

Lumpy Durable Consumption Demand and the Limited Ammunition of Monetary Policy

Alisdair McKay

Federal Reserve Bank of Minneapolis

Johannes F. Wieland*

U.C. San Diego, NBER

June 26, 2020

Abstract

The prevailing neo-Wicksellian view holds that the central bank's objective is to track the natural rate of interest (r^*), which itself is largely exogenous to monetary policy. We challenge this view using a fixed-cost model of durable consumption demand, in which expansionary monetary policy prompts households to accelerate the timing of lumpy durable adjustments. This yields an intertemporal trade-off in aggregate demand as encouraging households to adjust today leaves fewer households acquiring durables going forward. Interest rates must be kept low to support demand going forward, so accommodative monetary policy today reduces r^* in the future. We show that this mechanism is quantitatively important in explaining the persistently low level of real interest rates and r^* after the Great Recession.

JEL Classification: E21, E43, E52

*We are grateful to Adrien Auclert, Robert Barsky, David Berger, Jeff Campbell, Adam Guren, Jim Hamilton, Christopher House, Rohan Kekre, Emi Nakamura, Valerie Ramey, Jón Steinsson, Ludwig Straub, Stephen Terry, Joe Vavra, Venky Venkateswaran, Tom Winberry, Christian Wolf, and seminar participants at Boston University, University of Michigan, Harvard University, the Federal Reserve Banks of Boston, Minneapolis, San Francisco, and St. Louis, the 2019 ASSA meetings, U.C. Berkeley, Chicago Booth, LSE, IIES, SED 2019, and the 2019 NBER Summer Institute (ME). Charles Fries provided excellent research assistance. Wieland is grateful to the Federal Reserve Bank of Chicago for its hospitality while completing parts of this paper. The views expressed here are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of Minneapolis, the Federal Reserve Bank of Chicago, or the Federal Reserve System.

1 Introduction

When entering a recession, the first tool in the arsenal of macroeconomic policymakers is to lower interest rates. Lower real interest rates encourage businesses to invest and consumers to spend, which bolsters aggregate demand. An important component of this monetary transmission mechanism is to stimulate purchases of durable goods, which are particularly sensitive to interest rates (e.g., Mishkin, 1995; Barsky et al., 2003; Sterk and Tenreyro, 2018). In this paper we argue that stimulating demand for durable goods has additional consequences. As monetary stimulus increases the stock of durables today, there is less need to acquire durable goods in the future, all else equal. Monetary policy therefore raises aggregate demand today by borrowing demand from the future. To compensate for the weakness in aggregate demand going forward, the central bank must keep real interest rates low. That is, monetary policy stimulus has a side effect of reducing the real natural rate of interest (r^*) in subsequent periods.

This interaction between monetary policy and r^* is very different from the prevailing neo-Wicksellian view (Woodford, 2003) that r^* is largely exogenous to monetary policy and the central bank aims to manipulate the policy rate to track r^* .¹ In contrast, we argue that monetary policy has a powerful impact on the future evolution of r^* through the intertemporal shifting of aggregate demand.

We show that this intertemporal shifting is an important piece of the monetary transmission mechanism in a heterogeneous agent New Keynesian model in which households accumulate durable consumption goods subject to fixed adjustment costs. Households optimally follow an (S,s) policy, making lumpy durable purchases as their existing durable stock drifts down and hits an adjustment threshold. Expansionary monetary policy shifts the adjustment thresholds, accelerating adjustments by those who were close to an adjustment threshold. For instance, low interest rates may prompt some households to accelerate the purchase of a new car. In the subsequent periods, they no longer need to purchase a car as they have already done so. As a result, aggregate demand is weaker in periods following the stimulus.

The dynamics of demand for durable goods create a propagation mechanism that makes changes in real interest rates very persistent. Low interest rates in the past cause low interest rates in the

¹See Woodford (2003, p. 49): “In Wicksell’s view, price stability depended on keeping the interest rate controlled by the central bank in line with the natural rate determined by real factors (such as the marginal product of capital). [...] Wicksell’s approach is a particularly congenial one for thinking about our present circumstances [...]”

present, which cause low interest rates in the future. We use our model to construct a forecast for the evolution of interest rates following the Great Recession, in which the Federal Reserve engaged in massive countercyclical monetary stimulus. Based on information through 2012Q4, our model predicts a path of interest rates that largely tracks the path that came to pass over the next seven years. The model predicts liftoff from the ELB in 2015Q4 and predicts low levels of short-term and 5-year interest rates in 2019Q4 just as in the data. The slow normalization of interest rates reflects a persistent decline in r^* . We isolate the contribution of intertemporal shifting to the path of r^* and show that it is quantitatively important in explaining the large drop and, especially, the slow normalization of r^* .

In recent years, the low level of interest rates has received a lot of attention (Summers, 2015; Laubach and Williams, 2016). These low rates are generally thought to reflect secular phenomena such as demographic changes, slow trend productivity growth, and the rise in income inequality (Eggertsson et al., 2019; Auclert and Rognlie, 2018; Straub, 2018). Our results, however, demonstrate that cyclical forces can have large and very persistent effects on the natural rate of interest, and that these forces have contributed substantially to the low interest rates over the last decade.

A fixed-cost model is a natural modeling approach to capture the lumpiness of durable adjustments in the micro-data. However, the nature of the adjustment costs we include in our model is also central to our main findings. The logic of our argument can actually be reversed in models with higher-order adjustment costs, which is a common formulation in the literature (e.g. Christiano et al., 2005). In these models, lagged investment is also a state variable. Low interest rates today stimulate investment today, which lowers the marginal cost of investment in the future. This effect works against the intertemporal shifting effect we highlight, whereby higher investment today increases the future capital stock, which reduces marginal benefit from investing in the future.² Thus, if higher order adjustment costs are large enough, low interest rates today may even raise future aggregate demand. Higher order adjustment costs help DSGE models to match certain features of the aggregate response of durable demand to interest rates. However, they are at odds with the micro data that shows lumpy adjustments in consumer durables and business investment.

We show that our model is consistent with both the adjustment process at the micro level and the aggregate response of durable demand to interest rates. In particular, we show that the

²Because Barsky et al. (2007) abstract from higher-order adjustment costs, intertemporal shifting also occurs in their model, but they do not emphasize or quantify this feature.

impulse response of aggregate durable spending to a monetary policy shock is similar to what we estimate using the Romer and Romer (2004) shocks. Notably, our point estimates for the responses of GDP, aggregate durable expenditure, and the extensive margins of car and housing adjustments all show clear reversals consistent with intertemporal shifting. Turning to cross-sectional evidence, anticipated changes in sales tax rates create incentives for intertemporal substitution similar to changes in interest rates (Correia et al., 2013). We exploit this observation to make use of cross-sectional evidence from Baker et al. (2019) on the response of auto sales to anticipated sales tax changes at the state level. Again, the response of auto sales shows a clear reversal with cumulative sales returning to zero shortly after the sales tax change. Again, our model tracks this impulse response quite closely.

The timing of durable purchases in standard fixed-cost models is highly sensitive to intertemporal incentives (see House, 2014). Reiter, Sveen, and Weinke (2013) argue that this property implies a counterfactually large investment response to monetary stimulus in a New Keynesian model with a relatively standard (S,s) model of investment demand. We show that including two particular ingredients in the model is important to match the empirical evidence mentioned above. Without these ingredients the model-implied response of durable demand to interest rates is an order of magnitude larger than our empirical benchmarks. First, operating costs are a component of the user cost of durables that is not sensitive to interest rates, which limits the shift in the (S,s) adjustment thresholds. Second, match-quality shocks introduce inframarginal adjustments, which reduces the mass of households near the adjustment thresholds. We use micro-data on durable adjustments to estimate the frequency of match-quality shocks. Our model also includes information rigidities in the style of Carroll et al. (2018) and Auclert, Rognlie, and Straub (2020), which allows it to match the delayed responses to monetary policy shocks that are often observed in aggregate data.

Our analysis of intertemporal shifting effects leads us to focus on the effect of current interest rates on future demand. The literature on forward guidance analyzes the reverse link: the effect of future interest rates on current demand. Standard monetary models predict implausibly large output responses to forward guidance (Del Negro, Giannoni, and Patterson, 2015; Carlstrom, Fuerst, and Paustian, 2015). In response to this “forward guidance puzzle,” one strand of literature has developed alternative models in which non-durable consumption is less forward looking (McKay, Nakamura, and Steinsson, 2016; Werning, 2015; Farhi and Werning, 2017; Gabaix, 2016; Angeletos

and Lian, 2018). We show that including durable goods subject to fixed costs weakens forward guidance considerably. The opportunity cost of accelerating the adjustment is the foregone interest over the time period that the adjustment is brought forward. As a result, contemporaneous interest rates are especially relevant for the timing decision, whereas forward rates are much less powerful in shifting aggregate demand. These findings can help resolve the forward guidance puzzle. Moreover, the limited power of forward guidance contributes to the large decline in r^* during and after the Great Recession because the expected path of interest rates had to move more to stabilize the economy.

Our work builds on a growing literature that models aggregate demand using rich micro-foundations for household consumption that are disciplined by micro-data.³ Most of this literature focuses its attention on the determination of non-durable consumption and abstracts from consumer durables. Our interest in durable goods is motivated by the fact that they are more sensitive to monetary policy and more cyclical than non-durable consumption. For example, durable goods account for two thirds of the response of total consumption to an identified monetary shock (see Section 3). Our partial-equilibrium household decision problem builds on Berger and Vavra (2015) adding match-quality shocks and operating costs to lower the interest elasticity of durable demand, as well as sticky information to delay the demand response, so that both are in line with our empirical benchmarks.

Our analysis is made possible by the recent advances in the computation of heterogeneous agent macro models, specifically we make use of and extend the powerful sequence-space Jacobian techniques developed by Auclert et al. (2019, 2020). We make two technical contributions. We show how to implement the Kalman filter to recover the shocks that generated the aggregate time series data using only impulse response functions and not relying on a state space representation of the model. We then extend this filtering algorithm to incorporate the ELB constraint, by allowing for a sequence of anticipated monetary news shocks. Second, we show how r^* can be immediately calculated from the impulse response functions of the model without solving an auxiliary flexible-price model as is typically done in DSGE models.

Contemporaneous work by Mian, Straub, and Sufi (2019) argues that intertemporal shifting of demand can also occur following monetary accommodation due to the accumulation of household

³See Guerrieri and Lorenzoni (2017), McKay and Reis (2016), Gornemann, Kuester, and Nakajima (2016), McKay, Nakamura, and Steinsson (2016), Auclert (2019), Kaplan, Moll, and Violante (2018), Werning (2015) among others.

debt. The mechanism we focus on is a different one. To illustrate the difference, non-homothetic preferences are central to the mechanism described by Mian et al. so that the reduction in demand from indebted households is not offset by an increase in demand from their creditors. In contrast, our mechanism works through durable holdings with homothetic preferences.

A recent strand of literature has analyzed how past interest rates affect the power of monetary policy (Berger, Milbradt, Tourre, and Vavra, 2018; Eichenbaum, Rebelo, and Wong, 2018). These papers argue that the prevalence of fixed-rate mortgages in the United States leads to less refinancing activity in response to an interest rate reduction if past interest rates were low because the interest rates on existing mortgages are already low. The intertemporal shifting of demand that we describe is conceptually different from this mechanism: we argue that demand is weak going forward not that policy is less effective. Intertemporal shifting does not occur in the context of refinancing because a mortgagor that refinances to a lower fixed-rate mortgage will continue to enjoy higher disposable income on an ongoing basis and need not reduce their consumption in the future.

The paper is organized as follows: Section 2 presents our model of demand for durable consumption goods; Section 3 shows that the model aligns with the evidence on the response of durable spending and output to changes in real interest rates, discusses the empirical evidence for intertemporal shifting effects, and explains the roles of match-quality shocks and operating costs; Section 4 describes the general equilibrium model with sticky wages; Section 5 documents the implications for the monetary transmission mechanism: intertemporal shifting and weaker forward guidance; Section 6 shows that these features of the monetary transmission mechanisms have important implications for the dynamics of interest rates during and after the Great Recession; Section 7 concludes.

2 Model of Durable Demand

We begin with the household's partial equilibrium decision problem, which forms the demand side of the model. Later we will embed this demand block into a sticky-wage monetary model.

2.1 Household's Problem

Households consume non-durable goods, c , and a service flow from durable goods, s . Household $i \in [0, 1]$ has preferences given by

$$E_{i0} \int_{t=0}^{\infty} e^{-\rho t} u(c_{it}, s_{it}) dt. \quad (1)$$

The service flow from durables is generated from the household's stock of durable goods d_{it} . For the most part we have $s_{it} = d_{it}$, but we will complicate this relationship below. The expectation is individual-specific due to information frictions, which we also describe below.

Households hold a portfolio of durables and liquid assets denoted a_{it} . When a household with pre-existing portfolio (a_{it}, d_{it}) adjusts its durable stock, it reshuffles its portfolio to (a'_{it}, d'_{it}) subject to the payment of a fixed cost such that

$$a'_{it} + p_t d'_{it} = a_{it} + (1 - f)p_t d_{it}, \quad (2)$$

where p_t is the relative price of durable goods in terms of non-durable goods, and $f p_t d_{it}$ is a fixed cost proportional to the value of the durable stock. Liquid savings pay a safe real interest rate r_t . The household is able to borrow against the value of the durable stock up to a loan-to-value (LTV) limit λ

$$a_{it} \geq -\lambda(1 - f)p_t d_{it}. \quad (3)$$

Borrowers pay real interest rate $r_t + r_t^b$, where r_t^b is an exogenous borrowing spread.

The stock of durables depreciates at rate δ . Following Bachmann, Caballero, and Engel (2013), a fraction χ of depreciation must be paid immediately in the form of maintenance expenditures. This maintenance reduces the drift rate of the durable stock so we have

$$\dot{d}_{it} = -(1 - \chi)\delta d_{it}, \quad (4)$$

where a dot over a variable indicates a time derivative. The household must also pay a flow cost of operating the durable stock equal to νd_{it} . Broadly speaking these operating costs reflect expenditures such as fuel, utilities, and taxes.⁴

⁴We assume that operating costs do not scale with the relative durable price (in contrast to maintenance costs), but our results are essentially identical if we operating costs are proportional to the relative price.

When a household does not adjust its durable stock, its liquid assets evolve according to

$$\dot{a}_{it} = r_t a_{it} + r_t^b a_{it} I_{\{a_{it} < 0\}} - c_{it} + y_{it} - (\chi \delta p_t + \nu) d_{it}. \quad (5)$$

Household income, y_{it} , is given by

$$y_{it} = Y_t z_{it} \quad (6)$$

where Y_t is aggregate income and z_{it} is the household's idiosyncratic income share, which we later interpret as idiosyncratic labor productivity. $\ln z_{it}$ follows the Ornstein-Uhlenbeck process

$$d \ln z_{it} = \rho_z \ln z_{it} dt + \sigma_z d\mathcal{W}_{it} + (1 - \rho_z) \ln \bar{z} dt, \quad (7)$$

where $d\mathcal{W}_{it}$ is a standard Brownian motion, $\rho_z < 0$ controls the degree of mean reversion of the income process, σ_z determines the variance of the income process, and \bar{z} is a constant such that $\int z_{it} di = 1$.

We allow for the possibility that households may occasionally adjust their durables because their existing durables are no longer a good match for them. These match-quality shocks are meant to capture unmodeled life events that leave the household wanting to adjust for reasons other than income fluctuations and depreciation. For example, a job offer in a distant city may prompt the household to move houses. Or a growing family may require a larger car. We assume that a household is in a good match when it adjusts its durables, but over time the match quality can break down according to a Poisson process with intensity θ . Specifically, there is a state q_{it} that takes a value 1 when the household adjusts its durables and drops to zero with intensity θ . The service flow is

$$s_{it} = q_{it} d_{it}. \quad (8)$$

In equilibrium, households with bad matches will adjust their durable stocks immediately. These match-quality shocks are therefore a source of inframarginal adjustments.

Households have incomplete information about the aggregate state of the economy as in Mankiw and Reis (2002). Each household updates its information with i.i.d. probability Ξ . As in Carroll et al. (2018) and Auclert, Rognlie, and Straub (2020), we assume that households always know their idiosyncratic states and current income. They also learn of current real interest rate when they hit the borrowing constraint and they learn the current price of durables when they make

an adjustment. These assumptions ensure that households never violate the borrowing constraint. These information frictions allow the model to generate the hump-shaped response of durable and non-durable expenditure to monetary shocks. However, they only play a minor role in our main results as we show in Section 6.5.

In summary, the household maximizes (1) subject to (2), (3), (4), (5), (6), (7), and (8), taking as given exogenous paths for Y_t , r_t , p_t , and r_t^b . To compute the steady state of the model we use the continuous-time methods described in Achdou et al. (2017).

2.2 Distribution and Aggregate Quantities

We use the policy functions from the household's problem and the distribution of idiosyncratic state variables to construct aggregate quantities for a population of households. The individual state variables are the distance from the borrowing limit, $k = a + \lambda d$, the durable stock d , and idiosyncratic productivity z . The distribution over these variables is denoted $\Phi_t(k, d, z)$. In steady state, the exogenous variables are constant, $r_t = \bar{r}$, $r_t^b = \bar{r}^b$, $Y_t = \bar{Y}$, $p_t = \bar{p}$, and the steady state distribution over individual states is stationary and denoted by $\bar{\Phi}(k, d, z)$.

Aggregate nondurable expenditure is the sum of nondurable consumption and operating costs,

$$N_t = \int c_t(k, d, z) d\Phi_t(k, d, z) + \nu \int d d\Phi_t(k, d, z),$$

where $c_t(k, d, z)$ is the policy function for nondurable consumption, which varies with time as the policy rule depends on the current and future values of the aggregate variables Y_t , r_t , r_t^b , and p_t .

Aggregate durable expenditure is the sum of net durable expenditures from adjustments, including the fixed costs of adjustment, and maintenance costs

$$X_t = \int \lim_{dt \rightarrow 0} \frac{prob_{t,t+dt}(k, d, z)}{dt} (d_t^*(k, d, z) - (1 - f)d) d\Phi_t(k, d, z) + \chi \int d d\Phi_t(k, d, z)$$

where $prob_{t,t+dt}(k, d, z)$ is the probability that a household with individual state variables (k, d, z) will make an adjustment between t and $t+dt$, and $d_t^*(k, d, z)$ is the optimal durable stock conditional on adjusting. Since we integrate over changes in durable stocks at the household level, X_t reflects purchases of durables net of sales of durables. Our definition of X_t is therefore consistent with the construction of durable expenditure in the national accounts, in which transactions of used durables across households are netted out.

2.3 Calibration of the Household Problem

We set

$$u(c, s) = \frac{\left[(1 - \psi)^{\frac{1}{\xi}} c^{\frac{\xi-1}{\xi}} + \psi^{\frac{1}{\xi}} s^{\frac{\xi-1}{\xi}} \right]^{\frac{\xi(1-\sigma)}{\xi-1}} - 1}{1 - \sigma}.$$

ξ is the elasticity of substitution between durables and non-durables. There is a range of estimates for this elasticity. Many housing and durable models choose an elasticity of 1, which is consistent with estimates in Ogaki and Reinhart (1998) and the near-constant expenditure share of housing in the NIPA (Davis and Ortalo-Magné, 2011). However, housing and rent expenditure shares have trended up in the AHS, Census, and CEX survey data (Albouy et al., 2016), which is more in line with an elasticity below one. Further, studies allowing for non-homothetic preferences tend to estimate below-unitary elasticities (Pakoš, 2011; Davidoff and Yoshida, 2013; Albouy et al., 2016). We choose an elasticity of $\xi = 0.5$, which is at the lower end of these estimates. Choosing a lower value is conservative for intertemporal shifting in that the benefits of accelerating a durable adjustment are smaller.

We set $\sigma = 4$ implying an EIS of 1/4. This is at the lower end of the range typical in the literature. We need a low EIS to match the small response of non-durable consumption to monetary policy shocks we measure in the data (see Section 3). The durable exponent ψ is set to match the average ratio of the nominal values of the total durable stock (durable goods and private residential structures) and annual non-durable consumption (non-durable goods and services excluding housing) from 1970 to 2019.

Our calibration captures a broad notion of durables, which includes residential housing, autos, and appliances among other goods, as in Berger and Vavra (2015). While these goods differ in important respects, such as their depreciation rate and the probability of adjustment, they are all long-lasting and illiquid and purchases are lumpy and infrequent, features we stress in our analysis. Following this broad notion, our depreciation rate δ is the annual durable depreciation divided by the total durable stock in the BEA Fixed Asset tables, again averaged from 1970 to 2019. While 73% of the value of the total durable stock consists of residential housing, this component accounts for 23% of the total depreciation owing to the low depreciation rate of structures relative to cars and appliances. This explains why non-housing durables account for the majority (64%) of spending on durables and are thus important in the determination of aggregate demand.

The fixed cost of adjustment, in combination with the depreciation rate, is a key determinant of the probability of adjustment of the durable stock. We set the fixed cost to target a weighted average of the annual adjustment probabilities of individual durable goods. The three components of the average are the probability that a household moves to a new dwelling (15% per year as reported by Bachmann and Cooper, 2014), makes a significant addition or repair to their current dwelling (2.5% in the PSID), or acquires a new or used car (29.6% in the CEX). We attach a weight of 0.9 to the sum of housing moves and additions and repairs, and a weight of 0.1 to cars, based on the relative value of the housing stock and the car stock in the BEA fixed asset tables. This yields an annual adjustment probability of 0.19. Note that we cannot simply sum the probabilities of adjustments across durable goods, since this would overstate the liquidity of the households' total durable position in our model.

Since our objective is to explain the behavior of real interest rates during and after the Great Recession, we believe it is most appropriate to use the years 1991-2007 as a benchmark for interest rates, when the economy entered a period of low and stable inflation. The average, ex post, real federal funds rate in terms of non-durables over this period is equal to $\bar{r} = 1.5\%$. The steady state borrowing spread r^b is set to 1.7% based on the difference between the 30-year mortgage rate and the 10-year treasury rate over the same period.

Our calibration of the income process uses the estimates from Floden and Lindé (2001) converted to a continuous-time Ornstein-Uhlenbeck process. We set $\rho_z = \log(0.9136)$ and $\sigma_z = 0.2158$.

The discount rate ρ is set to match the average liquid financial asset holdings net of mortgage and auto loans to annual GDP ratio since 1970 of 0.87.⁵ The borrowing limit is set to $\lambda = 0.8$ in line with a 20% down payment requirement.

To calibrate the level of maintenance costs we use NIPA data on housing and car maintenance expenditures. In the housing output table, intermediate goods and services consumed amount to 52% of depreciation. We add to this the PCE on household maintenance. Turning to cars, PCE on motor vehicle maintenance and repair amounts to 47% of depreciation. We sum these three categories and divide by (total) durable depreciation to arrive at $\chi = 0.35$.⁶

⁵We use the same set of assets as McKay, Nakamura, and Steinsson (2016) and Guerrieri and Lorenzoni (2017): Currency and checkable deposits, time and savings deposits, MMF shares, Treasury securities, agency securities, municipal securities, corporate and foreign bonds, and mutual fund shares.

⁶ χ is not a convex combination of 52% and 47% because the numerator includes household maintenance and denominator includes depreciation from durables other than cars and housing, such as appliances and electronics, for which we assume there are no maintenance expenditures. Note that housing repair and maintenance are part of

Turning to operating costs, taxes on the housing sector, PCE on household utilities, and PCE on fuel oil and other fuels (excluding motor vehicle fuels) amount to 4.1% of the value of the housing stock. For cars, we find that PCE on motor vehicle fuels, lubricants, and fluids amounts to 22% of the value of the stock of vehicles. We sum the operating costs for cars and housing and divide by the total durable stock to obtain $\nu = 0.048$.

We set the probability that agents update their information set to $\eta = 0.167$, so that the expected time between updates is six quarters. This is in the middle of the range reported by Coibion and Gorodnichenko (2012, table 3), who find values between 0.11 and 0.24 among professional forecasters, households and firms.

2.4 Estimating the Arrival Rate of Match-Quality Shocks

The intensity of the match-quality shock, θ , does not have a natural data counterpart that lends itself to calibration. We therefore estimate this parameter using PSID data and a variant of the structural estimation method developed by Berger and Vavra (2015) that allows for match-quality shocks. We only provide a brief overview here, with details relegated to Appendix B.

The estimation method centers on matching the probability of a durable adjustment as a function of the “durable gap” $\omega_{it} \equiv d_{it}^* - d_{it}$, where d_{it}^* is the optimal durable stock based on the current state variables. Specifically, d^* is the level of durables the household would adjust to if it were forced to adjust immediately. Intuitively, in a fixed-cost model the probability of adjustment should be greater the larger is the absolute gap, since the benefit of adjusting the durable stock is larger.

