STAFF REPORT
No. 622

Lumpy Durable Consumption Demand and the Limited Ammunition of Monetary Policy

February 2021

Alisdair McKay
Federal Reserve Bank of Minneapolis

Johannes F. Wieland
*University of California, San Diego
and NBER*

The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

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

February 12, 2021

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 purchases of durable goods. This yields an intertemporal trade-off in aggregate demand as encouraging households to increase durable holdings 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

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., Barsky et al., 2003; Erceg and Levin, 2006; 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

---

[1]See 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 [...]."

stimulus.

The dynamics of demand for durable goods create a propagation mechanism that makes changes in real interest rates very persistent. 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 effective lower bound (ELB) in 2015Q4 and predicts low levels of 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 to 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, an increasing convenience yield for safe assets, and the rise in income inequality (Eggertsson et al., 2019; Del Negro et al., 2017; Auclert and Rognlie, 2018; Straub, 2018). Our results 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. However, our perspective is fully compatible with the view that secular phenomena have played a role in the decline in interest rates over a longer time horizon.

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, a common formulation in which adjustment costs are increasing in the rate of change of investment (e.g. Christiano et al., 2005). With higher-order adjustment costs, 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 durable stock, which reduces marginal benefit from investing in the future. 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 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 estimates for the responses of GDP, aggregate durable expenditure, and the extensive margins of car and housing adjustments all show reversals consistent with intertemporal shifting. While the extent of these reversals is not precisely estimated, in most cases the point estimates show complete reversals with the cumulative change in activity eventually returning to zero. 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. 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 extended with a relatively standard (S,s) model of investment demand. We show that including two particular ingredients in our 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, shocks to the quality of the match between a household and its durable stock introduce inframarginal adjustments, which reduce 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. (2020). This friction helps the model match the delayed responses to monetary policy shocks that are often observed in aggregate data, as shown by Auclert, Rognlie, and Straub (2020).

Our work builds on a growing literature that models aggregate demand using rich micro-foundations for household consumption that are disciplined by micro-data.[2] Most of this literature focuses its attention on the determination of nondurable 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 nondurable consumption. 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, Bardóczy, Rognlie, and Straub (2019). 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.

Intertemporal shifting effects of monetary policy have appeared in the literature previously. Leamer (2007, 2009) informally argues that monetary policymakers should take account of intertemporal shifting effects in housing. Kreamer (2019) shows that minimizing inefficient sectoral fluctuations in a stylized model of durable demand requires monetary policy to take into account intertemporal shifting effects. Our contribution is to provide evidence of intertemporal shifting effects, quantify them in a rich model disciplined by micro

---

[2]See 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.

4

and macro data, and examine its positive implications for the dynamics of real interest rates.

Contemporaneous work by Mian, Straub, and Sufi (2020) argues that intertemporal shifting of demand can also occur following monetary accommodation due to the accumulation of household 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 creates path dependence in the power of monetary policy. An interest rate cut is more powerful in spurring refinancing and stimulating the economy if interest rates have been high in the past and homeowners have high rates on their existing mortgages. While monetary accommodation today makes the economy less sensitive to future stimulus, it persistently raises the level of demand because refinancing redistributes disposable income to high-MPC households on an ongoing basis. Instead, we emphasize that stimulus today leads to weaker demand in the future by shifting the timing of durable expenditure forward.

The paper is organized as follows: Section 2 presents our model of durable demand; 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 intertemporal shifting implications for the monetary transmission mechanism; Section 6 shows that this feature of the monetary transmission mechanism has 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 nondurable goods, c, and a service flow from durable goods, s. Household $i \in [0, 1]$ has preferences given by

$$E_{i0} \int_{t=0}^{\infty} e^{-\rho t} u\left(c_{it}, s_{it}\right) \mathrm{d}t. \tag{1}$$

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

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

where p_t is the relative price of durable goods in terms of nondurable 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

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

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

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

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.

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

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

Household income, y_{it}, is given by $y_{it} = Y_t z_{it}$, 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

$$\mathrm{d}\ln z_{it} = -\rho_z \left(\ln z_{it} - \ln \bar{z}\right) \mathrm{d}t + \sigma_z \,\mathrm{d}\mathcal{W}_{it}^z, \tag{6}$$

where $\mathrm{d}\mathcal{W}_{it}^z$ is a standard Brownian motion, ρ_z 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 $\int z_{it} \,\mathrm{d}i = 1$.

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 it durables and drops to zero with intensity θ. The service flow is $s_{it} = q_{it} d_{it}$. 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 Poisson intensity Ξ. As in Carroll et al. (2020) and Auclert, Rognlie, and Straub (2020), we assume that households always know their idiosyncratic states and current income. They also learn the current real interest rate when they hit the borrowing constraint and they learn the

7

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 nondurable expenditure to monetary shocks.

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 the population of households. The individual state variables are liquid assets a, the durable stock d, and idiosyncratic productivity z. The distribution over these variables is denoted $\Phi_t(a, d, z)$. In steady state, the prices 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 $\bar{\Phi}(a, d, z)$. To compute the steady state of the model we use the continuous-time methods described in Achdou et al. (2017).

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

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

where $prob_{t,t+dt}(a, d, z)$ is the probability that a household with individual state variables (a, d, z) will make an adjustment between t and $t + dt$, and $d_t^*(a, d, z)$ 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 nondurables. Estimates of this elasticity range from substantially below 1 to around 1 (Ogaki and Reinhart, 1998; Davis and

Ortalo-Magné, 2011; Pakoš, 2011; 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. A low EIS allows the model to match the small response of nondurable consumption to monetary policy shocks we measure in the data (see Section 3). The importance of durables, 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 nondurable consumption (nondurable 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.

We set the fixed cost, f, 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 the 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.

Table 1: Calibration of the Model

Parameter	Description	Value	Source
Parameters of the Household's Problem			
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 Section 2.3
$\bar{r}$	Real interest rate	0.015	Annual real Fed. Funds Rate
$\bar{r}^b$	Borrowing spread	0.017	Mortgage-Treas. 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 Section 2.3
ν	Operating cost	0.048	See Section 2.3
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	Rate of information updating	0.667	Coibion and Gorodnichenko (2012)
General Equilibrium Parameters			
G/Y	Steady state govt share	0.2	Convention
ζ	Inverse durable supply elasticity	0.049	See Section 4.3
κ	Phillips curve slope	0.48	See Section 4.3
ρ_r	Real rate persistence	0.60	Estimated over 1991-2007
ρ_η	Real rate shock persistence	1.55	Estimated over 1991-2007
ϕ_π	Real rate response to inflation	0.79	Estimated over 1991-2007
ϕ_y	Real rate response to output gap	0.75	Estimated over 1991-2007
ρ_G	Non-household demand persistence	0.90	See Section 4.3
ρ_g	Productivity growth persistence	0.77	See Section 4.3
ρ_{r^b}	Borrowing spread persistence	0.63	Estimated over 1991-2007

Since our objective is to explain the behavior of real interest rates during and after the Great Recession, we use the years 1991-2007 as a benchmark for interest rates; a period when the economy entered a period of low and stable inflation. The average, ex post, real federal funds rate in terms of nondurables 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.

Based on the estimates of Floden and Lindé (2001), we set $\rho_z = -\log(0.9136)$ and $\sigma_z = 0.2158$. We set the discount rate ρ to match the average liquid financial asset holdings net of mortgage and auto loans to annual GDP ratio over 1970-2019 of 0.87.[3] 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. Total maintenance costs are the sum of intermediate goods and services consumed in the housing output table, the PCE on household maintenance, and the PCE on motor vehicle maintenance and repair. We divide these costs by total durable depreciation to arrive at $\chi = 0.35$.

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 intensity with which agents update their information set to an annual rate of $\Xi = 2/3$, so that the expected time between updates is six quarters. An expected information rigidity of six quarters is in the middle of the range implied by the estimates of Coibion and Gorodnichenko (2012, table 3) who find quarterly updating frequencies between 0.11 and 0.24 among professional forecasters, households and firms.

---

[3]Liquid assets are defined as in McKay et al. (2016) and Guerrieri and Lorenzoni (2017).

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 estimate this parameter using PSID data and the structural estimation method developed by Berger and Vavra (2015) modified to allow 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" $\omega_{it} \equiv d^*_{it} - d_{it}$, as well as the density of durable gaps, where d^*_{it} is the optimal durable stock based on the current state variables. 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^{*model}_{it} and the current durable stock d^{model}_{it} are known. In the data, we only observe current durable holdings d^{data}_{it} directly. We infer data gaps using a set of observables $\mathbf{Z}^{data}_{it}$, and the model-implied relationship between them and the optimal durable stock, $d^{*data}_{it} = F^{model}(\mathbf{Z}^{data}_{it})$, where F^{model} is the model's mapping from the observables to d^*.

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.

Figure 1 plots the model- and data-implied hazards of adjustment conditional on the durable gap at the optimal parameter estimate, $\theta = 0.1575$.[4] 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 80% of the variation in the hazard rate. The adjustment probability for small gaps is substantial in both the model and the data, suggesting that the match-quality shock is quantitatively important.[5]

---

[4]Appendix Figure A.1 plots the density of durable gaps.

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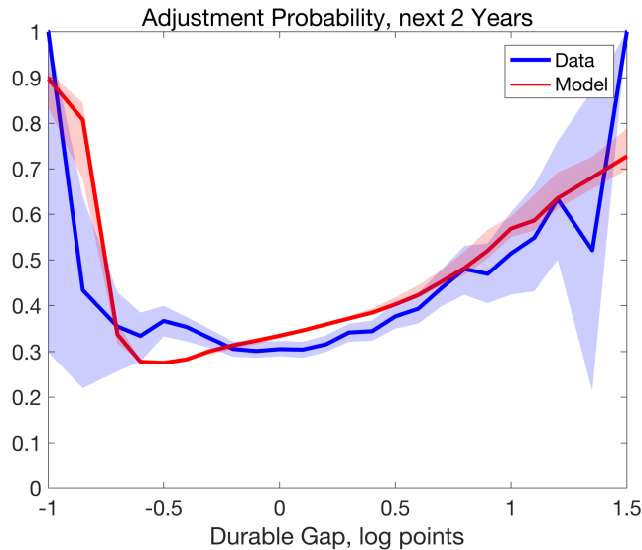
Figure 1: Hazard of adjustment conditional on the durable gap $\omega = d^* - d$, where d^* is the optimal durable choice conditional on adjusting and d the initial durable stock. Shaded areas are 95% confidence bands.

Our estimate implies that 75% of adjustments in steady state are due to the match-quality process.

3 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.

3.1 Evidence from Identified Monetary Shocks

We use monetary shocks from Romer and Romer (2004), extended by Wieland and Yang (2017), for 1969Q1-2006Q4, 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 are informative about the plausibility of intertemporal shifting effects in the monetary transmission mechanism and also serve as empirical benchmarks for the model.

Our outcome variables are log real GDP per capita; the log of total real durable expenditure per capita, $x_t = \ln(X_t)$; and the extensive margin of durable purchases. 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.[6]

We estimate the impulse responses using local projections,

$$z_{t+h} = \alpha_h + \sum_{m=0}^{M} \beta_{h,m}\epsilon_{t-m} + \sum_{l=1}^{L} \gamma_{h,l}z_{t-l} + \delta_h t + \eta_{t,h}, \qquad h = 0, ..., H, \tag{7}$$

where z_t is the outcome variable, ϵ_t is the Romer-Romer monetary shock, and t is a time trend.[7] The impulse response function is the sequence $\{\beta_{h,0}\}_{h=0}^{H}$. Standard errors are Newey-West. In our baseline specification we chose the lag length $M = L = 16$ quarters. We normalize the Romer shock to yield a 25 basis point decline in the real interest rate on impact. The decline in the real interest rate persists for 3-4 years (see Appendix Figure A.3).

The top-left panel of Figure 2 plots our estimated impulse response function for log GDP. It displays a hump-shaped increase that peaks at 0.30% after 9 quarters. For the next 9 quarters, GDP declines and it undershoots its trend before a gradual return to steady state. Both the initial positive peak and the subsequent negative trough are statistically

---

[6] Appendix C details the construction of the variables used in this analysis. We 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 1980-2006.

