Department of Economics

Copenhagen Business School

Working paper 8-2019

Consumption Dynamics under Time-Varying Unemployment Risk

Karl Harmenberg
Erik Öberg

Consumption Dynamics under Time-Varying Unemployment Risk[*]

Karl Harmenberg[†]   Erik Öberg[‡]

July 27, 2019

Abstract

Private consumption demand falls in response to increased unemployment risk during a recession, as households increase their precautionary savings and postpone irreversible durable investments. The postponement effect is seven times as large as the precautionary-savings effect in a calibrated buffer-stock savings model. In consequence, anticipation of future unemployment risk is more important than realized unemployment shocks in accounting for durable expenditure dynamics during recessions, while the opposite is true for nondurables. The importance of anticipation of future unemployment risk also means that having many 'hand-to-mouth' households, who do not respond to changes in income risk, significantly dampens the demand response for durables to an adverse labor market shock. We find that the model elasticities of durable and nondurable expenditures with respect to unemployment risk are close to what we estimate in micro survey data.

1 Introduction

With the advent of heterogeneous-agent models for business-cycle analysis, a recent literature has emphasized the role of income risk as a key driver of fluctuations in consumption demand: in recessions, idiosyncratic income risk increases (Storesletten et al., 2004, Guvenen et al., 2014), which induces households to increase their buffer stock of liquid assets (Challe and Ragot, 2016, Krueger et al., 2016, McKay, 2017). The resulting fall in consumption demand may increase overall business-cycle volatility (Ravn and Sterk, 2017, Den Haan et al., 2018) and warrant recalibration of optimal countercyclical fiscal and monetary policies as a response (McKay and Reis, 2017, Kekre, 2018, Berger et al., 2018a). In this paper, we quantitatively assess the importance of fluctuations in perceived income risk for consumption dynamics during recessions.

To this end, we investigate the magnitude and mechanisms of the expenditure demand response to an adverse labor-market shock, using a calibrated model of household consumption-saving decisions. A key value added of our study is that we separate between expenditures on nondurable and durable goods, motivated by the fact that consumer durables make up a large share of fluctuations in aggregate consumption demand.[1]

The adverse labor-market shock affects beliefs about the likelihood of becoming unemployed ex ante, and generates more unemployment spells ex post. Our main finding is that for durable expenditures, the ex-ante belief channel is significantly more important than the ex-post realized-shocks channel, while the opposite is true for nondurable expenditures. The difference arises due to an amplified durable expenditure response to the increase in income risk associated with the increased likelihood of becoming unemployed. We show that this finding has important implications for the role of household liquidity in shaping aggregate consumption dynamics. Finally, using micro survey data, we estimate expenditure responses to fluctuations in unemployment risk, and show that these responses are closely in line with the model, both for durables and nondurables.

Specifically, we extend a standard buffer-stock savings model to include both nondurable and durable consumption, with durable purchases subject to a non-convex adjustment cost as in Grossman and Laroque (1990) and Berger and Vavra (2015). Because of the adjustment cost, it is costly for households to frequently adjust their durable stock, which generates lumpy purchase behavior consistent with the data. Due to prudent preferences and a potentially binding credit constraint, nondurable expenditures respond to heightened income risk through a standard precautionary-savings motive (Kimball, 1990, Carroll and Kimball, 2001). In addition, the adjustment cost generate a real-options motive for households to postpone durable expenditures when facing heightened income risk (Bernanke, 1983, Dixit and Pindyck, 1994).

We estimate the model processes for employment and income and calibrate the preference and technology parameters to match sample moments in Italian micro data (which is the data we use for the empirical

---

[1]See Stock and Watson (1999). In the first year of the Great Recession, US car expenditures fell by \$107 billion, i.e, by approximately \$330 per capita (Dupor et al., 2018).

investigation). We then compute, in partial equilibrium, the aggregate impulse-response functions (IRFs) to a large and persistent increase in the job-separation rate in line with the Italian experience during the recent Eurozone crisis. The shock raises the probability of becoming unemployed ex ante and depresses income ex post as more households become unemployed.

We decompose the expenditure response into the impact of a change in household beliefs and the impact of realizations of the unemployment shocks. The *belief channel* accounts for 75 percent of the total expenditure response for durables, while only 35 percent for nondurables. The difference stems entirely from the adjustment costs involved in durable purchases. In a corresponding flexible-adjustment model, the belief channel is as important for durables as it is for nondurables.

The increase in the likelihood of becoming unemployed depresses expected income, the first moment of the income process, and raises income risk, the higher-order moments. To investigate the source of the amplified belief-channel response for durables, we perform a second decomposition into the contribution of the first moment and the higher-order moments of the income process. The amplified belief-channel response for durables stems entirely from an amplified response to the increase in income risk, which again is due to the adjustment cost. The risk-channel response is raised by a factor 7 in the model compared to the flexible-adjustment model. In contrast, the expected-income channel response explains most of the belief-channel response for nondurables. Our results suggest that the real-option effect for durables is signficantly more important than the precautionary-savings effect for nondurables for shaping consumption dynamics during recessions.

By varying the model discount factor, we trace out how the expenditure response depends on households' holdings of liquid assets. For both nondurables and durables, the realized-shocks channel response is monotonically decreasing in household asset holdings, as closer to a binding credit constraint, the expenditure response to income shocks increases. In contrast, the belief-channel response is U-shaped. When households have a lot of liquid assets and are well insured, they do not respond to the risk component of the adverse labor market shock; when households face a binding credit constraint, they behave 'hand-to-mouth' and do not respond to beliefs about their future income at all. The belief-channel response is maximized for intermediate values, where households are simultaneously forward-looking and not well insured. Because the total response for durables is primarily driven by the risk-channel response, having many households close to the credit constraint dampens the total expenditure response to the adverse labor-market shock. For nondurables, the total response is primarily driven the realized-shocks channel response, and the total response is therefore close to monotonically increasing in the number of households close to the credit constraint.

For computing the aggregate impulse-response function, we develop a novel method for simulating heterogenous-agent models with permanent-income shocks, which greatly reduces the computational cost of computing aggregate statistics. We prove that for simulating aggregate statistics, one does not need to

keep track of the distribution of permanent income, but can instead simulate the model using distorted *permanent-income neutral probabilities* instead. As a result, the poor convergence properties associated with the fat tail in the income distribution (which results from the permanent income shocks) are bypassed.

We corroborate our findings with reduced-form evidence from Italian micro survey data. Based on observable household characteristics, we estimate household-level fluctuations in the probability of becoming unemployed. We then estimate the effect of a change in the unemployment probability on purchases of medium-sized durables (motor vehicles and large furniture) and nondurable goods separately.[2] We document that durable purchases respond strongly to fluctuations in unemployment probability, with a semi-elasticity around 2.5, while the semi-elasticity for nondurable expenditures is small and statistically indistinguishable from 0. We relate these estimates to our model by matching the idiosyncratic unemployment probability distribution of the model to the survey data and running the same regressions on model-generated data. Although not targeted in the calibration, the expenditure responses to fluctuations in unemployment probability growth are close to the empirical estimates, both for durable and nondurable goods. The fact that these model elasticities are close to the data both for nondurable and durable goods corroborates the relative strength of the precautionary-savings motive (which is active for both nondurable and durable goods) and the real-option motive (which is only active for durable goods) of our model.

Related litterature.   Our paper adds to the growing literature emphasizing the role of heterogeneity and incomplete markets in business-cycle consumption dynamics. One strand of this literature has emphasized that limited insurance raises the consumption response to transitory income shocks (the 'marginal propensity to consume') with implications for the transmission of aggregate shocks (Kaplan and Violante, 2014, Carroll et al., 2017, Kaplan et al., 2018, Auclert, 2018, Berger et al., 2018b). Another strand has emphasized that limited insurance raises the consumption response to fluctuations in idiosyncratic income risk (Challe and Ragot, 2016, Krueger et al., 2016, McKay, 2017, Den Haan et al., 2018, Heathcote and Perri, 2018). Both strands have almost exclusively focused on nondurable consumption alone. We contribute by showing that the relative strengths of these two channels are different for durable expenditures and nondurable expenditures, with risk being more important for durable expenditures.

By so doing, our paper connects to the recent literature that has shown that durable expenditures respond differently from nondurable expenditures to stimulus policy. In a model framework similar to ours, Berger and Vavra (2015) show that, due to the adjustment costs involved in durable purchases, durable expenditure dynamics are state dependent: the responsiveness to stimulus policy is lower in recession than in expansions. McKay and Wieland (2019) show that, in contrast to nondurables, durable expenditures respond

---

[2]Benito (2006) similarly uses panel data to estimate expenditure repsonses to unemployment risk. To the best of our knowledge, previous micro-level studies on the relation between expenditures and income risk have only used cross-sectional variation, see Carroll et al. (2003) for nondurable expenditures and Eberly (1994), Foote et al. (2000) and Bertola et al. (2005) for durable expenditures.

more strongly to contemporaneous changes in the interest rate compared to changes in the future (forward guidance). Sterk and Tenreyro (2018) argue that the resdistributive wealth effects of an accommodative monetary policy shock induces substitution from nondurables to durables. Another strand has emphasized that durables are more interest-rate sensitive than nondurables and studied the implications for monetary policy, see, e.g, Erceg and Levin (2006), Barsky et al. (2007) and Basu and De Leo (2017).

Our investigation builds on an earlier literature studying investment decisions under non-convex adjustment costs, going back to the seminal contributions of Arrow et al. (1951), Grossman and Laroque (1990) and Dixit and Pindyck (1994). Chetty and Szeidl (2016) show that non-convex adjustment costs in themselves smoothen the response to first-moment income shocks. Hassler (1996) shows that it is transitory, rather than permanent, shocks to income risk that may generate large investment responses through raising the option value of postponement. A large literature has investigated the role of non-convex adjustment costs for business-cycle dynamics in quantative models of firm investment, see, e.g., Bloom (2009), Gilchrist et al. (2014), Bachmann and Bayer (2013), Bloom et al. (2018). For consumer durables, Druedahl (2015) also highlight that non-convex adjustment costs raise the durable expenditure response to fluctuations in income risk. Bayer et al. (2019) show that the precautionary-savings motive induces substitution from illiquid to liquid assets in response to an increase in income risk, which in turn affect investment dynamics. We contribute by showing that the non-convex adjustment cost has a quantitatively large effect on the demand response of consumer durables to fluctuations in income risk, that this primarily stems from the real-options effect, and that this is consistent with micro data evidence.

Finally, our results accord well with the evidence presented by Carroll and Dunn (1997) and Dupor et al. (2018). Carroll and Dunn (1997) show that unemployment expectations are a good predictor of fluctuations in aggregate spending on consumer durables. Dupor et al. (2018) trace out the sources of the 2008 U.S. fall in car purchases and find that the fall in house prices can only account for twenty percent of the fall, attributing the remainder to falling household income expectations.

Dispostion. Our analysis proceeds as follows. In Section 2, we describe the model setup, parametric assumptions and the calibration procedure. In Section 3, we analyze the consumption response to the job-separation rate shock and investigate the underlying mechanisms. In Section 4, we corroborate the model predictions with reduced-form evidence using micro survey data. Section 5 concludes the paper.

2 Model

Our model extends a standard buffer-stock model (see, e.g., Carroll (1997)) to separate between nondurable and durable consumption goods with an adjustment cost for the durable good. We first describe the general formulation of the household optimization problem. Then, we specify and estimate the household income

and employment processes and calibrate the model parameters.

2.1 The household problem

The economy is populated by a continuum of households on the unit interval. All households are ex-ante identical but ex-post different due to different realizations of the shock processes. Time is infinite and discrete and each period corresponds to a quarter. Each household solves the problem

$$\max_{\{C_{it},D_{it},B_{it}\}_{t=0}^{\infty}} \quad E_0 \sum_{t=0}^{\infty} \beta^t u(C_{it}, D_{it})$$

$$\text{s.t.} \quad C_{it} + D_{it} + qB_{it} \leqslant \Upsilon(Y_{it}, n_{it}) + (1-\delta)D_{it-1} + B_{it-1} - A(D_{it}, D_{it-1}), \quad (1)$$

$$B_{it} \geqslant -\chi\left[(1-\delta)D_{it} - A(0, D_{it-1})\right], \quad (2)$$

$$C_{it}, D_{it} \geqslant 0. \quad (3)$$

We now describe each compoment of this decision problem.

Preferences. Households have time-additive homothetic preferences over the consumption of nondurable goods C_{it} and the consumption of durable goods D_{it}. They have a subjective discount factor $\hat{\beta}$. To maintain a stationary income distribution (discussed below), households die with probability ω. The effective discount factor is thus $\beta = (1-\omega)\hat{\beta}$.

Prices and technology. Households can purchase nondurable and durable goods at a unitary relative price and invest in a risk free liquid asset B_{it} at a price q.[3] The durable good depreciates by δ each period.

Income process. In each period t, households receive income Υ_{it} that depends on their employment status n_{it} and their earnings potential Y_{it}. Y_{it} is a continuous variable while n_{it} takes the value of 0 if the household is unemployed and 1 if the household is employed. The earnings potential Y_{it} and the employment status n_{it} are stochastic and their processes are specified in Subsection 2.2 below.

Market frictions. Households face two market frictions. First, adjusting the durable stock is associated with an adjustment cost $A(D_{it}, D_{it-1})$. We follow Grossman and Laroque (1990) and assume that, conditional on adjusting, the adjustment cost $A(D_{it}, D_{it-1})$ is linear in the stock of durable goods prior to the adjustment decision and zero otherwise:

$$A(D_{it}, D_{it-1}) = \begin{cases} 0 & \text{if } D_{it} = (1-\delta)D_{it-1}, \\ hD_{it-1} & \text{if } D_{it} \neq (1-\delta)D_{it-1}. \end{cases} \quad (4)$$

---

[3] Under homothetic preferences, a constant relative price different from unity would only linearly rescale the decision functions.

When households readjust, they can only recover a fraction of the value of their previous investments into their durable stock, making previous investments into the durable stock *partially irreversible.* For durables, the main sources of partial irreversibility are likely to be their illiquidity and the prevalence of rebate prices in the market for used goods.[4] The adjustment-cost parameter h can be interpreted as the average resale loss when selling the replaced stock D_{it-1} in the second hand market.[5] With this interpretation of the adjustment cost, D_{it} should be interpreted as the stock of a "typical" durable good, e.g., cars, and not the total stock composed of different durable goods. We do not expect households to sell off their current stock of furniture when they purchase a new car, and vice versa. This interpretation is important when we choose the moments to which we calibrate our model, as described in Section 2.2.

There are other frictions in the household's adjustment of the stock of durables, which are less well described by the assumption of a linear adjustment cost. For example, car buyers in Italy need to pay a flat fee to the Public Automobile Registry (Bertola et al., 2005). The buyer's fee is independent of the value of the car or any other characteristics of the household, and is better described as a fixed cost. The time spent purchasing a durable good is better modeled as a cost proportional to labor income. Although we cannot test it directly, our best guess is that such costs are small in relation to the resale loss, especially for medium-sized durable goods such as motor vehicles and large furniture.