Gaps are easily computed in the model, since both the optimal durable choice d_{it}^{*model} and the current durable stock d_{it}^{model} are known. In the data, we only observe current durable holdings d_{it}^{data} directly. We infer data gaps using a set of observables \mathbf{Z}_{it}^{data} , and the model-implied relationship between them and the optimal durable stock, $d_{it}^{*data} = F^{model}(\mathbf{Z}_{it}^{data})$, where F^{model} is the model’s mapping from the observables to d^* . We then minimize the distance between the hazard rate of adjustment in model and data, $h^{model}(\omega) - h^{data}(\omega)$ where $h(\omega)$ is the probability of adjusting given a gap ω , as well as the distance between the distribution of gaps, $f^{model}(\omega) - f^{data}(\omega)$ where $f(\omega)$ is the density at gap ω . Like Berger and Vavra (2015), we allow for measurement error in intermediate goods and services consumed.

Table 1: Calibration of the Full Model

Parameter	Name	Value	Source
<i>Parameters of the Household's Problem</i>			
ρ	Discount factor	0.096	Net Assets/GDP = 0.87
σ	Inverse EIS	4	See Section 3
ψ	Durable exponent	0.581	d/c ratio = 2.64
ξ	Elas of substitution	0.5	See text
\bar{r}	Real interest rate	0.015	Annual real FFR
\bar{r}^b	Borrowing spread	0.017	Mortgage T-Bill spread
δ	Depreciation rate	0.068	BEA Fixed Asset
f	Fixed cost	0.194	Ann. adjustment prob = 0.19
θ	Intensity of match-quality shocks	0.158	See Section 2.4
χ	Required maintenance share	0.35	See text
ν	Operating cost	0.048	See text
ρ_z	Income persistence	-0.090	Floden and Lindé (2001)
σ_z	Income st. dev.	0.216	Floden and Lindé (2001)
λ	Borrowing limit	0.8	20% Down payment
Ξ	Prob. of information updating	0.167	Coibion and Gorodnichenko (2012)
<i>General Equilibrium Parameters</i>			
G/Y	Steady state govt share	0.2	Convention
ζ	Inverse durable supply elasticity	0.049	See Section 4
κ	Phillips curve slope	0.48	See Section 4
ρ_r	Real rate persistence	-0.47	Estimated over 1991-2007
ϕ_π	Real rate response to inflation	0.43	Estimated over 1991-2007
ϕ_y	Real rate response to output gap	0.66	Estimated over 1991-2007
ρ_g	Non-household demand persistence	-0.90	See text
ρ_{r^b}	Borrowing spread persistence	-0.63	Estimated over 1991-2007

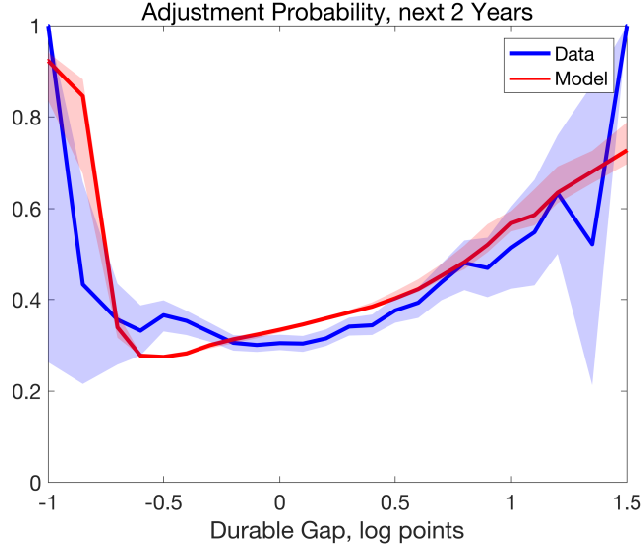


Figure 1: Hazard of adjustment conditional on the durable gap $\omega = d^* - d$, where d^* is the optimal durable choice conditional on adjusting and d the initial durable stock. Shaded areas are 95% confidence bands.

\mathbf{Z}_{it}^{data} .

The arrival rate of the match-quality shock, θ , is primarily identified by the hazard of adjustment for small gaps. Intuitively, at small gaps the household should be relatively far from an adjustment threshold due to the fixed cost, whereas adjustments at large gaps likely reflect that the household crossed an adjustment threshold. Most of the adjustments at small gaps are therefore attributed to the match-quality shock. As we vary θ , we re-calibrate the discount rate ρ , the fixed cost f , and the preference for durables ψ to match the same targets for net assets, the probability of adjustment, and the durable-stock-to-non-durable-consumption ratio.

Our set of observables from the PSID, \mathbf{Z}_{it}^{data} , are net liquid assets a_{it} , the value of the durable stock d_{it} , and annualized consumption expenditures over the following wave $\bar{c}_{i,t,t+2}$. We use direct information from the PSID on durable adjustments. For housing we use the answer to whether the household has moved along with the recorded date of moving, and whether the household made a major addition or repair to the home. Similarly, if any car acquisition occurred after the previous wave, then we record this as a car adjustment. Just like our total adjustment frequency, we weight a housing adjustment by 0.9 and car adjustments by 0.1.

Figure 1 plots the model- and data-implied hazards of adjustment conditional on the durable gap at the optimal parameter estimate, $\theta = 0.1575$. The bootstrapped 95% confidence band is

[0.157,0.159], based on sampling households from the PSID with replacement. The model accounts well for the upward-sloping hazard and explains 79% of the variation in the hazard rate. Note that the probability of adjustment is substantial in both model and data for even small gaps, suggesting that the match-quality shock is important in matching the data.⁷ Indeed, our estimate implies that 75% of all adjustments are due to the match-quality process. This split between match-quality shocks and adjustments due to depreciation or income fluctuations roughly coincides with the stated reasons for moving house in the March CPS Supplement (see Appendix Table 3).

3 The Response of Durable Demand to Monetary Policy

In many fixed-cost models, the timing of durable adjustments is very sensitive to intertemporal incentives (see House, 2014), which poses a challenge in modeling the monetary transmission mechanism because durable demand is excessively responsive to monetary policy (see Reiter et al., 2013). In this section we show that our model predicts a response of durable demand to interest rates that is in line with several empirical benchmarks. The model is consistent with the magnitude of the spending response as well as the subsequent reversal of this spending that is the hallmark of intertemporal shifting effects. We identify changes in real interest rates in two ways. We begin with identified monetary policy shocks before turning to quasi-experimental evidence from changes in state sales tax changes.

3.1 Evidence from Identified Monetary Shocks

We use monetary shocks from Romer and Romer (2004), extended by Wieland and Yang (2017), for 1969Q1-2007Q4, as a source of exogenous variation in real interest rates. We estimate the impulse responses of several outcome variables to these shocks. The impulse response functions we estimate serve two purposes in our analysis. First, they are a source of information about the plausibility of intertemporal shifting effects in the monetary transmission mechanism. Second, they serve as empirical benchmarks against which we can compare the model.

Our outcome variables are the log of total durable expenditure, $x_t = \ln(X_t)$, and the extensive

⁷We estimate larger adjustment probabilities at small gaps than Berger and Vavra (2015) do because we follow a different approach to identifying adjustments in the data. Berger and Vavra exclude durable adjustments smaller than 20% of the value of the durable stock in part to filter out idiosyncratic moves across location.

margin of durable purchases.⁸ To measure the extensive margin, we construct time series for the fraction of the population moving residence each year and the fraction of the population buying a car each quarter using micro-data from the PSID and the CEX, respectively.⁹ We express these outcomes in terms of the cumulative response to better show the extent of intertemporal shifting. We estimate the impulse responses using local projections that condition on 16 lags and a time trend. The standard errors account for serial correlation.

We also estimate the response of the real interest rate in terms of non-durable goods, r_t , aggregate income, Y_t , and the relative durable price, p_t . These impulse responses show how the prices households face react in response to the monetary policy shock, which we will use below when comparing the model to the data. To estimate these impulse responses we make use the equivalence result from (Plagborg-Møller and Wolf, 2019) who show that VARs and local projections yield the same impulse response up to the horizon of included lags (16 quarters in our case). The benefit of using a VAR is that it generates smoother impulse responses beyond 16 quarters, which is useful when feeding these paths into the household decision problem below.¹⁰ The VAR includes 16 lags of the monetary shock, the real interest rate, log GDP, and the relative price of durables, in which the monetary shock is ordered first. The VAR also includes a time-trend and the standard errors are block-bootstrapped.

The top-left panel of Figure 2 displays an economically sizable and statistically significant decline in the annualized, ex post, real federal funds rate in terms of non-durables given a one-standard deviation expansionary monetary policy shock.¹¹ In the center-left panel, the relative durable price displays a small increase but remains close to the steady state.¹²

The bottom-left panel shows that GDP increases with a lag. After 16 quarters the response of GDP turns from positive to negative, indicating that a reversal in demand is taking place.

⁸Appendix C details the construction of the variables used in this analysis.

⁹We construct the annual time series for the probability of moving to a different residence using PSID data from 1969-1997 following Bachmann and Cooper (2014). Bachmann and Cooper (2014) show that the moving probability from the PSID is in line with the shorter time series from the CPS March Supplement and the AHS. For the probability of buying a car we use CEX data from 1980-2007. Appendix Figure A.1 shows that CEX car expenditures aggregate well to NIPA expenditures. See Appendices C.2 and C.3 for details.

¹⁰We do not use the VAR approach for the cumulative durable spending outcomes, because constructing cumulative impulse response functions by aggregating estimates of the level impulse response function from the VAR is inefficient, which is reflected in large standard error bands.

¹¹We normalize the Romer shock to yield a 25 basis point decline in the real interest rate on impact.

¹²Appendix Figure A.2 shows that while the relative durable price is close to flat, both non-durable and durable prices increase with some lag following an expansionary monetary policy shock.

The graphs for *cumulative* real durable expenditure, the *cumulative* moving probability, and the *cumulative* car acquisition probability provide further evidence for intertemporal shifting. The cumulative durable expenditure (top-right panel) peaks at 14% after 15 quarters. In other words, durable expenditures are above normal for the first 15 quarters. However, durable expenditures subsequently come in below normal, as evidenced by the decline in cumulative durable expenditure. By the end of the impulse response horizon, cumulative durable expenditures have fallen to around one third of their original peak. Thus, the point estimate implies that two thirds of the increase in durable expenditures was subsequently reversed. Of course, as in the case of GDP, there is considerable uncertainty about the size of this reversal.

The responses of the *cumulative* moving probability and the *cumulative* car acquisition provide suggestive evidence that this reversal is driven by the extensive margin of adjustment. In particular, for both adjustment probabilities we see a peak in the cumulative impulse response function at around 12-16 quarters, followed by a near-complete reversal. This suggests that households accelerate durable purchases following an expansionary monetary shock, and that these adjustments are subsequently missing.

3.2 Evaluating the Model

We now evaluate the model’s ability to match the durable spending response to identified monetary policy shocks. Our primary focus is on the magnitude of this response or, in other words, the interest elasticity of durable demand. Limiting this sensitivity poses a challenge and motivates some of the ingredients we include in our model.

To evaluate the model, we feed the mean impulse response function for the real interest rate, r_t , aggregate income, Y_t , and the relative durable price, p_t , into the household decision problem of Section 2 starting in steady state. We assume that these paths come as a surprise at $t = 0$, but become subsequently known to agents whenever they update their information set. All paths return to steady state after 32 quarters and stay there subsequently.¹³

The left panel of Figure 3 plots the model-implied cumulative durable expenditures against the data. In our full model, the peak real durable response is 14.5% compared to 11.5% in the data.

¹³Formally, we compute the Jacobians of the endogenous variables with respect to r_t , Y_t , and p_t , and multiply them with the perfect foresight paths of these variables. Because this effectively linearizes the model around the aggregate state variables, the uncertainty around the estimated impulse response functions does not affect the impulse response function from our model.

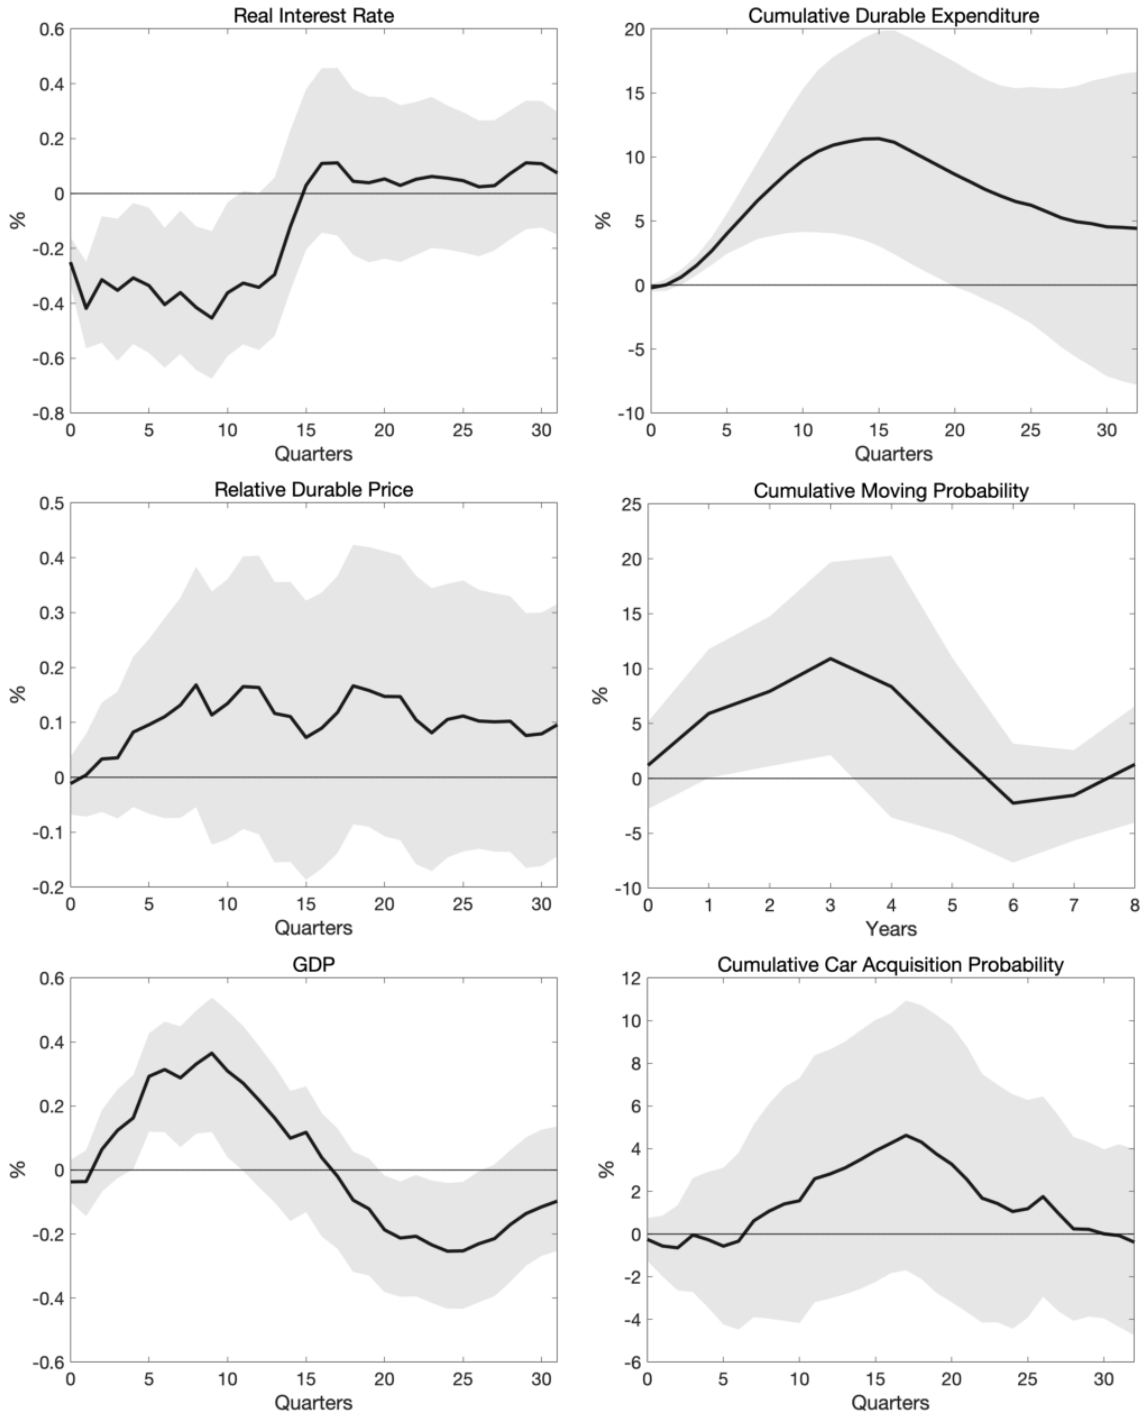


Figure 2: Impulse response function of the real interest rate in terms of non-durables (top-left panel), the relative durable price (center-left), real GDP (bottom-left), cumulative real durable expenditure (top-right), cumulative probability of moving house (center-right), and cumulative probability of buying a car (bottom-right), to a one-standard-deviation Romer and Romer monetary policy shock.

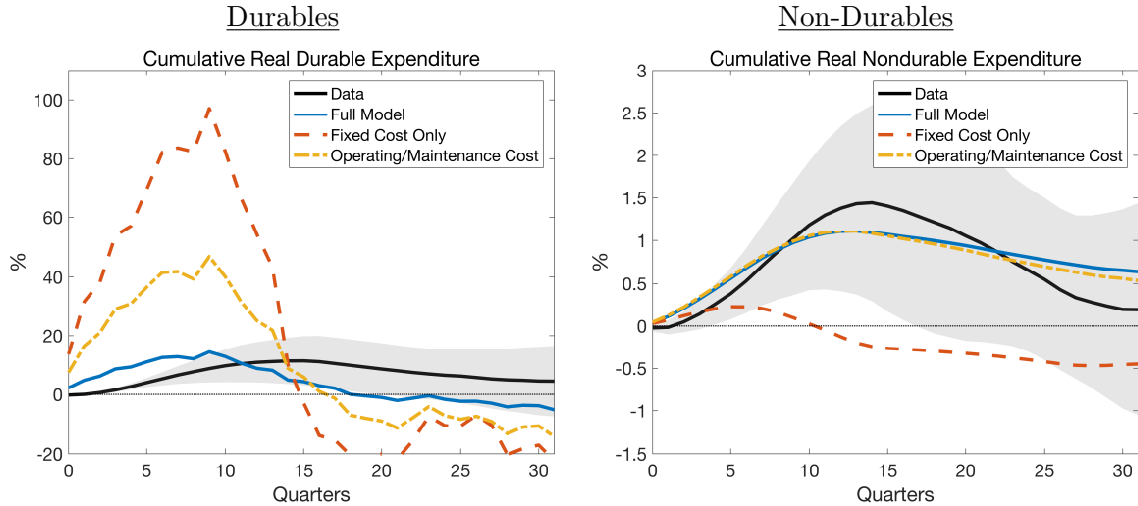


Figure 3: Model simulations feeding in observed paths for (Y_t, r_t, p_t) . The panels show the cumulative response of durable expenditures (left) and non-durable spending (right). Each panel shows the model with only fixed costs, the model with fixed costs and operating costs, and the full model.

The full model response thus falls well within the 95% confidence interval of the peak response, [3,20%]. In addition, the model produces a reversal in durable demand within the confidence band.

To highlight that operating/maintenance costs and match-quality shocks are necessary for this success, Figure 3 makes two additional model comparisons.¹⁴ First, we compute the cumulative durable demand response in a model that abstracts from match-quality shocks but includes the operating and maintenance costs from our full model (“Operatings/Maintenance costs”). This model yields a counterfactually large peak cumulative durable demand response of 47%, four times larger than in the data. Second, we compute the cumulative durable demand response model that abstracts from both match-quality shocks and operating/maintenance costs (“Fixed Cost Only”). This model predicts that cumulative durable demand peaks at 97%, 8 times more than their peak response in the data.¹⁵

While the magnitude of the durable expansion in the full model is consistent with the data, it does occur earlier than in the data. The main determinant of how gradual cumulative durable expenditure increases is the information rigidity. Without information frictions, durable expen-

¹⁴For both of these alternative models, we re-calibrate the discount rate ρ , the fixed cost f , and the durable preference ψ to match our empirical targets for the net liquid asset/GDP ratio, the frequency of adjustment, and the durable-stock-to-nondurable-consumption ratio.

¹⁵Models with a high interest elasticity of durable demand, like the fixed-cost-only model, will show very little fluctuation in the durables real interest rate in equilibrium (Barsky et al., 2007). Appendix Figure A.3 plots the impulse response for the durables real interest rate, which displays an economically sizable, persistent, and statistically significant decline following an expansionary monetary policy shock.

diture in the first quarter is already as large as the peak cumulative response in our benchmark model. While a larger degree of information rigidity would allow us to get closer to the data on this dimension, it does pull us outside the range of values estimated in the literature. Ultimately, we show in Section 6.5 that the degree of information rigidity does not have major quantitative implications for the importance of intertemporal shifting, which suggests that matching the exact timing of durable expenditure in the data is not critical for our purposes.

In the right panel of Figure 3 we show that the full model also provides a good match with the dynamics of cumulative non-durable expenditures. The model with operating/maintenance costs performs about as well, whereas the fixed-cost-only model predicts too much substitution from non-durable to durable spending. The relatively small non-durable spending response in the data motivated our choice of a relatively low elasticity of intertemporal substitution in Section 2. Note that cumulative non-durable consumption increases by 1.43%, eight times less than the durable response. As durable expenditure makes up 22.9% of total consumption excluding rent, the durable component accounts for 70% of the total increase of consumption expenditure following a monetary shock.

Why do operating costs and match-quality shocks reduce the sensitivity of durable demand to interest rate changes? These ingredients play a key role in limiting the sensitivity of the extensive margin of durable demand to intertemporal prices. Match-quality shocks are a source of infra-marginal adjustments. We target a certain probability of adjustment in total and by associating more of these adjustments with the match-quality shock, fewer are attributed to households that have hit an (S, s) band. Therefore including match-quality shocks means there are fewer households near the adjustment thresholds that can be induced to accelerate their adjustments by monetary policy.¹⁶ To demonstrate this role of match-quality shocks, Figure 4 shows the density of durable holdings in the fixed-cost model and the full model for a particular levels of net liquid assets and idiosyncratic income. The vertical dashed lines show the adjustment thresholds. The mass of households near the adjustment threshold is substantially smaller in the full model.

Similarly, operating costs stabilize the extensive margin of demand because they are a component of the user cost of durables that is not sensitive to interest rates. Including operating

¹⁶The logic of how match-quality shocks affect the extensive margin response of durable demand has antecedents in the literature on price setting (see Golosov and Lucas, 2007; Midrigan, 2011; Nakamura and Steinsson, 2010; Alvarez et al., 2016).

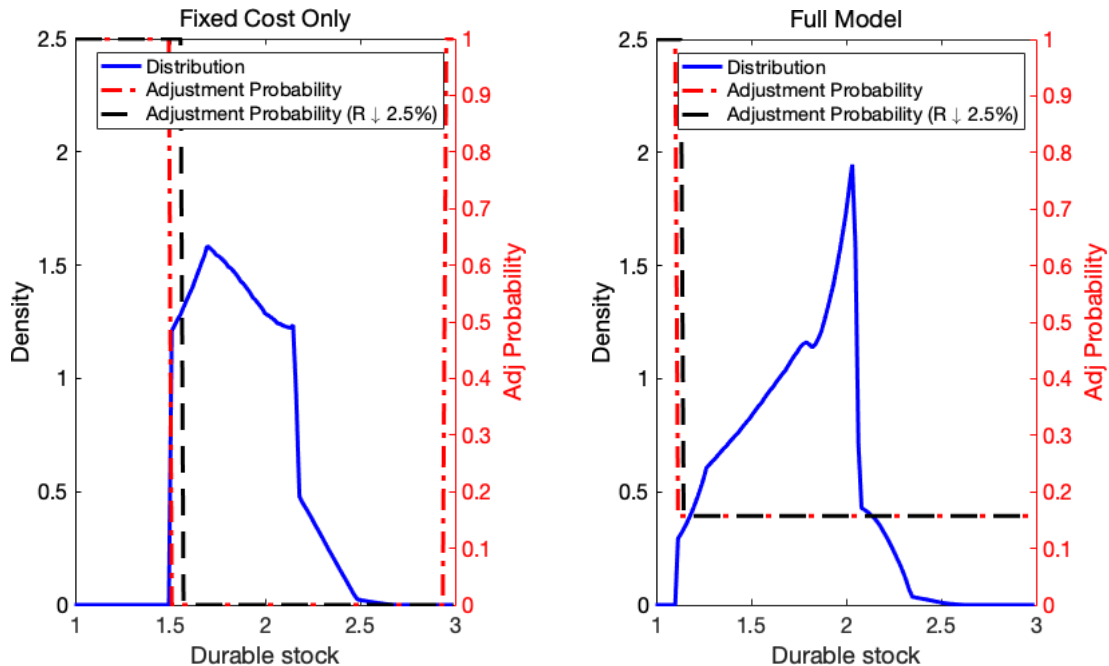


Figure 4: Distribution of households over durable holdings. The distributions shown are conditional on net liquid assets of 1.99 and gross income of 0.93. Net liquid assets are measured as $a - \lambda d$. The adjustment thresholds are shown at their steady state positions and in response to a 2.5 p.p. cut in the interest rate for the current quarter.

costs therefore stabilizes the user cost and therefore durable demand. In Figure 4, the adjustment thresholds shift to the right in response to monetary stimulus, but much less so in the full model, which includes operating costs.