[7] For the adjustment probabilities we also include a squared time trend as these time series display a distinct U-shaped pattern.
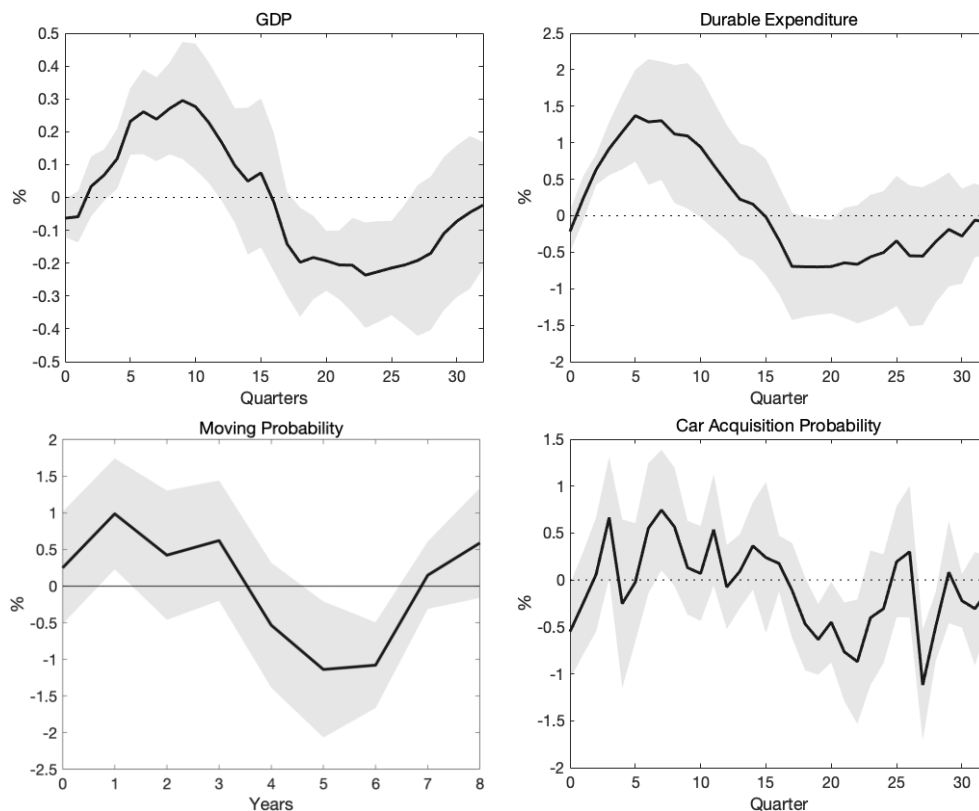
Figure 2: Impulse response function of real GDP (top-left panel), real durable expenditure (top-right), probability of moving house (bottom-left), and probability of buying a car (bottom-right) to a Romer and Romer monetary policy shock. Shaded areas are 95% confidence bands.

significant at conventional levels. Our estimates suggest that while monetary policy is able to stimulate economic activity in the short run, it comes at the cost of weaker activity in the medium run. This pattern is consistent with monetary policy borrowing demand from the future. Appendix D documents that this result is robust to excluding the deterministic trend, including fewer lagged terms, and restricting the sample to the post-Volcker period.

The panels for real durable expenditure, the moving probability, and the car acquisition probability in Figure 2 provide further evidence for such intertemporal shifting. In each case we observe a statistically significant positive response to expansionary monetary policy for the first three years followed by a statistically significant contraction. The volatility in the impulse response function for car acquisitions stems in part from the greater sampling variability in this series. These results are consistent with the hypothesis that households accelerate durable purchases following an expansionary monetary shock, but these adjustments
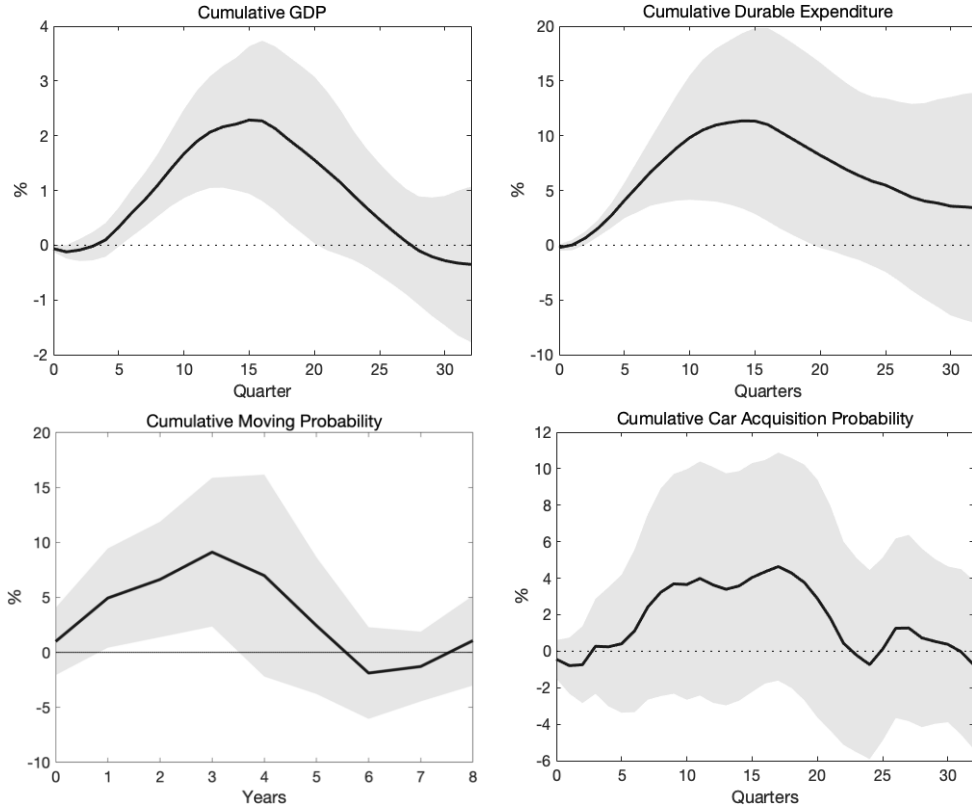
15

Figure 3: Impulse response function of cumulative real GDP (top-left panel), cumulative real durable expenditure (top-right), cumulative probability of moving house (bottom-left), and cumulative probability of buying a car (bottom-right) to a Romer and Romer monetary policy shock. Shaded areas are 95% confidence bands.

are subsequently missing.

In Figure 3 we plot estimates for the cumulative impulse response functions of GDP, durable expenditure, and the extensive margin. These correspond to the integral under the impulse response functions in Figure 2, which reveals the extent to which the initial increase in demand is later reversed. The point estimates of the cumulative impulse response functions are consistent with a complete reversal of GDP and the extensive margin, as well as a near-complete reversal of durable expenditure. The near complete reversal of durable expenditure is consistent with long run monetary neutrality, which implies the durable stock will be unaffected in the long run. While the confidence bands on our estimates generally allow for an incomplete reversal, note that from Figure 2 we can reject the hypothesis that no reversal takes place.

We estimate impulse response functions for longer horizons than are typical in the literature, which allows us to observe the intertemporal shifting effects. Typically impulse response functions are reported for up to 16 quarters (see Ramey, 2016, for a survey).[8] An exception is Jordà et al. (2020) who estimate a long-horizon effect of monetary policy using an identification scheme based on "trilemma" shocks. They find GDP is below normal for 12 years following a contractionary shock. Their shock implies that the real interest rate is also elevated for 12 years. In our estimates, the intertemporal shifting effect of monetary policy only becomes visible once the real rate normalizes, which occurs after 3-4 years for the Romer and Romer (2004) shock.

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 estimate the effect of a Romer-Romer shock on the real interest rate, r_t, aggregate income, Y_t, and the relative durable price, p_t, and feed the mean impulse response functions into the household problem of Section 2 starting in steady state.[9] We assume that these paths come as a surprise at $t = 0$ and then become known to agents when they update their information set. These variables return to steady state after 32 quarters.

The left panel of Figure 4 plots the model-implied cumulative durable expenditures against the data. In our model, the peak real durable response is 11.4%, the same as the peak real durable response in the data. 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 4 makes two additional model comparisons.[10] First, we compute the

---

[8]Tenreyro and Thwaites (2016) estimate the response of GDP to Romer and Romer (2004) shocks up to 20 quarters out. Their impulse response functions also display a reversal consistent with what we find.

[9]See Appendix E for details and estimates of the impulse response functions for r_t and p_t. For log GDP we use the impulse response function plotted in the top-left panel of Figure 2.

[10]For 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,
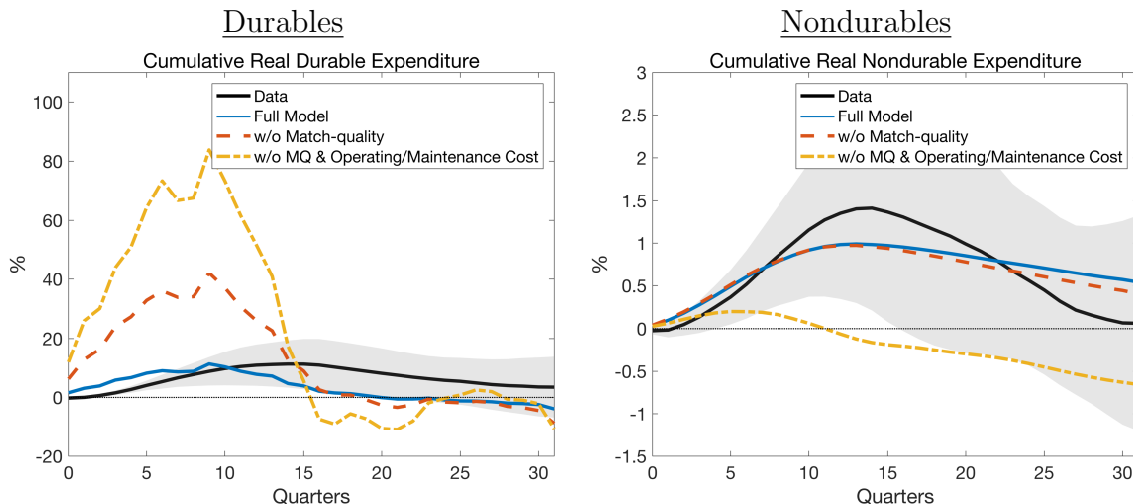
Figure 4: Cumulative response of durable (left) and nondurable (right) expenditure to a simulated monetary policy shock. Model simulations feeding in the estimated impulse responses for (Y_t, r_t, p_t). Each panel shows the full model, the model without the match-quality shocks, and the model without match-quality and operating and maintenance costs.

cumulative durable demand response in a model that abstracts from match-quality shocks but includes the operating and maintenance costs from our full model ("w/o match-quality shocks"). This model yields a counterfactually large peak cumulative durable demand response of 42%, four times larger than in the data. Second, we compute the response in a model that further abstracts from operating and maintenance costs ("w/o MQ and operating/maint. costs"). This model predicts that cumulative durable demand peaks at 84%, seven times more than their peak response in the data.

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 gradually cumulative durable expenditure increases is the information rigidity. Without information frictions, durable expenditure 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 our main results concern a long-horizon forecast for which the degree of information rigidity does not have major quantitative implications so matching the exact timing of durable expenditure in the data is not critical for our purposes.

and the durable-stock-to-nondurable-consumption ratio.

In the right panel of Figure 4 we show that the full model also provides a good match to the dynamics of cumulative nondurable expenditures. The model without match-quality shocks performs about as well, whereas the model abstracting from both match-quality shocks and operating/maintenance costs predicts too much substitution from nondurable to durable spending. The relatively small nondurable spending response in the data motivated our choice of a relatively low elasticity of intertemporal substitution in Section 2.

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 inframarginal 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.[11] 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 costs therefore stabilizes the user cost and therefore durable demand.

We now show that the willingness of households to shift the timing of their durable adjustments also aligns well with the observed extensive margin responses for cars and housing. For this analysis, we consider two different calibrations that interpret durables more narrowly as either cars or 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.[12] As above, we simulate the impulse response for the extensive margin by feeding

---

[11]The 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).