Our second market friction is that households face a credit constraint, given by Equation (2). Specifically, households cannot borrow more than a fraction χ of the pledgable part of their next period stock of durable goods $[(1-\delta)D_{it} - A(0, D_{it})]$.

Recursive formulation. For interpretative and computational purposes, it is useful to recast the household problem on recursive form. We collect all exogenous state variables in $\mathcal{S}$ and impose that these follow a Markov process. Denote the ingoing values of the durable stock, prior to depreciation, by D and the ingoing value of the liquid asset stock by B. Denote the choice of nondurable consumption, durables and liquid assets by C, D' and B'. Define $V_{NA}(\cdot)$ as the value function conditional on not adjusting the stock of durables, $V_A(\cdot)$ as the value function conditional on adjusting the stock of durables and the collateral requirement

---

[4]Such equilibria could be explained by, e.g., asymmetric information (Akerlof, 1970).

[5]See Gavazza and Lanteri (2018) for a model of endogenous illiquidy of new car purchases through equilibrium dynamics in the secondary market.

parameter $\hat{\chi} \equiv \chi(1 - \delta - h)$. The recursive representation is then

$$
\begin{aligned}
V_{NA}(B, D; \mathcal{S}) = \quad &\max_{C, B'} u(C, D') + \beta EV(B', D'; \mathcal{S}') && (5) \\
\text{s.t.} \quad &D' = (1 - \delta)D, \\
&C + qB' \leqslant \Upsilon(Y_{\mathcal{S}}, n_{\mathcal{S}}) + B, \\
&B' \geqslant -\hat{\chi}D', \\
&C \geqslant 0, \\
V_A(B, D; \mathcal{S}) = \quad &\max_{C, B', D'} u(C, D') + \beta EV(B', D'; \mathcal{S}') && (6) \\
\text{s.t.} \quad &C + qB' + D' \leqslant \Upsilon(Y_{\mathcal{S}}, n_{\mathcal{S}}) + (1 - \delta - h)D + B, \\
&B' \geqslant -\hat{\chi}D', \\
&C, D' \geqslant 0, \\
V(B, D; \mathcal{S}) = \quad &\max\{V_{NA}(B, D; \mathcal{S}), V_A(B, D; \mathcal{S})\}. && (7)
\end{aligned}
$$

and, while not stated, also subject to the laws of motion for the state vector $\mathcal{S}$.

Given this formulation, a solution to the household problem is a collection of policy functions g_C^{NA}, g_B^{NA} and a value function V_{NA} that solve (5), policy functions g_C^A, g_B^A, g_D^A and a value function V_A that solve (6), and a value function V that, given V_{NA}, V_A, solves (7).

2.2 Calibration

We estimate the model processes and calibrate parameters to match moments from the Italian Survey of Household Income and Wealth (SHIW) for the years 1998-2014 as well as to a selected set of Italian macro statistics. All estimated continuous processes are discretized, with the details provided in the Appendix A. We use the SHIW for calibration as we in Section 4 will test the model predictions against reduced-form regression evidence from the same dataset.

Income process. Households receive their earnings potential, Y_{it}, if employed and a fraction of this earnings potential, bY_{it}, if unemployed:

$$
\Upsilon(Y_{it}, n_{it}) = Y_{it}(n_{it} + b(1 - n_{it})).
$$

The replacement rate is linear in the earnings potential Y_{it}, mainly for tractability. We set $b = 0.45$ to match the average replacement rate in Italy estimated by Martin (1996).

Earnings potential process. Following a large macroeconomic literature, we assume that the evolution of the earnings potential Y_{it} follows from the realizations of a permanent income shock η_{it} and a transitory income shock ϵ_{it}. Both are identically and independently distributed across time and households:

$$Y_{it} = Z_{it}e^{\epsilon_{it}}, \qquad\qquad \epsilon_{it} \sim N(0, \sigma_\epsilon), \qquad\qquad (8)$$

$$Z_{it} = Z_{it-1}e^{\eta_{it}}, \qquad\qquad \eta_{it} \sim N(0, \sigma_\eta). \qquad\qquad (9)$$

Although the earnings potential process features a unit root, the positive death probability ω guarantees the existence of a stationary earnings potential distribution.

Our specification of the earnings potential process provides a parsimonious and tractable approximation of the income processes faced by the households in the SHIW data. However, with our specified income process, we ignore two potentially important features of the income risks that the households are facing.

First, unemployment spells may affect earnings dynamics in the long run, implying the income losses associated with an unemployment shock are larger than the income loss during the unemployment spell (Jacobson et al., 1993), with a magnitude that may vary over the business cycle (Davis and von Wachter, 2011). In our model, the employment status and permanent earnings potential are independently distributed from each other, implying that we may underestimate the increase in income risk associated with entering a recession.

Second, we impose that the shocks to the log earnings potential are normally distributed. Guvenen et al. (2016) have shown in US administrative data that long-run changes in earnings are characterized by fat tails. To the extent that this is carries over to the Italian setting, the shocks would have to be drawn from a process with excess kurtosis to match this fact. However, to make our model comparable to the vast majority of the literature on consumption dynamics, we opt for a log normal earnings potential shock distribution.

We estimate the parameters of the income process using the SHIW sample used for the regressions in Section 4, with the additional restriction that the household head remains employed between any two consecutive waves. As described in Section 4, the data provide annual estimates at a biennial frequency. Denote $y_{4,it}$ as the log of the total household income over the last four quarters in quarter t. We retrieve the residuals $\Delta y_{4,it}^{res}$ from regressing the two-year log growth $y_{4,it} - y_{4,it-8}$ on sex, education and region interacted with a four-degree polynomial of age and year fixed effects. We assume that $\Delta y_{4,it}^{res}$ follows the log of the process described by (8)-(9) at the biennial frequency. The biennial model moments are then identified by

$$\sigma_{\epsilon,\text{biennial}}^2 = -\text{Cov}(\Delta y_{4,it}^{res}, \Delta y_{4,it-8}^{res}),$$

$$\sigma_{\eta,\text{biennial}}^2 = (\text{Var}(\Delta y_{4,it}^{res}) - 2\sigma_{\epsilon,\text{biennial}}^2).$$

After retrieving estimates of $\sigma^2_{\epsilon,\text{biennial}}, \sigma^2_{\eta,\text{biennial}}$, we rescale them to a quarterly frequency by setting

$$\sigma^2_{\epsilon} = \sigma^2_{\epsilon,\text{biennial}},$$
$$\sigma^2_{\eta} = \sigma^2_{\eta,\text{biennial}}/8.$$

The parameter values are reported in Table 1. We note that the parameter values are similar to several of the estimates provided in Krueger et al. (2010), which surveys the estimation of income processes across several countries.

Employment process. The process for employment status n_{it} is governed by a job-finding probability λ_t and a job-separation probability ζ_t. We are interested in estimating the demand response to a labor market recession and therefore we allow these probabilities to depend on the aggregate state, which we assume can take two values, expansion and recession, denoted with $\Theta \in \{E, R\}$:

$$\lambda_t = \lambda(\Theta_t) \tag{10}$$

$$\zeta_t = \zeta(\Theta_t). \tag{11}$$

Specifically, we aim for a process that captures the unemployment dynamics in Italy during the recent Eurozone crisis. In the Appendix C, we estimate the job-finding and job-separation rates for Italy 1998-2013 using OECD data and find that between the starting year of that recession, 2011, and 2013, the last year of observation, the quarterly job-separation rate increased from 1.01 to 2.08 percent (an relative increase of about 100 percent), while there was no significant movement in the job-finding rate. Accordingly, we impose that our process satisfies $\lambda(R)/\lambda(E) = 1$ and $\zeta(R)/\zeta(E) = 2$ and that the unconditional expectations $E(\lambda)$ and $E(\zeta)$ equal the average job-finding and job-separation probability in the period 1998-2013. The resulting parameter values are reported in Table 1.

We assume that the aggregate state Θ_t shifts between expansion and recession following a Markov process with transition matrix T:

$$T = \begin{pmatrix} p_{EE} & 1 - p_{EE} \\ 1 - p_{RR} & p_{RR} \end{pmatrix}.$$

Based on one observation alone, we cannot infer the ex-ante probability and the expected length of the Eurozone crisis. Therefore, we set p_{EE} and p_{RR} to match the average length of recessions (9 quarters) and the share of total time spent in recessions (23%) in Italy for the period 1948-2016. We use recession indicators constructed by the Economic Cycle Research Institute, which dates recessions for a wide range of countries using methods similar to those of the NBER Business Cycle Dating Committee.

Potential earnings process	Value	Target moments	Value	Data
σ_ϵ	0.158	$-\text{Cov}(\Delta y^{res}_{4,it}, \Delta y^{res}_{4,it-8})$	0.158	SHIW 1998-2014
σ_η	0.073	$\frac{1}{8}(\text{Var}(\Delta y^{res}_{4,it}) - 2\sigma^2_\epsilon)$	0.073	SHIW 1998-2014
Job-finding process	*Value*	*Target moments*	*Value*	*Data*
$\lambda(E)$	0.140	Mean find. rate 1998-2013	0.140	OECD
$\lambda(R)$	0.140	Find. rate(2013)/Find. rate(2011)	~ 1	OECD
Job-separation process	*Value*	*Target moments*	*Value*	*Data*
$\zeta(E)$	0.012	Mean sep. rate 1998-2013	0.015	OECD
$\zeta(R)$	0.025	Sep. rate(2013)/Sep. rate(2011)	~ 2	OECD

Table 1: *Calibrated parameter values for the income process.*

Preference and technology parameters. The utility function is CRRA with a constant elasticity of substitution between durable and nondurable goods:

$$U(C_{it}, D_{it}) = \frac{1}{1-\sigma}\left(\left[\alpha C_{it}^{\frac{\gamma-1}{\gamma}} + (1-\alpha)D_{it}^{\frac{\gamma-1}{\gamma}}\right]^{\frac{\gamma}{\gamma-1}}\right)^{1-\sigma}.$$

We make this assumption because of tractability. Together with an unemployment benefit that is linear in earnings potential Y_{it} and a credit constraint and an adjustment cost that are both linear in durable assets, the functional form of the utility function gives that the household policy functions are linear in the permanent earnings potential Z_{it}. The permanent earnings potential Z_{it} can thus be eliminated as a state variable in the computation of the household value and policy functions (see Appendix A).

We assume a unitary intratemporal elasticity of substitution between the two goods, i.e., $\gamma = 1$, which means that the preference specification is Cobb-Douglas.[6] There is scarce evidence on the intratemporal elasticity of substitution between nondurable and durable goods at the micro level, taking adjustment frictions associated with durable goods into account. Using time series data, Ogaki and Reinhart (1998) estimate that the intratemporal elasticity of substitution between durable and nondurable goods is approximately 1.

We set the relative risk aversion/inverse intertemporal elasticity of substitution, σ, to 2, in the middle of the standard range $[1, 5]$. The quarterly real interest rate, $1/q - 1$, is 1 percent and the mortality rate, ω, is set to generate an average working life of 50 years.

The rest of the parameters are calibrated to match moments from the SHIW data. For the discount factor β, we target the mean level of bond holdings in the model against the mean level of net financial assets, both normalized by total household income. For the credit constraint parameter $\hat{\chi}$, we target the share of households with negative net financial wealth.

---

[6]In this regard, we follow Barsky et al. (2007) and Berger and Vavra (2015).

Parameters	Value	Target	Empirical	Model
Discount rate, β	0.97	Mean norm. fin. assets	1.72	1.76
Nondurable weight, α	0.93	Mean norm. dur. assets	1.24	1.26
Depreciation rate, δ	0.038	Mean yearly purch. freq.	0.14	0.14
Adjustment cost, h	0.23	Mean norm. purchase size	1.20	1.22
Collateral constraint, $\hat{\chi}$	0.72	Share neg. fin. wealth	0.11	0.09

Table 2: *The calibrated preference and technology parameters of the model. The targeted moments are computed from the SHIW sample used for the regressions in Section 4. For moments that concern the stock and flow of durable goods, we use the stock and flow of motor vehicles in the data.*

As discussed in Section 2, D_t in the model has the interpretation of a household's stock of a typical durable good. Hence, we do not calibrate the model to match moments of households' total stock of durable goods. Instead, we calibrate against moments of households' total stock of motor vehicles. The Cobb-Douglas weight on nondurables α, the depreciation rate δ and the adjustment cost parameter h are set to match the mean level of wealth in motor vehicles normalized by quarterly income, the mean yearly purchase frequency of motor vehicles and the mean purchase size of motor vehicles normalized by quarterly income, respectively.

As shown in Table 2, the calibration yields moments that closely match the data. In the ergodic distribution of the calibrated model, about 2 percent of the households have a level of liquidity (the sum of the bonds and the collateral value of durable assets) of less than 10 percent of their permanent earnings potential, i.e., few households face the risk of a binding credit constraint. This is roughly consistent with the SHIW data, in which 4 percent of the households report that they have been denied credit or abstained from applying for credit in anticipation of being denied.

2.3 Computation

We solve for the household decision functions by value function iteration. In Appendix A, we describe the computational procedure in detail.

For the simulations, we develop a novel method which greatly reduces the number of households needed in each simulation to obtain accurate estimates of the aggregate dynamics. A "naive" simulation method is to simulate many households and sum up their individual consumption and savings. The naive simulation method is in practice costly, since the permanent-income distribution features a fat tail. One single observation with extremely high permanent income can shift the aggregates significantly, so the naive simulation method requires many independent households. We prove in Appendix B that there is an easy way to reduce the number of observations required by simulating the economy using distorted *permanent-income-neutral probabilities*.

In the simulation code, the only adjustment required is to replace the probability weights $\{p_i^\eta\}$ of the different values for the permanent-income shock η by the perturbed values $\{p_i^\eta \eta_i\}$. If there is a ten percent

probability of a doubling of permanent income, then the perturbed probability is twenty percent. The perturbed probability captures that although only ten percent of households receive the large positive shock, twenty percent of households weighted by ex-post permanent income received the large positive shock. In Appendix B, we show that by using the permanent-income neutral probabilities when simulating the model, we do not need to keep track of the fat-tailed permanent-income distribution which makes simulating the model computationally cheap.

3   Model Results

In this section, we study the demand response to the increase in the job-separation rate that Italy experienced during the recent Eurozone crisis. This shock affects households' beliefs about future income *ex ante* as well as the realized evolution of income *ex post*. Our aim is to assess the relative importance of these channels for durables and nondurables separately. Furthermore, the belief channel reflects changes in expected income as well as in income risk (the higher-order moments of future income). In a second step, we decompose the total belief response into these two components.

