Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk*

Christian Bayer[†], Ralph Luetticke[‡], Lien Pham-Dao[§] and Volker Tjaden[¶]

May 22, 2017

Abstract

Households face large income uncertainty that varies substantially over the business cycle. We examine the macroeconomic consequences of these variations in a model with incomplete markets, liquid and illiquid assets, and a nominal rigidity. Heightened uncertainty depresses aggregate demand as households respond by hoarding liquid "paper" assets for precautionary motives, thereby reducing both illiquid physical investment and consumption demand. This translates into output losses. We document the empirical response of portfolio liquidity and aggregate activity to surprise changes in idiosyncratic income uncertainty and find both to be quantitatively in line with our model. The welfare consequences of uncertainty shocks and of the policy response thereto depend crucially on a household's asset position.

Keywords: Incomplete Markets, Nominal Rigidities, Uncertainty Shocks.

JEL-Codes: E22, E12, E32

*The paper was circulated before as "Household Income Risk, Nominal Frictions, and Incomplete Markets." We would like to thank the editor and four anonymous referees, Thomas Hintermaier, Andreas Kleiner, Keith Küster, Alexander Kriwoluzky and seminar participants in Bonn, Birmingham, Hamburg, Madrid, Mannheim, Konstanz, the EEA-ESEM 2013 in Gothenburg, the SED meetings 2013 in Seoul, the SPP meeting in Mannheim, the Padova Macro Workshop, the 18th Workshop on Dynamic Macro in Vigo, the NASM 2013 in L.A., the VfS 2013 meetings in Düsseldorf, the 2014 Konstanz Seminar on Monetary Theory and Monetary Policy, the ASSET 2014 in Aix-en-Provence and SITE 2015 for helpful comments and suggestions. Special thanks go to Moritz Kuhn, Ulrike Steins and Ariel Mecikovsky for helping us with the merging of SCF and SIPP waves, respectively. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FTP/2007-2013) / ERC Grant agreement no. 282740. The views expressed in this paper do not reflect the Bundesbank's views or the views of the Eurosystem as a whole. The replication files for this paper are available at http://www.erc-hump-uni-bonn.net/code-and-data/.

†Department of Economics, Universität Bonn. Address: Adenauerallee 24-42, 53113 Bonn, Germany. E-mail: christian.bayer@uni-bonn.de (corresponding author).

‡Department of Economics, University College London, 30 Gordon Street, London, WC1H 0AX, UK.

§Deutsche Bundesbank, Research Data and Service Centre, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main, Germany

¶Blue Yonder GmbH, Ohiostraße 8, 76149 Karlsruhe, Germany.

1 Introduction

The Great Recession has brought about a reconsideration of the role of uncertainty in business cycles. Increased uncertainty has been documented and studied in various markets. However, uncertainty with respect to household income stands out in its size and importance. Shocks to household income are large and exhibit systematic changes over the business cycle. The seminal work by Storesletten et al. (2001) estimates that for the U.S., the variance of persistent shocks to disposable household income almost doubles in recessions relative to expansions.[1]

The starting point of this paper is that such swings in the riskiness of household income lead to systematic variations in the propensity to consume and to a rebalancing of household portfolios if asset markets are incomplete and assets differ in their liquidity. In such a setup, households use precautionary savings, and structure their portfolios to smooth consumption. We quantify the aggregate consequences of precautionary savings and portfolio turnovers in response to shocks to household income risk by means of a dynamic stochastic general equilibrium model. In our model, households have access to two types of assets to smooth consumption. They can either hold liquid (low return) nominal bonds or invest in illiquid, high-dividend-paying physical capital. This asset structure allows us to disentangle savings and physical investment and obtain fluctuations in aggregate demand.[2] To generate aggregate output effects from these fluctuations, we augment this incomplete markets framework in the tradition of Bewley (1980), Huggett (1993), and Aiyagari (1994) by sticky prices à la Rotemberg (1982).

We model the illiquidity of physical capital by a transaction cost. Consequently, households trade capital only from time to time, very similar to Kaplan and Violante (2014) and Kaplan et al. (2016), who follow the tradition of Baumol (1952) and Tobin (1956) in modeling the portfolio choice between liquid and illiquid assets.

In this economy, when idiosyncratic income risk increases, individually optimal asset holdings rise and consumption demand declines. Importantly, households also rebalance their portfolios toward the liquid asset because it provides better consumption smoothing. They take into account that tapping the illiquid asset more often, as increased income risk is realized, creates costs that drive down the effective return (after trading costs) of the illiquid asset. This flight to liquidity is reminiscent of the observed patterns of the share of liquid assets in the portfolios of U.S. households during the Great Recession

[1] Work by Guvenen et al. (2014b) documents that changes in individual labor income become left skewed in recessions.

[2] In a standard Aiyagari (1994) economy, where all savings are in physical capital, an increase in savings does not lead to a fall in total demand (investment plus consumption) because savings increase investments one-to-one.

Figure 1: Liquid assets relative to illiquid assets on household balance sheets

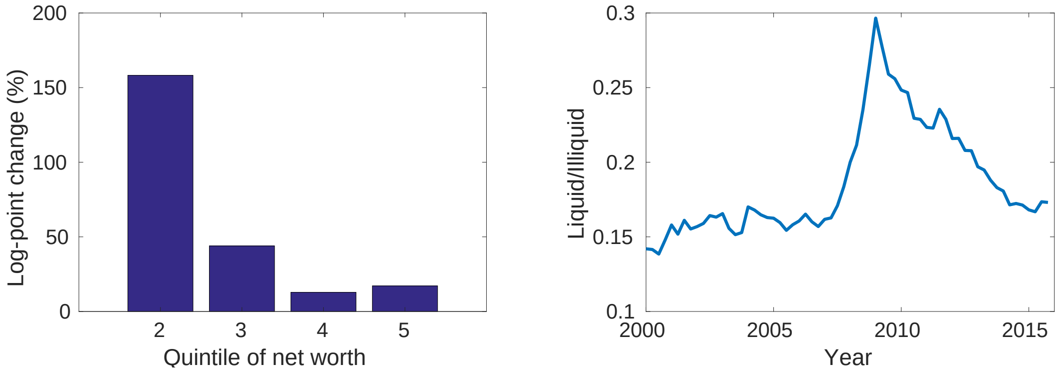

(a) Cross-section: 2010 vs. 2007 (b) Aggregate: 2000 - 2016

Notes: Net liquid assets are composed of money market, checking, savings and call accounts, as well as corporate and government bonds and T-bills net of credit card debt. All other assets net of all other debt make up illiquid wealth.

(a) Change between 2007 and 2010 by quintile of the wealth distribution. Based on the Survey of Consumer Finances. Change in the average ratio of liquid to illiquid assets within the four last quintiles of net worth; see Section 5 for details. The first quintile does not consistently hold positive amounts of liquid and/or illiquid assets.

(b) Ratio of liquid to illiquid assets on household balance sheets based on quarterly data from the Flow of Funds; for details, see Appendix E.1.

(see Figure 1). According to the 2010 Survey of Consumer Finances, the share of liquid assets in household portfolios increased relative to 2007 across all wealth quintiles, with the strongest relative increase for the lower middle class (see Figure 1 (a)). Also in the aggregate, we see a substantial increase in portfolio liquidity around the crisis (see Figure 1 (b)). This increase in portfolio liquidity, we find to be a general response to estimated shocks to idiosyncratic income uncertainty. In our model, this portfolio rebalancing toward liquid paper reinforces, through a reduction in physical investment, the decline in consumption demand caused by higher uncertainty. Consequently, aggregate demand declines even more strongly than consumption, and investment and consumption co-move.

Quantitatively, we find the following: a one standard deviation increase in household income risk decreases aggregate activity by 0.2% and investment by 0.6%, both in the data and in our model. This is about the effect that Fernández-Villaverde et al. (2015) report for a two standard deviation fiscal-policy-uncertainty shock. At the zero lower bound, when neither monetary nor fiscal policy stabilizes the economy, the output loss is almost 2%. In addition to the aggregate consequences, an uncertainty shock has rich distributional consequences, as the price of and return on capital fall more than the return on liquid assets upon an uncertainty shock. We use our model to estimate the welfare consequences of these distributional effects. Our welfare calculations imply that households rich in illiquid physical capital lose the most as capital returns fall strongly in times of high income risk. At the same time, their large but illiquid wealth helps little to smooth consumption. Households rich in liquid assets, by contrast, even though they might hold less total wealth, are much better insured and do not suffer from lower capital returns as much; hence, their welfare losses are smaller.

Our model allows us to assess the importance of systematic monetary and fiscal policy for the stabilization of the economy in response to uncertainty shocks. Aggressive monetary policy can stabilize the economy by cutting interest rates on liquid assets and pushing household portfolios back toward illiquid investments. Expansionary fiscal policy instead supplies the economy with the additional liquid assets demanded by the private sector. Thus, both policies can be used effectively for aggregate stabilization.

Yet, they have different welfare consequences. To understand the consequences of various systematic policy responses, we compare three regimes: first, a regime that corresponds to our baseline calibration of fiscal and monetary policy; second, a regime with perfect stabilization through monetary policy; and third, a regime in which fiscal policy perfectly stabilizes. We find that a one standard deviation increase in household income risk depresses welfare equivalent to 25 basis points of lifetime consumption on average. However, there is a large heterogeneity. Well-insured, wealthy households suffer substantially less from the increase in uncertainty. For them, the equilibrium changes in prices are more important. Therefore, households rich in nominal assets suffer from stabilizing monetary policy as it drives down their asset returns. For the same reason, households rich in real assets like stabilization through fiscal policy. It crowds out investment and keeps capital returns high.

The remainder of the paper is organized as follows. Section 2 starts off with a review of the related literature. Section 3 develops our model, and Section 4 discusses the solution method. Section 5 introduces our estimation strategy for the income process and explains the calibration of the model. Section 6 presents the numerical results. Section 7 provides empirical evidence for an increase in the liquidity of household portfolios after uncertainty shocks. Section 8 concludes. An appendix provides details on the properties of the value and policy functions, the numerics, the estimation of the uncertainty process from income data, and further robustness checks.

2 Related Literature

Our paper contributes to the recent literature that explores empirically and theoretically the aggregate effects of time-varying uncertainty. The seminal paper by Bloom (2009) discusses the effects of time-varying (idiosyncratic) productivity uncertainty on firms' factor demand, exploring the idea and effects of time-varying real option values of investment. This paper has triggered a stream of research that explores under which conditions such variations have aggregate effects.[3]

A more recent branch of this literature investigates the aggregate implications of uncertainty shocks beyond their transmission through investment and has also broadened the sources of uncertainty studied. The first papers in this vein highlight nonlinearities in the New Keynesian model, in particular the role of precautionary price setting.[4] Fernández-Villaverde et al. (2015), for example, look at a medium-scale DSGE model à la Smets and Wouters (2007). They find that at the zero lower bound output drops by more than 1.5% after a two standard deviation shock to the volatility of taxes if a countervailing fiscal policy response is ruled out. Off the ZLB the drop reduces to 0.2%.[5] In a similar framework, Basu and Bundick (2012) highlight the labor market's response to uncertainty about aggregate TFP and time preferences. They argue that, if uncertainty increases, the representative household will want to save more and consume less. Then, with King et al. (1988) preferences, the representative household will also supply more labor, which in a New Keynesian model depresses output through a "paradox of toil." When labor supply increases, wages and hence marginal costs for firms fall. This increases markups when prices are sticky, which finally depresses demand for consumption and investment, and a recession follows. Overall, they find aggregate effects similar to those in Fernández-Villaverde et al. (2015), in particular at the zero lower bound.

While our paper also focuses on precautionary savings, it differs substantially in the transmission channel. We are agnostic about the importance of the "paradox of toil," because it crucially relies on a wealth effect in labor supply. We therefore assume Greenwood et al. (1988) preferences to eliminate any direct impact of uncertainty on

[3] To name a few: Arellano et al. (2012), Bachmann and Bayer (2013), Christiano et al. (2010), Chugh (2012), Di Tella (2016), Gilchrist et al. (2014), Narita (2011), Panousi and Papanikolaou (2012), Schaal (2012), and Vavra (2014) have studied the business cycle implications of a time-varying dispersion of firm-specific variables, often interpreted as and used to calibrate shocks to firm risk, propagated through various frictions: wait-and-see effects from capital adjustment frictions, financial frictions, search frictions in the labor market, nominal rigidities, balance sheets, and agency problems.

[4] With sticky prices, firms target a higher markup the more uncertain the future aggregate price level.

[5] Born and Pfeifer (2014) report an output drop of 0.025% for a similar model and policy risk shock under a slightly different calibration. Regarding TFP risk, they find hardly any aggregate effect.

labor supply to isolate the demand channel of precautionary savings instead.[6] Moreover, since we focus on *idiosyncratic* income uncertainty, we can identify the uncertainty process outside the model from the Survey of Income and Program Participation (SIPP).

This focus on household income risk and the response of precautionary savings links our paper to the burgeoning literature on heterogeneous agent New-Keynesian models, in particular to Ravn and Sterk (2013) and Den Haan et al. (2015). Both highlight the importance of idiosyncratic unemployment risk. In their setups, households face unemployment risk in an incomplete markets model with labor market search and nominal frictions. Both papers differ from ours in the asset market setup and the shocks considered. Ravn and Sterk (2013) look at a setup with government bonds as a means of savings. They then study a joint shock to job separations and the share of long-term unemployed. This increases income risk and hence depresses aggregate demand because of higher precautionary savings. They find that such first moment shocks to the labor market can be significantly propagated and amplified through this mechanism.

Den Haan et al. (2015) consider a model with money and equity instead, where equity is not physical capital as in our model, but is equated with vacancy-ownership. In addition, they assume wage rigidity. As in our model, poorer households, in their model the unemployed, are the marginal holders of money, the low-return asset, as they effectively discount the future more. When unemployment goes up, demand for money increases. This in turn leads to deflation, pushing up real wages because nominal wages are assumed to be sticky. This has a second-round effect on money demand. Because the labor intensity of production cannot be adjusted, higher real wages depress the equity yield on existing and newly formed vacancies, which then induces portfolio adjustments by households toward money, amplifying the deflation and the related output drop.

To some extent, our transmission mechanism shares this feature but additionally highlights the importance of liquidity. Households increase their precautionary savings in conjunction with a portfolio adjustment toward the liquid asset, because its services in consumption smoothing become more valuable to households. We find that the liquidity effect is important in our model in which the labor intensity of production can be adjusted.

With respect to the broader literature on New-Keynesian incomplete markets models we share with Gornemann et al. (2012) the focus on the distributional consequences of systematic monetary and fiscal policy responses (in our case to uncertainty shocks). While they look at labor markets, we focus on household portfolios and find them to

[6] Similarly, in a search model, higher uncertainty about match quality might translate into longer search and more endogenous separation. A priori the labor supply response is hence ambiguous.

be important for the distributional consequences. We share this portfolio focus in an incomplete markets economy with sticky prices with Kaplan et al. (2016), who discuss the transmission of monetary policy,[7] and with Guerrieri and Lorenzoni (2011), who model the effect of a credit crunch.[8]

3 Model

We model an economy composed of a firm sector, a household sector and a government sector. Firms are either perfectly competitive intermediate goods producers or final goods producers that face monopolistic competition, producing differentiated final goods out of homogeneous intermediate inputs. Price setting for these goods is subject to a pricing friction à la Rotemberg (1982). Households supply labor and capital, and own all final goods producers, absorbing their rents. The government sector runs both a fiscal authority and a monetary authority. The fiscal authority levies a time-constant labor income tax, emits government bonds and adjusts expenditures both to business cycle conditions and to stabilize debt in the long run. The monetary authority sets the nominal interest rate on government bonds according to a Taylor rule.

The household sector is subdivided into two types of agents: workers and entrepreneurs. The transition between both types is stochastic. Both rent out physical capital. Only workers supply labor. The efficiency of a worker's labor evolves randomly, which exposes the worker-household to labor-income risk. Entrepreneurs do not work, but earn all pure rents that arise in our economy from the convex costs of capital adjustment and, more important, from monopolistic competition in final goods.

All households self-insure against the income risks they face by saving. They can save in the form of a liquid nominal asset (bonds) and a less liquid physical asset (capital). Liquidity is understood in the spirit of Kaplan and Violante's (2014) model of wealthy hand-to-mouth consumers, in which trading illiquid capital is costly. We model this cost as a utility cost that can be considered as a summary of the time, pecuniary, and pure utility costs of liquidating an illiquid asset such as a house, for example.[9] Liquid assets, i.e., bonds in our model economy, on the other hand can be traded without any costs.

Both bonds and capital pay a return. The monetary authority sets the nominal return on bonds. The illiquid asset pays a real dividend that is determined by the rental rate of

[7] Luetticke (2015) builds upon the present paper and in a parallel work discusses the transmission of monetary policy shocks.

[8] Further examples of the New-Keynesian incomplete markets literature are Auclert (2015); Challe and Ragot (2016); McKay et al. (2016); McKay and Reis (2016); Werning (2015), all of which, however, build on a single-asset framework.

[9] Kaplan and Violante (2014) find that physical transaction costs and utility costs yield similar results for the portfolio problem.

capital paid by the intermediate-good-producing sector on a perfectly competitive rental market. This intermediate-good-producing sector combines labor and capital services into intermediate goods and sells them to the entrepreneurs.

3.1 Households

To be more specific, there is a continuum of ex-ante identical households of measure one, indexed by i. Households are infinitely lived, have time-separable preferences with time-discount factor β, and derive felicity from consumption c_{it} and leisure. They obtain income from supplying labor, n_{it}, renting out capital, k_{it}, and interest income on bonds, b_{it}. Whenever a household wants to adjust its holdings of capital, it needs to pay some felicity cost χ_{it} that is an i.i.d. draw from a logistic distribution.[10] Holdings of bonds have to be above an exogenous debt limit $\underline{B}$, and holdings of capital have to be non-negative.

A household's labor income $w_t h_{it} n_{it}$ is composed of the aggregate wage rate, w_t, the household's hours worked, n_{it}, and its idiosyncratic labor productivity, h_{it}, which evolves according to the following process:

$$h_{it} = \frac{\tilde{h}_{it}}{\int \tilde{h}_{it} di} \tag{1}$$

$$\tilde{h}_{it} = \begin{cases} \exp\left(\rho_h \log \tilde{h}_{it-1} + \epsilon_{it}^h\right) & \text{with probability } 1-\zeta \text{ if } h_{it-1} \neq 0 \\ 1 & \text{with probability } \iota \text{ if } h_{it-1} = 0 \\ 0 & \text{else} \end{cases} \tag{2}$$

where scaling individual $\tilde{h}_{it}$ by its cross-sectional average, $\int \tilde{h}_{it} di$, makes sure that average worker productivity is constant. The shocks ϵ_{it}^h to productivity are normally distributed with time-varying variance as given by

$$\begin{aligned} \sigma_{h,t}^2 &= \bar{\sigma}_h^2 \exp s_t \\ s_{t+1} &= \rho_s s_t + \epsilon_t^s \\ \epsilon_t^s &\sim \mathcal{N}\left(-\frac{\sigma_s^2}{2(1+\rho_s)}, \sigma_s^2\right), \end{aligned} \tag{3}$$

i.e., at time t households observe a change in the variance of shocks that drive next period's productivity. In words, we assume that idiosyncratic productivity normally

[10] This assumption on the distribution of adjustment costs yields closed-form solutions for expected adjustment costs given the value of adjustment.

evolves according to a log AR(1) process with time-varying variance.[11] With probability ζ households become entrepreneurs ($h = 0$).[12] With probability ι an entrepreneur returns to the labor force. An entrepreneurial household obtains a fixed share of the pure rents, Π_t, in the economy (from monopolistic competition and creation of capital). We assume that the claim to the pure rent cannot be traded as an asset.

With respect to leisure and consumption, households have Greenwood-Hercowitz-Huffman (GHH) preferences and maximize the discounted sum of felicity:

$$E_0 \max_{\{c_{it}, n_{it}, \Delta k_{it}\}} \sum_{t=0}^{\infty} \beta^t u\left[c_{it} - h_{it} G(n_{it})\right] - \mathbb{I}_{\Delta k_{it} \neq 0} \chi_{it}, \tag{4}$$

where χ_{it} is the utility cost of adjustment and $\mathbb{I}_{\Delta k_{it} \neq 0}$ is an indicator function that takes value one if a household adjusts its holdings of physical capital and zero otherwise. The maximization is subject to the budget constraints described further below. The felicity function u takes the form of a constant relative risk aversion (CRRA) with risk aversion parameter $\xi > 0$,

$$u(x_{it}) = \frac{1}{1-\xi} x_{it}^{1-\xi},$$

where $x_{it} = c_{it} - h_{it} G(n_{it})$ is household i's composite demand for goods consumption c_{it} and leisure. Goods consumption bundles varieties j of differentiated goods according to a Dixit-Stiglitz aggregator:

$$c_{it} = \left(\int c_{ijt}^{\frac{\eta-1}{\eta}} dj \right)^{\frac{\eta}{\eta-1}}.$$

Each of these differentiated goods is offered at price p_{jt}, so that for the aggregate price

[11] We assume that uncertainty fluctuations are exogenous partly for analytical clarity. It is likely that some of the fluctuations in uncertainty we observe in the data reflect endogenous responses such as increases in unemployment.

[12] Note that the notation is somewhat counterintuitive: the no-labor state will be a high income state. The assumption of how to allocate pure rents is borrowed from Romei (2015). Attaching the rents in the economy to an exogenously determined group of households instead of distributing it with the factor incomes for capital or labor has the advantage that the factor prices and thus factor supply decisions remain the same as in any standard New-Keynesian framework. The assumption of stochastic transitions to the entrepreneurial state, $h = 0$, can be thought of as a household inventing a new version of a differentiated product that replaces an older existing version of that product (keeping the mass of products constant). We assume that the innovation is drastic such that the old product version disappears and plays no role in price setting. The innovating household then focuses exclusively on the production of this product and can no longer supply any additional labor. The incumbent household returns to the labor force with median productivity, i.e., $\tilde{h}_{it} = 1$.

level, $P_t = \left(\int p_{jt}^{1-\eta} dj\right)^{\frac{1}{1-\eta}}$, the demand for each of the varieties is given by

$$c_{ijt} = \left(\frac{p_{jt}}{P_t}\right)^{-\eta} c_{it}.$$

The disutility of work, $h_{it}G(n_{it})$, determines a household's labor supply given the aggregate wage rate, w_t, and a labor income tax, τ, through the first-order condition:

$$h_{it}G'(n_{it}) = (1-\tau)w_t h_{it}, \tag{5}$$

such that all households supply the same number of hours $n_{it} = N_t$. Total effective labor input, $\int n_{it}h_{it}di$, is hence also equal to N_t because $\int h_{it}di = 1$.

Assuming a constant Frisch elasticity of aggregate labor supply $1/\gamma$,

$$G(N_t) = \frac{1}{1+\gamma}N_t^{1+\gamma}, \quad \gamma > 0,$$

we simplify the expression for the composite consumption good x_{it}:

$$x_{it} = c_{it} - h_{it}G(N_t) = c_{it} - \frac{(1-\tau)w_t h_{it} N_t}{1+\gamma}, \tag{6}$$

where we make use of the constant elasticity assumption that yields:

$$h_{it}G(N_t) = \frac{h_{it}G'(N_t)N_t}{1+\gamma} = \frac{(1-\tau)w_t h_{it} N_t}{1+\gamma}.$$

We weight the disutility of work by productivity h_{it} to simplify the exposition, because changes in the cross-sectional distribution of incomes directly measure changes in productivity. There is no endogenous reaction of hours worked. Thus, shocks to idiosyncratic productivity move income one-for-one and we will – as a shorthand notation – call the risk households face regarding their productivity "income risk" and the shocks to h "income shocks," accordingly. Since the Frisch elasticity is constant, any weighting of the disutility G is irrelevant as long as the distribution of *log incomes* is treated as a state, because the disutility of labor is always a constant fraction of labor income.

Since trading capital is subject to a utility loss we need to consider two versions of the household's budget constraint. When the household participates in the capital market, the budget constraint reads:

$$c_{it} + b_{it+1} + q_t k_{it+1} = b_{it}\frac{1+R(b_{it}, R^b_{t-1})}{\pi_t} + (q_t + r_t)k_{it} + (1-\tau)(w_t h_{it} N_t + \mathbb{I}_{h=0}\Pi_t),$$
$$k_{it+1} \geq 0, b_{it+1} \geq \underline{B},$$

where b_{it} is real bond holdings, $\underline{B}$ is an exogenous borrowing constraint, k_{it} is the amount of illiquid assets, q_t is the price of these assets, r_t is their dividend, R is the nominal interest rate on bonds, which depends on the portfolio position of the household and the central bank's interest rate R^b_{t-1}, which is set one period before, and $\pi_t = \frac{P_t - P_{t-1}}{P_{t-1}}$ is realized inflation.

We assume as in Kaplan et al. (2016) that there is a wasted intermediation cost, $\overline{R}$, when households resort to unsecured borrowing. Therefore, we specify:

$$R(b_{it}, R^b_{t-1}) = \begin{cases} R^b_{t-1} & \text{if } b_{it} \geq 0 \\ R^b_{t-1} + \overline{R} & \text{if } b_{it} < 0. \end{cases} \tag{7}$$

This assumption creates a mass of households with zero unsecured credit but with the possibility to borrow, though at a penalty rate.

Substituting the expression $c_{it} = x_{it} + \frac{(1-\tau)w_t h_{it} N_t}{1+\gamma}$ for consumption, we obtain:

$$x_{it} + b_{it+1} + q_t k_{it+1} = b_{it}\frac{1+R(b_{it}, R^b_{t-1})}{\pi_t} + (q_t + r_t)k_{it} + (1-\tau)\left(\frac{\gamma}{1+\gamma} w_t h_{it} N_t + \mathbb{I}_{h=0}\Pi_t\right),$$
$$k_{it+1} \geq 0, \quad b_{it+1} \geq \underline{B}. \tag{8}$$

All households that decide not to participate in the capital market obtain dividends, but may only adjust their bond holdings. For them, the budget constraint simplifies to:

$$x_{it} + b_{it+1} = b_{it}\frac{1+R(b_{it}, R^b_{t-1})}{\pi_t} + r_t k_{it} + (1-\tau)\left(\frac{\gamma}{1+\gamma} w_t h_{it} N_t + \mathbb{I}_{h=0}\Pi_t\right), \quad b_{it+1} \geq \underline{B}. \tag{9}$$

Note that we assume that households replace depreciated capital through maintenance every period such that the dividend, r_t, is the net return on capital.