[12]The 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
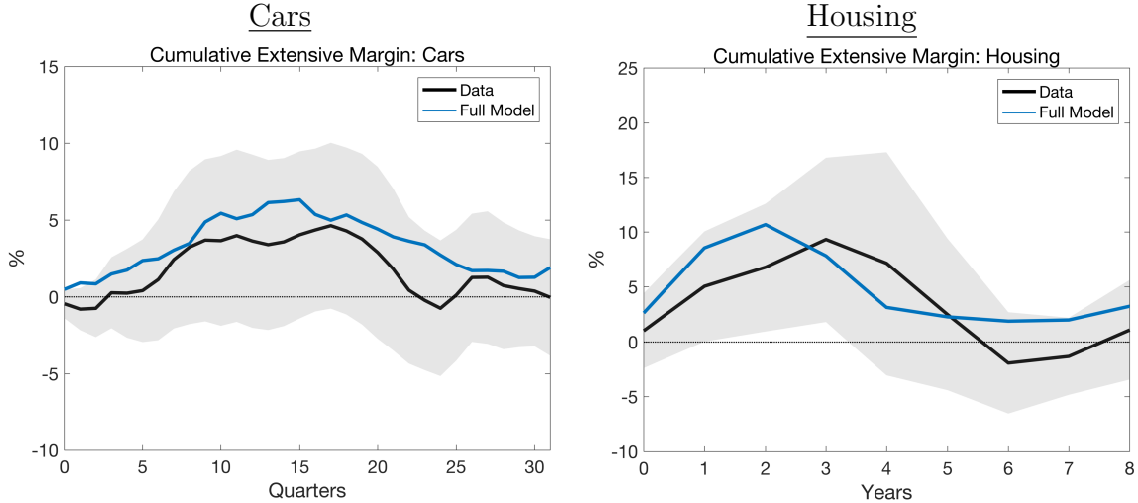
Figure 5: Cumulative response of the extensive margin of durable adjustment to a simulated monetary policy shock. Model simulations feeding in estimated impulse responses for (Y_t, r_t, p_t). The panels show the model for cars (left) and the model for housing (right) and the empirical estimates from Section 3.1.

the empirical impulse responses of Y_t, r_t, and p_t into the model. The left panel of Figure 5 shows that the extensive margin for car adjustments in our 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. The right panel shows the results for the housing model. Our 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.

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 durable expenditure through time. Empirical studies of anticipated VAT or sales

---

bonds. For Figure 5 we feed in the estimated impulse response of the relative price of cars in the car model and the estimated impulse response of relative price of housing in the housing model.
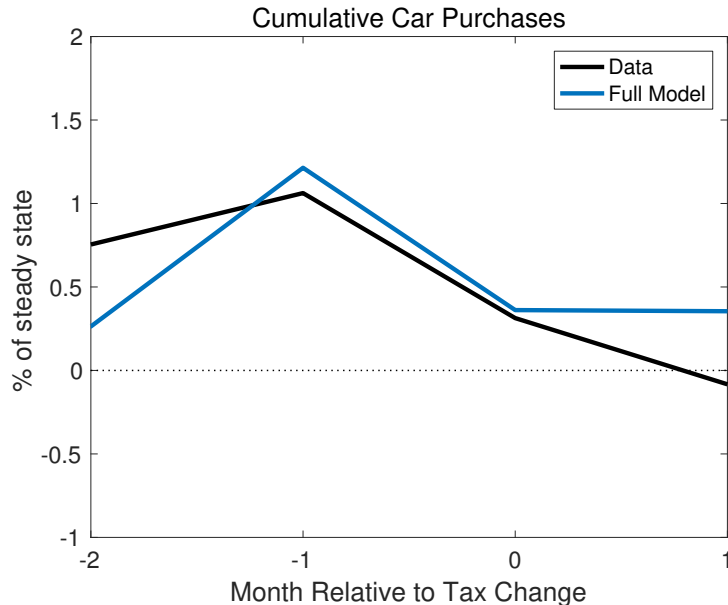
Figure 6: Cumulative car acquisitions for an anticipated sales tax increase in Baker et al. (2019) versus the model. We convert the tax elasticity in Baker et al. (2019) to an interest rate elasticity by dividing by 12.

tax changes 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). 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), who analyze the response of auto sales to anticipated changes in state sales tax rates.[13] 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. Such a sales tax change 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 $12.7/12 = 1.1$.

Using the model calibration that interprets durables as cars, we calculate the response of

---

[13]See their Table 3, Col. 1.

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) report that sales tax increases are known 2-3 months in advance for referenda and 6-9 months in advance for legislated changes and our assumption is the mid-point of these horizons. As shown in Figure 6, our model produces a peak elasticity of 1.2 and a subsequent reversal that tracks the estimate from Baker et al. quite closely.

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.

4.1   Production, Labor Supply, and Aggregate Supply

In designing the supply-side of our model we have a number of objectives. First, one theory of the decline in nondurable consumption and durable expenditure in the Great Recession is that households revised down their expectations for income growth (see De Nardi et al., 2011; Dupor et al., 2018). A reduction in income leads households to reduce their consumption expenditure through standard consumption-smoothing logic, but also to reduce their target durable stocks, which leads to an abrupt decline in durable expenditure. To incorporate these effects, we allow for shocks to aggregate productivity, which affect income expectations. A second objective in designing the model is to avoid abrupt changes in potential output. Our main results offer a positive explanation of how monetary policy was conducted during and after the Great Recession. In practice, policy makers typically view potential output as evolving in a smooth manner. In keeping with that view, we have designed the model so changes in productivity are serially correlated so that there can be a large change in expectations for permanent income without an abrupt change in realized productivity. For the same reason we also eliminate short-run wealth effects on labor supply, which we do in a manner inspired by Jaimovich and Rebelo (2009).[14] Third, we introduce labor supply into

---

[14]Cesarini et al. (2017) provide evidence for small short-run wealth effect on labor supply based on lottery earning in Sweden. Galí et al. (2012) estimate weak short-tun wealth effects on labor supply in aggregate data.

the model in a way that does not alter the household decision problem studied in Sections 2 and 3. This means the model validation performed in Section 3 remains applicable to the general equilibrium model. Achieving this requires an additively separable variant of the Jaimovich and Rebelo preferences. Fourth, we introduce nominal frictions in the model to generate a conventional New Keynesian Phillips curve. Our objective here is that the Phillips curve should be standard to make clear that our results are driven by the demand side of the model. For this purpose, we adapt the standard sticky-wage environment developed by Erceg et al. (2000) to allow for uninsured idiosyncratic labor productivity and productivity growth.

We now turn to the specifics of our assumptions. 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 aggregate labor supply. Productivity follows the process $\mathrm{d} \ln Z_t = g_t$ where g_t follows the Ornstein-Uhlenbeck process $\mathrm{d} g_t = -\rho_Z g_t \, \mathrm{d}t + \sigma_Z \, \mathrm{d} \mathcal{W}_t^Z$, where $\mathcal{W}_t^Z$ is a standard Brownian motion. Changes to productivity are permanent and the innovations are serially correlated. An innovation to this process can be interpreted as a shock to income expectations that takes some time to fully materialize.

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

$$E_0 \int_{t=0}^{\infty} e^{-\rho t} \left[u\left(c_{it}, s_{it}\right) - \Omega_t \int_0^1 v(n_{ijt}) \, \mathrm{d}j \right] \mathrm{d}t \tag{8}$$

$$\Omega_t \equiv \int_0^1 u_c\left(c_{it}, s_{it}\right) z_{i,t} Z_t \, \mathrm{d}i$$

where Ω_t is a time-varying preference shifter in the spirit of Jaimovich and Rebelo (2009). While Jaimovich-Rebelo preferences limit short-run wealth effects through a Greenwood, Hercowitz, and Huffman (1988) formulation, our desire for additive separability leads us to a different approach, which is to incorporate the productivity-weighted average marginal utility of consumption into Ω_t. Jaimovich-Rebelo preferences are compatible with a balanced growth path because the disutlity of labor eventually scales with consumption, which is proportional to productivity on the balanced growth path, implying the disutility of labor supply keeps pace with the return to work. We achieve a similar outcome by scaling Ω_t by Z_t directly.

Labor supply is determined by a set of unions as described below and 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.

Aggregate labor supply is given by

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

where

$$l_{jt} = \int_0^1 z_{it} n_{ijt} \, \mathrm{d}i.$$

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 obtain a standard New Keynesian Phillips curve through sticky nominal wages. A continuum of unions set the nominal piece rate, $\tilde{W}_{jt} = W_{jt}/Z_t$, of each type of labor. Specifying wages as a piece rate implies that nominal prices inherit the stickiness of wages, without us needing to incorporate additional pricing frictions. The union maximizes the equally-weighted utility of the households subject to a Rotemberg-style adjustment cost of $\frac{\Psi}{2} \Omega_t L_t \left(\mu_{jt} \right)^2$, where Ψ is a parameter that controls the strength of the nominal rigidity and μ_{jt} is the growth rate of $\tilde{W}_{jt}$ such that $\mathrm{d}\ln \tilde{W}_{jt} = \mu_{jt} \, \mathrm{d}t$.[15] Among union workers supplying type j, all labor is equally rationed, $n_{ijt} = l_{jt}$. 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 F presents the union's problem in detail and shows that the linearized symmetric equilibrium gives rise to the following Phillips curve

$$\dot{\pi}_t = \rho \pi_t - \kappa \left(\frac{Y_t - \bar{Y}_t}{\bar{Y}_t} \right), \tag{9}$$

---

[15]We model the Rotemberg adjustment cost as a utility cost for the union rather than as a resource cost based on the arguments in Eggertsson and Singh (2019).

where $\pi_t = \frac{d\ln P_t}{dt}$ and $\bar{Y}_t$ is potential output. This is a continuous-time version of the standard New Keynesian Phillips curve.

We now turn to the supply of durable goods. The durable good is produced by a representative firm using the production function,

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

where X_t is the production of durables, M_t is the input of the nondurable good, and υ 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

$$p_t = (1-\zeta)^{-1}\upsilon^{-\frac{1}{1-\zeta}}\left(\frac{X_t}{Z_t\bar{H}}\right)^{\frac{\zeta}{1-\zeta}}. \tag{10}$$

Thus, $(1-\zeta)/\zeta$ is the supply elasticity of the durable good.

The final good is used for several purposes including nondurable 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 Ornstein-Uhlenbeck process given by $dr_t^b = -\rho_{r^b}(r_t^b - \bar{r}^b)\,dt + \sigma_{r^b}\,d\mathcal{W}_t^{r^b}$. Appendix F shows the market-clearing conditions of the model.

4.2 Government

Monetary policy is governed by a standard interest rate rule,

$$\dot{r}_t = -\rho_r(r_t - \bar{r}) + \phi_\pi \pi_t + \phi_y \frac{Y_t - \bar{Y}_t}{\bar{Y}_t} + \eta_t^r, \tag{11}$$

where the first term captures interest rate smoothing, the second and third terms capture the endogenous monetary policy response, and η_t^r is an exogenous term that follows $d\eta_t^r = -\rho_\eta \eta_t^r\,dt + \sigma_\eta\,d\mathcal{W}_t^\eta$.

Fiscal policy consists of a constant debt policy,

$$A_t = \int_0^1 a_{it}\,dt = \bar{A}. \tag{12}$$

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

$$y_{it} = (Y_t - \bar{\tau}_t)\, z_{it} \tag{13}$$

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.[16] We assume that G_t follows an Ornstein-Uhlenbeck process, $\mathrm{d}\ln G_t = -\rho_G (\ln G_t - \ln \bar{G})\, \mathrm{d}t + \sigma_G\, \mathrm{d}\mathcal{W}_t^G$. 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 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 we estimate in response to monetary shocks (see Appendix Figure A.3). An elastic supply of

---

[16]The 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.

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 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.[17]

The slope of the Phillips curve is 0.48. 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 one needs to divide by 16 to compare to a quarterly discrete-time model, which yields a slope of 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), $\phi_\pi = 0.79$, $\phi_y = 0.75$, and $\rho_\eta = 1.55$ (equivalent to a quarterly persistence of 0.68).[18] 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. Similarly, we set $\rho_g = 0.77$ equivalent to a quarterly persistence of 0.83 such that an innovation to productivity achieves 90% of its long-run level within three years. This speed of convergence strikes a balance between our focus on cyclical developments and a smooth evolution of potential output.