We now show that the willingness of households to shift the timing of their durable adjustments in our full model also aligns well with the observed extensive margin responses for cars and housing. We consider two different calibrations to tailor the model to cars and housing, respectively. The primary difference in the calibrations is the depreciation rate. Housing structures depreciate at a much slower rate, 2% per year, while cars depreciate at a much higher rate, 20% per year, than the value-weighted durable stock.¹⁷ As above, we simulate the impulse response for the extensive margin by feeding the empirical impulse responses of Y_t , r_t , and p_t into the model.¹⁸ Figure 5 has

¹⁷The probability of adjustment is also higher for cars (7.4% quarterly) than for housing (15% annually), and households own more housing wealth ($d/c = 1.92$) than car wealth ($d/c = 0.201$). We recalibrate the discount rate ρ , the fixed cost f , and the durable exponent ψ to match these targets, as well as a net-liquid-asset-to-GDP ratio of 0.92 for housing and 1.31 for cars. When we include match-quality shocks, 75% of all adjustments will come from the match-quality process, which is the same fraction as in our estimated model for all durables. This requires $\theta = 0.12$ for housing and $\theta = 0.22$ for cars. We only subtract the collateralized loans from liquid assets for the durable we calibrate to. We also allow for a higher borrowing spread $r^b = 0.03$ in the car model based on the average spread of four-year car loans with five-year treasury bonds.

¹⁸We use a different relative price response: for housing we use the relative price of housing to non-durable

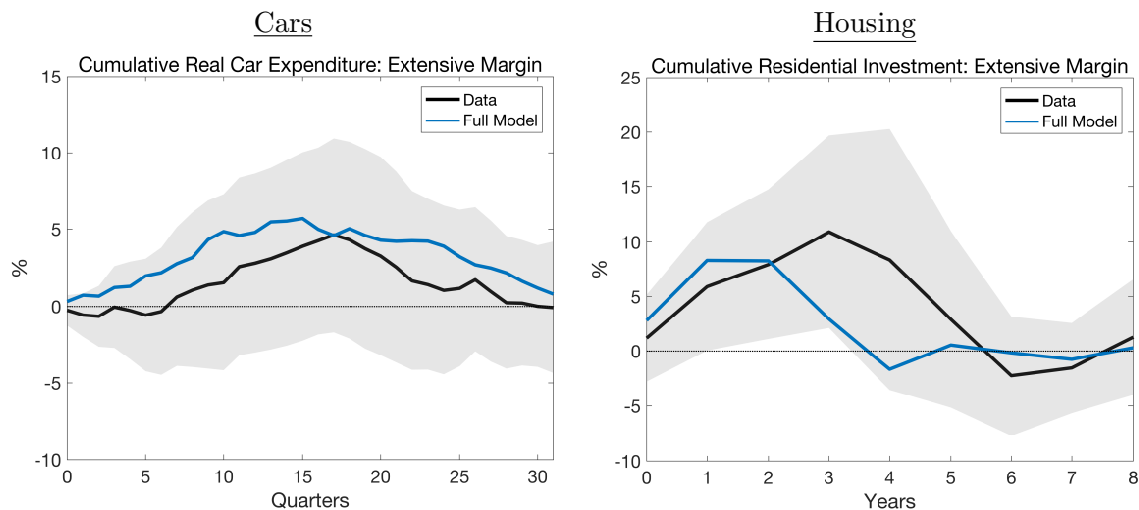


Figure 5: Model simulations feeding in observed paths for (Y_t, r_t, p_t) . The panels show the cumulative response of the extensive margin in the model for cars (left panel) and the model for housing (right panel) in the full model against the data.

the results. The left panel shows that the extensive margin for car adjustments in our full model accords well with the data. In both model and data, the extensive margin response is initially positive but then fully reverses. The peak cumulative response is only slightly higher than in the data.

In the right panel of Figure 5 we plot the cumulative extensive margin for the housing model. Our full model again accords well with the data. It matches the peak cumulative response and it predicts a reversal very similar to what is observed in the data.

From this analysis, the willingness of households to shift the timing of durable adjustments in the full model appears to be consistent with this evidence from identified monetary policy shocks.

3.3 Quasi-Experimental Evidence

We now evaluate the model’s ability to fit evidence from quasi-experimental variation in real interest rates. An advantage of this analysis is that it focuses on a narrower set of economic mechanisms than does the preceding analysis of monetary policy shocks because in this case only real interest rates are affected.

There is extensive quasi-experimental evidence that variation in intertemporal prices can shift consumption, whereas for cars we use the relative price of cars. Appendix Figure A.4 shows that these relative prices behave similarly to the relative price for all durables.

durable expenditure through time. Empirical studies of anticipated VAT or sales tax consistently estimate increases in household durable expenditures followed by complete or near-complete reversals (Cashin and Unayama, 2011; D’Acunto et al., 2016; Baker et al., 2019). Intuitively, one can interpret the anticipated price increase from a VAT or sales tax increase as a low real interest rate, which is why such policies are termed “unconventional fiscal policy” (Correia et al., 2013). Just as in our model, the low real interest rate in the data induces households to pull forward durable expenditure. Similarly, one can interpret the Cash-for-Clunkers program as a temporary low price of acquiring durables and hence a low real interest rate. Mian and Sufi (2012) show that this program similarly caused an acceleration of durable expenditures followed by an almost complete reversal. That low real interest rates pull forward durable expenditures lends credence to our emphasis on intertemporal shifting effects in the monetary policy transmission mechanism.

We now show that our model is quantitatively consistent with the estimates in Baker et al. (2019),¹⁹ who analyze the response of auto sales to anticipated changes in state sales tax rates. We focus on this evidence because anticipated sales tax changes are a broad-based change in incentives—as opposed to Cash for Clunkers, which was targeted—but at the same time exploits regional variation. Baker et al. estimate a cumulative 12.7 percent increase in monthly auto sales leading up to a 1 percentage point increase in sales tax. This implies an annualized 12 percent decrease in the real interest rate for cars in the month before the tax increase so the elasticity of the extensive margin of auto sales to interest rates is about $12.7/12 = 1.1$.

In our model calibrated to cars, we calculate the response of the extensive margin to a one month drop in the real interest rate with a magnitude of 1% annualized. We assume that the real rate drop is known 5 months ahead of implementation. Baker et al. (2019) note that sales tax increases are known 2-3 months in advance for referenda and 6-9 months in advance for legislated changes. Our choice is the mid-point. As shown in Figure 6, our model produces a peak elasticity of 1.2 and a subsequent reversal that tracks their point estimate quite closely and is well inside the confidence intervals.²⁰

¹⁹See their Table 3, Col. 1.

²⁰Baker et al. do not provide confidence bands for the cumulative effects, so we construct them assuming zero correlation between the individual estimates.

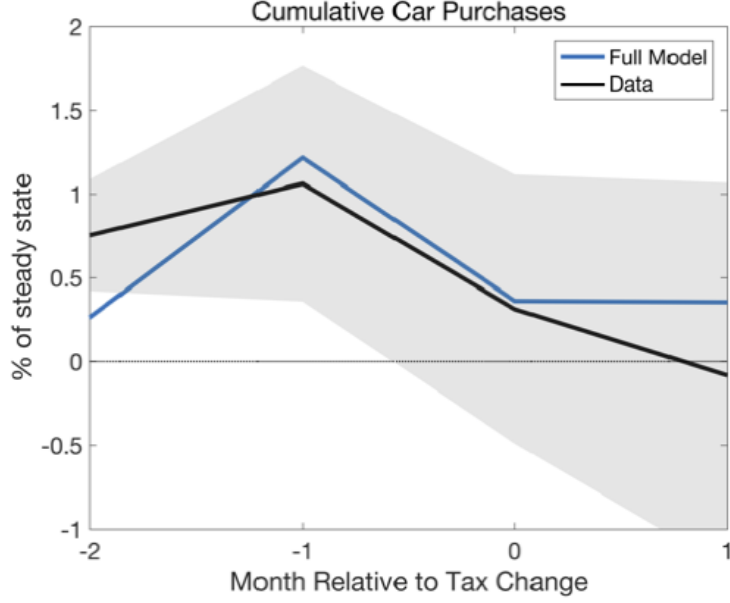


Figure 6: Cumulative car acquisitions for an anticipated sales tax increase in Baker et al. (2019) versus the full model. We convert the tax elasticity in Baker et al. (2019) to an interest rate elasticity by dividing by 12.

4 General Equilibrium Model

So far we have focused on the household problem, which serves as the demand side of our model. We now specify the supply side and market clearing conditions. With a full general equilibrium model we can then study monetary experiments that differ from the identified monetary policy shocks and determine the importance of intertemporal shifting for the dynamics of interest rates during and after the Great Recession.

4.1 Labor Supply, Production, and Aggregate Supply

We adapt the standard sticky-wage environment developed by Erceg et al. (2000) to allow for uninsured idiosyncratic labor productivity. Each household i supplies a continuum of differentiated labor of type $j \in [0, 1]$, with hours denoted n_{ijt} . We extend the household preferences with an additively separable disutility of labor supply

$$E_0 \int_{t=0}^{\infty} e^{-\rho t} \left[u(c_{it}, s_{it}) - \bar{u}_{c,t} \int_0^1 v(n_{ijt}) dj \right] dt. \quad (9)$$

where we use the modifier on the disutility of labor, $\bar{u}_{c,t} = \int_0^1 \frac{\partial u(c_{it}, s_{it})}{\partial c_{it}} di$, to eliminate wealth effects on labor supply.²¹ Labor supply is determined by a set of unions as described below so the household takes labor supply and labor income as given. As the disutility of labor is additively separable and labor income is outside the household's control, the decision problem we analyzed in the previous sections is unchanged.

Final goods are produced with a technology that is linear in labor, $Y_t = Z_t L_t$, where Z_t is the exogenous level of productivity and L_t is an aggregate of labor supply given by

$$L_t = \left(\int_0^1 l_{jt}^{\frac{\varphi-1}{\varphi}} dj \right)^{\frac{\varphi}{\varphi-1}},$$

where

$$l_{jt} = \int_0^1 z_{it} n_{ijt} di.$$

We now interpret z_{it} as idiosyncratic labor productivity. In this formulation, each household faces uninsurable risk to their productivity z_{it} , but face the same (relative) exposure to each variety of labor j . The final good is produced by a representative firm. Prices are flexible and equal to nominal marginal cost: $P_t = W_t/Z_t$, where W_t is the price index associated with the aggregator L_t . The real wage is then $W_t/P_t = Z_t$. We assume that the movements in Z_t are permanent, $d \ln Z_t = \epsilon_t^Z$, so that ϵ_t^Z has the interpretation of a permanent income shock.

The final good is used for several purposes including non-durable consumption, an input into durable production, and government consumption. In addition, we interpret the spread between the borrowing and saving interest rates as reflecting an intermediation cost. We assume the intermediation cost follows an exogenous autoregressive process given by $\dot{r}_t^b = -\rho_{rb} r_t^b + \epsilon_t^{rb}$. Appendix E shows the market-clearing conditions of the model.

We obtain an upward-sloping Phillips curve through sticky nominal wages. A continuum of unions set the wage, W_{jt} , of each type of labor. The union maximizes the equally-weighted utility of the households subject to a Rotemberg-style adjustment cost of $\frac{\Phi}{2} \bar{u}_{c,t} L_t (\mu_{jt})^2$, where Φ is a parameter that controls the strength of the nominal rigidity and μ_{jt} is the growth rate of W_{jt} such that $d \ln W_{jt} = \mu_{jt} dt$. Among union workers supplying type j , all labor is equally rationed,

²¹We do so for computational convenience. The wealth effect in labor supply introduces an additional loop in finding the market-clearing prices and quantities since it creates a feedback from distribution of marginal utility to the real interest rate in the policy rule. In unreported results we have found that this wealth effect is not quantitatively important for output dynamics.

$n_{ijt} = l_{jt}$. In a symmetric equilibrium, all workers supply L_t units of labor and each household receives real, pre-tax income of $z_{it}Y_t$.

Appendix E presents the union's problem and shows that the log-linearized symmetric equilibrium gives rise to the following Phillips curve

$$\dot{\pi}_t = \rho\pi_t - \kappa \left(\frac{Y_t - \bar{Y}_t}{\bar{Y}_t} \right), \quad (10)$$

where $\pi_t = \frac{d \ln P_t}{dt}$ and \bar{Y}_t is potential output.

The durable good is produced by a unit measure of perfectly competitive firms using the production function,

$$X_t = v Z_t^\zeta M_t^{1-\zeta} \bar{H}^\zeta,$$

where X_t is the production of durables, M_t is the input of the non-durable good, and v is a constant. The constant flow \bar{H} of land is made available and sold by the government at a competitive price. Z_t enters the production function here in a manner that is "land-augmenting" so that the long-run relative price of durables is unaffected by permanent TFP shocks. The first order conditions of this problem lead to a relative price of

$$p_t = (1 - \zeta)^{-1} v^{-\frac{1}{1-\zeta}} \left(\frac{X_t}{Z_t \bar{H}} \right)^{\frac{\zeta}{1-\zeta}}, \quad (11)$$

where X_t is aggregate durable goods production. Thus, $(1 - \zeta)/\zeta$ is the supply elasticity of the durable good.

4.2 Government

Monetary policy is governed by a standard interest rate rule,

$$\dot{r}_t = \rho_r(r_t - \bar{r}) + \phi_\pi \pi_t + \phi_y \frac{Y_t - \bar{Y}_t}{\bar{Y}_t} + \epsilon_t^r, \quad (12)$$

where the first term captures interest rate smoothing, the second and third terms the endogenous monetary policy response, and ϵ_t^r is an exogenous shock.

Fiscal policy consists of a constant debt policy,

$$A_t = \int_0^1 a_{it} dt = \bar{A}. \quad (13)$$

We assume that the government levies taxes proportional to z_{it} where the tax rate $\bar{\tau}_t$ is set to satisfy the government budget constraint so we have

$$y_{it} = (Y_t - \bar{\tau}_t) z_{it} \tag{14}$$

and the period-by-period government budget constraint is

$$\bar{\tau}_t = r_t \bar{A} + G_t$$

where G_t is an exogenous level of government consumption.²² We assume that it follows an autoregressive process in logs, $d \ln G_t = \rho_g (\ln G_t - \ln \bar{G}) + \epsilon_t^G$. In our analysis, government consumption will stand in for changes in demand that originate outside the household sector and we will at times refer to “non-household demand.”

4.3 Calibration of the General Equilibrium Model

We set v so as to normalize the relative price of durables to one in steady state. We calibrate the inverse supply elasticity of durable goods to $\frac{\zeta}{1-\zeta} = 0.049$. Our choice of ζ reflects land’s share in the production of durables, which we calculate as follows. Residential investment is on average 36% of broad durable consumption expenditures (NIPA Table 1.1.5, 1969-2007). New permanent site structures account for 58% of residential investment (NIPA Table 5.4.5). Davis and Heathcote (2007) report that 11% of sales of new houses reflect the value of land. Therefore payments for new land amount to a little over 2% of the expenditure on durables. However, Davis and Heathcote (2007) also report that the existing stock of housing is paired with more valuable land and land accounts for 36% of the value of the housing stock, which is substantially larger than the 11% share in new housing. In our model, durables trade at a single price so there is no distinction between the cost of creating new durables and the value of the stock. We therefore take the mid-point of 11% and 36%, which implies that payments to land account for 5% of expenditure on durables.

An elastic supply of durable goods is consistent with the small relative price response in Figure 2. An elastic supply of durable goods also finds some support from Goolsbee (1998) and House and Shapiro (2008) who present evidence on the response of capital goods prices to policies that stimulate investment demand. House and Shapiro find little evidence of a price response and argue

²²The government also raises a small amount of revenue from selling land. In steady state this amounts to 0.5% of GDP. For computational convenience we assume this revenue finances an independent stream of spending.

for a high supply elasticity. Goolsbee argues for less elastic supply in general, but for the categories of goods that also serve as consumer durables (autos, computers, and furniture) he finds little price response.

The slope of the Phillips curve is 0.48. Note that the slope of the Phillips curve is expressed in terms of the change in annualized inflation for a unit of the output gap per year so we would need to divide by 16 to compare to a quarterly discrete-time model, which yields a slope of $0.48/16 = 0.03$. That value is squarely in the middle of empirical estimates (Mavroeidis et al., 2014).

We estimate the monetary policy rule from 1991-2007, since there is no significant trend in the real rate over this period. This yields $\rho_r = -0.60$ (equivalent to a quarterly persistence of 0.86), $\phi_\pi = 0.79$ and $\phi_y = 0.75$.²³ We also estimate the process for the borrowing spread over 1991-2007, which yields $\rho_{r,b} = -0.63$ and is equivalent to a quarterly persistence of 0.85. We set the persistence of the non-household demand shock $\rho_g = -0.42$ equivalent to a quarterly persistence of 0.9. We deliberately choose a value at the lower end of the persistence spectrum typically estimated for demand shocks, since a more persistent shock naturally has more persistent effects on r^* . This is a conservative choice, since we emphasize the prolonged low levels of r^* after the Great Recession.

4.4 Solving the Model

Following Auclert et al. (2019), we translate the partial equilibrium Jacobians of the model into general equilibrium Jacobians by determining the endogenous prices that satisfy the market clearing conditions (11) and (13). See appendix A for details. For a model that is linear in aggregate states and aggregate shocks, our solution is equivalent to the impulse response functions obtained from a stochastic linear rational expectations solver (see Boppart, Krusell, and Mitman, 2018).

Our analysis assumes that the economy's dynamics are linear the aggregate states but can be non-linear in idiosyncratic states in the style of perturbation approaches to solving heterogeneous-agent models (e.g. Reiter, 2009). During the Great Recession, the economy hit the ELB, which creates a kink in the response of the interest rates with respect to the state of the economy. As we discuss later, we will incorporate the ELB with a sequence of monetary news shocks, which captures the effect of the ELB on the expected path of rates. We have also investigated the sensitivity of our results to another form of nonlinearity. Existing work by Berger and Vavra (2015) argues that

²³The long-run responses are $\frac{\phi_\pi}{-\rho_r} = 1.26$ and $\frac{\phi_y}{-\rho_r} = 1.31$. Note that our estimated rule satisfies the conventional Taylor principle since it is specified in terms of a real rate.

monetary policy is less effective in recessions in models with lumpy durables. In Appendix I, we show our results are robust to reducing the effectiveness of monetary policy.

5 The Monetary Transmission Mechanism

The acceleration of durable purchases in response to monetary policy stimulus has two implications for the monetary transmission mechanism that we highlight. First, monetary policy shifts aggregate demand intertemporally, which means stimulus now reduces demand in the future. Second, forward guidance is less powerful than implied by standard New Keynesian models.

These aspects of the monetary transmission mechanism are captured by the sequence-space Jacobian of the output gap with respect to real interest rates, which we call it the “monetary transmission matrix” and denote it \mathcal{M} . The (i, j) element of \mathcal{M} gives the general equilibrium response of the output gap at a horizon of $i - 1$ quarters with respect to news about real interest rates at a horizon $j - 1$ quarters,²⁴

$$\mathcal{M} = \begin{pmatrix} \frac{d\hat{Y}_0}{dr_0} & \frac{d\hat{Y}_0}{dr_1} & \cdots \\ \frac{d\hat{Y}_1}{dr_0} & \frac{d\hat{Y}_1}{dr_1} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix}$$

where $\hat{Y}_t = \frac{Y_t - \bar{Y}_t}{\bar{Y}_t}$ is the output gap. The below-diagonal elements of \mathcal{M} show how the output gap responds to monetary stimulus in the past and therefore are informative about intertemporal shifting. The above-diagonal elements of \mathcal{M} show how the output gap responds to monetary stimulus in the future and therefore inform us about the effects of forward guidance. These features of \mathcal{M} play a key role in our analysis of r^* , which follows in Section 6.

5.1 An Illustration of Intertemporal Shifting

To show the intertemporal shifting effect of monetary stimulus we consider a simple experiment in which the central bank reduces the real interest rate by 1% (annualized) for the current quarter. Following that quarter the real interest rate returns to steady state and remains there forever. The initial cut comes as a surprise, but subsequently the path of the real interest rate is known with

²⁴See Auclert et al. (2019) for more discussion of sequence space Jacobians. In our continuous-time setting, we define \mathcal{M} as the response to a change in real interest rates that occurs at the given horizon and lasts for one quarter. This definition reflects our choice to solve and simulate the model at quarterly time intervals.

certainty. The first column of the monetary transmission matrix captures this impulse response function albeit with the sign reversed as the figure is showing a reduction in real interest rates.

The left panel of Figure 7 plots the impulse response function of output in the full model (solid blue line). Output expands on impact by 0.34%. But once the stimulus is removed at quarter $t = 1$, output falls below steady state by -0.16%. Subsequently, output gradually converges to steady state. The red dotted line shows that these dynamics are largely a reflection of durable expenditure. Intuitively, a reduction in the real interest rate reduces the opportunity cost of owning a durable. This induces households near the durable adjustment threshold to accelerate their durable purchases, which results in an increase in aggregate demand and output at $t = 0$. However, households that adjust at $t = 0$ are no longer interested in adjusting at $t = 1$ so durable demand falls below its steady state level at $t = 1$ and in subsequent quarters.

The intertemporal shifting of aggregate demand is thus an intuitive outcome of our model that stands in contrast to models without durable goods. The standard three-equation New Keynesian model is a useful benchmark to compare to. Not only is it familiar, but also its monetary transmission matrix is close to that of a version of our heterogeneous agent model without durables or sticky information, which is close to the incomplete-markets irrelevance case in Werning (2015) (see Figures A.5 and A.6 in the appendix). The dashed blue line in Figure 7 shows that the response of output in the three-equation model displays no reversal in demand following monetary accommodation.²⁵

In the right panel of Figure 7, we plot the corresponding *cumulative* increases in output, which are the integrals under the impulse response functions in the left panel. The solid blue line for the full model shows that half of the increase in output at $t = 0$ is subsequently reversed. The reversal is almost entirely accounted for by durable expenditure. In the long run, the distribution of durable holdings will return to the steady state distribution so there is a strong tendency for the purchases that are accelerated into date $t = 0$ to be offset by missing purchases going forward.²⁶

This example featured a sharp change in intertemporal prices between quarters 0 and 1, which

²⁵We simulate the three-equation model assuming an intertemporal elasticity of substitution equal to $\hat{\sigma}^{-1} = 1.38$, which implies that the (1, 1) element of the monetary transmission matrix matches the corresponding element in our full model.

²⁶The main reason that the increase in durable purchases is not fully reversed is that the monetary policy shock increases the number of adjustments that occur and each adjustment incurs adjustment costs that are counted as part of durable expenditure. The cumulative number of adjustments is not fully reversed because some households will be prompted to adjust by monetary policy and then soon after experience shocks that would induce another adjustment regardless of the monetary accommodation. In those situations monetary policy has created an extra adjustment.

