

PRECAUTIONARY SAVINGS, ILLIQUID ASSETS, AND THE AGGREGATE CONSEQUENCES OF SHOCKS TO HOUSEHOLD INCOME RISK

CHRISTIAN BAYER

Dept. of Economics, Universität Bonn

RALPH LUETTICKE

Dept. of Economics, University College London

LIEN PHAM-DAO

Research Data and Service Centre, Deutsche Bundesbank

VOLKER TJADEN

Blue Yonder GmbH

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. 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.

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. Storesletten, Telmer, and Yaron (2001) estimate that for the United

Christian Bayer: christian.bayer@uni-bonn.de

Ralph Luetticke: r.luetticke@ucl.ac.uk

Lien Pham-Dao: lien.pham-dao@bundesbank.de

Volker Tjaden: volker.tjaden@mailbox.org

The paper was circulated before as “Household Income Risk, Nominal Frictions, and Incomplete Markets.” We would like to thank the four anonymous referees, Thomas Hintermaier, Andreas Kleiner, Keith Kuester, 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 Los Angeles, 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 merge 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 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.econometricsociety.org/ecta/supmat/13601.Data_and_Programs.zip.

States the variance of persistent shocks to disposable household income almost doubles in recessions.¹

The starting point of the present paper is that households use precautionary savings and structure their portfolios to smooth consumption if asset markets are incomplete and assets differ in their liquidity. Therefore, in such a setting, swings in the riskiness of household income lead not only to systematic variations in the propensity to consume, but also to a rebalancing of household portfolios.

We quantify the aggregate consequences of precautionary savings and portfolio adjustments 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. Illiquidity is modeled by a transaction cost. As a result, households trade capital only from time to time.² This two-asset structure allows us to disentangle savings and physical investment, and thus obtain strong fluctuations in aggregate demand in response to household income risk.³ To generate aggregate output effects from demand fluctuations, we augment this incomplete markets framework in the tradition of [Bewley \(1980\)](#), [Huggett \(1993\)](#), and [Aiyagari \(1994\)](#) by sticky prices à la [Rotemberg \(1982\)](#).

In this economy, an increase in income risk makes households consume less and save more. In addition, and importantly, they rebalance their portfolios toward the liquid asset because it provides better consumption smoothing. They take into account that they will have to adjust the illiquid asset more often to keep consumption smooth and this drives down the effective return of the illiquid asset because transactions are costly. Thus, higher income risk leads to a flight to liquidity.

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; 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 panel (a). Also in the aggregate, we see a substantial increase in portfolio liquidity around the crisis; see panel (b). We find this increase in portfolio liquidity to be a general response to estimated shocks to idiosyncratic income uncertainty in the data. 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.

More generally, when we estimate the time-series behavior of household income risk from the Survey of Income and Program Participation and then use these estimates to understand the consequences of a rise in income risk, we find the following results: In the data, a 1 standard deviation increase in household income risk decreases aggregate activity by 0.2% and investment by 1% over the first year after the shock, and our model

¹Work by [Güvenen, Ozkan, and Song \(2014b\)](#) documents that changes in individual labor income become left skewed in recessions.

²This setup is similar to [Kaplan and Violante \(2014\)](#) and [Kaplan, Moll, and Violante \(2017\)](#) following the tradition of [Baumol \(1952\)](#) and [Tobin \(1956\)](#) in modeling the portfolio choice between liquid and illiquid assets.

³In 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. Net liquid assets are composed of money market, checking, savings accounts, and call accounts, as well as corporate and government bonds and Treasury bills (T-bills) net of credit card debt. All other assets net of all other debt make up net illiquid wealth. (a) The change between 2007 and 2010 by quintile of the wealth distribution based on the Survey of Consumer Finances. The change in the average ratio of liquid to illiquid assets within the four last quintiles of net worth; see Section 6 for details. The first quintile does not consistently hold positive amounts of liquid and/or illiquid assets. (b) The ratio of liquid to illiquid assets on household balance sheets is based on quarterly data from the Flow of Funds; for details, see Appendix G.1.

matches these numbers. At the zero lower bound, when neither monetary nor fiscal policy stabilizes the economy, our model suggests an output loss of 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 when uncertainty increases. 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 as much from lower capital returns; 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 1 standard deviation increase in household income risk depresses welfare equivalent to 27 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.

In exploring portfolio adjustment as a new channel through which uncertainty affects real activity, our paper adds to the recent literature that explores the aggregate effects of time-varying uncertainty. In particular, Bloom's (2009) paper on the effects of time-varying (idiosyncratic) productivity uncertainty on firms' factor demand through the real option value of irreversible investment has triggered a stream of research on the aggregate effects of variations in firm-level productivity risk.⁴

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.⁵ Fernández-Villaverde, Guerrón-Quintana, Kuester, and Rubio-Ramírez (2015), for example, look at a medium-scale dynamic stochastic general equilibrium (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 2 standard deviation shock to the volatility of taxes. Off the zero lower bound, the drop reduces to 0.2%.⁶ In a similar framework, Basu and Bundick (2017) highlight how price stickiness can generate co-movement of consumption and investment after a decline in consumption demand driven by a shock to demand uncertainty. Overall, they find aggregate effects similar to those in Fernández-Villaverde et al. (2015).

We focus on idiosyncratic instead of aggregate income uncertainty and abstract from the labor supply effects of income risk by assuming Greenwood, Hercowitz, and Huffman (1988) preferences. This paper thereby isolates the precautionary savings and, in particular, the portfolio channel of income uncertainty. The focus on idiosyncratic 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 (2017) and Den Haan, Rendahl, and Riegler (2017). Both papers highlight the importance of uninsurable idiosyncratic unemployment risk in amplifying *first moment shocks* in labor search models. In contrast, we look at second moment shocks. More importantly, the two papers differ from ours in their asset market setup, assuming that all assets are perfectly liquid such that the portfolio reallocation we highlight is absent by definition.

With respect to the broader literature on New Keynesian incomplete markets models, we share with Gornemann, Kuester, and Nakajima (2012) a focus on the distributional consequences of systematic monetary and fiscal policy (here in response to uncertainty shocks). While they highlight labor market effects, we focus on portfolios. We share this focus with Kaplan, Moll, and Violante (2017), who discuss the transmission of monetary policy,⁷ and with Guerrieri and Lorenzoni (2017), who model the effect of a credit crunch.⁸

⁴To name a few, Arellano, Bai, and Kehoe (2012), Bachmann and Bayer (2013), Christiano, Motto, and Rostagno (2014), Chugh (2016), Di Tella (2018), Gilchrist, Sim, and Zakrajšek (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.

⁵With sticky prices, firms target a higher markup the more uncertain the future aggregate price level.

⁶Born 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.

⁷Luetticke (2017) builds on the framework of our paper to discuss the transmission of monetary policy.

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

The remainder of the paper is organized as follows. Section 2 estimates changes in household income risk and their effects on aggregates and household portfolios. Section 3 provides an intuition for these findings. Section 4 develops our quantitative model, and Section 5 discusses the solution method. Section 6 explains the calibration of the model. Section 7 presents the numerical results. Section 8 concludes. Appendices are supplied in the Supplemental Material (Bayer, Luetticke, Pham-Dao, and Tjaden (2019)), available in a supplementary file on the journal website <http://econometricsociety.org/ecta/supmat/13601.pdf>. They provide 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. EMPIRICAL EVIDENCE

To analyze the aggregate effects of shocks to household income risk, we first need to identify these shocks. For this purpose, we employ data from the Survey of Income and Program Participants (SIPP), covering the time period 1984–2013, and estimate a process for household income and its shock distribution similar to Storesletten, Telmer, and Yaron (2001, 2004), who estimate differences in income risk between recessions and booms. Different from their approach, we do not restrict ourselves to a generic business cycle relationship, but instead want to estimate a sequence of shocks to income risk first, and then to study their effect on household portfolios and a manifold of aggregate variables.

The central idea of Storesletten, Telmer, and Yaron's approach is to identify differences in the variance of persistent income shocks over time by comparing different cohorts of households at a given age. Like differences in the growth rings of a tree, the variance of income within a cohort memorizes the variances of shocks the cohort faced in the past to the extent that income is persistent. Since households of different cohorts accumulated income shocks at different times, differences across cohorts in terms of their within-cohort variance can identify the evolution of income-shock variances over time. Take for example a pure unit-root income process. In this example, the income variance of a cohort at a given point in time is simply the sum of all the variances of income shocks that cohort went through. Averaging the increase in income variances over all working-age cohorts between two adjacent quarters gives an estimate of the average income risk in that period. Below unit-root persistence in income, transitory income shocks, sampling uncertainty, and the persistence of income risk complicate the estimation, but the procedure we lay out below follows this basic intuition.

2.1. *Estimating Income Risk Over Time*

2.1.1. *Income Process*

Since the focus of this paper is on private self-insurance, our income measure is household labor income after taxes and transfers. The SIPP data are originally available at monthly frequency and represent individual-level income data. We aggregate these data to the household level and to quarterly frequency, restricting the data to households whose head is at least 30 and below 56 years of age. We generate household labor income by summing over household head and spouse, and impute taxes and transfers using TAXSIM.

We assume that the labor income of a household after taxes and transfers is composed of a transitory, a persistent, a household-fixed, and a deterministic component, that is,

income y of household i in quarter t is given by

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

where c defines a cohort by the quarter when a household head turns 30, $f(o_{it})$ measures the effect of observable household characteristics o_{it} , τ_{it} is a moving average (MA(1)) transitory shock or measurement error, μ_i is a household fixed effect, and h_{it} is a persistent component.

2.1.2. Risk Process

We assume transitory shocks and fixed effects to be homoscedastic, while the variance $\sigma_{\varepsilon,t}^2$ of the shocks ε_{it}^h to the persistent component, h_{it} , evolves slowly according to a log-autoregression (AR(1)) process around a quadratic time trend:

$$\begin{aligned} \sigma_{\varepsilon,t}^2 &= \bar{\sigma}_\varepsilon^2 \exp(s_t + t\theta_1 + t^2\theta_2), \\ s_{t+1} &= \rho_s s_t + \varepsilon_t^s, \\ \varepsilon_t^s &\sim \mathcal{N}\left(-\frac{\sigma_s^2}{2(1+\rho_s)}, \sigma_s^2\right). \end{aligned} \tag{2}$$

2.1.3. Autocovariance Structure of Residual Income

In a first step, we estimate $f(o_{it})$ by ordinary least squares (OLS) to remove this non-stochastic part and work with residual income, $\tau_{it} + h_{it} + \mu_i$. We then generate short panels of residual income and its first two lags for each household in a sample quarter $t = 1, \dots, T$. From these data sets, we calculate $ac_{0,j}(c, t)$, $j = 0, 1, 2$, that is, the sample variance ($j = 0$) and first two autocovariances ($j = 1, 2$) of residual income for cells defined by survey quarter, t , and c , the quarter when a household head turned 30.

