas the non-parametric approach estimates $p = 72$ basis coefficients. The additional degrees of freedom come at the cost of precision. Nevertheless, even with this substantially wider uncertainty, the posterior is centered near zero and clearly excludes the large negative consumption responses predicted by FIRE models.

Second, the basis function approach provides a natural way to encode economic priors without imposing a specific structural model. The two-region smooth structure is a minimal restriction: as shown in Figure 9, it holds for a standard FIRE model, a sticky expectations model, a finite-horizon planning model, and a model with myopia. By centering the prior on the FIRE model and allowing the data to pull the posterior away from it, the estimation allows us to assess how strongly the data favor a Jacobian different from the FIRE benchmark. As we show here, the data clearly reject the FIRE model and are consistent with the structural model estimate of $\theta^R = 1$.

Finally, the approach connects to a broader literature on non-parametric estimation in macroeconomics. The use of basis-function approximations to estimate flexible functions while imposing minimal smoothness restrictions has a long tradition in semi-parametric econometrics (Newey, 1997; Chen, 2007), and Barnichon and Brownlees (2019) apply this idea to macroeconomic impulse responses by using B-spline smoothing of local projections. The Bayesian formulation provides a natural way to quantify uncertainty over the estimated Jacobian and its implied consumption responses.

6. Conclusion

Across ten externally identified macroeconomic shocks and two estimation strategies, we show household consumption does not respond directly to interest rates at any horizon. The headline implication is for the modeling of household behavior in heterogeneous-agent New Keynesian models: Even when the consumer block is calibrated to match the best available microdata on intertemporal MPCs (Fagereng, Holm and Natvik, 2021), the resulting model substantially overstates the direct interest-rate channel relative to the macro data. Heterogeneity and borrowing constraints alone—the dampening mechanisms emphasized in Kaplan, Moll and Violante (2018)—are not enough. To rationalize the empirical impulse responses, an additional friction must be layered on top: in our structural estimation it takes the form of extremely sticky expectations over the interest rate (Carroll et al., 2020); in our non-parametric estimation it takes the form of a Jacobian with all elements near zero.

Our findings align with a long literature documenting weak interest rate sensitivity in household consumption. The near-rationality literature (Akerlof and Yellen, 1985; Cochrane, 1989) shows that small deviations from optimal consumption-saving behavior generate negligible utility losses, suggesting households may rationally choose not to optimize along the interest rate margin. This result finds anecdotal support in Choi (2022), who documents that popular financial advice provides no guidance on adjusting consumption in response to interest rates—suggesting most households do not consider this margin operationally relevant. Our structural and non-parametric estimates formalize this intuition: the data require household consumption to be essentially invariant to interest rate changes.

If household consumption does not respond directly to interest rates, then investment likely carries the weight of monetary policy's effect on aggregate demand. The labor income channel that operates in our framework is itself downstream of firm production and hiring decisions, which in standard models respond to the cost of capital and to demand for capital goods. The transmission of monetary policy to consumption, in this view, runs from interest rates to investment to employment and labor income, and only then to household spending. This finding suggests the importance of future research that better incorporates firm behavior into New Keynesian models (Winberry et al., 2025).

Our results also bear on the forward guidance puzzle. A large literature, beginning with McKay, Nakamura and Steinsson (2016) and including Gabaix (2020), Angeletos and Lian (2018), Farhi and Werning (2019), and Woodford (2019), has sought to dampen the implausibly large consumption response to forward guidance by attenuating how households respond to news about future interest rates. Our empirical evidence suggests households do not react much to interest rate changes at any horizon, which is inconsistent with the large forward-guidance effects predicted by standard New Keynesian models.

Several caveats are in order. Our framework abstracts from a richer asset structure—in particular, the distinction between liquid and illiquid assets and the specific role of housing and mortgage debt emphasized in recent empirical work (Wong, 2021; Cloyne, Ferreira and Surico, 2020; Foulis et al., 2026). How a two-asset extension would interact with our results is not completely clear, and we leave it to future research. A second concern is that our identification relies on the MIT shock linearization, which abstracts from movements in risk premia potentially correlated with the shocks we use. We do not believe this would overturn our central finding, and a treatment that takes risk premia seriously is a useful direction for further work.

References

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

Alvarez, Fernando, Luigi Guiso, and Francesco Lippi. 2012. "Durable Consumption and Asset Management with Transaction and Observation Costs." *American Economic Review*, 102(5): 2272–2300.

Angeletos, George-Marios, and Chen Lian. 2018. "Forward Guidance without Common Knowledge." *American Economic Review*, 108(9): 2477–2512.

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

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2020. "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model." National Bureau of Economic Research NBER Working Paper 26647.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2024. "The Intertemporal Keynesian Cross." *Journal of Political Economy*, 132(12): 4068–4121.

Barnichon, Régis, and Christian Brownlees. 2019. "Impulse Response Estimation by Smooth Local Projections." *The Review of Economics and Statistics*, 101(3): 522–530.

Barnichon, Régis, and Geert Mesters. 2023. "A Sufficient Statistics Approach for Macro Policy." *American Economic Review*, 113(11): 2809–2845.

Barsky, Robert B., and Eric R. Sims. 2011. "News Shocks and Business Cycles." *Journal of Monetary Economics*, 58(3): 273–289.

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

Bayer, Christian, Benjamin Born, and Ralph Luetticke. 2024. "Shocks, Frictions, and Inequality in US Business Cycles." *American Economic Review*, 114(5): 1211–1247.

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

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

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

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

Carlstrom, Charles T., Timothy S. Fuerst, and Matthias Paustian. 2015. "Inflation and Output in New Keynesian Models with a Transient Interest Rate Peg." *Journal of Monetary Economics*, 76: 230–243.

Carroll, Christopher D. 1997. "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)." National Bureau of Economic Research Working Paper 6298.

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

Chen, Xiaohong. 2007. "Large Sample Sieve Estimation of Semi-Nonparametric Models." In *Handbook of Econometrics*. Vol. 6B, , ed. James J. Heckman and Edward E. Leamer, Chapter 76, 5549–5632. Elsevier.

Chernozhukov, Victor, and Han Hong. 2003. "An MCMC Approach to Classical Estimation." *Journal of Econometrics*, 115(2): 293–346.

Chib, Siddhartha, Minchul Shin, and Anna Simoni. 2018. "Bayesian Estimation and Comparison of Moment Condition Models." *Journal of the American Statistical Association*, 113(524): 1656–1668.

Chodorow-Reich, Gabriel, Plamen T. Nenov, and Alp Simsek. 2021. "Stock Market Wealth and the Real Economy: A Local Labor Market Approach." *American Economic Review*, 111(5): 1613–1657.

Choi, James J. 2022. "Popular Personal Financial Advice versus the Professors." *Journal of Economic Perspectives*, 36(4): 167–192.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 113(1): 1–45.