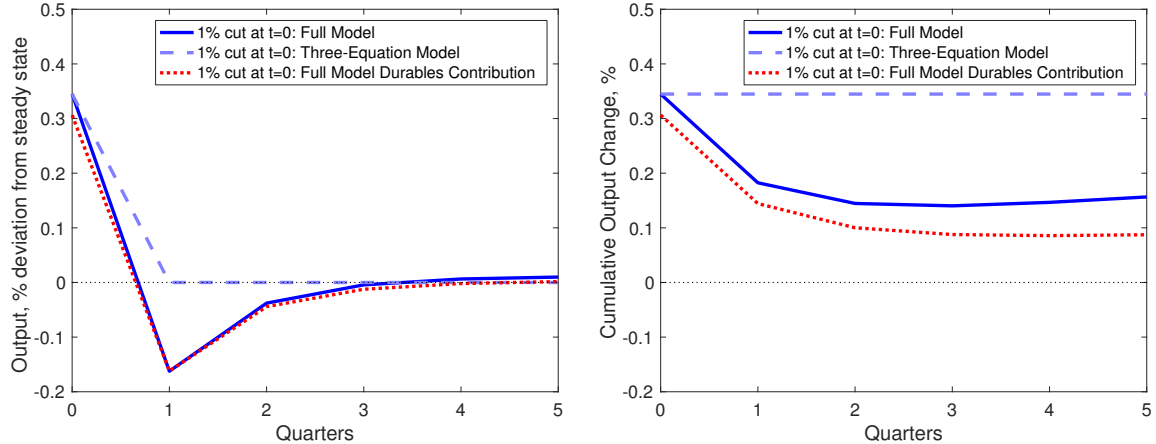


Figure 7: The left panel shows the percentage change in output following a 1% reduction in the real interest rate at time $t = 0$ in the full model with durables (solid line) compared to the standard New Keynesian model (dashed line). The right panel shows the corresponding cumulative change in output, i.e., the integral of the impulse response function in the left panel. The three-equation model has its elasticity of intertemporal substitution calibrated to match the initial output response to contemporaneous interest rate changes in the full model.

makes intertemporal shifting of aggregate demand obvious to see. In this sense, the example has more in common with the quasi-experimental evidence discussed in Section 3.3 (such as anticipated sales tax changes), than the monetary shocks in Section 3, which are very persistent so accommodation is only gradually reduced thereby obscuring the reversal.

5.2 The Power of Forward Guidance

We now turn our attention from the first column of the monetary transmission matrix to the first row. The first row shows the power of forward guidance because it captures the response of output at date t to news about real interest rates at future dates. Figure 8 plots the change in \hat{Y}_t in response to a reduction in r_{t+s} where s varies along the horizontal axis. The real interest rate is expected to fall by 1% (annualized) for one quarter. The negative slope in the figure shows that changes in future real interest rates have a smaller effect on current output than do contemporaneous interest rates. This stands in contrast to the standard consumption Euler equation that is embedded in the three-equation model, which implies that real interest rates at any horizon have the same effect on current output.

Why is it that durable goods demand dampens the power of forward guidance? In choosing to accelerate an adjustment, a household is giving up the interest on savings that it would have earned

until the adjustment would have otherwise occurred. Therefore, the opportunity cost of accelerating a durable purchase includes the interest foregone *over the period between the new purchase date and the original purchase date.*²⁷ The households that are induced to adjust their durables by monetary policy are those that were already near the adjustment threshold and therefore likely to adjust soon even in the absence of stimulus. Therefore, the opportunity cost of adjusting the timing is primarily affected by short-term interest rates. By contrast, forward interest rates at horizons beyond the time the household was already planning to adjust do not materially change the incentives to accelerate the timing of a purchase.

By this logic, the extensive margin of durable demand is particularly sensitive to short-term interest rates and less sensitive to forward interest rates. As a result, lumpy durables are an effective way to mitigate the “forward guidance puzzle” that arises in standard New Keynesian models. These models predict forward guidance is implausibly powerful because promises of lower interest rates arbitrarily far in the future have the same output effect today as contemporaneous interest rate cuts (Del Negro, Giannoni, and Patterson, 2015; McKay, Nakamura, and Steinsson, 2016).

Information rigidities are another force in the model that makes current output less sensitive to news about future interest rates. However, this force applies regardless of the horizon at which the interest rate changes. The dashed line in Figure 8 shows the model without sticky information. Forward guidance is more powerful in this case but so too are contemporaneous rates. The negative slope in the figure is roughly the same indicating the power of forward guidance diminishes with the horizon of the guidance in a similar fashion whether or not the model includes the information rigidities.

To be clear, forward guidance still has important effects on output in our model; they are simply much weaker than those implied by the standard three-equation model.

6 Interest Rates During and After the Great Recession

We now turn our attention to evolution of the real interest rate, r , and the natural rate of interest, r^* , during and after the Great Recession. The intertemporal shifting effects in our model imply that

²⁷For a household that needs to borrow to finance the durable, the opportunity cost includes the additional interest paid by pulling the purchase forward.

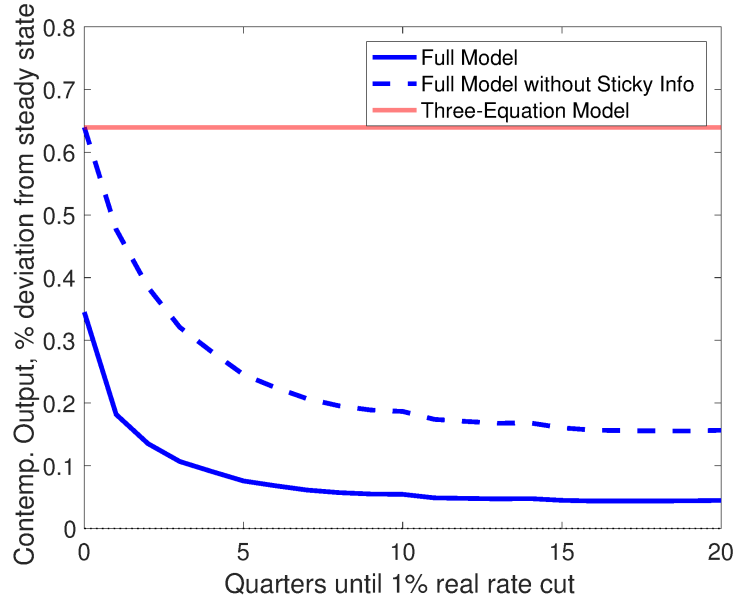


Figure 8: Output responses at date 0 to real interest rate cuts of 1% (annualized) that last for one quarter at horizons of 0 to 20. Full model refers to our model with lumpy durables and information rigidities. The three-equation model has its elasticity of intertemporal substitution calibrated to match the initial output response to contemporaneous interest rate changes in the model without information rigidities.

interest rate changes are very persistent. The model therefore predicts that the accommodative monetary policy during the Great Recession will be followed by low interest rates for many years.

6.1 The Great Recession Through the Lens of Our Model

We use a filtering approach to extract the shocks that account for the aggregate time series during the Great Recession. We seek to match four aggregate time series from 1991-2019:²⁸ the output gap, \hat{Y}_t , constructed using the CBO's estimate of potential output; the change in the durable expenditure share (relative to potential GDP), $s_t^x - s_{t-1}^x$ where $s_t^x = \frac{p_t x_t}{Y_t}$; the demeaned ex-ante real interest rate, $r_t - \bar{r}$, based on the Federal Funds Rate net of average nondurable inflation from 1991-2007; and the demeaned spread of the 30-year mortgage rate over the ten-year treasury yield,

²⁸We chose 1991 as a starting date for three reasons. First, it is sufficiently distant from the Great Recession that the initial state of the economy should have little effect on the dynamics of the economy during the Great Recession. Second, the real rate displays no trend from 1991 through 2007, which side-steps issues for how to detrend the real rate. Third, the persistence of inflation is small and statistically insignificant after 1991, and this is a key input the determination of non-durable inflation expectations and thus the ex ante real rate. (The persistence is much higher before 1991.)

$r_t^b - \bar{r}^b$.²⁹ We will extract four shocks from these series: the permanent productivity shock ϵ_t^Z (equivalent to a permanent income shock), the non-household demand shock ϵ_t^G ,³⁰ the monetary policy shock ϵ_t^r , and the shock to the borrowing spread $\epsilon_t^{r^b}$.

We use a novel filtering algorithm that we describe in detail in Appendix F. For each of the shocks we construct the impulse response functions of $\{\hat{Y}, \Delta s^x, r - \bar{r}, r^b - \bar{r}^b\}$, which are reported in Appendix H. We then proceed recursively: at each date t , we create a forecast for $\{\hat{Y}_t, \Delta s_t^x, r_t - \bar{r}, r_t^b - \bar{r}^b\}$ based on all the previous shocks we have filtered. We then solve for the shocks at date t that explain the difference between the data observed at date t and our forecast. We again make use of the assumption that the economy’s dynamics are linear in the shocks to perform this calculation. Specifically, the forecast for the data is a convolution of the previous shocks and the impulse response functions and solving for the date t shocks requires inverting a matrix of the impact response of each data series with respect to each shock. We initialize this procedure in 1991 assuming the economy is in steady state. This assumption has negligible effects on the shocks we filter during the Great Recession. This filtering method is equivalent to the Kalman filter (and also the Kalman smoother) given that there is no measurement error and the initial state is known with certainty (see Appendix F). The benefit of this approach to filtering is that it relies only on impulse response functions and does not require a state transition matrix, which is not readily available for a heterogeneous agent model in which the state includes a distribution.

The ELB was an important constraint on monetary policy during the Great Recession, and we explicitly incorporate it into our filtering procedure. In measuring the ex ante short-term real interest rate, we assume expected inflation is constant. Given this assumption, the ELB on the Federal Funds rate directly translates into an effective lower bound for the ex ante real rate, $r \geq \underline{r} = -2.5$.³¹ As the measured real interest rate never violates this constraint, our filtering algorithm naturally imposes it through realizations of monetary policy shocks. However, as these shocks are expected to dissipate over time, it is possible for the expected to violate the constraint. To ensure this does not happen, we incorporate monetary news shocks. At any point in time t for horizon $h > 1$ we calculate the extent to which the path violates the ELB, $\mathbb{E}_t r_{t+h}^{news} = \max\{\mathbb{E}_t r_{t+h} - \underline{r}, 0\}$.

²⁹We demean the series using the mean from 1991-2007, so as to not incorporate the downward-trend in the real rate during the Great Recession.

³⁰In this section, we label the government spending shock “non-household demand shock” since its role is to account for the residual output gap that cannot be explained by the shocks to households.

³¹The bound is equal to the ELB on the Federal Funds rate, equal to 0.15%, net of the average nondurable inflation from 1991-2007, which is equal to 2.65%.

We then append these shocks $\{\mathbb{E}_t r_{t+h}^{news}\}_{h=1}^H$ to our forecast. Because the monetary news shocks affect the variables we target, $\{\hat{Y}_t, \Delta s_t^x, r_t - \bar{r}, r_t^b - \bar{r}^b\}$, we must also update our inference on the shocks $\epsilon_t^Z, \epsilon_t^G, \epsilon_t^r, \epsilon_t^{r^b}$.³² The updated set of shocks and the monetary news shocks imply a new forecasted path for the real rate. We again check whether it violates the ELB. If it does, we keep iterating on this procedure until the ELB constraint is satisfied.

The shocks are uniquely identified as they imply very different impulse response functions for $\{\hat{Y}, \Delta s^x, r - \bar{r}, r^b - \bar{r}^b\}$. A permanent decline in productivity causes a durable overhang, with a large reduction in durable spending, a negative output gap, and a reduction in the real rate as the central bank accommodates. A negative non-household demand shock causes a negative output gap along with an increase in durable spending to potential GDP as accommodative monetary policy stimulates durable expenditure. A contractionary monetary policy shock causes a negative output gap and reduces durable spending relative to potential GDP, accompanied by an increase in the real rate. Finally, a shock to the borrowing spread is easily identified as it is the sole source of variation in r^b .

In Figure 9 we plot the filtered shocks scaled by their impact on the contemporaneous output gap. Our procedure identifies large negative productivity shocks during the Great Recession owing to the persistent weakness in durable spending along with low real interest rates. A permanent decline in productivity is akin to an overbuilding shock in the sense that the economy now has more durables than it would like. With some delay our filter also identifies a fall in non-household demand as the productivity shocks are not sufficient to explain the decline in the output gap. While borrowing spread shocks are prevalent in the run-up to the financial crisis they have little effect on the output gap. Monetary policy shocks tend to be slightly positive throughout the Great Recession, as the other shocks predict a decline in real interest rates that is not feasible due to the ELB constraint.

6.2 The Slow Normalization of Interest Rates after the Great Recession

Figure 10 shows that our model predicts a very slow normalization of interest rates after the Great Recession. In the left panel we plot a forecast of the real interest rate in the model against the data. To construct the figure we use the filtered shocks up to 2012Q4 and then use the model to

³²The effect of the zero lower bound is therefore conveniently captured by the Jacobians of $\{\hat{Y}, \Delta s^x, r - \bar{r}, r^b - \bar{r}^b\}$ with respect to current and expected future real interest rates.

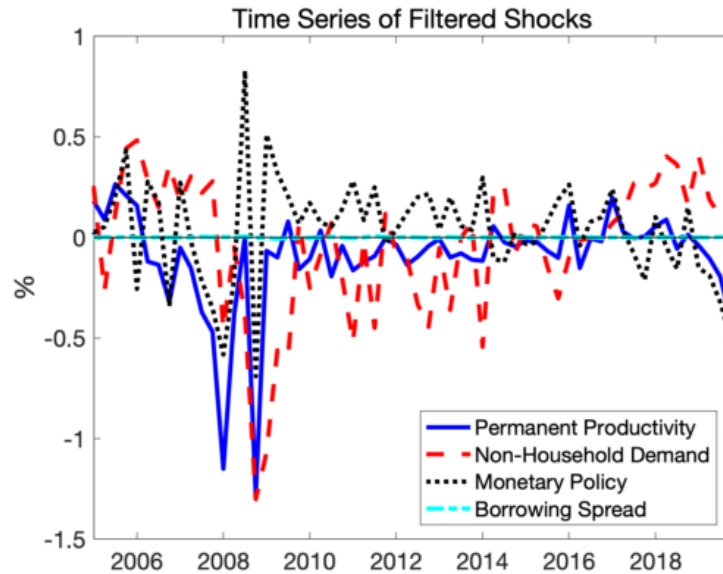


Figure 9: Aggregate shocks filtered such that the model exactly replicates the aggregate time series of the output gap, the real rate of interest, and the change in durable spending as a fraction of potential output. In the figure, the shocks are scaled based on their contemporaneous effect on the output gap.

forecast real interest rates through 2019Q4. The model predicts that the ELB will continue to bind until 2015Q4, when lift-off did in fact occur. Even after the lift-off, the normalization occurs slowly. The model predicts the real interest rate will remain 2.7 percentage points below steady state in 2017Q4 and 1.8 percentage points below steady state in 2019Q4, which closely tracks the interest rate path that came to be realized. The model fits the short-term real interest rate exactly up to 2012Q4 by construction because the short rate is one of the series we match in our filtering.

We use the beginning of 2013 as a benchmark for our forecasts because the taper tantrum of May 2013 reflected the first intentions of the Federal Reserve to begin the tightening cycle. But the model predicts a very persistent ELB episode even early on in the Great Recession. In the right panel of Figure 10, we plot the 5-year real interest rate from the model, computed using the expectations hypothesis of the term structure, against the 5-year TIPS yield. The long-term real rate in the model drops dramatically at the onset of the Great Recession because agents expect a prolonged zero interest rate policy. For example, as of 2009, agents expect the zero lower bound to bind for the following 18 quarters. The TIPS yield displays a much more gradual decline, as financial markets underestimated the duration of the zero interest rate policy.³³ As of 2012Q4, our

³³For example, in 2010 Eurodollar futures implied the federal funds rate would lift off after one to two years (see p.53 of Federal Reserve, 2010) while in the end it did not occur until the end of 2015.

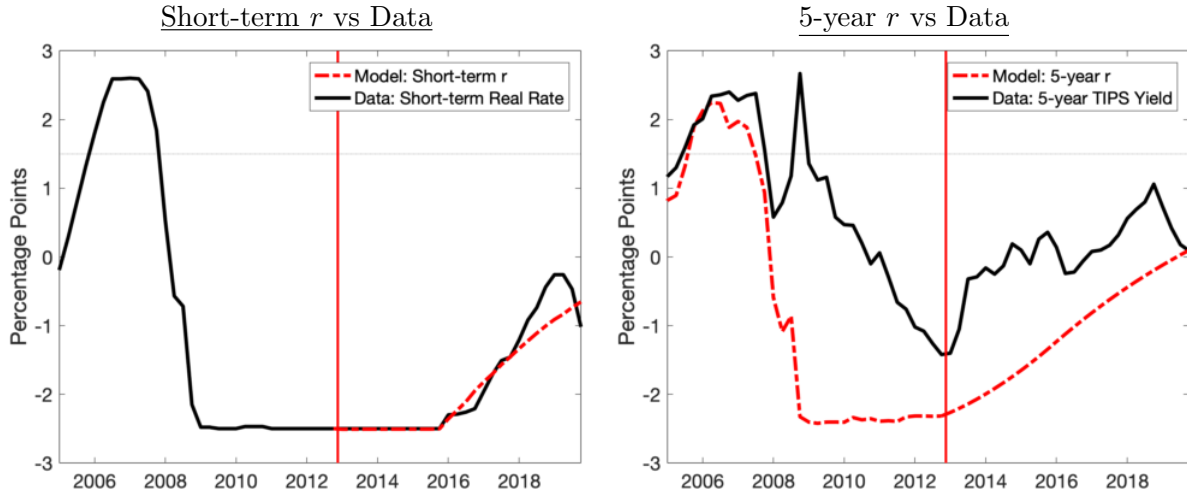


Figure 10: Left panel: short-term (contemporaneous) real interest rate from the model (forecasted after 2012Q4), and ex-ante real interest rate in the data. Right panel: The 5-year real interest rate (forecasted after 2012Q4), as well as the 5-year TIPS yield. The dashed line is the steady state real interest rate, equal to 1.5%

model predicts a slow normalization of long-term real rates to -1.4 percentage point below steady state in 2019Q4, which is broadly in line with the data.

According to our simulation, cyclical factors occurring before 2013 generate the persistently low levels of interest rates during and after the Great Recession and the late lift-off of interest rates in December 2015. This suggests that one does not need to appeal to secular forces, such as demographics, to explain the behavior of real interest rates over this period.

As we show below, the slow normalization of real interest rates reflects a slow normalization of the natural rate of interest, r^* . In turn, r^* normalizes slowly primarily because low interest rates themselves reduce future demand thereby bringing about low interest rates again in subsequent periods thereby propagating low rates forward in circular fashion.

6.3 Definition of r^* and the Role of Intertemporal Shifting

We define r^* as the real interest rate that is consistent with a zero output gap. To implement this definition we must account for the fact that the current output gap depends not just on the contemporaneous interest rate but also on expectations of future real interest rates. Therefore, at date t we seek a path for real interest rates going forward that is consistent with a zero output gap at t and an expectation that the output gap will remain zero going forward.

We can calculate r^* by building on ideas from the sequence space approach to heterogeneous agent models (see Auclert et al., 2019). Define $\vec{Y}_t \equiv (\hat{Y}_t, \mathbb{E}_t \hat{Y}_{t+1}, \mathbb{E}_t \hat{Y}_{t+2}, \dots)'$ as the vector of time t forecasts of output gaps at all future dates, $\vec{r}_t \equiv (r_t - \bar{r}, \mathbb{E}_t r_{t+1} - \bar{r}, \mathbb{E}_t r_{t+2} - \bar{r}, \dots)'$ as the vector of time t forecasts of real interest deviations at all future dates, and define $\eta_t \equiv (Z_t, G_t - \bar{G}, r_t^b - \bar{r}_t^b)'$ as the vector of exogenous stochastic processes. In Appendix G we show using a first order approximation around steady state that

$$\vec{Y}_t = \mathcal{M}\vec{r}_t + \mathcal{Q}\eta_t + \mathcal{D}(\Phi_t - \bar{\Phi}), \quad (15)$$

where Φ_t is the distribution of households over idiosyncratic states. In this context, Φ_t can be interpreted as a vector that gives a discrete representation of the distribution as in the Reiter (2009) method of solving heterogeneous agent models. \mathcal{M} is the monetary transmission matrix we introduced in Section 5 in which the (i, j) element gives the sensitivity of the output gap at horizon $i - 1$ to real interest rate changes at horizon $j - 1$. The matrix \mathcal{Q} contains the impulse response functions of shocks to η on the output gap, with the k^{th} column corresponding to the k^{th} shock and row i corresponding to horizon $i - 1$. The matrix \mathcal{D} captures how changes in the distribution of idiosyncratic states affect the output gap.

\vec{r}_t^* is defined as the vector that sets $\vec{Y}_t = 0$. Using equation (15) we have

$$\vec{r}_t^* = -\mathcal{M}^{-1} \left(\mathcal{Q}\eta_t + \mathcal{D}(\Phi_t - \bar{\Phi}) \right). \quad (16)$$

\vec{r}_t^* is a function of the state of the economy—as is clear from the equation above—and the past behavior of real interest rates affects \vec{r}_t^* through the distribution of individual states Φ_t .

r^* is often defined in terms of a flexible-price equilibrium, which would correspond to a situation in which \vec{r}_t is set to \vec{r}_t^* not only in the present and future, but also in the past. We call this the “full-accommodation” scenario. However, the Federal Reserve was not able to fully accommodate the shocks hitting the economy in the Great Recession due to the ELB. The fact that interest rates did not fall as much as the full-accommodation r^* did implies that there was less monetary stimulus and therefore less intertemporal shifting of demand. This difference, captured in evolution of Φ_t , then implies that r^* is higher under the actual history of monetary policy than under full accommodation. Therefore we consider an “actual-accommodation” scenario in which we simulate the evolution of the state economy using the real interest rate expectations we filter from the data. In both scenarios, \vec{r}_t^* is calculated from equation (16). The difference between them is the history of interest rates and therefore the value of Φ that enters this equation.

In Appendix G we show that we can compute r^* using the matrices \mathcal{M} and \mathcal{Q} as well as the filtered shocks and real rate expectations, assuming that we start in steady state at $t = 0$,

$$\begin{aligned} \bar{r}_t^* = & - \sum_{k=0}^{t-1} \mathcal{M}_{[1+t-k, 1+t-k..]}^{-1} \mathcal{M}_{[1+t-k, 1..t-k]} \left[\mathbb{E}_k \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} - \mathbb{E}_{k-1} \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} \right] \\ & - \sum_{k=0}^t \mathcal{M}_{[1+t-k, 1+t-k..]}^{-1} \mathcal{Q}_{[1+t-k, \dots]} \epsilon_k^\eta, \end{aligned} \quad (17)$$

where $\epsilon_t^\eta \equiv (\epsilon_t^Z, \epsilon_t^G, \epsilon_t^{r^b})'$. Importantly, we do not need to infer the infinite-dimensional state from the data to compute r^* . We do assume that we start in steady state at $t = 0$, but the importance of this assumption vanishes over time. In particular, making this assumption in 1991 has negligible consequences for our inference about r^* during the Great Recession.

To understand the intuition behind equation (17), start with the second term that involves the exogenous shocks ϵ_k^η and set $k = 0$. The matrix $\mathcal{Q}_{[1+t, \dots]}$ tells us how this shock at $t = 0$ affects the output gap from time t onward. The operation $-\mathcal{M}_{[1+t, \dots, 1+t..]}^{-1}$ then determines how interest rates from t onward must move to close this output gap. We repeat this idea for all shocks that have occurred from date 0 to date t and sum the implied interest rate movements. This is the contribution of the fundamental shocks to r^* at date t .

The first term in equation (17) is the contribution of intertemporal shifting. It captures how past interest rate movements affect r^* going forward. To understand the intuition, again start with $k = 0$. The term in square brackets tells us the news about the real rate path up to $t - 1$ that arrived at date 0.³⁴ The $\mathcal{M}_{[1+t, \dots, 1..t]}$ matrix multiplying this term converts these news into changes in the output gap for time t onward. And the $-\mathcal{M}_{[1+t, \dots, 1+t..]}^{-1}$ operation then determines how interest rates from t onward must move to close these output gaps. The sum then adds up all these intertemporal shifting effects from past real rate news.

Our concepts of full accommodation and actual accommodation differ only in the intertemporal shifting term. Under actual accommodation, the past interest rate expectations are the ones we filter from the data. Under full accommodation, the historical interest rates are those that close current and expected future output gaps.

³⁴News occur either because the real rate itself is subject to an exogenous shock or because it responds endogenously to one of the other shocks.