These empirical autocovariances equal their theoretical counterparts $\omega_{0,j}(c, t)$, $j = 0, 1, 2$, up to sampling error. Substituting in the variances for the various terms in equation (2) yields, for the theoretical autocovariances,

$$\begin{aligned} \omega_{0,0}^2(c, t) &= (1 + \rho_\tau) \sigma_\tau^2 + \sigma_{\mu,c}^2 + \sigma_h^2(c, t), \\ \omega_{0,1}^2(c, t) &= \rho_\tau \sigma_\tau^2 + \sigma_{\mu,c}^2 + \rho_h \sigma_h^2(c, t - 1), \\ \omega_{0,2}^2(c, t) &= \sigma_{\mu,c}^2 + \rho_h^2 \sigma_h^2(c, t - 2), \end{aligned} \tag{3}$$

where $\sigma_h^2(c, t)$ is the variance of the persistent income component of cohort c at quarter t . This itself evolves slowly and accumulates income-risk shocks according to

$$\begin{aligned} \sigma_h^2(c, t) &= \sum_{j=c}^t \rho_h^{2(t-j)} \sigma_{\varepsilon,j}^2 = \bar{\sigma}_\varepsilon^2 \sum_{j=c}^t \rho_h^{2(t-j)} \exp(s_j + j\theta_1 + j^2\theta_2), \\ s_t &= \rho_s s_{t-1} + \varepsilon_t^s. \end{aligned} \tag{4}$$

2.1.4. Estimator

Equations (3) and (4) allow us to formulate a quasi-maximum-likelihood estimator for the parameters of interest ($\rho_h, \rho_s, \rho_\tau, \bar{\sigma}_\varepsilon^2, \sigma_\tau^2, \sigma_\mu^2, \sigma_s$) along with the *sequence of shocks to income risk* ε_i^s . Approximating the sampling noise by Gaussian error terms, the (quasi-)log-likelihood function of our model is given by

$$-2 \log L = \sum_{(c,t) \in S} \psi(c, t)' \Sigma(c, \tau)^{-1} \psi(c, t) + \sum_{j \in T} (\varepsilon_j^s)^2 / \sigma_s^2 + \#T \log \sigma_s^2, \quad (5)$$

where

$$\psi(c, t) = \begin{pmatrix} ac_{0,0}(c, t) - \omega_{0,0}(c, t) \\ ac_{0,1}(c, t) - \omega_{0,1}(c, t) \\ ac_{0,2}(c, t) - \omega_{0,2}(c, t) \end{pmatrix} \quad (6)$$

is the difference between theoretical and empirical autocovariances and, as such, is a function of the model parameters and the sequence of income-risk shocks. The matrix $\Sigma(c, t)$ is the corresponding variance–covariance matrix of $\psi(c, t)$ resulting from sampling uncertainty. Note that Σ is cell specific because differences in the income risk between cells lead to differences in sampling uncertainty regarding these income risks. We estimate the matrices $\Sigma(c, t)$ by block-bootstrapping the microdata clustered at the (c, t) cell level, that is, preserving the cell and autocorrelation structure of the data. The set S captures all cohort–quarter pairs we observe, that is, 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, that is, 1976Q1–2013Q1. Note that since we estimate the (autoco-)variances within a cohort–age cell, we control for anything that is common across households in this cell such as average wages, average hours, or average taxes.⁹

Our methodology extends Storesletten, Telmer, and Yaron’s (2001, 2004) method of moments estimator for income risk, reformulating it as a quasi-maximum-likelihood estimator.¹⁰ More fundamentally, it follows the idea of pseudo-panels as pioneered by Deaton (1985), that is, we treat each short panel of residual income and its two lags as an independent data set and aggregate the data in terms of second moments to the cohort level. These aggregated data are the unit of observation on which we estimate the nonlinear model for the laws of motion for income and shock variances. Further details can be found in Appendix F.

2.1.5. Estimation Results

Table I presents the estimation results. Income is persistent with a quarterly autocorrelation of ($\rho_h = 0.98$), which is slightly below the corresponding annual autocorrelation that Storesletten, Telmer, and Yaron (2004) report. The baseline persistent income risk is with $\bar{\sigma}_\varepsilon = 0.06$ comparable to their numbers.

⁹We force $\sum_{j \in T} \varepsilon_j^s = 0$. For 1959Q1–1975Q4 we set the shocks ε_i^s to zero because, as the persistence of h is below 1, income has limited memory. In turn, risk shocks occurring long before the first observation of income, that is, long before 1983Q4, have very little impact on any empirical variance and are hence weakly identified, and the estimate of σ_s becomes biased.

¹⁰The parameters could also be estimated using such a method of moments approach. For example, one could use a two-step procedure. In a first step, one minimizes the first sum in (5) searching over the parameters of the income process. The second step then uses the residuals from this first step and fits the autoregressive process for income risk $\sigma_h^2(c, t)$.

TABLE I
PARAMETER ESTIMATES^a

ρ_h	ρ_s	ρ_τ	$\tilde{\sigma}_\varepsilon$	σ_τ	σ_μ	σ_s
0.98	0.84	0.34	0.06	0.12	0.27	0.54
(0.06)	(0.06)	(0.01)	(0.03)	(0.00)	(0.01)	(0.10)

^aAll parameters correspond to quarterly frequency of the data. Bootstrapped standard errors using a wild bootstrap are given in parentheses; see Appendix F.3 for details. A quadratic time trend in the variance of persistent income shocks is included (not reported). The estimate for the average uncertainty $\tilde{\sigma}_\varepsilon$ includes the average time–trend effect for 1983–2013.

The estimated variability of income risk is large. The standard deviation of s , $\sigma_s = 0.54$, implies that on average roughly every 10 years there is an income-risk shock that triples income risk, which is equivalent to a 2 standard deviation shock. The persistence of income risk ($\rho_s = 0.84$) is in the range of typical business cycle fluctuations. Figure 2(a) displays the estimated series of persistent income risk and panel (b) displays the estimated sequence of shocks, ε_t^s , to income risk together with their confidence bounds. As one can see, income risk is low late in a boom and typically increased at the onset of a recession. The Great Recession stands out in size of income risk.

2.2. Responses to Shocks to Income Risk

We use the estimated sequence of shocks to household income risk, $\{\varepsilon_t^s\}_{t=1976Q1,\dots,2013Q1}$, to estimate their aggregate repercussions and the effects they have on the portfolios households hold. We focus on the post-Volcker disinflation era and discard all estimated shocks before 1983Q1, as structural breaks in monetary policy may impact results. Another reason is that the shocks before the start of the SIPP sample are not well identified because ρ_h is below 1.

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, aggre-




FIGURE 2.—Estimated level of household income risk over time. Quarterly 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. Left panel: Estimated standard deviation of persistent income shocks for the period 1979–2013. National Bureau of Economic Research (NBER) recession dates are in gray. Right panel: Shocks to income risk. Bootstrapped 1 standard deviation confidence bounds are shaded in gray. The dotted vertical line shows the first quarter for which we have observations from the SIPP.

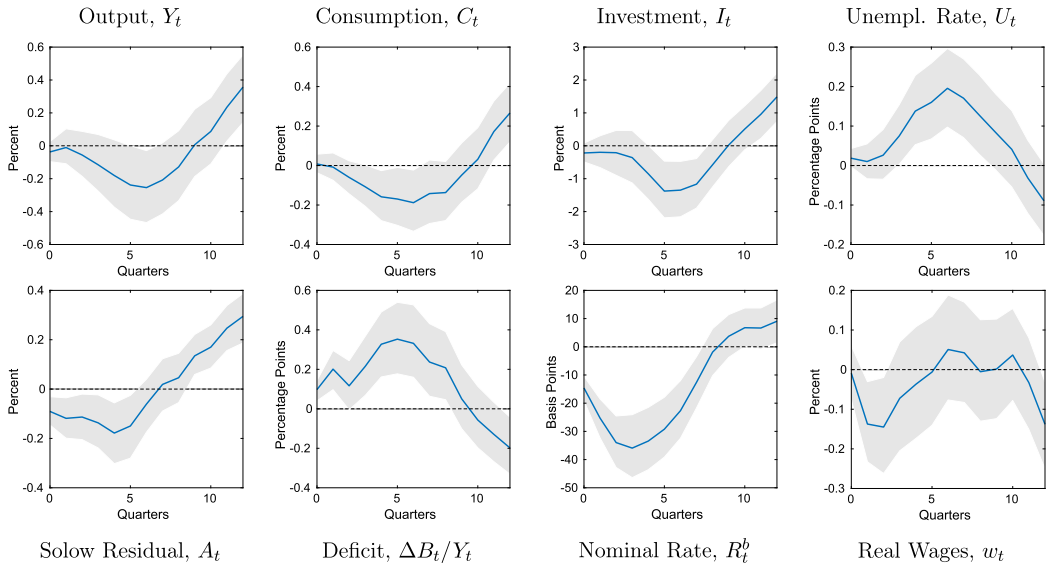


FIGURE 3.—Empirical response to household income-risk shock. Estimated response of \mathbf{X}_{t+j} , $j = 0, \dots, 12$, where $\mathbf{X}_t = [Y_t, C_t, I_t, A_t, \Delta B_t/Y_t, w_t, R_t^b, U_t]$, to the estimated shocks to household income risk, ε_t^s . The regressions control for the lagged state 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 are shown in gray (block bootstrap).

gate output falls. Investment declines particularly strongly, as households rebalance their portfolios toward liquid assets, while the (equilibrium) return premium on illiquid assets goes up. Combining the income-risk series with cross-sectional information on household portfolios from the Survey of Consumer Finances (SCF), we find that the increase in the liquidity of household portfolios is particularly concentrated among the relatively poor.

2.2.1. Aggregate Response

Figure 3 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 . As a robustness check, we have estimated impulse responses using a different ordering of variables, controlling also for the contemporaneous response of aggregates. Results are similar and can be found in Appendix H.2.