Christiano, Lawrence J., Mathias Trabandt, and Karl Walentin. 2011. "DSGE Models for Monetary Policy Analysis." In *Handbook of Monetary Economics*. Vol. 3A, , ed. Benjamin M. Friedman and Michael Woodford, Chapter 7, 285–367. Amsterdam:Elsevier.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary Policy when Households Have Debt: New Evidence on the Transmission Mechanism." *The Review of Economic Studies*, 87(1): 102–129.

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

Del Negro, Marco, Marc P. Giannoni, and Christina Patterson. 2023. "The Forward Guidance Puzzle." *Journal of Political Economy Macroeconomics*, 1(1): 43–79.

Di Maggio, Marco, Amir Kermani, and Kaveh Majlesi. 2020. "Stock Market Returns and Consumption." *The Journal of Finance*, 75(6): 3175–3219.

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

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

Farhi, Emmanuel, and Iván Werning. 2019. "Monetary Policy, Bounded Rationality, and Incomplete Markets." *American Economic Review*, 109(11): 3887–3928.

Fernald, John G. 2014. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Federal Reserve Bank of San Francisco Working Paper 2012-19.

Foulis, Angus K., Jonathon Hazell, Atif R. Mian, and Belinda Tracey. 2026. "How Do Interest Rates Affect Consumption? Household Debt and the Role of Asset Prices." National Bureau of Economic Research NBER Working Paper 34911.

Francis, Neville, Michael T. Owyang, Jennifer E. Roush, and Riccardo DiCecio. 2014. "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.

Gabaix, Xavier. 2020. "A Behavioral New Keynesian Model." *American Economic Review*, 110(8): 2271–2327.

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

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

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. 2007. "The U.S. Treasury Yield Curve: 1961 to the Present." *Journal of Monetary Economics*, 54(8): 2291–2304.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2021. "What Do Data on Millions of U.S. Workers Reveal about Lifecycle Earnings Dynamics?" *Econometrica*, 89(5): 2303–2339.

Hagedorn, Marcus, Jinfeng Luo, Iourii Manovskii, and Kurt Mitman. 2019. "Forward Guidance." *Journal of Monetary Economics*, 102: 1–23.

Hall, Robert E. 1988. "Intertemporal Substitution in Consumption." *Journal of Political Economy*, 96(2): 339–357.

Havránek, Tomáš. 2015. "Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting." *Journal of the European Economic Association*, 13(6): 1180–1204.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2017. "Optimal Tax Progressivity: An Analytical Framework." *The Quarterly Journal of Economics*, 132(4): 1693–1754.

Hebden, James, and Fabian Winkler. 2021. "Impulse-Based Computation of Policy Counterfactuals." Board of Governors of the Federal Reserve System Finance and Economics Discussion Series 2021-042.

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

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

Kim, Jae-Young. 2002. "Limited Information Likelihood and Bayesian Analysis." *Journal of Econometrics*, 107(1–2): 175–193.

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

Leeper, Eric M., Alexander W. Richter, and Todd B. Walker. 2012. "Quantitative Effects of Fiscal Foresight." *American Economic Journal: Economic Policy*, 4(2): 115–144.

McKay, Alisdair, and Christian K. Wolf. 2023. "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" *Econometrica*, 91(5): 1695–1725.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson. 2016. "The Power of Forward Guidance Revisited." *American Economic Review*, 106(10): 3133–3158.

Mertens, Karel, and Morten O. Ravn. 2011. "Understanding the Aggregate Effects of Anticipated and Unanticipated Tax Policy Shocks." *Review of Economic Dynamics*, 14(1): 27–54.

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

Nakamura, Emi, and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics*, 133(3): 1283–1330.

Newey, Whitney K. 1997. "Convergence Rates and Asymptotic Normality for Series Estimators." *Journal of Econometrics*, 79(1): 147–168.

Newey, Whitney K., and Daniel McFadden. 1994. "Large Sample Estimation and Hypothesis Testing." In *Handbook of Econometrics*. Vol. 4, , ed. Robert F. Engle and Daniel L. McFadden, Chapter 36, 2111–2245. Elsevier.

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

Ramey, Valerie A. 2011. "Identifying Government Spending Shocks: It's All in the Timing." *Quarterly Journal of Economics*, 126(1): 1–50.

Ramey, Valerie A. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics*. Vol. 2, , ed. John B. Taylor and Harald Uhlig, 71–162. Elsevier.

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

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

Winberry, Thomas, Adrien Auclert, Matthew Rognlie, and Ludwig Straub. 2025. "New Keynesian Economics with Household and Firm Heterogeneity." National Bureau of Economic Research, Inc NBER Working Paper 34611.

Wong, Arlene. 2021. "Refinancing and the Transmission of Monetary Policy to Consumption." Princeton University, Department of Economics Working Paper 2021-57.

Woodford, Michael. 2019. "Monetary Policy Analysis when Planning Horizons are Finite." *NBER Macroeconomics Annual*, 33(1): 1–50.

Online Appendix

A. Sticky expectations

Following Carroll et al. (2020), we model households as having sticky expectations over the future paths of aggregate variables. Each period, a fraction $(1-\theta^X)$ of households update their beliefs about the path of aggregate variable $X \in \{R, Y^L\}$, while the remaining fraction θ^X retain their beliefs from the previous period. The probability of updating is independent over time, across households, and across X.

To compute the sticky-expectations-adjusted Jacobians, we use the recursive formula from Auclert, Rognlie and Straub (2020). Let $\mathcal{J}^{o,i,\text{FIRE}}$ denote the full-information rational expectations Jacobian relating output o to input i, and let $\mathcal{J}^{o,i}$ denote the sticky-expectations-adjusted Jacobian. The element $\mathcal{J}^{o,i}_{t,s}$ gives the response at horizon t to a shock at horizon s. The adjustment formula (dropping the X from θ^X for readability) is:

$$\mathcal{J}^{o,i}_{t,s} = \begin{cases} \mathcal{J}^{o,i,\text{FIRE}}_{t,0} & s=0,\ \text{column 0 not sticky } (X = Y^L), \\ (1-\theta)\,\mathcal{J}^{o,i,\text{FIRE}}_{t,0} + \displaystyle\sum_{j=0}^{t-1} \theta^{j+1}(1-\theta)(1+r)^{j+1}\,\mathcal{J}^{o,i,\text{FIRE}}_{t-j-1,0} & s=0,\ \text{column 0 sticky } (X = R), \\ (1-\theta)\,\mathcal{J}^{o,i,\text{FIRE}}_{t,s} & s>0,\ t=0, \\ (1-\theta)\,\mathcal{J}^{o,i,\text{FIRE}}_{t,s} + \theta\,\mathcal{J}^{o,i}_{t-1,s-1} & s>0,\ t>0. \end{cases} \tag{36}$$

Equivalently, the $s=0$, column 0 sticky case admits the simpler recursion

$$\mathcal{J}^{o,i}_{t,0} = \theta(1+r)\,\mathcal{J}^{o,i}_{t-1,0} + (1-\theta)\,\mathcal{J}^{o,i,\text{FIRE}}_{t,0}, \qquad \mathcal{J}^{o,i}_{0,0} = (1-\theta)\,\mathcal{J}^{o,i,\text{FIRE}}_{0,0}. \tag{37}$$

The $s>0$ cases follow Auclert, Rognlie and Straub (2020) directly: when a household updates at date $t_u \leq s$, it learns the path of the input before the shock arrives and from then on behaves as a fully-informed household with information dated t_u periods later; weighting by the i.i.d. update probability $(1-\theta)\theta^{t_u}$ and collecting terms gives the recursion in the last row of (36).