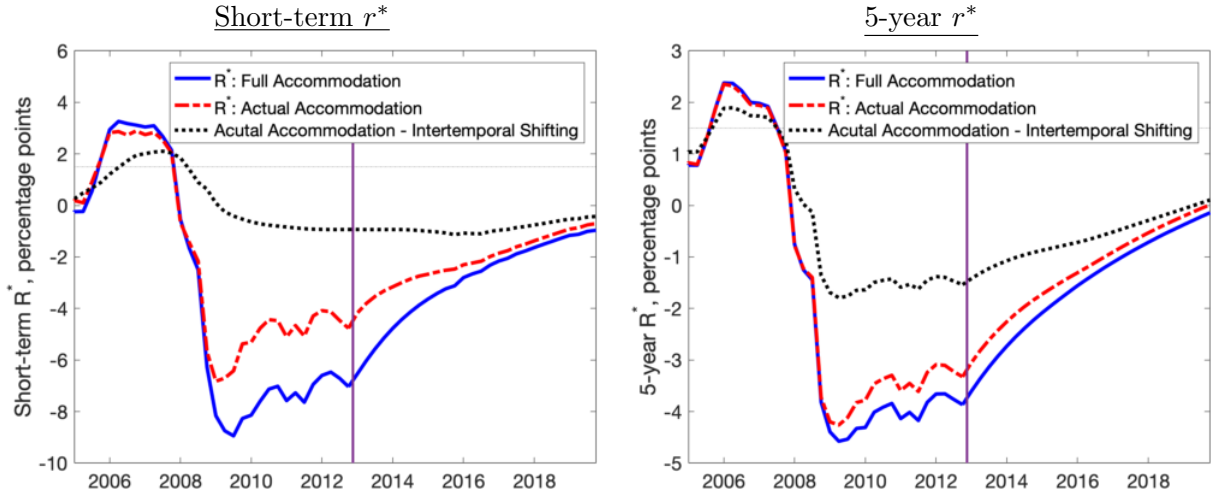


Figure 11: Time series of the short-term natural rate of interest (left panel) and the 5-year natural rate of interest (right panel) under the full-accommodation and actual-accommodation definitions. The dotted line shows the contribution of intertemporal shifting effects of previous real interest rates to the actual-accommodation r^* . We stop incorporating new shocks after 2012Q4, so the figure plots the expected path of the natural rate from 2013Q1 onward. The dotted horizontal line is the steady state real interest rate, equal to 1.5%.

Equation (17) shows that the below-diagonal elements of \mathcal{M} are key to determine the importance of intertemporal shifting. In our model, these elements are positive because lower real interest rates today borrow aggregate demand from the future. In a standard New Keynesian model these elements are zero, as past interest rates have no effect on the current or future output gaps. In DSGE models with higher-order adjustment costs these elements can even be negative, as stimulating investment today reduces the cost of investment tomorrow.

6.4 r^* During and After the Great Recession

The shocks that occurred during the Great Recession imply a substantial and very persistent decline in the natural rate of interest, which drives the model's forecast of low real interest rates. In the left panel of Figure 11 we plot the short-term natural rate of interest implied by our model—that is the first element of \tilde{r}_t^* for each date t . We show both the full-accommodation concept of r^* as well as the actual-accommodation concept in which we condition on the actual history of real interest rates. The former is the standard definition of r^* , while the latter is more useful for understanding the persistently low interest rates that were observed after the Great Recession. Both imply a sharp fall in r^* in 2008 and 2009. This collapse in r^* is not surprising as there was indeed a large negative

output gap despite low interest rates. The key result in the figure is that r^* recovers very gradually after the recession. In 2015Q4, when the federal funds rate lifted off from the ELB, the actual-accommodation r^* is forecast to be 4.0 percentage points below steady state. Even in 2019Q4, the actual-accommodation r^* is 2.2 percentage points below steady state. The full-accommodation r^* is even lower because it features lower real interest rates during the recession and therefore even more intertemporal shifting effects.

The persistent decline in r^* reflects the intertemporal shifting effect of low interest rates. The figure plots the contribution of intertemporal shifting to the actual-accommodation r^* . In 2015Q4, intertemporal shifting accounts for 2.6 percentage points of the decline in r^* . By the end of the forecast period, nearly all of the decline in r^* predicted by the model is due to intertemporal shifting.

The contribution of intertemporal shifting is nearly constant from 2009 through 2019. This reflects two features. First, because of the zero lower bound constraint, the real interest rate path is also almost flat over this period, which determines the amount of intertemporal shifting. Second, intertemporal shifting implies that the zero lower bound will be binding for a long time since low interest rates yesterday cause low interest rates today, which in turn cause low interest rates in the future.

In the right panel of Figure 11 we show that the 5-year r^* , computed from expectations of short-term rates using the expectations hypothesis, displays a corresponding large drop and slow normalization. The 5-year r^* is 2.9 percentage points below steady state in 2015Q4 and 1.5 percentage points below steady state in 2019Q4. Intertemporal shifting largely accounts for the very persistent behavior of these long-term rates.³⁵ Thus, our models also predicts that medium to long-term rates are significantly impacted by intertemporal shifting long after the end of the Great Recession.

These results show the decline in real interest rates during the Great Recession were themselves a key reason that the short-term and 5-year r^* remained low in subsequent periods. This finding stands in contrast to the prevailing neo-Wicksellian view that r^* is largely exogenous to monetary policy, and that policy must simply track this r^* path.

³⁵To compute the intertemporal shifting contribution to an H -horizon long-term rate, we solve for the intertemporal shifting contribution to the short rate at $t, t + 1, \dots, t + H$ and average them in accordance with the expectations hypothesis.

6.5 Comparison with Other Models

We repeat our filtering and r^* calculation with two alternative models in order to show (i) which features of the model are important to the results and (ii) that there is nothing mechanical in our procedure that drives the results.

We begin by considering a version of the full model without information rigidities. For comparability, we filter the data again and infer the set of shocks compatible with the full-information model. Figure 12 shows that under full information r^* does not fall as far as in the baseline model with information rigidities. This makes sense as information rigidities dampen the power of monetary policy stimulus (see Figure 8) so real interest rates must move by more to close the output gap. But both models predict a very gradual normalization of real interest rates in the forecast conditional on 2012Q4 information. This also makes sense as when we make the long-horizon forecasts there are few individuals who have not yet updated their information sets so information rigidities are not an important consideration.

We next consider the standard New Keynesian model. This model differs both in that forward guidance is more powerful and in that there are no intertemporal shifting effects. Of course, the three-equation model does not include durable expenditure, non-household demand shocks, or borrowing spreads so we cannot repeat our analysis directly in that model. Instead we conduct the following thought experiment: what would happen if the monetary transmission mechanism in the full-information model worked just as the monetary transmission mechanism in the three-equation model? To implement this experiment, we use \mathcal{M} from the three-equation model and \mathcal{Q} from the full-information model.³⁶ In this case the below diagonal elements of \mathcal{M} are all zero, and the diagonal and above-diagonal elements are all identical.

In this case, the r^* path returns to steady state much more quickly than in models with durables. This makes sense in light of the lack of intertemporal shifting effects in the standard New Keynesian model. Because the path of interest rates is expected to rise more rapidly in the three-equation model, long-term interest rates are higher (see the right panel). Even though forward guidance is more powerful in this model, the fact that there is less forward guidance stimulus implies a lower path of short-term interest rates during the recession is needed in the three-equation model as compared to the full-information model.

³⁶We set the EIS in the three-equation model to match the (1,1) element of \mathcal{M} in the two models.

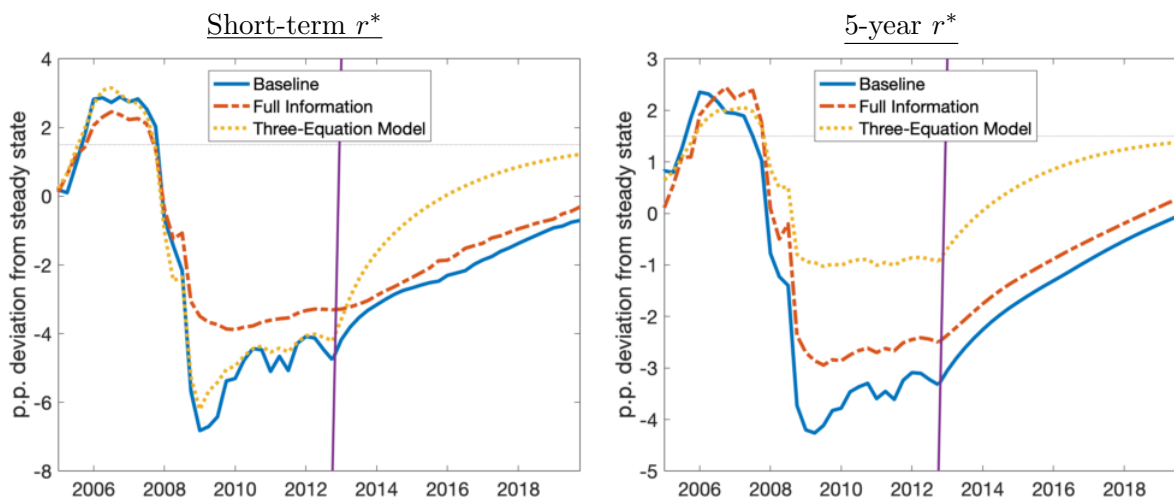


Figure 12: Actual accommodation r^* in our baseline model (blue solid line) against the model with full information (red dashed line) and the standard New Keynesian model (yellow dash-dotted line). The left panel plots the short-term r^* and the right panel the five-year r^* based on the expectations hypothesis. In each case we filter the data given the assumptions about monetary policy, so each model matches exactly the time series for the output gap, the contemporaneous real interest rate, the change in durable expenditure share to potential GDP, and the borrowing spread. From 2013Q1 onward the reported r^* is a forecast based on the history of shocks up to 2012Q4. The dotted horizontal line is the steady state real interest rate, equal to 1.5%.

The monetary transmission mechanism affects what shocks are filtered from the data. Because the standard New Keynesian model lacks strong internal propagation mechanisms that deliver persistently low real interest rates, it repeatedly infers that large negative shocks keep hitting the economy from 2008-2013. In contrast, our benchmark model infers a set of large shocks from 2008-2009, which then propagate into persistently low interest rates through intertemporal shifting and relatively weak forward guidance.

The fact that the standard New Keynesian model predicts a much more rapid normalization of real interest rates shows that there is nothing mechanical about our analysis that necessarily implies a slow normalization of r^* after the Great Recession. Rather, it is because the models with durables have a strong internal propagation of the reduction in aggregate demand that r^* remains persistently low.

7 Conclusion

We develop a fixed-cost model that is suitable to analyze the monetary transmission mechanism. The model is broadly consistent with both the microeconomic lumpiness of durable adjustments while at the same time consistent with the aggregate response of the economy to changes in interest rates.

Our model predicts that real interest rates will remain low for many years following the Great Recession in line with the realized path of real interest rates. While secular forces may well explain the decline of real interest rates starting in the 1980s, our model provides an alternative explanation for the particularly low interest rates in the 2010s.

The persistence of real interest rates in the model reflects intertemporal shifting of demand for durables. Low real interest rates shift the adjustment thresholds for durables, which induces households to pull forward durable adjustments. As these durable purchases are missing in subsequent periods, interest rates must be kept low to sustain aggregate demand. We show that the standard New Keynesian model, in which there is no such intertemporal shifting, is unable to explain the persistence of low real interest rates after the Great Recession.

The view we put forward here, in which r^* responds quite strongly to changes in monetary policy, contrasts with the neo-Wicksellian paradigm, in which r^* is generally considered to be independent of monetary policy. One implication of our analysis is that monetary accommodation has a side effect of reducing future policy space by reducing r^* towards the lower bound implied by the ELB. While our analysis focuses on a positive description of the economy, the intertemporal shifting of demand we highlight may have important implications for optimal monetary policy.

References

- ACHDOU, Y., J. HAN, J.-M. LASRY, P.-L. LIONS, AND B. MOLL (2017): “Income and wealth distribution in macroeconomics: A continuous-time approach,” Tech. rep., National Bureau of Economic Research.
- ALBOUY, D., G. EHRLICH, AND Y. LIU (2016): “Housing demand, cost-of-living inequality, and the affordability crisis,” Tech. rep., National Bureau of Economic Research.
- ALVAREZ, F., H. LE BIHAN, AND F. LIPPI (2016): “The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach,” *American Economic Review*, 106, 2817–51.
- ANGELETOS, G.-M. AND C. LIAN (2018): “Forward guidance without common knowledge,” *American Economic Review*, 108, 2477–2512.
- AUCLERT, A. (2019): “Monetary policy and the redistribution channel,” *American Economic Review*, 109, 2333–67.
- AUCLERT, A., B. BARDÓCZY, M. ROGNLIE, AND L. STRAUB (2019): “Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models,” Tech. rep., National Bureau of Economic Research.
- AUCLERT, A. AND M. ROGNLIE (2018): “Inequality and aggregate demand,” Tech. rep., National Bureau of Economic Research.
- AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2020): “Micro Jumps, Macro Humps: monetary policy and business cycles in an estimated HANK model,” Tech. rep., National Bureau of Economic Research.
- BACHMANN, R., R. J. CABALLERO, AND E. M. ENGEL (2013): “Aggregate implications of lumpy investment: new evidence and a DSGE model,” *American Economic Journal: Macroeconomics*, 5, 29–67.
- BACHMANN, R. AND D. COOPER (2014): “The Ins and Arounds in the US Housing Market,” Working paper.
- BAKER, S. R., L. KUENG, L. MCGRAHAN, AND B. T. MELZER (2019): “Do Household Finances Constrain Unconventional Fiscal Policy?” *Tax Policy and the Economy*, 33, 1–32.
- BARSKY, R., C. L. HOUSE, AND M. KIMBALL (2003): “Do Flexible Durable Goods Prices Undermine Sticky Price Models?” Tech. rep., National Bureau of Economic Research.
- BARSKY, R. B., C. L. HOUSE, AND M. S. KIMBALL (2007): “Sticky-price models and durable goods,” *American Economic Review*, 97, 984–998.
- BERGER, D. AND J. VAVRA (2015): “Consumption dynamics during recessions,” *Econometrica*, 83, 101–154.
- BERGER, D. W., K. MILBRADT, F. TOURRE, AND J. VAVRA (2018): “Mortgage Prepayment and Path-Dependent Effects of Monetary Policy,” Tech. rep., National Bureau of Economic Research.
- BOPPART, T., P. KRUSELL, AND K. MITMAN (2018): “Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative,” *Journal of Economic Dynamics and Control*, 89, 68–92.
- CARLSTROM, C. T., T. S. FUERST, AND M. PAUSTIAN (2015): “Inflation and output in New

- Keynesian models with a transient interest rate peg,” *Journal of Monetary Economics*, 76, 230–243.
- CARROLL, C. D., E. CRAWLEY, J. SLACALEK, K. TOKUOKA, AND M. N. WHITE (2018): “Sticky expectations and consumption dynamics,” Tech. rep., National Bureau of Economic Research.
- CASHIN, D. AND T. UNAYAMA (2011): “The intertemporal substitution and income effects of a VAT rate increase: Evidence from Japan,” *Discussion papers from Research Institute of Economy, Trade and Industry*.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113, 1–45.
- COIBION, O. AND Y. GORODNICHENKO (2012): “What can survey forecasts tell us about information rigidities?” *Journal of Political Economy*, 120, 116–159.
- COIBION, O., Y. GORODNICHENKO, L. KUENG, AND J. SILVIA (2017): “Innocent Bystanders? Monetary policy and inequality,” *Journal of Monetary Economics*, 88, 70–89.
- CORREIA, I., E. FARHI, J. P. NICOLINI, AND P. TELES (2013): “Unconventional fiscal policy at the zero bound,” *American Economic Review*, 103, 1172–1211.
- CZAJKA, J. L., J. E. JACOBSON, AND S. CODY (2003): “Survey estimates of wealth: A comparative analysis and review of the Survey of Income and Program Participation,” *Soc. Sec. Bull.*, 65, 63.
- DAVIDOFF, T. AND J. YOSHIDA (2013): “Estimating consumption substitution between housing and non-housing goods using macro data,” Tech. rep., Working paper.
- DAVIS, M. A. AND J. HEATHCOTE (2007): “The price and quantity of residential land in the United States,” *Journal of Monetary Economics*, 54, 2595–2620.
- DAVIS, M. A., A. LEHNERT, AND R. F. MARTIN (2008): “The Rent-price ratio for the aggregate stock of owner-occupied housing,” *Review of Income and Wealth*, 54, 279–284.
- DAVIS, M. A. AND F. ORTALO-MAGNÉ (2011): “Household expenditures, wages, rents,” *Review of Economic Dynamics*, 14, 248–261.
- DEL NEGRO, M., M. P. GIANNONI, AND C. PATTERSON (2015): “The Forward Guidance Puzzle,” Working paper.
- D’ACUNTO, F., D. HOANG, AND M. WEBER (2016): “The effect of unconventional fiscal policy on consumption expenditure,” Tech. rep., National Bureau of Economic Research.
- EGGERTSSON, G. B., N. R. MEHROTRA, AND J. A. ROBBINS (2019): “A model of secular stagnation: Theory and quantitative evaluation,” *American Economic Journal: Macroeconomics*, 11, 1–48.
- EICHENBAUM, M., S. REBELO, AND A. WONG (2018): “State Dependent Effects of Monetary Policy: The Refinancing Channel,” Tech. rep., National Bureau of Economic Research.
- ERCEG, C. J., D. W. HENDERSON, AND A. T. LEVIN (2000): “Optimal monetary policy with staggered wage and price contracts,” *Journal of monetary Economics*, 46, 281–313.
- FARHI, E. AND I. WERNING (2017): “Monetary policy, bounded rationality, and incomplete markets,” Tech. rep., National Bureau of Economic Research.

- FEDERAL RESERVE (2010): “Report to the FOMC on Economic Conditions and Monetary Policy: Tealbook A,” Tech. rep., Board of Governors of the Federal Reserve System, september 15. <https://www.federalreserve.gov/monetarypolicy/files/FOMC20100921tealbooka20100915.pdf>.
- FLODEN, M. AND J. LINDÉ (2001): “Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance?” *Review of Economic dynamics*, 4, 406–437.
- GABAIX, X. (2016): “A behavioral New Keynesian model,” Tech. rep., National Bureau of Economic Research.
- GOLOSOV, M. AND R. E. LUCAS (2007): “Menu costs and Phillips curves,” *Journal of Political Economy*, 115, 171–199.
- GOOLSBEE, A. (1998): “Investment Tax Incentives, Prices, and the Supply of Capital Goods*,” *The Quarterly Journal of Economics*, 113, 121–148.
- GORNEMANN, N., K. KUESTER, AND M. NAKAJIMA (2016): “Doves for the rich, hawks for the poor? Distributional consequences of monetary policy,” CEPR Discussion Paper No. DP11233.
- GUERRIERI, V. AND G. LORENZONI (2017): “Credit crises, precautionary savings, and the liquidity trap,” *The Quarterly Journal of Economics*, 132, 1427–1467.
- HAMILTON, J. D. (1994): *Time series analysis*, Princeton, NJ: Princeton University Press.
- HOUSE, C. L. (2014): “Fixed costs and long-lived investments,” *Journal of Monetary Economics*, 68, 86–100.
- HOUSE, C. L. AND M. D. SHAPIRO (2008): “Temporary investment tax incentives: Theory with evidence from bonus depreciation,” *American Economic Review*, 98, 737–68.
- KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): “Monetary policy according to HANK,” *American Economic Review*, 108, 697–743.
- LAUBACH, T. AND J. C. WILLIAMS (2016): “Measuring the natural rate of interest redux,” *Business Economics*, 51, 57–67.
- LI, G., R. F. SCHOENI, S. DANZIGER, AND K. K. CHARLES (2010): “New Expenditure Data in the PSID: Comparisons with the CE,” *Monthly Labor Review*, 133, 29–39.
- MANKIW, N. G. AND R. REIS (2002): “Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve,” *The Quarterly Journal of Economics*, 117, 1295–1328.
- MAVROEIDIS, S., M. PLAGBORG-MØLLER, AND J. H. STOCK (2014): “Empirical evidence on inflation expectations in the New Keynesian Phillips Curve,” *Journal of Economic Literature*, 52, 124–88.
- MCKAY, A., E. NAKAMURA, AND J. STEINSSON (2016): “The power of forward guidance revisited,” *American Economic Review*, 106, 3133–58.
- MCKAY, A. AND R. REIS (2016): “The role of automatic stabilizers in the US business cycle,” *Econometrica*, 84, 141–194.
- MIAN, A., L. STRAUB, AND A. SUFI (2019): “Indebted Demand,” Harvard University working paper.
- MIAN, A. AND A. SUFI (2012): “The effects of fiscal stimulus: Evidence from the 2009 cash for

- clunkers program,” *The Quarterly journal of economics*, 127, 1107–1142.
- MIDRIGAN, V. (2011): “Menu costs, multiproduct firms, and aggregate fluctuations,” *Econometrica*, 79, 1139–1180.
- MISHKIN, F. S. (1995): “Symposium on the monetary transmission mechanism,” *Journal of Economic perspectives*, 9, 3–10.
- NAKAMURA, E. AND J. STEINSSON (2010): “Monetary Non-neutrality in a Multisector Menu Cost Model,” *The Quarterly Journal of Economics*, 125, 961–1013.
- OGAKI, M. AND C. M. REINHART (1998): “Measuring intertemporal substitution: The role of durable goods,” *Journal of political Economy*, 106, 1078–1098.
- PAKOŠ, M. (2011): “Estimating intertemporal and intratemporal substitutions when both income and substitution effects are present: the role of durable goods,” *Journal of Business & Economic Statistics*, 29, 439–454.
- PLAGBORG-MØLLER, M. AND C. K. WOLF (2019): “Local projections and VARs estimate the same impulse responses,” *Unpublished paper: Department of Economics, Princeton University*, 1.
- REITER, M. (2009): “Solving heterogeneous-agent models by projection and perturbation,” *Journal of Economic Dynamics and Control*, 33, 649–665.
- REITER, M., T. SVEEN, AND L. WEINKE (2013): “Lumpy investment and the monetary transmission mechanism,” *Journal of Monetary Economics*, 60, 821 – 834.
- ROMER, C. D. AND D. H. ROMER (2004): “A new measure of monetary shocks: Derivation and implications,” *American Economic Review*, 94, 1055–1084.
- STERK, V. AND S. TENREYRO (2018): “The transmission of monetary policy through redistributions and durable purchases,” *Journal of Monetary Economics*, 99, 124–137.
- STRAUB, L. (2018): “Consumption, Savings, and the Distribution of Permanent Income,” Harvard University Working Paper.
- SUMMERS, L. H. (2015): “Demand Side Secular Stagnation,” *American Economic Review*, 105, 60–65.
- WERNING, I. (2015): “Incomplete markets and aggregate demand,” Tech. rep., National Bureau of Economic Research.
- WIELAND, J. F. AND M.-J. YANG (2017): “Financial dampening,” Tech. rep., National Bureau of Economic Research.
- WOODFORD, M. (2003): *Interest and Prices*, Princeton, NJ: Princeton University Press.

A Computational Appendix

We solve the model building on the routines available from Benjamin Moll's website <http://www.princeton.edu/~Emoll/HACTproject.htm> and described in Achdou et al. (2017).

A.1 Steady state

Define $k = a + \lambda d$ as the distance from the borrowing limit. Construct tensor grids over the state variables (k, d, z) . Then the steady state policy function is constructed as follows:

1. Start with an initial guess of the value function $v(k, d, z)$ and the value conditional on making an adjustment $v^*(k, d, z)$.
2. Solve for the optimal consumption and saving decisions when not adjusting. Compute v_k both as a forward difference v_k^f and as a backward difference v_k^b . At the boundaries of v_k^f and v_k^b impose that the drift of k is zero. Invert $v_k(k, d, z) = U_c(c, d)$ to solve for $c^f(k, d, z)$ and $c^b(k, d, z)$, and the corresponding drift of k , $s^f(k, d, z)$ and $s^b(k, d, z)$. Finally, let $c^0(k, d, z)$ be the consumption consistent with zero drift. Pick among the candidates based on the following rule:
 - (a) If $s^f < 0$ and $s^b < 0$ pick c^b, s^b .
 - (b) If $s^f > 0$ and $s^b > 0$ pick c^f, s^f .
 - (c) If $s^f < 0$ and $s^b > 0$ pick c^0, s^0 .
 - (d) If $s^f > 0$ and $s^b < 0$ pick the candidate that yields a larger value for the Hamiltonian.

Using the solution, compute the felicity function $u(c, d)$.

3. Construct the transition matrix A based on the endogenous drifts of k and the exogenous drifts and shocks to d, z . See Achdou et al. (2017) for details.
4. The HJB equation can now be written as $\min\{\rho v - u - Av, v - v^*\} = 0$, and solved using an LCP solver for v . We use Yuval Tassa's solver <http://www.mathworks.com/matlabcentral/fileexchange/20952>.
5. Compute optimal choice of d' conditional on adjusting and the corresponding $v^* = \max_{d'} v(k', d', z)$, where $k' = \lambda d' + k - (\lambda + f)d$.