Upon a 1 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 total factor productivity (TFP) series (Fernald (2012)) falls as well. One explanation could be that upon a decrease in aggregate demand, markups go up as they do in New Keynesian models and this is captured as a decrease in measured TFP (see Hall (1989)). Real wages fall slightly and the unemployment rate goes up by 0.2 percentage points. The government seems to react systematically by making use of stabilizing monetary and fiscal policy—government deficits go up by 0.25 percentage points of gross domestic product (GDP) over the first year and the nominal

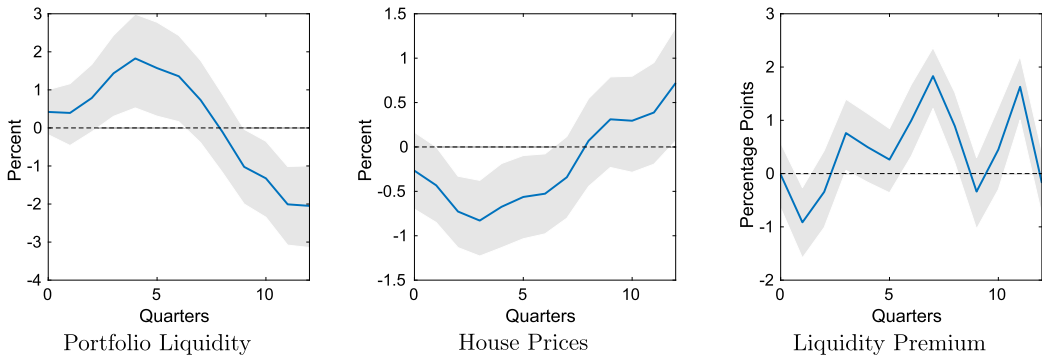


FIGURE 4.—Response of household portfolios, house prices, and the liquidity premium to household income-risk shock. Estimated response of the liquidity of household portfolios, the price of houses (Case–Shiller S&P Index), and the difference between the return on housing and the nominal rate (liquidity premium) to income risk using local projections. The set of control variables is as in Figure 3. Bootstrapped 66% confidence bounds are shown in gray (block bootstrap).

return on 3-month T-bills goes down on average by 35 basis points (annualized) over the first four quarters after the shock. After roughly 10 quarters, the recessionary effect of the income-risk shock becomes expansionary and output, consumption, and investment overshoot their trends.

The decline in investment—despite a decrease in interest rates—finds its repercussions in household balance sheets; see Figure 4. 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.¹¹

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 50% on average; see Kaplan and Violante (2014)). Hence any change in house prices directly affects portfolio liquidity. However, as house prices, measured by the Case–Shiller index, fall only by 1% after an increase in uncertainty (see Figure 4), 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, measured as the rent plus price increase of houses relative to the 3-month T-bill return, increases relatively quickly after the shock to household income risk; see again Figure 4.¹²

¹¹Kaplan, Moll, and Violante (2017) use a very similar taxonomy to split assets into liquid and illiquid. The reason for treating equities as illiquid is that most equities are held in the form of pension funds. Equity shares held directly play a role only above the 85th wealth percentile, but even these are often closely held equities 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.

¹²We 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 Standard and Poors (S&P) national house price index. Rents are imputed on the basis of the consumer price index (CPI) for rents of primary residences, fixing the rent–price ratio in 1981Q1 to 4%.

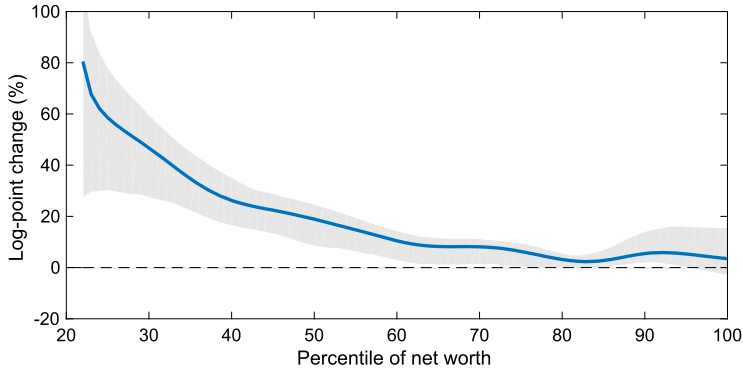


FIGURE 5.—Response of portfolio liquidity to household income risk by wealth percentile. Estimated log difference of the liquid to illiquid ratio of household portfolios across the wealth distribution in response to a 1 standard deviation shock to household income risk. Income-risk shocks are identified from the SIPP. Portfolio composition is estimated from the SCF years 1983–2013, only households with at least two adults, and households where the head is between 30 and 55 years of age. Bootstrapped 66% confidence bands are shown in gray.

2.2.2. Response by Wealth Group

As our theoretical explanation focuses on heterogeneity among households as a result of uninsurable risk, it will have rich cross-sectional implications for households' responses to income risk that go beyond average household portfolios and the differential changes in the return on liquid and illiquid assets.

To provide evidence along this dimension, we use the waves 1983–2013 of the Survey of Consumer Finances (SCF). The data sets contain detailed information on household balance sheets. In line with our treatment of the SIPP data, which we use to estimate income risk, we restrict the sample to households whose 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.

For each wave, t , of the SCF, we estimate a function that maps the percentile rank, prc , of a household in the total wealth distribution into liquid, $\lambda^{\text{LI}}(\text{prc}, t)$ and illiquid, $\lambda^{\text{LI}}(\text{prc}, t)$, asset holdings by a local-linear regression; see Appendix G for details. Using this function, we then calculate the average ratio of liquid to illiquid assets of a household at a given percentile in the wealth distribution $\lambda(\text{prc}, t) = \frac{\lambda^{\text{LI}}(\text{prc}, t)}{\lambda^{\text{LI}}(\text{prc}, t)}$.

We calculate the average shock in the year preceding an SCF wave, $\bar{\varepsilon}_{t-1}^s$, and regress the liquidity ratio of all percentiles (above the 20th) on the shock, an intercept, and a linear time trend, following the idea of a local projection here, too:

$$\lambda(\text{prc}, t) = \gamma_0(\text{prc}) + \gamma_1(\text{prc})t + \gamma_2(\text{prc})\bar{\varepsilon}_{t-1}^s + \zeta.$$

Figure 5 shows the coefficients, $\gamma_2(\text{prc})$, of the uncertainty shock for this regression. Again, we bootstrap the confidence bands. The poorer is a household, the stronger is its increase in liquidity holdings.

3. A SIMPLE EXPOSITORY PARTIAL EQUILIBRIUM MODEL

To provide intuition on why households want to increase the liquidity of their portfolios upon an increase in income uncertainty, we commence with a stylized three-period model

of income uncertainty and portfolio decisions. We present our quantitative general equilibrium model in Section 4. The three-period model is meant for illustrative purposes and focuses on the effect of uncertainty on asset demand without discussing aggregate effects.

Households hold an endowment $y > 0$ in period 1, which they can either consume, c_1 , or invest. When investing, they have to decide between illiquid assets, k_1 , and liquid assets, b_1 . The liquid asset pays a zero net return in period 2 (a storage technology). The illiquid asset pays a positive net return, but only in period 3. In period 2, half of the households obtain high income, $y^H = y + \sigma$, $\sigma > 0$, and half obtain low income, $y^L = y - \sigma \geq 0$. They can again invest in a liquid zero-net-return asset in positive amounts, but they can neither borrow against the illiquid asset nor sell it. There is no endowment in period 3.

With their consumption and savings decisions, households maximize the sum over period felicities, $u(c_t) = \frac{c_t^{1-\xi}}{1-\xi}$, that is, we abstract from discounting. Since in period 2 all uncertainty is resolved, consumption in period 3 will only depend on income in period 2. Let $c_t^{H,L}$ denote consumption in period t after income is high or low in period 2, respectively. Since all uncertainty is revealed by then, the household splits resources evenly between periods 2 and 3 if no borrowing constraint binds. When the constraint binds, the household does not save in period 2, and hence consumes all income and liquid assets in period 2, being left with only illiquid asset income in period 3. If gross returns on the illiquid asset, R^k , are not too high, the household will not be constrained in the high income state. Moreover, households will only hold liquid assets despite the higher return on the illiquid asset if they expect to be constrained in the low income case. Therefore, we focus on this case (see Appendix A for details), such that

$$c_2^H = \frac{1}{2}(b_1 + y + \sigma + R^k k_1), \quad c_3^H = c_2^H,$$

$$c_2^L = b_1 + y - \sigma, \quad c_3^L = R^k k_1.$$

Keeping investment and savings decisions in period 1 fixed, this means that an increase in income risk, σ , leaves consumption, c_3^L , unchanged because the borrowing constraint binds in period 2, but it increases consumption c_3^H because the household saves some of its income from period 2 in the high income state. Therefore, average consumption in period 3 rises with income risk σ . Conversely, consumption in period 2 falls more in the bad state in σ than it rises in the good state and average consumption in period 2 falls in σ . In other words, an increase in risk shifts consumption *ceteris paribus* from period 2, when capital is illiquid, to period 3, when all assets become liquid. Anticipating this shift, the household has an incentive to undo it by rebalancing its portfolio in period 1, increasing the share of liquid assets. Importantly, this comes on top of the precautionary motives that lead to an increased demand for liquid assets also in a setup without any illiquid investment option.

This implies that the demand for *liquid* assets increases more strongly in uncertainty when households can invest in an illiquid asset, than it does when they cannot.

PROPOSITION 1: Define $b_1^*(\sigma)$ and $k_1^*(\sigma)$ as the optimal liquid and illiquid asset holdings. Define $\tilde{b}_1(\sigma)$ as the liquid asset holdings of a household that does not have the option to invest in an illiquid asset. Now, suppose income uncertainty is large enough such that $b_1^*(\sigma) > 0$ and the returns on the illiquid investment are positive but not too large, that is, $1 < R^k < (1 + R^k \frac{\xi-1}{\xi})^\xi$. Then $\frac{\partial b_1^*}{\partial \sigma} > \frac{\partial \tilde{b}_1}{\partial \sigma} > 0 > \frac{\partial k_1^*}{\partial \sigma}$, that is, liquid asset holdings increase in σ and they increase more than in a model where all assets are liquid, while illiquid asset holdings decrease.

See Appendix A for the proof.

From a macroeconomic perspective, this shift in portfolios becomes important when changes in the demand for liquid and illiquid assets impact the demand for final goods differently. In our general equilibrium model in the next section, this is the case. A decrease in the demand for illiquid assets leads to an immediate decline in physical investment. An increase in the demand for liquid assets leads to an increase in goods demand only when the suppliers of these assets—households and the government—decide to use the extra funds. In turn, when output is demand determined, because prices are sticky, this decrease in demand due to higher income uncertainty leads to a decrease in output. Portfolio rebalancing exacerbates this demand-driven downturn, all of which is in line with our empirical findings in the previous section.