The $s=0$ rows differ for income and interest rates, reflecting how the contemporaneous shock interacts with the household's information set. For labor income, households observe the income they have received and consume out of it regardless of whether they update expectations of the future income path; column 0 is therefore not subject to sticky expectations and equals its FIRE counterpart. For the period-0 portfolio return, by contrast, an unobserved date-0 revaluation continues to affect the household's wealth even before they notice: a household that does not update until date $t_u = j+1$ has been carrying an unrecognized wealth perturbation that compounds at the steady-state rate, so the effective shock at the moment of awakening is $(1+r)^{j+1}$ times the date-0 revaluation. Their forward-looking response from t_u onward inherits this scaling, generating the $(1+r)^{j+1}$ factor in (36) and the $\theta(1+r)$ growth term in (37). Salience plays a role here: the cash arrival from a labor-income shock is highly visible, while a revaluation of financial wealth need not be (For survey evidence on how often households check their financial assets, see Alvarez, Guiso and Lippi (2012)).

B. Income process

The idiosyncratic earnings process follows the two-component jump-drift specification of Kaplan, Moll and Violante (2018). Log earnings are the sum of two independent Ornstein–Uhlenbeck processes with Poisson jumps: a transitory component that arrives frequently (roughly once every 12.5 quarters) and mean-reverts quickly (half-life of about one quarter), and a persistent component that arrives rarely (roughly once every 38 years) and mean-reverts slowly (half-life of about 18 years). Each component evolves in continuous time according to

$$dz_{j,it} = -\beta_j z_{j,it}\, dt + \varepsilon_{j,it}\, dN_{j,it}, \qquad \varepsilon_{j,it} \sim \mathcal{N}(0, \sigma_j^2), \tag{38}$$

where $dN_{j,it}$ is a Poisson counter with arrival rate λ_j and the jump size is i.i.d. Gaussian. The six parameters $(\lambda_j, \beta_j, \sigma_j)_{j=1,2}$ are estimated by Kaplan, Moll and Violante (2018) via simulated method of moments to match eight annual earnings moments from Social Security Administration data, as reported in Guvenen et al. (2021).

We discretize this process onto a 33-state grid (3 transitory × 11 persistent). Both household groups use the same discretized earnings process, differing only in their discount factors. We apply a progressive tax transformation $y^{\text{after-tax}} = \left(y^{\text{pre-tax}}\right)^{1-\tau}$, where $\tau = 0.181$ is the tax progressivity parameter from Heathcote, Storesletten and Violante (2017); this compresses the after-tax earnings distribution relative to the pre-tax distribution. We thank Bence Bardóczy for helpful advice on implementing this discretization.

C. Alternative structural models

We consider two alternative models of household expectations formation beyond sticky expectations: finite horizon planning and myopia. Both models modify the FIRE Jacobians to capture different departures from full rationality, and both preserve the smooth-within-regions structure that motivates the basis function approach in Section 5.

Finite horizon planning Following Woodford (2019), we model households as having finite planning horizons. At date t, a fraction ρ^s of households rationally anticipate shocks at horizon s; the remaining fraction $1 - \rho^s$ do not incorporate information about horizon-s shocks until later. The parameter $\rho \in [0, 1]$ governs how rapidly the fraction of forward-looking households decays with the forecast horizon. When $\rho = 1$, all households have infinite planning horizons (FIRE); when $\rho = 0$, no household looks beyond the current period.

The finite-horizon-planning-adjusted Jacobian $\mathcal{J}^{o,i,\text{FHP}}$ relates to the FIRE Jacobian $\mathcal{J}^{o,i,\text{FIRE}}$ through:

$$\mathcal{J}^{o,i,\text{FHP}}_{t,s} = \begin{cases} \mathcal{J}^{o,i,\text{FIRE}}_{t,0}, & \text{if } s = 0, \\ \rho^s \, \mathcal{J}^{o,i,\text{FIRE}}_{t,s} + (1 - \rho^s) \, \mathcal{J}^{o,i,\text{FHP}}_{t-1,s-1}, & \text{if } s > 0 \text{ and } t > 0. \end{cases} \tag{39}$$

The first column ($s = 0$) is unchanged: all households observe current-period shocks. For $s > 0$, the response at horizon (t, s) is a weighted average of the immediate response of the fraction ρ^s who anticipated the shock, and the lagged response of the fraction $1 - \rho^s$ who only learned about it in the previous period.

Myopia The myopia model captures households who respond to news about future shocks as if those shocks were arriving immediately. A fraction $\mu \in [0, 1]$ of households are myopic: they respond to all future shocks using the column-0 (immediate-arrival) Jacobian. The remaining fraction $1 - \mu$ have full rational expectations. The myopia-adjusted Jacobian is:

$$\mathcal{J}^{o,i,\text{myopic}}_{t,s} = (1 - \mu) \, \mathcal{J}^{o,i,\text{FIRE}}_{t,s} + \mu \, \mathcal{J}^{o,i,\text{FIRE}}_{t-s,0}, \tag{40}$$

for $t \geq s$, and zero otherwise (myopic households do not respond in advance). When $\mu = 0$, the model reduces to FIRE; when $\mu = 1$, all households are fully myopic and treat anticipated shocks as if they had already occurred.

Both models generate Jacobians that are smooth within the causal ($t \geq s$) and anticipatory ($t < s$) regions, with a discontinuity along the diagonal—the minimal regularity structure imposed by the basis function approach.

D. Impulse responses and posterior distributions for all shocks

This appendix presents detailed results for each of the ten structural shocks used in our analysis. For each shock, we show three figures: (1) the estimated impulse responses for consumption, income, the interest rate, and portfolio returns; (2) the marginal posterior distributions for the sticky expectations parameters θ^Y

and θ^R; and (3) a comparison of the empirical consumption impulse response with model-based simulations under different assumptions about expectations formation (FIRE, sticky expectations for both variables, and sticky expectations for only one variable).