---

[17]An alternative explanation for the muted relative durable price response to monetary policy shocks is nominal rigidities in the durables sector. Both formulations should have similar implications for durable demand provided that they generate similar dynamics for relative prices.

[18]The 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.

4.4 Solving the Model

We use the sequence-space approach of Auclert et al. (2019) to solve for the aggregate dynamics of the economy. The first step in this method is to compute the partial equilibrium response of aggregate variables to changes in prices and aggregate shocks under perfect foresight. To do so, we approximate the value functions and policy rules at discrete time steps as is standard practice (see Achdou et al., 2017). We use quarterly time steps. We then arrive at a set of partial equilibrium Jacobians in which the (i, j) elements relate changes in aggregate quantities at a horizon of i quarters to changes in prices or shocks at a horizon of j quarters. We then translate the partial equilibrium Jacobians of the model into general equilibrium Jacobians by solving for the endogenous prices that satisfy the market clearing conditions (10) and (12). See appendix A for details. This procedure assumes that the economy's dynamics are linear in the aggregate states but allows for nonlinear policy rules with respect to idiosyncratic states. The solution is equivalent to the impulse response functions obtained from a perturbation approach such as the Reiter (2009) method (see Boppart, Krusell, and Mitman, 2018).

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. 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.[19] We have also investigated the robustness of our results to nonlinear aggregate dynamics. We return to both of these issues below.

5 The Monetary Transmission Mechanism

The acceleration of durable purchases in response to monetary policy stimulus implies that monetary policy shifts demand intertemporally. This feature of the monetary transmission mechanism is captured by the sequence-space Jacobian of the output gap with respect to real interest rates. We call it the "monetary transmission matrix" and denote it by $\mathcal{M}$. The (i, j) element of $\mathcal{M}$ gives the general equilibrium response of the output gap at a horizon of

---

[19]Our method of incorporating the ELB through news shocks is an implementation of the method described by Holden and Paetz (2012).

$i - 1$ quarters with respect to news about real interest rates at a horizon $j - 1$ quarters,

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

where $\hat{Y}_t = \frac{Y_t - \bar{Y}_t}{\bar{Y}_t}$ is the output gap. The discrete-time nature of this matrix reflects the quarterly time steps on which we approximate the solution of the model. 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.

To illustrate how $\mathcal{M}$ captures the intertemporal shifting of aggregate demand, in the left panel of Figure 7 we plot columns one, five, and nine of our model's $\mathcal{M}$ multiplied by -0.01. These show the impulse response function of the output gap to a surprise 1% (annualized) real interest rate cut that occurs in the current quarter (column 1), four quarters from now (column 5), and eight quarters from now (column 9). In each case, output expands in the run-up to the interest rate cut, but once stimulus is removed, output falls below steady state. The initial increase in output primarily reflects the low user cost of durables. While the real interest rate falls only at one date, the user cost falls before that date due the anticipated increases in the relative price of durables. The subsequent decline in output reflects missing durable demand as durable adjustments are shifted earlier in time.

To show that the intertemporal shifting of aggregate demand is a consequence of durable demand, the right panel of Figure 7 plots the same columns in a model with only nondurable goods. The nondurables model contains the same set of ingredients as our full model (e.g., idiosyncratic risk, sticky information) but it abstracts from durable consumption. We calibrate the intertemporal elasticity of substitution, σ^{-1}, to match the peak output effect in our full model in response to a 1% reduction in the real interest rate that lasts for 9 quarters $(t = 0$ to $t = 8)$.[20] In this model, the output responses to interest rate cuts are positive at all horizons so there is no intertemporal shifting. The increasing output effect for more distant interest rate cuts reflects the fact that more households will have learned of the shock. The spike in output when the interest rate cut occurs reflects the redistribution from creditors

---

[20]The durable share in utility is set to zero, $\psi = 0$, rendering $\delta, \xi, f, \theta, \nu$ irrelevant. The borrowing limit is set to $-\lambda$ times the 25^{th} percentile of durable holdings in our full model. The parameter ρ is set to match the same net asset to GDP ratio as in the full model. All other parameters are identical to the full model.
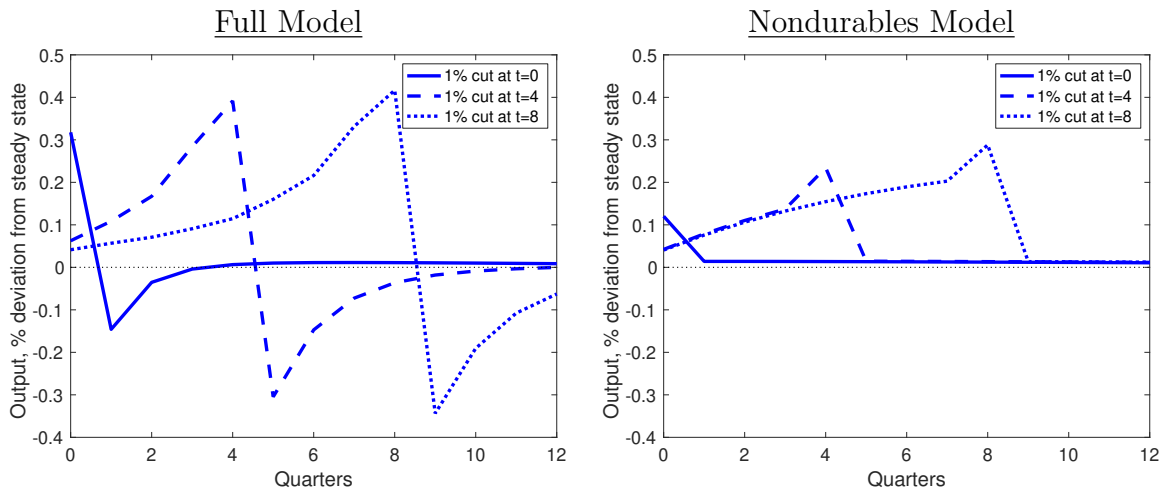
Figure 7: Percentage change in output following a 1% reduction in the real interest rate at horizons $t = 0$ (solid line), $t = 4$ (dashed line), and $t = 8$ (dotted line). The left panel shows the effect in the full model, and the right panel shows the effect in the nondurables model.

to debtors (see Auclert, 2019).

Several recent studies explore the role of marginal propensities to consume in amplifying changes in demand that are directly generated by monetary policy (e.g., Kaplan et al., 2018; Auclert et al., 2020). While our model does well in matching the aggregate response of durable expenditure, nondurable consumption, and the extensive margin to monetary policy shocks (see Section 3), its average marginal propensity to spend is below typical empirical estimates registering 0.19 per quarter for total expenditure and 0.04 for nondurable expenditure. The low marginal propensity to spend reflects the fact that households in the model can easily borrow against the value of their durables. We have found that tightening the borrowing limit raises the marginal propensity to spend but has little impact on the intertemporal shifting effects we focus on.

6 Interest Rates During and After the Great Recession

We now turn our attention to the 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 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 the 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:[21] 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 $s_t^x = \frac{p_t x_t}{\bar{Y}_t}$; 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, $r_t^b - \bar{r}^b$.[22]

We will extract quarterly aggregates of the innovations to the exogenous processes $\mathcal{W}_t^Z$, $\mathcal{W}_t^G$, $\mathcal{W}_t^\eta$, and $\mathcal{W}_t^{r^b}$ from the data. They correspond to the productivity shock, the non-household demand shock,[23] the monetary policy shock, and the shock to the borrowing spread.

We use a novel filtering algorithm that we describe in detail in Appendix G. 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 I. 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 innovations at date t that explain the difference between the data observed at date t and our forecast. We 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. The

---

[21]We 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 nondurable inflation expectations and thus the ex ante real rate. (The persistence is much higher before 1991.)

[22]We 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.

[23]In 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.

specific initialization date 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 G). 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, $r \geq \underline{r} = -2.5$.[24] 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 path of rates to violate the constraint. To ensure this does not happen, we incorporate monetary news shocks. At any point in time t for horizon $h \geq 1$ we calculate the extent to which the path violates the ELB, $\mathbb{E}_t r_{t+h}^{news} = \max\{\underline{r} - \mathbb{E}_t r_{t+h}, 0\}$. We then introduce $\{\mathbb{E}_t r_{t+h}^{news}\}_{h=1}^{H}$ as news shocks about the future path of the real interest rate. 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 other shocks. 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

---

[24]The 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%.

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. This spread can account for a divergence between interest rates on Treasuries and household debt due to, for instance, convenience yield.

In Figure 8 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.[25] 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 Rates after the Great Recession

Figure 9 shows our main result: our model predicts a very slow normalization of interest rates after the Great Recession. We plot the model's forecast of the real interest rate against the data. To construct the figure we use the filtered shocks up to 2012Q4 and then use the model to 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 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 real interest rate exactly up to 2012Q4 by construction because it is one of the series we match in our filtering. We use the beginning of 2013 as a benchmark for our forecasts

---

[25]The spread between mortgages rates and Treasuries only displayed a short-lived spike. This contrasts with the behavior of convenience yields, as measured by the spread between corporate bonds and US Treasuries, which were particularly elevated in the years following the Great Recession (e.g. Del Negro et al., 2017). The decline in real interest rates faced by households during and after the Great Recession implies factors beyond convenience yields are needed to explain the low level of policy rates over this period.
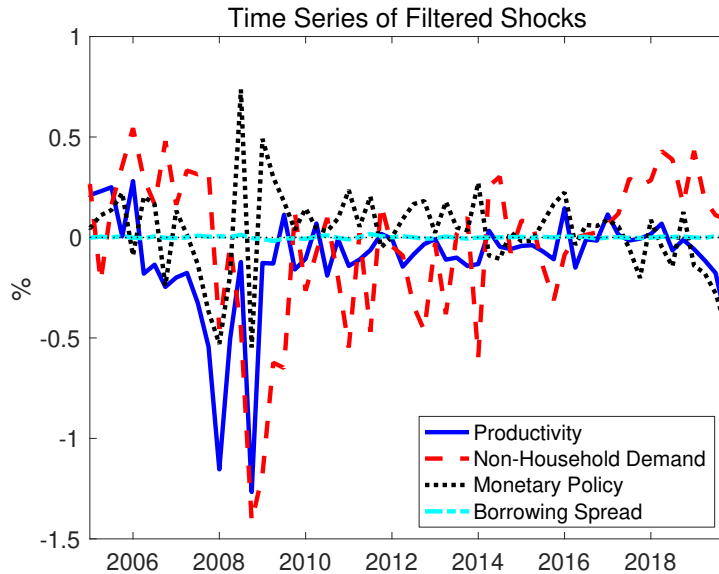
Figure 8: Aggregate shocks filtered from the output gap, the real interest rate, the change in durable spending as a fraction of potential output, and the mortgage-Treasury spread. The shocks are scaled based on their contemporaneous effect on the output gap.

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.

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.

The intertemporal shifting of durable demand in the model is key to the slow normalization of real interest rates following the Great Recession. One way to show this is the following thought experiment: what would have happened if the monetary transmission mechanism in the model worked just as the monetary transmission mechanism in the nondurables model? To implement this experiment, we replace the $\mathcal{M}$ matrix in the full model with the $\mathcal{M}$ matrix from the nondurables model described in Section 5. We then filter the data and produce a new forecast for the real interest rate after 2012Q4. Figure 9 shows that the nondurables $\mathcal{M}$ matrix predicts a much more rapid normalization of real interest rates.[26] This is because

_____

[26]The results are very similar if we instead use the entire nondurables model for filtering and producing the real rate forecasts. We prefer to isolate the contribution of the $\mathcal{M}$ matrix as using the full nondurables
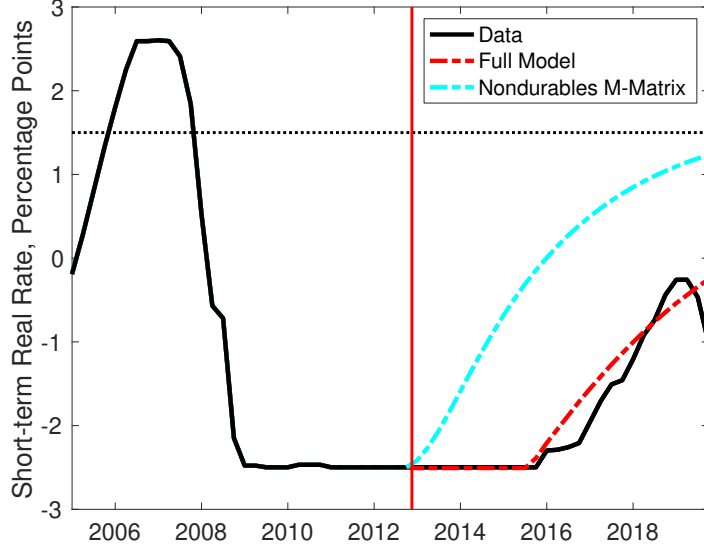
Figure 9: Real interest rate from the model (forecasted after 2012Q4), and ex-ante real interest rate in the data. The dotted horizontal line is the steady state real interest rate, equal to 1.5%.

under the nondurables $\mathcal{M}$ matrix, monetary policy does not need to keep interest rates low to counteract intertemporal shifting effects from past monetary stimulus.