4. QUANTITATIVE MODEL

To understand the quantitative importance of portfolio rebalancing in response to changes in income uncertainty, we build a dynamic model of heterogeneous households with incomplete markets, time variations in income risks, and sticky prices. The economy is 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, issues government bonds, and adjusts expenditures in response to deviations of government debt from its long-run target as well as in response to business cycle conditions. The monetary authority sets the nominal interest rate on government bonds according to a Taylor rule.

4.1. Households

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, but only workers supply labor. The efficiency of a worker's labor evolves randomly, exposing worker households to labor-income risk. Entrepreneurs do not work, but earn all pure rents in our economy. All households self-insure against the income risks they face by saving in a liquid nominal asset (bonds) and a less liquid physical asset (capital). Trading illiquid capital is costly as in Kaplan and Violante (2014).

To be specific, there is a continuum of ex ante identical households of measure 1, 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} , from renting out capital, k_{it} , and from interest on bonds, b_{it} . Whenever a household adjusts its holdings of capital, it needs to pay some felicity cost, χ_{it} , that is an independent and identically distributed (i.i.d.) draw from a logistic distribution.¹³ Holdings of bonds have to be above an exogenous debt limit \underline{B} , and holdings of capital have to be nonnegative.

¹³Kaplan and Violante (2014) find that physical transaction costs and utility costs yield similar results for the portfolio problem. Assuming a logistic distribution of adjustment costs yields closed-form solutions for expected adjustment costs given the value of adjustment.

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} . In line with our empirical specification, we assume that productivity evolves according to a log-AR(1) process with time-varying volatility and a fixed probability of transition between the worker and the entrepreneur state,

$$\tilde{h}_{it} = \begin{cases} \exp(\rho_h \log \tilde{h}_{it-1} + \varepsilon_{it}^h) & \text{with probability } 1 - \zeta \text{ if } h_{it-1} \neq 0, \\ 1 & \text{with probability } \iota \text{ if } h_{it-1} = 0, \\ 0 & \text{else,} \end{cases} \tag{7}$$

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

$$\begin{aligned} \sigma_{h,t}^2 &= \bar{\sigma}_h^2 \exp s_t, \\ s_{t+1} &= \rho_s s_t + \varepsilon_t^s, \\ \varepsilon_t^s &\sim \mathcal{N}\left(-\frac{\sigma_s^2}{2(1 + \rho_s)}, \sigma_s^2\right), \end{aligned} \tag{8}$$

that is, at time t , households observe a change in the variance of shocks that drive the next period’s productivity. In words, we assume that idiosyncratic productivity normally evolves according to a log AR(1) process with time-varying variance.¹⁴ With probability ζ , households become entrepreneurs ($h = 0$). With probability ι , an entrepreneur returns to the labor force with median productivity. An entrepreneurial household obtains a fixed share of the pure rents, Π_t , in the economy (from monopolistic competition and creation of capital).¹⁵ We assume that the claim to the pure rent cannot be traded as an asset.

This modeling strategy, the introduction of an exogenous entrepreneur state, serves two purposes. First and foremost, it solves the problem of the allocation of pure rents without distorting factor returns and without introducing another tradable asset—an issue in any heterogeneous agent New Keynesian model.¹⁶ Second, we use the entrepreneur state—as a transitory state in which incomes are typically extremely high—to match the

¹⁴For simplicity, we abstract from transitory income shocks and permanent income differences in the model. We assume, partly for analytical clarity, that uncertainty fluctuations are exogenous. Likely some of the fluctuations in uncertainty in the data reflect endogenous responses through, say, unemployment as in [Ravn and Sterk \(2017\)](#) or, beyond the labor income focus of our paper, a change in the insurance offered by financial markets as in [Brunnermeier and Sannikov \(2016\)](#).

¹⁵Note that the notation with $h = 0$ for the entrepreneur state is somewhat counterintuitive: the entrepreneur state will be a high income state. 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 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.

¹⁶The 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. Allocating pure rents exogenously is not the only way to allocate them without distorting factor returns, but it is the only way to avoid the introduction of a third asset, which then would need to be priced. If one is willing to assume that pure and capital rents come in illiquid form, as in

wealth distribution following the idea by [Castaneda, Diaz-Gimenez, and Rios-Rull \(1998\)](#). The entrepreneur state does not change the asset returns or investment opportunities available to households.

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

$$E_0 \max_{\{c_{it}, n_{it}, \Delta k_{it}\}} \sum_{t=0}^{\infty} \beta^t u[cn - G(h_{it}, n_{it})] - \mathbb{I}_{\Delta k_{it} \neq 0} \chi_{it}, \quad (9)$$

where χ_{it} is the utility cost of adjustment and $\mathbb{I}_{\Delta k_{it} \neq 0}$ is an indicator function that takes value 1 if a household adjusts its holdings of physical capital and 0 otherwise. The assumption of GHH preferences simplifies the numerical analysis substantially and allows us to abstract from the labor supply effects of uncertainty.¹⁷ The maximization is subject to the budget constraints described further below. The felicity function u exhibits a constant relative risk aversion (CRRA) with risk aversion parameter $\xi > 0$,

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

where $x_{it} = c_{it} - G(h_{it}, n_{it})$ is household i 's composite demand for goods consumption, c_{it} , and leisure, and G measures the disutility from work. 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 level, $P_t = (\int p_{jt}^{1-\eta} dj)^{\frac{1}{1-\eta}}$, 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, $G(h_{it}, n_{it})$, determines a household's labor supply given the aggregate wage rate, w_t , and a labor income tax, τ , through the first-order condition:

$$\frac{\partial G(h_{it}, n_{it})}{\partial n_{it}} = (1 - \tau)w_t h_{it}. \quad (10)$$

Assuming that G has a constant elasticity with respect to (w.r.t.) n , $\frac{\partial G(h_{it}, n_{it})}{\partial n_{it}} = (1 + \gamma) \frac{G(h_{it}, n_{it})}{n_{it}}$ with $\gamma > 0$, we can simplify the expression for the composite consumption good,

[Kaplan, Moll, and Violante \(2017\)](#), pure rents can be priced using the rate of return on capital. However, this approach requires not only assets (claims on physical capital and claims on future pure rents) to be illiquid, but also their corresponding asset income.

¹⁷[Basu and Bundick \(2017\)](#) (their appendix Figure D.6) show that the assumption of [Greenwood, Hercowitz, and Huffman \(1988\)](#) preferences leads to a smaller recessionary response to uncertainty than under [King, Plosser, and Rebelo \(1988\)](#) preferences. On the other hand, [Auclert and Rognlie \(2017\)](#) show that under [Greenwood, Hercowitz, and Huffman \(1988\)](#) preferences, monetary and fiscal multipliers tend to be larger than under [King, Plosser, and Rebelo \(1988\)](#) preferences. This ought to be taken into account when interpreting our results.

x_{it} , by making use of the first-order condition (10):

$$x_{it} = c_{it} - G(h_{it}, n_{it}) = c_{it} - \frac{(1 - \tau)w_t h_{it} n_{it}}{1 + \gamma}. \tag{11}$$

When the Frisch elasticity of labor supply is constant, the disutility of labor is always a constant fraction of labor income. Therefore, in both the budget constraint of the household and its felicity function, only after-tax income enters and neither hours worked nor productivity appears separately.

This implies that we can assume $G(h_{it}, n_{it}) = h_{it} \frac{n_{it}^{1+\gamma}}{1+\gamma}$ without further loss of generality as long as we treat the empirical distribution of income as a calibration target. This functional form simplifies the household problem, as h_{it} drops out from the first-order condition and all households supply the same number of hours, $n_{it} = N(w_t)$. Total effective labor input, $\int n_{it} h_{it} di$, is hence also equal to $N(w_t)$ because $\int h_{it} di = 1$. This means that we can read off productivity risk directly from the estimated income risk and treat both interchangeably. Correspondingly, as a shorthand notation, we will call the risk households face regarding their productivity *income risk* and call the shocks to h *income shocks*.

The households optimize subject to their budget constraint,

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

$$k_{it+1} \geq 0, b_{it+1} \geq \underline{B},$$

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, $\pi_t = \frac{P_t}{P_{t-1}}$ is realized inflation, and R is the nominal gross interest rate on bonds, which depends on the portfolio position of the household and the central bank’s interest rate R_t^b , which is set one period before. All households that decide not to participate in the capital market ($k_{it+1} = k_{it}$) still obtain dividends and can adjust their bond holdings. Depreciated capital has to be replaced for maintenance, such that the dividend, r_t , is the net return on capital.

We assume that there is a wasted intermediation cost, \bar{R} , when households resort to unsecured borrowing and we specify

$$R(b_{it}, R_t^b) = \begin{cases} R_t^b & \text{if } b_{it} \geq 0, \\ R_t^b + \bar{R} & \text{if } b_{it} < 0. \end{cases}$$

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

$$x_{it} + b_{it+1} + q_t k_{it+1} = b_{it} \frac{R(b_{it}, R_t^b)}{\pi_t} + (q_t + r_t)k_{it} + (1 - \tau) \left(\frac{\gamma}{1 + \gamma} w_t h_{it} N_t + \mathbb{I}_{h_{it}=0} \Pi_t \right), \tag{12}$$

$$k_{it+1} \geq 0, b_{it+1} \geq \underline{B}.$$

Since a household’s saving decision will be some nonlinear function of that household’s wealth and productivity, inflation, π_t , and accordingly aggregate real bond holdings, B_{t+1} ,

will be functions of the joint distribution, Θ_t , of (b, k, h) in 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 (2).

Three functions thus characterize the household's problem: 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:

$$\begin{aligned}
 &V_a(b, k, h; \Theta, R^b, s) \\
 &= \max_{k', b'_a} u[x(b, b'_a, k, k', h)] + \beta \text{EV}(b'_a, k', h', \Theta', R^{b'}, s'), \\
 &V_n(b, k, h; \Theta, R^b, s) \\
 &= \max_{b'_n} u[x(b, b'_n, k, k, h)] + \beta \text{EV}(b'_n, k, h', \Theta', R^{b'}, s'), \\
 &\text{EV}(b', k', h; \Theta, R^b, s) \\
 &= E_{\chi', h', s'} \{ \max[V_a(b', k', h'; \Theta', R^{b'}, s') - \chi', V_n(b', k', h'; \Theta', R^{b'}, s')] \}.
 \end{aligned} \tag{13}$$

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

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

such that the probability to adjust is given by

$$\nu^*(b', k', h'; \Theta', s') := F_\chi[V_a(b', k', h'; \Theta', R^{b'}, s') - V_n(b', k', h'; \Theta', R^{b'}, s')], \tag{14}$$

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 B and C.