Since a household's saving decision will be some non-linear function of that household's wealth and productivity, the price level, P_t, and accordingly aggregate real bond holdings, B_{t+1}, will be functions of the joint distribution Θ_t of (b_t, k_t, h_t). This makes Θ_t a state variable of the household's planning problem. This distribution evolves as a result of the economy's reaction to shocks to uncertainty that we model as in (3).

With this setup, the dynamic planning problem of a household is characterized by: The value function V_a for the case where the household adjusts its capital holdings, the value function V_n for the case in which it does not adjust, and the expected envelope value, EV, over both options:

$$\begin{aligned}
V_a(b,k,h;\Theta,s) =& \max_{k',b'_a} u[x(b,b'_a,k,k',h)] + \beta EV(b'_a,k',h,\Theta',s) \\
V_n(b,k,h;\Theta,s) =& \max_{b'_n} u[x(b,b'_n,k,k,h)] + \beta EV(b'_n,k,h,\Theta',s) \\
EV(b',k',h;\Theta,s) =& E_{\chi',h',s'}\left\{\max\left[V_a(b',k',h';\Theta',s') - \chi', V_n(b',k',h';\Theta',s')\right]\right\}
\end{aligned} \tag{10}$$

Expectations about the continuation value are taken with respect to all stochastic processes (productivity, adjustment costs, and uncertainty), and hence depend on the current stochastic states.

Conditional on paying the adjustment cost, the household will choose a portfolio that trades off the higher liquidity of bonds against the higher return that illiquid assets pay (in equilibrium). The value of liquidity stems from smoother consumption. We denote the optimal consumption policies for the adjustment and non-adjustment cases as x_a^* and x_n^*, the bond holding policies as b_a^* and b_n^*, and the capital investment policy as k^*.

The household will pay the fixed cost to adjust its portfolio if and only if

$$V_a(b',k',h';\Theta',s') - \chi' \geq V_n(b',k',h';\Theta',s'),$$

such that the probability to adjust is given by

$$\nu^*(b',k',h';\Theta',s') := F_\chi\left[V_a(b',k',h';\Theta',s') - V_n(b',k',h';\Theta',s')\right], \tag{11}$$

where F_χ is the cumulative distribution function of χ. We assume this distribution to be logistic, so that the EV term has a closed-form expression given $V_{a,n}$. Details on the properties of the value functions and policy functions (differentiable and increasing in total resources), the first-order conditions, and the algorithm we employ to calculate the policy functions can be found in Appendices A and B.

3.2 Intermediate Goods Producers

Intermediate goods are produced with a constant returns to scale production function:

$$Y_t = N_t^\alpha K_t^{(1-\alpha)}.$$

Let MC_t be the relative price at which the intermediate good is sold to entrepreneurs.

The intermediate-good producer maximizes profits,

$$MC_t Y_t = MC_t N_t^{\alpha} K_t^{(1-\alpha)} - w_t N_t - (r_t + \delta) K_t,$$

but it operates in perfectly competitive markets, such that the real wage and the user costs of capital are given by the marginal products of labor and capital:

$$w_t = \alpha MC_t \left(K_t/N_t\right)^{1-\alpha}, \qquad r_t + \delta = (1-\alpha) MC_t \left(N_t/K_t\right)^{\alpha}. \tag{12}$$

3.3 Price Setting

Final goods producers differentiate the intermediate good and set prices. We assume price adjustment costs á la Rotemberg (1982). For tractability, we assume that price setting is delegated to a mass-zero group of households (managers) that are risk neutral and compensated by a share in profits. They do not participate in any asset market. Under this assumption, price setting is carried out with a time-constant discount factor. Managers maximize the present value of real profits given the demand for good j,

$$y_{jt} = \left(p_{jt}/P_t\right)^{-\eta} Y_t, \tag{13}$$

and quadratic costs of price adjustment, i.e., they maximize:

$$E_0 \sum_{t=0}^{\infty} \beta^t Y_t \left\{ \left(\frac{p_{jt}}{P_t} - MC_t\right) \left(\frac{p_{jt}}{P_t}\right)^{-\eta} - \frac{\eta}{2\kappa} \left(\log \frac{p_{jt}}{p_{jt-1}}\right)^2 \right\}. \tag{14}$$

From the corresponding first-order condition for price setting, it is straightforward to derive the Phillips curve:

$$\log(\pi_t) = \beta E_t \left[\log(\pi_{t+1}) \tfrac{Y_{t+1}}{Y_t}\right] + \kappa \left(MC_t - \tfrac{\eta-1}{\eta}\right), \tag{15}$$

where π_t is the gross inflation rate, $\pi_t := \frac{P_t}{P_{t-1}}$, and MC_t are the real marginal costs. The price adjustment then creates real costs $\frac{\eta}{2\kappa} Y_t \log(\pi_t)^2$.

Since managers are a measure-zero group in the economy, all profits – net of price adjustment costs – go to the entrepreneur households (whose $h = 0$). In addition, these households also obtain profit income from adjusting the aggregate capital stock. They can transform I_t consumption goods into ΔK_{t+1} capital goods (and back) according to the transformation function:

$$I_t = \tfrac{\phi}{2} \left(\Delta K_{t+1}/K_t\right)^2 K_t + \Delta K_{t+1}.$$

Since they are facing perfect competition in this market, entrepreneurs will adjust the stock of capital until the following first-order condition holds:

$$q_t = 1 + \phi \Delta K_{t+1}/K_t. \tag{16}$$

3.4 Government

The government operates a monetary and a fiscal authority. The monetary authority controls the nominal interest rate on liquid assets, while the fiscal authority issues government bonds to finance deficits and adjusts expenditures to stabilize debt in the long run and output in the short run.

We assume that monetary policy sets the nominal interest rate on bonds following a Taylor (1993)-type rule with interest rate smoothing:

$$\frac{R^b_{t+1}}{\bar{R}^b} = \left(\frac{R^b_t}{\bar{R}^b}\right)^{\rho_{R^b}} \left(\frac{\pi_t}{\bar{\pi}}\right)^{(1-\rho_{R^b})\theta_\pi}. \tag{17}$$

The coefficient $\bar{R}^b \geq 0$ determines the nominal interest rate in the steady state. The coefficient $\theta_\pi \geq 0$ governs the extent to which the central bank attempts to stabilize inflation around its steady-state value: the larger θ_π the stronger is the reaction of the central bank to deviations from the inflation target. When $\theta_\pi \to \infty$, inflation is perfectly stabilized at its steady-state value. $\rho_{R^b} \geq 0$ captures interest rate smoothing.

We assume that the government issues bonds according to the rule:

$$\frac{B_{t+1}}{\bar{B}} = \left(\frac{B_t R^b_t/\pi_t}{\bar{B}\bar{R}^b/\bar{\pi}}\right)^{\rho_B} \left(\frac{\pi_t}{\bar{\pi}}\right)^{-\gamma_\pi} \left(\frac{T_t}{\bar{T}}\right)^{-\gamma_T}, \tag{18}$$

using the revenues to finance government expenditures, $G_t = \tau(w_t N_t + \Pi_t) - \frac{R^b_t}{\pi_t} B_t + B_{t+1}$. Hence the primary deficit is given by

$$G_t - T_t = B_t \frac{R^b_t/\pi_t}{\bar{R}^b/\bar{\pi}} \left[\left(\frac{B_t R^b_t/\pi_t}{\bar{B}\bar{R}^b/\bar{\pi}}\right)^{\rho_B - 1} \left(\frac{\pi_t}{\bar{\pi}}\right)^{-\gamma_\pi} \left(\frac{T_t}{\bar{T}}\right)^{-\gamma_T} - \frac{\bar{R}^b}{\bar{\pi}} \right]. \tag{19}$$

The coefficient ρ_B captures whether and how fast the government seeks to repay its outstanding obligations $B_t R^b_t/\pi_t$. For $\rho_B < 1$ the government actively stabilizes government debt, for $\rho_B = 1$ debt will only be stable if the monetary authority violates the Taylor principle ($\theta_\pi < 1$) such that higher inflation reduces the real rate on bonds even in the long run ("fiscal dominance").[13] The coefficient γ_π captures the cyclicality of the deficit:

[13] Note that the price level is determinate in our model when the central bank violates the Taylor principle not only when $\rho_B = 1$ but also for smaller values of ρ_B, because the economy is non-Ricardian

for $\gamma_\pi = 0$ the deficit does not react to inflation (and thereby to output, hence it is acyclical), for $\gamma_\pi > 0$ the deficit is countercyclical, for $\gamma_\pi < 0$ it is procyclical. The parameter γ_T measures the response of the deficit to tax income.

3.5 Goods, Bonds, Capital, and Labor Market Clearing

The labor market clears at the competitive wage given in (12). The bond market clears, whenever the following equation holds:

$$B_{t+1} = B^d(\Theta_t; q_t, \pi_t, R^b_{t+1}) := E\left[\nu^* b^*_a + (1-\nu^*) b^*_n\right], \tag{20}$$

where ν^*, b^*_a, b^*_n are functions of $(b, k, h; q_t, \pi_t, R^b_{t+1})$, and expectations are taken w.r.t. the distribution $\Theta_t(b, k, h)$. Equilibrium requires the total *net* amount of bonds the household sector demands, B^d, to equal the supply of government bonds. In gross terms there are more liquid assets in circulation as some households borrow up to $\underline{B}$.

Last, the market for capital has to clear:

$$\begin{aligned} q_t &= 1 + \phi \frac{K_{t+1} - K_t}{K_t}, \\ K_{t+1} &= K^d(\Theta_t; q_t, \pi_t, R^b_{t+1}) := E[\nu^* k^* + (1-\nu^*) k], \end{aligned} \tag{21}$$

where the first equation stems from competition in the production of capital goods, the second equation defines the aggregate supply of funds from households – both those that trade capital, $\nu^* k^*$, and those that do not, $(1-\nu^*)k$. Again ν^* and k^* are functions of $(b, k, h; q_t, \pi_t, R^b_{t+1})$. The goods market then clears due to Walras' law, whenever labor, bonds, and capital markets clear.

3.6 Recursive Equilibrium

A *recursive equilibrium* in our model is a set of policy functions $\{x^*_a, x^*_n, b^*_a, b^*_n, k^*, \nu^*\}$, value functions $\{V_a, V_n, EV\}$, pricing functions $\{r, w, \pi, q, R^b\}$, aggregate capital and labor supply functions $\{K, N\}$, distributions Θ_t over individual asset holdings and productivity, and a perceived law of motion Γ, such that

1. Given $\{V_a, V_n\}$, Γ, prices, and distributions, the policy functions $\{x^*_a, x^*_n, b^*_a, b^*_n, k^*, \nu^*\}$ solve the households' planning problem, and given the pol-

since households value real government debt because they can use it to smooth consumption. For any given path of the nominal interest rate, there is hence only one path of the inflation rate that clears the bond market.

icy functions $\{x_a^*, x_n^*, b_a^*, b_n^*, k^*, \nu^*\}$, prices and distributions, the value functions $\{V_a, V_n\}$ are a solution to the Bellman equations (10).

2. The labor, the final goods, the bond, the capital and the intermediate good markets clear, and interest rates on bonds are set according to the central bank's Taylor rule, i.e., (12), (15), (20), and (21) hold.

3. The actual and the perceived law of motion Γ coincide, i.e., $\Theta' = \Gamma(\Theta, s')$.

4 Numerical Implementation

The dynamic program (10) and hence the recursive equilibrium is not computable, because it involves the infinite dimensional object Θ_t. We discretize the distribution Θ and represent it by its histogram, a finite dimensional object.

4.1 Solving the Household's Planning Problem

In solving for the household's policy functions, we apply an endogenous gridpoint method as originally developed in Carroll (2006) and extended by Hintermaier and Koeniger (2010), iterating over the first-order conditions. We start with a guess for the adjustment probabilities and use (11) to update the adjustment probabilities until convergence. In each iteration we check for concavity of the value functions and find that the value functions are concave on the entire domain on which we solve them, i.e., we operate a special case of the algorithm suggested by Fella (2014). Details on the algorithm can be found in Appendices A.4 and B.

We approximate the idiosyncratic productivity process by a discrete Markov chain with 26 states. We obtain the time-varying transition probabilities for this Markov chain using the method proposed by Tauchen (1986).[14]

4.2 Aggregate Fluctuations

Even though the histogram is a finitely dimensional object, it is still highly dimensional, which makes it difficult to apply standard techniques to solve for a competitive equilibrium with aggregate risk. We deal with this issue by using (and comparing) two alternative approaches to tackle the problem. Both give similar results.

[14] We solve the household policies for 80 points on the grid for bonds and on the grid for capital using four-times log-scaled grids. We experimented with changing the number of grid points without a noticeable impact on results.

Our baseline approach builds on and extends Reiter (2009) and solves for aggregate dynamics by first-order perturbation around the stationary equilibrium without aggregate shocks. What we add to Reiter's method is to approximate the three-dimensional distribution Θ by a distribution that has a fixed copula and time-varying marginals. To reduce the dimensionality of the value function and its derivatives, we approximate them by a sparse polynomial around their stationary equilibrium solutions.

Alternatively, we solve for a Krusell-Smith equilibrium, assuming households forecast prices by using aggregate bond and capital, B_t and K_t, as state variables along with the lagged interest rate set by the central bank R_t^B, the uncertainty state s_t and $E_t(h_{it})$. The quality of the approximation is reasonable, but since we need to include five aggregate state variables, we cannot use a high resolution for the aggregate states. The result is that the forecasting functions of households only arrive at roughly 99% $\mathbf{R}^2$ statistics. For the Krusell-Smith method, we approximate the stochastic volatility process by Tauchen's method and 5 states for income risk. We use 3 nodes each for aggregate bonds, capital holdings, the lagged interest rates, and $E_t h$. Details on both methods can be found in Appendix C.

5 Calibration

We calibrate the model to the U.S. economy. The aggregate data used for calibration spans 1983 to 2015 (post Volcker-disinflation). One period in the model refers to a quarter of a year. The choice of parameters as summarized in Tables 1 to 4 is explained next. We present the parameters as if they were individually chosen in order to match a specific data moment, but all calibrated parameters are determined jointly of course.

5.1 Household Income Process

We estimate the process for productivity directly from micro data. Given our preference specification, productivity and income follow the same process after controlling for aggregate time effects. We use an approach similar to Storesletten et al. (2004) to estimate variations in income uncertainty. However, we do not preimpose a business cycle relationship. We focus on log household labor income *after taxes and transfers*, because it is the remaining fluctuations in income that households need to self-insure against. In the data, we model income as composed of a transitory, a persistent, a household-fixed and a deterministic component. Income of household i in quarter t is given by:

$$\begin{aligned} \log y_{it} &= f(o_{it}) + \tau_{it} + h_{it} + \mu_i, \\ h_{it} &= \sum_{s=c}^{t} \rho_h^{t-s} \epsilon_{is}^h, \\ \tau_{it} &= \epsilon_{it}^\tau + \rho_\tau \epsilon_{it-1}^\tau, \\ \epsilon_{it}^\tau \sim \mathcal{N}(0, \sigma_\tau^2), &\qquad \epsilon_{it}^h \sim \mathcal{N}(0, \sigma_{h,t}^2), \qquad \mu_i \sim \mathcal{N}(0, \sigma_\mu^2), \end{aligned} \tag{22}$$

where c is the quarter of labor market entry, $f(o_{it})$ measures the effect of observable household characteristics o_{it}, τ_{it} is an MA(1) transitory shock or measurement error, and μ_i is an individual fixed effect. Regarding the variances, we assume that the variances of transitory shocks and fixed effects are time constant while the variance $\sigma_{h,t}^2$ of the persistent component, h_{it}, evolves slowly according to an AR(1) process as in (3) after controlling for a linear-quadratic time trend.

We use data from the Survey of Income and Program Participation (SIPP) from 1984 to 2013 for households whose head is at least 30 and below 56 years of age to identify the parameters of the income process and the uncertainty shocks, ϵ_t^s. From the monthly and individual SIPP data, we generate household labor income (husband and wife) after taxes and transfers on a quarterly basis (imputing taxes using TAXSSIM), remove the explainable part $f(o_{it})$ based on education and age information, and calculate variances and the first two autocovariances at quarterly frequency for a cell of households defined by age and time (both in quarters). This way, our empirical strategy follows Storesletten et al. (2004), i.e., we fit our statistical model to the observable autocovariances of income for an age-time cell. Different to their approach, we set up a quasi-maximum likelihood estimator that takes into account the sampling error of the variance-autocovariance estimates for the age-year cells obtained from bootstrapping, and estimate the series of uncertainty shocks that maximize the model's likelihood. In other words, we select shocks, parameters of the income process, and the variances, $(\rho_h, \rho_\tau, \rho_s, \bar{\sigma}_p, \sigma_\tau, \sigma_\mu, \sigma_s)$, to minimize both the deviations of theoretical from empirical variances/autocovariances and structural shocks, (ϵ_t). Further details can be found in Appendix D. For simplicity, we disregard transitory income shocks and fixed effects in our business cycle model. Table 1 presents the estimation results. The estimated series of persistent income uncertainty is displayed in Figure 2. As one can see, the Great Recession, for example, was accompanied by a drastic increase in the uncertainty of household income.

Figure 2: Estimated level of household income risk over time

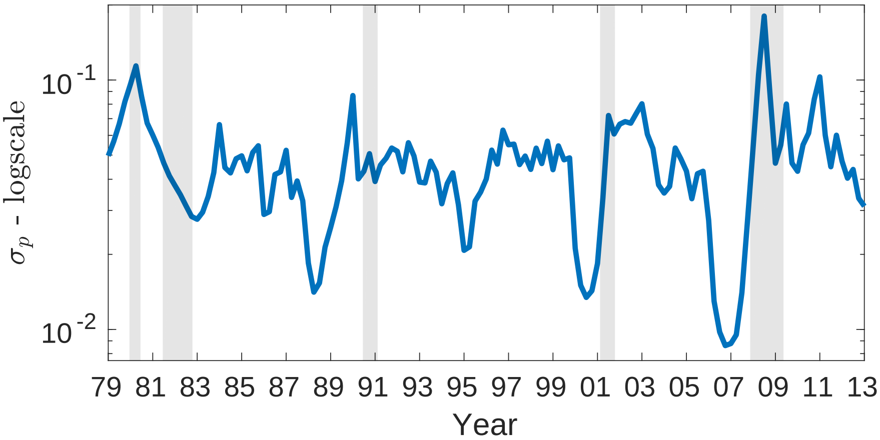

Notes: Estimated variance of persistent income shocks for the period 1979 to 2013. Data come from the SIPP files 1984-2013 for after tax household level income. Only households with at least two married adults, the oldest male being age 30-55, are admitted. Household income is the sum of the incomes of the oldest male and female in a household. NBER recession dates in gray.

Table 1: Estimated parameters for the household income process

Parameter	Value	Description
ρ_h	0.98	Persistence of household labor income
$\bar{\sigma}$	0.06	Average STD of shocks to household labor income
ρ_s	0.84	Persistence of the income-innovation variance, σ_h^2
σ_s	0.54	Conditional STD (log scale) of σ_h^2

Notes: All values are reported at quarterly frequency. For details on the estimation see Appendix D.

5.2 Technology and Preferences

While we can estimate the income process directly from the data, all other parameters are calibrated within the model. Table 3 summarizes our calibration with respect to non-government parameters. In detail, we choose the parameter values as follows.

5.2.1 Intermediate, Final, and Capital Goods Producers

We parameterize the production function of the intermediate good producer according to the U.S. National Income and Product Accounts (NIPA). In the U.S. economy the income share of labor is about 2/3. Accounting for profits we hence set $\alpha = 0.73$.

To calibrate the parameters for the monopolistic competition, we use standard values for markup and price stickiness that are widely employed in the New Keynesian literature. The Phillips curve parameter κ implies an average price duration of 4 quarters (in the equivalent Calvo setting), assuming flexible capital at the firm level. The steady-state marginal costs, $exp(-\mu) = 0.95$, imply a markup of 5%.

We calibrate the adjustment cost of capital, $\phi = 12.5$, to match an investment to output volatility of 4.5 conditional on a TFP shock (see Appendix J).

5.2.2 Households

For the felicity function, $u = \frac{1}{1-\xi} x^{1-\xi}$, we set the coefficient of relative risk aversion, $\xi = 4$, as in Kaplan and Violante (2014). Given that we have effectively set the idiosyncratic elasticity of labor supply in line with low micro-estimates to zero, the chosen value for the inverse Frisch elasticity of labor supply, $\gamma = 1$, reflects the fact that estimates for the aggregate inverse elasticity typically range between 0.5 and 1 (Chetty et al., 2011).

We calibrate the probability of leaving the entrepreneurial state to 1/16 per quarter following the numbers that Guvenen et al. (2014a) report on the probability of dropping out of the top 1% income group in the U.S. (25% p.a.). The fraction of households in the entrepreneurial state, and hence the probability of entering this state, is calibrated to match the average Gini coefficient of total net worth in our SCF sample (78%).

The waves 1983-2013 of the Survey of Consumer Finances (SCF) contain the information we need to calibrate the other household parameters. In line with our treatment of the SIPP data, which we use to estimate income risks, we restrict the sample to households whose household head is between 30 and 55 years of age and married, and to households with at least two adults. This sample selection not only makes the wealth and income data comparable, but also limits the compositional effects of demographic change.

Table 2: Moments targeted in calibration

Targets	Model	Data	Source	Parameter
Mean illiquid assets (K/Y)	2.86	2.86	NIPA	Discount factor
Mean liquidity (B/K)	0.09	0.09	SCF	Mean adj. costs
2nd quintile liquidity (b/k)	0.33	0.33	SCF	Variance adj. costs
Gini total wealth	0.78	0.78	SCF	Fraction of entrepreneurs
Fraction borrowers	0.16	0.16	SCF	Borrowing penalty

The time-discount factor, β, and the distribution of costs (pinned down by its mean and variance) of asset market participation, F_χ, are jointly calibrated to match the *average ratios* of liquid and illiquid assets to output and the *portfolio liquidity of the poor*. In particular, we target the average portfolio liquidity of the second wealth quintile.

We equate illiquid assets to all capital goods at current replacement values in the NIPA tables (1983-2015) because all illiquid assets in our model are both productive and produced. Because they are not productive assets in the NIPA sense, we disregard non-housing, durable consumption goods. This implies for the total value of illiquid assets relative to nominal GDP a capital-to-output ratio of 286% and an annual real return for illiquid assets of 4.5%. We fix the aggregate supply of government bonds, B_t, so as to match the average ratio of aggregate net liquid to net illiquid assets in the Survey of Consumer Finances (average 1983-2013: 9%).[15] We consider all deposits, money market accounts, and bonds net of credit card debt as liquid assets. All other assets in the SCF and all non-credit-card debts are considered illiquid as in Kaplan et al. (2016). Since we abstracted from consumer durables, we also disregard car wealth and auto loans in these calculations.

[15] This number is relatively close to the ratio of average U.S. federal debt held by domestic private agents relative to capital of 10.5%.

Table 3: Calibrated parameters: Firms and households

Parameter	Value	Description	Target
Households			
β	0.98	Discount factor	see Table 2
ξ	4	Relative risk aversion	Kaplan and Violante (2014)
γ	1	Inverse of Frisch elasticity	Chetty et al. (2011)
μ_χ	64125	Participation utility costs	see Table 2
σ_χ	22500	Participation utility costs	see Table 2
$\overline{R}$	10.5%	Borrowing penalty	see Table 2
Intermediate Goods			
α	0.70	Share of labor	Income share of labor of 2/3
δ	1.35%	Depreciation rate	NIPA: Fixed assets
Final Goods			
κ	0.09	Price stickiness	Mean price duration of 4 quarters
μ	0.05	Markup	5% markup
Capital Goods			
ϕ	12.5	Capital adjustment costs	Relative investment volatility of 4.5

The empirical distribution of portfolio liquidity sheds light on how state-dependent liquidation decisions are, i.e., whether the logistic distribution for adjustment costs has a high or a low variance. For this purpose, we calculate for each household in the SCF waves since 1983 the amount of net liquid and net illiquid assets. We then estimate for each survey year a function that maps the percentile rank of a household in the total wealth distribution into liquid and illiquid asset holdings by a local-linear regression; see Appendix E for details. Using these two functions, we then calculate the average relation of liquid to illiquid assets of a household at a given percentile in the wealth distribution. Figure 3 shows the average holding of liquid relative to illiquid assets over the period 1983-2013 in the SCF and implied by the model. Note that only households above the 20th percentile have typically non-negligible amounts of positive illiquid asset holdings net of illiquid mortgage debt in every year of the sample. In the data, richer households hold a smaller fraction of their wealth in liquid form.[16]

Our model produces this downward sloping curve, too. The intuition is that households hold bonds because they provide better short-term consumption smoothing than

[16] Despite this fact, the Gini coefficient of liquid wealth is larger than the Gini of illiquid wealth.