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

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

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.

model requires us to also change the filtering procedure as that model cannot be used to interpret data on durable expenditure.

The sequence space allows for a convenient method to solve for r^*. Define $\vec{\hat{Y}}_t \equiv (\hat{Y}_t, \mathbb{E}_t\hat{Y}_{t+1}, \mathbb{E}_t\hat{Y}_{t+2}, \cdots)'$ as the vector of time t forecasts of output gaps at all future dates, $\vec{r}_t \equiv (r_t - \bar{r}, \mathbb{E}_t r_{t+1} - \bar{r}, \mathbb{E}_t r_{t+2} - \bar{r}, \cdots)'$ as the vector of time t forecasts of real interest deviations at all future dates, and define $\eta_t \equiv (g_t, G_t - \bar{G}, r_t^b - \bar{r}^b)'$ as the vector of exogenous stochastic processes. In Appendix H we show using a first order approximation around steady state that

$$\vec{\hat{Y}}_t = \mathcal{M}\vec{r}_t + \mathcal{Q}\eta_t + \mathcal{D}(\Phi_t - \bar{\Phi}), \tag{14}$$

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 impulse responses in $\mathcal{Q}$ are calculated under the assumption of no change in real interest rates. 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 $\vec{\hat{Y}}_t = 0$. Using equation (14) we have

$$\vec{r}_t^* = -\mathcal{M}^{-1}\left(\mathcal{Q}\eta_t + \mathcal{D}\left(\Phi_t - \bar{\Phi}\right)\right). \tag{15}$$

$\vec{r}_t^*$ 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.

In Appendix H 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 $t = 0$,

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

$$-\sum_{k=0}^{t}\mathcal{M}_{[1+t-k..,1+t-k..]}^{-1}\mathcal{Q}_{[1+t-k,..]}\left[\mathbb{E}_k\left(\eta_k\right) - \mathbb{E}_{k-1}\left(\eta_k\right)\right]. \tag{16}$$

36

Importantly, we do not need to infer the high-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 (16), start with the second term that involves the surprise in the exogenous processes η_k and set $k = 0$. The matrix $Q_{[1+t,..]}$ tells us how a shock at time 0 affects the output gap from time t onward. The operation $-\mathcal{M}^{-1}_{[1+t..,1+t..]}$ 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 (16) 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.[27] The $\mathcal{M}_{[1+t..,1..t]}$ matrix multiplying this term converts these news into changes in the output gap for time t onward. And the $-\mathcal{M}^{-1}_{[1+t..,1+t..]}$ 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.[28]

Equation (16) shows that the below-diagonal elements of $\mathcal{M}$ 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.

---

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

[28]In DSGE models, r^* is often defined in terms of a flexible-price equilibrium. In our equation (16) this corresponds to past interest rates being set equal to the natural rate of interest,

$$\mathbb{E}_k \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} = \mathbb{E}_k \begin{pmatrix} r^*_k \\ \vdots \\ r^*_{t-1} \end{pmatrix}.$$

In contrast, we compute r^* based on the realized and expected real interest rate path at each date.
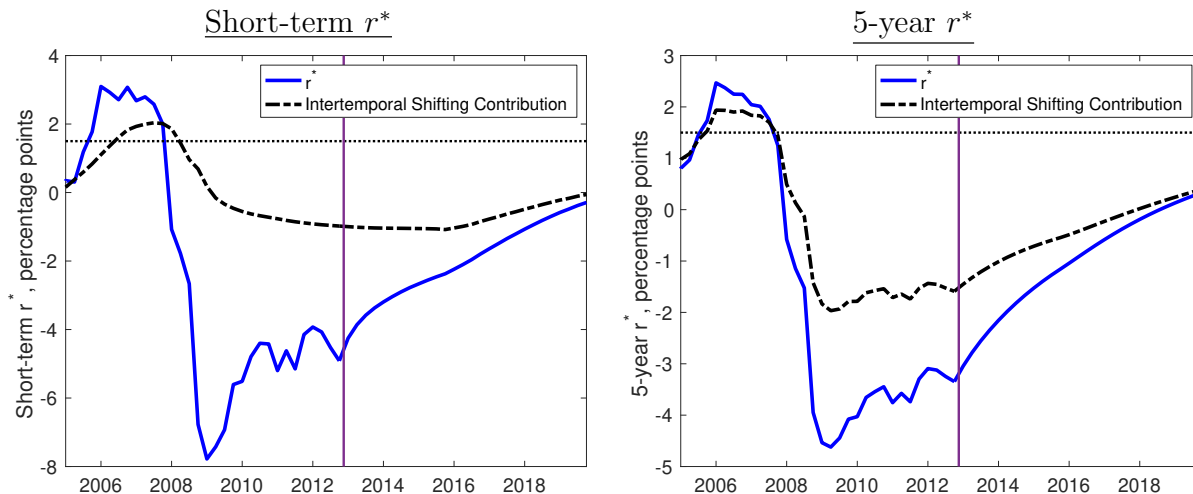
Figure 10: Time series of the short-term natural rate of interest (left panel) and the 5-year natural rate of interest (right panel). The dashed black line shows the contribution of intertemporal shifting effects of previous real interest rates to r^*. Forecast as of 2012Q4. The dotted horizontal line is the steady state real interest rate, equal to 1.5%.

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 10 we plot the short-term natural rate of interest implied by our model—that is the first element of $\vec{r}_t^*$ for each date t. There is 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, r^* is forecast to be 3.9 percentage points below steady state. Even in 2019Q4, r^* is 1.9 percentage points below steady state.

(S, s) models can generate nonlinear dynamics that make durable expenditure more sensitive to stimulus in a boom than in a recession (e.g. Berger and Vavra, 2015). In Appendix J, we investigate the sensitivity of our results to nonlinear aggregate dynamics such as this. From this investigation we conclude that our analysis may understate the decline in r^* during the Great Recession, but the nonlinearities do not have important effects for the medium-run forecasts of r^* that are our main results.

The figure plots the contribution of intertemporal shifting to 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 for two reasons. First, because of the zero lower bound constraint, the real interest rate path is also almost flat over much of this period and this determines the amount of intertemporal shifting. Second, even after lift off, the model predicts a gradual normalization of interest rates so the contribution of intertemporal shifting only recedes gradually.

In the right panel of Figure 10 we show the 5-year r^*, computed from expectations of short-term rates using the expectations hypothesis. The 5-year r^* displays a corresponding large drop and slow normalization. The 5-year r^* is 2.7 percentage points below steady state in 2015Q4 and 1.2 percentage points below steady state in 2019Q4. Intertemporal shifting largely accounts for the very persistent behavior of these long-term rates.[29] Thus, our model 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 r^* remained low in subsequent periods.

In Figure 11, we decompose the path for the short-term natural rate of interest into contributions from each shock. The red dashed line shows that the real interest rate must fall by almost 7 percentage points to offset the effect of the productivity shock on the output gap. This calculation incorporates both the direct effects of the productivity shock as well as the intertemporal shifting effects from offsetting past productivity shocks. As Figure 11 shows, the productivity shock is particularly important in explaining the persistently low level of r^*. This is because the productivity shock causes persistent weakness in aggregate demand as households adjust to a lower durable stock. Persistent weakness in demand leads to persistently low interest rates, which in turn generate persistent intertemporal shifting effects leading to continued low real interest rates.

Due to the binding ELB, real interest rates did not decline as much as r^* did. With less

---

[29]To compute the intertemporal shifting contribution to an H-horizon long-term rate, we solve for the intertemporal shifting contribution to the short rate at $t, t + 1, ..., t + H$ and average them in accordance with the expectations hypothesis.
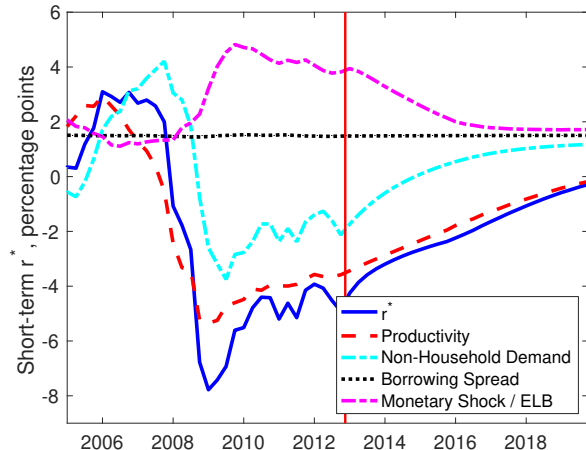
Figure 11: Time series of the short-term natural rate of interest implied by each shock. The blue line is the level of r^*, the red line is the level of r^* implied by the productivity shock, the cyan line the is the level of r^* implied by the non-household demand shock, the black line is the level of r^* implied by the borrowing spread shock, and the magenta line is the level of r^* implied by the monetary policy shock and the effective lower bound. Forecasts as of 2012Q4.

monetary accommodation in the recession, there were weaker intertemporal shifting effects leading to a higher path of r^*. The magenta line in Figure 11 shows the contribution to r^* from past monetary policy that was tighter than r^*.

This figure shows that the nature of the shocks causing a recession can be an important determinant of the future evolution of r^*. A drop in the desired durable stock, such as the one caused by the productivity shock, can cause a persistent decline in r^*. In contrast, recessions induced by tight monetary policy will be followed by an above-normal natural real rate of interest as durable demand shifts into the future.

7   Conclusion

We develop a fixed-cost model of durable consumption demand 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

40

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. In this manner, intertemporal shifting effects lead to a slow normalization of r^* following 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 r^* towards the lower bound implied by the ELB thereby reducing future policy space. 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," .

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.

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," Federal Reserve Bank of Boston Research Department Working Paper No. 14-3.

Baker, S. R., L. Kueng, L. McGranahan, 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.

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.

Carroll, C. D., E. Crawley, J. Slacalek, K. Tokuoka, and M. N. White (2020): "Sticky expectations and consumption dynamics," *American Economic Journal: Macroeconomics*, 12, 40–76.

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*.

Cesarini, D., E. Lindqvist, M. J. Notowidigdo, and R. Östling (2017): "The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries," *The American Economic Review*, 107, 3917–3946.

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.

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.

DE NARDI, M., E. FRENCH, AND D. BENSON (2011): "Consumption and the great recession," Tech. rep., National Bureau of economic research.

DEL NEGRO, M., D. GIANNONE, M. P. GIANNONI, AND A. TAMBALOTTI (2017): "Safety, liquidity, and the natural rate of interest," *Brookings Papers on Economic Activity*, 2017, 235–316.

DUPOR, B., R. LI, M. MEHKARI, AND Y.-C. TSAI (2018): "The 2008 US Auto Market Collapse," FRB St. Louis Working Paper No. 2020-004.

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.

EGGERTSSON, G. B. AND S. R. SINGH (2019): "Log-linear Approximation versus an Exact Solution at the ZLB in the New Keynesian Model," *Journal of Economic Dynamics and Control*, 105, 21–43.

EICHENBAUM, M., S. REBELO, AND A. WONG (2018): "State Dependent Effects of Monetary Policy: The Refinancing Channel," .

ERCEG, C. AND A. LEVIN (2006): "Optimal monetary policy with durable consumption goods," *Journal of Monetary Economics*, 53, 1341–1359.

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.

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.

GALÍ, J., F. SMETS, AND R. WOUTERS (2012): "Unemployment in an estimated New Keynesian model," *NBER macroeconomics annual*, 26, 329–360.