4.2. 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,

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

but 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:

$$w_t = \alpha \text{MC}_t (K_t/N_t)^{1-\alpha}, \quad r_t + \delta = (1 - \alpha) \text{MC}_t (N_t/K_t)^\alpha. \quad (15)$$

4.3. Price Setting

Final-goods producers differentiate the intermediate good and set prices. We assume price adjustment costs à 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, managers maximize the present value of real profits given the demand for good j ,

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

and quadratic costs of price adjustment, that is, they maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t Y_t \left\{ \left(\frac{p_{jt}}{P_t} - \text{MC}_t \right) \left(\frac{p_{jt}}{P_t} \right)^{-\eta} - \frac{\eta}{2\kappa} \left(\log \frac{p_{jt}}{p_{j,t-1}} \right)^2 \right\}, \quad (16)$$

with a time constant discount factor.¹⁸ From the corresponding first-order condition for price setting, it is straightforward to derive the Phillips curve,

$$\log(\pi_t) = \beta E_t \left[\log(\pi_{t+1}) \frac{Y_{t+1}}{Y_t} \right] + \kappa \left(\text{MC}_t - \frac{\eta - 1}{\eta} \right), \quad (17)$$

where π_t is the gross inflation rate, $\pi_t := \frac{P_t}{P_{t-1}}$, and MC_t is 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 mass-zero group in the economy, their consumption does not show up in any resource constraint and 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

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

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

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

¹⁸The choice of the discount factor has relatively little impact on results. Given that we calibrate to a zero inflation steady state and approximate the aggregate dynamics linearly, only the steady-state value of the discount factor matters in the manager's problem. Our baseline sets the discount factor of managers and households to be equal. We tried setting the managers' time preference rate to the interest rate on bonds and to the median discount factor of entrepreneurs as alternatives. The results are robust.

4.4. *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:

$$\frac{R_{t+1}^b}{\bar{R}^b} = \left(\frac{R_t^b}{\bar{R}^b}\right)^{\rho_R} \left(\frac{\pi_t}{\bar{\pi}}\right)^{(1-\rho_R)\theta_\pi} \tag{18}$$

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 is θ_π , the stronger is the reaction of the central bank to deviations from the inflation target. When $\theta_\pi \rightarrow \infty$, inflation is perfectly stabilized at its steady-state value; $\rho_R \geq 0$ captures interest rate smoothing.

We assume that the government issues bonds according to the rule (cf. Woodford (1995))

$$\frac{B_{t+1}}{\bar{B}} = \left(\frac{B_t R_t^b / \pi_t}{\bar{B} \bar{R}^b / \bar{\pi}}\right)^{\rho_B} \left(\frac{\pi_t}{\bar{\pi}}\right)^{-\gamma_\pi} \left(\frac{\mathcal{T}_t}{\bar{\mathcal{T}}}\right)^{-\gamma_\mathcal{T}}, \tag{19}$$

using tax revenues, $\mathcal{T}_t = \tau(w_t N_t + \Pi_t)$, to finance government consumption, G_t , and interest on debt. The coefficient ρ_B captures whether and how fast the government seeks to repay its outstanding obligations, $B_t R_t^b / \pi_t$. For $\rho_B < 1$, the government actively stabilizes real government debt, and for $\rho_B = 1$, the government rolls over all outstanding debt including interest. The coefficients γ_π and $\gamma_\mathcal{T}$ capture the cyclicity of debt issuance: for $\gamma_\pi = \gamma_\mathcal{T} = 0$, new debt does not actively react to tax revenues and inflation, but only to the value of outstanding debt; for $\gamma_\pi > 0 > \gamma_\mathcal{T}$, debt is countercyclical; for $\gamma_\pi < 0 < \gamma_\mathcal{T}$, debt is procyclical.

4.5. *Goods, Bonds, Capital, and Labor Market Clearing*

The labor market clears at the competitive wage given in (15). The bond market clears whenever the equation

$$B_{t+1} = B^d(\Theta_t; R_t^b, s_t; q_t, \pi_t) := E[\nu^* b_a^* + (1 - \nu^*) b_n^*], \tag{20}$$

holds, where ν^* , b_a^* , and b_n^* are functions of the states $(b, k, h; R_t^b, s_t)$, of current prices q_t and π_t , and of expectations of future prices. Expectations in the right-hand-side expression 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,

$$q_t = 1 + \phi \frac{K_{t+1} - K_t}{K_t}, \tag{21}$$

$$K_{t+1} = K^d(\Theta_t; R_t^b, s_t; q_t, \pi_t) := E[\nu^* k^* + (1 - \nu^*) k],$$

where the first equation stems from competition in the production of capital goods and the second equation defines the aggregate supply of funds from households—both those that trade capital, ν^*k^* , and those that do not, $(1 - \nu^*)k$. Again ν^* and k^* are functions of the state variables $(\Theta_t; R_t^b, s_t)$, and current and expected future prices. The goods market then clears due to Walras' law whenever labor, bonds, and capital markets clear.

4.6. 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 the following statements hold:

(i) 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. Additionally, given the policy 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 (13).

(ii) 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, that is, (15), (17), (20), and (21) hold.

(iii) The actual and the perceived law of motion Γ coincide, that is, $\Theta' = \Gamma(\Theta, s')$.

4.7. Price-Level Determinacy

Since our economy is non-Ricardian, price-level determinacy depends not only on monetary policy alone, but also on fiscal policy (see Bénassy (2005), Leith and von Thadden (2008), for a treatment in an overlapping generations (OLG) framework, and Linnemann (2006) for the case of distortionary taxation) because the demand for bonds does not increase one-for-one in outstanding real government debt, even abstracting from the real effects of inflation. Thereby, the demand for government bonds creates a “nominal anchor.”

Given expected future inflation, a nominal interest rate, and a wealth distribution, households demand a real amount B^d of bonds. Any change in the price level scales the real liquid wealth holdings of all households as well as real government debt. Yet the change in real liquid wealth does not lead to a proportional increase in the demand for liquid wealth because households want to hold a certain portfolio structure and precautionary savings. Conversely, there is a “Pigou effect” on the demand for goods (Pigou (1943)) and only one price level clears the bond market. This is important as it constitutes the key difference from a representative agent model. Hagedorn (2016) provides a discussion of the special case of an interest rate peg and a government debt rule that allows the government to stabilize nominal government debt and shows that the price level is determinate then. We sketch the idea in Appendix D, where we also show that indeterminacy may still arise if the government overly aims at returning to the steady-state level of real debt. Conversely, if the fiscal policy does not stabilize real debt at all ($\rho_B \geq 1$), the monetary authority needs to violate the Taylor principle ($\theta_\pi < 1$) such that higher inflation reduces the real rate on bonds even in the long run so as to keep real debt stable.

5. NUMERICAL IMPLEMENTATION

The dynamic program (13) and, hence, the recursive equilibrium are not computable, because computation involves the infinite-dimensional object Θ_t . We discretize the distribution Θ_t and represent it by its histogram, a finite-dimensional object.

5.1. Solving the Household's Planning Problem

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).¹⁹

In solving for the household's policy functions, we apply an endogenous grid-point 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 (14) 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, that is, we operate a special case of the algorithm suggested by Fella (2014). Details on the algorithm can be found in Appendices B.4 and C.

5.2. 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.

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, with very similar results. Details on both methods can be found in Appendix E.

6. CALIBRATION

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

6.1. Technology and Preferences

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

6.1.1. 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 $\alpha = 0.73$.

¹⁹We solve the household policies for 80 points on the grid for bonds and on the grid for capital using log-scaled grids. We experimented with changing the number of grid points without a noticeable impact on results.

TABLE II
CALIBRATED PARAMETERS: FIRMS AND HOUSEHOLDS

Parameter	Value	Description	Target
<i>Households</i>			
β	0.9795	Discount factor	See Table III
ξ	4	Relative risk aversion	Standard value
γ	1	Inverse of Frisch elasticity	Chetty, Guren, Manoli, and Weber (2011)
μ_x	63,563	Participation utility costs	See Table III
σ_x	22,500	Participation utility costs	See Table III
\bar{R}	11%	Borrowing penalty	See Table III
<i>Intermediate Goods</i>			
α	0.70	Share of labor	Income share of labor of 2/3
δ	1.35%	Depreciation rate	NIPA: Fixed assets
<i>Final Goods</i>			
κ	0.09	Price stickiness	Mean price duration of 4 quarters
η	20	Elasticity of substitution	5% markup
<i>Capital Goods</i>			
ϕ	11.4	Capital adjustment costs	Relative investment volatility of 4.5

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 four quarters (in the equivalent Calvo setting), assuming flexible capital at the firm level. The steady-state marginal costs, $\frac{\eta-1}{\eta} = 0.95$, imply a markup of 5%.

We calibrate the adjustment cost of capital, $\phi = 11.4$, to match an investment-to-output volatility of 4.5 conditional on a TFP shock (see Appendix I).

6.1.2. Households

For the felicity function, $u = \frac{1}{1-\xi}x^{1-\xi}$, we set the coefficient of relative risk aversion, $\xi = 4$. 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)).

For the labor-income process, we use the estimated coefficients for the persistent component of after-tax household income from Section 2; see Table I. Because taxes are linear in our model, pre-tax and after-tax incomes are proportional, and our estimator takes out average tax rate changes by controlling for year effects. We calibrate the probability of leaving the entrepreneurial state to 1/16 per quarter following the numbers that Guvenen, Kaplan, and Song (2014a) report on the probability of dropping out of the top 1% income group in the United States (see their Table 2; roughly 25% per annum (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 time-discount factor, β , and the distribution of costs (pinned down by its mean and variance) of asset market participation, F_x , 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

TABLE III
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
Second 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

and produced. Because they are not productive assets in the NIPA sense, we disregard non-housing durable consumption goods. For the total value of illiquid assets relative to nominal GDP this implies a capital-to-output ratio of 286% and an annual real return for illiquid assets of 4.5%.²⁰ We use the Survey of Consumer Finances (SCF) to estimate the liquidity of household portfolios as described in Section 2. 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 (average 1983–2013: 9%).²¹ 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, Moll, and Violante (2017). Since we abstracted from consumer durables, we also disregard car wealth and auto loans here.