Figure 3: Average holdings of liquid assets relative to illiquid assets by wealth quintile

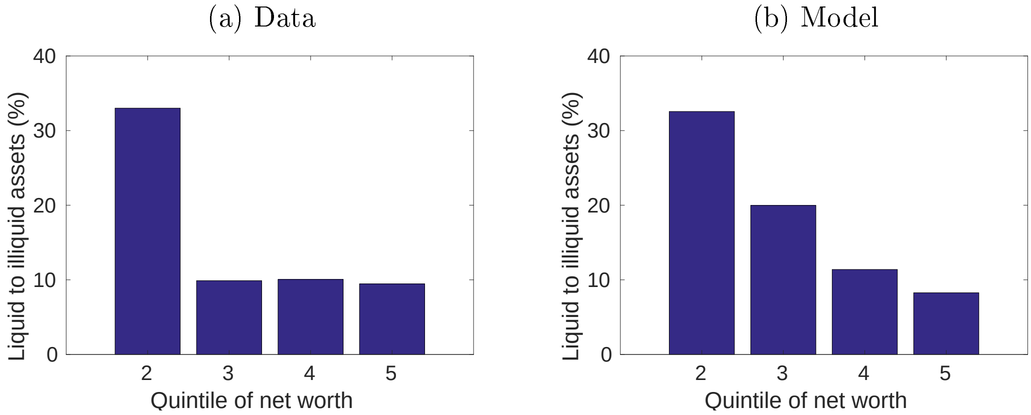

Notes: Estimated net liquid asset holdings relative to estimated net illiquid assets by quintile of the net wealth distribution. Average over the estimates from the SCFs 1983-2013. We select only households composed of at least two adults whose head is between 30 and 55 years of age. Estimation by a local linear estimator with a Gaussian kernel and a bandwidth of 0.05. Relative holdings below the 21st wealth percentile are not reported, because the net illiquid asset holdings can be zero and net liquid holdings negative.

capital and that this value of liquidity decreases in the amount of bonds a household holds. Furthermore, for richer households a larger share of income comes from capital and is hence not subject to labor income risk. Therefore, richer households, which typically hold both more bonds and more capital (even relative to their income), hold less liquid portfolios. Table 2 summarizes how we match our targets from the wealth distribution. The calibrated adjustment costs imply an average adjustment frequency of 5.3% per quarter that increases for households with unbalanced portfolios to up to 15% probability of adjustment.[17] The average adjustment frequency is close to the implied adjustment probability in Kaplan and Violante (2014).

Of course, ours is a highly stylized treatment of the financial sector in assuming that no physical investment is directly financed by the issuance of liquid assets. This stylized view, however, is motivated by the data. The Flow of Funds show (Z1-Table L.213) that roughly 73% of all corporate equities held or issued in the U.S. are either not publicly traded (11%) or are held by agents other than households or depository institutions (62%) (important holders are mutual funds (24%), the rest of the world (16%), pension funds (12%)).[18] On top of that, some of the corporate equities held by households will be held

[17]We provide robustness checks on the distribution of adjustment costs in Appendix L and find our results to robust.

[18]Roughly 50% of the mutual funds are directly held by households; the rest are mostly held by pension funds.

Table 4: Calibrated parameters: Government

Parameter	Value	Description	Target
Monetary Policy			
$\bar{R}^b$	1.0062	Nominal Rate	2.5% p.a.
$\bar{\pi}$	1.00	Inflation	0% p.a.
θ_π	1.25	Reaction to inflation deviations	Standard value
ρ_{R^b}	0.8	Persistence in Taylor Rule	Standard value
Fiscal Policy			
ρ_B	0.98	Reaction to debt deviations	Discount factor
γ_π	1.5	Reaction to inflation deviations	Empirical deficit response
γ_T	0.75	Reaction to tax deviations	Empirical deficit response
τ	0.3	Labor tax rate	$G/Y = 20\%$

by large holders (e.g., company founders) who impact the asset price when transacting. Even more extreme is the distribution of corporate bond holdings, of which more than 80% are held outside households and depository institutions (Flow of Funds, Z1-Table L.223).

5.2.3 Monetary Policy

We set the coefficients of the Taylor rule to standard values commonly used for New Keynesian models. The coefficient θ_π describes the reaction of the nominal interest rate to deviations of inflation from the steady state and ρ_{R^b} captures persistence in the nominal interest rate. We set $\theta_\pi = 1.25$ and $\rho_{R^b} = 0.8$. We set steady-state inflation to zero. The steady-state nominal interest rate is therefore equal to the real rate, which we set to 2.5% (annual). In order to match the fraction of indebted households, we add a wedge between the lending and the borrowing rate of $\overline{R} = 10.5\%$ (annual).

5.2.4 Fiscal Policy

Government spending evolves according to a fiscal rule similar to Bi et al. (2013). Government spending reacts to deviations of debt from the steady state, ρ_B, and to deviations of inflation, γ_π, and taxes, γ_T, from the steady state. We set $\rho_B = \beta$, such that the speed of repayment of government debt matches the time-preference rate of households. We set $\gamma_\pi = 1.5$, and $\gamma_T = 0.75$ to match the empirical government deficit response to a one standard deviation shock to household income risk (cf. Figure 9). We choose the tax rate and government expenditures such that they account for 20% of output in the steady state, implying a tax rate of 30%.

6 Quantitative Results

Having determined the parameters of the model, we can quantify the aggregate effects of shocks to household income risk in an environment of incomplete markets and a portfolio choice between an illiquid real and a liquid paper asset. We will compare these to the effects of shocks to household income risk in the data in Section 7 and find them to be very similar in size.

6.1 Household Portfolios and the Individual Response to Income Risk

We first describe the individual household response in partial equilibrium to clarify the mechanics of a shock to household income risk (or for short: uncertainty shock). For that purpose, we fix prices and expectations at their steady-state level and solve for the household decisions by discretizing the uncertainty process. We also solve the model without uncertainty shocks and use this variant to obtain the stationary cross-sectional distribution of households in liquid and illiquid assets and income. We then order households along the net worth dimension, and estimate from the policy functions of the model with uncertainty shocks average consumption as well as liquid and illiquid asset holdings by net worth using a local-linear regression technique. We compare a situation when uncertainty is at its average value to an increase by one standard deviation (a variance increase of 54%). Given that the planning problem of households uses the actual uncertainty process but fixed prices, this identifies the average partial equilibrium effect of uncertainty.

Figure 4 presents the results. For all households, consumption declines. As income risk goes up, households want to save more and they want to do this in liquid form. In fact, those households that decide to adjust their portfolios sell illiquid capital in exchange for liquid assets. Therefore, the liquidity of portfolios increases across all wealth groups.

Figure 5 shows the general equilibrium response of portfolio liquidity and consumption across the wealth distribution, where we allow prices to adjust and expectations to be consistent with equilibrium. In equilibrium, wage incomes fall and pure profits increase. Therefore, poorer households use some of their liquidity to smooth consumption, and compared to the partial equilibrium response, their liquidity increase is muted. On the other hand, some rich entrepreneur households see a temporary increase in income, and

Figure 4: Partial equilibrium response – Change in individual policy upon an income risk shock keeping prices and expectations constant at steady-state values

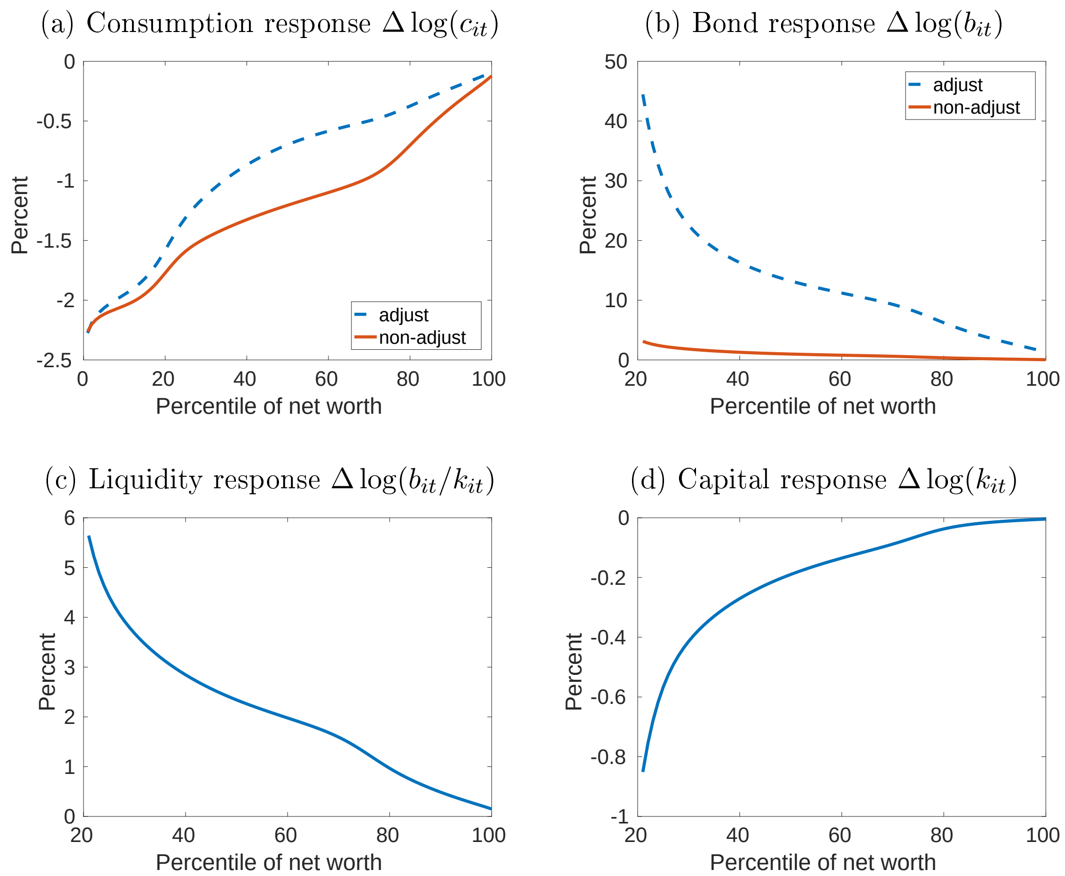

Response of individual consumption and asset demand policies at constant prices and price expectations to a one standard deviation increase in income risk. Policies by wealth percentile are estimated using a local linear regression technique with a Gaussian kernel and a bandwidth of 0.05. The figures compare the estimated function at average risk and at a one standard deviation increase, which is equal to a 54% increase in the variance of income shocks.

Top Panels: Conditional on adjustment decisions.

Bottom Panels: Average response over both adjusters and non-adjusters.

Figure 5: General equilibrium response – Change in the liquidity of household portfolios and consumption

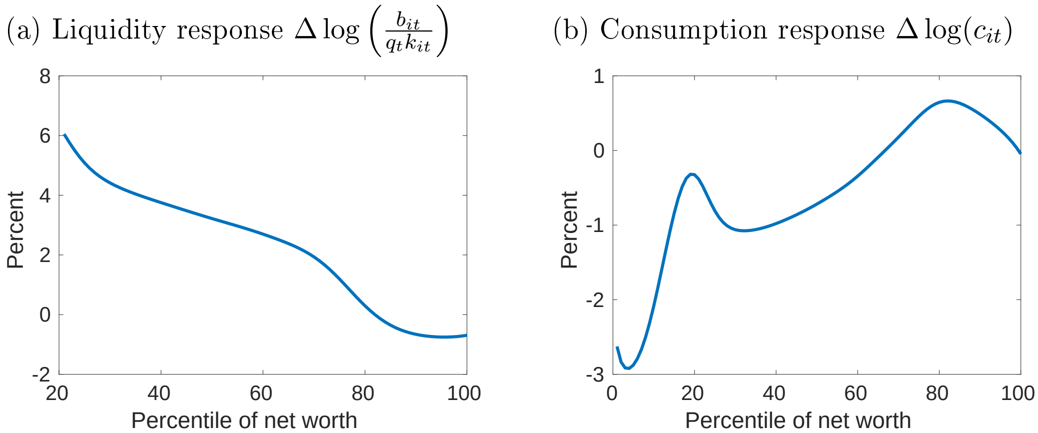

Notes: Change in the distribution of liquidity and consumption at all percentiles of the wealth distribution after 3 years at equilibrium prices and price expectations after a one standard deviation shock to income risk. The liquidity of the portfolios is averaged using frequency weights from the simulated wealth distribution and reported conditional on a household falling into the x-th wealth percentile. The left-hand panel shows the change in portfolio liquidity; the right-hand panel shows the consumption response. As with the data, we use a Gaussian kernel-weighted local linear smoother with bandwidth 0.05.

invest in liquid assets. Interestingly, this picture resembles what we found in Figure 1, in which the increase in the liquidity of the portfolios is strongest for the lower middle class. Only the magnitude of changes in the liquidity of household portfolios during the Great Recession was much more dramatic. However, our estimates for the income process also suggest that the Great Recession saw not a one standard deviation but more than a five times stronger increase in income uncertainty. For a one standard deviation shock to income uncertainty, we find quantitatively more comparable results to what our model produces; see Section 7.

6.2 Aggregate Consequences of Shocks to Household Income Risk

6.2.1 Main Findings

As Figure 4 shows, upon an increase in income risk, the demand for consumption and capital simultaneously falls. Given that output is partly demand determined, output, wages, and dividends need to fall in equilibrium. Figure 6 displays the impulse responses of aggregate output and its components, real bond holdings and the capital stock as well as asset prices and returns for our baseline calibration. The assumed monetary

Figure 6: Aggregate response to household income risk shock

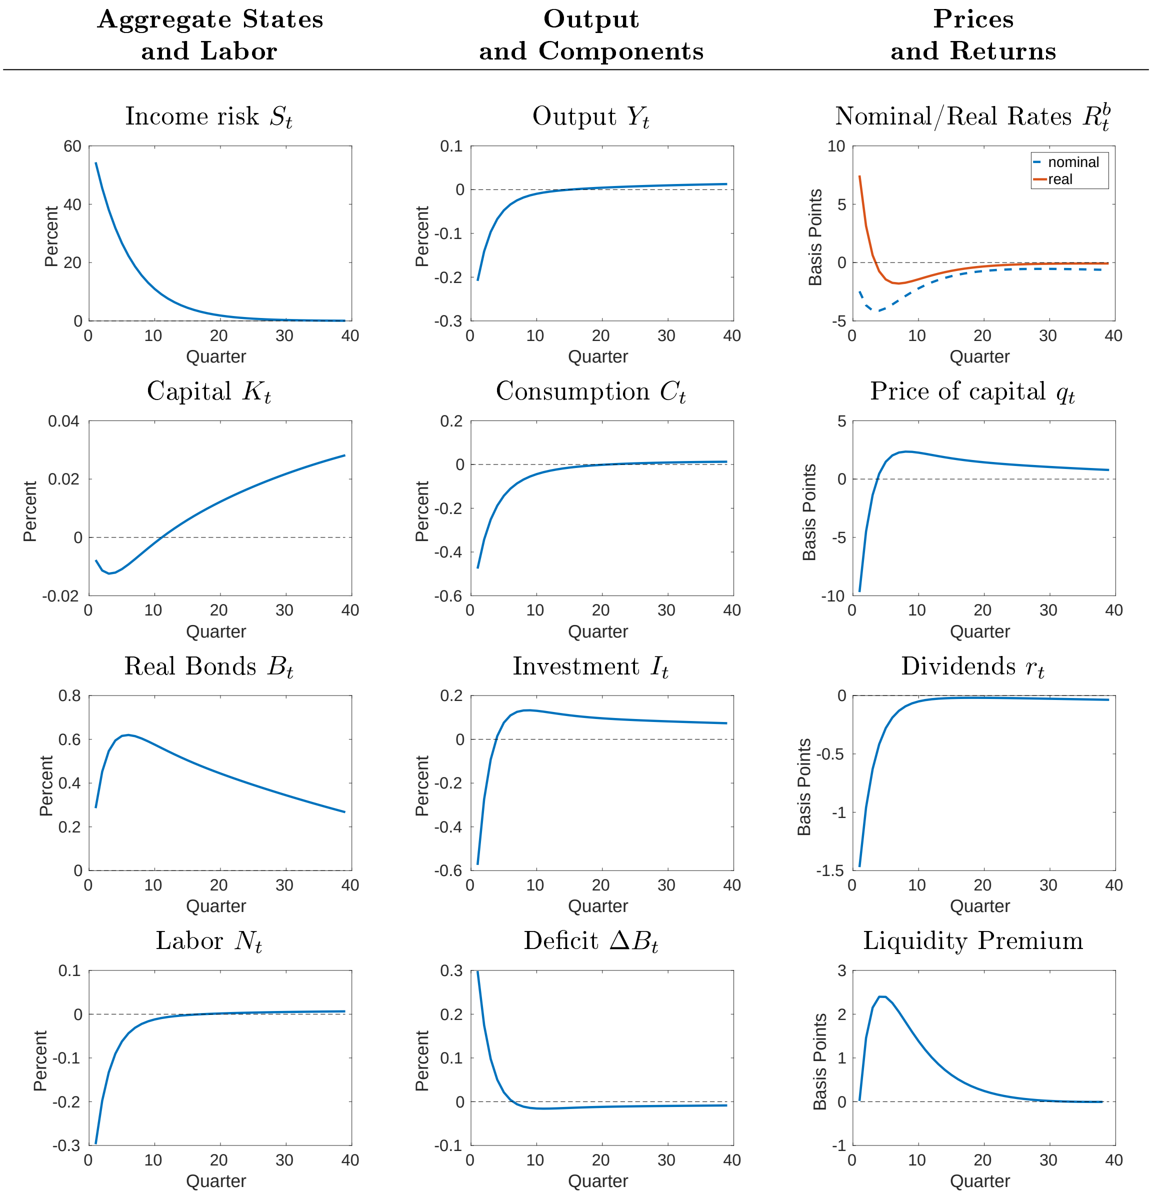

Notes: Liquidity Premium: $\frac{E_t q_{t+1} + r_t}{q_t} - \frac{R^b_t}{E_t \pi_{t+1}}$.

Impulse responses to a one standard deviation increase in the variance of income shocks. All rates (dividends, interest, liquidity premium) are *not* annualized.

policy reacts to the uncertainty-induced deflation by cutting rates. Fiscal policy expands government expenditure. After a one standard deviation increase in the variance of idiosyncratic income shocks, output drops on impact by 0.2% and only recovers after 12 quarters. Consumption falls even more and remains subdued for roughly 20 quarters. Investment on impact sees the sharpest decline of all aggregates - three times stronger than output.

The output drop in our model results from households increasing their precautionary savings in conjunction with a portfolio adjustment toward the liquid asset. In times of high uncertainty, households dislike illiquid assets because of their limited use for short-run consumption smoothing. Consequently, the price of capital decreases on impact. Since the demand for the liquid asset is a demand for paper and not for (investment) goods, demand for both consumption and investment goods falls. At the same time the central bank cuts interest rates on bonds, which stabilizes the demand for illiquid assets. Despite an increase in the quantity of bonds, the liquidity premium, i.e., the return difference between illiquid and liquid assets, increases.

Quantitatively, we find that fluctuations in household income risk explain a significant fraction (17%) of the business cycle in terms of standard deviations; see Appendix G. This increases markedly when the zero lower bound on the nominal interest rate binds. The output loss is ten times larger, roughly 2%, at the zero lower bound without monetary or fiscal stabilization; see Appendix H.

6.2.2 Stabilization Policy

There are two ways the government can stabilize the economy in our setup: by cutting rates on bonds to shift asset demand from liquid to illiquid assets, i.e., by monetary policy, or by increasing the supply of government bonds, i.e., through fiscal policy. Our baseline calibration is a mix of the two following the empirical results in Section 7.

To obtain a better understanding of the differences between the two policy options, we next consider two extreme scenarios: One where monetary policy reacts very strongly to inflation, $\theta_\pi = 100$, but fiscal policy does not at all, $\gamma_\pi = 0$, $\gamma_T = 0$, and an alternative scenario, where monetary policy keeps a nominal interest rate peg, $\theta_\pi = 0$, and fiscal policy reacts strongly to inflation, $\gamma_\pi = 100$.[19]

[19]The price level is determinate despite the fact that fiscal policy is stabilizing the value of debt, because government bonds provide liquidity services. Given some expected real interest rate, households for a given wealth distribution demand a certain amount of bonds. The government bond supply rule then pins down the inflation rate in the current period as long as $\rho_B + \gamma_\pi \neq 0$ (essentially this is the "Pigou effect"; see Leith and von Thadden, 2008, on price level determinacy in non-Ricardian economies). Since the wealth distribution shifts with inflation, the actual condition for determinacy requires $\rho_B + \gamma_\pi$

Figure 7: Aggregate response to household income risk shock with stabilization policy

(a) **Monetary Stabilization**

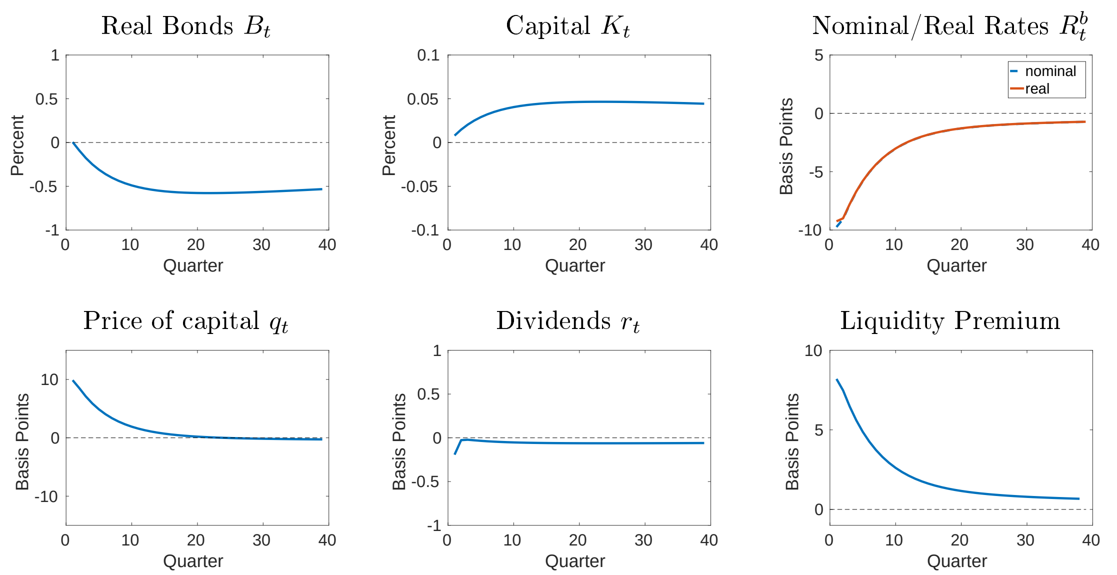

(b) **Fiscal Stabilization**

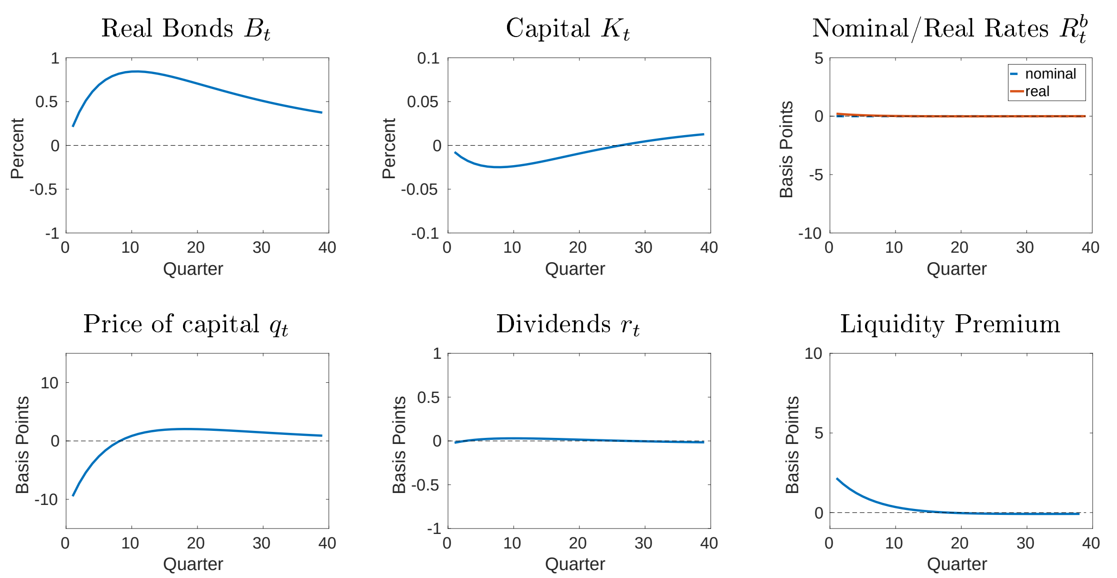

Notes: Liquidity Premium: $\frac{E_t q_{t+1} + r_t}{q_t} - \frac{R_t^b}{E_t \pi_{t+1}}$.
Impulse responses to a one standard deviation increase in the variance of income shocks. All rates (dividends, interest, liquidity premium) are *not* annualized.
Top Panels: Taylor-rule coefficient on inflation is set to $\theta_\pi = 100$, and fiscal policy does not respond to inflation $\gamma_\pi = 0$ and taxes $\gamma_T = 0$.
Bottom Panels: Fiscal policy strongly responds to inflation $\gamma_\pi = 100$, and the Taylor-rule coefficient on inflation is zero, $\theta_\pi = 0$.

Both regimes successfully stabilize inflation and output at their steady-state levels. Yet, they still see a drop in consumption, as households want to increase their precautionary savings. The results for the key variables that differ across the regimes are depicted in Figure 7. Under monetary stabilization, the central bank increases the liquidity premium by lowering the interest rate until the excess supply of goods at steady-state inflation is eliminated. The lower interest rate spurs both investment and consumption. In the fiscal policy case, the government instead supplies the liquid assets the household demands until the increase in the liquidity premium is eliminated as the return on liquid assets rises when they are more abundant. The increased savings of households are thus held in the form of government bonds used to finance government expenditures in the fiscal stabilization case. The welfare consequences for the different groups of households vary in the two regimes due to their different implications for the price of and return on liquid and illiquid assets. Monetary stabilization drives down the return on liquid assets and increases the price of capital. Fiscal stabilization by contrast increases the return on liquid assets but lowers the price of capital.

6.3 An Extension with Asset-Backed Securities

One potential limitation of our model is that it cannot handle mortgages, which are an important way for the banking sector to create liquid assets out of illiquid investments. To capture the effect of mortgages, we assume that any newly created capital good is partly credit financed, i.e., any investment creates an illiquid asset and a liquid asset-backed security at the same time. Let ζ being the number of bonds issued as mortgages per unit of capital. We assume that the fraction ζ cannot be adjusted by the household sector. However, households can buy back the mortgages originated from "their" illiquid assets.