To investigate how the adjustment cost for durable goods affects aggregate consumption dynamics, we compare the simulation outcome of our model to that of a corresponding model without adjustment costs. The *flexible-adjustment model* is identical to the baseline model in all dimensions except that we set the adjustment cost h to 0 and recalibrate the preference and technology parameters to match the same moments as targeted in the baseline model. Without this recalibration, the wealth distribution in the flexible-adjustment model would differ significantly from the wealth distribution of the baseline model (and the data), which would obscure a comparison between the two. The recalibrated parameters are shown in Table 6 in Appendix D.

3.1   Impulse responses

We analyze the impulse-response functions (IRFs) to the recessionary job-separation rate shock. We populate the economy with a continuum of households, and let the recession shock hit the economy at t = 0. Prior to feeding the recession shock, we simulate the economy for a long time with the aggregate state drawn from its stochastic law of motion. After these burn-in periods, the economy enters into an expansion which lasts 30 quarters (the average length of an expansion). At t = 0, the economy enters into a recession that lasts 9 quarters (the average length of a recession) after which it reverts to the expansion regime for 30 quarters and then once more evolves stochastically. Since the consumption response depends on the stochastic history leading up to the experiment, we repeat the experiment several times and report the average response.[7]

---

[7]The simulation results converge quickly due to our method of using permanent-income neutral probabilities, outlined in Appendix B. We populate each simulation economy with 1000 households and repeat the experiment 320 times. The burn-in

The IRFs are shown in Figure 1. The recession periods are indicated by the shaded area. The curves are responses expressed as percentages of expenditures in the pre-recession quarter $t = -1$. The left column shows the response in the baseline model with the adjustment cost and the right column shows the response in the flexible-adjustment model. The top row shows the response of durable expenditures and the middle row shows the cumulative response of durable expenditures, which also approximates the change in the stock of durable holdings. The bottom row shows the response of nondurable expenditures.

Adjustment costs amplify the total demand response for durables. The black solid IRF is the total demand response of each aggregate variable. We start describing the flexible-adjustment model in the right column. The job-separation rate shock triggers a large drop in durable expenditures in the first recession period $t = 0$. Durable expenditures revert back in period $t = 1$ to a slightly lower level compared to the pre-recession period $t = -1$ and then spikes upward when the economy leaves the recession. Looking at the cumulative response, this pattern translates into an immediate fall in period $t = 0$ and then a gradual decline over the recession. This is the same pattern as we see for nondurables in the bottom row. Unless the credit constraint binds, our preference specification (preferences over a constant-returns-to-scale aggregator) implies that households will hold the ratio of durable to nondurable consumption constant in the flexible-adjustment model, and the cumulative response of durable goods expenditures to any shock is approximately a scaled version of the response of nondurables (scaled by the depreciation rate δ).[8] As discussed in the previous section, our calibration renders that few households face a binding credit constraint.

As depicted in the left column, our baseline model produces a close-to-identical response for nondurables. For durables, we see that compared to the flexible-adjustment model, the job-separation rate shock triggers a larger fall in durable expenditures in the first recession period $t = 0$, and this response is also more persistent. The cumulative response amounts to a fall of 110 percentage points at the end of the recession, compared to a fall of 85 percentage points in the flexible-adjustment model. With adjustment costs, the demand response to the recessionary job-separation rate shock is thus amplified by approximately 35 percent.

Adjustment costs amplify the belief channel and dampen the realized-shocks channel. To understand the mechanisms behind the aggregate consumption-expenditure responses, we decompose each response into two components, an ex-ante *belief channel* and an ex-post *realized-shocks channel*:

$$\text{Total Consumption Response} \approx \text{Belief Channel} + \text{Realized-Shocks Channel}.$$

period lasts for 200 quarters.

[8]The first order conditions give $\frac{u_{D,t}}{u_{C,t}} = \frac{r+\delta}{1+r}$ which determines the ratio C_t/D_t if $u(C_t, D_t) = U(v(C_t, D_t))$ with v being constant returns to scale.

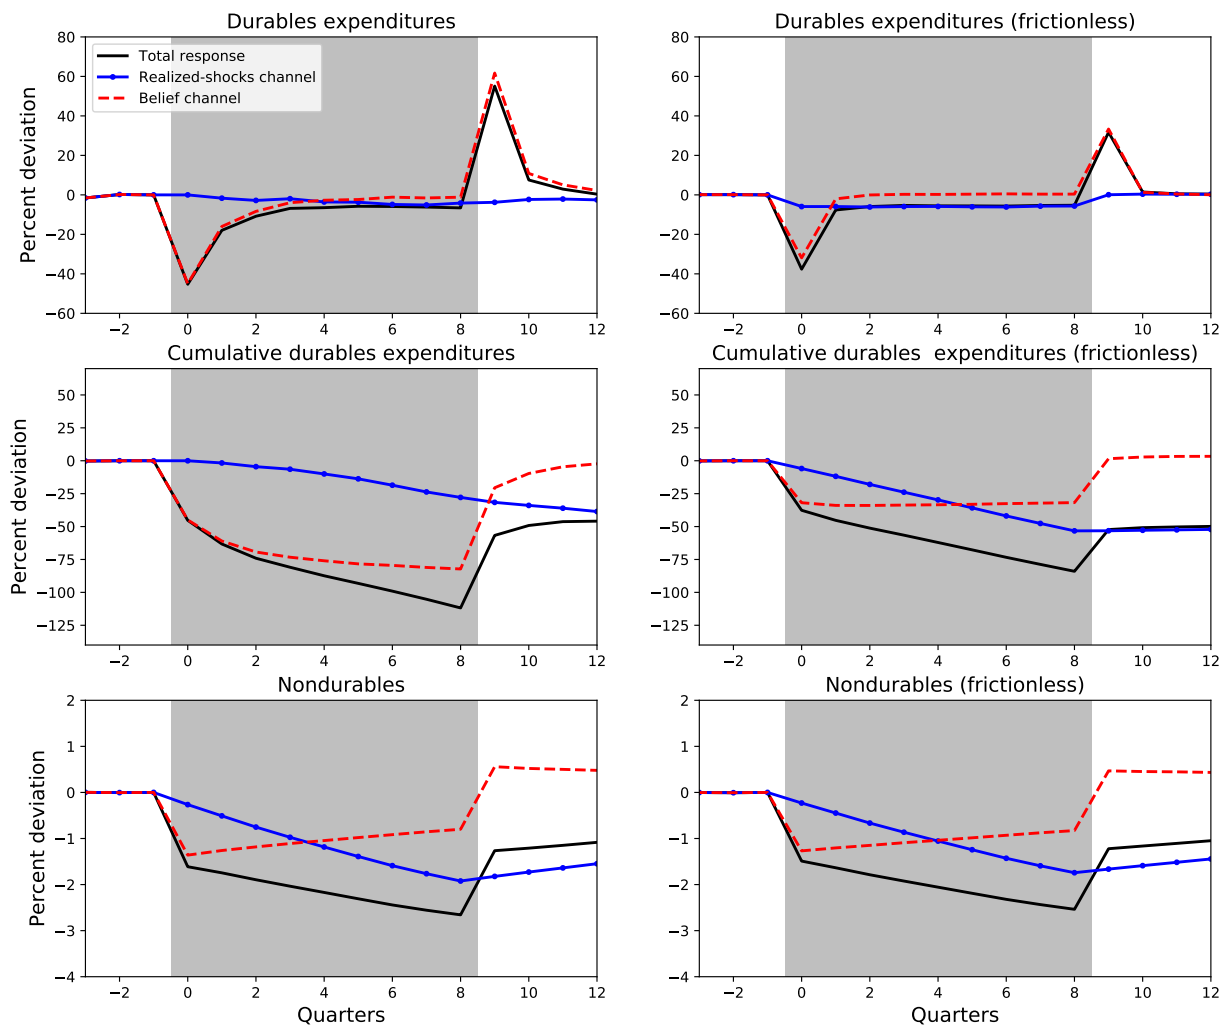***Figure 1:*** *Impulse-response functions to a recession shock. The recession periods are indicated by the shaded area. The left-hand column depicts the outcome in the baseline model, the right-hand column depicts the outcome in the flexible-adjustment model. The percentage deviation refers to the percentage deviation compared to the expenditure level in period −1, before the economy enters the recession.*

The realized-shocks channel captures the mechanical effect of more households becoming unemployed during the recession. It answers the question: "What is the consumption response if the economy is in a recession but households believe they are in an expansion?". We compute the realized-shocks channel by having households employ the expansion decision function while feeding them shocks drawn from the recession distribution. The belief channel captures the change in behavior in response to the job-separation rate shock. It answers the opposite question: "What is the consumption response if the economy is in an expansion but households believe they are in an recession?". We compute the belief channel by having the households employ the recession decision function while feeding them shocks drawn from the expansion distribution. There are also interaction effects between the two channels, but in practice the decomposition is near exact.

In Figure 1, the blue dotted curve is the realized-shocks channel and the red dashed curve is the belief channel. Starting again with the flexible-adjustment model in the right column, the shock triggers an immediate belief-channel response in durable expenditures in the first recession period $t = 0$, which directly reverts back in period $t = 1$. The realized-shocks response is persistent and near constant during the recession, as the higher job-separation rate means that more households become unemployed over time. Translated into the cumulative response, the belief channel is stronger in the short run, but is rather quickly (after four quarters) surpassed by the realized-shocks channel. The pattern is the same for nondurables.

The mechanisms are different in the baseline model. As seen in the top row, both the larger fall on impact and the higher persistence of the total response stems from the belief-channel response. In contrast, the realized-shocks channel response is smaller compared to the flexible-adjustment model in each period during the recession. This reflects that the response is smoother in the baseline model; we see that the expenditure fall through the realized-shocks channel also persists after the recession has ended. In terms of the cumulative response, the belief-channel response is approximately 2.5 times as large in the baseline model compared to the flexible-adjustment model at the end of the recession, and it accounts for 75 percent of the total response. The realized-shocks channel response is approximately half of the response in the flexible-adjustment model. The inclusion of the adjustment costs thus dampens the role of realized shocks and amplifies the role of beliefs in shaping the demand response of durable expenditures.

Adjustment costs amplify the belief channel through raising the response to income risk. In Figure 1, the ex-ante belief channel reflects the response both to the change in expected income (the first moment of the income process) and the change in perceived income risk (the higher-order moments). Here, we decompose the belief channel into these two components:

$$\text{Belief Channel} \approx \text{Expected-Income Channel} + \text{Income-Risk Channel}.$$

16

We compute the expected-income channel by first computing the net-present value of the fall in income associated with entering a recession. We then subtract this value, in a lump-sum fashion, from the households' assets and compute the consumption response, without feeding the recession shock and having households employ the expansion decision functions. We similarly compensate the households with a lump-sum transfer upon leaving the recession. We compute the risk channel by the following setup: the shocks are drawn from the expansionary distribution, households believe that they are in a recession, but as they enter the recession they are lump-sum compensated for the fall in the net-present value of their future income stream. By holding the net-present value of expected income constant, this approximates the consumption response through the risk channel.

The decomposition works well when households do not face a significant risk of the credit constraint binding. If the households face a significant risk of a binding credit constraint, the response to a lump-sum transfer given today will not be the same as the response to a transfer of the same net-present value given in the future. Close to a binding credit constraint, our decomposition overestimates the response of the expected-income channel, as the marginal propensity to consume is larger here, and underestimates the risk channel. However, as discussed in the previous section, our calibration renders that most households do not face a substantial risk of a binding credit constraint.

We show the decomposition of the belief channel in Figure 2. In the flexible-adjustment model, the risk channel plays only a minor role in accounting for the belief-channel response. Most of the response comes through the change in expected income. Looking at the response of cumulative durable expenditures and nondurable expenditures, we see that the shock triggers an immediate decline in consumption that persists over the whole recession, in line with consumption-smoothing motives. The response through the risk channel, stemming from the households' precautionary-savings motive, is small.

Turning to the baseline model in the left column, we see that the amplified belief-channel response stems from the amplified income-risk channel. The expected-income channel response is smoother in the baseline model, but at the end of the recession, the fall in cumulative durable expenditure is approximately the same as in the flexible-adjustment model. The risk channel, however, produces a 50 percentage point drop in cumulative durable expenditures in the baseline model, compared to a 7 percentage point drop in the flexible-adjustment model. The inclusion of the adjustment cost amplifies the response to the increase in income risk by approximately a factor of 7.

In sum, in our model, the demand response for durable goods to the job-separation rate shock is primarily driven by changes in household beliefs regarding their future income. The belief response stems primarily from the response to the increase in perceived income risk. This contrasts with the response of nondurables, for which the realized income shocks account for most of the demand response at the end of the recession
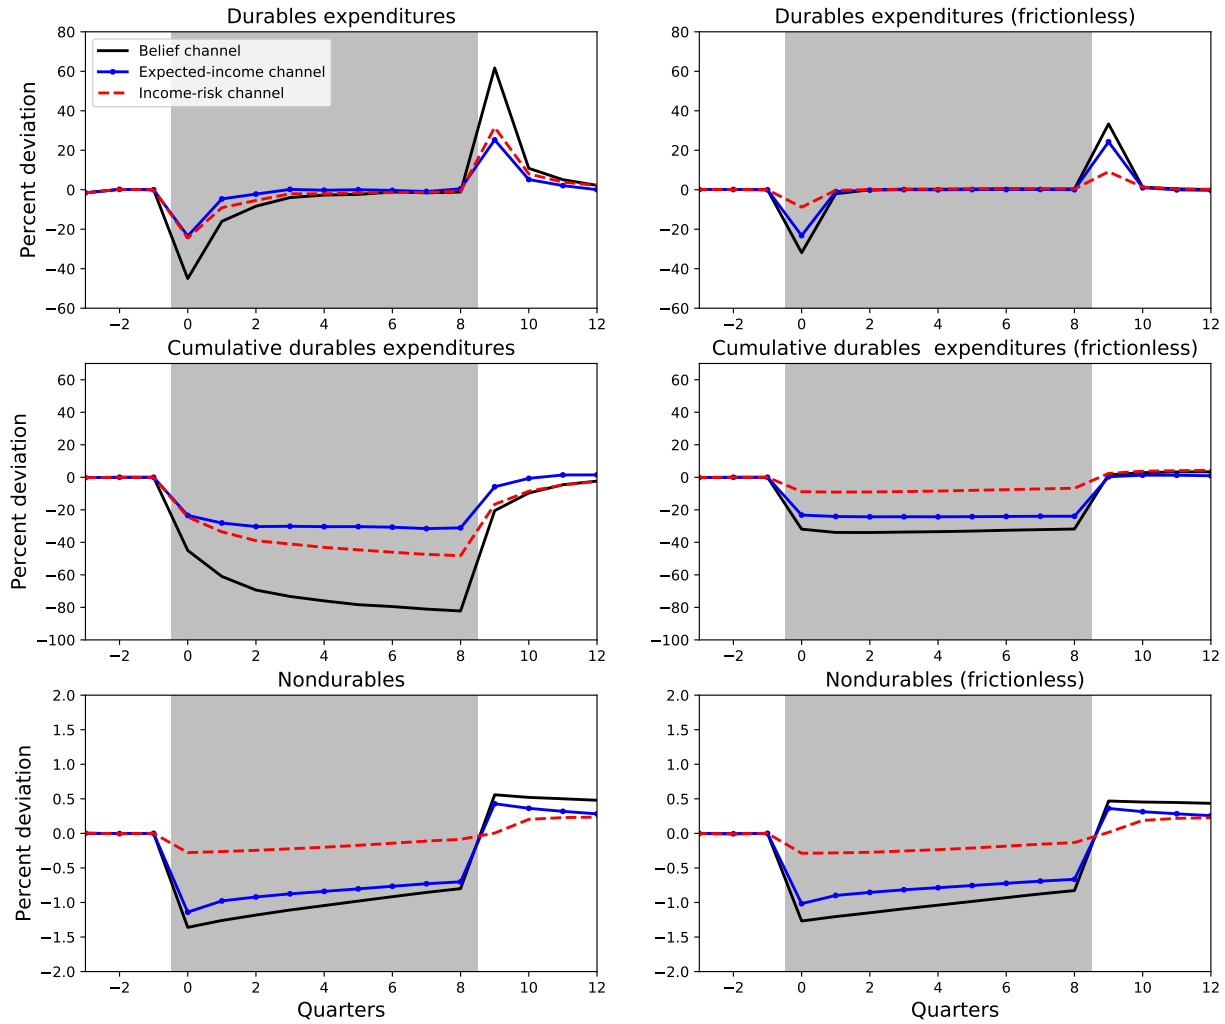
Figure 2: *Decomposition of the belief channel in the impulse-response functions to a recession shock. The recession periods are indicated by the shaded area. The left-hand column depicts the outcome in the baseline model, the right-hand column depicts the outcome in the flexible-adjustment model. The percentage deviation refers to the percentage deviation compared to the expenditure level in period −1, before the economy enters the recession.*

and for which income risk plays only a small role. The difference originates from the adjustment cost associated with durable purchases. In the next subsection, we investigate how the adjustment cost shapes the consumption response by scrutinizing the household decision functions.