The empirical distribution of portfolio liquidity sheds light on how state-dependent liquidation decisions are, that is, whether the logistic distribution for adjustment costs has a high or a low variance. Figure 6 shows the average holding of liquid relative to illiquid assets over the period 1983–2013 in the SCF and implied by the model. Portfolio liquidity is estimated using a local linear regression as described in Section 2. Note that only households above the 20th percentile have typically nonnegligible 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.²²

Our model produces this downward sloping curve, too. The intuition is that households hold bonds because they provide better short-term consumption smoothing than 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 III summarizes how we match our targets from the wealth distribution. The calibrated adjustment costs imply an average adjustment frequency of 5.6% per quarter that increases for households with unbalanced portfolios to up to 14% probability of adjustment.²³ The average adjustment frequency is close to the implied adjustment probability in Kaplan and Violante (2014). The maximum adjustment cost a household is willing

²⁰We calibrate to the capital-to-GDP ratio in NIPA instead of using the illiquid assets-to-GDP ratio from the Flow of Funds (roughly 3.3) because the latter contains land as a nonproduced asset, debt assets turned illiquid when held in pension funds, and foreign illiquid assets, but they lack foreign owned capital used in production. We view it as important that our model replicates the production structure of the economy.

²¹This number is relatively close to the ratio of average U.S. federal debt held by domestic private agents relative to capital of 10.5%.

²²Despite this fact, the Gini coefficient of liquid wealth is larger than the Gini of illiquid wealth.

²³We provide robustness checks on the distribution of adjustment costs in Appendix L and find our results to be robust.




FIGURE 6.—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. The average is over the estimates from the SCFs for 1983–2013. We select only households composed of at least two adults whose head is between 30 and 55 years of age. Estimation is 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 can be negative.

to pay is equivalent to 6.3% of one quarter's consumption on average. The top 10% of households in terms of willingness to pay for adjusting their portfolios are willing to pay a felicity cost equivalent to 21% of a quarter's consumption.²⁴

Of course, it is a highly stylized treatment of the financial sector to assume 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 United States are either not publicly traded (11%) or held by agents other than households or depository institutions (in total 62% of all equities; out of these, 24% are held by mutual funds, 16% by the rest of the world, 12% by pension funds, and the remaining 12% by all other sectors). Importantly, the non-debt assets issued by these non-household and non-bank sectors are typically less

TABLE IV
CALIBRATED PARAMETERS: GOVERNMENT

Parameter	Value	Description	Target
<i>Monetary Policy</i>			
\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	Standard value
ρ_R	0.80	Inertia in Taylor rule	Standard value
<i>Fiscal Policy</i>			
ρ_B	0.86	Reaction to debt	Autocorrelation of government debt
γ_π	1.5	Reaction to inflation	Deficit response to uncertainty
γ_T	0.5075	Reaction to tax rev.	Standard deviation of deficits
τ	0.30	Labor tax rate	$G/Y = 20\%$

²⁴The consumption equivalents, ce , are calculated by solving $u(x_a^*) - u[(1 - ce)x_a^*] = V_a - V_n$, where the right-hand side is the maximal adjustment costs a household is willing to pay.

liquid.²⁵ On top of that, some of the corporate equities held by households will be held 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).

6.1.3. *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 captures persistence in the nominal interest rate. We set $\theta_\pi = 1.25$ and $\rho_R = 0.8$. 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).²⁶ To match the fraction of indebted households, we add a wedge between the lending and the borrowing rate of $\bar{R} = 11\%$ (annual).

6.1.4. *Fiscal Policy*

Government spending evolves according to a fiscal rule similar to Woodford (1995) or Bi, Leeper, and Leith (2013). 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%. We estimate the persistence of government debt by the autocorrelation of government debt in the United States, $\rho_B = 0.86$. We calibrate γ_π by matching the estimated peak response in primary surpluses after an uncertainty shock after four quarters, and calibrate γ_T such that the model with TFP shocks replicates the volatility of primary surpluses relative to GDP in the data.

7. QUANTITATIVE RESULTS

Having determined the parameters of the model, we can quantify the aggregate effects of shocks to household income risk in our model. They turn out to be very similar in size to what we found empirically in Section 2.

7.1. *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. 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 risk 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 from the policy functions of the model with risk shocks, estimate 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 of income risk by 1 standard deviation (an increase in the variance of income shocks of 54%). Given

²⁵Roughly 50% of the mutual funds are directly held by households; the rest are mostly held by pension funds.

²⁶The zero steady-state inflation assumption is equivalent to assuming any nonzero steady-state inflation rate together with perfect indexation to this rate (see, e.g., Basu and Bundick (2017)). Under indexation, the steady-state inflation rate drops out of the price-adjustment costs and hence the Phillips curve.




FIGURE 7.—Partial equilibrium response: Change in individual policy at constant prices and expectations. The response of individual consumption and asset demand policies at constant prices and price expectations to a 1 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 1 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.

that the so-specified planning problem of households uses the actual uncertainty process but fixed prices, this identifies the average partial equilibrium effect of uncertainty.

Figure 7 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 8 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 invest in liquid assets. This picture resembles what we found in Section 2. The increase in the liquidity of the portfolios is strongest for the lower middle class.

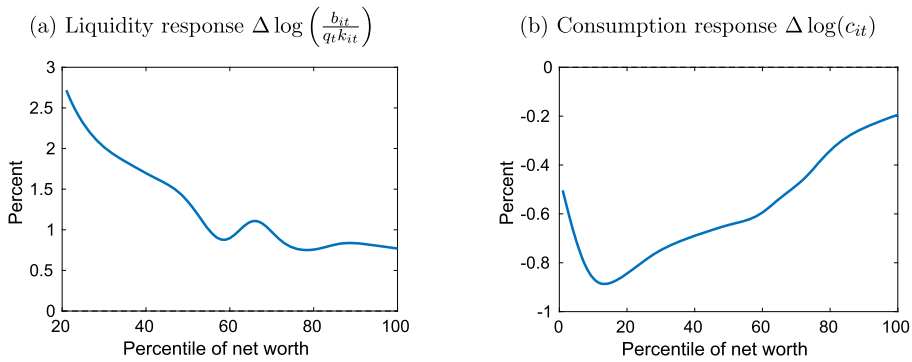


FIGURE 8.—General equilibrium response: Change in the liquidity of household portfolios and consumption. *Notes:* The change in the distribution of liquidity and consumption at all percentiles of the wealth distribution after two quarters at equilibrium prices and price expectations after a 1 standard deviation shock to income risk. The liquidity of the portfolios is averaged using frequency weights from the simulated wealth distribution and is 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.

7.2. Aggregate Consequences of Shocks to Household Income Risk

7.2.1. Main Findings

As Figure 7 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 employment, and dividends need to fall in equilibrium. Figure 9 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 policy reacts to the uncertainty-induced deflation by cutting rates. Fiscal policy expands government expenditure. After a 1 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: almost five times stronger than output.

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 typically show 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.

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, that is, the return difference between illiquid and liquid assets, increases. Quantitatively, we find that fluctuations in household income risk explain a significant fraction (21%) of the business cycle in terms of standard deviations; see Appendix I.

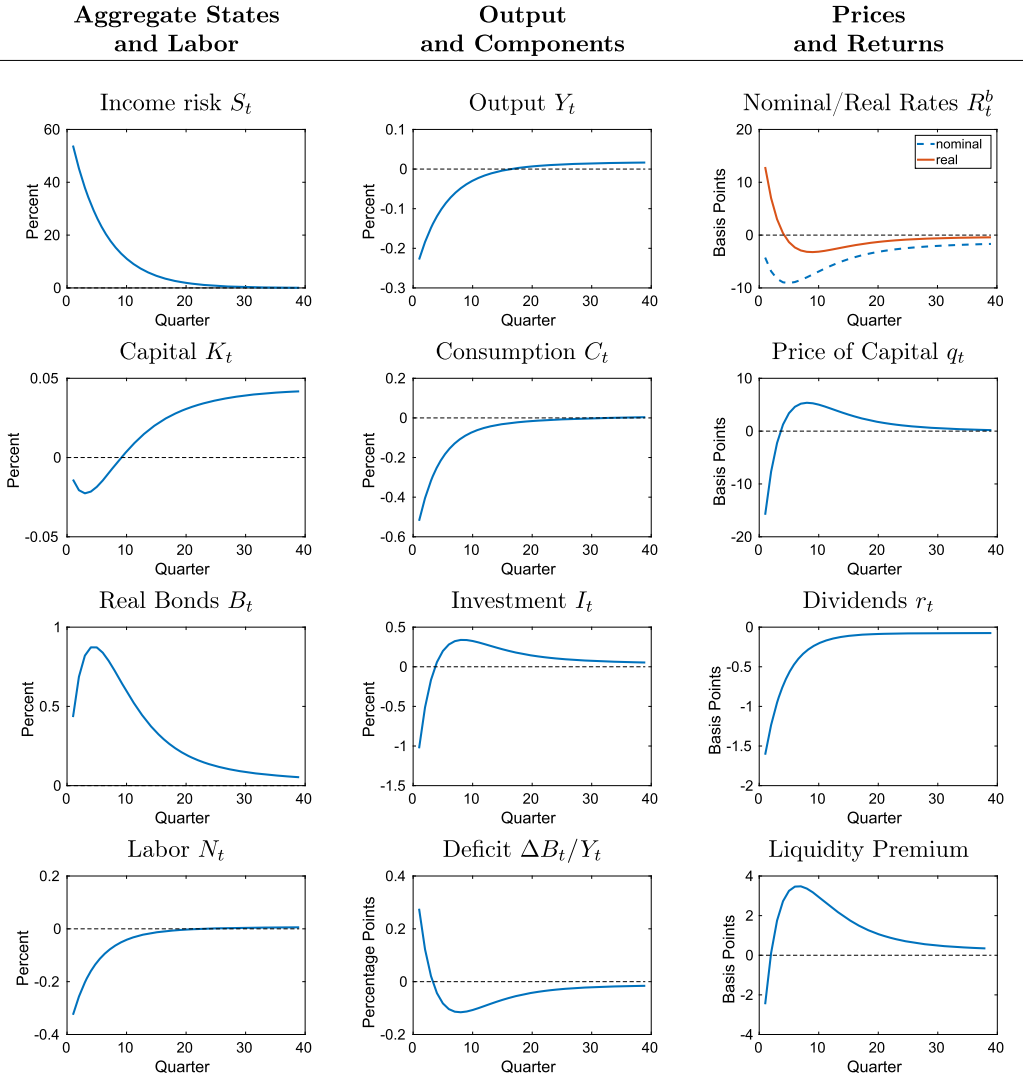


FIGURE 9.—Aggregate response to household income-risk shock. *Notes:* The liquidity premium is $\frac{E_t q_{t+1} + r_t}{q_t} - \frac{R_t^b}{E_t \pi_{t+1}}$. Impulse responses to a 1 standard deviation increase in the variance of income shocks. All rates (dividends, interest, liquidity premium) are *not* annualized.