For this purpose, we model a wedge $\underline{R}$ between the borrowing and lending rate on mortgages due to the costs of intermediation. The rate paid to lenders is the central bank's interest rate R_t^b, and the rate borrowers pay is $R_t^b + \underline{R}$. This means that, net of the payments on the asset-backed securities, the dividend stream from illiquid assets decreases to $r_t = F_K(K, N) - \delta - \zeta(R_t^b + \underline{R})$. However, also the price of illiquid assets decreases, which now is $q_t = 1 - \zeta + \phi \frac{K_{t+1} - K_t}{K_t}$ in equilibrium, because the producer of each unit of illiquid assets can sell ζ units of asset-backed securities for each unit of the illiquid capital good in addition to that good itself.

To allow households to adjust the extent to which they effectively draw their mort-

to be somewhat larger.

gages, we assume that the borrowing wedge does not apply to securities backed by assets the household itself owns. This means that the household now faces three interest rates on liquid assets:

$$R(b_{it}, R^b_{t-1}) = \begin{cases} R^b_{t-1} & \text{if } b_{it} \geq \zeta k_{it} \\ R^b_{t-1} + \underline{R} & \text{if } 0 < b_{it} < \zeta k_{it} \\ R^b_{t-1} + \overline{R} & \text{if } b_{it} < 0. \end{cases} \tag{23}$$

The highest interest rate applies to unsecured borrowing $b < 0$. An intermediate interest rate applies if the household buys back securities originated from the illiquid asset it owns $0 < b < \zeta k$, i.e., that the household saves by paying back a mortgage. The lowest interest rate applies when the household accumulates liquid assets beyond those it has originated.

The bond market equilibrium condition then reads:

$$\begin{aligned} &\zeta K_{t+1} + B_{t+1} \qquad (24) \\ = &\int\int\int_{b \geq \underline{B}} \left[\nu^* b^*_a(b,k,h;q_t,\pi_t,R^b_{t+1}) + (1-\nu^*) b^*_n(b,k,h;q_t,\pi_t,R^b_{t+1})\right] d\Theta_t(b,k,h), \end{aligned}$$

where ζK_{t+1} is the amount of asset-backed securities in circulation. The market clearing condition for illiquid assets remains unchanged.

We have re-calibrated the amount of government debt to keep the average portfolio liquidity unchanged when not counting securities held by the issuer. The ratio of mortgage liabilities of households to their net worth in the Flow of Funds (Table Z1-B.101) is roughly 10%. We target this number and a fraction of original mortgage debt to be paid back of 50%. Therefore, we set $\zeta = 20\%$ and calibrate $\underline{R} = 0.5\%$ p.a.

Appendix I shows the impulse response to an uncertainty shock in this model extension. Compared to our baseline scenario, the recessionary impact of uncertainty is larger, because the re-balancing of portfolios implies a decline in the supply of liquid assets as households reduce the stock of capital.

6.4 Redistributive and Welfare Effects

So far we have described how the individual reaction to an increase in uncertainty has aggregate repercussions. Since these aggregate consequences affect asset prices, dividends, wages, and entrepreneurial incomes differently, our model predicts that not all agents (equally) lose from the decline in consumption upon an uncertainty shock. For example, if capital prices fall, those agents who have high productivity and hence are rich in human capital, but hold little physical capital, could actually gain from the uncertainty shock.

These agents are net savers. They increase their holdings of physical capital and can now do so more cheaply.

To quantify and understand the relative welfare consequences of the uncertainty shock and of systematic policy responses, we calculate the difference in expected value EV after a one standard deviation increase in uncertainty relative to its steady-state value for all (b, k, h) triples.[20] To put this number into perspective we normalize by the expected discounted felicity stream from consumption and leisure given (b, k, h) that a household expects. This way, we can calculate how much larger lifetime consumption would need to be to compensate a household for the effect of the uncertainty shock. This consumption equivalent takes the form:

$$\begin{aligned} \mathtt{CE}(b,k,h) &= \left[\frac{EV(b,k,h;\Theta^{ss},\sigma_s) - EV(b,k,h;\Theta^{ss},0)}{EU(b,k,h)} + 1\right]^{1/(1-\xi)} - 1, \quad (25) \\ EU(b,k,h) &= \sum_{t=0}^{\infty} \beta^t u(x_t^*), \end{aligned}$$

where Θ^{ss} is the steady-state distribution and the sequence x_t^* results from optimal decisions of households using stationary equilibrium policies.

Table 5 provides the consumption equivalents for both baseline and stabilization policies. The average welfare loss (one-sided) from the uncertainty shock is 0.25% of lifetime consumption. Table 5 shows how much larger or smaller the losses are across population groups. What confounds results somewhat is that households with low labor income mechanically gain from an increase in the variance of shocks to productivity h, because expected productivity growth is positively related to uncertainty for low productivity and negatively related for high productivity households.[21] Also, entrepreneurs profit from the uncertainty shock as markups and thus profits go up. Therefore, it is particularly useful to look at the differences in welfare effects across groups, keeping the other characteristics constant; see the "Median" rows in Table 5.

In general, welfare losses are substantially more pronounced for those households with few asset holdings. In fact, comparing the average welfare loss across the policy regimes, we find that the numbers are very similar. It follows that the main source of

[20]For this purpose, we calculate the value functions iterating backward given the equilibrium price and uncertainty paths after an uncertainty shock, which we obtain by linearization using Reiter's procedure. We check with the Krusell-Smith variant for our baseline and find virtually the same results.

[21]The conditional expected productivity growth is $g(\sigma_{h,t}, \log \tilde{h}_{it}) := E_t \exp(\Delta \log \tilde{h}_{it+1}) \frac{\int \tilde{h}_{it}}{\int \tilde{h}_{it+1}} = \exp[(\rho_h - 1)\log \tilde{h}_{it} + 0.5\sigma_{h,t}^2] \frac{\int \tilde{h}_{it}}{\int \tilde{h}_{it+1}}$. Across all households expected productivity growth is zero. The cross derivative is negative $\frac{\partial^2 g}{\partial \sigma_{h,t} \partial \tilde{h}_{it}} = g(\sigma_{h,t}, \log \tilde{h}_{it})(\sigma_{h,t}[1 - (\int \tilde{h}_{it})^{-\rho_h^2})](\rho_h - 1) < 0$.

Table 5: Welfare effects of household income risk shock

(a) Policy regime: Baseline

	Quintiles of bond holdings					Percentiles of human capital			
	1.	2.	3.	4.	5.	0-33	33-66	66-99	Entr.
Conditional	0.04	-0.11	-0.06	-0.00	0.14	0.14	-0.04	-0.09	0.91
Median	-0.10	-0.06	-0.03	-0.01	0.06	0.19	-0.04	-0.16	1.72

	Quintiles of capital holdings					Average CE: -0.25
	1.	2.	3.	4.	5.	
Conditional	-0.06	-0.08	-0.05	0.01	0.16	
Median	-0.20	-0.12	-0.04	0.03	0.13	

(b) Policy regime: Monetary stabilization

	Quintiles of bond holdings					Percentiles of human capital			
	1.	2.	3.	4.	5.	0-33	33-66	66-99	Entr.
Conditional	0.08	-0.10	-0.04	0.01	0.07	0.16	-0.03	-0.09	-0.06
Median	-0.08	-0.06	-0.01	0.01	-0.01	0.20	-0.02	-0.14	0.14

	Quintiles of capital holdings					Average CE: -0.24
	1.	2.	3.	4.	5.	
Conditional	-0.04	-0.07	-0.04	0.01	0.15	
Median	-0.16	-0.10	-0.02	0.03	0.14	

(c) Policy regime: Fiscal stabilization

	Quintiles of bond holdings					Percentiles of human capital			
	1.	2.	3.	4.	5.	0-33	33-66	66-99	Entr.
Conditional	0.06	-0.11	-0.07	-0.00	0.13	0.15	-0.03	-0.09	0.22
Median	-0.09	-0.06	-0.04	-0.01	0.05	0.19	-0.04	-0.16	0.18

	Quintiles of capital holdings					Average CE: -0.23
	1.	2.	3.	4.	5.	
Conditional	-0.06	-0.08	-0.05	0.01	0.18	
Median	-0.20	-0.12	-0.04	0.04	0.17	

Notes: Welfare costs of a one standard deviation increase in income risk in terms of consumption equivalents (CE) as defined in (25) in % minus the population average for each regime. "Conditional" refers to integrating out the missing dimensions, whereas "Median" refers to median asset holdings of the respective other assets.
(b) Policy coefficients are: $\theta_\pi = 100$, $\gamma_\pi = 0$, $\gamma_T = 0$.
(c) Policy coefficients are: $\theta_\pi = 0$, $\gamma_\pi = 100$, $\gamma_T = 0.75$.

Figure 8: Welfare effects of household income risk shock with stabilization

(a) Monetary Stabilization (b) Fiscal Stabilization

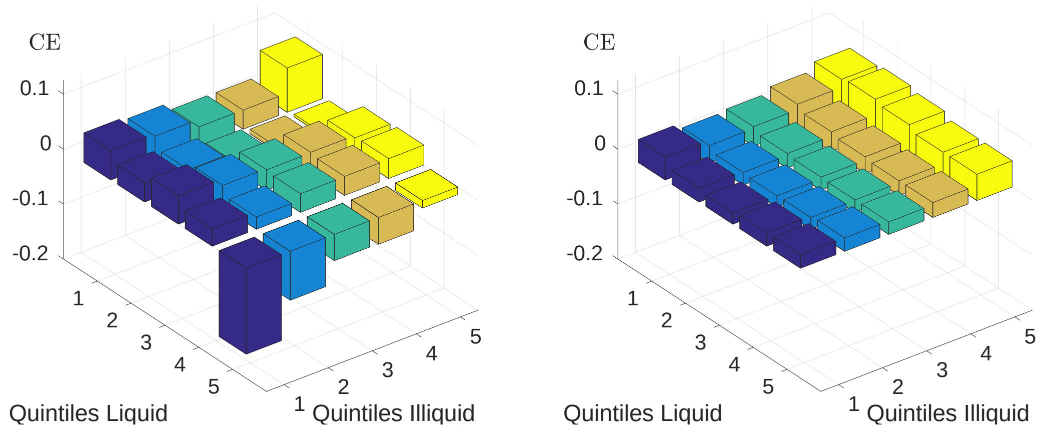

Notes: Welfare costs of a one standard deviation increase in income risk relative to the baseline policy specification in terms of consumption equivalents (`CE`) as defined in (25) integrating out h_{it} (excluding entrepreneurs). The graphs refer to the conditional expectations of `CE` with respect to the joint distribution of liquid and illiquid assets for workers.
Left Panel: Taylor-rule coefficient on inflation is set to $\theta_\pi = 100$, and fiscal policy does not respond to inflation $\gamma_\pi = 0$ and taxes $\gamma_T = 0$.
Right Panel: Fiscal policy strongly responds to inflation $\gamma_\pi = 100$, and the Taylor-rule coefficient on inflation is zero, $\theta_\pi = 0$.

welfare losses is the lack of idiosyncratic insurance. The (one-sided) welfare costs of the aggregate downturn itself is less important and on the order of 0.02% of lifetime consumption (baseline minus fiscal stabilization).

Notwithstanding, the stabilization policies have sizable distributional consequences. Fiscal stabilization benefits particularly those who hold a lot of illiquid wealth as it stabilizes the dividend payments from these assets. Across liquid asset holdings the welfare benefits are relatively evenly distributed; see Figure 8 (b). Stabilization through monetary policy, by contrast, redistributes from households with particularly liquid portfolios to households with very little total wealth. Focusing on households with close to no illiquid wealth (1st quintile), we observe that the relative welfare gains amount to as much as 10 basis points of lifetime consumption for indebted households (1st quintile) to minus 20 basis points for households in the top quintile of liquid wealth; see Figure 8 (a). Households that hold large amounts of illiquid wealth also benefit under monetary stabilization from stable markups and hence relatively stable dividends. Yet, even households with balanced portfolios, which are rich in both liquid and illiquid assets, lose from monetary stabilization relative to the baseline because in equilibrium their asset returns fall.

7 Empirical Evidence

How does our model relate to the data? With the estimation of the income process as described in Section 5.1, we obtain a sequence of shocks, ϵ_t^s, to household income risk. We can use this sequence $\{\epsilon_t^s\}_{t=1976Q1...2013Q1}$ to assess whether the predictions of our model are borne out by the data. Since the shocks before the SIPP samples start are not well identified and as structural breaks in monetary policy may impact results, we focus on the post-Volcker disinflation era and discard estimated shocks before 1983Q1.

We first estimate the effect of household income risk on aggregate economic activity and average household portfolios. As we will see, upon an increase in income risk, aggregate output falls with a particularly strong decline in investment, as households rebalance their portfolios toward liquid assets, while the (equilibrium) return premium on illiquid assets goes up. We then combine the uncertainty series with cross-sectional information on household portfolios from the Survey of Consumer Finances (SCF), and find that the relatively poor initially increase the liquidity of their portfolios more than proportionally.

7.1 Aggregate Response

Figure 9 shows the response of aggregate variables to an increase in household income risk. We estimate the response by local projections using the shock series we identified from the SIPP data while controlling for lagged aggregate variables, lagged income risk and a time trend. In line with the specification of our model, we assume that the realized uncertainty shock in time $t + 1$ is observed at time t. Alternatively to the local projection method, we have estimated impulse responses using structural VARs that give very similar results and are available in Appendix F.2.

Upon a one standard deviation increase in income risk, output falls by roughly 0.2% on average over the first year. The trough is reached six quarters after the shock with a 0.3% decrease in output. Consumption has very similar dynamics but goes down slightly less. Investment drops too, but its reaction is roughly five times as strong as the output reaction. The measured Solow residual from Fernald's TFP series (Fernald et al., 2012) falls as well, which is in line with any model of monopolistic competition, in which the Solow residual captures markup movements (Hall, 1989). The government seems to react systematically by making use of stabilizing monetary and fiscal policy – government

Figure 9: Empirical response to household income risk shock

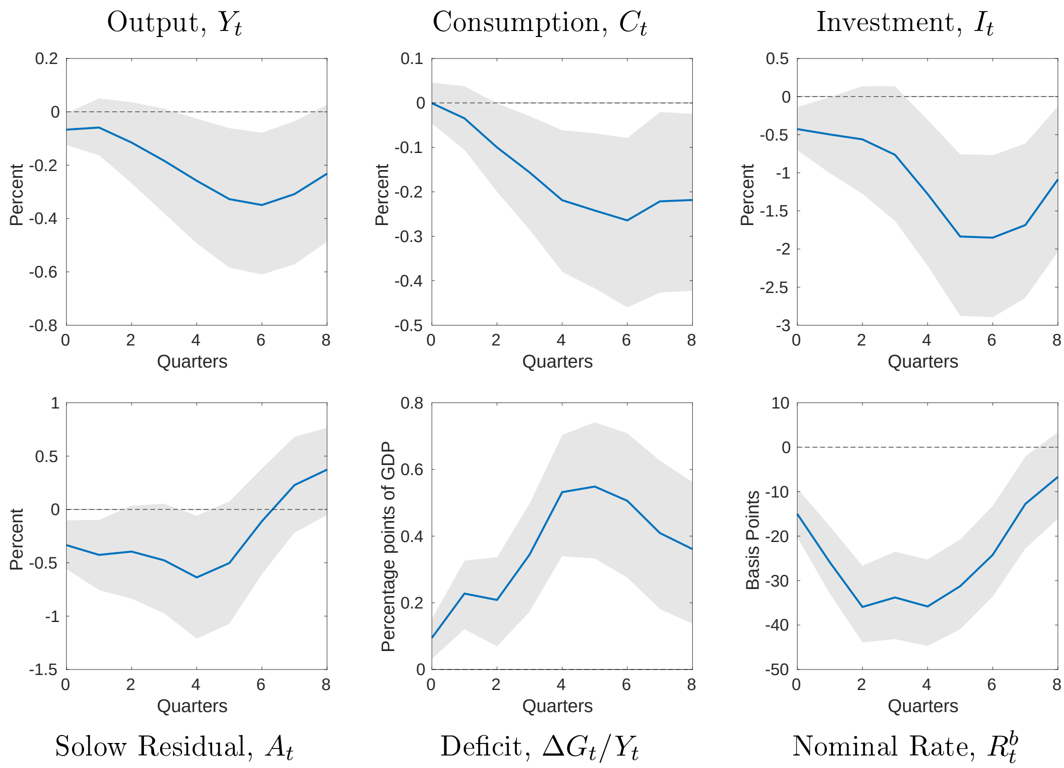

Estimated response of $\mathbf{X}_{t+j}, j = 0 \dots 8$, where $\mathbf{X}_t = [Y_t, C_t, I_t, A_t, \Delta B_t/Y_t, R_t^b]$, to the estimated shocks to household income risk, ϵ_t^s. The regressions control for the lagged state of the economy $\mathbf{X}_{t-1}$ and lagged levels of income risk s_{t-1}. The nominal rate is the 3-month T-bill rate. Bootstrapped 66% confidence bounds in gray (block bootstrap).

deficits go up by 0.3 percentage point of GDP over the first year and the nominal return on 3-month Treasury bills goes down on average by 35 basis points (annualized) over the first four quarters after the shock.

The decline in investment – despite a decrease in interest rates – finds its repercussions in household balance sheets; see Figure 10. The ratio of liquid-to-illiquid assets goes up after an increase in household income risk. We calculate this ratio from the Flow of Funds (Table Z1-B.101) by subsuming as liquid assets all deposits, cash, debt securities (including government bonds), and loans held directly, while we treat all other real and financial assets as illiquid.[22]

A part of the increase in the liquidity of household portfolios is driven by real house prices as houses make up the lion's share of the illiquid assets of households (close to

[22] Kaplan et al. (2016) use a very similar taxonomy to split assets into liquid and illiquid. The reason to treat equities as illiquid is that most equities are held in the form of pension funds. Equity shares held directly only play a role above the 85th wealth percentile, but even these are often closely held equities

Figure 10: Response of household portfolios, house prices and the liquidity premium to household income risk shock

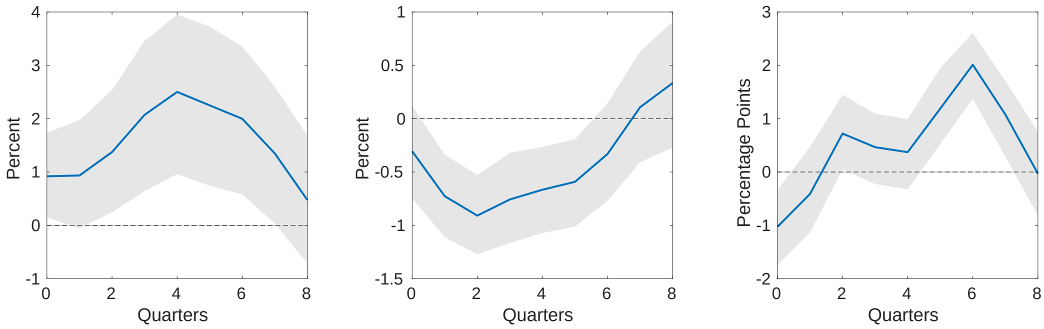

Portfolio Liquidity | House prices | Liquidity Premium

Estimated response of the liquidity of household portfolios, the price of houses (Case-Shiller S&P Index), and the difference between the return on housing and the nominal rate (Liquidity Premium) to income risk using local projections. The set of control variables is as in Figure 9. Bootstrapped 66% confidence bounds in gray (block bootstrap).

50% on average; see Kaplan and Violante, 2014). Hence any change in house prices directly affects portfolio liquidity. However, as house prices fall only by 1% after an increase in uncertainty, they can make up only for about a quarter of the increase in liquidity. The largest part of the increase in portfolio liquidity must therefore come from outright different reactions in the demand for liquid and illiquid assets. In fact, the return premium of houses over liquid assets increases relatively quickly after the shock to household income risk.[23]

Overall, we find very similar responses to uncertainty shocks in terms of the size of the peak response in the model and the data. The data show typically hump-shaped responses, which our model cannot generate because both government expenditures and investment can adjust on impact in the model, while they do so slowly in reality.

---

such as S-corporations, or other illiquid forms. Publicly traded equities, which a single household can sell without price impact, play a significant role in household portfolios only for a relatively small fraction of households and a small fraction of the aggregate capital stock.

[23]We proxy the liquidity premium by the realized return on housing (rent-price ratio in t plus realized growth rate of house prices in $t+1$) relative to the nominal rate. The house price we use is the Case-Shiller S&P national house price index. Rents are imputed on the basis of the CPI for rents of primary residences, fixing the rent-price ratio in 1981Q1 to 4%.

Figure 11: Response of portfolio liquidity to household income risk by wealth percentile

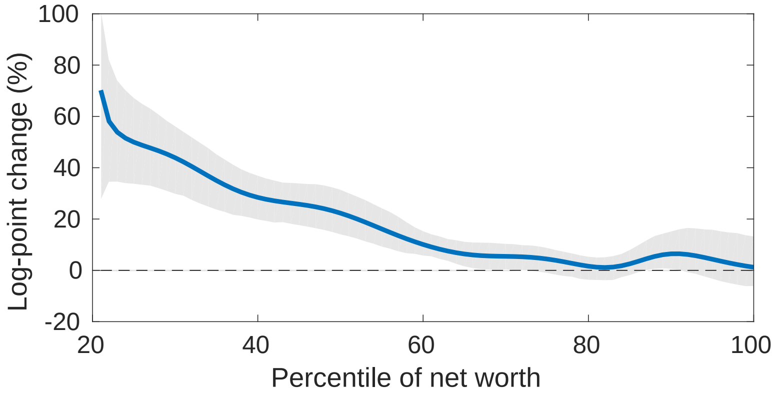

Estimated log difference of liquid-to-illiquid ratio of household portfolios across the wealth distribution in response to a one standard deviation shock to household income risk. Income risk shocks are identified from the SIPP. Portfolio composition is estimated from the SCF survey years 1983-2013, only household with at least two adults and the household head being between 30 and 55 years of age. Bootstrapped 66% confidence bands in gray, based on a non-parametric bootstrap.

7.2 Response by Wealth Group

Besides our model's predictions for aggregate economic variables, average household portfolios, and the differential changes in the return on liquid and illiquid assets, it also has rich cross-sectional implications for households' response to income risk.

To compare the model and the data in this dimension, we use the estimated liquidity ratio $\lambda(prc,t) = \frac{\lambda^{LI}(prc,t)}{\lambda^{IL}(prc,t)}$ by each percentile, prc, of net worth for each SCF survey year t, of which we use the time averages for the calibration of the model. Here, we use their variation over time; see Appendix E.2 on how they are constructed.

We calculate the average shock in the year preceding an SCF wave, $\bar{\epsilon}^s_{t-1}$, and regress the liquidity ratio of all percentiles (above the 20th) on the shock, an intercept and a linear time trend:

$$\lambda(prc,t) = \gamma_0(prc) + \gamma_1(prc)t + \gamma_2(prc)\bar{\epsilon}^s_{t-1} + \zeta,$$

i.e., we use a local projection technique. Figure 11 shows the coefficients, $\gamma_2(prc)$, of the uncertainty shock for this regression. Again, we use a block bootstrap to estimate confidence bands. The poorer a household, the stronger its increase in liquidity holdings. Compared to our model, the increase in liquidity for the poor is even stronger. One reason might be that poorer households, since they hold a mortgage on a house, are more highly leveraged in the data than households in our model.

8 Conclusion

This paper examines how variations in the riskiness of household income affect the macroeconomy through precautionary savings. For this purpose, we develop a novel and tractable framework that combines nominal rigidities and incomplete markets in which households choose portfolios of liquid paper and illiquid physical assets – thereby embedding incomplete markets with wealthy hand-to-mouth consumers in a New Keynesian setup. In this model, higher income risk triggers a flight to liquidity because the liquid asset is better suited for short-run consumption smoothing. This reduces not only consumption but also investment, and hence depresses economic activity.

Calibrating the model to match the evolution of uncertainty about household income in the U.S., we find that, in line with the empirical evidence, a one standard deviation spike in income uncertainty leads to losses in output, consumption, and investment. The decline in aggregate activity predicted by our model becomes sizable at the zero lower bound. This may help us to understand the severity of the Great Recession, for which we document a shift toward liquid assets across all percentiles of the U.S. wealth distribution.

The welfare effects of such uncertainty shocks crucially depend on a household's asset position and the stance of monetary and fiscal policy. Monetary policy that drastically lowers the return on liquid assets in times of increased uncertainty limits the negative welfare effects of uncertainty shocks but redistributes resources from liquid to illiquid asset holders – both of which are typically wealthy. Fiscal policy can similarly ameliorate the uncertainty shock by providing the liquid asset that households demand. This keeps the return on all assets high at the expense of lower physical capital and wages in the future.

Our model abstracts from a number of features that may make the impact of an uncertainty shock even stronger than what we highlight. The two most important features are the absence of a firm sector that has a motive to hold liquidity (see, e.g., Negro et al., 2017) and a richer model of banking and mortgage financing. Mortgages constitute the most important asset, besides government bonds, that banks use to create inside money. We have sketched the effect it has when households in times of high uncertainty demand less illiquid assets and hence write fewer mortgages, depressing the supply of liquid assets. Then not only does household's demand for liquid assets go up with uncertainty but also the supply of liquid assets from the private sector goes down. However, it would require going beyond our current model with two assets to obtain a detailed account of these effects. Similarly, in future work it might be interesting to investigate the effect of changes in higher moments of income shocks along the lines of Guvenen et al. (2014b).

References

Ahn, S., Kaplan, G., Moll, B., Winberry, T., and Wolf, C. (2017). When inequality matters for macro and macro matters for inequality. mimeo, U Chicago.

Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. Quarterly Journal of Economics, 109(3):659–684.

Arellano, C., Bai, Y., and Kehoe, P. (2012). Financial markets and fluctuations in uncertainty. Federal Reserve Bank of Minneapolis Research Department Staff Report 466.

Auclert, A. (2015). Monetary policy and the redistribution channel. mimeo, MIT.

Bachmann, R. and Bayer, C. (2013). Wait-and-see' business cycles? Journal of Monetary Economics, 60(6):704–719.

Basu, S. and Bundick, B. (2012). Uncertainty shocks in a model of effective demand. NBER Working Paper Series, 18420.