3.2 Mechanisms

The two key features of our model are the adjustment cost for durable purchases and the inability of households to fully insure against income risk. In this subsection, we study the household decision functions and investigate how these features shape the consumption demand response. We begin with a general illustration and description of the household decision functions, and use these insights to explain how adjustment costs affect the realized-shocks channel, the expected-income channel and the income-risk channel displayed in Figures 1 and 2. We then vary the model discount factor to investigate how the asset holdings of the households shape the expenditure responses.

Household decision functions. We summarize household decisions in Figure 3. In the upper panels, we illustrate the decision functions of a household that is currently employed in an expansion and thus faces a low job-separation probability. The axes consitute the two dimensions of the state space: the vertical axis is the household stock of durables and the horizontal axis is the household cash on hand (the sum of income and financial assets), both normalized by the permanent earnings potential Z_t. The gray dots in the upper left panel display a scatter plot of the ergodic distribution of households in the state space. The shaded regions indicate whether or not the household will adjust its stock of durable goods. In the dark gray "inactivity region", the household chooses not to adjust, while in the light gray "activity region", it chooses to adjust. The dashed red curves constitute the adjustment thresholds. The arrows show where the households would move in the next period in the absence of any exogenous shocks. The arrows are not drawn to scale; the blue solid curve indicates the next period cash on hand and durable goods that the household will hold in expectation conditional on making an adjustment. In the lower left panel, we overlay the inactivity region and adjustment target of an unemployed household in the expansion state. In the lower right panel, we overlay the inactivity region and adjustment target of an employed household in the recession state.

As seen from the two upper panels, the household decision rule over adjusting its durable stock can be viewed as a two-dimensional (S,s)-decision rule. The adjustment cost means that households only adjust their durable stock if the value of adjusting today is greater than the value of postponing the adjustment into the future. The households adjust when they are sufficiently far from their target consumption level. If the households are to the left of the leftward adjustment threshold, they downsize their durable stock to increase their stock of financial assets and consume nondurable goods. If the households are below the bottom adjustment threshold, they instead upgrade their durable stock. Inside the inactivity region, households

19
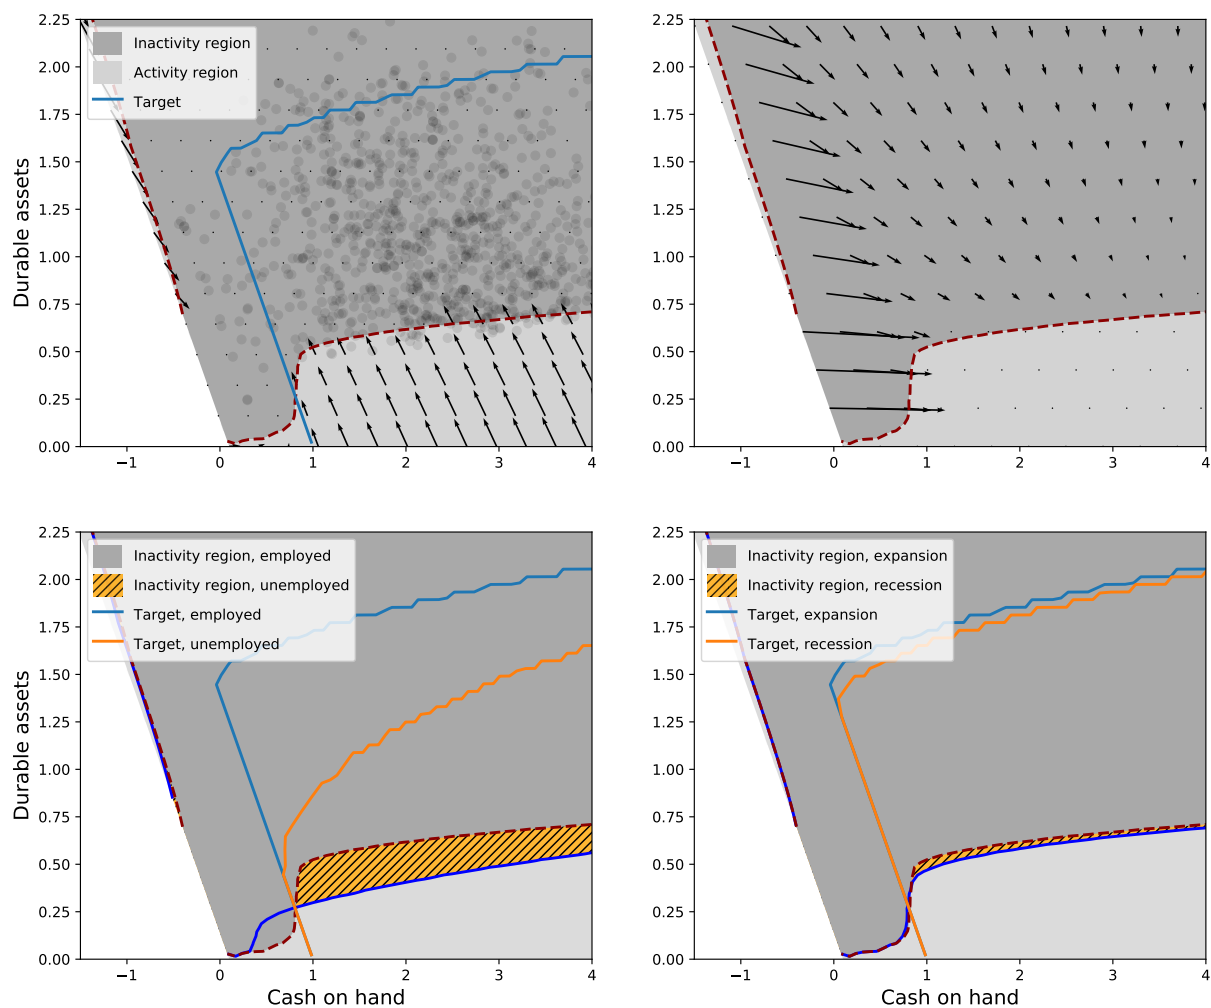
Figure 3: *Household decision functions. The axes are normalized by the permanent earnings potential. The upper-left panel illustrates the drift in the activity region for an employed household in the expansion state, alongside the adjustment target and a sample drawn from the stationary distribution of households. The upper-right panel illustrates the drift in the inactivity region for an employed household in the expansion state. The lower-left panel shows the difference in the inactivity region and the adjustment target between an employed and an unemployed household in the expansion state. The lower-right panel shows the difference in the inactivity region and the adjustment target between an employed household in the expansion state and an employed household in the recession state.*

pursue standard buffer-stock behavior while their durable stock depreciates over time, as indicated by the arrows in the upper-right panel. Households accumulate assets up to a certain point, their buffer-stock target, and deaccumulate assets if their wealth exceeds this target.

The leftward boundary of the state space has a negative slope since durable assets can be used as collateral, as seen in Equation (2). Around the intercept of the adjustment target and the inactivity-region threshold, the adjustment threshold sharply shifts up. This shift is due to the credit constraint: households to the left of the shift are constrained conditional on a durable purchase while the households to the right are unconstrained. When the households have very little liquid assets, it is preferable to save up more liquid assets in order to make a purchase closer to the unconstrained optimal size.

Adjustment costs and income shocks. The realized-shocks and the expected-income channel responses were smoother in the baseline model compared to the flexible-adjustment model in Figures 1 and 2. This difference can be understood in terms of the decision functions displayed in Figure 3.

Starting with the expected-income channel, households in the flexible-adjustment model that are sufficiently far from the binding credit constraint respond by immediately cutting consumption, both nondurables and durables, in accordance with permanent-income theory. In the baseline model with the adjustment cost, the same underlying force is present. Since the expected-income channel was computed by means of a lump-sum transfer, in Figure 3 the expected-income shock constitutes a leftward shift of the entire distribution.

At lower levels of cash on hand, the adjustment target and the adjustment threshold are both lower, leading to a fall in durable expenditures. The adjustment occurs when the households cross the adjustment threshold into the activity region. However, as seen in the upper-left panel of Figure 3, most households are far from the adjustment threshold and will thus not react immediately. The response is sluggish, as depreciation gradually brings all households down to the adjustment threshold.[9]

The realized-shocks channel response in the baseline model in Figure 1 is even more sluggish. First, households become unemployed gradually in the high separation-rate recession. Second, the expenditure response of the newly unemployed takes time. From the lower-left panel of Figure 3, we see that the decision functions react strongly to a realized separation shock. The inactivity region expands significantly and the shift in the adjustment target is large except for households with very low levels of cash on hand. However, only a fraction of the unemployed households are close to the adjustment threshold, where they will react immediately.

---

[9]That aggregate durable expenditure respond more sluggishly to aggregate income shocks when subject to non-convex adjustment frictions has been formalized in the context of a complete-markets model by Chetty and Szeidl (2016): If idiosyncratic risk is sufficiently high then the effect of aggregate shocks on aggregate consumption is *as if* the aggregate consumption decisions are made by a representative consumer with habit preferences in consumption. Our experiment shows that the insight carry over to our incomplete-markets environment.

Adjustment costs and the risk channel.   As shown in Figure 2, the adjustment cost greatly amplified the durable expenditure response to the increased in income risk associated with the job-separation rate shock. In the flexible-adjustment model, consumption expenditures respond to risk for precautionary-savings reasons, stemming from potentially binding borrowing constraints and/or prudent preferences (Kimball, 1990, Carroll and Kimball, 2001). This mechanism is present but weak in our calibration, as indicated by the small relative size of the risk channel in the nondurable expenditure response in the baseline model, as well as by the small relative size of the risk channel in the durable expenditure response in the flexible-adjustment model, seen in Figure 2.

The risk-channel response for durable purchases is larger in the baseline model because, with the adjustment cost, the households have an additional reason to cut expenditures in response to increased income risk — a "wait-and-see" motive (Dixit and Pindyck, 1994). The adjustment cost implies that households will, barring an extreme sequence of shocks, not downsize its durable stock. In consequence, a purchase implies commitment. Commitment becomes more costly when unemployment risk (and therefore income risk) increases, since the target level of consumption when unemployed differs from the target level when employed, as illustrated by the adjustment targets in the lower-left panel of Figure 3. The increased value of postponing the durable adjustment is reflected in the outward shift of the adjustment threshold in the lower-right panel of Figure 3. This outward shift also reflects the response to the change to expected income when unemployment risk increases, but as shown in Figure 2, the risk-channel response is the stronger effect. The outward shift of the adjustment threshold may look small in relation to the shift in response to a realized separation shock in the lower-left panel, but since the higher separation risk affects all households when entering the recession, it amounts to the large risk-channel response in Figure 2.

Moving from right to left in the cash-on-hand dimension in the lower-right panel of Figure 3, the outward shift of the adjustment threshold grows. This is because the distance between the target durable stocks of the employed and unemployed grows as cash on hand shrinks, as shown in the lower-left panel. At some point, however, the household's asset position is so small that an employed household cannot buy its unconstrained optimal size of durables even if becoming unemployed. Here, conditional on adjustment, the household will use all its available resources, regardless of its employment status. Since the purchase plan is not state contingent in this region, the adjustment threshold does not respond to the increase in unemployment risk.

Expenditure responses when varying households' liquid assets.   The analysis of the decision functions showed that the household-level response to unemployment risk varied with the asset position. This raises the question: how do the aggregate durable and nondurable expenditure responses to the job-separation rate shock depend on the asset holdings in the population? To investigate this, we vary the discount factor β and compute the response to the job-separation rate shock. By varying β, we generate variation in households' mean asset holdings while maintaining the other calibration moments close to their targets (see

Tables 7 and 8 in Appendix E).

In Figure 4, we show the aggregate responses, measured at the end of the recession, for 24 model economies with discount factors in the range of $[0.920, 0.973]$, where the baseline calibration is indicated by the vertical line. The left column displays the response of cumulative durable expenditures and nondurable expenditures for the baseline model, the right column displays the corresponding responses for the flexible-adjustment model. On the x axis, we have the mean level of *normalized liquidity*: the sum of households' financial assets and collateral value of their durable assets, where the normalization is with respect to permanent earnings potential. When this sum takes the value 0, the credit constraint is binding for all households. We display the total expenditure response together with the decomposition into the belief-channel and the realized-shocks channel responses. We do not display the decomposition of the belief-channel response into its risk-channel and expected-income channel responses, as this decomposition works less well when more households come close to a binding credit constraint.

Starting with the response of cumulative durable expenditures for the flexible-adjustment model in the upper-right corner, the realized-shocks channel response is larger than the belief-channel response across the entire spectrum. The response through the realized-shocks channel is monotonically decreasing in asset holdings, while the response of the belief channel is U-shaped.