Appendix J presents details on four model extensions that help us to understand the importance of certain aspects of the model. First, to understand the importance of portfolio rebalancing, we show that in an economy where capital is also liquid, the effects of income-risk shocks are substantially muted, and while output still falls, investment increases. Second, to show the importance of price stickiness, we solve the model for flexible prices. In that case, income risk has a small expansionary effect through investment, very similar to the monetary stabilization case we discuss in the next section. Inflation falls until the demand for goods equals supply, which itself does not depend on inflation. Hence, we see larger movements in inflation. These are large enough to let the interest rate, which is still given by the Taylor rule, fall sufficiently to undo the thirst for liquidity. Instead,

households then save more in illiquid assets and a small investment boom follows. Third, to understand the importance of stabilization policies, we look at a case where the interest rate is pegged and fiscal policy also does not stabilize. In this “crisis” case, the output loss is 10 times larger, roughly 2%. Fourth, to understand the potential role of endogenous liquidity, we study a version of the model where a banking sector generates liquid assets out of the illiquid investments of households, mortgaging a fixed fraction of the investment. Results are all very similar to our baseline. Further robustness checks regarding the calibration can be found in Appendix L, where we keep the model structure the same but vary the risk aversion, the Frisch elasticity, the manager’s discount factor, and the degree of state dependence in portfolio adjustment decisions.

7.2.2. *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, that is, by monetary policy, or by increasing the supply of government bonds, that is, through fiscal policy. Our baseline calibration is a mix of the two, following the empirical results in Section 2.

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, $\gamma_\pi = 0$, and an alternative scenario, where monetary policy keeps a nominal interest rate peg, $\theta_\pi = 0$, and fiscal policy reacts strongly to inflation, $\gamma_\pi = 100$.


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 10. Under monetary stabilization, the central bank increases the liquidity premium by lowering the interest rate on bonds until the excess supply of goods at steady-state inflation is eliminated. The lower interest rate spurs investment such that the capital stock increases. In the fiscal policy case, the government instead supplies the liquid assets that households demand 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.

7.3. *Redistributive and Welfare Effects*

Since the aggregate consequences of uncertainty shocks affect asset prices, dividends, wages, and entrepreneurial incomes differently, our model predicts that not all agents (equally) lose from the 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 1 standard deviation increase in uncertainty relative to its steady-state value for all

(a) Monetary Stabilization



(b) Fiscal Stabilization




FIGURE 10.—Aggregate response to household income-risk shock with stabilization policy. *Notes:* The liquidity premium is $\frac{E_t q_{t+1} + r_t}{q_t} - \frac{R_t^b}{E_t \pi_{t+1}}$. Impulse responses to a 1 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 $\gamma_\pi = 100$ ($\gamma_T = 0.5075$), and the Taylor-rule coefficient on inflation is zero, $\theta_\pi = 0$.

(b, k, h) triples.²⁷ 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

²⁷For 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.

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

$$\begin{aligned} CE(b, k, h) &= \left[\frac{EV(b, k, h; \Theta^{ss}, \sigma_s) - EV(b, k, h; \Theta^{ss}, 0)}{EU(b, k, h)} + 1 \right]^{1/(1-\xi)} - 1, \\ EU(b, k, h) &= \sum_{t=0}^{\infty} \beta^t u(x_t^*), \end{aligned} \quad (22)$$

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

Table V provides the consumption equivalents for both baseline and stabilization policies. The average welfare loss (one-sided) from the uncertainty shock is 0.27% of lifetime consumption. Table V 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.²⁸ 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 V.

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 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 0.04% of lifetime consumption (baseline minus fiscal stabilization).

Notwithstanding, 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 11(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 (first quintile), we observe that the relative welfare gains amount to as much as 10 basis points of lifetime consumption for indebted households (first quintile) to minus 10 basis points for households in the top quintile of liquid wealth; see Figure 11(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 more balanced portfolios, that are rich in liquid assets and illiquid assets, lose from monetary stabilization relative to the baseline because in equilibrium their asset returns fall.

²⁸The conditional expected productivity growth is $g(\sigma_{h,t}, \log \tilde{h}_{it}) := E_t \exp(\Delta \log \tilde{h}_{it+1}) \frac{f_{\tilde{h}_{it}}}{f_{\tilde{h}_{it+1}}} = \exp[(\rho_h - 1) \log \tilde{h}_{it} + 0.5 \sigma_{h,t}^2] \frac{f_{\tilde{h}_{it}}}{f_{\tilde{h}_{it+1}}}$. Across all households expected productivity growth is zero. The cross-derivative is negative: $\frac{\partial^2 g}{\partial \sigma_{h,t} \partial \tilde{h}_{it}} = g(\sigma_{h,t}, \log \tilde{h}_{it}) (\sigma_{h,t} [1 - (f_{\tilde{h}_{it}})^{-\rho_h^2}]) (\rho_h - 1) < 0$.

TABLE V
WELFARE EFFECTS OF HOUSEHOLD INCOME-RISK SHOCK^a

(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	0.13	0.14	-0.04	-0.10	1.03
Median	-0.08	-0.04	-0.01	0.02	0.08	0.20	-0.03	-0.15	1.88
	Quintiles of Capital Holdings					Average CE: -0.27			
	1	2	3	4	5				
Conditional	-0.06	-0.07	-0.04	0.02	0.15				
Median	-0.18	-0.10	-0.02	0.06	0.15				
(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	0.11	-0.07	-0.07	-0.05	0.09	0.17	-0.04	-0.10	0.07
Median	-0.11	-0.05	-0.05	-0.05	-0.02	0.20	-0.05	-0.15	0.21
	Quintiles of Capital Holdings					Average CE: -0.24			
	1	2	3	4	5				
Conditional	-0.02	-0.09	-0.06	-0.01	0.17				
Median	-0.20	-0.14	-0.04	0.04	0.19				
(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	0.07	-0.11	-0.08	-0.02	0.13	0.16	-0.04	-0.09	0.23
Median	-0.13	-0.08	-0.06	-0.03	0.05	0.19	-0.06	-0.18	0.18
	Quintiles of Capital Holdings					Average CE: -0.23			
	1	2	3	4	5				
Conditional	-0.06	-0.09	-0.06	0.01	0.18				
Median	-0.22	-0.14	-0.06	0.04	0.18				

^aWelfare costs of a 1 standard deviation increase in income risk in terms of consumption equivalents (CE) as defined in (22) in percent minus the population average for each regime. "Conditional" refers to integrating out the other individual states, whereas "Median" refers to median asset holdings/productivity in the other states. Panel (b) policy coefficients are $\theta_\pi = 100$, and $\gamma_\pi = 0$, $\gamma_T = 0$; panel (c) policy coefficients are $\theta_\pi = 0$, $\gamma_\pi = 100$, and $\gamma_T = 0.5075$.

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 tractable framework that combines nominal rigidities and incomplete markets in which households choose portfolios of liquid paper and illiquid physical assets. 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 United States, we find that, in line with the empirical evidence, a spike in income risk 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

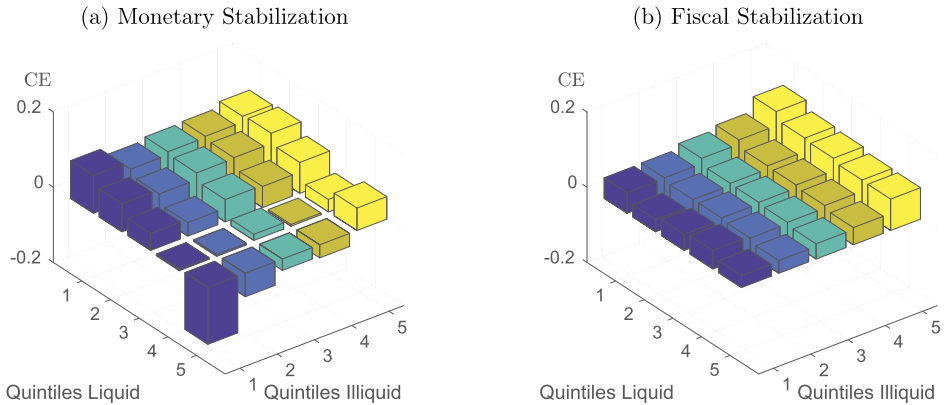


FIGURE 11.—Welfare effects of household income-risk shock with stabilization. *Notes:* Welfare costs of a 1 standard deviation increase in income risk relative to the baseline policy specification in terms of consumption equivalents (CE) as defined in (22) and integrating out h_{it} (excluding entrepreneurs). The graphs refer to the conditional expectations of CE with respect to the joint distribution of liquid and illiquid assets for workers. Left panel: the Taylor-rule coefficient on inflation is set to $\theta_\pi = 100$, and fiscal policy does not respond to inflation $\gamma_\pi = 0$ and $\gamma_\tau = 0$. Right panel: fiscal policy strongly responds to inflation $\gamma_\pi = 100$ ($\gamma_\tau = 0.5075$), and the Taylor-rule coefficient on inflation is zero, $\theta_\pi = 0$.

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 uncertainty shocks crucially depend on a household's asset position and the stance of monetary and fiscal policy. Monetary policy that lowers the return on liquid assets in times of increased uncertainty limits the negative aggregate consequences of uncertainty 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 future capital and wages.

This highlights the importance of the supply of liquid funds in the economy. We have extended our baseline model to incorporate inside money created by banks in exchange for mortgages. This extension sketches the effect it has when households in times of high uncertainty demand less illiquid assets and hence banks write fewer mortgages: When uncertainty increases, the supply of liquid assets is depressed, while demand for liquid assets is high. However, it would require going beyond our current model with two assets to obtain a detailed account of these effects.