Baumol, W. J. (1952). The transactions demand for cash: An inventory theoretic approach. Quarterly Journal of Economics, 66(4):pp. 545–556.

Bewley, T. (1980). Interest bearing money and the equilibrium stock of capital. mimeo.

Bi, H., Leeper, E., and Leith, C. (2013). Uncertain fiscal consolidations. Economic Journal, 123(566):31–63.

Bloom, N. (2009). The impact of uncertainty shocks. Econometrica, 77(3):623–685.

Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption inequality and partial insurance. American Economic Review, 98(5):1887–1921.

Born, B. and Pfeifer, J. (2014). Policy risk and the business cycle. Journal of Monetary Economics, 68:68–85.

Carroll, C. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. Economics Letters, 91(3):312–320.

Challe, E. and Ragot, X. (2016). Precautionary saving over the business cycle. Economic Journal, 126(590):135–164.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2011). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. American Economic Review, 101(3):471–475.

Christiano, L., Motto, R., and Rostagno, M. (2010). Financial factors in economic fluctuations. ECB Working Paper 1192.

Chugh, S. (2012). Firm risk and leverage-based business cycles. mimeo, University of Maryland.

Den Haan, W., Riegler, M., and Rendahl, P. (2015). Unemployment (fears) and deflationary spirals. CPER Discussion Paper DP10814.

Den Haan, W. J. (2010). Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents. Journal of Economic Dynamics and Control, 34(1):79–99.

Di Tella, S. (2016). Uncertainty shocks and balance sheet recessions. Journal of Political Economy, page forthcoming.

Fella, G. (2014). A generalized endogenous grid method for non-smooth and non-concave problems. Review of Economic Dynamics, 17(2):329 – 344.

Fernald, J. et al. (2012). A quarterly, utilization-adjusted series on total factor productivity. Federal reserve bank of San Francisco working paper, 19:2012.

Fernández-Villaverde, J., Guerrón-Quintana, P. A., Kuester, K., and Rubio-Ramírez, J. (2015). Fiscal volatility shocks and economic activity. American Economic Review, 105(11):3352–84.

Gilchrist, S., Sim, J. W., and Zakrajšek, E. (2014). Uncertainty, financial frictions, and investment dynamics. NBER Working Paper Series, 20038.

Gornemann, N., Kuester, K., and Nakajima, M. (2012). Monetary policy with heterogeneous agents. Federal Reserve Bank of Philadelphia, WP No. 12-21.

Greenwood, J., Hercowitz, Z., and Huffman, G. W. (1988). Investment, capacity utilization, and the real business cycle. American Economic Review, 78(3):402–417.

Guerrieri, V. and Lorenzoni, G. (2011). Credit crises, precautionary savings, and the liquidity trap. NBER Working Paper Series, 17583.

Guvenen, F., Kaplan, G., and Song, J. (2014a). The glass ceiling and the paper floor: Gender differences among top earners, 1981-2012. NBER Working Paper Series, 20560.

Guvenen, F., Ozkan, S., and Song, J. (2014b). The nature of countercyclical income risk. Journal of Political Economy, 122(3):621–660.

Hall, R. E. (1989). Invariance properties of Solow's productivity residual. NBER Working Paper Series, 3034.

Hintermaier, T. and Koeniger, W. (2010). The method of endogenous gridpoints with occasionally binding constraints among endogenous variables. Journal of Economic Dynamics and Control, 34(10):2074–2088.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. Journal of Economic Dynamics and Control, 17(5-6):953–969.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. American Economic Review, 95(1):161–182.

Kaplan, G., Moll, B., and Violante, G. L. (2016). Monetary policy according to HANK. NBER Working Paper Series, 21897.

Kaplan, G. and Violante, G. L. (2010). How much consumption insurance beyond self-insurance? American Economic Journal: Macroeconomics, 2(4):53–87.

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. Econometrica, 82:1199–1239.

King, R. G., Plosser, C. I., and Rebelo, S. T. (1988). Production, growth and business cycles: I. The basic neoclassical model. Journal of Monetary Economics, 21(2):195 – 232.

Krusell, P. and Smith, A. A. (1997). Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns. Macroeconomic Dynamics, 1(02):387–422.

Krusell, P. and Smith, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. Journal of Political Economy, 106(5):867–896.

Leith, C. and von Thadden, L. (2008). Monetary and fiscal policy interactions in a New Keynesian model with capital accumulation and non-Ricardian consumers. Journal of Economic Theory, 140(1):279–313.

Luetticke, R. (2015). Transmission of monetary policy with heterogeneity in household portfolios. mimeo, University of Bonn.

McKay, A., Nakamura, E., and Steinsson, J. (2016). The power of forward guidance revisited. American Economic Review, 106(10):3133–3158.

McKay, A. and Reis, R. (2016). The role of automatic stabilizers in the us business cycle. Econometrica, 84(1):141–194.

Narita, F. (2011). Hidden actions, risk-taking and uncertainty shocks. mimeo, University of Minnesota.

Negro, M. D., Eggertsson, G., Ferrero, A., and Kiyotaki, N. (2017). The great escape? a quantitative evaluation of the Fed's liquidity facilities. American Economic Review, 107(3):824–857.

Panousi, V. and Papanikolaou, D. (2012). Investment, idiosyncratic risk, and ownership. Journal of Finance, 67(3):1113–1148.

Ravn, M. O. and Sterk, V. (2013). Job uncertainty and deep recessions. mimeo, UCL.

Reiter, M. (2002). Recursive computation of heterogeneous agent models. mimeo, Universitat Pompeu Fabra.

Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. Journal of Economic Dynamics and Control, 33(3):649–665.

Reiter, M. (2010). Approximate and almost-exact aggregation in dynamic stochastic heterogeneous-agent models. Economics Series 258, Institute for Advanced Studies.

Romei, F. (2015). Need for (the right) speed: the timing and composition of public debt deleveraging. mimeo.

Rotemberg, J. J. (1982). Sticky prices in the united states. Journal of Political Economy, pages 1187–1211.

Santos, M. S. (1991). Smoothness of the policy function in discrete time economic models. Econometrica, 59(5):1365–1382.

Schaal, E. (2012). Uncertainty, productivity and unemployment in the great recession. Federal Reserve Bank of Minneapolis, mimeo.

Schmitt-Grohé, S. and Uribe, M. (2004). Solving dynamic general equilibrium models using a second-order approximation to the policy function. Journal of Economic Dynamics and Control, 28(4):755–775.

Smets, F. and Wouters, R. (2007). Shocks and frictions in us business cycles: A bayesian DSGE approach. American Economic Review, 97(3):586–606.

Stokey, N. and Lucas, R. (1989). Recursive Methods in Economic Dynamics. Harvard University Press.

Storesletten, K., Telmer, C. I., and Yaron, A. (2001). The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk. European Economic Review, 45(7):1311–1339.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Cyclical dynamics in idiosyncratic labor market risk. Journal of Political Economy, 112(3):695–717.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. Economics Letters, 20(2):177 – 181.

Taylor, J. B. (1993). Discretion versus policy rules in practice. In Carnegie-Rochester conference series on public policy, volume 39, pages 195–214. Elsevier.

Tobin, J. (1956). The interest-elasticity of transactions demand for cash. Review of Economics and Statistics, 38(3):pp. 241–247.

Vavra, J. (2014). Inflation dynamics and time-varying volatility: New evidence and an Ss interpretation. Quarterly Journal of Economics, 129(1):215–258.

Werning, I. (2015). Incomplete markets and aggregate demand. NBER Working Paper Series, 21448.

Winberry, T. (2016). A toolbox for solving and estimating heterogeneous agent macro models. mimeo, Chicago Booth.

A Dynamic Planning Problem with Two Assets and Fixed Adjustment Probabilities

The dynamic planning problem of a household in the model is characterized by two Bellman equations, V_a in the case where the household can adjust its capital holdings and V_n otherwise. We will first go through the problem with exogenous adjustment probabilities, as the first-order conditions of the model with adjustment decisions that describe portfolio and consumption choices turn out to be of the same structure as under given adjustment probabilities.

With fixed adjustment probabilities, the value functions are given by

$$\begin{aligned} V_a(b,k,h;R_b,\Theta,s) &= \max_{k',b'_a \in \Gamma_a} u[x(b,b'_a,k,k',h)] \\ &\quad + \beta\left[\nu EV_a(b'_a,k',h';R'_b,\Theta',s') + (1-\nu)EV_n(b'_a,k',h';R'_b,\Theta',s')\right] \\ V_n(b,k,h;R_b,\Theta,s) &= \max_{b'_n \in \Gamma_n} u[x(b,b'_n,k,k,h)] \qquad (26) \\ &\quad + \beta\left[\nu EV_a(b'_n,k,h';R'_b,\Theta',s') + (1-\nu)EV_n(b'_n,k,h';R'_b,\Theta',s')\right] \end{aligned}$$

where the budget sets are given by

$$\begin{aligned} \Gamma_a(b,k,h;R_b,\Theta,s) &= \Big\{k' \geq 0, b' \geq \underline{B} \mid q(\Theta,s)(k'-k) + b' \leq r(\Theta,s)k \qquad (27) \\ &\quad + \frac{1+R_b(\Theta,s)}{1+\pi(\Theta,s)}b + (1-\tau)\left(\frac{\gamma}{1+\gamma}w(\Theta,s)hN + \mathbb{I}_{h=0}\Pi(\Theta,s)\right)\Big\} \\ \Gamma_n(b,k,h;R_b,\Theta,s) &= \Big\{b' \geq \underline{B} \mid b' \leq +r(\Theta,s)k \qquad (28) \\ &\quad + \frac{1+R_b(\Theta,s)}{1+\pi(\Theta,s)}b + (1-\tau)\left(\frac{\gamma}{1+\gamma}w(\Theta,s)hN + \mathbb{I}_{h=0}\Pi(\Theta,s)\right)\Big\} \\ x(b,b',k,k',h;R_b,\Theta,s) &= \frac{\gamma}{1+\gamma}w(\Theta,s)hN + r(\Theta,s)k + \frac{(1+R_b)}{1+\pi(\Theta,s)}b \qquad (29) \\ &\quad - q(\Theta,s)(k'-k) - b' \end{aligned}$$

To save on notation, let Ω be the set of possible idiosyncratic state variables controlled by the household, let Z be the set of potential aggregate states, let $\Gamma_i : \Omega \to \Omega$ be the correspondence describing the feasibility constraints, and let $A_i(z) = \{(\omega, y) \in \Omega \times \Omega : y \in \Gamma_i(\omega, z)\}$ be the graph of Γ_i. Hence the states and controls of the household problem can be defined as

$$\Omega = \{\omega = (b,k) \in \mathbb{R}^2 : \underline{B} \leq b < \infty, 0 \leq k < \infty\} \tag{30}$$

$$z = \{h, \Theta, s\} \tag{31}$$

and the return function $F : A \to R$ reads:

$$F(\Gamma_i(\omega, z), \omega; z) = \frac{x_i^{1-\gamma}}{1-\gamma} \tag{32}$$

Define the value before the adjustment/non-adjustment shock is realized as

$$v(\omega, z) := \nu V_a(\omega, z) + (1-\nu) V_n(\omega, z).$$

Now we can rewrite the optimization problem of the household in terms of the definitions above in a compact form:

$$V_a(\omega, z) = \max_{y \in \Gamma_a(\omega, z)} [F(\omega, y; z) + \beta E v(y, z')] \tag{33}$$

$$V_n(\omega, z) = \max_{y \in \Gamma_n(\omega, z)} [F(\omega, y; z) + \beta E v(y, z')]. \tag{34}$$

Finally we define the mapping $T : C(\Omega) \to C(\Omega)$, where $C(\Omega)$ is the space of bounded, continuous and weakly concave functions.

$$\begin{aligned} (Tv)(\omega, z) &= \nu V_a(\omega, z) + (1-\nu) V_n(\omega, z) \\ V_a(\omega, z) &= \max_{y \in \Gamma_a(\omega, z)} [F(\omega, y; z) + \beta E v(y, z')] \\ V_n(\omega, z) &= \max_{y \in \Gamma_n(\omega, z)} [F(\omega, y; z) + \beta E v(y, z')]. \end{aligned} \tag{35}$$

A.1 Properties of Primitives

The following properties of the primitives of the problem obviously hold:

P 1. *Properties of sets* $\Omega, \Gamma_a(\omega, z), \Gamma_n(\omega, z)$

1. Ω *is a convex subset of* R^3.
2. $\Gamma_i(\cdot, z) : \Omega \to \Omega$ *is non-empty, compact-valued, continuous, monotone and convex for all* z.

P 2. *Properties of return function* F

F is bounded, continuous, strongly concave, C^2 differentiable on the interior of A, and strictly increasing in each of its first two arguments.

A.2 Properties of the Value and Policy Functions

Lemma 1. *The mapping T defined by the Bellman equation for v fulfills Blackwell's sufficient conditions for a contraction on the set of bounded, continuous and weakly concave functions $C(\Omega)$.*

a) It satisfies discounting.

b) It is monotonic.

c) It preserves boundedness (assuming an arbitrary maximum consumption level).

d) It preserves strict concavity.

Hence, the solution to the Bellman equation is strictly concave. The policy is a single-valued function in (b, k), and so is optimal consumption.

Proof. The proof proceeds item by item and closely follows Stokey and Lucas (1989) taking into account that the household problem in the extended model consists of two Bellman equations.

a) Discounting
Let $a \in R_+$ and the rest be defined as above. Then it holds that:

$$\begin{aligned}(T(v+a))(\omega,z) =&\nu \max_{y\in\Gamma_a(\omega,z)}[F(\omega,y,z)+\beta Ev(y,z')+a]\\ &+(1-\nu)\max_{y\in\Gamma_n(\omega,z)}[F(\omega,y,z)+\beta Ev(y,z')+a]\\ =&(Tv)(\omega,z)+\beta a\end{aligned}$$

Accordingly, T fulfills discounting.

b) Monotonicity
Let $g : \Omega \times Z \to R^2$, $f : \Omega \times Z \to R^2$ and $g(\omega, z) \geq f(\omega, z)$ $\forall \omega, z \in \Omega \times Z$, then it

follows that:

$$
\begin{aligned}
(Tg)(\omega, z) =& \nu \max_{y\in\Gamma_a(\omega,z)} [F(\omega, y, z) + \beta Eg(y, z')] \\
&+ (1-\nu) \max_{y\in\Gamma_n(\omega,z)} [F(\omega, y, z) + \beta Eg(y, z')] \\
\geq& \nu \max_{y\in\Gamma_a(\omega,z)} [F(\omega, y, z) + \beta Ef(y, z')] \\
&+ (1-\nu) \max_{y\in\Gamma_n(\omega,z)} [F(\omega, y, z) + \beta Ef(y, z')] \\
=& Tf(\omega, z)
\end{aligned}
$$

The objective function for which Tg is the maximized value is uniformly higher than the function for which Tf is the maximized value. Therefore, T preserves monotonicity.

c) Boundedness
From properties **P1** it follows that the mapping T defines a maximization problem over the continuous and bounded function $[F(\omega, y) + \beta Ev(y, z'))]$ over the compact sets $\Gamma_i(\omega, z)$ for $i = \{a, n\}$. Hence the maximum is attained. Since F and v are bounded, Tv is also bounded.

d) Strict Concavity
Let $f \in C''(\Omega)$, where C'' is the set of bounded, continuous, strictly concave functions on Ω. Since the convex combination of two strictly concave functions is strictly concave, it is sufficient to show that $T_i[C''(\Omega)] \subseteq C''(\Omega)$, where T_i is defined by

$$
T_i v = \max_{y\in\Gamma_i(\omega,z)} [F(\omega, y, z) + \beta Ev(y, z')], i \in \{a, n\}
$$

Let $\omega_0 \neq \omega_1, \theta \in (0, 1), \omega_\theta = \theta\omega_0 + (1-\theta)\omega_1$.
Let $y_j \in \Gamma_i(\omega_j, z)$ be the maximizer of $(T_i f)(\omega_j)$ for $j = \{0, 1\}$ and $i = \{a, n\}$, $y_\theta = \theta y_0 + (1-\theta)y_1$.

$$
\begin{aligned}
(T_i f)(\omega_\theta, z) \geq& [F(\omega_\theta, y_\theta, z) + \beta Ef(y_\theta, z')] \\
>& \theta[F(\omega_0, y_0, z) + \beta Ef(y_0, z'))] + (1-\theta)[F(\omega_1, y_1, z) + \beta Ef(y_1, z')] \\
=& \theta(Tf)(\omega_0, z) + (1-\theta)(Tf)(\omega_1, z)
\end{aligned}
$$

The first inequality follows from y_θ being feasible because of convex budget sets.

The second inequality follows from the strict concavity of f. Since ω_0 and ω_1 are arbitrary, it follows that $T_i f$ is strictly concave, and since f is arbitrary that $T[C''(\Omega)] \subseteq C''(\Omega)$.

□

Lemma 2. *The value function is C^2 and the policy function C^1 differentiable.*

Proof. The properties of the choice set **P1**, of the return function **P2**, and the properties of the value function proven in (1) fulfill the assumptions of Santos's (1991) theorem on the differentiability of the policy function. According to the theorem, the value function is C^2 and the policy function C^1 differentiable.
Note that strong concavity of the return function holds for CRRA utility, because of the arbitrary maximum we set for consumption. □

Lemma 3. *The total savings $S_i^* := b_i^*(\omega, z) + q(z)k_i^*(\omega, z)$ and consumption $c_i^*, i \in \{a, n\}$ are increasing in ω if $r(z)$ is positive. In the adjustment case, total savings and consumption are increasing in total resources $R_a(z) = [q(z) + r(z)]k + b\frac{1+R_b(z)}{1+\pi(z)}$ for any $r(z)$.*

Proof. Define $\tilde{v}(S, z) := \max_{\{b,k|b+q(z)k \leq S\}} Ev(b, k; z')$ and resources in the case of no adjustment $R_n = r(z)k + b\frac{1+R_b(z)}{1+\pi(z)}$. Since v is strictly concave and increasing, so is $\tilde{v}$ by the line of the proof of Lemma 1.d). Denote $\varphi(z) = (1 - \tau)\left(\frac{\gamma}{1+\gamma}w(z)hN + \mathbb{I}_{h=0}\Pi(z)\right)$. Now we can (re)write the planning problem as

$$V_a(b, k; z) = \max_{S \leq \varphi(z)+R_a} \left[u\left(\varphi(z) + [q(z) + r(z)]k + b\frac{1 + R_b(z)}{1 + \pi(z)} - S\right) + \beta\tilde{v}(S, z)\right]$$
$$V_n(b, k; z) = \max_{b' \leq \varphi(z)+R_n} \left[u\left(\varphi(z) + r(z)k + b\frac{1 + R_b(z)}{1 + \pi(z)} - b'\right) + \beta Ev(b', k; z')\right].$$

Due to differentiability we obtain the following (sufficient) first-order conditions:

$$\frac{\partial u\left(\varphi(z) + [q(z) + r(z)]k + b\frac{1+R_b(z)}{1+\pi(z)} - S\right)}{\partial c} = \beta\frac{\partial\tilde{v}(S, z)}{\partial S}$$
$$\frac{\partial u\left(\varphi(z) + r(z)k + b\frac{1+R_b(z)}{1+\pi(z)} - b'\right)}{\partial c} = \beta\frac{\partial v(b', k; z)}{\partial b'}. \tag{36}$$

Since the left-hand sides are decreasing in $\omega = (b, k)$, and increasing in S (respectively b'), and the right-hand side is decreasing in S (respectively b'), $S_i^* = \begin{cases} qk' + b' \text{ if } i = a \\ qk + b' \text{ if } i = n \end{cases}$

must be increasing in ω.
Since the right-hand side of (36) is hence decreasing in ω, so must be the left-hand side of (36). Hence consumption must be increasing in ω.
The last statement follows directly from the same proof. □

A.3 Euler Equations

Denote the optimal policies for consumption, for bond holdings and capital as $x_i^*, b_i^*, k^*, i \in \{a, n\}$ respectively. The first-order conditions for an inner solution in the (non-)adjustment case read:

$$k^* : \frac{\partial u(x_a^*)}{\partial x} q = \beta E \left[\nu \frac{\partial V_a(b_a^*, k^*; z')}{\partial k} + (1-\nu) \frac{\partial V_n(b_a', k'; z')}{\partial k} \right] \tag{37}$$

$$b_a^* : \frac{\partial u(x_a^*)}{\partial x} = \beta E \left[\nu \frac{\partial V_a(b_a^*, k^*; z')}{\partial b} + (1-\nu) \frac{\partial V_n(b_a^*, k^*; z')}{\partial b} \right] \tag{38}$$

$$b_n^* : \frac{\partial u(x_n^*)}{\partial x} = \beta E \left[\nu \frac{\partial V_a(b_n^*, k; z')}{\partial b} + (1-\nu) \frac{\partial V_n(b_n^*, k; z')}{\partial b} \right] \tag{39}$$

Note the subtle difference between (38) and (39), which lies in the different capital stocks k' vs. k in the right-hand side expressions.

Differentiating the value functions with respect to k and m, we obtain:

$$\frac{\partial V_a(b, k; z)}{\partial k} = \frac{\partial u[x_a^*(b, k; z)]}{\partial x} (q(z) + r(z)) \tag{40}$$

$$\frac{\partial V_a(b, k; z)}{\partial b} = \frac{\partial u[x_a^*(b, k; z)]}{\partial x} \frac{1 + R^b(z)}{1 + \pi(z)} \tag{41}$$

$$\frac{\partial V_n(b, k; z)}{\partial b} = \frac{\partial u[x_n^*(b, k; z)]}{\partial x} \frac{1 + R^b(z)}{1 + \pi(z)} \tag{42}$$

$$\begin{aligned} \frac{\partial V_n(b, k; z)}{\partial k} &= r(z) \frac{\partial u[x_n^*(b, k; z)]}{\partial x} \\ &+ \beta E \left[\nu \frac{\partial V_a[b_n^*(b, k; z), k; z']}{\partial k} + (1-\nu) \frac{\partial V^n[b_n^*(b, k; z), k; z']}{\partial k} \right] \\ &= r(z) \frac{\partial u[x_n^*(b, k; z)]}{\partial x} + \beta \nu E \frac{\partial u\{x_a^*[b_n^*(b, k; z), k; z], k; z'\}}{\partial x} (q(z') + r(z')) \\ &+ \beta(1-\nu) E \frac{\partial V_n\{[b_n^*(b, k; z), k; z], k; z'\}}{\partial k} \end{aligned} \tag{43}$$

such that the marginal value of capital in non-adjustment is defined recursively.

Now we can plug the second set of equations into the first set of equations and obtain the following Euler equations (in slightly shortened notation):

$$\frac{\partial u[x_a^*(b,k;z)]}{\partial x}q(z) = \beta E\left[\nu\frac{\partial u[x_a^*(b_a^*,k^*;z')]}{\partial x}[q(z')+r(z')]+(1-\nu)\frac{\partial V^n(b_a^*,k';z')}{\partial k'}\right] \tag{44}$$

$$\frac{\partial u[x_a^*(b,k;z)]}{\partial x} = \beta E\frac{1+R^b(z')}{1+\pi(z')}\left[\nu\frac{\partial u[x_a^*(b_a^*,k^*;z')]}{\partial x}+(1-\nu)\frac{\partial u[x_n^*(b_a^*,k';z')]}{\partial x}\right] \tag{45}$$

$$\frac{\partial u[x_n^*(b,k,;z)]}{\partial x} = \beta E\frac{1+R^b(z')}{1+\pi(z')}\left[\nu\frac{\partial u[x_a^*(b_n',k;z')]}{\partial x}+(1-\nu)\frac{\partial u[x_n^*(b_n^*,k;z')]}{\partial x}\right] \tag{46}$$

In words, when deciding between the liquid and the illiquid asset, the household compares the one-period return difference between the two assets $E\frac{1+R^b(z')}{1+\pi(z')} - E\frac{r(z')+q(z')}{q(z)}$ weighted with the marginal utility under adjustment and the probability of adjustment and the difference between the return in the no adjustment case, $E\frac{1+R^b(z')}{1+\pi(z')}\frac{\partial u[x_n^*(b_a^*,k';z')]}{\partial x}$, and the marginal value of illiquid assets when not adjusting $\frac{\partial V^n(b_a^*,k';z')}{\partial k'}$. The latter reflects both the utility derived from the dividend stream and the utility from occasionally selling the asset.

A.4 Algorithm

The algorithm we use to solve for optimal policies given the Krusell-Smith forecasting rules is a version of Hintermaier and Koeniger's (2010) extension of the endogenous grid method, originally developed by Carroll (2006).

It works iteratively until convergence of policies as follows: Start with some guess for the policy functions x_a^* and x_n^* on a given grid $(b,k)\in B\times K$. Define the shadow value of capital

$$\begin{aligned}\beta^{-1}\psi(b,k;z) :=& \nu E\left\{\frac{\partial u\{x_a^*[b_n^*(b,k,z),k;z']\}}{\partial x}[q(z')+r(z')]\right\} \\ &+(1-\nu)E\frac{\partial V_n[b_n^*(b,k,z),k;z']}{\partial k} \\ &= \nu E\left\{\frac{\partial u\{x_a^*[b_n^*(b,k,z),k;z']\}}{\partial x}[q(z')+r(z')]\right\} \\ &+(1-\nu)E\left\{\frac{\partial u\{x_n^*(b_n^*(b,k,z),k;z')]}{\partial x}r(z')\right\} \\ &+(1-\nu)E\left\{\psi[b_n^*(b,k,z),k;z']\right\}.\end{aligned} \tag{47}$$

Guess initially $\psi = 0$. Then

1. Solve for an update of x_n^* by standard endogenous grid methods using equation (46), and denote $b_n^*(b, k; z)$ as the optimal bond holdings without capital adjustment.