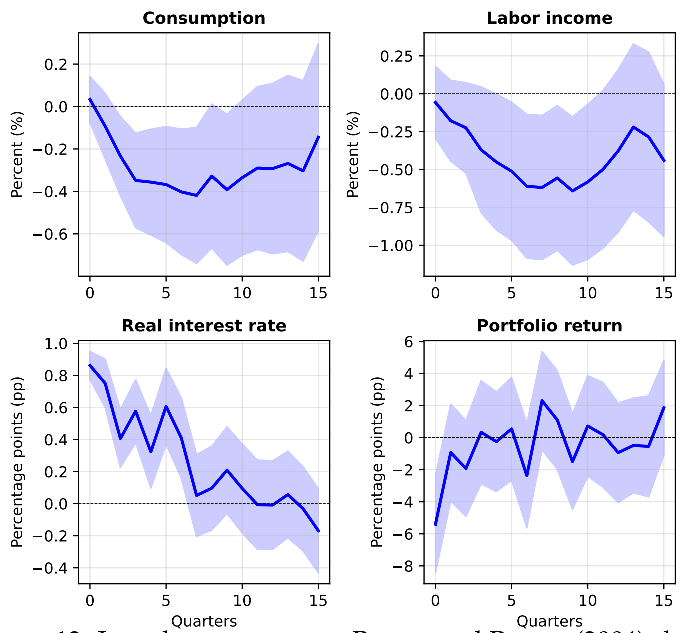

Figure 12: Impulse responses to Romer and Romer (2004) shock

(a) Interest rate

(b) Aggregate labor income

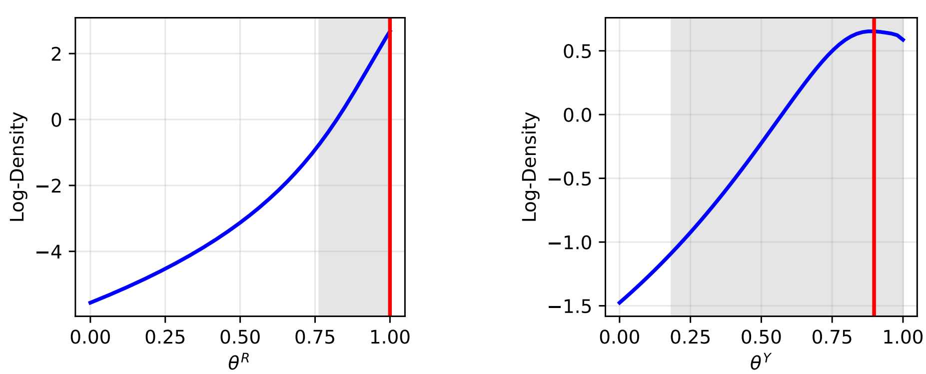

Figure 13: Marginal posterior distributions: Romer and Romer (2004) shock

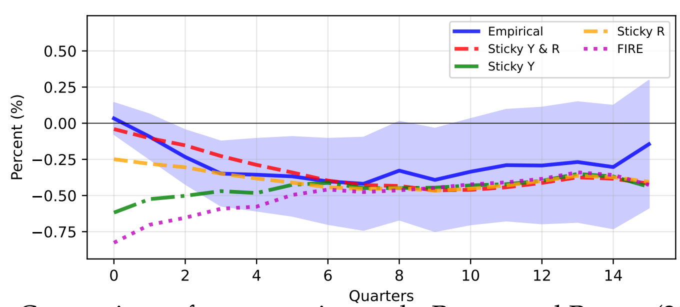

Figure 14: Comparison of consumption path: Romer and Romer (2004) shock

D.1. Romer and Romer (2004).

Figure 15: Impulse responses to Gertler and Karadi (2015) shock

(a) Interest rate

(b) Aggregate labor income

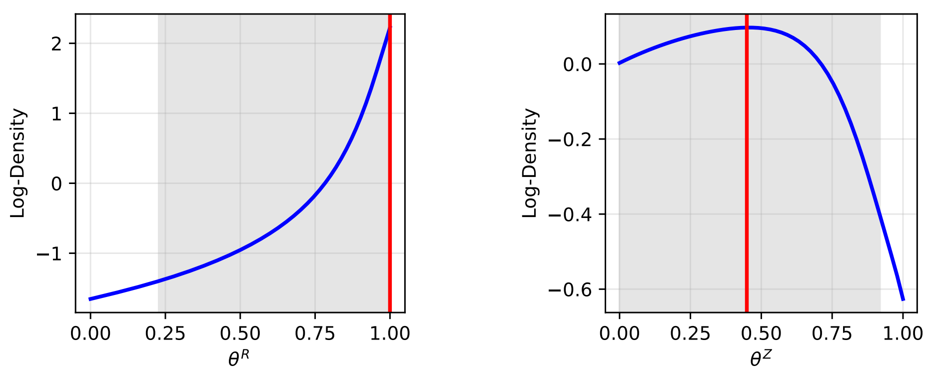

Figure 16: Marginal posterior distributions: Gertler and Karadi (2015) shock

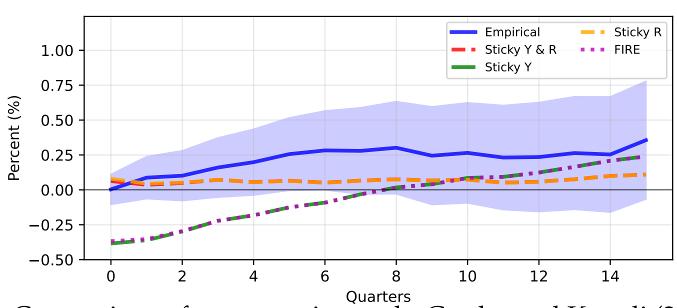

Figure 17: Comparison of consumption path: Gertler and Karadi (2015) shock

D.2. Gertler and Karadi (2015).