6. Repeat steps 1-5 until convergence.
7. To obtain the steady state distribution, convert the policy functions for k' and d' conditional on an adjustment to index form. Fractions of an index determine the weights we assign to each index.
8. Create a matrix $C^{noadj} = A - \text{diag}(\theta)$. Then set all the columns in C^{noadj} that correspond to adjustment points to zero. Define $A^{adj} = A - C^{noadj}$. This matrix contains the mass at adjustment points that needs to be reallocated to the nodes of the optimal k', d' . We assign this mass to the nodes surrounding k', d' based on the index fractions in the previous step. This yields a matrix of adjustments C^{adj} . The transition matrix is then $C = C^{noadj} + C^{adj}$.
9. Solve $0 = C\Phi$ for the steady state distribution Φ .

A.2 Jacobians

The k^{th} column of the Jacobians hold the impulse response function with respect to a shock $k - 1$ periods in the future. The dimension of the Jacobian described here will be $T \times T$. The procedure largely follows Auclert et al. (2019) but we need to compute the Jacobian numerically since our model features non-differentiable policy rules.

1. Start with a shock T periods in the future and solve the policy function backwards, by repeating steps 1-5 given the terminal condition $v_{T+1} = v$. Each iteration reduces t by dt . Continue until $t = 0$ is reached.
2. Take the whole sequence of policy functions from v_0 to v_T . Repeat steps 7-8 for each period of the IRF and record outcomes for each period. This yields the IRF for the last column (T) of the Jacobian. Note that the initial distribution Φ_0 requires a modification if $p_0 \neq 1$. The distribution of k needs to shift since $k = a - \lambda pd$ and a, d are fixed in that instant.
3. For each t , if the adjustment thresholds change, then all the mass in Φ_t that is in the new adjustment region must be immediately shifted to its new location using the procedure in step 8. Call the new distribution $\hat{\Phi}_t$. Then compute $\Phi_{t+dt} = \hat{\Phi}_t + C_t \hat{\Phi}_t dt$. Repeat this step until $t = T$.

4. Repeat the previous two steps using the sequence of policy functions for v_k to v_T followed by $k - 1$ periods of the steady state policy function v_{T+1} . This yields the IRF for the $T - k$ column.
5. Conduct this procedure for a shock to G, p, Y, r, r^b .
6. For the productivity shock only the initial distribution gets rescaled, so there is no need to compute a policy function backward.

A.3 General Equilibrium

Following Auclert et al. (2019) we compute the partial equilibrium Jacobians for all outcome variables given news at time 0 to one-time deviations to $r_s, r_s^b, G_s, Y_s, p_s$, with rows corresponding to the quarter in which the outcome is measured and columns corresponding to periods in which the deviation occurs. Since we express the model variables relative to productivity Z_t , the productivity shock causes a rescaling of the initial aggregate distribution, which we capture by a Jacobian with a single column. Using matrix algebra we can then solve for the impulse response functions in general equilibrium by incorporating the persistence of exogenous variables and the necessary endogenous price and income movements that satisfy (11) and (13).

B Estimation

The estimation strategy largely follows Berger and Vavra (2015). The details of the data selection and estimation algorithm below are meant to facilitate replication of our results.

B.1 Data

We use PSID data from 1999 through 2009.

B.1.1 Variables

Real non-durable consumption is nominal non-durable consumption in the PSID deflated by the BEA price index for non-durables (NIPA table 1.1.4). Nominal non-durable consumption is the sum of food expenditures, utility expenditures, home insurance, transportation expenditures, property taxes, health expenditures, child care expenditures, and education expenditures. We exclude any

loan or lease payments from transportation expenditures to align the definition of non-durables with our model.

Real durable holdings are the sum of real house values and real vehicle values. Real house values are reported nominal house values deflated by the OFHEO national house price index. For renters we convert rent to a house value using the national house-to-rent ratio from Davis et al. (2008) available at <http://www.aei.org/housing/land-price-indicators/>. The PSID records the net wealth of up to three vehicles per household. We sum these values, add total vehicle debt (detailed below), and deflate the sum with the BEA price index for motor vehicles (NIPA table 1.2.4).

Real liquid asset holdings are the sum of cash and deposit holdings, stock holdings, and bond holdings, deflated by the non-durables price index.

We construct net real liquid assets by subtracting real debt from housing and vehicles. Mortgage debt is directly reported and we deflate it using the non-durables price index. We construct existing vehicle debt from the initial loan amount on all three cars and subtract the number of payments made times the average payment amount. In less than 1% of cases this results in a negative debt value, in which case we set vehicle debt to zero.

Housing adjustments come from either moving or a significant addition or repair. The PSID records the month and year of the most recent move since either the last interview (pre-2003) or since January two years ago. If a move is recorded and the move falls after the previous interview, then we code it as a housing adjustment for the current wave; otherwise it is an adjustment in the previous wave. When the move falls in the month of the interview we break the tie based on whether the interview was in the first or second half of the month. For significant additions and repairs we record them as housing adjustments in the wave that they are reported.

Car adjustments are set to one if any one of the three reported cars has been acquired since the previous wave. This is the case if the most recent car's acquisition date is after the previous wave's interview date, or (if there is insufficient information using the date) a new car has been acquired less than three years ago and it was not reported in the previous wave.

The sample weight is the household weight in the PSID.

B.1.2 Sample selection

We only keep head of households since the data is reported at the household level. We drop heads of households 21 and younger, as well as households present for fewer than 3 waves. This selection helps with the estimation of household fixed effects. We drop households with zero durable holdings, or those with missing information on any variable. We winsorize all variables at the 5th and 95th percentile.

B.1.3 Household fixed effects

We demean durable holdings by the households average durable holdings over the sample. This accounts for permanent differences in tastes for durables across households, which are not part of the model. We also divide non-durable consumption, liquid asset holdings, and real debt holdings by a household's average non-durable consumption over the sample. This helps account for permanent differences in income, which are again not part of the model.

B.1.4 Consistency with national aggregates

We divide all variables by average non-durable consumption in the sample. We then multiply each scaled variable (durables, liquid assets, debt) by a factor so that the sample average aligns with national aggregates from the fixed asset tables (durable-to-non-durable-consumption ratio) and the flow of funds (liquid-asset-to-non-durable-consumption and debt-to-non-durable-consumption). The rescaling is necessary because the PSID collects data for 72% of non-durable expenditures on average (Li et al., 2010). Further, households appear to overestimate the value of their vehicles (Czajka et al., 2003).

B.2 Estimation Algorithm

The steps of the algorithm are as follows:

1. Solve the model for a given intensity of match quality shocks θ .
2. Forecast the probability of adjustment $P(a, d, y)$ for each initial state (a, d, y) over the next two years. Also forecast the average non-durable consumption expenditure (including operating

and maintenance costs) \bar{c} for each initial state (a, d, y) over the next two years. From the latter we obtain a steady-state distribution $G(a, d, \bar{c})$.

3. Regress the optimal durable stock d^* in the model on $a, a^2, d, \bar{c}, d/\bar{c}$ weighted using the steady-state distribution. The vector of estimated coefficient is β .
4. Add measurement to the model variables a, d, \bar{c} using three independent Gaussian quadratures, one for each variable. This yields a new distribution $\hat{G}(a, d, \bar{c})$ which includes measurement error. Note that the adjustment probabilities are based on the true underlying $G(a, d, \bar{c})$.
5. Compute gaps $\omega = d^* - d$ for each point in the distribution \hat{G} . Integrating over ω using \hat{G} yields the pdf $f(\omega)$ in the model. Similarly integrating the probability of adjustment $P(a, d, y)$ over ω using \hat{G} yields the hazard rate $h(\omega)$ in the model.
6. In the data combine reported a, d, \bar{c} and our estimates β to predict d^* and the durable gap $\omega = d^* - d$. Use the sample weights to compute $f(\omega)$ and the adjustment hazard $h(\omega)$.
7. Compute loss function $L = \sum_{\omega} w(\omega)[|f^{model}(\omega) - f^{data}(\omega)| + |h^{model}(\omega) - h^{data}(\omega)|]$ where the weight is $w(\omega) = \frac{1}{4}(f^{model}(\omega) + f^{data}(\omega))(h^{model}(\omega) + h^{data}(\omega))$. This weighting function attaches more weight to bins the greater the fraction of adjustments accounted for by that bin. Conversely, we attach little weight to regions in which both model and data predict few adjustments.
8. Repeat steps 4,5, and 7 using a range of values for the standard deviation of the measurement error. Then pick the value that results in the smallest loss in 7.
9. Repeat steps 1-8 using a range of values for θ . Pick the θ with the smallest loss in 8.
10. To construct standard errors, sample 1000 new datasets with replacement from the original dataset. Repeat steps 6 and 7 for each dataset, record the loss-minimizing value for θ and the associated density and hazard function from both data and model.

C Data Appendix

C.1 Variables for Impulse Response Functions

In this section we detail how we construct the variables for the empirical impulse response functions to monetary policy shocks in figures 2 and 3 of section 3.

The base series the set of variables in Table 2, which we download from the St Louis FRED database.

Variable Name	FRED Series Code
Population	B230RC0Q173SBEA
Income (GDP)	GDPC1
Federal Funds Rate	FEDFUNDS
Consumer Durable Expenditure	PCDG
Residential Investment	PRFI
Consumer Non-durable Expenditure	PCEND
Consumer Service Expenditure	PCES
Consumer Housing Services Expenditure	DHSGRC0
Durable Price Index	DDURRD3Q086SBEA
Residential Investment Price Index	B011RG3Q086SBEA
Non-durable Price Index	DNDGRG3M086SBEA
Services Price Index	DSERRG3M086SBEA
Services Price Index: Housing	DHUTRG3Q086SBEA
Consumer Expenditure: Motor Vehicles	DMOTRC1Q027SBEA
Motor Vehicles Price Index	DMOTRG3Q086SBEA
House Price Index	USSTHPI
Residential Investment: Permanent Site	A943RC1Q027SBEA
Residential Investment: Other	A863RC1Q027SBEA
Residential Investment Price Index: Other	A863RG3Q086SBEA

Table 2: Variable names and FRED series code.

To construct real durable and non-durable expenditure, we proceed as follows. The problem is one where we have two components of nominal expenditure $Y_t = X_{1t} + X_{2t}$ (e.g., durable expenditure

equals consumer durables plus residential investment), and their respective price indices P_{1t} and P_{2t} . We want to construct the price index P_t for Y_t .

We first construct the growth rate of nominal spending, $\Delta y_t = \Delta \ln(Y_t) = \ln(Y_t) - \ln(Y_{t-1})$, and of the price indices, Δp_{1t} and Δp_{2t} . Define the share of good 1 in nominal expenditure, $s_{1t} = \frac{X_{1t}}{Y_t}$. Then the growth rate of the aggregate price index is $\Delta p_t = s_{1,t-1}\Delta p_{1t} + (1 - s_{1,t-1})\Delta p_{2t}$, from which we can construct the aggregate price index P_t . The growth rate of real expenditure is $\Delta y_t - \Delta p_t$, from which we can construct aggregate real expenditure. We convert all real expenditure to per capita by dividing by population.

We follow this procedure for both durables (consumer durables plus residential investment) and non-durables (consumer non-durables plus services net of housing).

For the price series of residential investment and consumer services we make specific modifications. We separate residential investment into investment into new structures and other residential investment. For investment into new structures we use the FHFA national house price index to capture changes in the price of land as well as the price of the new structure. For other residential investment we use that respective price index from the BEA. The weights are based on nominal expenditures in new residential structures and other residential investment and calculated as above.

For consumption of services we remove housing services because housing services in the model are obtained from H_t and not counted in C_t . To do so we follow the same procedure as above for the housing and non-housing component of services. But rather than adding, we subtract the housing component, $Y_t = X_{2t} - X_{1t}$. The share of rent in nondurable expenditure is $s_{1t} = -\frac{X_{1t}}{Y_t}$. With these two modifications, we can calculate real expenditure and the price index as above.

The relative price series for durables is the price of durables divided by the price of non-durables and services. The real interest rate is defined in terms of non-durables. It is the federal funds rate net of realized non-durable inflation from this quarter to the next quarter.

C.2 PSID: Housing Adjustment Probability

This section details how we construct the time series for the probability of housing adjustment using the Panel Survey of Income Dynamics (PSID). We use these data to measure the extensive margin response of housing to monetary shocks in figure ?? of section 3.

C.2.1 Sample

We keep only people who are heads of household and those who are in the Survey Research Center (SRC) sample. We only use PSID data from 1969-1997 when the survey frequency is annual.

C.2.2 Adjustment Series

We use the moved since spring series to create a record of adjustments. If moved since spring is true, we record an adjustment for that year. If moved since spring is false, we record no adjustment for that year. If the value is missing, we stop counting until the next time an adjustment is recorded, then count from there.

Tenure Status Tenure status refers to knowing what type of housing (rent or own). We use the series from the PSID labeled own/rent or what. There are also values for “other” or “not applicable.” If the values are “other” or “not applicable” we mark the tenure status as missing.

To align with Bachmann and Cooper (2014), we set values to missing if the observation does not have a tenure status or is lag does not have a tenure status. For example, if their observation is in the year 1992, we will set the adjustment series to missing if we do not know whether the head of household was owning or renting in either 1991 or 1992.

C.2.3 Weights

To weight observations in the regression, we take the mean of the family weight provided associated with the head of household across all periods for which we have year to year housing adjustments (1969-1997).

C.2.4 Time Series

We create a time series of the probability of adjustment by aggregating the panel using the family weight.

C.3 CEX: Car acquisition probability

We use the consumer expenditure (CEX) survey from 1980-2017 to construct a quarterly time series of the probability of a household acquiring a car or truck (used or new). We use this time series in table ?? to measure the extensive margin response of auto purchases to monetary policy shocks.

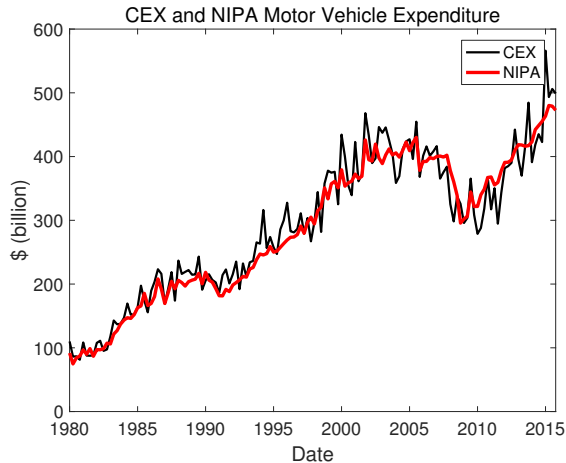


Figure A.1: Household motor vehicle expenditure from the personal consumption expenditure table in NIPA, and extrapolated from CEX micro-data based on NIPA definitions based on survey weights.

We download the pre-compiled files from the BLS for 1996 onwards. The CEX files before 1996 are from ICPSR. We only use the family and expenditure files. We clean the micro-data files following Coibion et al. (2017). We drop households with negative elderly expenditure, zero total expenditure, or zero food consumption. We sum expenditures that occur in the same month but are reported in different interviews as recommended by the BLS. We drop households that report more than three monthly expenditures per interview, since it is difficult to allocate the additional monthly expenditures. We also drop interviews with fewer than three months of data.

In the expenditure files we sum the UCC codes 450110 (new cars), 450210 (new trucks), 460110 (used cars), 460901 (used trucks). All expenditure series are net of trade-in value. This definition aligns with the BEA definition of motor vehicle expenditure. Using the household weights, total motor vehicle expenditure implied by the CEX tracks BEA personal consumption motor vehicle expenditure very well (Figure A.1).

We construct the probability of adjustment by setting an indicator equal to 1 whenever a household’s motor vehicle expenditures are positive, and aggregating the indicator using household weights.

C.4 CPS: Reasons for moving

Reasons for moving are from the March 2001 CPS supplement. We classify these reasons according to whether they broadly fall into life-cycle changes, which more closely correspond to our unmodeled

match-quality shocks, and sS-band reasons, which capture the traditional adjustments in fixed cost model from a depreciating durable stock.

We classify wanting a new/better home and cheaper housing is suggestive of durable holdings away from target and price-sensitivity in line with the traditional fixed cost model. These account for roughly a quarter of all adjustments, in line with our calibration of match-quality shocks. Changes in marital status or job-related moving more closely align with life-cycle changes that are less directly tied to interest rates. These account for roughly half of all transitions. The remaining quarter of moves are obviously classifiable in either category.

Table 3: What was ... main reason for moving?

	Frequency	sS-Band	Life-cycle
Change in marital status	0.84%	No	Yes
To establish own household	1.06%	No	Yes
Other family reason	1.92%	No	Yes
New job or job transfer	1.45%	No	Yes
To look for work or lost job	0.27%	No	Yes
To be closer to work	0.44%	?	?
Retired	0.08%	No	Yes
Other job-related reason	0.16%	No	Yes
Wanted to own home, not rent	1.41%	?	?
Wanted new or better house	2.47%	Yes	No
Wanted better neighborhood	0.55%	?	?
Cheaper housing	0.76%	Yes	No
Other housing reason	1.52%	?	?
Attend/leave college	0.43%	No	Yes
Change of climate	0.08%	No	Yes
Health reasons	0.19%	No	Yes
Total	13.64%	3.23%	6.48%

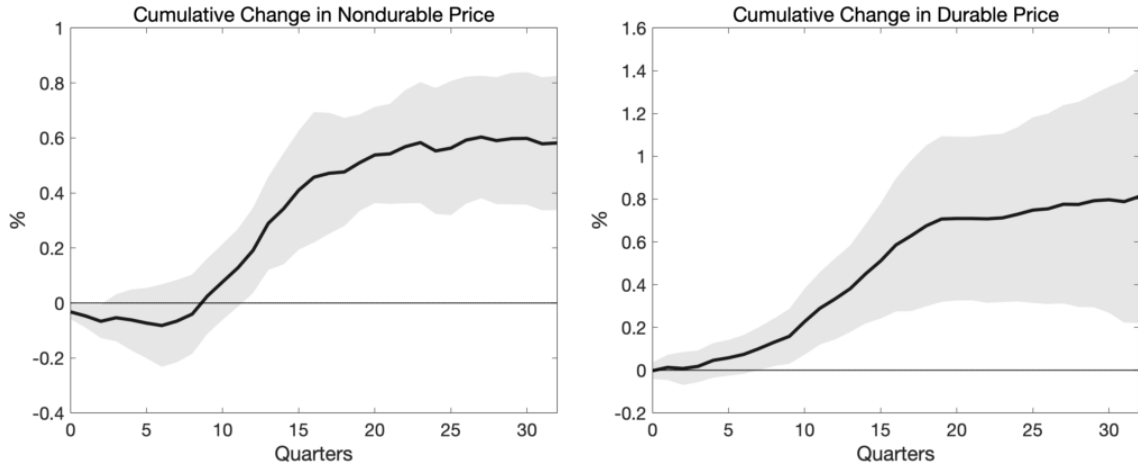


Figure A.2: Impulse response function for the level of non-durable prices (left panel) and durable prices (right panel) following a Romer and Romer (2004) monetary policy shock.

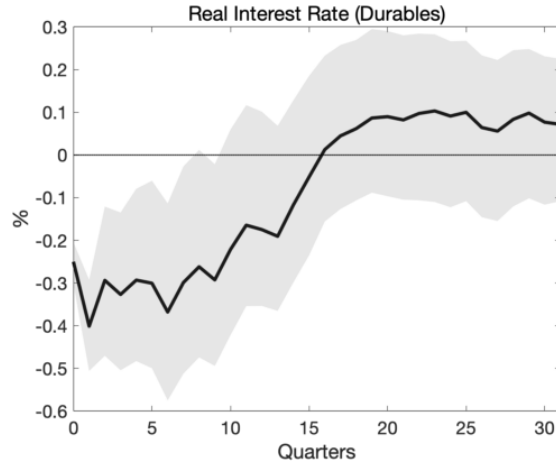


Figure A.3: Impulse response function for the real interest rate in terms of durables following a Romer and Romer (2004) monetary policy shock.

D Additional Results

D.1 Price levels: non-durables and durables

See Figure A.2.

D.2 Real rate in terms of durables

See Figure A.3.

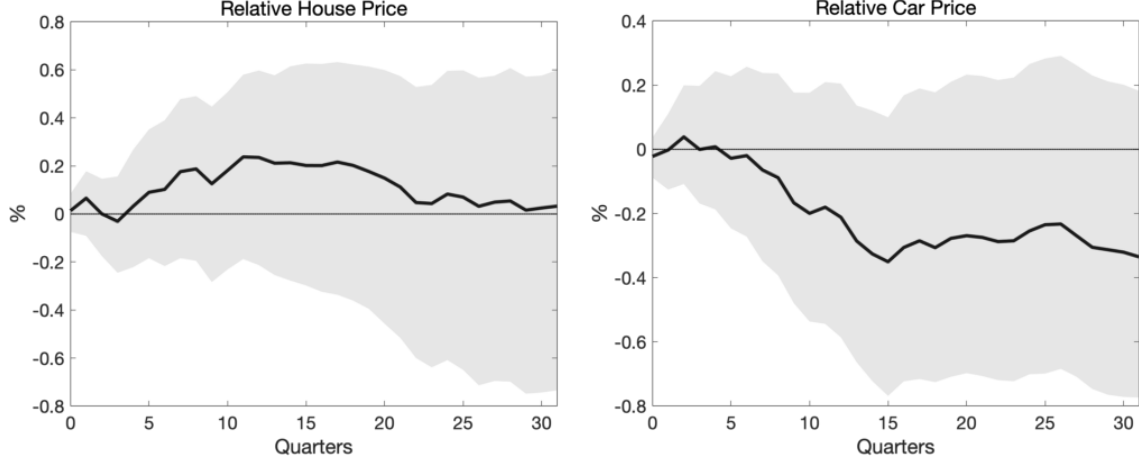


Figure A.4: Impulse response function for the price of residential investment relative to the price of non-durables and services (left panel) and the price of motor vehicles relative to the price of non-durables and services (right panel) following a Romer and Romer (2004) monetary policy shock.

D.3 Other relative prices: residential investment and cars

See Figure A.4.

D.4 Forward Guidance in Non-Durables Model

Figures A.5 and A.6 show the same experiments as Figures XX and YY in a version of our model without durables. As Figures A.5 and A.6 show, the non-durables model is similar to the three-equation New Keynesian model.

E Details of the General Equilibrium Model

E.1 The Labor Market

The labor demand curve of each labor type j is,

$$l_{jt} = L_t \left(\frac{W_{jt}}{W_t} \right)^{-\varphi}$$

so that the aggregate wage equal,

$$W_t = \left(\int_0^1 W_{jt}^{1-\varphi} dj \right)^{\frac{1}{1-\varphi}}$$

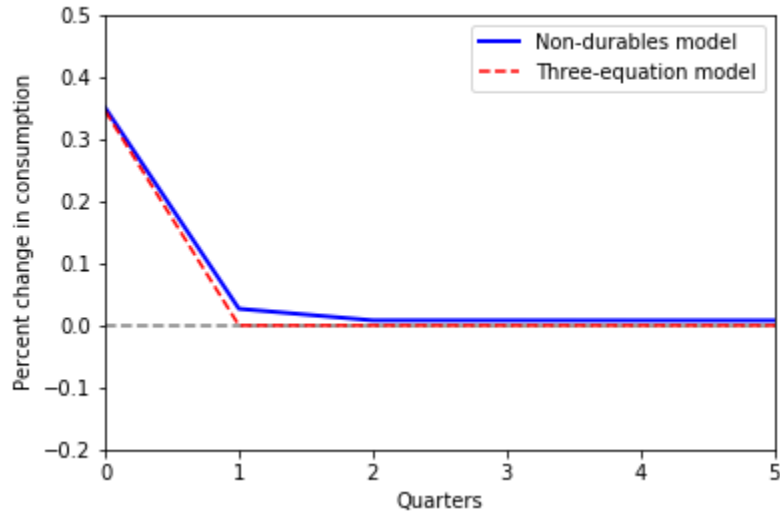


Figure A.5: Response of aggregate consumption to 1% interest rate cut at dates 0 in model with only non-durable consumption.

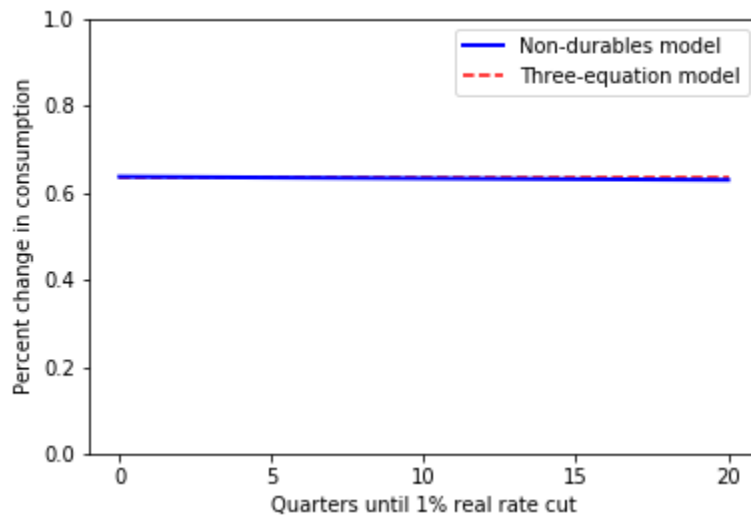


Figure A.6: Response of aggregate consumption to 1% interest rate cut at horizon given by the horizontal axis in model with only non-durable consumption.