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.

GREENWOOD, J., Z. HERCOWITZ, AND G. W. HUFFMAN (1988): "Investment, capacity utilization, and the real business cycle," *The American Economic Review*, 402–417.

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.

HOLDEN, T. AND M. PAETZ (2012): "Efficient Simulation of DSGE Models with Inequality Constraints," University of Surrey Discussion Paper 15/12.

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.

JAIMOVICH, N. AND S. REBELO (2009): "Can News about the Future Drive the Business Cycle?" *American Economic Review*, 99, 1097–1118.

JORDÀ, Ò., S. R. SINGH, AND A. M. TAYLOR (2020): "The long-run effects of monetary policy," Tech. rep., National Bureau of Economic Research.

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

KILIAN, L. (1998): "Small-sample confidence intervals for impulse response functions," *Review of economics and statistics*, 80, 218–230.

KREAMER, J. (2019): "Sectoral Heterogeneity and Monetary Policy," Forthcoming in American Economic Journal: Macroeconomics.

LAUBACH, T. AND J. C. WILLIAMS (2016): "Measuring the natural rate of interest redux," *Business Economics*, 51, 57–67.

LEAMER, E. E. (2007): "Housing is the business cycle," in *Economic Policy Symposium Proceedings-*, Federal Reserve Bank of Kansas City, 149–233.

——— (2009): "Homes and Cars: Why are the Cycles in Homes and Consumer Durables so Similar?" *The B.E. Journal of Economic Analysis & Policy*, 9.

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. R., L. STRAUB, AND A. SUFI (2020): "Indebted Demand," Working Paper 26940, National Bureau of Economic Research.

MIDRIGAN, V. (2011): "Menu costs, multiproduct firms, and aggregate fluctuations," *Econometrica*, 79, 1139–1180.

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," .

RAMEY, V. A. (2016): "Macroeconomic shocks and their propagation," in *Handbook of macroeconomics*, Elsevier, vol. 2, 71–162.

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.

SIMS, C. A. (1996): "Inference for multivariate time series models with trend," in *Originally presented at the August 1992 ASSA Meetings*.

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.

TENREYRO, S. AND G. THWAITES (2016): "Pushing on a string: US monetary policy is less powerful in recessions," *American Economic Journal: Macroeconomics*, 8, 43–74.

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.

ONLINE APPENDIX

A Computational Appendix

We solve the model building on the methods described in Achdou et al. (2017).

A.1 Steady state

Define $k = a + \lambda(1 - f)pd$ 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`.

46

5. Compute optimal choice of d' conditional on adjusting and the corresponding $v^* = \max_{d'} v(k', d', z)$, where $k' = k + (1-f)(1-\lambda)pd - (1-\lambda(1-f))pd'$.

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 $C^{noadj} = A - diag(\theta)$. 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

47

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 (10) and (12).

B   Estimation

The data selection and estimation strategy largely follows Berger and Vavra (2015).

B.1   Data

Observables. We use PSID data from 1999 through 2009. 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 $\bar{c}_{i,t,t+2}$.

Real nondurable consumption is nominal nondurable consumption in the PSID deflated by the BEA price index for nondurables (NIPA table 1.1.4). Nominal nondurable 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 nondurables 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 nondurables 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 nondurables 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. We weight a housing adjustment by 0.9 and car adjustments by 0.1.

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 5^{th} and 95^{th} percentile. The sample weight is the household weight in the PSID.

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 nondurable consumption, liquid asset holdings, and real debt holdings by a household's average nondurable consumption over the sample. This helps account for permanent differences in income, which are again not part of the model.

Consistency with national aggregates. We divide all variables by average nondurable 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-nondurable-consumption ratio) and the flow of funds (liquid-asset-to-nondurable-consumption and debt-to-nondurable-consumption). The rescaling is necessary because the PSID collects data for 72% of nondurable 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

1. Pick a given intensity of match quality shocks θ. Calibrate the discount rate ρ, the fixed cost f, and the preference for durables ψ to match the targets for net assets, the probability of adjustment, and the durable-stock-to-nondurable-consumption ratio.

2. Forecast the probability of adjustment $P(a, d, y)$ over the next two years. Also forecast

the average nondurable consumption expenditure $\bar{c}$ 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 error to the model variables $a, d, \bar{c}$ using three independent Gaussian quadratures. This yields a new distribution $\hat{G}(a, d, \bar{c})$ which includes measurement error.

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.

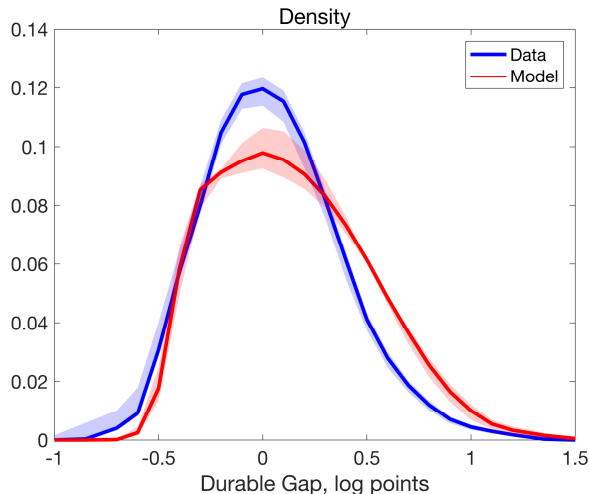Figure A.1 displays the density of gaps at our estimated parameters.

Figure A.1: Density of 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.

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, 3 and 4 of Section 3. We obtained the data from the the St Louis Fed FRED database. The variable identifiers are listed in Table 2.

To construct real durable and nondurable 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 indeces, Δp_{1t} and Δp_{1t}. 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.

Variable Name	FRED Series Code
Population	B230RC0Q173SBEA
Income (GDP)	GDPC1
Federal Funds Rate	FEDFUNDS
Consumer Durable Expenditure	PCDG
Residential Investment	PRFI
Consumer Nondurable Expenditure	PCEND
Consumer Service Expenditure	PCES
Consumer Housing Services Expenditure	DHSGRC0
Durable Price Index	DDURRD3Q086SBEA
Residential Investment Price Index	B011RG3Q086SBEA
Nondurable 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.

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 the associated 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 durables 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 nondurables and services. The real interest rate is defined in terms of nondurables. It is the

federal funds rate net of realized nondurable inflation over the next four quarters.

C.2 PSID: Housing Adjustment Probability

We use the Panel Survey of Income Dynamics (PSID) to construct a time series for the probability of housing adjustments. We use data from 1969-1997 when the survey frequency is annual. We keep only people who are heads of household and those who are in the Survey Research Center (SRC) sample.

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.

Following 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. 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 download pre-compiled files from the BLS for 1996 onwards and earlier raw files from ICPSR. We clean the micro-data files following Coibion et al. (2017).

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 implies by the CEX tracks BEA personal consumption motor vehicle expenditure very well.

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.

D Robustness of Impulse Response Functions

In each plot of Figure A.2 we compare our baseline impulse response function for GDP (blue line) against an alternative specification (red line). In Figure A.2a we drop the deterministic trend. This helps allay concerns that we are biasing the model towards stationarity (Sims, 1996). In Figure A.2b we include only four lags of the dependent variable and the monetary shock (vs 16 in the baseline) to address concerns that we may overfit the data. And in Figure A.2c we restrict the sample to the post-Volcker period, 1984-2016. Due to the shorter sample we reduce the lag length to four in that last case. For each of these three alternative specifications, the estimated response is close to our baseline estimates both in economic and statistical terms. In particular, all alternative specifications display an initial increase in GDP and subsequent reversal.

E Impulse Response Functions for r_t and p_t

We estimate the impulse response of the real interest rate in terms of nondurable goods, r_t and the relative durable price, p_t to a Romer-Romer monetary policy shock. For 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 in Section 3. Other than generating smoother response beyond 16 quarters, the impulse response functions estimated by the VAR are very similar to impulse response functions estimated by local projections.

We estimate two bivariate VARs, in which the monetary shock is ordered first. The second variable is respectively the real interest rate and the change in the relative price of durables. The VAR also includes a time-trend and the standard errors are block-bootstrapped and bias-corrected following Kilian (1998). Figure A.3 plots these impulse response functions.

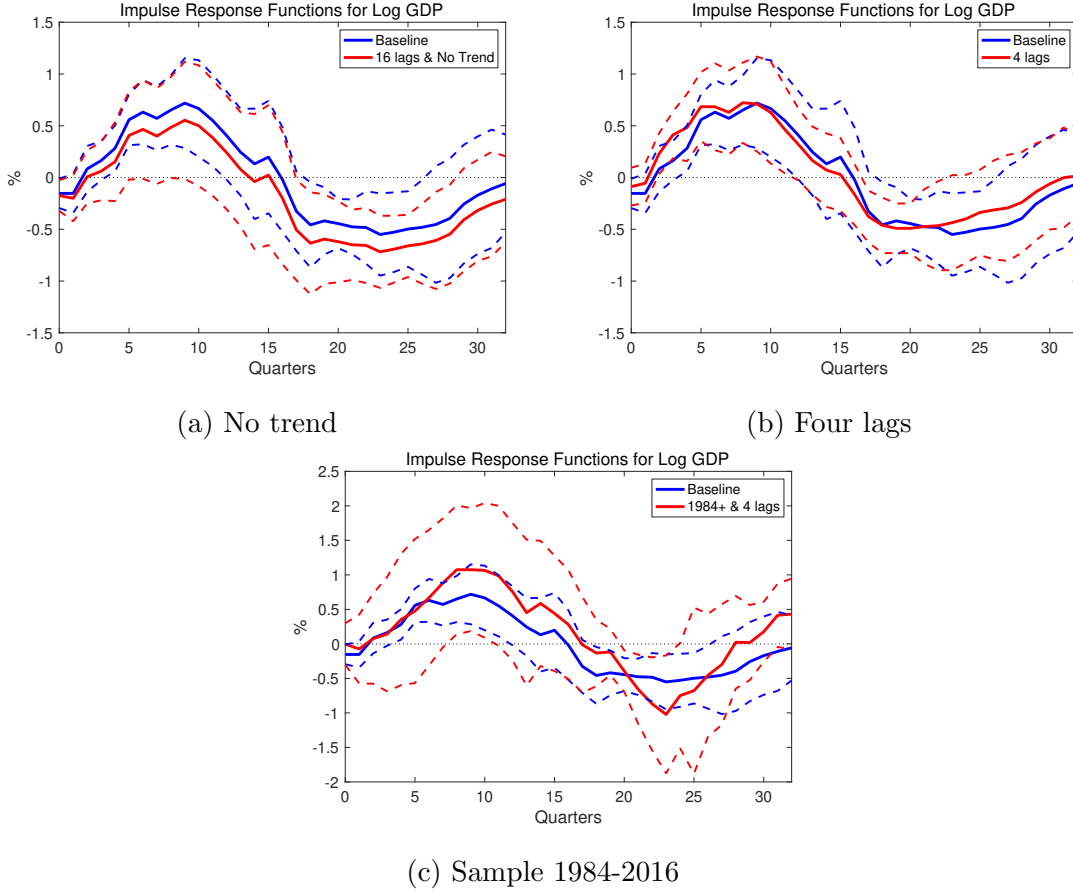(a) No trend

(b) Four lags

(c) Sample 1984-2016

Figure A.2: Robustness of the impulse response function for GDP estimated in Section 3. The blue line depicts the baseline specification and the red line the alternative specification. Dashed lines are 95% confidence intervals. In panel (a), the alternative specification drops the deterministic time trend. In panel (b), the alternative specification includes four lags of the dependent variable and the monetary policy shock (as opposed to 16 in the baseline). In panel (c), the alternative specification is estimated over 1984-2006 and with four lags.
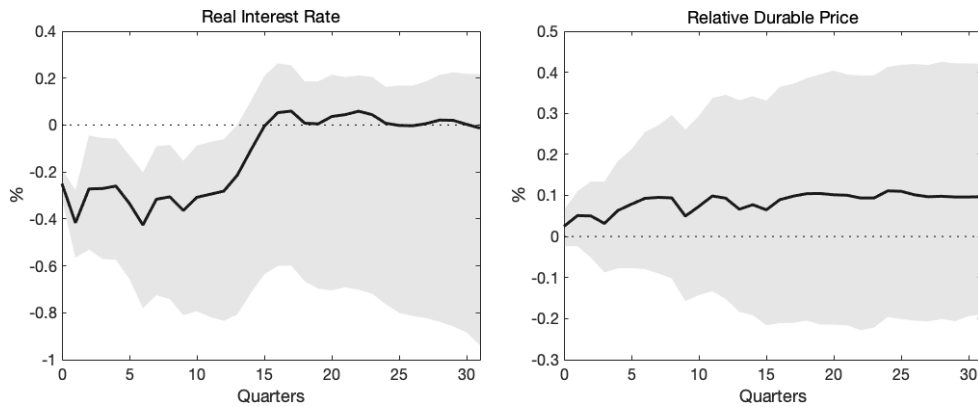

Figure A.3: Impulse response function of the real interest rate in terms of nondurables (left panel) and the relative durable price (right panel) to a Romer-Romer monetary policy shock.