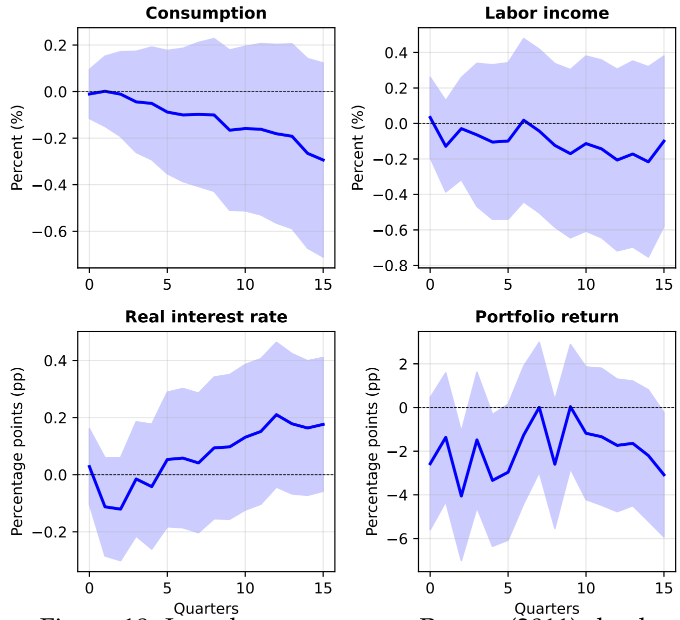

Figure 18: Impulse responses to Ramey (2011) shock

(a) Interest rate

(b) Aggregate labor income

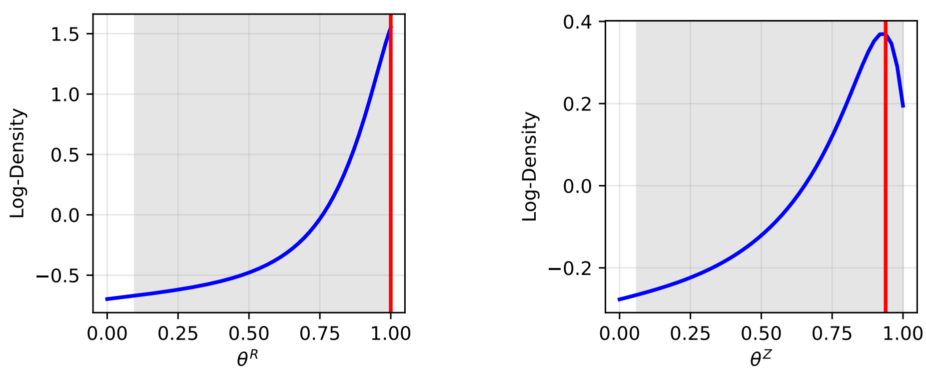

Figure 19: Marginal posterior distributions: Ramey (2011) shock

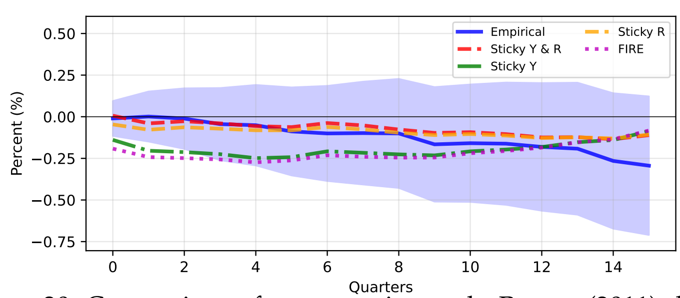

Figure 20: Comparison of consumption path: Ramey (2011) shock

D.3. Ramey (2011).

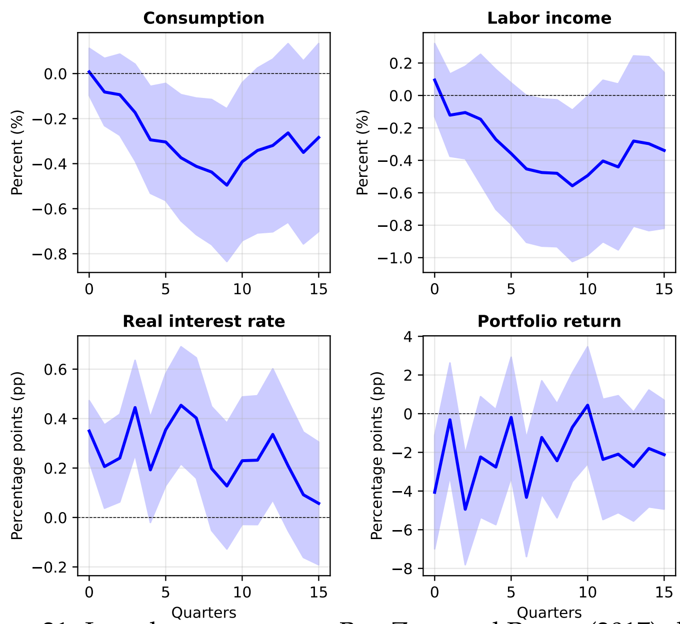

Figure 21: Impulse responses to Ben Zeev and Pappa (2017) shock

(a) Interest rate

(b) Aggregate labor income

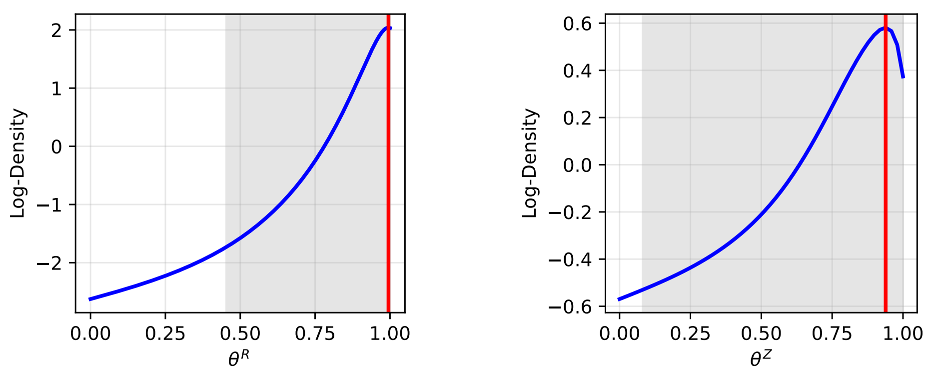

Figure 22: Marginal posterior distributions: Ben Zeev and Pappa (2017) shock

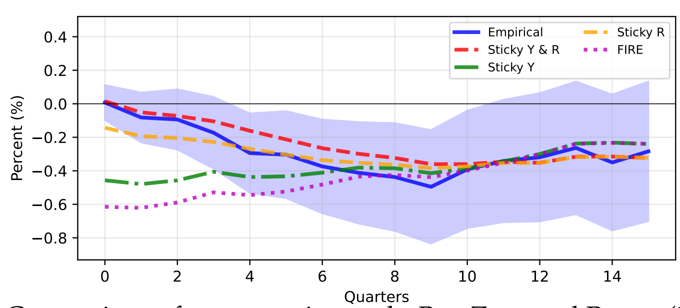

Figure 23: Comparison of consumption path: Ben Zeev and Pappa (2017) shock

D.4. Ben Zeev and Pappa (2017).