2. Find for every k' on-grid some (off-grid) value of $\tilde{b}_a^*(k'; z)$ such that combining (45) and (44) yields:

$$0 = \nu E \left\{ \frac{\partial u[x_a^*(\tilde{b}_a^*(k', z), k'; z')]}{\partial x} \left[\frac{q(z') + r(z')}{q(z)} - \frac{1 + R^b(z')}{1 + \pi(z')} \right] \right\} \tag{48}$$
$$+ (1 - \nu) E \left\{ \frac{\partial u[x_n^*(\tilde{b}_a^*(k', z), k'; z')]}{\partial x} \left[\frac{r(z')}{q(z)} - \frac{1 + R^b(z')}{1 + \pi(z')} \right] \right\} + (1 - \nu) E \left[\frac{\psi(\tilde{b}_a^*(k', z), k'; z')}{q(z)} \right]$$

N.B. that $E\psi$ takes the stochastic transitions in h' into account and does not replace the expectations operator in the definition of ψ. If no solution exists, set $\tilde{b}_a^* = \underline{B}$. Uniqueness (conditional on existence) of $\tilde{b}_a^*$ follows from the strict concavity of v.

3. Solve for total initial resources, by solving the Euler equation (45) for $\tilde{x}^*(k', z)$, such that:

$$\tilde{x}^*(k', z)$$
$$= \frac{\partial u}{\partial x}^{-1} \left\{ \beta E \frac{1 + R^b(z')}{1 + \pi(z')} \left[\nu \frac{\partial u\{x_a^*[b_a^*(k', z), k'; z']\}}{\partial x} + (1 - \nu) \frac{\partial u\{x_n^*[b_a^*(k', z), k'; z']}{\partial x} \right] \right\} \tag{49}$$

where the right-hand side expressions are obtained by interpolating $x_a^*(b_a^*(k', z), k', z')$ from the on-grid guesses $x_a^*(b, k; z)$ and taking expected values with respect to z'.

This way we obtain total non-human resources $\tilde{R}_a(k', z)$ that are compatible with plans $(b^*(k'), k')$ and a consumption policy $\tilde{\tilde{x}}_a^*(\tilde{R}_a(k', z), z)$ in total resources.

4. Since (consumption) policies are increasing in resources, we can obtain consumption policy updates as follows: Calculate total resources for each (b, k) pair $R_a(b, k) = (q+r)k + b\frac{1+R^b}{1+\pi}$ and use the consumption policy obtained before to update $x_a^*(b, k, z)$ by interpolating at $R_a(b, k)$ from the set $\left\{ (\tilde{\tilde{x}}_a^*(\tilde{R}_a(k', z), z), R_a(k', z)) \middle| k' \in K \right\}$.[24]

[24] If a boundary solution $\tilde{b}^*(\underline{B}) > \underline{B}$ is found, we use the "n" problem to obtain consumption policies for resources below $\tilde{b}^*(\underline{B})$.

5. Update ψ: Calculate a new value of ψ using (43), such that:

$$\begin{aligned}\psi^{new}(b,k,z) =& \beta\nu E\left\{\frac{\partial u\{x_a^*[b_n^*(b,k,z),k;z']\}}{\partial x}[q(z')+r(z')]\right\}\\ &+\beta(1-\nu)E\left\{\frac{\partial u\{x_n^*(b_n^*(b,k,z),k;z')]}{\partial x}r(z')\right\}\\ &+\beta(1-\nu)E\left\{\psi^{old}[b_n^*(b,k,z),k;z']\right\}. \end{aligned} \tag{50}$$

making use of the **updated** consumption policies.

B Dynamic Planning Problem with Two Assets and Logistic Distribution of Adjustment Costs

With logistically distributed adjustment costs, concavity of the value function is no longer guaranteed, because ν will depend on $(\omega; z)$. If the function EV in equation (10) is convex, then the policy functions will still be continuously differentiable and the value function twice differentiable because the prerequisites of Lemma 2 and 3 are still fulfilled.

Let $f(\chi)$ be the density function of the adjustment costs. Since $V_a \geq V_n$ we can write

$$EV(\omega;z) = V_n(\omega;z) + \int_0^{V_a(\omega;z)-V_n(\omega;z)} (V_a(\omega;z) - V_n(\omega;z) - \chi)f(\chi)d\chi \tag{51}$$

In turn if $f(\chi) > 0$ for all $\chi > 0$ (the adjustment cost distribution has unbounded support on $\mathbb{R}_+$), the derivative of EV w.r.t. ω takes the form

$$\frac{\partial EV}{\partial \omega} = \frac{\partial V_n}{\partial \omega} + \nu^*(\omega;z)\left[\frac{\partial V_a}{\partial \omega} - \frac{\partial V_n}{\partial \omega}\right]. \tag{52}$$

In words, first-order conditions of a model with fixed adjustment probabilities and a model with state dependent adjustment probabilities are the same. We make use of this fact and simply replace the state independent adjustment probability by a guess for an adjustment probability function in the algorithm described in Appendix A.4. We then update the adjustment probabilities by making use of the closed-form solution to the expected adjustment costs under the logistic distribution assumption for χ when calculating the value functions in iteration (n):

$$V^a_{(n)} = u(x^a_{(n)}) + \beta EV_{(n)}(b^a, k^a, h'; z) \tag{53}$$

$$V^n_{(n)} = u(x^n_{(n)}) + \beta EV_{(n)}(b^n, k, h'; z) \tag{54}$$

where

$$EV_{(n)} = \nu^*_{(n)} V^a_{(n)} + (1 - \nu^*_{(n)}) V^n_{(n)} - AC(\nu^*; \mu_\chi, \sigma_\chi) \tag{55}$$

where μ_χ and σ_χ are the mean and the scale of the logistic distribution

$$F(\chi) = \frac{1}{1 + \exp\{-\frac{\chi - \mu_\chi}{\sigma_\chi}\}}.$$

The adjustment probability can be updated after the two value functions have been calculated for a given $\nu^*(\omega, z)$ as

$$\nu^*_{(n+1)}(\omega, z) = F[V^a_{(n)}(\omega, z) - V^n_{(n)}(\omega, z)].$$

Given the new adjustment probabilities, consumption and savings policies can be determined again using the endogenous grid method. The expected conditional adjustment cost is given by

$$\begin{aligned} AC(\nu; \mu_\chi, \sigma_\chi) &= \int_0^{F^{-1}(\nu)} \chi dF(\chi) = \int_0^\nu F^{-1}(p) dp \\ &= \int_0^\nu \mu_\chi + \sigma_\chi [\log p - \log(1-p)] dp \\ &= \mu_\chi \nu + \sigma_\chi \left[\nu \log \nu + (1-\nu)\log(1-\nu)\right] \end{aligned}$$

Given that concavity of the value functions is not guaranteed, we check for monotonicity of the derivatives of the value function and for uniqueness of the optimal portfolio solution in the algorithm, implementing thereby a version of Fella's (2014) algorithm, and find that the solution turns out to be globally concave.

C Solving the Model with Aggregate Shocks

C.1 Local Approximation

Our model has a three-dimensional idiosyncratic state space with two endogenous states. We experimented with the grid size for liquid and illiquid asset holdings as well as for the process of productivity. Given that we focus on second moment changes we require $n_h = 26$ productivity states and find that with a log-spaced grid for assets results are no longer affected by grid size beyond $n_b = 80, n_k = 80$ points. This means that a full grid contains $n_b \times n_k \times n_h = 166,400$ points. This renders solving the model by perturbing the histogram and the value functions on a full grid infeasible such that we cannot apply a perturbation method without state-space reduction, such as in Reiter (2002).

Instead, we develop a variant of Reiter's (2009) method to solve heterogeneous agent models with aggregate risk. We represent the dynamic system as a set of non-linear difference equations, for which hold

$$E_t F\left(X_t, X_{t+1}, Y_t, Y_{t+1}\right) = 0,$$

where the set of control variables is $Y_t = (V_t, \frac{\partial V_t}{\partial b}, \frac{\partial V_t}{\partial k}, \tilde{Y}_t)$, i.e., value functions and their marginals with respect to k, b as well as some aggregate controls $\tilde{Y}_t$ such as dividends, wages, etc. The set of state variables $X_t = (\Theta_t, R_t^b, s_t)$ is given by the histogram Θ_t of the distribution over (b, k, h) and the aggregate states R_t^b, s_t. In principle, we can solve this system with Schmitt-Grohé and Uribe's (2004) method as argued in Reiter (2002), but in practice the state space is too rich and the solution becomes numerically infeasible and unstable.

Hence, we need to reduce the dimensionality of the system. We therefore first approximate value functions and their derivatives at all grid points around their value in the stationary equilibrium without aggregate risk, $V^{SS}(b, k, h)$, by a sparse polynomial $P(b, k, h)$ with parameters $\Omega_t = \Omega(R_t^b, \Theta_t, s_t)$. For example, we write the value function as

$$V(b, k, h; R_t^b, \Theta_t, s_t)/V^{SS}(b, k, h) \approx P(b, k, h)\Omega_t.$$

Note the difference to a global approximation of the functions for finding the stationary equilibrium without aggregate risk. Here, we only use the sparse polynomial to capture *deviations* from the stationary equilibrium values, cf. Ahn et al. (2017) and different from Winberry (2016) and Reiter (2009). We define the polynomial basis functions such that the grid points of the full grid coincide with the Chebyshev nodes for this basis.

In the system F, we then use the Bellman equation to obtain V_t from V_{t+1} on the full grid and then calculate the difference of Ω_t to the regression coefficients for the polynomial that fits $V_t(b,k,h)/V^{SS}(b,k,h)$.

This reduces the number of variables in the difference equation substantially, but leaves us still with too many state variables from the histogram at the full grid. Reiter (2010) and Ahn et al. (2017) suggest using state-space reduction techniques to deal with this issue. In continuous time, the state-space reduction can be done based on a Taylor expansion in time derivatives. In discrete time, there is no obvious basis for the state-space reduction.

Yet, using Sklar's Theorem and writing the distribution function in its copula form such that $\Theta_t = \mathcal{C}_t(F_t^b, F_t^k, F_t^h)$ with the copula $\mathcal{C}_t$ and the marginal distributions for liquid and illiquid assets and productivity $F_t^{b,k,h}$. Now fixing $\mathcal{C}_t = \mathcal{C}$ can break the curse of dimensionality, reducing the number of state variables from $n_b \times n_k \times n_h = 166,400$ to $n_b + n_k + n_h = 186$, as we now only need to perturb the marginal distributions.

Fixing the copula $\mathcal{C}$ to the one from the stationary distribution, the approximation does not impose any restriction on the stationary distribution when aggregate shocks are absent, such that the approximation then becomes exact. Therefore, it is less restrictive for the stationary state then assuming a parametric form for the distribution function. The copula itself is obtained by fitting a cubic spline to the stationary distribution of ranks in b, k, h.

The idea behind this approach is that given the economic structure of the model, prices only depend on aggregate asset demand and supply, as in Krusell and Smith (1998), and not directly on higher moments of the *joint* distributions Θ_t, Θ_{t+1}. Our approach imposes no restriction on how the marginal distributions change, i.e., how many more or less liquid assets the portfolios of the x-th percentile have. It only restricts the change in the likelihood of a household being among the x-percent richest in liquid assets to be among the y-percent richest in illiquid assets. We check whether the time-constant copula assumption creates substantial numerical errors and find none by comparing it to the Krusell and Smith (1998) solution. See Figure 12 for a comparison of the IRFs for our baseline calibration.

In addition, we calculate the $\mathbf{R}^2$ statistics for the estimate $\mathcal{C}(F_{t+1}^b, F_{t+1}^k, F_{t+1}^h)$ of distribution Θ_{t+1}:

$$\mathbf{R}^2 = 1 - \frac{\int \left[d\mathcal{C}(F_{t+1}^b, F_{t+1}^k, F_{t+1}^h) - d\Theta_{t+1}\right]^2}{\int \left[d\Theta_{t+1}\right]^2}$$

Table 6: Den Haan (2010) statistic

	Absolute error (in %) for			
	Price of Capital q_t	Capital K_t	Inflation π_t	Real Bonds B_t
Mean	0.04	0.14	0.04	0.99
Max	0.38	0.40	0.22	2.78

Notes: Differences in percent between the simulation of the linearized solution of the model and a simulation in which we solve for the actual intratemporal equilibrium prices in every period for $t = \{1, ..., 1.000\}$; see Den Haan (2010).

plugging in for $F_{t+1}^{b,k,h}$ the linearized solutions $H(F_t^{b,k,h}, R_t^B, s_t)$ and for $d\Theta_{t+1}$ the solution from iterating the histogram forward given the policy functions. This yields a measure of fit for our approximation of the distribution function by a fixed copula. Absent aggregate shocks, the measure is 100% by construction. Given the solution technique, the appropriateness of the fixed copula assumption is captured by the derivative $\frac{\partial \mathbf{R}^2}{\partial x_t}$ of the $\mathbf{R}^2$ statistics with respect to state variable x_t. We find that this derivative is roughly 0.00019% with respect to uncertainty, such that, extrapolating linearly, the $\mathbf{R}^2$ at 99.9999% remains extremely high after a one standard deviation increase in uncertainty (a shock of size 0.54).

Finally, we check the quality of the linearized solution (in aggregate shocks) by solving the household planning problem given the implied expected continuation values from our solution technique but solving for the actual intratemporal equilibrium, as suggested by Den Haan (2010). We simulate the economy over T=1.000 periods and calculate the differences between our linearized solution and the non-linear one. The maximum difference is 0.4% for the capital stock and 2.6% for bonds while the mean absolute errors are substantially smaller; see Table 6.

C.2 Krusell-Smith Equilibrium

Alternatively, we assume that households use forecasting rules to predict future prices on the basis of a restricted set of moments, as in Krusell and Smith (1997, 1998).

Specifically, these rules nowcast inflation, π_t, and capital price, q_t, and forecast the term $\left[\log(1+\pi_{t+1})\frac{Y_{t+1}}{Y_t}\right]$ in the Phillips curve. These rules are used when calculating the continuation values in the Bellman equation. We assume these functions to be log-linear in government debt, B_t, last period's nominal interest rate, R_t^b, the aggregate stock of capital, K_t, average h_{it} (denoted H_t below), and the uncertainty state, s_t (and s_{t+1} for the forecasting term).

We formulate the problem in terms of relative price nowcasts and inflation forecasts such that we have a description of the conditional distributions of all future prices households expect. Note also that it is sufficient to write the problem in terms of *price* nowcasts and the Phillips curve forecast, because given these, households can back out future state variables describing aggregate *quantities*, $\{K_{t+s}, B_{t+s}\}$, from the government's budget constraint and the capital supply function, and future nominal rates R^b_{t+s} from the Taylor rule .

In detail, this means that when households know $K_t, B_t, R^b_t, s_t, H_t$, they can back out markups from the Phillips curve (15) using the stipulated rules for inflation in t and the conditional inflation forecasts for $t+1$. Given this, they can calculate real wages and total output. In turn, they know future government debt, B_{t+1}, from the government's budget constraint (18). The future nominal interest rate, R^b_{t+1}, is pinned down by the Taylor rule (17). Finally, from the nowcast for capital prices (21) households can determine the next period's capital stock K_{t+1}. Using these model-implied forecasts for $K_{t+1}, B_{t+1}, R^b_{t+1}, H_{t+1}$, households can then forecast next period's inflation, capital prices, etc. conditional on shock realizations ad infinitum. The law of motion for average productivity is given analytically by

$$\begin{aligned}\log H_{t+1} &:= \log \int h_{it+1} = \frac{1}{2} var(\log h_{it+1}) \\ &= \rho_h^2 \frac{1}{2} var(\log h_{it}) + \frac{1}{2}\bar{\sigma}_h^2 \exp(s_t) = \rho_h^2 \log H_t + \frac{1}{2}\bar{\sigma}_h^2 \exp(s_t).\end{aligned}$$

Below are the functional forms we use in the nowcasts/ forecasts of prices. We let the coefficients depend on the uncertainty state (hat denotes deviations from steady state):

$$\begin{aligned}\log \pi_t &= \beta^1_\pi(s_t) + \beta^2_\pi(s_t)\log \hat{B}_t + \beta^3_\pi(s_t)\log \hat{K}_t \\ &+ \beta^4_\pi(s_t)\hat{R^b_t} + \beta^5_\pi H_t^{-1}, \quad (56)\\ \log q_t &= \beta^1_q(s_t) + \beta^2_q(s_t)\log \hat{B}_t + \beta^3_q(s_t)\log \hat{K}_t \\ &+ \beta^4_q(s_t)\hat{R^b_t} + \beta^5_q H_t^{-1}. \quad (57)\\ \left[\log(1+\pi_{t+1})\frac{Y_{t+1}}{Y_t}\right] &= \beta^1_{E\pi}(s_t) + \beta^2_{E\pi}(s_t)\log \hat{B}_t + \beta^3_{E\pi}(s_t)\log \hat{K}_t \\ &+ \beta^4_{E\pi}(s_t)\hat{R^b_t} + \beta^5_{E\pi} H_t^{-1} + \beta^6_{E\pi}(s_{t+1}). \quad (58)\end{aligned}$$

Whether these rules yield good nowcasts of prices depends on the asset-demand functions, $b^*_{a,n}$ and k^*. If these are sufficiently close to linear in human capital, h, and in non-human wealth, b and k, at the mass of Θ_t, B_t and K_t will suffice and we can expect approximate aggregation to hold. For our exercise, the four endogenous aggregate states – R_t^b, H_t, B_t, and K_t – and the aggregate stochastic state s_t are sufficient to describe the evolution of the aggregate economy.

Technically, finding the equilibrium is the same as in Krusell and Smith (1997), as we need to find market clearing prices within each period. Concretely, this means that the posited rules, (56) to (58), are used to solve for households' policy functions. Having solved for the policy functions conditional on the forecasting rules, we then simulate n independent sequences of economies for $t = 1, \ldots, T$ periods, keeping track of the actual distribution Θ_t. In each simulation, the sequence of distributions starts from the stationary distribution implied by our model without aggregate risk. We then calculate in each period t the optimal policies for market clearing inflation rates and capital prices assuming that households resort to the policy functions derived under rule (56) to (58) from period $t + 1$ onward. Having determined the market clearing prices, we obtain the next period's distribution Θ_{t+1}. In doing so, we obtain n sequences of equilibria. The first 250 observations of each simulation are discarded to minimize the impact of the initial distribution. We next re-estimate the parameters of (56) to (58) from the simulated data and update the parameters accordingly. By using $n = 20$ and $T = 750$, it is possible to make use of parallel computing resources and obtain $10,000$ equilibrium observations. Subsequently, we recalculate policy functions and iterate until convergence in the forecasting rules.

The posited rules, (56) to (58), approximate the aggregate behavior of the economy fairly well. The minimal within sample R^2 is above 99%. The forecast performance is not perfect because we need to force households to effectively approximate the process for $\log \int h$ by a three state Markov chain. This variable moves slowly and leads to small but persistent low frequency errors.

C.3 Comparison of Results

Figure 12 compares the impulse response function obtained from the Reiter-method like solution to the non-linear Krusell-Smith solution. The Krusell-Smith impulse response functions are generated by linearly interpolating the policy functions, setting the uncertainty state to exactly its expected path after a one standard deviation shock, i.e., they are obtained without simulation.

The impulse responses look qualitatively similar across the two methods. One should take the results and hence the differences, however, with a grain of salt, as we need to

Table 7: Laws of motion for Krusell and Smith

Price of Capital $\log q_t$					(R^2: 99.26)
	s_1	s_2	s_3	s_4	s_5
β_q^1	-7.26	-7.34	-7.42	-7.54	-7.77
β_q^2	-0.02	-0.03	-0.03	-0.03	-0.03
β_q^3	-0.21	-0.22	-0.21	-0.21	-0.20
β_q^4	0.05	0.06	0.06	0.06	0.07
β_q^5	7.66	7.66	7.66	7.66	7.66
Inflation $\log \pi_t$					(R^2: 99.75)
	s_1	s_2	s_3	s_4	s_5
β_π^1	-4.10	-4.17	-4.25	-4.35	-4.52
β_π^2	-0.07	-0.07	-0.06	-0.06	-0.06
β_π^3	-0.06	-0.06	-0.06	-0.06	-0.05
β_π^4	0.03	0.03	0.03	0.03	0.03
β_π^5	4.38	4.38	4.38	4.38	4.38
Expectation term $\left[\log(1+\pi_{t+1})\frac{Y_{t+1}}{Y_t}\right]$					(R^2: 99.63)
	s_1	s_2	s_3	s_4	s_5
$\beta_{E\pi}^1$	-3.19	-3.24	-3.30	-3.36	-3.45
$\beta_{E\pi}^2$	-0.06	-0.06	-0.06	-0.06	-0.06
$\beta_{E\pi}^3$	-0.04	-0.04	-0.04	-0.04	-0.04
$\beta_{E\pi}^4$	0.02	0.02	0.02	0.02	0.02
$\beta_{E\pi}^5$	3.39	3.39	3.39	3.39	3.39

[1] For readability all values are mutiplied by 100.

approximate the continuous aggregate states in the Krusell and Smith algorithm very coarsely with 3 grid points each for K_t, B_t, R_t^b, H_t and 5 grid points for s_t. In addition, we need to decrease the points on the idiosyncratic assets grids to 40 each, as the total number of nodes is with $n_b \times n_k \times n_h \times n_s \times n_R \times n_B \times n_K \times n_H \approx 16E(+6)$ already very large. This leads to an underestimation of the persistence of the uncertainty shock and the slow moving average idiosyncratic productivity, which decreases the aggregate effects.

Figure 12: Comparison of Krusell-Smith vs. Reiter method

Aggregate States and Labor	**Output and Components**	**Prices and Returns**

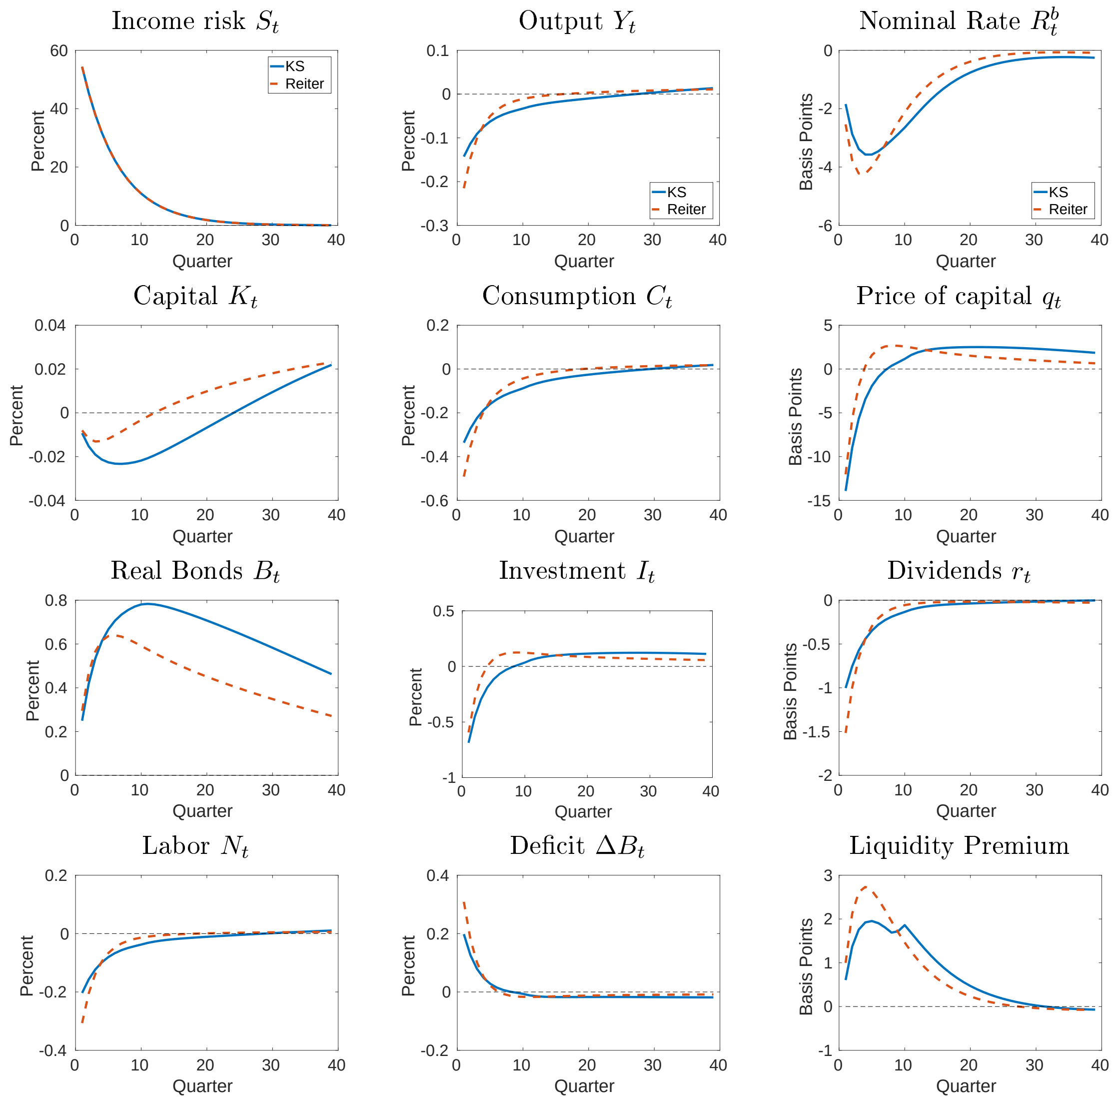

Notes: Liquidity Premium: $\frac{E_t q_{t+1} + r_t}{q_t} - \frac{R_t^b}{E_t \pi_{t+1}}$.
Impulse responses to a one standard deviation increase in the variance of income shocks. All rates (dividends, interest, liquidity premium) are *not* annualized.

D Estimation of the Stochastic Volatility Process for Household Income

D.1 Data

We estimate the income process based on the Survey of Income and Program Participation (SIPP) panels 1984, '85, '86, '87, 90', 91', '92, '93, '96, 2001, '04, and '08. We do not use the 1988 and 1989 surveys because of their known deficiencies due to the survey design and small sample size that resulted from budgetary constraints.

The SIPP panels provide monthly individual income data for up to three years (more in the 2008 survey) for each household member for each wave. The waves we use span the period 1983Q4 to 2013Q1. We constrain the sample to households with two married adults whose head is between 30 and 55 years of age and calculate for each household the labor income after taxes and transfers using NBER TAXSSIM. We aggregate income to quarterly frequency and restrict the sample to households that supply at least 260 hours of work (both spouses together) per quarter (50% of full-time).