The objective of the union is therefore,

$$\max_{\{\mu_{jt}\}} \int_{t=0}^{\infty} e^{-\rho t} \int_0^1 \left[u_c(c_{it}, d_{it}) \frac{W_{jt}}{P_t} l_{jt} - \bar{u}_{c,t} v(l_{jt}) - \frac{\Psi}{2} \bar{u}_{c,t} L_t \mu_{jt}^2 \right] di dt$$

subject to

$$\begin{aligned} d \ln W_{jt} &= \mu_{jt} dt \\ l_{jt} &= L_t \left(\frac{W_{jt}}{W_t} \right)^{-\varphi} \end{aligned}$$

The first order conditions of the union are,

$$\begin{aligned} \lambda_{jt} &= \Psi \mu_{jt} \bar{u}_{c,t} L_t \\ d\lambda_{jt} - \rho \lambda_{jt} dt &= -(1 - \varphi) \bar{u}_{c,t} \left(\frac{W_{jt}}{P_t} \right)^{1-\varphi} \left(\frac{W_t}{P_t} \right)^{\varphi} L_t dt - \varphi L_t \left(\frac{W_{jt}}{W_t} \right)^{-\varphi} \bar{u}_{c,t} v_L(l_{jt}) dt \end{aligned}$$

Imposing symmetry and market clearing

$$\begin{aligned} \Psi \mu_t \bar{u}_{c,t} \frac{Y_t}{Z_t} &= \lambda_t \\ d\lambda_t - \rho \lambda_t dt &= -(1 - \varphi) \bar{u}_{c,t} \frac{Y_t}{Z_t} dt - \varphi \bar{u}_{c,t} \frac{Y_t}{Z_t} v_L \left(\frac{Y_t}{Z_t} \right) dt \\ \pi_t &= \mu_t \end{aligned}$$

The non-linear Phillips curve is then,

$$d\pi_t = \left[\rho - \frac{d\bar{u}_{c,t}}{\bar{u}_{c,t}} - \frac{dY_t}{Y_t} + \frac{dZ_t}{Z_t} \right] \pi_t dt - \frac{\varphi - 1}{\Psi} \left[\frac{\varphi}{\varphi - 1} v_L \left(\frac{Y_t}{Z_t} \right) - 1 \right] dt$$

and log-linearized around the zero inflation state in which $Y_t = \bar{Y}_t = Z_t v_L^{-1} \left(\frac{\varphi - 1}{\varphi} \right)$,

$$d\pi_t = \rho \pi_t dt - \frac{\varphi}{\Psi} \eta \left(\frac{Y_t - \bar{Y}_t}{\bar{Y}_t} \right) dt$$

where $\frac{1}{\eta}$ is the Frisch elasticity. Letting $\kappa = \frac{\varphi \eta}{\Psi}$ gives (10).

E.2 Market Clearing

Non-durables market clearing:

$$Y_t = \int_0^1 c_{it} di + M_t + G_t + (r_t^b - r_t) \int_0^1 a_{it} I_{(a_{it} < 0)} di.$$

Durable goods market clearing:

$$X_t = \int_0^1 \left(\frac{dd_{it}}{dt} - \delta d_{it} \right) di + f \int_0^1 I_{d'_{it} \neq d_{it}} d_{it} + \nu \int_0^1 d_{it} di.$$

Bond market clearing:

$$\int_0^1 a_{it} di = A_t.$$

F Data Filtering Using the *MA* Representation

Heterogeneous-agent macro models are difficult to represent in a state-space form because the state of the model includes a distribution. For example, in the canonical Krusell-Smith model the distribution of wealth across households is part of the model's state. Because the state-space representation poses a challenge, authors sometimes prefer to work with impulse response functions as an alternative representation of the model's dynamics. These two representations are equivalent for a linear(ized) model (Boppart et al., 2018).

In this appendix we discuss how one can use the impulse response functions of the model to infer the shocks that generated the observed time series data. That is we implement a restricted version of the Kalman filter to recover the shocks. Our principal aim is to recover an estimate of the shocks that generated the data. One can also compute the likelihood of the data. (Auclert et al., 2019) describe an alternative method for computing the likelihood using the MA representation of the model that is not subject to the restrictions we impose. However, their likelihood calculation does not yield estimates of the shocks that generated the data.

We impose two restrictions on the standard Kalman filtering framework. First, we do not allow for measurement error in the observation equation. Second, we assume that either (a) the system is initially in steady state at the start of the sample or (b) the researcher knows the initial state with certainty and knows the transition path of the model back to steady state. If the system is stable, ie all the eigenvalues of the state transition matrix are inside the unit circle, the second restriction is not costly in situations where the researcher has a sufficient burn-in period at the start of the sample so that the effect of the initial state dissipates before the sample of interest begins. The meaning of a "sufficient" burn-in period depends on how persistent the effects of the initial state are.

Typically, researchers use the smoothed estimate of the shocks meaning the estimate of the shock at date t that conditions account of all data observed including data after date t . Under the restrictions above, the Kalman filter and the Kalman smoother are equivalent.

Lastly, we describe how to create counterfactuals using alternative shocks and/or alternative impulse response functions.

F.1 Filtering Algorithm

We will present the method starting from a state space representation so that we can compare our method to the Kalman filter.

Consider a dynamic system with a state space representation

$$X_t = AX_{t-1} + B\epsilon_t \quad (18)$$

$$Y_t = CX_t \quad (19)$$

where X is the state, ϵ is an i.i.d. mean-zero innovations and Y is the observed data. ϵ and Y are dimension $N \times 1$ and X is dimension $M \times 1$. A , B , and C are conformable matrices.

We assume that this internal description of the model is unknown to the researcher. Instead, the researcher has access to an external description of the system, i.e. impulse response functions. Let $R(\tau, i)$ be the response of Y_τ to a unit change in the i th element of ϵ_0 . The impulse responses are given by

$$R(\tau, i) = CA^\tau B\mathbf{1}_i,$$

where $\mathbf{1}_i$ is the standard basis vector in the i th dimension. $R(\tau, i)$ is a $N \times 1$ vector. Let $R(\tau)$ be a $N \times N$ matrix where the i th column is $R(\tau, i)$. Notice that

$$R(\tau) = CA^\tau B.$$

The researcher may also have access to an estimate of the effects of the initial state of the system $S(\tau) = CA^{\tau+1}X_{-1}$ for $\tau \geq 0$. In practice one may wish to assume that the system is initially in steady state so $S(\tau) = 0$ for all τ . For a stationary system, where $A^t \rightarrow 0$ as $t \rightarrow \infty$, the role of the initial state will diminish over time so if one has a sufficient burn-in period of data assuming the system starts in steady state will have limited effect on the results.

The researcher has data $\{Y_t\}_{t=0}^T$ and wishes to recover an estimate of $\{\epsilon_t\}_{t=0}^T$. The filtering then proceeds recursively as follows: Let $Q_0 = S(0)$. At date 0 solve (18) and (19) for

$$\epsilon_0 = (CB)^{-1} (Y_0 - CAX_{-1})$$

and notice that we can rewrite this as

$$\epsilon_0 = R(0)^{-1} (Y_0 - Q_0).$$

Now suppose that we have solved for $\{\epsilon_\tau\}_{\tau=0}^{t-1}$ and we wish to solve for ϵ_t . Let $Q_t = CAX_{t-1}$ and by repeated substitution of (18) we have

$$\begin{aligned} Q_t &= \sum_{\tau=0}^{t-1} CA^{t-\tau} B\epsilon_\tau + S(t) \\ &= \sum_{\tau=0}^{t-1} R(t-\tau)\epsilon_\tau + S(t). \end{aligned}$$

From (18) and (19) we then have

$$\epsilon_t = R(0)^{-1} (Y_t - Q_t). \quad (20)$$

F.2 Relationship to the Kalman Filter

Let $\hat{X}_{t|t-1}$ be the point estimate of X_t given information through $t-1$. The Kalman filter updates this estimate as (Hamilton, 1994, eq. 13.2.15)

$$\hat{X}_{t|t} = \hat{X}_{t|t-1} + P_{t|t-1} C' \left(CP_{t|t-1} C' \right)^{-1} \left(Y_t - C\hat{X}_{t|t-1} \right),$$

where $P_{t|t-1}$ is the covariance matrix of $\hat{X}_{t|t-1}$. Because we assume that the initial state (or rather its effects) is known and there is no measurement error, once Y_{t-1} is observed, ϵ_{t-1} is known and therefore the only reason $\hat{X}_{t|t-1}$ is uncertain is because of ϵ_t . Therefore $P_{t|t-1} = B\Sigma B'$ where Σ is the covariance matrix of ϵ . Plugging this in above we have

$$\begin{aligned} \hat{X}_{t|t} &= \hat{X}_{t|t-1} + B\Sigma B' C' \left(CB\Sigma B' C' \right)^{-1} \left(Y_t - C\hat{X}_{t|t-1} \right) \\ &= \hat{X}_{t|t-1} + \underbrace{B (CB)^{-1}}_{=\epsilon_t} \left(Y_t - C\hat{X}_{t|t-1} \right). \end{aligned}$$

Now notice that the update to $\hat{X}_{t|t-1}$ is just $B\epsilon_t$ so we have

$$\begin{aligned} \epsilon_t &= (CB)^{-1} \left(Y_t - C\hat{X}_{t|t-1} \right) \\ &= (CB)^{-1} \left(Y_t - CA\hat{X}_{t-1|t-1} \right) \end{aligned}$$

where the second line follows from Hamilton eq. 13.2.17. Using the logic above, X_{t-1} is known after Y_{t-1} is observed so $\hat{X}_{t-1|t-1} = X_{t-1}$ so the above equation becomes

$$\epsilon_t = R(0)^{-1} (Y_t - Q_t)$$

in the notation of our filtering algorithm, which is the same as (20).

With regard to the Kalman smoother, the smoothed estimate updates the filtered estimate according to (Hamilton eq. 13.6.16)

$$\hat{X}_{t|T} = \hat{X}_{t|t} + J_t \left(\hat{X}_{t+1|T} - \hat{X}_{t+1|t} \right)$$

where (Hamilton eq. 13.6.11)

$$J_t = P_{t|t} A' P_{t+1|t}^{-1}.$$

In our case there is no uncertainty over X_t conditional on information through date t so $P_{t|t}$ is a zero matrix. Therefore the smoother does not update the filtered estimates of the states and therefore does not update the estimate of the shocks.

F.3 Counterfactuals

Now suppose the researcher is interested in some counterfactual outcome. These counterfactuals could come from an alternative history of shocks or an alternative model represented by different impulse response functions. Let $\{\tilde{Y}_t\}$ denote the counterfactual outcomes, $\tilde{R}(\tau)$ denote the counterfactual impulse response functions, and $\{\tilde{\epsilon}_t\}$ denote the counterfactual shocks. We then have

$$\tilde{Y}_t = \sum_{\tau=0}^t \tilde{R}(t-\tau) \tilde{\epsilon}_\tau + S(t).$$

F.4 Incorporating the ELB

Suppose that the k^{th} element of Y is bounded below by \underline{r} . We define $\mathbb{E}_t U_{t+h}^j$ as the forecast for Y that including all available information up to time t from the j^{th} iteration of this algorithm. For $j = 0$ this is equal to,

$$\mathbb{E}_t U_{t+h}^0 = \sum_{\tau=0}^t R(t+h-\tau) \epsilon_\tau^0 + S(t+h).$$

where ϵ_τ^j are the set of shocks inferred in the j^{th} iteration, with $\epsilon_\tau^0 = R(0)^{-1} (Y_t - Q_t)$ as above.

We use this forecast to calculate the amount by which the ELB is expected to be violated by in the future,

$$\mathbb{E}_t \epsilon_{t+h}^{r,news} = \max\{\mathbb{E}_t U_{t+h}^j \mathbf{1}_k - \underline{r}, 0\}.$$

Note that $\epsilon_t^{r,news} = 0$ since $U_t = Y_t$ and Y_t satisfies the ELB. Denote these elements as a vector $\vec{\epsilon}_t^{r,news} = \mathbb{E}_t(0, \epsilon_{t+1}^{r,news}, \epsilon_{t+2}^{r,news}, \dots)$.

To improve the performance of this algorithm, we incorporate a fraction $\lambda \in (0, 1]$ of the discrepancy with the ELB into the news shocks,

$$\vec{r}_t^{j,news} = \vec{r}_t^{j-1,news} + \lambda \vec{\epsilon}_t^{r,news}$$

with $r_t^{0,news} = \vec{0}$.

Let $R^r(\tau, h, i)$ store the impulse response functions for these news shocks on outcome i . Specifically, column $h + 1$ is the IRF of outcome i to a news shock at horizon h . Then our forecast for period t is,

$$V_t^j = \sum_{\tau=0}^t R(t-\tau) \epsilon_\tau + \vec{r}_t^{j,news'} R^r(t) + S(t).$$

We must now update our inference of the shocks ϵ_τ , since $V_t \neq Y_t$,

$$\epsilon_\tau^j = \epsilon_\tau^{j-1} - R(0)^{-1} (U_t - V_t)$$

We use this new set of shocks to update our forecast,

$$\mathbb{E}_t U_{t+h}^j = \sum_{\tau=0}^t R(t+h-\tau) \epsilon_\tau^j + \vec{r}_t^{j,news'} R^r(t+h) + S(t+h).$$

We can then again check if these updated forecasts satisfy the ELB. If they do not, we keep iterating on the algorithm. If they do, we proceed to the next period.

G Derivation and Decomposition of r^*

G.1 Derivation of Equation (15)

Consider an abstract representation of our model expressed in discrete time steps corresponding to the time intervals on which we compute the model:

$$\begin{aligned}\hat{Y}_t &= \mathcal{Y}(h_t, \Phi_t) \\ h_t &= \mathcal{H}(\vec{r}_t, \eta_t) \\ \Phi_{t+1} &= \mathcal{T}(\Phi_t, h_t) \\ \eta_{t+1} &= \mathcal{F}(\eta_t, \epsilon_{t+1}^\eta).\end{aligned}$$

The first equation states that the output gap, \hat{Y}_t , is a function, \mathcal{Y} , of the household policy rules, h_t , and the distribution of households over individual states, Φ_t . In our model, a household chooses whether or not to adjust its durable stock and if so the level of durables, how much to consume in non-durables, and how much to save in liquid assets. All of these decision rules are contained in the collection h_t . The second equation states that the policy rules depend the vector of current and expected future real interest rates, $\vec{r}_t \equiv (r_t, \mathbb{E}_t r_{t+1}, \dots)'$, and the exogenous aggregate states, $\eta_t \equiv (Z_t, G_t, r_t^b)'$.³⁷ In Appendix G.3 we extend the analysis to allow for prices other than real interest rates to affect the decision rules, but we begin with a simpler formulation here for ease of exposition. The third equation shows how the distribution of individual states evolves as a function of the household decisions. In heterogeneous agent models, the evolution of the distribution depends on individual decisions as well as the stochastic process of idiosyncratic shocks. In our formulation, the effect of idiosyncratic shocks is embedded within the function \mathcal{T} . Finally, the fourth equation gives the law of motion for the exogenous aggregate states.

Current and future real rates affect the policy rules at t . Previous real interest rates do not affect the policy rules because the policy rules are conditional on individual states. However, past interest rates affect the output gap at t through their effect on the distribution of individual states

³⁷To see how such an equation is implied by the model, consider the Bellman equation associated with a standard consumption-savings problem:

$$V_t(x, z) = \max_{a'} \left\{ u(x - a') + \beta \mathbb{E}_{z'} V_{t+1}((1 + r_t)a + wz', z') \right\},$$

where x represents cash on hand. The value function V_t depends on the interest rate r_t and the value function V_{t+1} . Recursively, V_{t+1} depends on all future interest rates. Therefore the policy rules at t depend on all future interest rates.

Φ_t . For example, if low interest rates in the past caused households to stock up on durables, then this is reflected in the distribution of households over levels of durables.

We linearize the system around steady state:

$$\begin{aligned}\hat{Y}_t &= \mathcal{Y}_h h_t + \mathcal{Y}_\Phi (\Phi_t - \bar{\Phi}) \\ h_t &= \mathcal{H}_r \vec{r}_t + \mathcal{H}_\eta \eta_t \\ \Phi_{t+1} - \bar{\Phi} &= \mathcal{T}_\Phi (\Phi_t - \bar{\Phi}) + \mathcal{T}_h h_t \\ \eta_{t+1} &= \mathcal{F}_\eta \eta_t + \mathcal{F}_\epsilon \epsilon_{t+1}^\eta.\end{aligned}$$

As with Φ_t , h_t can be interpreted as a vector that gives a discrete representation of the decision rules as in the Reiter (2009) method. Using the linearized system, the forecast at date t of the output gap at date $t + s$ for $s \geq 0$ is given by

$$\mathbb{E}_t \hat{Y}_{t+s} = \mathcal{Y}_h \left(\mathcal{H}_r \mathbb{E}_t \vec{r}_{t+s} + \mathcal{H}_\eta \mathcal{F}_\eta^s \eta_t \right) + \sum_{k=0}^{s-1} \mathcal{Y}_\Phi \mathcal{T}_\Phi^{s-k-1} \mathcal{T}_h \left(\mathcal{H}_r \mathbb{E}_t \vec{r}_{t+k} + \mathcal{H}_\eta \mathcal{F}_\eta^k \eta_t \right) + \mathcal{Y}_\Phi \mathcal{T}_\Phi^s (\Phi_t - \bar{\Phi}),$$

where \mathcal{Y}_h is the partial Jacobian of \hat{Y}_t with respect to h_t and so on. As $\vec{r}_t \equiv (r_t, \mathbb{E}_t r_{t+1}, \mathbb{E}_t r_{t+2}, \dots)$ we can write $\mathbb{E}_t \vec{r}_{t+s} = S_s \vec{r}_t$ where S_s is a shift operator that chops off the first s elements of \vec{r}_t . Using this shift operator and rearranging yields

$$\mathbb{E}_t \hat{Y}_{t+s} = \left[\mathcal{Y}_h \mathcal{H}_r S_s + \sum_{k=0}^{s-1} \mathcal{Y}_\Phi \mathcal{T}_\Phi^{s-k-1} \mathcal{T}_h \mathcal{H}_r S_k \right] \vec{r}_t + \left[\mathcal{Y}_h \mathcal{H}_\eta \mathcal{F}_\eta^s + \sum_{k=0}^{s-1} \mathcal{Y}_\Phi \mathcal{T}_\Phi^{s-k-1} \mathcal{T}_h \mathcal{H}_\eta \mathcal{F}_\eta^k \right] \eta_t + [\mathcal{Y}_\Phi \mathcal{T}_\Phi^s] (\Phi_t - \bar{\Phi}). \quad (21)$$

This equation shows that the forecast of the output gap at $t + s$ is (to a first order approximation) a linear function of the expected real interest rate path, the exogenous states η_t , and the distribution Φ_t . Stacking equation (21) for $s \geq 0$ then yields equation (15) with the terms in square brackets forming the rows of \mathcal{M} , \mathcal{Q} , and \mathcal{D} , respectively.

G.2 Derivation of Equation (17)

As shown in the text, the solution for r^* for a given set of states η_r, Φ_t is,

$$\vec{r}_t^* = -\mathcal{M}^{-1} (\mathcal{Q} \eta_t + \mathcal{D} (\Phi_t - \bar{\Phi})).$$

To determine r^* as a function of past real rates and the exogenous states η , we solve out for

the endogenous state Φ_t . Solving the state backwards yields,

$$\begin{aligned}
\Phi_t - \bar{\Phi} &= \mathcal{T}_h h_{t-1} + \mathcal{T}_\Phi (\Phi_{t-1} - \bar{\Phi}) \\
&= \mathcal{T}_h \mathcal{H}_r \vec{r}_{t-1} + \mathcal{T}_h \mathcal{H}_\eta \eta_{t-1} + \mathcal{T}_\Phi (\mathcal{T}_h h_{t-2} + \mathcal{T}_\Phi (\Phi_{t-2} - \bar{\Phi})) \\
&= \sum_{k=0}^{t-1} \left(\mathcal{T}_\Phi^k \mathcal{T}_h \mathcal{H}_r \right) \vec{r}_{t-1-k} + \sum_{k=0}^{t-1} \left(\mathcal{T}_\Phi^k \mathcal{T}_h \mathcal{H}_\eta \right) \eta_{t-1-k}
\end{aligned}$$

with $\Phi_0 = \bar{\Phi}$.

We next show how to express the first two terms in terms of the matrices \mathcal{M} and \mathcal{Q} . Start with the term that captures the history of interest rates,

$$\begin{aligned}
\mathcal{D} \left[\sum_{k=0}^{t-1} \left(\mathcal{T}_\Phi^k \mathcal{T}_h \mathcal{H}_r \right) \vec{r}_{t-1-k} \right] &= \mathcal{Y}_\Phi \begin{pmatrix} I \\ \mathcal{T}_\Phi \\ \mathcal{T}_\Phi^2 \\ \vdots \end{pmatrix} \left[\sum_{k=0}^{t-1} \left(\mathcal{T}_\Phi^k \mathcal{T}_h \mathcal{H}_r \right) \vec{r}_{t-1-k} \right] \\
&= \mathcal{Y}_\Phi \sum_{k=0}^{t-1} \begin{pmatrix} (\mathcal{T}_\Phi^k) \\ (\mathcal{T}_\Phi^{k+1}) \\ (\mathcal{T}_\Phi^{k+2}) \\ \vdots \end{pmatrix} \mathcal{T}_h \mathcal{H}_r \vec{r}_{t-1-k}
\end{aligned}$$

To see the connection with the monetary transmission matrix, we split \mathcal{M} into two components, one capturing how the evolution of the state and the other the policy function,

$$\mathcal{M} = \begin{pmatrix} 0 \\ \mathcal{Y}_\Phi \mathcal{T}_h \mathcal{H}_r \\ \mathcal{Y}_\Phi \mathcal{T}_\Phi \mathcal{T}_h \mathcal{H}_r + [0 \quad \mathcal{Y}_\Phi \mathcal{T}_h \mathcal{H}_r] \\ \mathcal{Y}_\Phi \mathcal{T}_\Phi^2 \mathcal{T}_h \mathcal{H}_r + [0 \quad \mathcal{Y}_\Phi \mathcal{T}_\Phi \mathcal{T}_h \mathcal{H}_r] + [0 \quad 0 \quad \mathcal{Y}_\Phi \mathcal{T}_h \mathcal{H}_r] \\ \vdots \end{pmatrix} + \begin{pmatrix} \mathcal{Y}_h \mathcal{H}_r \\ [0 \quad \mathcal{Y}_h \mathcal{H}_r] \\ [0 \quad 0 \quad \mathcal{Y}_h \mathcal{H}_r] \\ \vdots \end{pmatrix}$$

For general $s = t + 1$ the term of past real rate expectations can then be split into a component involving interest rate innovations up to time $s - 1$ and one component involving expected interest

rates from s onward,

$$\begin{aligned} \mathcal{D} \left[\sum_{k=0}^s (\mathcal{T}_\Phi^k \mathcal{T}_h \mathcal{H}_r) \vec{r}_{t-1-k} \right] &= \sum_{k=0}^{s-1} \mathcal{M}_{[1+s-k, \dots, 1+s-k]} \left[\mathbb{E}_k \begin{pmatrix} r_k \\ \vdots \\ r_{s-1} \end{pmatrix} - \mathbb{E}_{k-1} \begin{pmatrix} r_k \\ \vdots \\ r_{s-1} \end{pmatrix} \right] \\ &\quad + \sum_{k=0}^{s-1} (\mathcal{M}_{[1+s-k, \dots, 1+s-k]} - \mathcal{M}_{[s-k, \dots, s-k]}) \mathbb{E}_k \vec{r}_s \end{aligned}$$

We take a similar approach for expressing the historical contribution of the exogenous states η . It will again be convenient to write the matrix \mathcal{Q} as the sum of the state component and the policy component,

$$\mathcal{Q} = \begin{pmatrix} 0 \\ \mathcal{Y}_\Phi \mathcal{T}_h \mathcal{H}_\eta \\ \mathcal{Y}_\Phi \mathcal{T}_\Phi \mathcal{T}_h \mathcal{H}_\eta + \mathcal{Y}_\Phi \mathcal{T}_h \mathcal{H}_\eta \mathcal{F}_\eta \\ \mathcal{Y}_\Phi \mathcal{T}_\Phi^2 \mathcal{T}_h \mathcal{H}_\eta + \mathcal{Y}_\Phi \mathcal{T}_\Phi \mathcal{T}_h \mathcal{H}_\eta \mathcal{F}_\eta + \mathcal{Y}_\Phi \mathcal{T}_h \mathcal{H}_\eta \mathcal{F}_\eta^2 \\ \vdots \end{pmatrix} + \begin{pmatrix} \mathcal{Y}_h \mathcal{H}_\eta \\ \mathcal{Y}_h \mathcal{H}_\eta \mathcal{F}_\eta \\ \mathcal{Y}_h \mathcal{H}_\eta \mathcal{F}_\eta^2 \\ \vdots \end{pmatrix}$$

Then we can express the historical contribution of the exogenous states η to the states solely

in terms of past shocks and the \mathcal{Q} matrix,

$$\begin{aligned}
\mathcal{D} \left[\sum_{k=0}^{t-1} \left(\mathcal{T}_{\Phi}^k \mathcal{T}_h \mathcal{H}_{\eta} \right) \eta_{t-1-k} \right] &= \mathcal{Y}_{\Phi} \begin{pmatrix} I \\ \mathcal{T}_{\Phi} \\ \mathcal{T}_{\Phi}^2 \\ \vdots \end{pmatrix} \left[\sum_{k=0}^{t-1} \left(\mathcal{T}_{\Phi}^k \mathcal{T}_h \mathcal{H}_{\eta} \right) \eta_{t-1-k} \right] \\
&= \mathcal{Y}_{\Phi} \sum_{k=0}^{t-1} \begin{pmatrix} (\mathcal{T}_{\Phi}^k) \\ (\mathcal{T}_{\Phi}^{k+1}) \\ (\mathcal{T}_{\Phi}^{k+2}) \\ \vdots \end{pmatrix} \mathcal{T}_h \mathcal{H}_{\eta} \eta_{t-1-k} \\
&= \sum_{k=0}^{t-1} (\mathcal{Q}_{[2+k, \dots]} - \mathcal{Q}_{[1+k, \dots]} \mathcal{F}_{\eta}) \eta_{t-1-k} \\
&= \sum_{k=0}^{t-1} \mathcal{Q}_{[2+k, \dots]} \eta_{t-1-k} - \sum_{k=0}^{t-1} \mathcal{Q}_{[1+k, \dots]} \mathcal{F}_{\eta} \eta_{t-1-k} \\
&= -\mathcal{Q} \mathcal{F}_{\eta} \eta_{t-1} + \sum_{k=0}^{t-1} \mathcal{Q}_{[2+k, \dots]} \epsilon_{t-1-k}^{\eta}
\end{aligned}$$

Substituting our solution for the state into the equation for r^* yields,

$$\begin{aligned}
\bar{r}_t^* &= -\mathcal{M}^{-1} (\mathcal{D}(\Phi_t - \bar{\Phi}) + \mathcal{Q}\eta_t) \\
&= -\mathcal{M}^{-1} \sum_{k=0}^{t-1} \mathcal{M}_{[1+t-k, \dots, 1..t-k]} \left[\mathbb{E}_k \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} - \mathbb{E}_{k-1} \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} \right] \\
&\quad - \mathcal{M}^{-1} \sum_{k=0}^{t-1} (\mathcal{M}_{[1+t-k, \dots, 1+t-k..]} - \mathcal{M}_{[t-k, \dots, t-k..]}) \mathbb{E}_k \bar{r}_t \\
&\quad - \mathcal{M}^{-1} \sum_{k=0}^{t-1} \mathcal{Q}_{[1+k, \dots]} \epsilon_{t-k}^{\eta}
\end{aligned}$$

This equation tells us that r^* is not just a function of the shocks (last line), but it can also vary with how past interest rates were set in the past (second line) and with past expectations of current and future rates (third line).