The realized-shocks channel response is monotonically decreasing in the mean asset holdings because households' marginal propensity to consume increases when approaching a binding credit constraint. To understand the U-shape in the belief-channel response, we begin by looking at the endpoints of the curve. For large asset holdings, the households are well insured and respond only to the expected-income component of the job-separation rate shock. At the other end, for small mean asset holdings most households are effectively 'hand-to-mouth' and do not respond to changed beliefs about the future. In between these two extremes, most households are forward looking but imperfectly insured, giving rise to a precautionary-savings response. The U-shape in the belief-channel response notwithstanding, the total response is primarily determined by the realized-shocks channel response. Therefore, quantitatively, the realized income shocks interacted with the degree of 'hand-to-mouth' behavior in the economy largely determines the aggregate response. This is consistent with many papers that have emphasized that the degree of hand-to-mouth behavior is a key determinant of the nondurable consumption response to aggregate shocks, see, e.g., Berger et al. (2018a), Kaplan et al. (2018) and Auclert et al. (2018).

As before, the response of nondurables in the flexible-adjustment model and the response of nondurables in the baseline model behave analogously to cumulative durable expenditures in the flexible-adjustment model. The pattern is different for cumulative durable expenditures in the baseline model. The realized-shocks channel response is still monotone in mean asset holdings and the belief-channel response is still U-shaped. However, in contrast to nondurable expenditures, the belief-channel response is larger than the
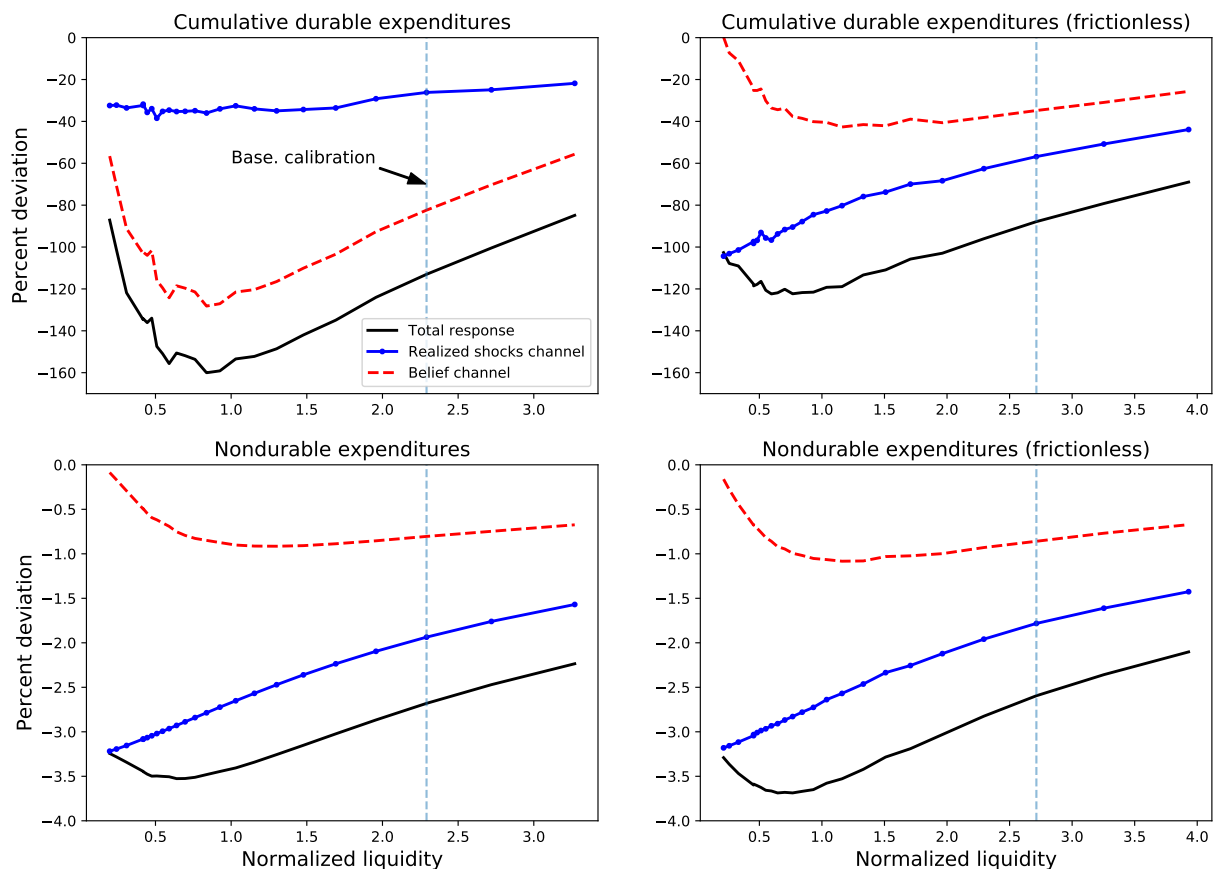
Figure 4: *Demand responses as functions of households' mean asset holdings. Each data point is the demand response in the final period of the recession for an economy populated with households with a given discount factor. The discount factor varies β from 0.920 to 0.973. On the horizontal axis, we show households mean normalized liquidity, defined as households' average sum of financial assets and collateralizable durable assets. On the vertical axis, we show the demand response in the final recession period as a percent deviation from the expenditure rate in the period prior to entering the recession.*

realized-shocks channel response across the entire spectrum.

The relative importance of the belief channel for durable expenditures implies that the total response of durable expenditures varies differently when changing the households' asset holdings. Around the baseline calibration, the belief-channel response for cumulative durable expenditures is more sensitive to changes in the mean liquid assets than the belief-channel response for nondurable expenditures. This belief-channel sensitivity does not translate into a larger sensitivity of the total response, however, since the realized-shocks channel response is less sensitive for durables compared to nondurables around this point. Close to the binding credit constraint, the U-shape of the risk-channel response generates a large U-shape for the total response of cumulative durable expenditures. The credit constraint short circuits forward-looking considerations and it therefore mutes the belief-channel response. For durable expenditures, in contrast with nondurable expenditures, the belief channel is the quantitatively most important channel. Therefore, high levels of 'hand-to-mouth' behavior significantly dampens the total response of durable expenditures to the job-separation rate shock.

4   Model Meets Data

In this section, we document expenditure responses, for both nondurable and durable goods, to fluctuations in the probability of becoming unemployed using the Italian Household Survey of Income and Wealth (SHIW). Thereafter, we investigate whether our calibrated consumption model can explain the empirical estimates.

In the empirical exercise, we will not be able to separate the expenditure response to the first moment and the higher moments of the labor market shock, as we did in using our theoretical model in the previous section. Our empirical estimates are instead a mix of both. However, from these estimates, we will still be able to partially infer the importance of real-options effect in accounting for the expenditure response. We can do this because we can estimate the expenditure response for durable and nondurable goods separately, and we know from the model that the amplification through the real-options effect is only active for durable goods.

4.1   Data description

The SHIW is administered by the Bank of Italy and surveys a representative sample of Italian households.[10] The survey is conducted biennally with a sample of 8000 households, half of which are part of a rotating panel with gradual replacement. We use the panel part of the survey for years 1998-2014.[11] To be part of our sample in year t, we require that 1) households are surveyed in three consecutive waves centered around

---

[10]The SHIW data have previously been used for documenting expenditure patterns in, e.g., Bertola et al. (2005), Jappelli and Pistaferri (2014) and Auclert (2018).

[11]The survey goes back earlier, but for these years, the definitions of the variables that we use are not consistent with the later waves.

year t, 2) the household head is aged 25-60, 3) the household head is employed in year t, 4) the household head is part of the labor force (employed or unemployed) in year $t + 2$ and 5) the income of the household head is at least 40 percent of total household income in year t. We make the last restriction to ensure that an unemployment spell is a meaningful income loss for the household. In Appendix G, we check whether our results are robust to changing this cutoff value. With these restrictions, our final sample consist of $4,574$ observations.

The variables that we use consist of information on households purchases of nondurable and durable goods (motor vehicles and furniture), disposable income, financial and real wealth and also basic household and job characteristics. Information regarding state variables, such as wealth, age, household size etc., refers to that state at the end of the year. Information regarding flow variables, such as consumption expenditures and income, refers to the sum of the flow over the past year. The employment status of a household head in a year is defined as the employment status held for most part of that year. In Appendix F, we provide descriptive statistics and a detailed description of all variables used and more information on survey design and data quality.

4.2 Econometric framework

We seek to identify the expenditure response, both in terms of durable and nondurable goods, to fluctuations in the probability of becoming unemployed. Since durable purchases are lumpy, we estimate a probit model relating the probability of a durable purchase to the growth of unemployment risk. Let DP_{it} be an indicator variable that takes the value of 1 if a durable good has been purchased in year t, denote the probability of being unemployed in year $t + 2$ by $Uprob_{it}$ and define the operator $\Delta X_t = X_t - X_{t-2}$ for any variable X (recall that the SHIW surveys households at a biennial frequency). We estimate

$$P(DP_{it} = 1) = \Phi(\beta \Delta Uprob_{it} + \beta_x X_{it}), \tag{12}$$

where $\Phi()$ is the cumulative normal distribution, and $\mathbf{X}_{it}$ is a collection of relevant covariates.

For nondurables, expenditures are always positive, and we estimate the effect of unemployment probability growth on log expenditure growth Δc_t using a standard OLS specification:

$$\Delta c_t = \alpha_0 + \alpha \Delta Uprob_{it} + \alpha_x X_{it} + \epsilon_t. \tag{13}$$

The collection of covariates $\mathbf{X}_{it}$ are included to identify the effect of fluctuations in the unemployment probability on expenditures. Before we present and discuss these covariates, we describe the measurement of $Uprob_{it}$.

Measuring the likelihood of unemployment To retrieve a measure of the unemployment probability, we condition the probability of being unemployed in year $t + 2$ on characteristics observed at time t. Specifically, our strategy is to estimate a probit model,

$$P(U_{it+2} = 1) = \Phi(\beta_t \mathbf{Z}_{it}), \tag{14}$$

where U_{it+2} is an indicator that takes the value of 1 if the household head is unemployed in year $t + 2$ and 0 otherwise, and $\mathbf{Z}_{it}$ is a vector of covariates. We estimate this model separately for each year in the sample.

The covariates in $\mathbf{Z}_{it}$ are chosen to partition the labor market into fairly narrow categories, which are subject to idiosyncratic shocks and/or common shocks with category-specific treatment intensity. The variables in $\mathbf{Z}_{it}$ are occupation, industry of employment, superregion of residence (north, center and south/islands), region of residence (the 20 administrative regions of Italy), bins of population size of the town area in which the household resides, education level, five-year age bins, sex, marital status and household size. The premise for selecting these variables is that basic job characteristics, such as industry or region, are plausible sources of idiosyncratic demand shocks. In addition, basic household characteristics, such as age and household size, are plausible determinants of selection into unemployment.

We do not directly regress the employment status on $\mathbf{Z}_{it}$. Instead, we retrieve factors $\mathbf{F}_t$ from the covariates $\mathbf{Z}_{it}$ and use the associated factor projections $\mathbf{f}_{it}$ for the estimation of the unemployment probability:

$$P(U_{it+2} = 1) = \Phi(\beta_t \mathbf{f}_{it}). \tag{15}$$

The motivation for this approach is that when estimating (15), the yearly samples contain 1,150.5 observations on average, while the total number of categories formed by the set $\mathbf{Z}_{it}$ is 88.[12] If we directly regress employment status on $\mathbf{Z}_{it}$, we will likely overfit the data. By regressing on the factor projections, we reduce the dimensionality while simultaneously keeping as much of the variation in the underlying independent variables as possible.

We retrieve the factors by multiple correspondence analysis, which is analogous to principal component analysis but adapted to suit categorical variables (Greenacre, 2007). In the choice of the number of factors used, there is a tension between how much of the variation in the underlying covariates the factors explain and explanatory power when using the factors for estimating (15). As a baseline, we choose the number of factors that maximizes the mean of McFadden's adjusted R^2 from the regressions, which gives us 6 factors. We check whether our estimates are robust to using up to 12 factors in Appendix G.

Having estimated individual-specific unemployment probability for all years in the sample, we recover

---

[12]Since we run this regression for each of the 8 biennial sample years between 1998 and 2012, the total sample size is 9,204 observations. This is larger than the sample of 4,574 observations used to estimate (12) and (13), since here, we only need to impose that households are part of two but not three consecutive waves of the survey.

the growth of unemployment probability by taking the first difference:

$$\Delta Uprob_{it} = P(U_{it+2} = 1|f_{it}) - P(U_{it} = 1|f_{it-2}). \tag{16}$$

Identification. To what extent is the variation in $\Delta Uprob_{it}$ exogenous in the estimation of (12) and (13)? The variation in $\Delta Uprob_{it}$ stems from two sources. First, the same set of covariates can predict a different unemployment rate in year t than in year $t-2$, creating variation in $\Delta Uprob_{it}$ while holding the household characteristics constant. For example, if there is a regional-industry specific adverse demand shock in year t that disproportionately raises the unemployment rate for young households heads in that region-industry in year $t+2$, young households belonging to that region and industry will experience a higher unemployment probability in year t than they did in year $t-2$. Second, household characteristics may change between any two waves in the sample. For example, which age bin the household head belongs to will mechanically change for some households in the sample. We only seek to capture the first type of variation, as changes in household characteristics may very well affect expenditure decisions through other channels than by changing the probability of becoming unemployed. Therefore, the control set $\mathbf{X}_{it}$ includes the same job and household characteristics variables used to predict $Uprob_{it}$.[13] In explaining our results below, we refer to these control variables as *Characteristics variables*.

Moreover, at the macroeconomic level, if a significant share of $\Delta Uprob_{it}$ arises from aggregate business cycle shocks, $\Delta Uprob_{it}$ may be correlated with changes in prices, in particular the interest rate and the relative price of durables, which have a separate effect on consumption decisions. Therefore, our set of control variables $\mathbf{X}_{it}$ includes time fixed effects. Local variation in the relative price for durables is more difficult to control for as time-region fixed effects would absorb a significant amount of the variation in $\Delta Uprob_{it}$. However, an adverse shock that disproportionately reduces the demand for durables should lead to a decline in their relative price. Our estimates for the effect of unemployment probability growth on the expenditure on durables (nondurables) should in that case be interpreted as lower (upper) bound of the true partial equilibrium effect.

Finally, changes in the probability of becoming unemployed are likely correlated with other changes in the households' financial position. If parts of those households with heightened probability of becoming unemployed also experience slower wage growth, unemployment probability growth is correlated with income growth. We therefore control for log income growth $\Delta \log y_{it}$. As we document below, changes to the probability of becoming unemployed also has an autoregressive component, meaning that current changes are likely correlated with consumption and savings decisions in the past, which affect the current period state variables. We therefore control for quartiles of net total assets and net financial assets in the previous period, TA_{t-2} and FA_{t-2}, as well as durable goods assets in the current period net of any purchases or

---

[13]The regions and superregions are perfectly collinear, however, and therefore we only include the regions in $\mathbf{X}_{it}$.

sells in this period, D_t, all of which are measued as a fraction of previous period income y_{t-2}. Finally, an increase in $\Delta Uprob_{it}$ could affect the consumption decision through a tightened credit constraint, as financial institutions are likely to be less inclined to lend to a household with worse employment prospects. To control for this indirect feedback mechanism, $\mathbf{X}_{it}$ also includes a dummy variable that indicates whether the household has faced a binding credit constraint. The dummy variable takes value 1 if the household has been denied or discouraged from applying for credit in year t.[14] In explaining our results below, we refer to all these control variables as *Financial variables*.