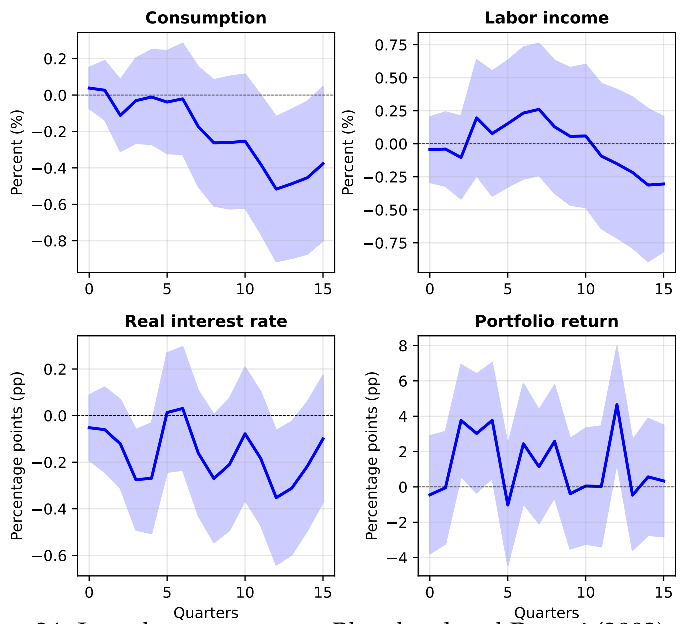

Figure 24: Impulse responses to Blanchard and Perotti (2002) shock

(a) Interest rate (b) Aggregate labor income

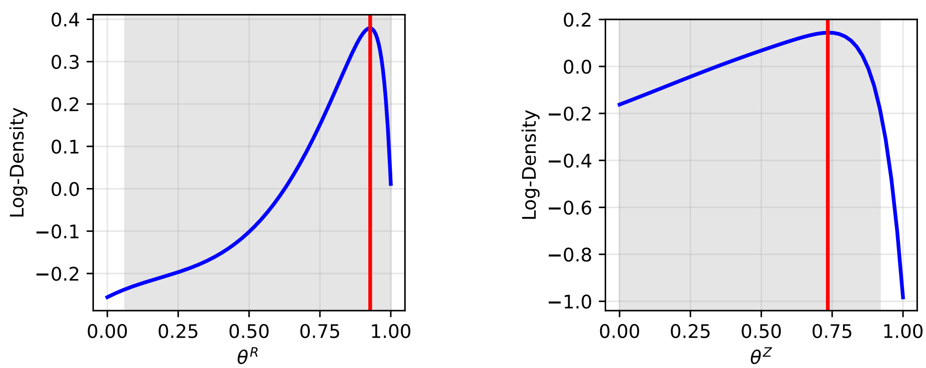

Figure 25: Marginal posterior distributions: Blanchard and Perotti (2002) shock

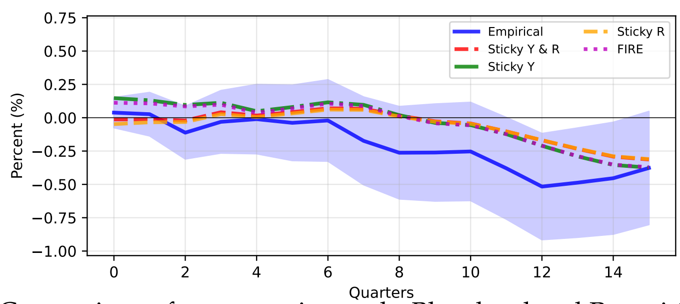

Figure 26: Comparison of consumption path: Blanchard and Perotti (2002) shock

D.5. Blanchard and Perotti (2002).

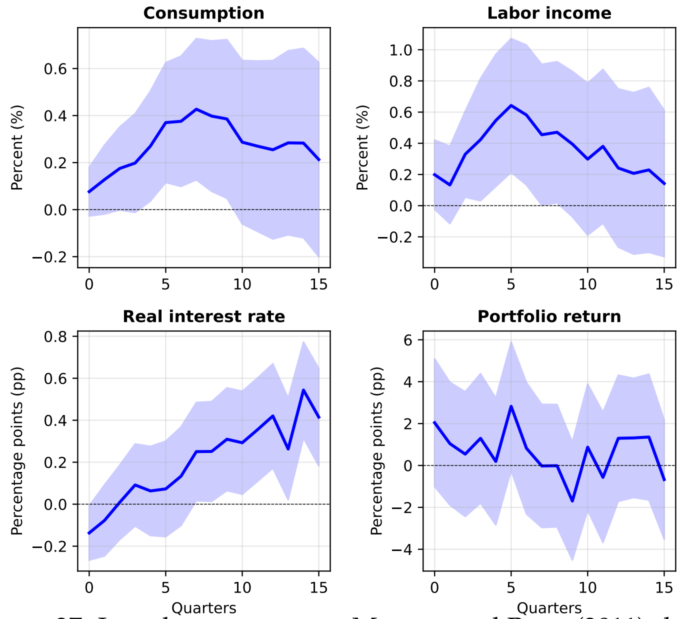

Figure 27: Impulse responses to Mertens and Ravn (2011) shock

(a) Interest rate

(b) Aggregate labor income

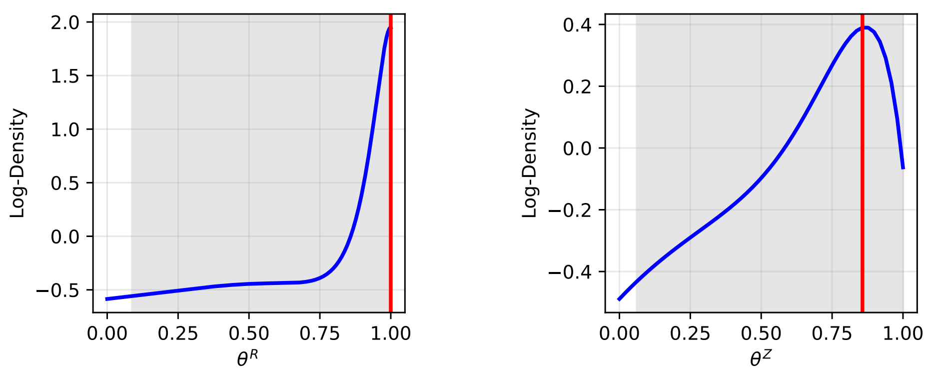

Figure 28: Marginal posterior distributions: Mertens and Ravn (2011) shock

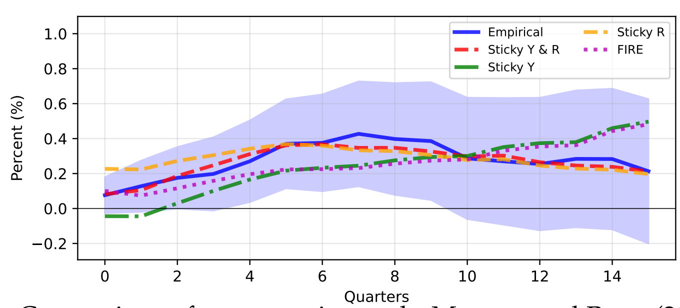

Figure 29: Comparison of consumption path: Mertens and Ravn (2011) shock

D.6. Mertens and Ravn (2011).