F   Details of the General Equilibrium Model

F.1   The Labor Market

The labor demand curve of each labor type j is,

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

where the aggregate wage is equal to

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

The union's problem can be stated in terms of piece rates $\tilde{W}_{jt} = W_{jt}/Z_t$

$$\max_{\{\mu_{jt}\}} \int_{t=0}^{\infty} e^{-\rho t} \int_0^1 \left[u_c(c_{it}, s_{it}) \frac{\tilde{W}_{jt} Z_t}{P_t} l_{jt} z_{it} - \Omega_t v(l_{jt}) - \frac{\Psi}{2} \Omega_t L_t \mu_{jt}^2 \right] di\, dt$$

subject to

$$\mathrm{dln}\, \tilde{W}_{jt} = \mu_{jt}\, dt$$

$$l_{jt} = L_t \left(\frac{\tilde{W}_{jt}}{\tilde{W}_t} \right)^{-\varphi}.$$

Using the definition of Ω_t, we can rewrite the objective as

$$\max_{\{\mu_{jt}\}} \int_{t=0}^{\infty} e^{-\rho t} \Omega_t \left[\frac{\tilde{W}_{jt}}{P_t} l_{jt} - v(l_{jt}) - \frac{\Psi}{2} L_t \mu_{jt}^2 \right] dt$$

To analyze the union's problem, treat $q_{jt} \equiv \ln \tilde{W}_{jt}$ as the state and μ_{jt} as the control. The Hamiltonian is

$$H = \Omega_t \left[\frac{e^{q_{jt}}}{P_t} L_t \left(\frac{e^{q_{jt}}}{\tilde{W}_t} \right)^{-\varphi} - v \left(L_t \left(\frac{e^{q_{jt}}}{\tilde{W}_t} \right)^{-\varphi} \right) - \frac{\Psi}{2} L_t \mu_{jt}^2 \right] + \lambda_{jt} \mu_{jt},$$

where λ_{jt} is the co-state. The necessary conditions for optimality are

$$\lambda_{jt} = \Psi \Omega_t L_t \mu_{jt}$$

$$\mathrm{d}\lambda_{jt} - \rho \lambda_{jt}\, dt = -(1-\varphi)\Omega_t \frac{\tilde{W}_{jt}}{P_t} l_{jt}\, dt - \varphi \Omega_t v'\left(l_{jt} \right) l_{jt}\, dt.$$

Imposing symmetry and the relationships $P_t = \tilde{W}_t$ and $Y_t = Z_t L_t$ yields the nonlinear Phillips curve

$$d\pi_t = \left[\rho\, dt - \frac{dY_t}{Y_t} - \frac{d\Omega_t}{\Omega_t} \right] \pi_t - \frac{(\varphi - 1)}{\Psi} \left[\frac{\varphi}{\varphi - 1} v'(L_t) - 1 \right] dt.$$

Linearizing around a zero inflation steady state in which $\bar{L} = v'^{-1}\left(\frac{\varphi-1}{\varphi}\right)$ and $\bar{Y}_t = Z_t\bar{L}$ yields

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

where $1/\eta$ is the Frisch elasticity. Letting $\kappa = \frac{(\varphi-1)\eta}{\Psi}$ gives (9).

F.2   Market Clearing

Nondurables 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:

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

G   Data Filtering Using the MA Representation

In this appendix we implement a restricted version of the Kalman filter to recover aggregate shocks. We impose three 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 this 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. Third, we assume that there are at least as many states as there are observables. Under these restrictions, the Kalman smoother coincides with the Kalman filter.

The Filtering Algorithm. Consider a dynamic system with a state space representation

$$X_t = AX_{t-1} + B\epsilon_t \tag{17}$$

$$Y_t = CX_t \tag{18}$$

where X is the state, ϵ is a vector of 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 will require that CB is invertible, which requires that there are at least as many states as there are observables.

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 ith 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 ith dimension. $R(\tau, i)$ is a $N \times 1$ vector. Let $R(\tau)$ be a $N \times N$ matrix where the ith 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 $A^t \to 0$ as $t \to \infty$, 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 (17) and (18) 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 (17) we have $Q_t = \sum_{\tau=0}^{t-1} R(t-\tau)\epsilon_\tau + S(t)$. From (17) and (18) we then have

$$\epsilon_t = R(0)^{-1}(Y_t - Q_t). \tag{19}$$

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.

59

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

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

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

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

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

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

in the notation of our filtering algorithm, which is the same as (19).

Incorporating the ELB. Let $\mathcal{R}$ be a $T \times T$ matrix that maps a path of output gaps into a path of real interest rates that satisfy the monetary rule. The Phillips curve and monetary policy response to inflation is embedded inside $\mathcal{R}$. With the ELB, we have

$$\overrightarrow{r}_t = \max\left\{\mathcal{R}\left(\mathcal{M}\overrightarrow{r} + \mathcal{Q}\epsilon_t + \overrightarrow{\hat{Y}}_{t|t-1}\right) + S(\eta_{t-1}, \epsilon_t), \underline{r}\right\}, \tag{20}$$

where $\overrightarrow{r}_t \equiv (r_t, \mathbb{E}_t r_{t+1}, \cdots, \mathbb{E}_t r_{t+T-1})'$, $\mathcal{Q}$ maps the current shock to a path of output gaps under constant real rates, $\overrightarrow{\hat{Y}}_{t|t-1}$ is the forecast of output gaps given past shocks and monetary news, and S is a function that captures the effect of the exogenous term in the interest rate rule. Given the $N \times 1$ data vector Y_t, we solve a system of $N + T$ equations such that $Y_t = CB\epsilon_t + Q_t$ as above and such that (20) holds where the unknowns are the elements of ϵ_t and $\overrightarrow{r}_t$. We solve this system iteratively using partial updates.

60

H Derivation and Decomposition of r^*

Derivation of Equation (14). Consider an abstract representation of our model expressed in discrete time steps corresponding to the time intervals on which we compute the model:

$$\hat{Y}_t = \mathcal{Y}(h_t, \Phi_t)$$

$$h_t = \mathcal{H}(\vec{r}_t, \eta_t)$$

$$\Phi_{t+1} = \mathcal{T}(\Phi_t, h_t)$$

$$\eta_{t+1} = \mathcal{F}(\eta_t, \epsilon_{t+1}^\eta).$$

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 nondurables, 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, $\vec{r}_t \equiv (r_t, \mathbb{E}_t r_{t+1}, \cdots)'$, and the exogenous aggregate states, $\eta_t \equiv (g_t, G_t - \bar{G}, r_t^b - \bar{r}_t^b)'$. We extend the analysis to allow for prices other than real interest rates to affect the decision rules below, 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 where ϵ_{t+1}^η is the vector of innovations to the aggregate stochastic processes, which are uncorrelated.[30]

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 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

---

[30]For example, the first element is $\sigma_Z \left(\mathcal{W}_{t+1}^Z - \mathcal{W}_t^Z \right)$.

durables.

We linearize the system around steady state:

$$\hat{Y}_t = \mathcal{Y}_h h_t + \mathcal{Y}_\Phi (\Phi_t - \bar{\Phi})$$

$$h_t = \mathcal{H}_r \vec{r}_t + \mathcal{H}_\eta \eta_t$$

$$\Phi_{t+1} - \bar{\Phi} = \mathcal{T}_\Phi (\Phi_t - \bar{\Phi}) + \mathcal{T}_h h_t$$

$$\eta_{t+1} = \mathcal{F}_\eta \eta_t + \epsilon_{t+1}^\eta.$$

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 $\vec{r}_t \equiv (r_t, \mathbb{E}_t r_{t+1}, \mathbb{E}_t r_{t+2}, \cdots)$ 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

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

$$+ \left[\mathcal{Y}_h \mathcal{H}_\eta \mathcal{F}_\eta^s + \sum_{k=0}^{s-1} \mathcal{Y}_\Phi \mathcal{T}_\Phi^{s-k-1} \mathcal{T}_h \mathcal{H}_\eta \mathcal{F}_\eta^k \right] \eta_t + \left[\mathcal{Y}_\Phi \mathcal{T}_\Phi^s \right] (\Phi_t - \bar{\Phi}).$$

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 (14) with the terms in square brackets forming the rows of $\mathcal{M}$, $\mathcal{Q}$, and $\mathcal{D}$, respectively.

Derivation of Equation (16). As shown in the text, the solution for r^* for a given set of states η_r, Φ_t is,

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

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,

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

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,

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

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,

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

$$+ \sum_{k=0}^{s-1}(\mathcal{M}_{[1+s-k..,1+s-k..]} - \mathcal{M}_{[s-k..,s-k..]})\mathbb{E}_k\vec{r}_s$$

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,

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

$$= \sum_{k=0}^{t-1}\mathcal{Q}_{[2+k..,..]}\eta_{t-1-k} - \sum_{k=0}^{t-1}\mathcal{Q}_{[1+k,..]}\mathcal{F}_\eta\eta_{t-1-k} = -\mathcal{Q}\mathcal{F}_\eta\eta_{t-1} + \sum_{k=0}^{t-1}\mathcal{Q}_{[2+k..,..]}\epsilon^\eta_{t-1-k}$$

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

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

$$- \mathcal{M}^{-1}\sum_{k=0}^{t-1}(\mathcal{M}_{[1+t-k..,1+t-k..]} - \mathcal{M}_{[t-k..,t-k..]})\mathbb{E}_k\vec{r}_t - \mathcal{M}^{-1}\sum_{k=0}^{t-1}\mathcal{Q}_{[1+k,..]}\epsilon^\eta_{t-k}$$

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

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,

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

$$- \mathcal{M}^{-1} \sum_{k=0}^{t-1} (\mathcal{M}_{[1+t-k..,1+t-k..]} - \mathcal{M}_{[t-k..,t-k..]}) \mathbb{E}_k \vec{r}_t^* - \mathcal{M}^{-1} \sum_{k=0}^{t-1} \mathcal{Q}_{[1+k,..]} \epsilon_{t-k}^{\eta}$$

Taking expectations of this r^* vector yields

$$\mathcal{M}_{[1+s..,1+s..]} \mathbb{E}_{t-s} \vec{r}_t^* = -\sum_{k=0}^{t-s} \mathcal{M}_{[1+t-k..,1..t-k]} \left[\mathbb{E}_k \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} - \mathbb{E}_{k-1} \begin{pmatrix} r_k \\ \vdots \\ r_{t-1} \end{pmatrix} \right]$$

$$- \sum_{k=0}^{t-s-1} (\mathcal{M}_{[1+t-k..,1+t-k..]} - \mathcal{M}_{[t-k..,t-k..]}) \mathbb{E}_k \vec{r}_t^* - \sum_{k=s}^{t} \mathcal{Q}_{[1+k,..]} \epsilon_{t-k}^{\eta}$$

We can now write expectations recursively,

$$\mathbb{E}_{t-s} \vec{r}_t^* = \mathbb{E}_{t-s-1} \vec{r}_t^* - \mathcal{M}_{[1+s..,1+s..]}^{-1} \mathcal{M}_{[1+s..,1..s]} \left[\mathbb{E}_{t-s} \begin{pmatrix} r_{t-s} \\ \vdots \\ r_{t-1} \end{pmatrix} - \mathbb{E}_{t-s-1} \begin{pmatrix} r_{t-s} \\ \vdots \\ r_{t-1} \end{pmatrix} \right]$$

$$- \mathcal{M}_{[1+s..,1+s..]}^{-1} \mathcal{Q}_{[1+s,..]} \epsilon_{t-s}^{\eta}$$