We then estimate the predictable part of log household income, based on age and education dummies and a linear quadratic term in age for each education level. Furthermore, we control for time effects, ethnicity and the number of dependent children. The residuals from this regression form the basis of our subsequent analysis. We eliminate the top-bottom 0.5% of the residuals from each age-quarter cell to remove outliers.

Then, we construct a sequence of quarterly panels containing, for each household in the panel, current residual income and two lags thereof. We use these data to calculate for each quarter and age expressed in quarters the variance and the first two autocovariances of residual income. We estimate the sampling variance-covariance of the empirical variance-autocovariance estimates for each quarter and age cell by bootstrapping, where we stratify by age and quarter.

D.2 Income Process and Theoretical Autocovariances

Recall that after removing the deterministic income part $f(o_{i,a,t})$, residual income takes the form (see equation (22)):

$$\begin{aligned}
\log y^*_{it} &= \tau_{it} + h_{it} + \mu_i, \\
h_{it} &= \sum_{s=c}^{t} \rho_h^{t-s} \epsilon^h_{is}, \\
\tau_{it} &= \epsilon^\tau_{it} + \rho_\tau \epsilon^\tau_{it-1}, \\
\epsilon^\tau_{it} \sim \mathcal{N}(0, \sigma^2_\tau), \qquad & \epsilon^h_{i,t} \sim \mathcal{N}(0, \sigma^2_{h,t}), \qquad \mu_i \sim \mathcal{N}(0, \sigma^2_\mu),
\end{aligned} \tag{59}$$

where c is the quarter in which the household starts to accumulate shocks (i.e., the quarter in which the household turns 30), τ is an MA(1) transitory shock or measurement error, and μ is an individual fixed effect. Regarding the variances, we assume that the variance of the persistent component evolves slowly according to an AR(1) process, while we assume the variances of transitory shocks and fixed effects are time constant. The variance of income shocks evolves according to

$$\begin{aligned}
\sigma^2_{h,t} &= \bar{\sigma}^2_h \exp s_t \\
s_t &= \rho_s s_{t-1} + \epsilon^s_t \\
\epsilon^s_t &\sim \mathcal{N}\left(-\frac{\sigma^2_s}{2(1+\rho_s)}, \sigma^2_s\right),
\end{aligned} \tag{60}$$

in addition to a linear quadratic time trend (θ_1, θ_2).

This implies that the variance of residual income, $\log y^*_{i,t}$, for a household of cohort c is given by

$$\sigma^2_{c,t}(\log y^*) = (1+\rho_\tau)\sigma^2_\tau + \sigma^2_{\mu,c} + \sigma^2_{c,t}(h), \tag{61}$$

$$\sigma^2_{c,t}(h) = (\bar{\sigma}^2_p + \theta_1 t + \theta_2 t^2) \sum_{j=c}^{t} \rho_h^{2(t-j)} \exp s_j, \tag{62}$$

$$s_t = \rho_s s_{t-1} + \epsilon^s_t. \tag{63}$$

For the autocovariances, we obtain

$$ac_c(\log y^*_{it}, \log y^*_{i,t-1}) = \rho_\tau \sigma^2_\tau + \sigma^2_{\mu,c} + \rho_h \sigma^2_{c,t-1}(h), \tag{64}$$

$$ac_c(\log y^*_{it}, \log y^*_{i,t-2}) = \sigma^2_{\mu,c} + \rho_h^2 \sigma^2_{c,t-2}(h). \tag{65}$$

D.3 Estimation

Our estimation strategy uses these theoretical (autoco-)variances and their sample counterparts $(\omega^0_{c,t}, \omega^1_{c,t}, \omega^2_{c,t})$ to construct a quasi-maximum likelihood estimator, i.e., treating sampling error as normally distributed. Let ψ denote the sampling error, then we have

$$\psi_0(c,t) = \omega^0_{c,t} - \sigma^2_{c,t}(\log y^*) \tag{66}$$

$$\psi_1(c,t) = \omega^1_{c,t} - ac_c(\log y^*_{it}, \log y^*_{i,t-1}) \tag{67}$$

$$\psi_2(c,t) = \omega^2_{c,t} - ac_c(\log y^*_{it}, \log y^*_{i,t-2}). \tag{68}$$

We estimate the covariance matrix of ψ, $\Sigma_\psi(c,t)$, by bootstrapping age-quarter strata. With these terms at hand, we can specify the log pseudo-likelihood as

$$-2\log L = \sum_{(c,t)\in S} \psi'(c,t)\Sigma_\psi(c,t)^{-1}\psi(c,t) + \#T \log\sigma^2_s + \sum_{t\in T} \epsilon^2_t/\sigma^2_s, \tag{69}$$

where S is the set of all cohort-quarter pairs we observe, i.e., the cohorts 1959Q1 - 2013Q1 (denoted by the quarter they turn 30) between 1983Q4 and 2013Q1 and T is the set of quarters for which we estimate shocks, i.e., 1976Q1-2013Q1. We force $\sum_{t\in T}\epsilon_t = 0$.

We directly estimate the shock series, ϵ_t, together with the parameters for the persistent income shocks $(\rho_h, \rho_s, \bar{\sigma}_p, \sigma_s)$, the transitory and permanent part, $(\sigma_\tau, \sigma_\mu, \rho_\tau)$, and the time trend (θ_1, θ_2). However, since the data contain only limited information on shocks far before the sample starts, we set all shocks eight years before the first sample year (i.e., before 1976Q1) to their unconditional mean, i.e., to zero, and exclude them from the calculation of the likelihood. Eight years correspond roughly to the half-life of income shocks $(\log(1/2)/\log\rho_h)$ and hence twice the half-life of deviations in income variances.

D.4 Results

Table 8 summarizes the parameter values of interest. Persistent shocks to idiosyncratic income have a quarterly autocorrelation of $\rho_h = 0.98$ and an average variance of $\bar{\sigma}^2 = 0.059$. For shocks to the variance of persistent income shocks, we estimate a quarterly autocorrelation of $\rho_s = 0.84$ and a shock to the income variance has a standard deviation of 0.54, i.e., a one standard deviation shock increases the variance of income shocks by 54%.

In theory, we could track the history of the variance of persistent income shocks, $\sigma^2_{p,t}$, back to the year when the oldest cohorts at the start of the survey in 1984 started to

Table 8: Parameter Estimates

ρ_h	ρ_s	$\bar{\sigma}$	σ_s	
0.979	0.837	0.059	0.544	
(0.001)	(0.058)	(0.009)	(0.060)	
ρ_τ	σ_τ^2	σ_μ^2	θ_1	θ_2
0.339	0.013	0.073	3.396	-4.372
(0.006)	(0.000)	(0.001)	(0.889)	(1.352)

Notes: Asymptotic standard errors in paranthesis. Information matrix estimated by the Hessian of the log-likelihood function. Time trend parameters are estimated coding the time 1959Q1-2013Q1 as $-1, \ldots, 1$. The estimate for the average uncertainty $\bar{\sigma}$ includes the average time-trend effect for 1983-2013.

accumulate shocks. This way we could obtain a time series for income uncertainty going back to 1959. However, for the time long before the actual sample starts, the shocks are weakly identified because $\rho_h < 1$. Thus, we set all shocks before 1976 to zero and do not try to estimate these shocks. Figure 13 (b) plots the shock series and asymptotic one standard deviation confidence bounds. In fact, for any quarter before the first survey observation the confidence bounds for the shocks become fairly wide.

Figure 13: Household income risk 1979-2013

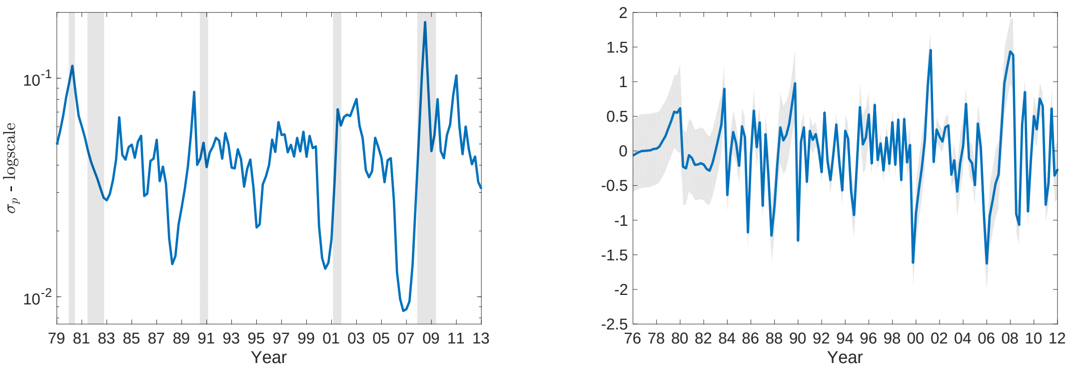

(a) Level (b) Shocks

Notes: Left panel: Risk level, log scale, NBER recession dates shaded in gray. Right panel: Shocks to income risk, one standard deviation confidence bounds shaded in gray.

E Wealth Distribution, Asset Classes, and Other Aggregate Variables

E.1 Data from the Flow of Funds

We can map our definition of liquid assets to the quarterly Flow of Funds (FoF), Table Z1. The financial accounts report the aggregate balance sheet of the U.S. household sector (including nonprofit organizations serving households) and are used in our analysis to quantify changes in the aggregate ratio of net liquid to net illiquid assets on a quarterly basis. Net liquid assets are defined as total currency and deposits, money market fund shares, various types of debt securities (Treasury, agency- and GSE-backed, municipal, corporate and foreign), loans (as assets), and total miscellaneous assets net of consumer credit, depository institution loans n.e.c., and other loans and advances.

Net illiquid wealth is composed of real estate at market value, life insurance reserves, pension entitlements, equipment and nonresidential intellectual property products of non-profit organizations, proprietors' equity in non-corporate business, corporate equities, mutual fund shares subtracting home mortgages as well as commercial mortgages.

E.2 Data from the Survey of Consumer Finances

We use eleven waves of the Survey of Consumer Finances (SCF, 1983-2013) to calibrate our model and to compare the cross-sectional implications of our model with the data.

Net liquid assets are classified as all households' savings and checking accounts, call and money market accounts (incl. money market funds), certificates of deposit, all types of bonds (such as savings bonds, U.S. government bonds, Treasury bills, mortgage-backed bonds, municipal bonds, corporate bonds, foreign and other tax-free bonds), and private loans net of credit card debt.

All other assets are considered to be illiquid. Most households hold their illiquid wealth in real estate and pension wealth from retirement accounts and life insurance policies. Furthermore, we identify business assets, other non-financial and managed assets and corporate equity in the form of directly held mutual funds and stocks as illiquid, because a large share of equities owned by private households is not publicly traded nor widely circulated (see Kaplan et al., 2016). From gross illiquid asset holdings, we subtract all debt except for credit card debt.

We exclude cars and car debt from the analysis altogether. What is more, we exclude from the analysis households that hold massive amounts of credit card debt such that their net liquid assets are below minus one month of average household income - the debt limit we use in our model. Moreover, we exclude all households whose equity in illiquid assets is below the negative of one average annual income. This excludes roughly 5% of U.S. households on average from our analysis and amounts to a debt limit on unsecured debt of 9,273 US$ in 2013, for example. Table 9 displays some key statistics of the distribution of liquid and illiquid assets in the population.

We estimate the asset holdings at each percentile of the net worth distribution by running a local linear regression that maps the percentile rank in net worth into the net liquid and net illiquid asset holdings. In detail, let LI_{it}, IL_{it} be the value of liquid and illiquid assets of household i in the SCF of year t, respectively. Let ω_{it} be its sample weight. Then we first sort the households by total wealth $(LI + IL)$ and calculate the percentile rank of a household i as $\mathbf{prc}_{it} = \sum_{j<i} \omega_{jt} / \sum_j \omega_{jt}$. We then run for each percentile, $prc = 0.01, 0.02, \ldots 1$, a local linear regression. For this regression, we calculate the weight of household i as $w_{it} = \sqrt{\phi(\frac{\mathbf{prc}_{it} - prc}{h})\omega_{it}}$, where ϕ is the probability density function of a standard normal, and $h = 0.05$ is the bandwidth. We then estimate the liquid and illiquid asset holdings at percentile prc at time t as the intercepts $\lambda^{LI,IL}(prc, t)$ obtained from the weighted regressions for year t:

$$w_{it} LI_{it} = \lambda^{LI}(prc, t) w_{it} + \beta^{LI}(prc, t)(\mathbf{prc}_{it} - prc) w_{it} + \zeta_{it}^{LI} \quad (70)$$

$$w_{it} IL_{it} = \lambda^{IL}(prc, t) w_{it} + \beta^{IL}(prc, t)(\mathbf{prc}_{it} - prc) w_{it} + \zeta_{it}^{IL}, \quad (71)$$

where ζ is an error term.

We can use these estimates for example to calculate average portfolio liquidity at time t as $\sum_{prc} \lambda^{LI}(prc, t) / \sum_{prc} \lambda^{IL}(prc, t)$. Figure 14 compares the percentage deviations of these average portfolio liquidity measures from their long-run mean and to those obtained from the FoF data for the years 1983 to 2013. The figure reveals that both data sources capture the very similar changes in the liquidity ratio over time.

However, it is important to note that the SCF, like many comparable surveys on wealth, systematically underestimates gross financial assets, and consequently, the average liquid to illiquid assets ratio in the FoF is roughly 20%, about twice as high as the one in the SCF. This is because households are more likely to underreport their financial wealth and especially deposits and bonds due to a larger number of potential asset items. In contrast, they tend to overestimate the value of their real estate and equity (compare also Table C.1. in Kaplan et al., 2016).

Table 9: Household portfolio composition:
Survey of Consumer Finances 1983-2013
Married households with head between 30 and 55 years of age

Moments	Model	Data
Fraction with $b < 0$	0.16	0.16
Fraction with $k > 0$	0.88	0.91
Fraction with $b \leq 0$ and $k > 0$	0.12	0.15
Gini liquid wealth	0.84	0.89
Gini illiquid wealth	0.79	0.79
Gini total wealth	0.78	0.78

Notes: Averages over the SCFs 1983-2013 using the respective cross-sectional sampling weights. Households whose liquid asset holdings fall below minus half quarterly average income are dropped from the sample. Ratios of liquid to illiquid wealth are estimated by first estimating local linear functions that map the percentile of the wealth distribution into average liquid and average illiquid asset holdings for each year, then averaging over years and finally calculating the ratios.

E.3 Other Aggregate Data

In Section 7, we depicted the impulse response functions of the log of real GDP, real personal consumption, real private investment, and real government consumption. These variables are taken from the national accounts data provided by the Federal Reserve Bank of St. Louis (Series: GDPMC1, PCECC96, GPDIC1, GCEC1) and data on government deficits from the NIPA tables for the U.S. (Table 3.1, BEA).

Data on house prices, Treasury bond returns and the liquidity premium stem from the same source. We use the secondary market rate of the 3-month Treasury bill (DTB3) as a measure for the short-term nominal interest rate. House prices are captured by the Case-Shiller S&P U.S. National Home Price Index (CSUSHPINSA) divided by the all-items CPI (CPIAUCSL). The liquidity premium we construct from nominal house prices, the CPI for rents, and the rate on 3-month Treasuries. We measure the liquidity premium as the excess realized return on housing. This is composed of the rent-price-ratio in t plus the quarterly growth rate of house prices in $t+1$, over the nominal return

Figure 14: Percentage deviation of portfolio liquidity from mean in SCF and FoF

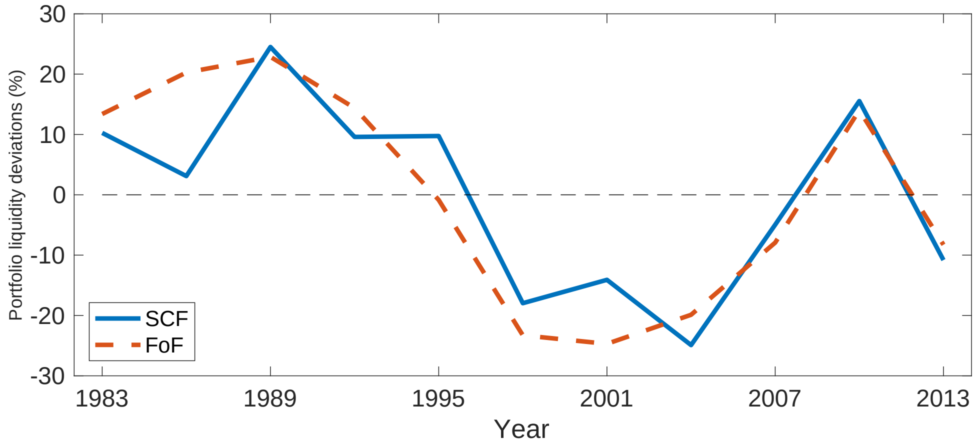

on riskless 3-month Treasury bills R_t^b (converted to a quarterly rate):

$$LP_t = \frac{R_{h,t}}{H_t} + \frac{H_{t+1}}{H_t} - (1 + R_t^b)^{\frac{1}{4}}. \tag{72}$$

Rents are imputed on the basis of the CPI for rents on primary residences paid by all urban consumers (CUSR0000SEHA) fixing the rent-price-ratio in 1981Q1 to 4%.

The Solow residual series we use is taken from the latest version (date of retrieval 2016-12-21) of Fernald's raw TFP series (Fernald et al., 2012). We construct an index from the reported growth rates and use the log of this index.

F Details on the Empirical Estimates of the Response to Shocks to Household Income Risk

F.1 Local Projection Method

In Figure 9 of section 7 we presented impulse response functions based on local projections (see Jordà, 2005). This method does not require the specification and estimation of a vector autoregressive model for the true data generating process. Instead, in the spirit of multi-step direct forecasting, the impulse responses of the endogenous variables X at time $t + j$ to uncertainty shocks, ϵ_t^s, at time t are estimated using horizon-specific single regressions, in which the endogenous variable shifted ahead is regressed on the current normalized uncertainty shock ϵ_t^s, a time trend, and controls $\mathbf{X}_{t-1}$. These controls are specified as the uncertainty level s_{t-1}, the return on T-bills R^b_{t-1} and the log of GDP Y_{t-1}, consumption C_{t-1}, of investment I_{t-1}, and of government expenditures G_{t-1}:

$$X_{t+j} = \beta_j + \beta_{j,0}\epsilon^s_t + \beta_{j,1}t + \gamma_j \mathbf{X}_{t-1} + \nu_{t+j},\ j = 0...8 \tag{73}$$

Hence, the impulse response function $\beta_{j,0}$ is just a sequence of projections of X_{t+j} in response to the shock ϵ^s_t, local to each forecast horizon $j = 0...8$. We focus on the post-Volcker disinflation era and use aggregate time series data from 1983Q1 to 2015Q4.
An important assumption made for employing the local projection method, which directly regresses the shocks on the endogenous variable of interest, is that the identified uncertainty shocks ϵ^s_t obtained from SIPP data are purely exogenous and orthogonal to all other structural shocks ν_{t+j} in the economy.

While this method allows for an identification that is fully consistent with our model, where all uncertainty fluctuations are exogenous, this identification strategy is arguably not very conservative. Therefore, we present additional evidence based on SVARs in the following.

F.2 Structural Vector Autoregressions

As a robustness check, we estimate structural VAR models, where the structural shocks to household income risk, ϵ^s_t, are identified using zero restrictions on the contemporaneous response of other variables to uncertainty shocks while income uncertainty itself is allowed to react to contemporaneous changes in all other variables.

To obtain the impulse response functions for income uncertainty and output, we estimate a structural VAR with two variables using Cholesky decomposition. For constructing the impulse response functions of all other variables of interest, we include one extra variable at a time in the SVAR, and assume that it responds within the period to changes in output. We include four lags of all variables and detrend the quarterly data using the HP-filter with weight 1600. The data source and estimation period are identical to those for local projections.

Figure 15 reports the results from this exercise. Overall, the SVAR-identified responses to uncertainty shocks look fairly similar to the ones we obtained using local projections. Compared to local projections, the confidence bounds appear tighter, and the responses of the nominal interest rate as well as the house price are slightly weaker.

Figure 15: Empirical response to household income risk shock

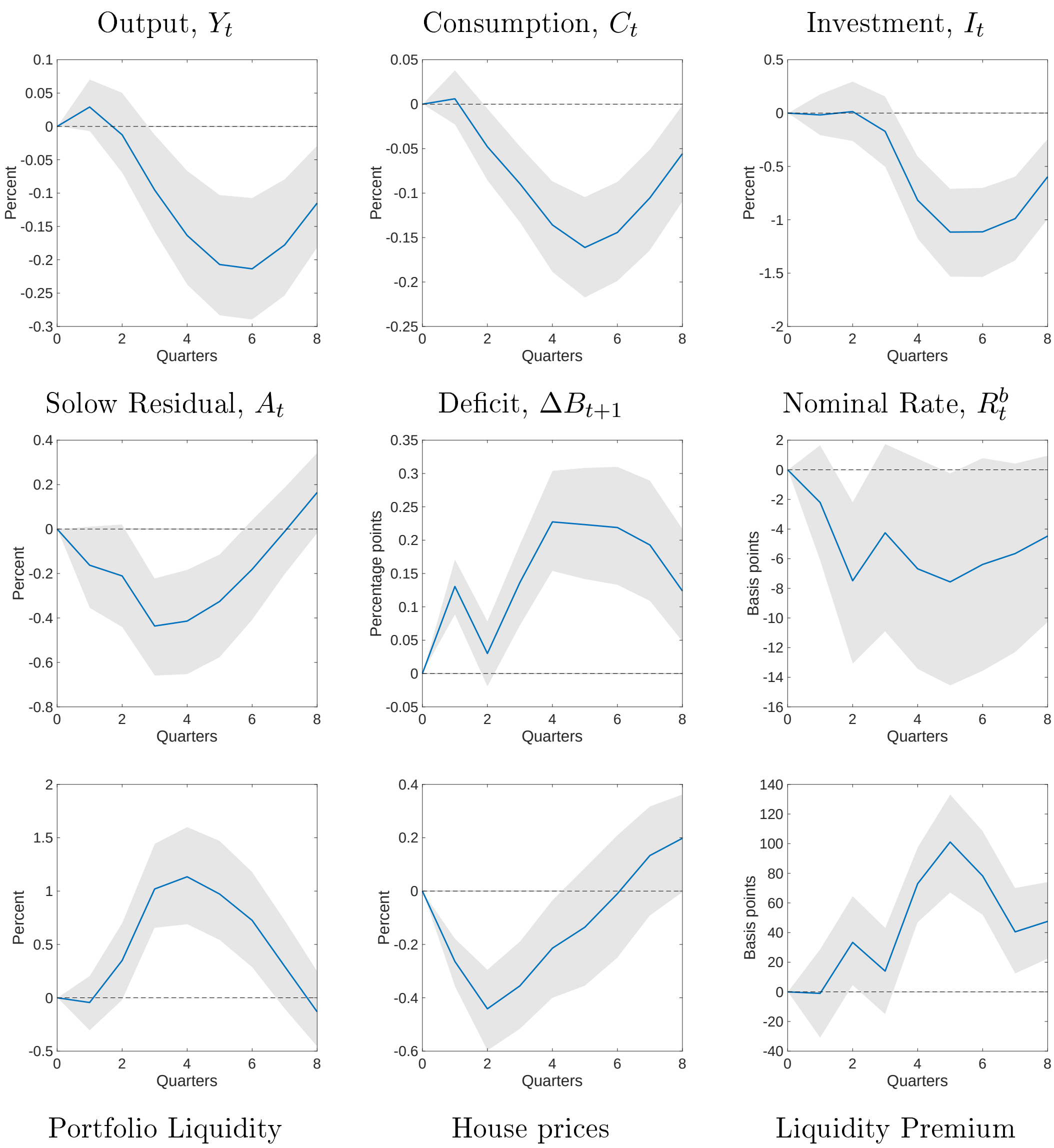

Estimated response of $X_{t+j}, j = 0 \ldots 8$, to shocks to household income risk. The impulse response functions are generated by estimating a three-variable SVAR with four lags that includes output, income risk and the displayed third variable of interest. Uncertainty shocks are identified by ordering income risk last in a Cholesky decomposition of the variance-covariance matrix of the reduced form VAR. Bootstrapped 66% confidence bounds shaded in gray (based on a parametric bootstrap).

G Unconditional Business Cycle Statistics

Table 10 reports unconditional business cycle statistics for the quarterly U.S. data we use in the empirical sections and for our model.

Table 10: Business cycle statistics data/model

	GDP	C	I	Deficit
	Time series standard deviation of … (in percent)			
Data	1.38	0.98	6.28	1.33
Model TFP	1.37	0.94	6.19	1.74
Model Uncertainty	0.23	0.54	0.67	0.30
	Correlation with GDP			
Data	1.00	0.92	0.92	-0.76
Model TFP	1.00	0.79	0.98	-0.84
Model Uncertainty	1.00	1.00	0.92	-0.07

Notes: Real GDP, Consumption (C), Investment (I) in logs. Net government savings (deficit) as a fraction of GDP. All data are HP-filtered with $\lambda = 1600$. Model refers to the baseline model with TFP or income risk shocks only.

H Response of the Model without Stabilization

Figure 16 shows the impulse responses for our baseline calibration with an interest rate peg and no fiscal stabilization.