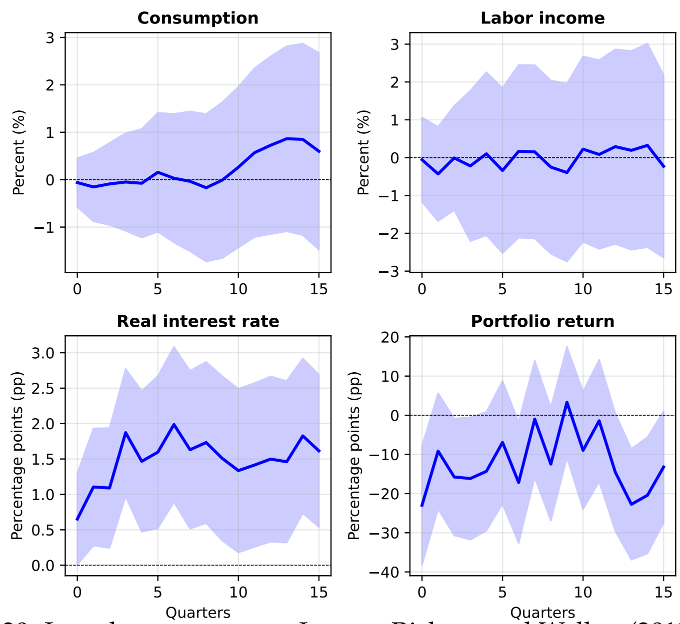

Figure 30: Impulse responses to Leeper, Richter, and Walker (2012) shock

(a) Interest rate

(b) Aggregate labor income

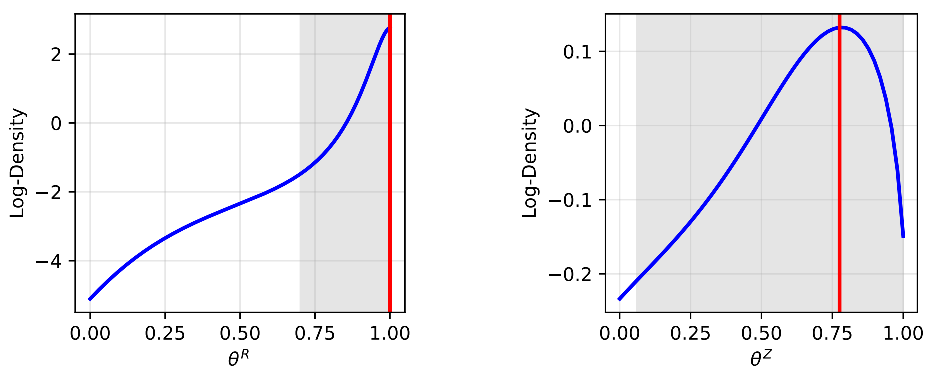

Figure 31: Marginal posterior distributions: Leeper, Richter, and Walker (2012) shock

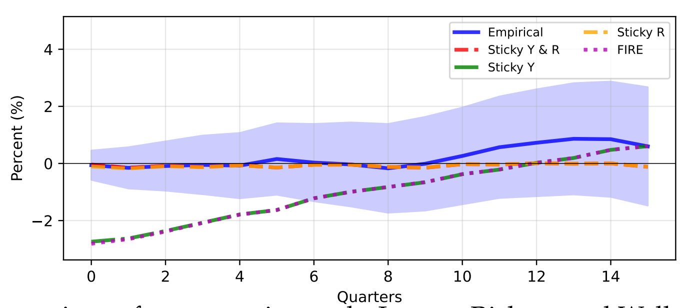

Figure 32: Comparison of consumption path: Leeper, Richter, and Walker (2012) shock

D.7. Leeper, Richter, and Walker (2012).

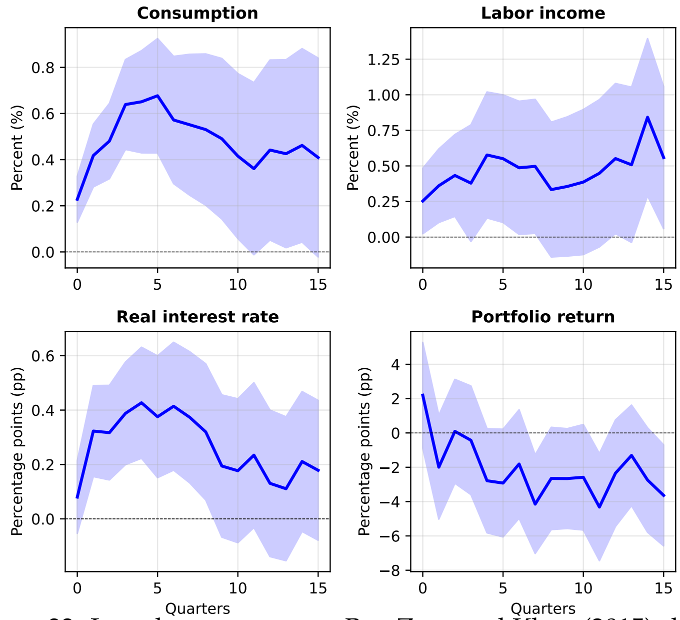

Figure 33: Impulse responses to Ben Zeev and Khan (2015) shock

(a) Interest rate

(b) Aggregate labor income

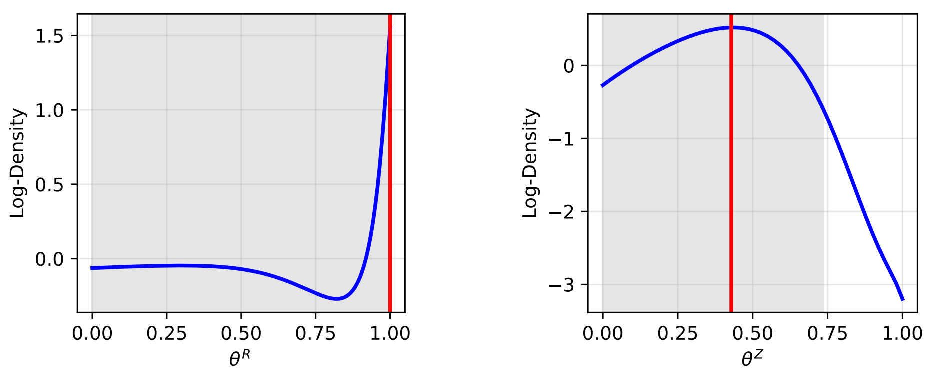

Figure 34: Marginal posterior distributions: Ben Zeev and Khan (2015) shock

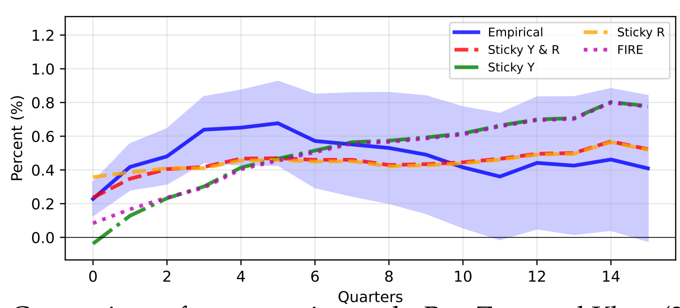

Figure 35: Comparison of consumption path: Ben Zeev and Khan (2015) shock

D.8. Ben Zeev and Khan (2015).