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

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

$$- \sum_{k=0}^{t} \mathcal{M}_{[1+k..,1+k..]}^{-1} \mathcal{Q}_{[1+k,..]} \epsilon_{t-k}^{\eta}$$

Extension with More Endogenous Prices. Consider the expanded system:

$$\hat{Y}_t = \mathcal{Y}(h_t, \Phi_t)$$

$$h_t = \mathcal{H}(\vec{r}_t, \vec{w}_t, \eta_t)$$

$$\Phi_{t+1} = \mathcal{T}(\Phi_t, h_t)$$

$$\eta_{t+1} = \mathcal{F}(\eta_t, \epsilon_{t+1}^{\eta})$$

$$0 = \mathcal{P}(h_t, \Phi_t),$$

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

$$0 = \left[\mathcal{P}_\Phi \sum_{k=0}^{s-1} \left(\mathcal{T}_\Phi^k \mathcal{T}_h \mathcal{H}_r S_{s-k-1} \right) + \mathcal{P}_h \mathcal{H}_r S_s \right] \vec{r}_t + \left[\mathcal{P}_\Phi \sum_{k=0}^{s-1} \left(\mathcal{T}_\Phi^k \mathcal{T}_h \mathcal{H}_w S_{s-k-1} \right) + \mathcal{P}_h \mathcal{H}_w S_s \right] \vec{w}_t$$

$$+ \left[\mathcal{P}_\Phi \sum_{k=0}^{s-1} \left(\mathcal{T}_\Phi^k \mathcal{T}_h \mathcal{H}_\eta \rho_\eta^{s-k-1} \right) + \mathcal{P}_h \mathcal{H}_\eta \rho_\eta^s \right] \eta_t + \mathcal{P}_\Phi \mathcal{T}_\Phi^s (\Phi_t - \bar{\Phi})$$

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 $\vec{w}_t$

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

Forecasting the output gap as before, we arrive at an analogous expression to our simpler

case without endogenous prices:

$$\mathbb{E}_t \hat{Y}_{t+s} = \left[\mathcal{Y}'_\Phi \sum_{k=0}^{s-1} \left(\mathcal{T}'^k_\Phi \mathcal{T}_h \mathcal{H}'_r S_{s-k-1} \right) + \mathcal{Y}_h \mathcal{H}'_r S_s \right] \vec{r}_t$$

$$+ \left[\mathcal{Y}'_\Phi \sum_{k=0}^{s-1} \left(\mathcal{T}'^k_\Phi \mathcal{T}_h \mathcal{H}'_\eta \mathcal{F}^{s-k-1}_\eta \right) + \mathcal{Y}_h \mathcal{H}'_\eta \mathcal{F}^s_\eta \right] \eta_t + \mathcal{Y}'_\Phi \mathcal{T}'^s_\Phi (\Phi_t - \bar{\bar{\Phi}}),$$

where we have redefined the matrices as follows:

$$\mathcal{Y}'_\Phi = \mathcal{Y}_\Phi - \mathcal{Y}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P$$

$$\mathcal{T}'_\Phi = \mathcal{T}_\Phi - \mathcal{T}_h \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{D}_P$$

$$\mathcal{H}'_r = [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{Q}_P] \mathcal{H}_r$$

$$\mathcal{H}'_\eta = [I - \mathcal{H}_w \mathcal{N}_P^{-1} \mathcal{Q}_P] \mathcal{H}_\eta.$$

Stacking these equations for $s \geq 0$ yields

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

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

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

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.

I Model Impulse Response Functions

Figure A.4 plots the model impulse response functions.

J Robustness to Aggregate Nonlinearities

In this appendix we investigate the sensitivity of our results to allowing for nonlinear aggregate dynamics. We do so by conducting a version of our analysis in the fully nonlinear
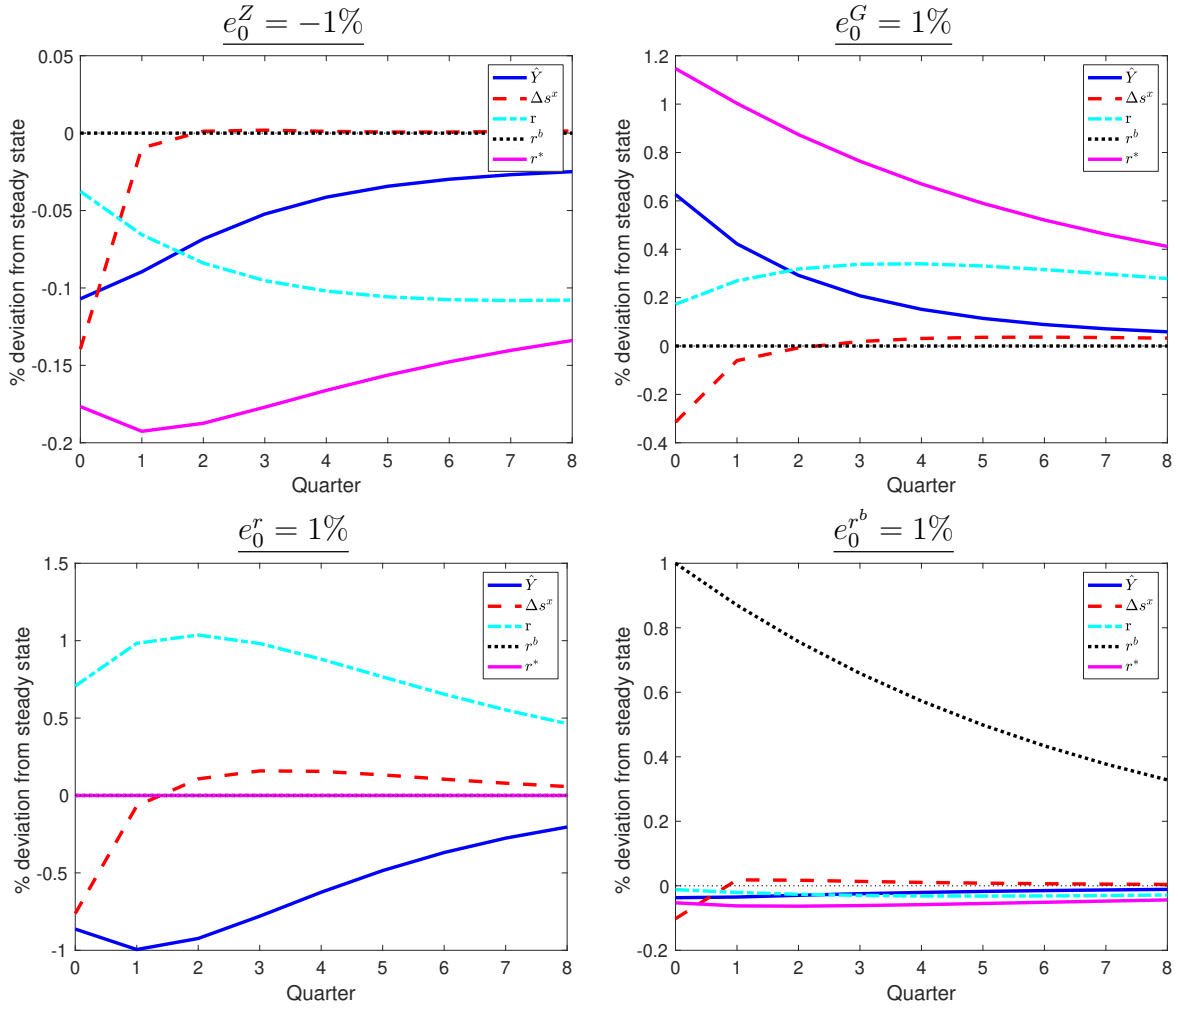
Figure A.4: Impulse response functions for the output gap $\hat{Y}$, the change in the durable expenditure share relative to potential GDP Δs^x, the real interest rate r, the borrowing spread r^b, and the contemporaneous natural rate of interest r^* following a shock to productivity e^Z, non-household demand e^G, the monetary policy rule e^r, and the borrowing spread e^{r^b}.

version of our model. Solving the nonlinear model is substantially more difficult than solving the linear model and moreover our filtering approach relies on linearity. Therefore, in this robustness analysis we conduct a somewhat simpler exercise and we use the full-information version of the model to ease the computational burden. In particular, we perform our inference procedure on a one-time shock. We proceed in the following steps:

1. Feed in a one-time permanent reduction of productivity of -16% into the nonlinear model and find the paths for the real interest rate, the relative durable price, and aggregate income that are consistent with market clearing. We obtain a 2.8% output gap and an 78% percent drop in the endogenous adjustment probability on impact.

2. Use the *linear* model to filter the model generated data and find the implied path of r^* as we do in the main text.

3. Feed the r^* path from the previous step into the nonlinear model and find the paths of the relative durable price and aggregate income that are consistent with market clearing.

Denote the vector of output gaps from step 3 as $\hat{Y}^{NL}$. If the output gap from step 3 is close to zero, then our procedure for inferring r^* using the linear approximation to the model is accurate because it is indeed the path for interest rates that is needed to close the output gap, which is the definition of r^*. To gauge how important the error is for our calculations we use the linear model to convert the residual output gap to an adjustment to the r^* path of $\mathcal{M}^{-1}\hat{Y}^{NL}$. Pre-multiplying by $\mathcal{M}^{-1}$ gives the change in real interest rates that would be needed to close the residual output gap.

Figure A.5 plots r^* and $r^* + \mathcal{M}^{-1}\hat{Y}^{NL}$. Because monetary policy is less powerful in the recession, the linear model initially underestimates how much the real interest rate needs to fall to close the output gap. However, after 10 quarters the two real rate paths follow a very similar pattern. This suggests our baseline analysis likely underestimates the drop in r^* during the Great Recession, but the forecast of persistently low interest rates is robust to state dependence.

There are two reasons that our results are fairly robust to the state dependence implied by the nonlinear model. First, while monetary policy is indeed less powerful during the recession
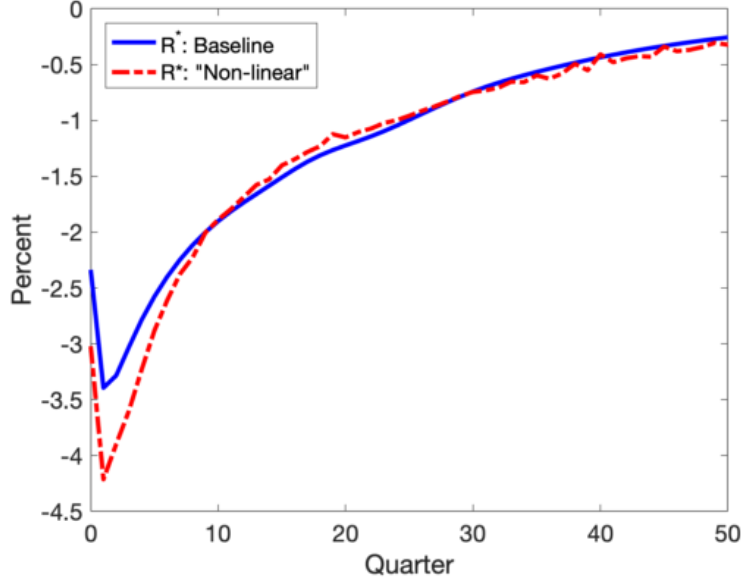
Figure A.5: Path for r^* following a permanent drop in productivity in the full information model. The blue line depicts the r^* path when we use the linear model for filtering the impulse response functions and calculating the corresponding r^*. We feed this path into the nonlinear model, calculate the residual output gap $\hat{Y}^{NL}$, and convert it into r^* space using $\mathcal{M}^{-1}\hat{Y}^{NL}$. The red line depicts the r^* from the linear model plus this residual, $\mathcal{M}^{-1}\hat{Y}^{NL}$.

in the nonlinear model, this effect dissipates rather quickly. Second, even if monetary policy is persistently less powerful there are two offsetting effects on our calculation of r^*. When we filter the data we use the observed movement in r when we infer what shock hit the economy. If we overestimate the power of monetary policy we overestimate the size of the shock the central bank is reacting to. However, the movement in interest rates that is needed to counter a given shock is underestimated. These two considerations partially offset each other.