We next solve out for these expectations of current and future rates by assuming that they are set to close all output gaps from time t onwards, consistent with the definition of r^* . Thus, the

expectations of future rates are now superscripted with a star,

$$\begin{aligned}\bar{r}_t^* &= -\mathcal{M}^{-1} \sum_{k=0}^{t-1} \mathcal{M}_{[1+t-k, \dots, 1..t-k]} \left[\mathbb{E}_k \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} - \mathbb{E}_{k-1} \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} \right] \\ &\quad - \mathcal{M}^{-1} \sum_{k=0}^{t-1} (\mathcal{M}_{[1+t-k, \dots, 1+t-k..]} - \mathcal{M}_{[t-k, \dots, t-k..]}) \mathbb{E}_k \bar{r}_t^* \\ &\quad - \mathcal{M}^{-1} \sum_{k=0}^{t-1} \mathcal{Q}_{[1+k, \dots]} \epsilon_{t-k}^\eta\end{aligned}$$

Taking expectations of this r^* vector yields

$$\begin{aligned}\mathcal{M}_{[1+s, \dots, 1+s..]} \mathbb{E}_{t-s} \bar{r}_t^* &= - \sum_{k=0}^{t-s} \mathcal{M}_{[1+t-k, \dots, 1..t-k]} \left[\mathbb{E}_k \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} - \mathbb{E}_{k-1} \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} \right] \\ &\quad - \sum_{k=0}^{t-s-1} (\mathcal{M}_{[1+t-k, \dots, 1+t-k..]} - \mathcal{M}_{[t-k, \dots, t-k..]}) \mathbb{E}_k \bar{r}_t^* \\ &\quad - \sum_{k=s}^t \mathcal{Q}_{[1+k, \dots]} \epsilon_{t-k}^\eta\end{aligned}$$

For substituting out past expectations of current and future real rates, it is particularly convenient to express the updating of expectations recursively,

$$\begin{aligned}\mathbb{E}_{t-s} \bar{r}_t^* &= \mathbb{E}_{t-s-1} \bar{r}_t^* - \mathcal{M}_{[1+s, \dots, 1+s..]}^{-1} \mathcal{M}_{[1+s, \dots, 1..s]} \left[\mathbb{E}_{t-s} \begin{pmatrix} r_{t-s} \\ \vdots \\ r_{t-1} \end{pmatrix} - \mathbb{E}_{t-s-1} \begin{pmatrix} r_{t-s} \\ \vdots \\ r_{t-1} \end{pmatrix} \right] \\ &\quad - \mathcal{M}_{[1+s, \dots, 1+s..]}^{-1} \mathcal{Q}_{[1+s, \dots]} \epsilon_{t-s}^\eta\end{aligned}$$

Repeated substitution of the expectation updating into the r^* equation then yields the formula

in the text,

$$\begin{aligned} \bar{r}_t^* &= - \sum_{k=0}^{t-1} \mathcal{M}_{[1+t-k..,1+t-k..]}^{-1} \mathcal{M}_{[1+t-k..,1..t-k]} \left[\mathbb{E}_k \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} - \mathbb{E}_{k-1} \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} \right] \\ &\quad - \sum_{k=0}^t \mathcal{M}_{[1+k..,1+k..]}^{-1} \mathcal{Q}_{[1+k,..]} \epsilon_{t-k}^\eta \end{aligned}$$

Full accommodation uses r^{*FA} in the history, actual accommodation uses the realized r .

Under full accommodation we can further simplify to arrive at a recursive expression,

$$\begin{aligned} \bar{r}_t^{*FA} &= \mathbb{E}_{t-1} \bar{r}_t^{*FA} - \mathcal{M}^{-1} \mathcal{Q} \epsilon_t^\eta \\ &= S_1 \bar{r}_{t-1}^{*FA} - \mathcal{M}^{-1} \mathcal{Q} \epsilon_t^\eta \end{aligned}$$

G.3 Extension with More Endogenous Prices

Consider the expanded system:

$$\begin{aligned} \hat{Y}_t &= \mathcal{Y}(h_t, \Phi_t) \\ h_t &= \mathcal{H}(\vec{r}_t, \vec{w}_t, \eta_t) \\ \Phi_{t+1} &= \mathcal{T}(\Phi_t, h_t) \\ \eta_{t+1} &= \mathcal{F}(\eta_t, \epsilon_{t+1}^\eta) \\ 0 &= \mathcal{P}(h_t, \Phi_t), \end{aligned}$$

where w_t is a vector of prices (other than real interest rates) at date t and $\vec{w}_t \equiv (w_t, \mathbb{E}_t w_{t+1}, \dots)'$.

The second equation therefore allows for other prices besides interest rates to affect household policy rules. $\mathcal{P}(h_{t+s}, \Phi_{t+s}) = 0$ gives the market clearing conditions for the prices in w_{t+s} . If w_{t+s} is a vector of prices, then \mathcal{P} is a vector-valued function. The prices in w_t can include tax rates and the \mathcal{P} can include government budget constraints or fiscal rules that set the tax rate.

Now let's take \vec{r}_t as given and solve for the resulting \vec{w}_t . Proceeding as with the forecast of the

output gap we have (to a first order approximation) the market clearing conditions at $t + s$ are

$$\begin{aligned}
0 = & \left[\mathcal{P}_\Phi \sum_{k=0}^{s-1} \left(\mathcal{T}_\Phi^k \mathcal{T}_h \mathcal{H}_r S_{s-k-1} \right) + \mathcal{P}_h \mathcal{H}_r S_s \right] \vec{r}_t \\
& + \left[\mathcal{P}_\Phi \sum_{k=0}^{s-1} \left(\mathcal{T}_\Phi^k \mathcal{T}_h \mathcal{H}_w S_{s-k-1} \right) + \mathcal{P}_h \mathcal{H}_w S_s \right] \vec{w}_t \\
& + \left[\mathcal{P}_\Phi \sum_{k=0}^{s-1} \left(\mathcal{T}_\Phi^k \mathcal{T}_h \mathcal{H}_\eta \rho_\eta^{s-k-1} \right) + \mathcal{P}_h \mathcal{H}_\eta \rho_\eta^s \right] \eta_t \\
& + \mathcal{P}_\Phi \mathcal{T}_\Phi^s (\Phi_t - \bar{\Phi})
\end{aligned}$$

Stacking this equation for $s \geq 0$ yields

$$\vec{0} = \mathcal{M}_P \vec{r}_t + \mathcal{N}_P \vec{w}_t + \mathcal{Q}_P \eta_t + \mathcal{D}_P (\Phi_t - \bar{\Phi})$$

Solve this for \vec{w}_t

$$\vec{w}_t = -\mathcal{N}_P^{-1} [\mathcal{M}_P \vec{r}_t + \mathcal{Q}_P \eta_t + \mathcal{D}_P (\Phi_t - \bar{\Phi})] \quad (22)$$

Forecasting the output gap as before:

$$\begin{aligned}
\mathbb{E}_t \hat{Y}_{t+s} &= \mathcal{Y}_\Phi \mathbb{E}_t (\Phi_{t+s} - \bar{\Phi}) + \mathcal{Y}_h \mathbb{E}_t h_{t+s} \\
&= \mathcal{Y}_\Phi \mathbb{E}_t (\Phi_{t+s} - \bar{\Phi}) + \mathcal{Y}_h \mathcal{H}_r \mathbb{E}_t \vec{r}_{t+s} + \mathcal{Y}_h \mathcal{H}_w \mathbb{E}_t \vec{w}_{t+s} + \mathcal{Y}_h \mathcal{H}_\eta \mathbb{E}_t \eta_{t+s} \\
&= [\mathcal{Y}_\Phi - \mathcal{Y}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P] [\mathcal{T}_\Phi \mathbb{E}_t (\Phi_{t+s-1} - \bar{\Phi}) + \mathcal{T}_h \mathbb{E}_t h_{t+s-1}] \\
&\quad + \mathcal{Y}_h [I - \mathcal{Y}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{M}_P] \mathcal{H}_r \mathbb{E}_t \vec{r}_{t+s} + \mathcal{Y}_h [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{Q}_P] \mathcal{H}_\eta \mathbb{E}_t \eta_{t+s} \\
&= [\mathcal{Y}_\Phi - \mathcal{Y}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P] [\mathcal{T}_\Phi \mathbb{E}_t (\Phi_{t+s-1} - \bar{\Phi}) + \mathcal{T}_h \mathcal{H}_r \mathbb{E}_t \vec{r}_{t+s-1} + \mathcal{T}_h \mathcal{H}_w \mathbb{E}_t \vec{w}_{t+s-1} + \mathcal{T}_h \mathcal{H}_\eta \mathbb{E}_t \eta_{t+s-1}] \\
&\quad + \mathcal{Y}_h [I - \mathcal{Y}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{M}_P] \mathcal{H}_r \mathbb{E}_t \vec{r}_{t+s} + \mathcal{Y}_h [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{Q}_P] \mathcal{H}_\eta \mathbb{E}_t \eta_{t+s} \\
&= [\mathcal{Y}_\Phi - \mathcal{Y}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P] \{ [\mathcal{T}_\Phi - \mathcal{T}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P] \mathbb{E}_t (\Phi_{t+s-1} - \bar{\Phi}) \\
&\quad + \mathcal{T}_h [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{M}_P] \mathcal{H}_r \mathbb{E}_t \vec{r}_{t+s-1} + \mathcal{T}_h [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{Q}_P] \mathcal{H}_\eta \mathbb{E}_t \eta_{t+s-1} \} \\
&\quad + \mathcal{Y}_h [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{M}_P] \mathcal{H}_r \mathbb{E}_t \vec{r}_{t+s} + \mathcal{Y}_h [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{Q}_P] \mathcal{H}_\eta \mathbb{E}_t \eta_{t+s} \\
&= \left\{ [\mathcal{Y}_\Phi - \mathcal{Y}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P] \sum_{k=0}^{s-1} \left([\mathcal{T}_\Phi - \mathcal{T}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P]^k \mathcal{T}_h [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{M}_P] \mathcal{H}_r S_{s-k-1} \right) + \mathcal{Y}_h [I - \mathcal{Y}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{M}_P] \mathcal{H}_r S_s \right\} \\
&\quad + \left\{ [\mathcal{Y}_\Phi - \mathcal{Y}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P] \sum_{k=0}^{s-1} \left([\mathcal{T}_\Phi - \mathcal{T}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P]^k \mathcal{T}_h [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{Q}_P] \mathcal{H}_\eta \mathcal{F}_\eta^{s-k-1} \right) + \mathcal{Y}_h [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{Q}_P] \mathcal{H}_\eta \mathcal{F}_\eta^s \right\} \\
&\quad + [\mathcal{Y}_\Phi - \mathcal{Y}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P] [\mathcal{T}_\Phi - \mathcal{T}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P]^s (\Phi_t - \bar{\Phi})
\end{aligned}$$

Redefine the matrices as follows:

$$\begin{aligned}
\mathcal{Y}'_{\Phi} &= \mathcal{Y}_{\Phi} - \mathcal{Y}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P \\
\mathcal{T}'_{\Phi} &= \mathcal{T}_{\Phi} - \mathcal{T}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P \\
\mathcal{H}'_r &= [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{Q}_P] \mathcal{H}_r \\
\mathcal{H}'_{\eta} &= [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{Q}_P] \mathcal{H}_{\eta}
\end{aligned}$$

Then we arrive at an analogous expression to our simpler case without endogenous prices.

$$\begin{aligned}
\mathbb{E}_t \hat{Y}_{t+s} &= \left[\mathcal{Y}'_{\Phi} \sum_{k=0}^{s-1} \left(\mathcal{T}'_{\Phi}{}^k \mathcal{T}_h \mathcal{H}'_r S_{s-k-1} \right) + \mathcal{Y}_h \mathcal{H}'_r S_s \right] \vec{r}_t \\
&+ \left[\mathcal{Y}'_{\Phi} \sum_{k=0}^{s-1} \left(\mathcal{T}'_{\Phi}{}^k \mathcal{T}_h \mathcal{H}'_{\eta} \mathcal{F}_{\eta}^{s-k-1} \right) + \mathcal{Y}_h \mathcal{H}'_{\eta} \mathcal{F}_{\eta}^s \right] \eta_t \\
&+ \mathcal{Y}'_{\Phi} \mathcal{T}'_{\Phi}{}^s (\Phi_t - \bar{\Phi})
\end{aligned}$$

and stacking these equations for $s \geq 0$ yields

$$\vec{\hat{Y}}_t = \mathcal{M} \vec{r}_t + \mathcal{Q} \eta_t + \mathcal{D} (\Phi_t - \bar{\Phi})$$

For the decomposition we use a similar approach of substituting out for \vec{w}_t . We can build the decomposition iteratively

$$\Phi_t = \sum_{k=0}^{t-1} \mathcal{T}'_{\Phi}{}^k \mathcal{T}_h \mathcal{H}'_r \vec{r}_{t-1-k} + \sum_{k=0}^{t-1} \mathcal{T}'_{\Phi}{}^k \mathcal{T}_h \mathcal{H}'_{\eta} \eta_{t-1-k}$$

Since the addition of endogenous prices leads to identical expressions up to a redefinition of the matrices, the same steps as in the simpler case can be followed.

H Model Impulse Response Functions

Figure A.7 plots the model impulse response functions.

I Robustness to the Scale of \mathcal{M}

Equation (17) shows that for a given history of news about real interest rates, a re-scaling of the monetary news matrix \mathcal{M} does not impact the importance of intertemporal shifting. The scale of

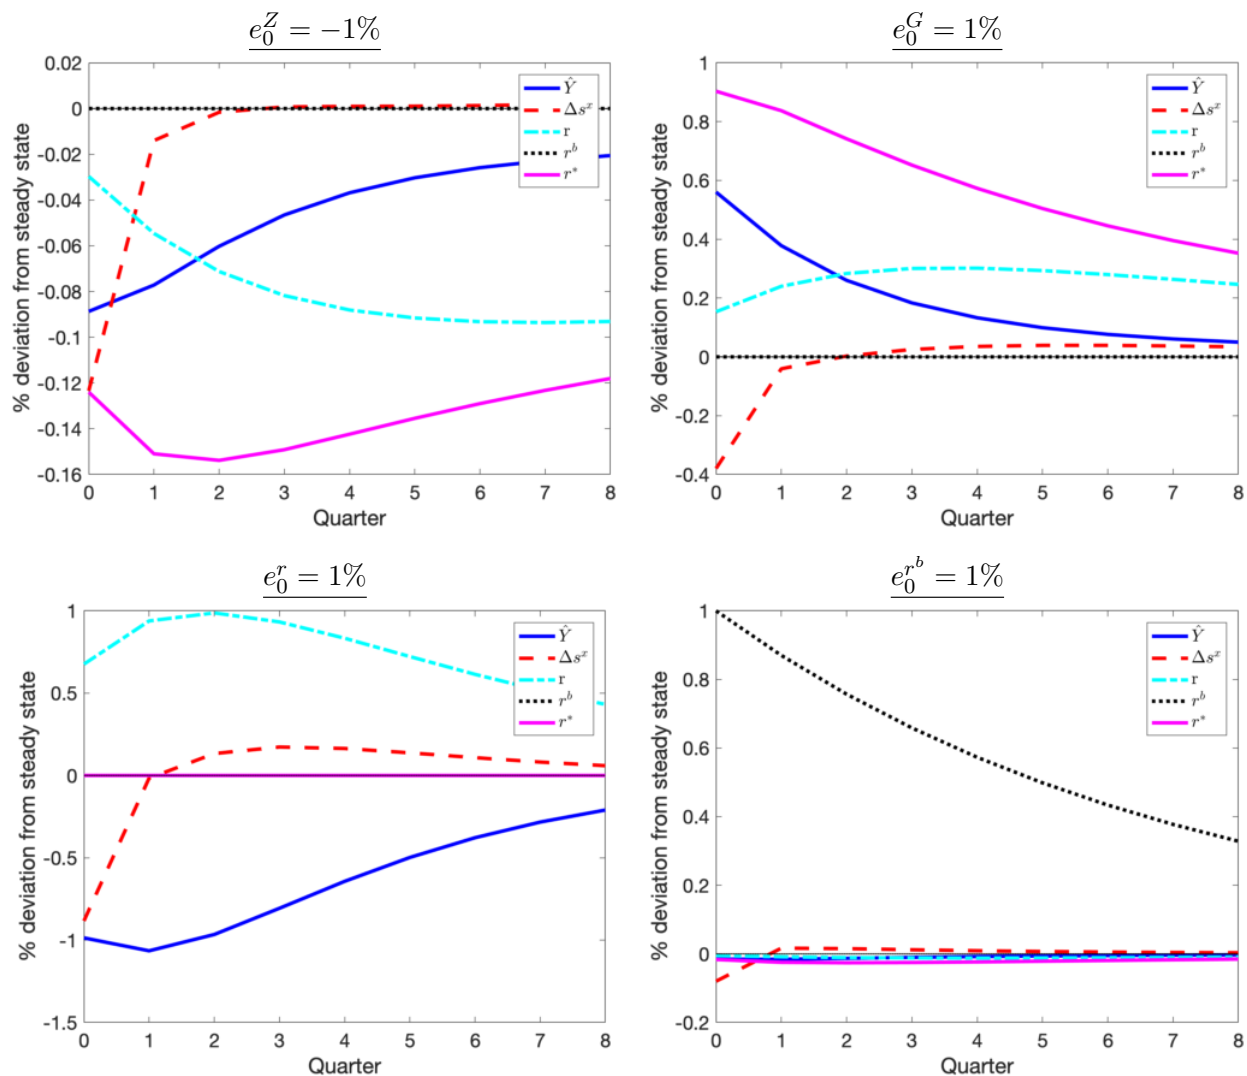


Figure A.7: Impulse response functions for the output gap \hat{Y} , the change in the durable expenditure share relative to potential GDP Δs^x , the real interest rate r , the borrowing spread r^b , and the contemporaneous natural rate of interest r^* following a shock to productivity e^Z , non-household demand e^G , the monetary policy rule e^r , and the borrowing spread e^{r^b} .

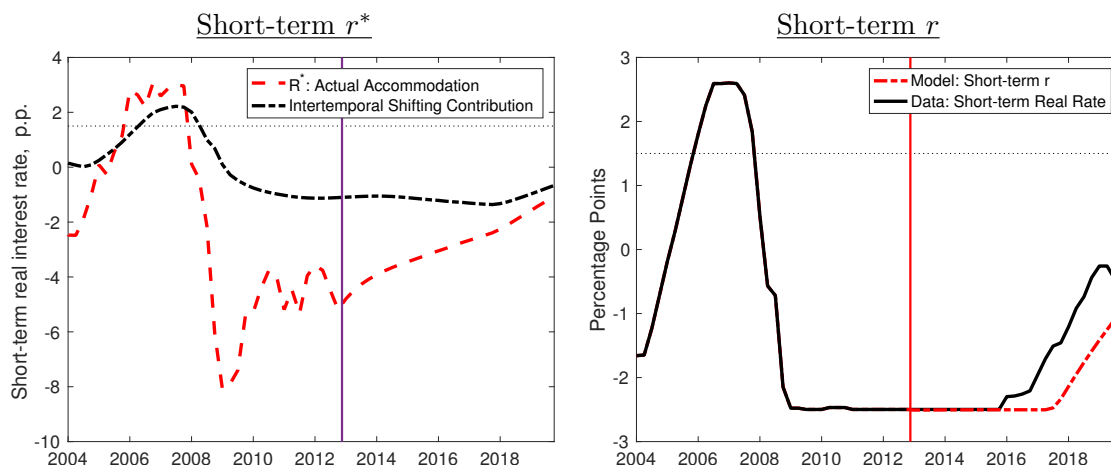


Figure A.8: The left panel shows the time series of the short-term r^* and the contribution of intertemporal shifting, when monetary policy is only half as effective at stimulating aggregate demand than in our baseline model. The right panel plots the short-term real rate in the model against the data. In both panels we stop incorporating new shocks after 2012Q4, so the figures plot the expected path of the natural rate from 2013Q1 onward.

\mathcal{M} , however, affects the shocks that we filter from the data. As a result, our analysis is robust, but not fully insensitive, to the scale of \mathcal{M} .

This robustness is reassuring in light of the possibility of state dependence in the power of monetary policy. Berger and Vavra (2015) show that durable demand is less sensitive to stimulus in a recession. Because we linearize with respect to the aggregate state variables, our analysis does not incorporate state dependence.

To evaluate the potential consequences of state dependence for our analysis, we examine a case where monetary policy is only half as effective at closing the output gap as we assume in our benchmark calculation. Formally, we scale down \mathcal{M} by a factor of $1/2$.

The left panel of Figure A.8 shows the r^* path similar to that in our baseline although the decline in r^* is larger and more persistent. The intertemporal shifting contribution is even more similar to our baseline model. The right panel shows a slower lift-off from the zero lower bound. r^* falls more when monetary policy is less powerful because it takes a larger movement in interest rates to close an output gap of a given size. However, the difference in the results is not very large because our filtering algorithm infers that the shocks hitting the economy were smaller as the observed movement in real interest rates is inferred to be driven by smaller shocks.