Relative to the existing literature on uncertainty, our model introduces portfolio rebalancing of households as a new channel to explain how an increase in uncertainty can lead to a recession. In modeling this new channel, we abstract from a richer model of firms and therefore from at least two important channels through which uncertainty impacts economic activity: first, from a “wait-and-see” channel that leads to a freeze in investment activity (see Bloom (2009), Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012), Bachmann and Bayer (2013)); second, from a cost-of-finance channel (see, e.g., Gilchrist, Sim, and Zakrajšek (2014)), in which increased uncertainty magnifies the financial frictions firms face. In terms of our model, both channels imply that the effective return on illiquid funds declines, a situation that should amplify the household sector's flight to liquidity. In addition, a richer model of the firm sector would allow for introducing liquidity concerns for the firm sector, too (see, e.g., Del Negro, Eggertsson, Ferrero, and Kiyotaki

(Del Negro et al. (2017))). Changes in aggregate uncertainty come on top of the idiosyncratic uncertainty changes we highlight and likely go hand in hand with them, affecting aggregate outcomes through similar precautionary savings channels (Basu and Bundick (2017)) or other more complementary ones (Fernández-Villaverde et al. (2015)). 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, Ozkan, and Song (2014b).

REFERENCES

- AIYAGARI, S. R. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109 (3), 659–684. [256]
- ARELLANO, C., Y. BAI, AND P. KEHOE (2012): “Financial Markets and Fluctuations in Uncertainty,” Federal Reserve Bank of Minneapolis, Research Department Staff Report 466. [258]
- AUCLERT, A. (2015): “Monetary Policy and the Redistribution Channel,” Report, MIT. [258]
- AUCLERT, A., AND M. ROGNLIE (2017): “A Note on Multipliers in NK Models With GHH Preferences,” Report, Stanford. [269]
- BACHMANN, R., AND C. BAYER (2013): “Wait-and-See” Business Cycles? *Journal of Monetary Economics*, 60 (6), 704–719. [258,287]
- BASU, S., AND B. BUNDICK (2017): “Uncertainty Shocks in a Model of Effective Demand,” *Econometrica*, 85 (3), 937–958. [258,269,279,288]
- BAUMOL, W. J. (1952): “The Transactions Demand for Cash: An Inventory Theoretic Approach,” *Quarterly Journal of Economics*, 66 (4), 545–556. [256]
- BAYER, C., R. LUETTICKE L. PHAM-DAO, AND V. TJADEN (2019): “Supplement to ‘Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk,’” *Econometrica Supplemental Material*, 87, <https://doi.org/10.3982/ECTA13601>. [259]
- BÉNASSY, J.-P. (2005): “Interest Rate Rules, Price Determinacy and the Value of Money in a non-Ricardian World,” *Review of Economic Dynamics*, 8 (3), 651–667. [274]
- BEWLEY, T. (1980): “Interest Bearing Money and the Equilibrium Stock of Capital,” Report. [256]
- BI, H., E. LEEPER, AND C. LEITH (2013): “Uncertain Fiscal Consolidations,” *Economic Journal*, 123 (566), 31–63. [279]
- BLOOM, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77 (3), 623–685. [258,287]
- BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. TERRY (2012): “Really Uncertain Business Cycles,” NBER Working Paper Series, 18245. [287]
- BORN, B., AND J. PFEIFER (2014): “Policy Risk and the Business Cycle,” *Journal of Monetary Economics*, 68, 68–85. [258]
- BRUNNERMEIER, M. K., AND Y. SANNIKOV (2016): “The I Theory of Money,” NBER Working Paper Series, (22533). [268]
- CARROLL, C. (2006): “The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems,” *Economics Letters*, 91 (3), 312–320. [275]
- CASTANEDA, A., J. DIAZ-GIMENEZ, AND J.-V. RIOS-RULL (1998): “Exploring the Income Distribution Business Cycle Dynamics,” *Journal of Monetary Economics*, 42 (1), 93–130. [269]
- CHALLE, E., AND X. RAGOT (2016): “Precautionary Saving Over the Business Cycle,” *Economic Journal*, 126 (590), 135–164. [258]
- CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (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. [276]
- CHRISTIANO, L. J., R. MOTTO, AND M. ROSTAGNO (2014): “Risk Shocks,” *American Economic Review*, 104 (1), 27–65. [258]
- CHUGH, S. K. (2016): “Firm Risk and Leverage-Based Business Cycles,” *Review of Economic Dynamics*, 20 (Supplement C), 111–131. [258]
- DEATON, A. (1985): “Panel Data From Time Series of Cross-Sections,” *Journal of Econometrics*, 30 (1-2), 109–126. [261]
- DEL NEGRO, M., G. EGGERTSSON, A. FERRERO, AND N. KIYOTAKI (2017): “The Great Escape? A Quantitative Evaluation of the Fed’s Liquidity Facilities,” *American Economic Review*, 107 (3), 824–857. [288]
- DEN HAAN, W. J., P. RENDAHL, AND M. RIEGLER (2017): “Unemployment (Fears) and Deflationary Spirals,” *Journal of the European Economic Association* (forthcoming). [258]
- DI TELLA, S. (2018): “Uncertainty Shocks and Balance Sheet Recessions,” *Journal of Political Economy* (forthcoming). [258]

- FELLA, G. (2014): "A Generalized Endogenous Grid Method for non-smooth and non-concave Problems," *Review of Economic Dynamics*, 17 (2), 329–344. [275]
- FERNALD, J. (2012): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," *Federal Reserve Bank of San Francisco, WP*, 19. [263]
- FERNÁNDEZ-VILLAVARDE, J., P. A. GUERRÓN-QUINTANA, K. KUESTER, AND J. RUBIO-RAMÍREZ (2015): "Fiscal Volatility Shocks and Economic Activity," *American Economic Review*, 105 (11), 3352–3384. [258, 288]
- GILCHRIST, S., J. W. SIM, AND E. ZAKRAJŠEK (2014): "Uncertainty, Financial Frictions, and Investment Dynamics," NBER Working Paper Series, 20038. [258,287]
- GORNEMANN, N., K. KUESTER, AND M. NAKAJIMA (2012): "Monetary Policy With Heterogeneous Agents," Federal Reserve Bank of Philadelphia, WP No. 12-21. [258]
- GREENWOOD, J., Z. HERCOWITZ, AND G. W. HUFFMAN (1988): "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, 78 (3), 402–417. [258,269]
- GUERRIERI, V., AND G. LORENZONI (2017): "Credit Crises, Precautionary Savings, and the Liquidity Trap," *Quarterly Journal of Economics*, 132 (3), 1427–1467. [258]
- GUVENEN, F., G. KAPLAN, AND J. SONG (2014a): "The Glass Ceiling and the Paper Floor: Gender Differences Among Top Earners, 1981–2012," NBER Working Paper Series, 20560. [276]
- GUVENEN, F., S. OZKAN, AND J. SONG (2014b): "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122 (3), 621–660. [256,288]
- HAGEDORN, M. (2016): "A Demand Theory of the Price Level," CEPR Discussion Paper, (11364). [274]
- HALL, R. E. (1989): "Invariance Properties of Solow's Productivity Residual," NBER Working Paper Series, 3034. [263]
- HINTERMAIER, T., AND W. KOENIGER (2010): "The Method of Endogenous Gridpoints With Occasionally Binding Constraints Among Endogenous Variables," *Journal of Economic Dynamics and Control*, 34 (10), 2074–2088. [275]
- HUGGETT, M. (1993): "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17 (5-6), 953–969. [256]
- KAPLAN, G., AND G. L. VIOLANTE (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82, 1199–1239. [256,264,267,277]
- KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2017): "Monetary Policy According to HANK," *American Economic Review* (forthcoming). [256,258,264,269,277]
- KING, R. G., C. I. PLOSSER, AND S. T. REBELO (1988): "Production, Growth and Business Cycles: I. The Basic Neoclassical Model," *Journal of Monetary Economics*, 21 (2), 195–232. [269]
- LEITH, C., AND L. VON THADDEN (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. [274]
- LINNMANN, L. (2006): "Interest Rate Policy, Debt, and Indeterminacy With Distortionary Taxation," *Journal of Economic Dynamics and Control*, 30 (3), 487–510. [274]
- LUETTICKE, R. (2017): "Transmission of Monetary Policy With Heterogeneity in Household Portfolios," Report, University College London. [258]
- MCKAY, A., AND R. REIS (2016): "The Role of Automatic Stabilizers in the US Business Cycle," *Econometrica*, 84 (1), 141–194. [258]
- MCKAY, A., E. NAKAMURA, AND J. STEINSSON (2016): "The Power of Forward Guidance Revisited," *American Economic Review*, 106 (10), 3133–3158. [258]
- NARITA, F. (2011): "Hidden Actions, Risk-Taking and Uncertainty Shocks," Report, University of Minnesota. [258]
- PANOUSI, V., AND D. PAPANIKOLAOU (2012): "Investment, Idiosyncratic Risk, and Ownership," *Journal of Finance*, 67 (3), 1113–1148. [258]
- PIGOU, A. C. (1943): "The Classical Stationary State," *Economic Journal*, 53 (212), 343–351. [274]
- RAVN, M. O., AND V. STERK (2017): "Job Uncertainty and Deep Recessions," *Journal of Monetary Economics*, 90 (Supplement C), 125–141. [258,268]
- REITER, M. (2009): "Solving Heterogeneous-Agent Models by Projection and Perturbation," *Journal of Economic Dynamics and Control*, 33 (3), 649–665. [275]
- ROMEI, F. (2015): "Need for (the Right) Speed: The Timing and Composition of Public Debt Deleveraging," Report. [268]
- ROTEMBERG, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 1187–1211. [256, 267,272]
- SCHAAL, E. (2012): "Uncertainty, Productivity and Unemployment in the Great Recession," Report, Federal Reserve Bank of Minneapolis. [258]

- SMETS, F., AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97 (3), 586–606. [258]
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2001): "The Welfare Cost of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk," *European Economic Review*, 45 (7), 1311–1339. [255,259,261]
- (2004): "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112 (3), 695–717. [259,261]
- TAUCHEN, G. (1986): "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20 (2), 177–181. [275]
- TAYLOR, J. B. (1993): "Discretion versus Policy Rules in Practice," in *In Carnegie-Rochester Conference Series on Public Policy*, Vol. 39. Elsevier, 195–214. [273]
- TOBIN, J. (1956): "The Interest-Elasticity of Transactions Demand for Cash," *Review of Economics and Statistics*, 38 (3), 241–247. [256]
- VAVRA, J. (2014): "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *Quarterly Journal of Economics*, 129 (1), 215–258. [258]
- WERNING, I. (2015): "Incomplete Markets and Aggregate Demand," NBER Working Paper Series, 21448. [258]
- WOODFORD, M. (1995): "Price-Level Determinacy Without Control of a Monetary Aggregate," *Carnegie-Rochester Conference Series on Public Policy*, 43 (Supplement C), 1–46. [273,279]

Co-editor Giovanni L. Violante handled this manuscript.

Manuscript received 2 July, 2015; final version accepted 31 May, 2018; available online 22 June, 2018.