Figure 16: Aggregate response to household income risk shock without stabilization

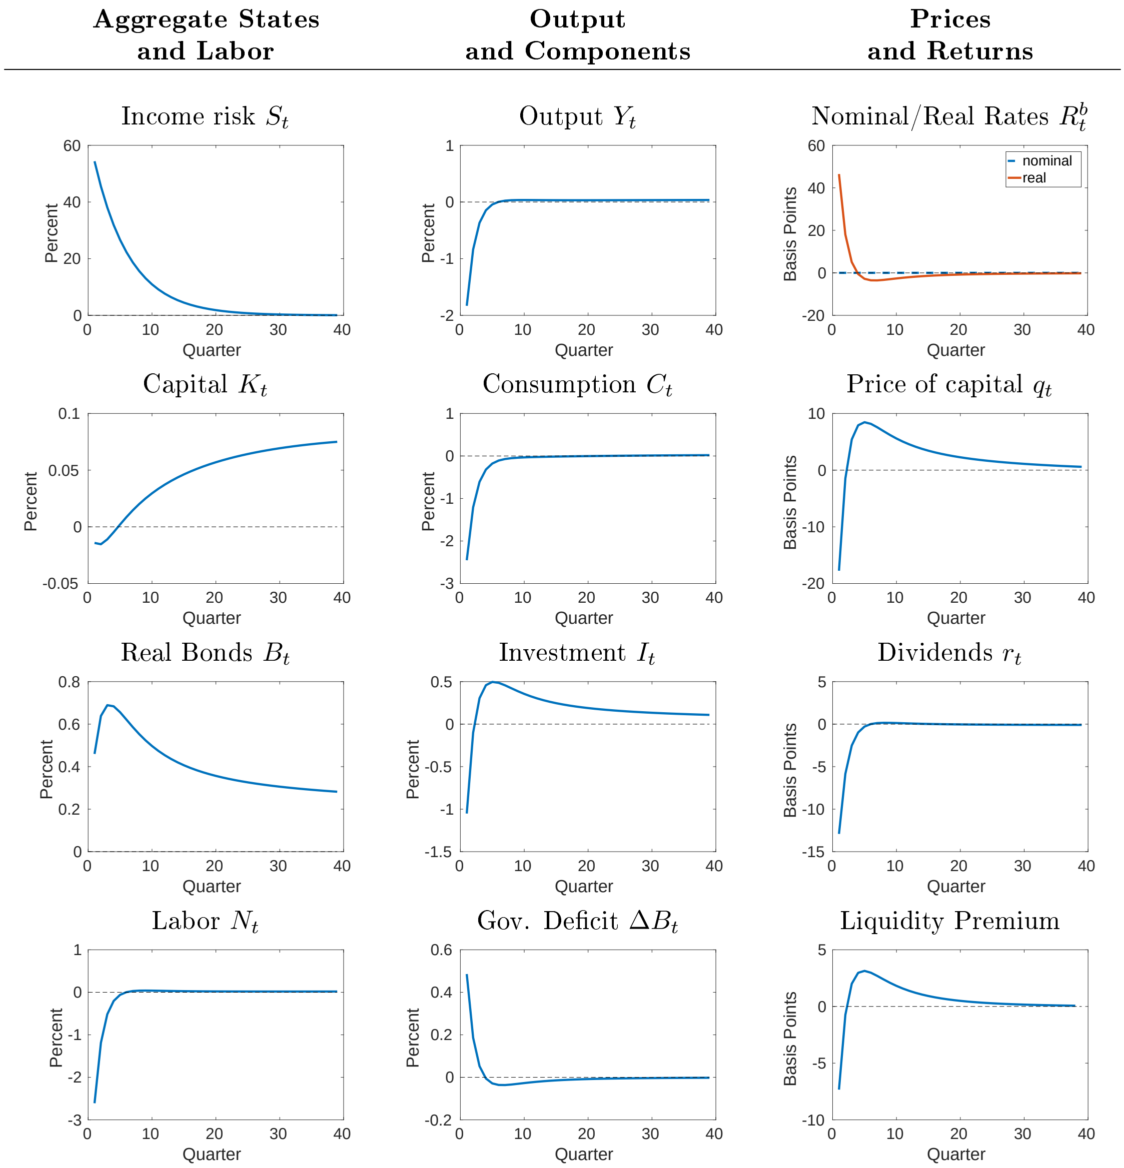

Notes: Liquidity Premium: $\frac{Eq_{t+1}+r_t}{q_t} - \frac{R_t^b}{E_t \pi_{t+1}}$.
Impulse responses to a one standard deviation increase in the variance of income shocks. All rates (dividends, interest, liquidity premium) are *not* annualized. All stabilization policy parameters are set to zero; $\theta_\pi = 0$, $\gamma_\pi = 0$, $\gamma_T = 0$, and $\rho_B = 1$.

I Response of the Model with Asset-Backed Securities

Figure 17 shows the impulse responses for our baseline calibration with ABS.

Figure 17: Aggregate response to household income risk with asset-backed securities

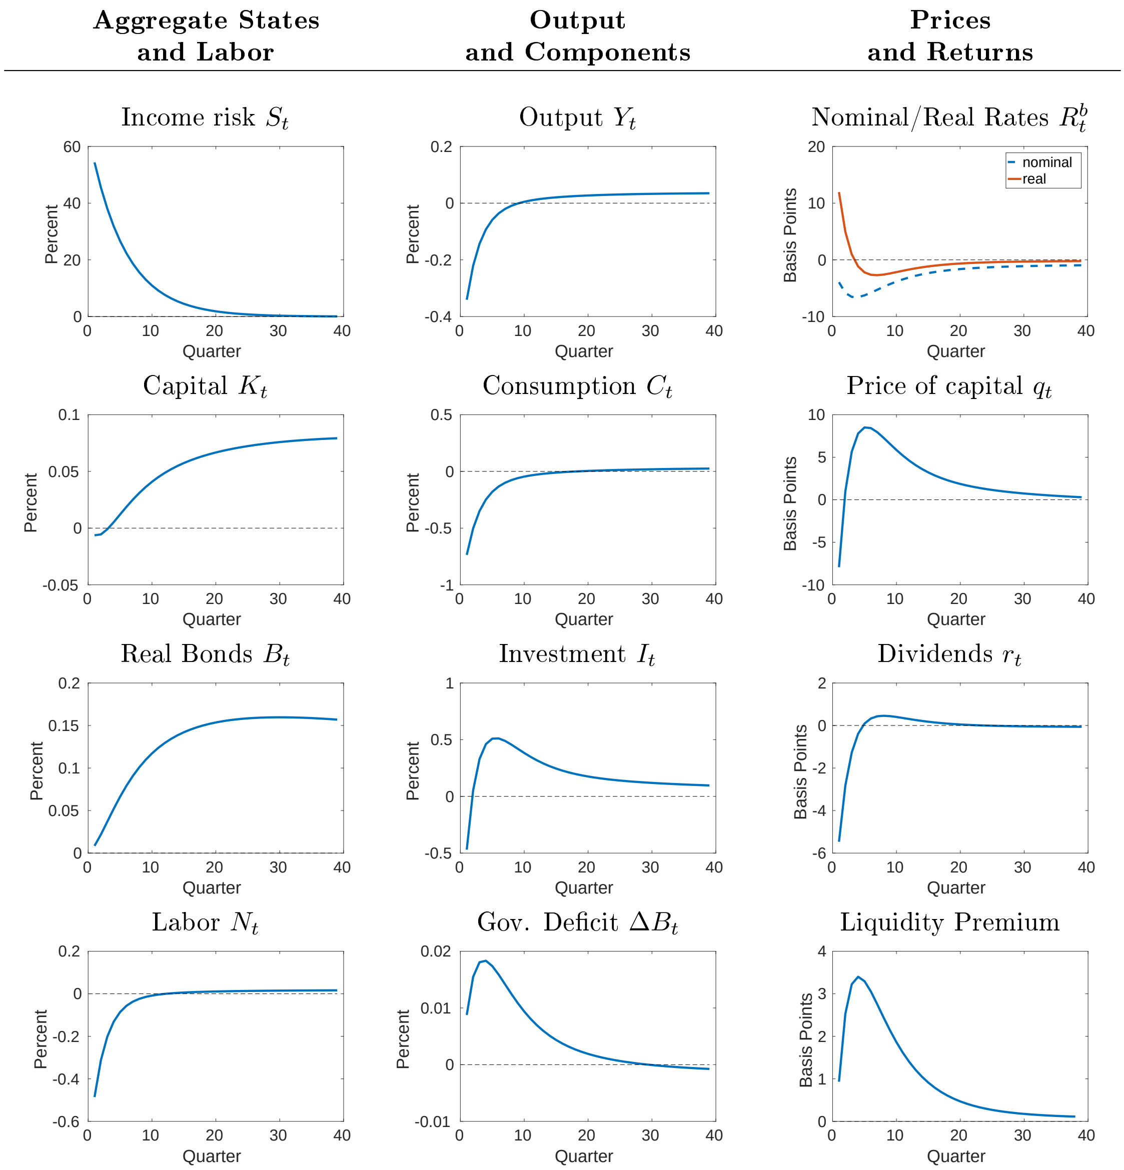

Notes: Liquidity Premium: $\frac{Eq_{t+1}+r_t}{q_t} - \frac{R^b_t}{E_t\pi_{t+1}}$.
Impulse responses to a one standard deviation increase in the variance of income shocks. All rates (dividends, interest, liquidity premium) are *not* annualized.

J Response of the Model to TFP Shocks

We generate the IRFs by solving the model with constant idiosyncratic income risk but time-varying total factor productivity in production, such that $Y_t = A_t F(K_t, L_t)$, where A_t is total factor productivity and follows an AR(1) process in logs. We assume a persistence of 0.95 and a standard deviation of 0.0075.

Figure 18: Aggregate response to a TFP shock

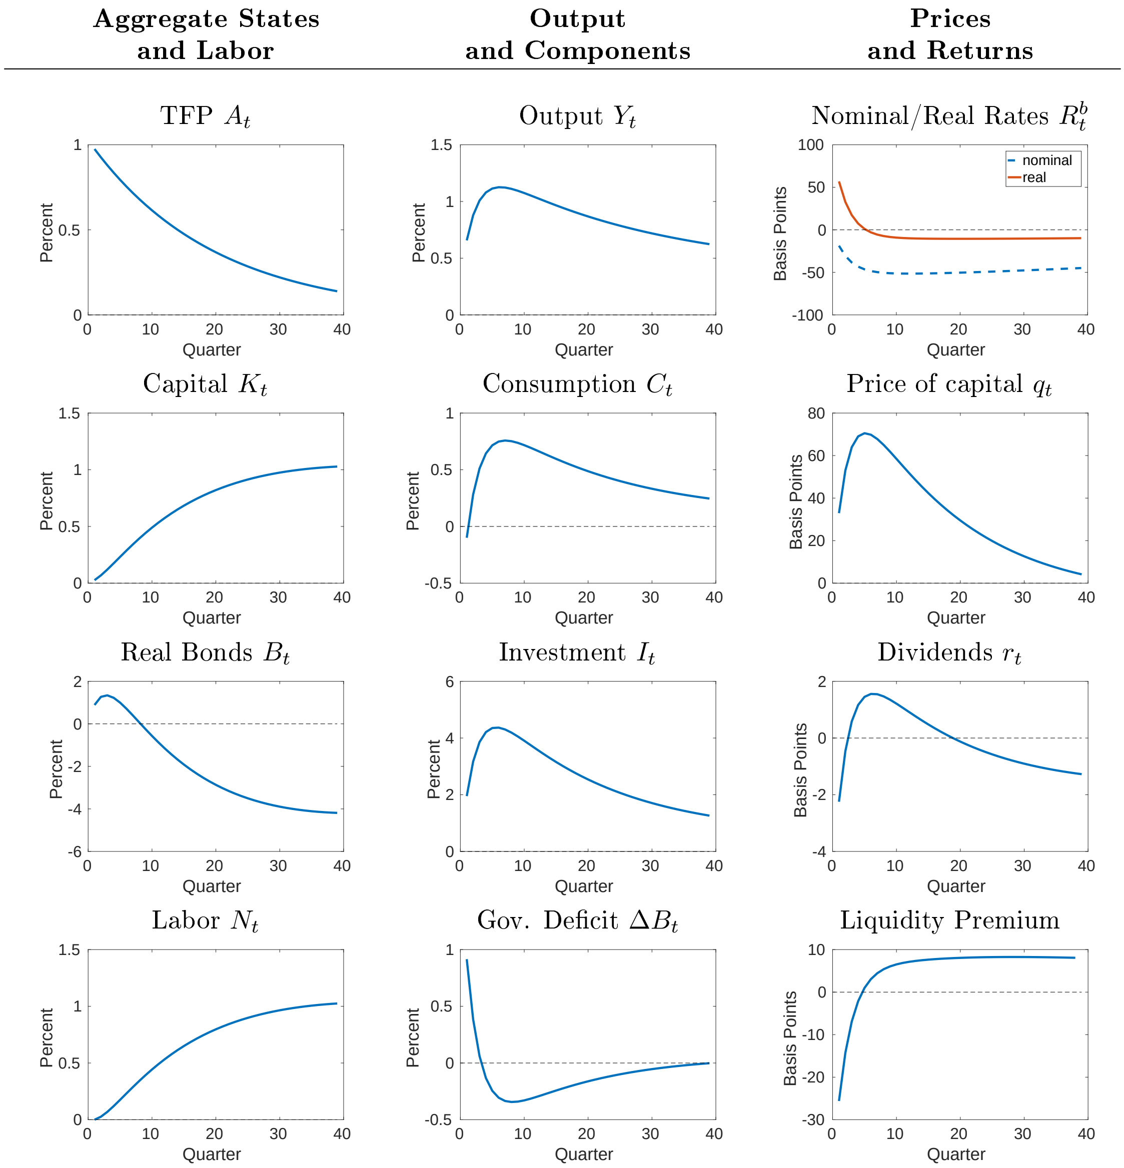

Notes: Liquidity Premium: $\frac{E_t q_{t+1} + r_t}{q_t} - \frac{R_t^b}{E_t \pi_{t+1}}$.
Impulse responses to a one standard deviation increase in TFP. All rates (dividends, interest, liquidity premium) are *not* annualized.

K Individual Consumption Responses to Persistent and Transitory Income Shocks

In order for the model to provide a useful framework for welfare analysis, it is important that the model replicates the empirical evidence on consumption responses to persistent and transitory income shocks (in partial equilibrium). For this purpose, we consider the average consumption elasticity to a persistent increase in income and an increase in liquid assets proportional to income (transitory income shock). These two elasticities are key to understanding the consumption smoothing behavior of an incomplete markets model; see Kaplan and Violante (2010) and Blundell et al. (2008). Table 11 provides these statistics for our model.

The model replicates the fact that transitory income shocks are well insured, while persistent income shocks are much less insured. Given the below unit-root autocorrelation of persistent income, our model predicts persistent income shocks to be somewhat better insured in comparison to assuming permanent shocks.

Table 11: Consumption smoothing in model and data

Elasticity of consumption to transitory and persistent income shocks		
	Data	Model
Transitory income change	0.05	0.05
Persistent income change	0.43	0.40

Data correspond to Kaplan and Violante (2010).

L Robustness Checks

For the risk aversion parameter, the Frisch elasticity of labor supply, and the frequency of price adjustment, we take standard values from the literature as there is no direct counterpart in the data. To account for this calibration strategy, we check the robustness of our findings with respect to the assumed parameter values. We do so by varying one of the parameters at a time while recalibrating to match the moments of Table 2 by adjusting the discount factor, the mean and variance of the distribution of adjustment costs, the fraction of entrepreneurs, and the borrowing penalty.

We find our results are qualitatively robust to all the considered parameter variations.

The impulse response functions for output, consumption, investment, and the liquidity premium are displayed in Figure 19. When we reduce the risk aversion households do not decrease investment demand as much as in the baseline and conversely the liquidity premium increases less. The illiquidity of capital is less important to households. An increase in the inverse Frisch elasticity is very similar to an increase in risk aversion. As can be seen from the household budget constraint when labor supply is maximized out, the lower the inverse Frisch elasticity, the less do the resources the household has for composite consumption fluctuate with productivity h. The recalibration of the illiquidity of capital only partially offsets this, because the return movements through central bank policy become relatively more important when households are effectively less affected by changes in income risk (either because they are less risk averse or better insured through the labor market).

When changing the stickyness of prices, the effect on output changes, but not the portfolio response by households. When prices become almost perfectly flexible, the response of output halfs, while making prices more sticky doubles the output response. The equilibrium liquidity premium is unaffected by price stickyness.

As a second robustness check, we vary the utility costs of portfolio adjustment. First, we make the adjustment probability more reactive to the value gained from adjustment by lowering the variance of the logistic distribution from which households draw the adjustment cost. Second, we consider a case of almost fixed adjustment probabilities by increasing the variance of the logistic distribution drastically. Third, we lower the mean of the logistic distribution such that the average adjustment probability goes up to 20% (and the average portfolio liquidity falls). All three cases show results qualitatively similar to our baseline; see Figure 20.

Making adjustment more state dependent yields quantitatively very similar results, the investment response is only slightly muted. When adjustment probabilities are fully exogenous, the investment response is almost twice as large. When the illiquid asset is more liquid, the effect of a shock to income risk becomes stronger in the short run but also shorter lived. The economic intuition seems to be the following: When the illiquid asset is very liquid, the demand for liquid assets becomes smaller but also less elastic to the return differences between the assets. Therefore, the central bank's intervention that cuts rates can stabilize less. The supply of liquid assets itself becomes more important, but with a small stock of outside liquidity, the same relative growth in government bonds stabilizes aggregate demand less. As soon as the stock of liquid funds, however, has increased sufficiently, households start to invest into illiquid assets again.

Figure 19: Robustness A: Aggregate response to household income risk shock

Risk Aversion* **Frisch Elasticity*** **Price Stickyness**

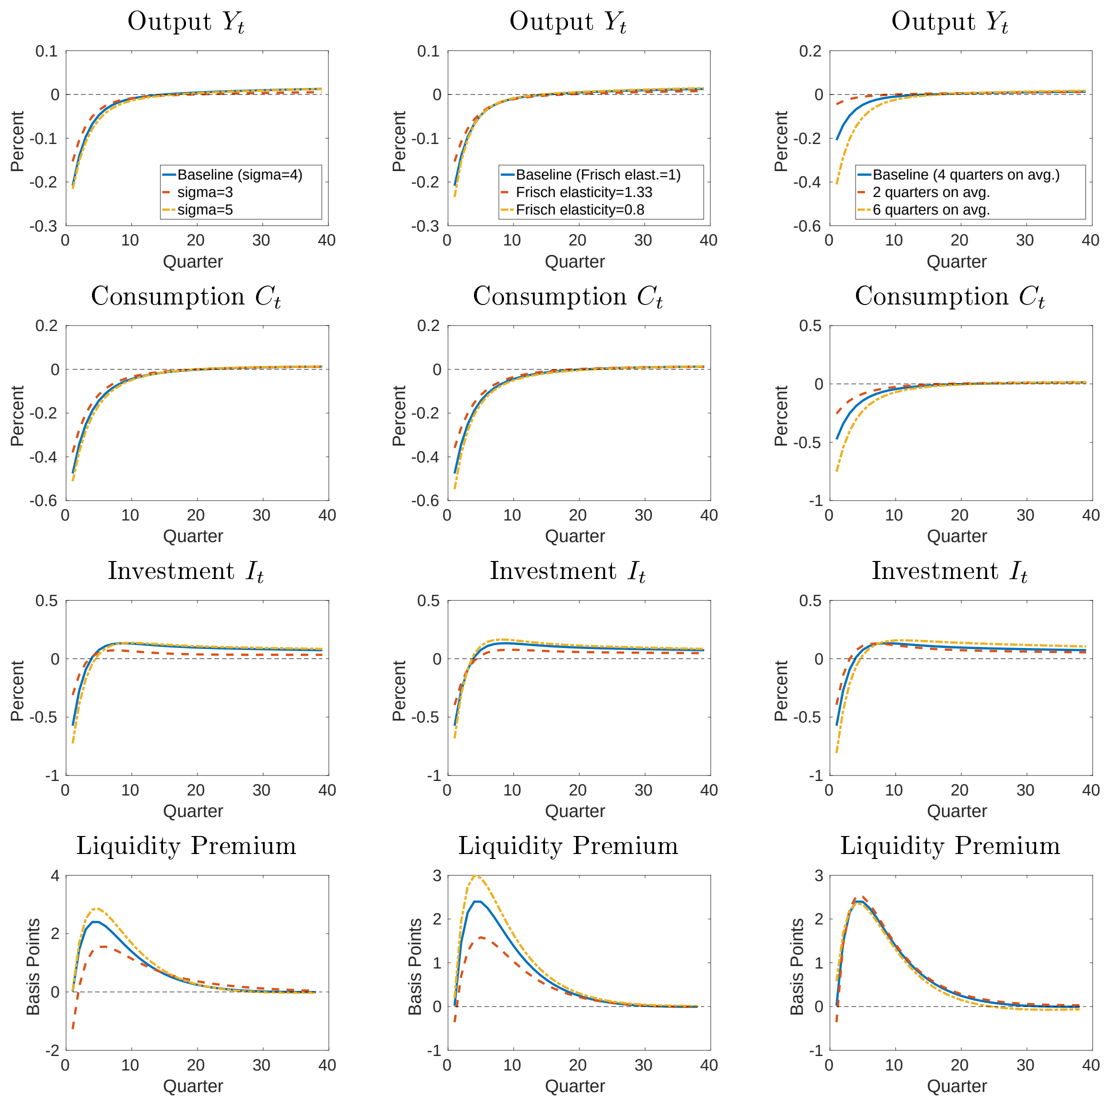

Notes: Liquidity Premium: $\frac{Eq_{t+1}+r_t}{q_t} - \frac{R^b_t}{\pi_{t+1}}$.

Impulse responses to a one standard deviation increase in the variance of income shocks. All rates (dividends, interest, liquidity premium) are *not* annualized.
* Recalibrated to match the moments of Table 2 by adjusting the discount factor, the mean and variance of the distribution of adjustment costs, the fraction of entrepreneurs, and the borrowing penalty.

Figure 20: Robustness B: Aggregate response to household income risk shock

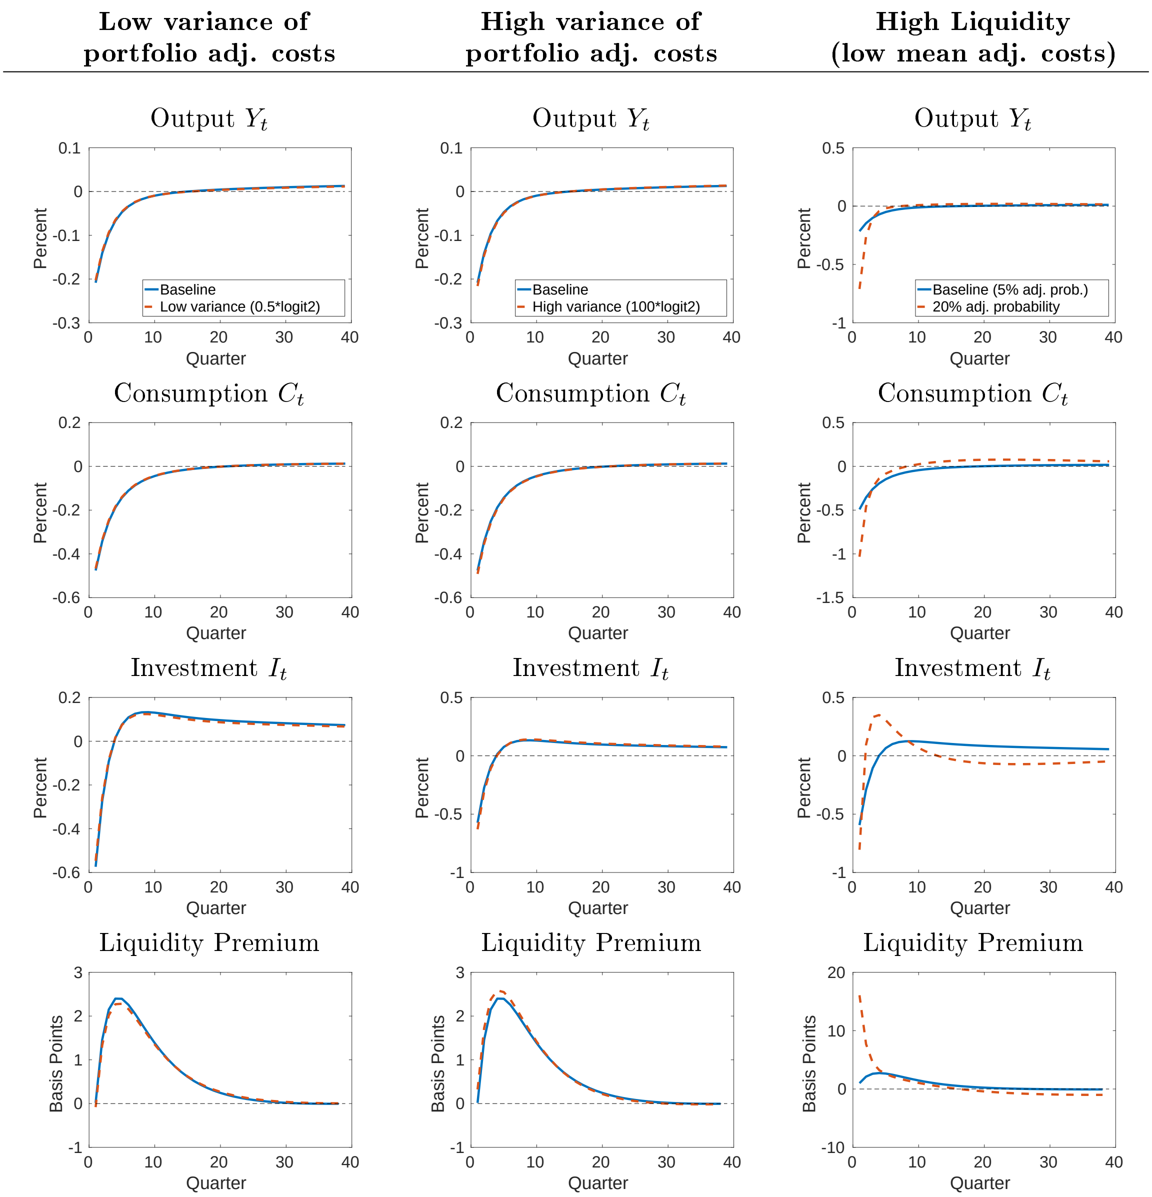

Notes: Liquidity Premium: $\frac{E_t q_{t+1} + r_t}{q_t} - \frac{R_t^b}{E_t \pi_{t+1}}$.

Impulse responses to a one standard deviation increase in the variance of income shocks. All rates (dividends, interest, liquidity premium) are *not* annualized.