4.3 Results

Table 3 shows the results from estimating Equations (12) and (13). In columns (1)-(3), we show the average marginal effect on the log of nondurable expenditure growth that results from estimating Equation (13). Our preferred specification is column (3), as the inclusion of all the control variables makes it more likely that the variation in $\Delta Uprob_{it}$ is exogenous to consumption decisions in period t. The estimation is in OLS and the numbers are straightforwardly interpreted. A 1 percent increase in unemployment probability for all households in the sample is associated with a fall of nondurable expenditures by 0.3 percent, but we cannot distinguish the effect from zero.

In columns (4)-(6) of Table 3, we show the estimated effect on the probability of having purchased a durable good that results from estimating equation (12). The right-hand side variable refers to purchases of goods either from the furniture or the motor vehicles category or both. The coefficients show the average marginal effect, normalized by the unconditional purchase probability, and should thus be interpreted as the aggregate marginal effect on the number of purchases if all households in the sample are subject to a 1 percent increase in unemployment probability growth, i.e., they show the semi-elasticity of durable purchases to unemployment probability growth. For our preferred specification in column (6), a 1 percentage point increase in unemployment probability growth for all households is associated with an estimated fall of aggregate durable purchases by 2.4 percent.

In columns (7)-(12) of Table 3, we show estimates when replacing the left-hand side variable with the probability of purchasing items in each of the two durable goods categories separately. The point estimates for both durable good categories are close to the estimate for the combined measure, although somewhat larger for motor vehicles and somewhat smaller for furniture. Unsurprisingly, the precision of the estimates are lower compared to the regression for total durables, especially so for motor vehicles, as the number of households that purchase items from each category seperately is significantly lower than the number of households that purchase items from any of the two categories; see the descriptive statistics in Table 10 in Appendix F.

---

[14]We define a household to be discouraged when the interviewee has answered that someone in the household has thought about applying for credit, but later changed his/her mind in anticipation that the application would be denied.

	(1) Nondurables	(2) Nondurables	(3) Nondurables	(4) Durables	(5) Durables	(6) Durables
ΔUprob_t	0.56*	-0.15	-0.30	-2.15***	-2.19***	-2.42***
	(0.31)	(0.27)	(0.26)	(0.83)	(0.83)	(0.82)
Char. variables	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	No	Yes	Yes	No	Yes	Yes
Financial variables	No	No	Yes	No	No	Yes
$\bar{R}^2$	0.02	0.28	0.33	0.06	0.06	0.07
N	4574	4574	4574	4574	4574	4574

	(7) Furniture	(8) Furniture	(9) Furniture	(10) Vehicles	(11) Vehicles	(12) Vehicles
ΔUprob_t	-1.93*	-2.23**	-2.36**	-3.21*	-2.50	-2.83
	(1.02)	(1.03)	(1.02)	(1.83)	(1.85)	(1.84)
Char. variables	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	No	Yes	Yes	No	Yes	Yes
Financial variables	No	No	Yes	No	No	Yes
$\bar{R}^2$	0.06	0.06	0.07	0.05	0.06	0.09
N	4574	4574	4574	4574	4574	4574

Table 3: *Regression results from estimating* (12) *and* (13). *For columns (4)-(12), the left-hand side variable is an indicator of whether a durable good has been purchased in the last 12 months. The coefficients show the average marginal effect normalized by the unconditional purchase probability, estimated with a probit model. In columns (1)-(3), the left-hand side variable is the log growth of expenditures on nondurable goods, estimated by OLS. For explanation of the control variables, we refer to the text in Section 4.2. *, **, *** indicate that the coefficients are significant at the 10%, 5% and 1% level, respectively.*

The effect of higher unemployment probability is substantially lower for nondurable goods than for durables. In Appendix G, we confirm that this also holds for a number of robustness checks. We take it as evidence that households respond to an increase in the unemployment probability primarily by cutting down on durable rather than nondurable expenditures and now turn to whether this prediction is quantitatively matched by the calibrated consumption-savings model.

4.4   Regressions on model-generated data

To run the same cross-sectional regressions on model-generated data as on the SHIW data, we augment the model with idiosyncratic fluctuations in unemployment probability growth. In the baseline model in Section 2 there is rich heterogeneity in the employment histories of the households, but since all household face the same job-separation and job-finding probability at each point in time, there is no cross-sectional variation in households' probability of becoming unemployed in the future. Moreover, having aggregate shocks to the job-separation probability in the model is redudant here, as this variation would be absorbed in the time fixed effects. For the purpose of running regression (12) and (13) on model-generated data, we therefore

Employment process	Value	Target moment	Value	Data
σ_ζ	0.22	$\mathrm{Sd}(\Delta\mathrm{Uprob}_{it})$	0.023	SHIW 1998-2014
ρ_ζ	0.91	$\mathrm{Corr}(\Delta\mathrm{Uprob}_{it}, \Delta\mathrm{Uprob}_{it+8})$	-0.28	SHIW 1998-2014

Table 4: *Calibrated parameter values for the job-separation probability process.*

replace the job-finding and job-separation probability processes in Equations (10) and (11) with

$$\lambda_{it} = \lambda, \tag{17}$$

$$\log \zeta_{it} = \log \bar{\zeta} + \rho_\zeta \log \zeta_{it-1} + \sigma_\zeta \epsilon_{it}^\zeta, \qquad\qquad \epsilon_{it}^\zeta \sim N(0,1). \tag{18}$$

That is, we assume a constant job-finding probability λ and a time-varying job-separation probability ζ_{it}, which follows an $AR(1)$ in logs.[15]

We make the admittedly strong assumption of holding the job-finding probability constant as the SHIW data is not sufficiently rich to identify the time variations in job-separation and job-finding probabilities separately. By holding the cross-sectional job-finding probability constant, we keep the fluctuations in labor market risk that household face at the micro level consistent with the aggregate risk that households face in the business cycle experiment in Section 3.

As before, we calibrate these process so that the average job-finding and job-separation probabilities, $\bar{\zeta}$ and λ, matches the average job-separation probabilities and the average job-finding probabilities for Italy in the period 1998-2013, using aggregate labor market statistics from the OECD. We calibrate ρ_ζ and σ_ζ to match the variation in $\Delta\mathrm{Uprob}_{it}$ that we have estimated in the SHIW data. More specifically, we match the standard deviation and the autocorrelation of the change in the unemployment probability growth, $\Delta\mathrm{Uprob}_{it}$ with their model counterparts. In the data, Uprob_{it} corresponds to the conditional probability of being unemployed for "most part" of the year ranging from quarter $t+5$ to quarter $t+8$. In the model, we define Uprob_{it} as the probability of being unemployed for at least two quarters in any of the quarters between $t+5$ and $t+8$. The calibrated parameter values and the target moments are reported in Table 4.

Constructing a cross-sectional sample We simulate the model to produce a panel sample of approximately the same size as the data sample and run the same regressions on this model-generated data. We repeat the exercise 1,000 times and report the average coefficients from running these regressions. The exact number of observations varies for each regression, as the number of realized unemployment spells is stochastic and we exclude the households that do become unemployed (see below). A period in the model is a quarter while the panel from the SHIW is biennial. To conform with the SHIW, we construct a biennial panel sample

---

[15]Since ϵ_{it}^ζ is normal, ζ_{it} is log-normal with support $(0, \infty)$. This violates the economic restriction that $\zeta_{it} \in [0,1]$. However, this theoretical inconsistency has no practical implications as ζ_{it} never exceeds unity in the estimated and discretized process.

	Data			**Simulated model**	
	(1)	(2)	(3)	(4)	(5)
	Durables	Vehicles	Nondurables	Durables	Nondurables
ΔUprob$_t$	-2.42***	-2.83	-0.30	-2.47**	-0.45***
	(0.82)	(1.84)	(0.26)	(1.11)	(0.06)
R^2	0.07	0.09	0.33	0.56	0.75
N	4574	4574	4574	$\sim$ 4550	$\sim$ 4550

Table 5: *Regression results from estimating* (12) *and* (13) *on actual data and model-generated data. For columns (1), (2) and (4), the left-hand side variable is an indicator of whether a durable good has been purchased in the last 12 months. The coefficients show the average marginal effect normalized by the unconditional purchase probability, estimated with a probit model. For the columns (3) and (5), the left-hand side variable is the log growth of expenditures on nondurable goods, estimated by OLS. For the estimates on actual data, *, **, *** indicate that the coefficients are significant at the 10%, 5% and 1% level, respectively.*

from the model simulations, definining variables in the same way they are defined in the SHIW data. From now on, the time index t refers to a given year.

We use the same regression specifications on the model-generated data as those used on actual data, that is (12) and (13). We exclude households that are unemployed when running the regression, as we did in the empirical exercise. Hence, a higher level of unemployment probability growth corresponds to what we referred to as the belief-channel response in Section 3, reflecting both lower expected income and higher income risk, but no direct income loss. The set of control variables includes log income growth, quartiles of net previous period net financial assets and current-period durable goods assets period prior to any purchase or sell, both of which are normalized by previous period income.

4.5   Model regression results

In Table 5, we show the results from estimating the response of expenditures to fluctuations in the unemployment probability together with the empirical regression results from our preferred spefications in column (3), (6) and (12) of Table 3. For the durable goods specifications in column (1), (2) and (4) of Table 5, we display the estimated average marginal effects, normalized by the unconditional purchase probability. In columns (3) and (5), we show the average marginal effect on the log of expenditure growth.

In the simulated model, the number of durable good purchases falls by 2.47 percent in response to a 1 percent increase in unemployment probability growth. This is very close to the corresponding empirical estimates of -2.42 for total durables and -2.83 for vehicle purchases. The parameters of the model durable goods were calibrated to match purchases of vehicles in the data. The standard errors for the model estimates are lower than the standard errors for vehicle purchases in the data, but still on the same order of magnitude. For nondurable goods, the model estimate of the response to unemployment probability growth is -0.45, somewhat larger than but still close to the empirical point estimate of -0.30.

What do we learn from comparing these numbers? The estimates from the simulated model are close to the empirical estimates, both for nondurable and durable goods. Since our model is calibrated to match unconditional moments in the data, there are no mechanical reasons why the model regression results, which capture correlations between individual-level growth rates of the relevant variables, should conform to the empirical regression results. For nondurables, we know that the expenditure growth respond to increase unemployment probability both through the expected income channel and due the precautionary-savings mechanism in response to the heightened income risk. For durables, there is also the real-options effect in response to the increased income risk. The fact that model elasticities in response to increased unemployment probability growth are close to the empirical elasticities for both nondurable and durable goods makes it unlikely that the strength of the real-options mechanism for durables is overestimated in our model. We take the close correspondence of the model and empirical estimates as evidence that the model captures a key dimension of the micro-level consumption dynamics, and that the model can be used to infer partial equilibrium consumption responses to aggregate labor market shocks and evaluating counterfactuals, as done in Section 3.

5   Concluding Remarks

In this paper, we have showed that a calibrated buffer-stock savings model predicts that in response to an adverse labor market shock, durable expenditures react primarily through changes in ex-ante beliefs about the likelihood of becoming unemployed, while nondurables react primarily through the ex-post realizations of unemployment spells. The difference is generated by the non-convex adjustment costs involved in durable purchases, which raise the response to the income-risk component of the belief-channel response. Because the ex-ante belief-channel response dominates the realized-shocks channel response for durables, having more credit-constrained 'hand-to-mouth' households may significantly dampen the overall durable demand response to the shock. We have also documented that the responses of our model to an increase in the probability of becoming unemployed is in line with reduced-form evidence using micro survey data, both for nondurable and durable goods, validating that our model can be used to estimate demand responses and perform counterfactuals.

Our model emphasizes the role of households' beliefs about their income processes in shaping aggregate consumption dynamics. In so doing, we have assumed rational expectations, but household belief formation may deviate from this benchmark. A burgeoning literature has started to investigate how households form beliefs about their income processes (Guvenen, 2007, Guvenen et al., 2014, Druedahl and Jorgensen, 2018, Rozsypal and Schlafmann, 2019). From the perspective of our model, further investigations within this research agenda, especially how the belief-formation process varies over the business cycle, seems key to understanding aggregate dynamics.

A recent literature has highlighted that for nondurables, having more credit-constrained 'hand-to-mouth' households may amplify the general-equilibrium feedback effects of aggregate shocks (Auclert et al., 2018, Kaplan et al., 2018), by raising the aggregate-demand response to the equilibrium changes in household income. Our model implies that for durables, the demand response to worsened labor market conditions is significantly dampened when more households are credit constrained, as these households do not respond through the powerful belief channel. Going forward, more research is needed on the role of hand-to-mouth households in shaping equilibrium feedback effects in models that feature both nondurable and durable consumption goods.

References

Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3):488.

Arrow, K. J., Harris, T., and Marschak, J. (1951). Optimal Inventory Policy. *Econometrica*, 19(3):250.

Auclert, A. (2018). Monetary Policy and the Redistribution Channel. *American Economic Review*, Forthc.

Auclert, A., Rognlie, M., and Straub, L. (2018). The Intertemporal Keynesian Cross. *Mimeo.*

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

Barsky, R. B., House, C. L., and Kimball, M. S. (2007). Sticky-Price Models and Durable Goods. *American Economic Review*, 97(3):984–998.

Basu, S. and De Leo, P. (2017). Should Central Banks Target Investment Prices? *Mimeo.*

Bayer, C., Lütticke, R., Pham-Dao, L., and Tjaden, V. (2019). Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk. *Econometrica*, 87(1):255–290.

Benito, A. (2006). Does job insecurity affect household consumption? *Oxford Economic Papers*, 58(1):157–181.

Berger, D., Dew-Becker, I., Milbrandt, K., Schmidt, L., and Takahashi, Y. (2018a). Layoff risk, the welfare cost of business cycles, and monetary policy. *Mimeo.*

Berger, D., Guerrieri, V., Lorenzoni, G., and Vavra, J. (2018b). House Prices and Consumer Spending. *The Review of Economic Studies*, 85(3):1502–1542.

Berger, D. and Vavra, J. (2015). Consumption Dynamics During Recessions. *Econometrica*, 83(1):101–154.

Bernanke, B. S. (1983). Irreversibility, uncertainty and cyclical investment. *Quarterly Journal of Economics*, 98(1):85–106.

Bertola, G., Guiso, L., and Pistaferri, L. (2005). Uncertainty and consumer durables adjustment. *Review of Economic Studies*, 72(4):973–1007.

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica*, 77(3):623–685.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. (2018). Really Uncertain Business Cycles. *Econometrica*, 86(3):1031–1065.

Carroll, C. D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *Quarterly Journal of Economics*, 112(1):1–55.