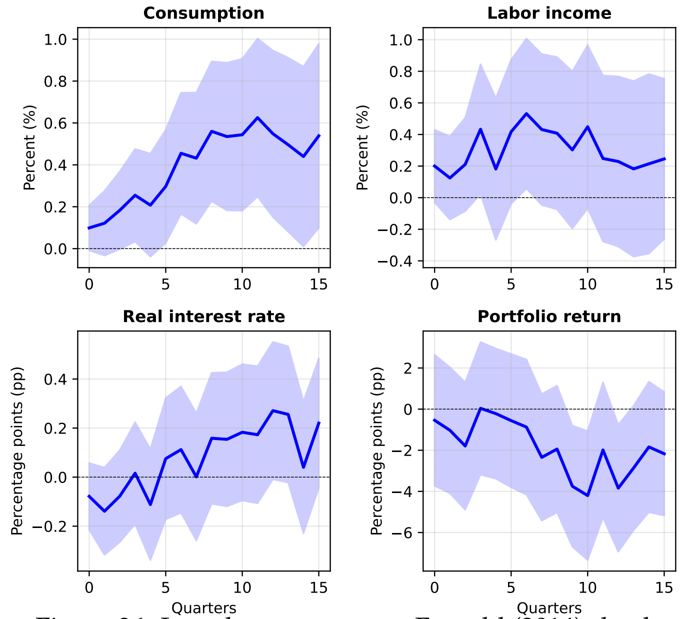

Figure 36: Impulse responses to Fernald (2014) shock

(a) Interest rate

(b) Aggregate labor income

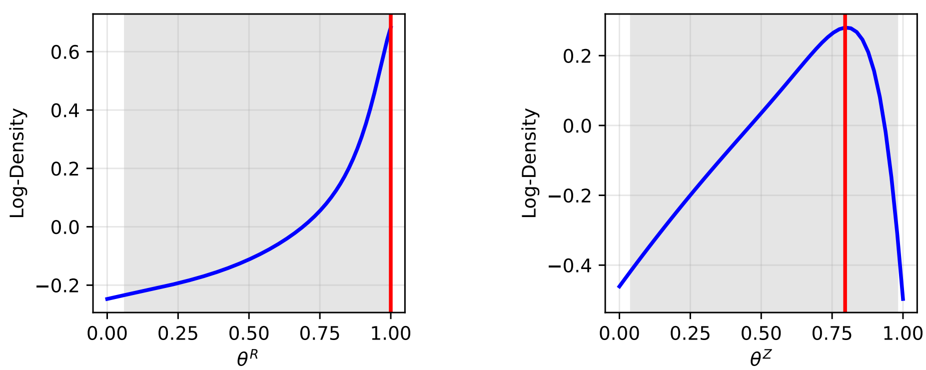

Figure 37: Marginal posterior distributions: Fernald (2014) shock

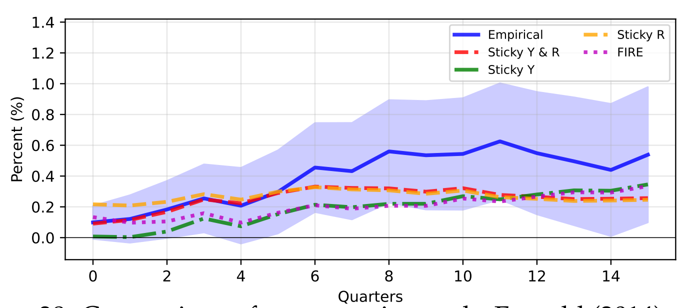

Figure 38: Comparison of consumption path: Fernald (2014) shock

D.9. Fernald (2014).

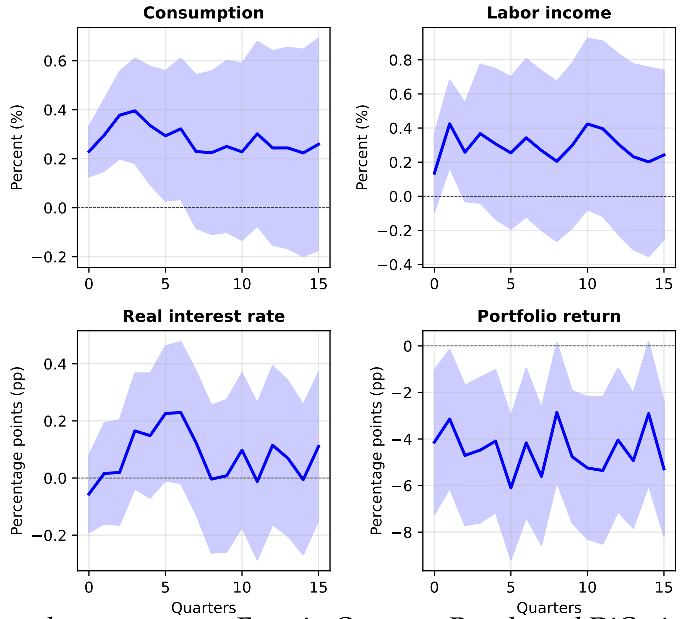

Figure 39: Impulse responses to Francis, Owyang, Roush, and DiCecio (2014) shock

(a) Interest rate

(b) Aggregate labor income

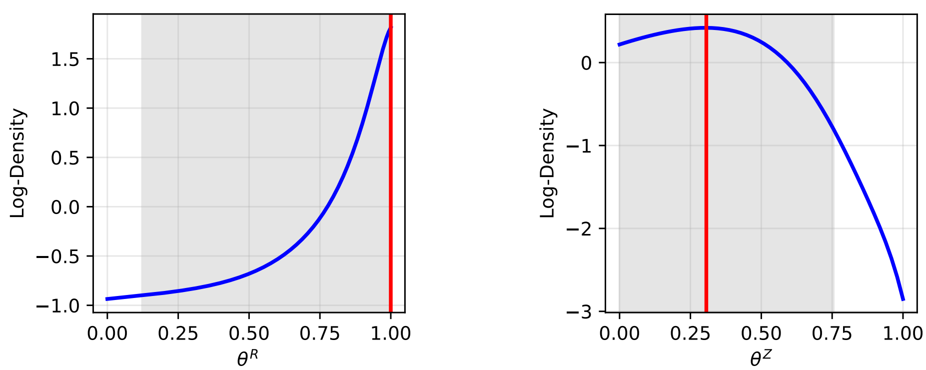

Figure 40: Marginal posterior distributions: Francis, Owyang, Roush, and DiCecio (2014) shock

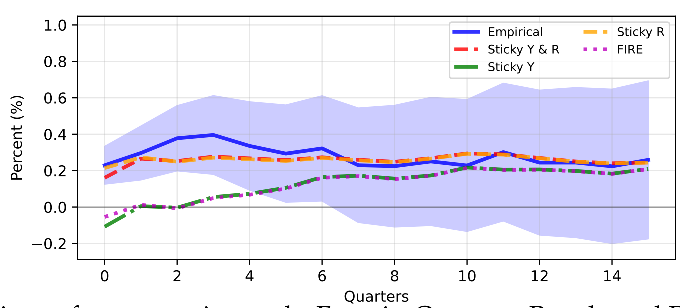

Figure 41: Comparison of consumption path: Francis, Owyang, Roush, and DiCecio (2014) shock

D.10. Francis, Owyang, Roush, and DiCecio (2014).