Carroll, C. D. and Dunn, W. E. (1997). Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets. *NBER Macroeconomics Annual*, 12:165–230.

Carroll, C. D., Dynan, K., and Krane, S. (2003). Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets. *Review of Economics and Statistics*, 85:586–604.

Carroll, C. D. and Kimball, M. S. (2001). Liquidity Constraints and Precautionary Saving. *Mimeo*.

Carroll, C. D., Slacalek, J., Tokuoka, K., and White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3):977–1020.

Carvalho, V. M. and Grassi, B. (2019). Large Firm Dynamics and the Business Cycle. *American Economic Review*, Forth.

Challe, E. and Ragot, X. (2016). Precautionary Saving Over the Business Cycle. *The Economic Journal*, 126(590):135–164.

Chetty, R. and Szeidl, A. (2016). Consumption Commitments and Habit Formation. *Econometrica*, 84(2):855–890.

Davis, S. J. and von Wachter, T. M. (2011). Recessions and the Cost of Job Loss. *Mimeo*.

Den Haan, W. J., Rendahl, P., and Riegler, M. (2018). Unemployment (Fears) and Deflationary Spirals. *Journal of the European Economic Association*, 16(5):1281–1349.

Dixit, A. K. and Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press.

Druedahl, J. (2015). Business Cycle Fluctuations in the Demand for Consumer Durables. *Mimeo*.

Druedahl, J. and Jorgensen, T. H. (2018). Can Consumers Distinguish Persistent from Transitory Income Shocks? *Mimeo*.

Dupor, B., Li, R., Saif, M. M., and Tsai, Y.-C. (2018). The 2008 U.S. Auto Market Collapse. *Mimeo*.

Eberly, J. C. (1994). Adjustment of consumers' durables stocks: Evidence from automobile purchases. *Journal of Political Economy*, 102(3):403–436.

Elsby, M. W. L., Hobijn, B., and ahin, A. (2013). Unemployment Dynamics in the OECD. *Review of Economics and Statistics*, 95(2):530–548.

Erceg, C. and Levin, A. (2006). Optimal monetary policy with durable consumption goods. *Journal of Monetary Economics*, 53(7):1341–1359.

Foote, C., Hurst, E., and Leahy, J. (2000). Testing the (S, s) Model. *American Economic Review: Papers and Proceedings*, 90(2):116–119.

Gabaix, X. (2011). The Granular Origins of Aggregate Fluctuations. *Econometrica*, 79(3):733–772.

Gavazza, A. and Lanteri, A. (2018). Credit Shocks and Equilibrium Dynamics in Consumer Durable Goods Markets. *Mimeo.*

Gilchrist, S., Sim, J., and Zakrajsek, E. (2014). Uncertainty, Financial Frictions, and Investment Dynamics. *Mimeo.*

Greenacre, M. (2007). *Correspondence Analysis in Practice.* Chapman & Hall/CRC, 2nd edition.

Grossman, S. J. and Laroque, G. (1990). Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods. *Econometrica*, 58(1):25.

Guvenen, F. (2007). Learning Your Earning: Are Labor Income Shocks Really Very Persistent? *American Economic Review*, 97(3):687–712.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2016). What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk? *Mimeo.*

Guvenen, F., Ozkan, S., and Song, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, 122(3):621–660.

Hassler, J. (1996). Variations in risk and fluctuations in demand: A theoretical model. *Journal of Economic Dynamics and Control*, 20(6-7):1115–1143.

Heathcote, J. and Perri, F. (2018). Wealth and Volatility. *The Review of Economic Studies*, 85(4):2173–2213.

Jacobson, L. S., Lalonde, R. J., and Sullivan, D. G. (1993). Earning losses of displaced workers. *American Economic Review*, 83(4):685–709.

Jappelli, T. and Pistaferri, L. (2014). Fiscal Policy and MPC Heterogeneity. *American Economic Journal: Macroeconomics*, 6(4):107–136.

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

Kaplan, G. and Violante, G. L. (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica*, 82(4):1199–1239.

Kekre, R. (2018). Unemployment Insurance in Macroeconomic Stablization. *Mimeo.*

Kimball, M. S. (1990). Precautionary Saving in the Small and in the Large. *Econometrica*, 58(1):53–73.

Krueger, D., Mitman, K., and Perri, F. (2016). Macroeconomics and Household Heterogeneity. In Taylor, J. and Uhlig, H., editors, *Handbook of Macroeconomics, Vol. 2.*

Krueger, D., Perri, F., Pistaferri, L., and Violante, G. L. (2010). Cross Sectional Facts for Macroeconomists. *Review of Economic Dynamics*, 13(1):1–14.

Martin, J. P. (1996). Measures of replacment rates for the purpose of international comparisons: A note. *OECD Economic Studies*, 26.

McKay, A. (2017). Time-varying idiosyncratic risk and aggregate consumption dynamics. *Journal of Monetary Economics*, 88:1–14.

McKay, A. and Reis, R. (2017). Optimal Automatic Stabilizers. *Mimeo.*

McKay, A. and Wieland, J. F. (2019). Lumpy Durable Consumption Demand and the Limited Ammunition of Monetary Policy. *Mimeo.*

Ogaki, M. and Reinhart, C. M. (1998). Measuring Intertemporal Substitution: The Role of Durable Goods. *Journal of Political Economy*, 106(5):1078–1098.

Ravn, M. O. and Sterk, V. (2017). Job uncertainty and deep recessions. *Journal of Monetary Economics*, 90:125–141.

Reed, W. J. (2001). The Pareto, Zipf and other power laws. *Economics Letters*, 74(1):15–19.

Rozsypal, F. and Schlafmann, K. (2019). Overpersistence Bias in Individual Income Expectations and its Aggregate Implications. *Mimeo.*

Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 15(2):127–148.

Sterk, V. and Tenreyro, S. (2018). The Transmission of Monetary Policy through Redistributions and Durable Purchases. *Journal of Monetary Economics*, 99:124–137.

Stock, J. H. and Watson, M. W. (1999). Business Cycle Fluctuations in US Macroeconomic Time Series. In *Handbook of Macroeconomics*, pages 3–64. Elsevier.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Cyclical Dynamics in Idiosyncratic Labor Market Risk. *Journal of Political Economy*, 112(3):695–717.

Tauchen, G. and Hussey, R. (1991). Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models. *Econometrica*, 59(2):371–396.

A Solution Method

In this appendix, we provide a description of how we solve the consumption model presented in Section 2. First, we describe how the consumption problem can be rewritten in a simplified form with fewer state variables. Second, we describe the solution method and how it is implemented in practice.

State space reduction in the recursive formulation. Due the combined assumptions of a linear replacement rate, a linear adjustment cost and preferences with constant relative risk aversion over a homothetic bundle of the two goods, the household problem can be normalized with respect to the permanent income state Z_t, similar to a standard buffer-stock model with only one good (see e.g. Carroll (1997)). In addition, since the transitory shocks have no dependence on past variables, ϵ can also be eliminated as a state variable. Finally, conditional on adjusting, one more state variable can be eliminated as the optimization problem only depends on the total available resources today. We describe each of these simplifications in turn.

Before describing the normalization with respect to Z_t, we make the variable substitution $\hat{B} = B + \hat{\chi}D$. This normalizes the borrowing constraint to $\hat{B} \geqslant 0$:

$$V_{NA}(\hat{B}, D; Z, \epsilon, n, \Theta) = \max_{C, \hat{B}^*} u(C, D^*) + \beta E V(\hat{B}', D'; Z', \epsilon', n', \Theta') \tag{19}$$

$$\text{s.t.} \quad \hat{B}' = \hat{B}^*$$

$$D' = D^*$$

$$D^* = (1 - \delta)D$$

$$C + q\hat{B}^* \leqslant Y(n + b(1 - n)) + \hat{B} - \hat{\chi}(1 - q(1 - \delta))D,$$

$$\hat{B}^*, C \geqslant 0,$$

$$Y = Z\epsilon,$$

$$V_A(\hat{B}, D; Z, \epsilon, n, \Theta) = \max_{C, D^*, \hat{B}^*} u(C, D^*) + \beta E V(\hat{B}', D'; Z', \epsilon', n', \Theta') \tag{20}$$

$$\text{s.t.} \quad \hat{B}' = \hat{B}^*$$

$$D' = D^*$$

$$C + (1 - q\hat{\chi})D^* + q\hat{B}^* \leqslant Y(n + b(1 - n)) + \hat{B} + (1 - \delta - h - \hat{\chi})D,$$

$$\hat{B}^*, C, D^* \geqslant 0,$$

$$Y = Z\epsilon,$$

$$V(\hat{B}, D; Z, \epsilon, n, \Theta) = \max\{V_{NA}(\hat{B}, D; Z, \epsilon, n, \Theta), V_A(\hat{B}, D; Z, \epsilon, n, \Theta)\}. \tag{21}$$

Now we normalize the household problem with respect to permanent income Z_t. We make the following

definitions:

- $\nu = VZ^{-(1-\sigma)}$ for variables $V = V, V_{NA}, V_A$

- $x = X/Z$ for any other variable X

- $x' = X'/Z'$ for any variable X'

Using that $u(\cdot)$ is homothetic, the recursive problem can be reformulated in terms of ν, x, x' without dependence of the state variable Z:

$$
\nu_{NA}(\hat{b}, d; \epsilon, n, \Theta) = \max_{c, \hat{b}^*} u(c, d^*) + \beta E \eta'^{1-\sigma} \nu(\hat{b}', d'; \epsilon', n', \Theta') \tag{22}
$$
$$
\text{s.t.} \quad \hat{b}' = \eta'^{-1} \hat{b}^*
$$
$$
d' = \eta'^{-1} d^*
$$
$$
d^* = (1 - \delta)d
$$
$$
c + q\hat{b}^* \leqslant \epsilon(n + b(1 - n)) + \hat{b} - \hat{\chi}(1 - q(1 - \delta))d,
$$
$$
b^*, c \geqslant 0,
$$
$$
\nu_A(\hat{b}, d; \epsilon, n, \Theta) = \max_{c, d^*, \hat{b}^*} u(c, d^*) + \beta E \eta'^{1-\sigma} \nu(\hat{b}', d'; \epsilon', n', \Theta') \tag{23}
$$
$$
\text{s.t.} \quad \hat{b}' = \eta'^{-1} \hat{b}^*
$$
$$
d' = \eta'^{-1} d^*
$$
$$
c + (1 - q\hat{\chi})d^* + q\hat{b}^* \leqslant \epsilon(n + b(1 - n)) + \hat{b} + ((1 - \delta)(1 - h) - \hat{\chi})d,
$$
$$
\hat{b}^*, c, d^* \geqslant 0,
$$
$$
\nu(\hat{b}, d; \epsilon, n, \Theta) = \max\{\nu_{NA}(\hat{b}, d; \epsilon, n, \Theta), \nu_A(\hat{b}, d; \epsilon, n, \Theta)\}. \tag{24}
$$

Second, ϵ can be eliminated as a state variable as it enters through the sufficient state variable $a = \epsilon(n + b^u(1 - n)) + \hat{b}$. Using this, we can write the problem as

$$v_{NA}(a,d;n,\Theta) = \max_{c,\hat{b}^*} u(c,(1-\delta)d) \tag{25}$$

$$+\beta E\eta'^{1-\sigma}v(\epsilon'(n'+b(1-n'))+\eta'^{-1}\hat{b}^*,\eta'^{-1}(1-\delta)d;n',\Theta')$$

$$\text{s.t.} \quad c+q\hat{b}^* \leqslant a-\hat{\chi}(1-q(1-\delta))d,$$

$$b^*,c \geqslant 0,$$

$$v_A(a,d;n,\Theta) = \max_{c,d^*,\hat{b}^*} u(c,d^*)+\beta E\eta'^{1-\sigma}v(\epsilon'(n'+b(1-n'))+\eta'^{-1}\hat{b}^*,\eta'^{-1}d^*;n',\Theta') \tag{26}$$

$$\text{s.t.} \quad c+(1-q\hat{\chi})d^*+q\hat{b}^* \leqslant a+((1-\delta)(1-h)-\hat{\chi})d,$$

$$\hat{b}^*,c,d^* \geqslant 0,$$

$$v(a,d;n,\Theta) = \max\{v_{NA}(a,d;\epsilon,n,\Theta),v_A(a,d;n,\Theta)\}. \tag{27}$$

Finally, note that conditional on adjusting, the only state variable is $w=a+((1-\delta)(1-h)-\hat{\chi})d$, such that

$$v_A(w;n,\Theta) = \max_{c,d^*,\hat{b}^*} u(c,d^*)+\beta E\eta'^{1-\sigma}v(\epsilon'(n'+b(1-n'))+\eta'^{-1}\hat{b}^*,\eta'^{-1}d^*;n',\Theta') \tag{28}$$

$$\text{s.t.} \quad c+(1-q\hat{\chi})d^*+q\hat{b}^* \leqslant w$$

$$\hat{b}^*,c,d^* \geqslant 0.$$

Given this formulation, a solution to the household problem is a collection of policy functions h_c^{NA}, h_b^{NA} and a value function v_{NA} that solve (26), policy functions h_c^A, h_b^A, h_d^A and a value function v_A that solve (28), and a value function v that, given v_{NA}, v_A, solves (27).

Computation. We solve the recursive problem by value function iteration. The convergence criterion is specified in terms of the distance between two consecutive value functions under the sup norm. We set the criterion to 10^{-5} when extracting decision rules used for the simulations in Sections 3 and 4, and to 10^{-2} for the calibration.

We discretize the processes for the income shocks. For the income shocks, we use Hermite-Gauss polynomials with five states. For the model augmented with idiosyncratic fluctuations in labor market transition rates analyzed in Section 4, we discretize the process for the transition rates using the method by Tauchen and Hussey (1991) with seven states.

The grids for the endogenous states a, d, w are linear up to a cutoff value and exponential in a sparse grid above the cutoff value. Given a value function v^i, we solve the expectation over the future value function by linear interpolation. Then, we compute v_{NA}^{i+1}, v_A^{i+1} by nonlinear optimization methods using

several initial guesses. We compute ν_A^{i+1} first for a given level of total assets and then interpolate ν_A^{i+1} for the two-dimensional state space by cubic interpolation. Finally, we retrieve ν^{i+1} by taking the maximum of $\{\nu_{NA}^{i+1}, \nu_A^{i+1}\}$.

B   Simulation Method

In this appendix, we describe the simulation method for computing the impulse-response functions in Section 3. Our simulation method greatly reduces the number of agents needed to estimate the evolution of aggregate variables precisely and applies to the entire class of consumption-savings models with permanent-income shocks where the decision functions can be normalized by permanent income (for other examples, see, e.g., Carroll (1997) and, more recently, McKay (2017) and Carroll et al. (2017)). To the best of our knowledge, this simulation method is new, and we therefore describe it in some detail. We start by intuitively describing why simulating this class of models is tedious. We then describe our simulation method which uses *permanent-income neutral probabilities*. Lastly, we prove that the law of motion for aggregates in the auxilliary model that employs permanent-income neutral probabilities is the same as the law of motion of the original model.

B.1   What is the problem?

Consider our consumption-savings model described by equations (19)-(21). Without loss of generality, we here suppress the employment state n and the aggregate state Θ, as their law of motion is entirely exogenous and aggregating across employment states is trivial. After normalizing our consumption-savings model with the permanent earnings potential Z (for short, we refer to this object as "permanent income" in this appendix) following the outline in Appendix A, the state variables are a and d, where a and d are normalized cash-on-hand and durable assets. To retrieve the non-normalized decision functions, the additional state variable is permanent income Z. The dynamics in the normalized cash-on-hand and durable asset dimensions are well behaved: convergence to the ergodic steady state is quick and the state space is bounded (if the shock distributions are bounded). The difficulties lie with the dynamics of the permanent-income dimension Z. Convergence to the ergodic steady state in this dimension is slow and accurately estimating total income requires a very large number of households.

As an example, we simulated the evolution of earnings potential for 100,000 households over 1,000 periods using our calibrated parameter values. We repeated the simulation 1,000 times and recorded the mean earnings potential of the population. We report the results in the histogram in Figure 5. Although the simulations most of the time yield a mean permanent income close to the true value of 1.0, substantial deviations still occur. 15.3% of simulations yielded a simulation error large than 1.0% and 1.6% of simulations yielded a simulation error larger than 2.0%. These variations are substantial and noticeable, on the same
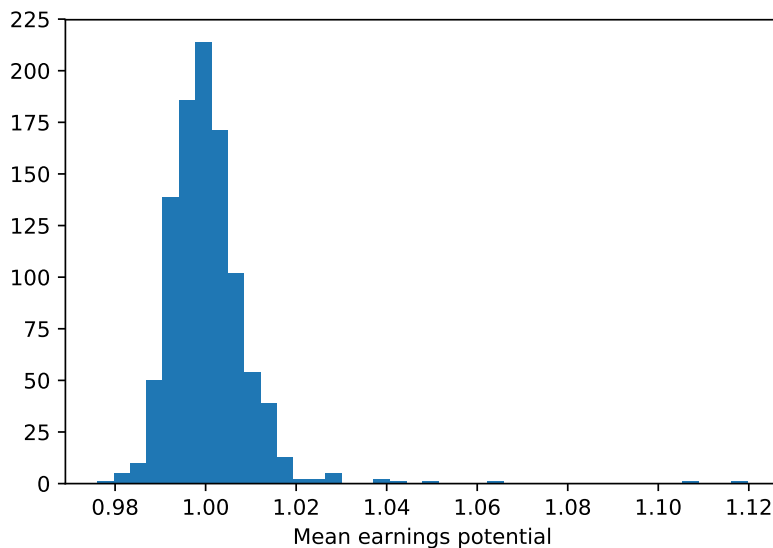
Figure 5: *Histogram of the mean earnings potential for 1,000 draws of 100,000 households after 1,000 periods.*

order of magnitude as a business-cycle shock.

That simulation errors are so large is not surprising: the random growth model with perpetual-youth households generates a stationary distribution with a Pareto tail (Reed, 2001). Intuitively, fluctuations in aggregate income, and therefore aggregate consumption, are to a large extent determined by shocks to the proverbial Bill Gates and Jeff Bezos of the model economy. The granularity literature explores to what extent idiosyncratic shocks at the individual household or firm level can generate aggregate fluctuations, essentially exploiting the poor convergence properties (Gabaix, 2011, Carvalho and Grassi, 2019).

B.2 Using permanent-income neutral probabilities

The simulation imprecision in the permanent-income dimension is a nuisance. We show that there is a way around keeping track of the permanent-income dimension when simulating aggregates. The idea is to simulate the evolution of "units of permanent income" instead of simulating the evolution of households. Keeping track of the permanent-income weighted distribution amounts to simulating the model with distorted probabilities which we call permanent-income neutral probabilities, in analogy with the (distorted) risk-neutral probabilities employed in asset pricing.

We write $X = (a, d) \subset \mathbf{X} = \mathbf{A} \times \mathbf{D}$ for the normalized assets of a household. The distribution of households over the state space $\mathbf{X} \times \mathbf{Z}$ is denoted by $\mu_{\mathbf{X} \times \mathbf{Z}}$ where $\mu_{\mathbf{X} \times \mathbf{Z}}(\mathcal{X} \times \mathcal{Z})$ reads "the mass of households

in the (measurable) subset $\mathcal{X} \times \mathcal{Z}$ of the state space".[16] The nondurable consumption of a household with state (X, Z) is given by $h_c(X)Z$. In other words, the consumption decision scales with Z. To compute total nondurable consumption, we integrate nondurable consumption over the distribution of households:

Total consumption $= \int_{\mathbf{X} \times \mathbf{Z}} h_c(X) Z d\mu_{\mathbf{X} \times \mathbf{Z}}$

Define the permanent-income weighted distribution by $\tilde{\mu}_{\mathbf{X}}(\mathcal{X}) = \int_{\mathcal{X} \times \mathbf{Z}} Z d\mu_{\mathbf{X} \times \mathbf{Z}}$ for all measurable sets $\mathcal{X} \subset \mathbf{X}$. $\tilde{\mu}_{\mathbf{X}}(\mathcal{X})$ reads "the mass of permanent-income-weighted purchasing power in the (measurable) subset $\mathcal{X}$ of the state space $\mathbf{X}$". For simulating aggregates, we do not need to know the full joint distribution over all state variables $\mu_{\mathbf{X} \times \mathbf{Z}}$, it suffices to know the permanent-income weighted distribution $\tilde{\mu}_{\mathbf{X}}$. For example, total nodurable consumption expenditure is given by $\int h_c(X) d\tilde{\mu}_{\mathbf{X}}$ where h_c is the normalized consumption function.

If we can simulate the law of motion for $\tilde{\mu}_{\mathbf{X}}$ directly and use this measure to compute aggregates, we bypass keeping track of the permanent-income distribution, which greatly improves the precision of the simulations. Once having solved for the decision functions using the method outlined in Appendix A, this can be done using the following steps:

1. Replace the probability distribution for the permanent-income shocks η, $\{p_i^\eta\}$, with the permanent-income neutral probability distribution $\{\tilde{p}_i\} = \{\eta_i p_i^\eta\}$. Note that since $E\eta = 1$, the weights $\{\eta_i p_i^\eta\}$ do indeed constitute a probability distribution.

2. Simulate the model using the permanent-income neutral probabilities, only keeping track of the normalized asset distibution.

3. Compute aggregate nondurable consumption as the simulated mean of normalized nondurable consumption, and similarly for liquid assets and durable assets.

In the next subsection, we prove the equivalence between simulating aggregate variables this way and simulating the model in a standard fashion.

As for the IRFs in Section 3, we populate the economy with $1,000$ households and simulate the evolution of the aggregates after a burn-in period of 200 quarters, and repeat the experiment 320 times. In order to gain the same precision when simulating without using permanent-income neutral probabilities, earlier experiments suggest that you need at least a burn-in period of 1000 quarters and repeat the experiment at least $24,000$ times.

B.3   Proof

Here, we prove the equivalence between the evolution of aggregates for the original model and the evolution of aggregates for the auxilliary model with permanent-income neutral probabilities. As before, without

---

[16]We use bold-font letters $(\mathbf{X}, \mathbf{A}, \mathbf{D})$ for the entire state space, caligraphic letters $(\mathcal{X}, \mathcal{A}, \mathcal{D})$ for subsets of the corresponding state space, and standard letters (X, a, d) for points in the state space.

loss of generality, we here suppress the employment state n and the aggregate state Θ. Consider the solution to the decision problem described by (26)-(28). The solution is characterized by the decision functions $c = h_c(X), \hat{b}^* = h_b(X), d^* = h_d(X)$. Write $h(X) = [h_b(X), h_d(X)]$. For a given realization of the permanent and transitory income shocks η', ϵ', the stochastic environment maps X, Z to a new state, $(X, Z) = (a, d, Z) \mapsto (b^*/\eta' + \epsilon', d^*/\eta', Z\eta')$. The distribution of individuals over states $(X, Z) \in \mathbf{X} \times \mathbf{Z}$ is given by the measure $\mu_{\mathbf{X} \times \mathbf{Z}}$ (where $\mathbf{X} \subset \mathbb{R}^2$ and $\mathbf{Z} \subset \mathbb{R}$ are endowed with the Borel σ-algebra). The distribution of households over normalized savings and durable assets is given by $\mu_{\mathbf{X}}(\mathcal{X}) = \int_{\mathcal{X} \times \mathbf{Z}} d\mu_{\mathbf{X} \times \mathbf{Z}}$, where $\mathcal{X}$ is a measurable subset of $\mathbf{X} = \mathbf{A} \times \mathbf{D}$. The permanent-income weighted distribution $\tilde{\mu}_{\mathbf{X}}$ is defined as earlier, $\tilde{\mu}_{\mathbf{X}}(\mathcal{X}) = \int_{\mathcal{X} \times \mathbf{Z}} Z d\mu_{\mathbf{X} \times \mathbf{Z}}$. Note that the permanent-income weighted distribution is a sufficient statistic for computing aggregates.

We now show how to get a law of motion for $\tilde{\mu}_{\mathbf{X}}$ by using the permanent-income neutral probabilities. Although the computation of the law of motion is heavy on notation, in the end it amounts to little more than accounting. To compute the mass of households in a subset $\mathcal{X}'$ tomorrow, we sum over all the possible shock realizations and see how many households are in the pre-image of the decision function.

First we compute the distribution $\mu_{\mathbf{X}' \times \mathbf{Z}'}$ of households in the next period. For a given pair of shocks η', ϵ' we have $(a, d, Z) \mapsto (a', d', Z') = (h_b(a, d)/\eta' + \epsilon', h_d(a, d)/\eta', \eta' Z)$. Write $p_i^\eta = P(\eta' = \eta_i)$ and $p_j^\epsilon = P(\epsilon' = \epsilon_j)$. The measure of households with state (X', Z') in the set $\mathcal{X}' \times \mathcal{Z}'$ is then given by

$$
\begin{aligned}
\mu_{\mathbf{X}' \times \mathbf{Z}'}(\mathcal{X}' \times \mathcal{Z}') &= \sum_i \sum_j p_i^\eta p_j^\epsilon \mu_{\hat{\mathbf{B}} \times \mathbf{D} \times \mathbf{Z}}\left[(\eta_i(\mathcal{X}' - (\epsilon_j, 0)) \times \mathcal{Z}'/\eta_i \right] \\
&= \sum_i \sum_j p_i^\eta p_j^\epsilon \mu_{\mathbf{X} \times \mathbf{Z}}\left[h^{-1}(\eta_i(\mathcal{X}' - (\epsilon_j, 0)) \times \mathcal{Z}'/\eta_i \right].
\end{aligned}
$$

This expression describes the law of motion of the distribution of households given the shocks η_i, ϵ_j and the decision function h.

The distribution of households in normalized asset space $\mathbf{X}$, is given by

$$
\begin{aligned}
\mu_{\mathbf{X}'}(\mathcal{X}') &= \int_{\mathcal{X}' \times \mathbf{Z}'} d\mu_{\mathbf{X}' \times \mathbf{Z}'} \\
&= \sum_i \sum_j p_i^\eta p_j^\epsilon \int_{h^{-1}(\eta_i(\mathcal{X}' - (\epsilon_j, 0))) \times \mathbf{Z}} d\mu_{\mathbf{X} \times \mathbf{Z}} \\
&= \sum_i \sum_j p_i^\eta p_j^\epsilon \mu_{\mathbf{X}}\left[h^{-1}(\eta_i(\mathcal{X}' - (\epsilon_j, 0))) \right] \\
&= E_\eta\left[E_\epsilon\left[\mu_{\mathbf{X}}\left[h^{-1}(\eta_i(\mathcal{X}' - (\epsilon_j, 0))) \right] \right] \right]. \quad (29)
\end{aligned}
$$

This distribution is not a sufficient statistic for computing aggregates since we need to weigh the importance of a household by its permanent income.

The law of motion for the permanent-income weighted distribution can also be computed analogously.

$$
\begin{aligned}
\tilde{\mu}_{\mathbf{X'}}(\mathcal{X}') &= \int_{\mathcal{X'}\times\mathbf{Z'}} Z'd\mu_{\mathbf{X'}\times\mathbf{Z'}} \\
&= \sum_i \sum_j p_i^\eta p_j^\epsilon \int_{h^{-1}\left(\eta_i(\mathcal{X'}-(\epsilon_j,0))\right)\times\mathbf{Z}} (\eta_i Z)d\mu_{\mathbf{A}\times\mathbf{D}\times\mathbf{Z}} \\
&= \sum_i \sum_j (p_i^\eta \eta_i) p_j^\epsilon \tilde{\mu}_{\mathbf{A}\times\mathbf{D}}\left[h^{-1}(\eta_i(\mathcal{X'}-(\epsilon_j,0)))\right] \\
&= \tilde{\mathsf{E}}_\eta\left[\mathsf{E}_\epsilon\left[\tilde{\mu}_{\mathbf{X}}\left[h^{-1}(\eta_i(\mathcal{X'}-(\epsilon_j,0)))\right]\right]\right]
\end{aligned}
\tag{30}
$$

where $\tilde{\mathsf{E}}_\eta$ is computed using the permanent-income neutral probabilities $\{\tilde{p}_i\} = \{p_i\eta_i\}$.

Comparing (30) and (29), we see that the law of motion for the weighted measure is the same as the law of motion for the unweighted measure, except that the probability distribution for the permanent shock is replaced by the permanent-income-neutral distribution $\{\tilde{p}\} = \{p_i\eta_i\}$.

B.4   Comments

Note that our method of simulating aggregates extends to the entire class of models in which all decision functions can be normalized with permanent income. For recent models within this class, see, e.g., Carroll et al. (2017) and McKay (2017). Also note that the method can be straightforwardly extended for simulations of higher order moments, such as variances, skewness, kurtosis, etc.

C   Estimation of Aggregate Labor Market Transition Rates

We estimate the quarterly separation and job-finding rates using the method developed by Elsby et al. (2013), which extends Shimer (2012), using annual data for the Italian unemployment rate, grouped by the duration of the unemployment spell. A key advantage of the method is that it is robust to temporal aggregation bias, as the rates are inferred from an underlying continuous-time process. We retrieve the data from the OECD for the period 1984-2014, allowing us to estimate the quarterly rates for the period 1984-2013.

Let t denote a quarter. To estimate the quarterly job finding rate f_t, define $F_t^{<d}$ as the probability that an unemployed worker exits unemployment within d quarters. $F_t^{<d}$ is estimated from

$$
F_t^{<d} = 1 - \frac{u_{t+d} - u_{t+d}^{<d}}{u_t}
$$

with an associated outflow rate given by

$$
f_t^{<d} = -\log(1 - F_t^{<d})/d
$$
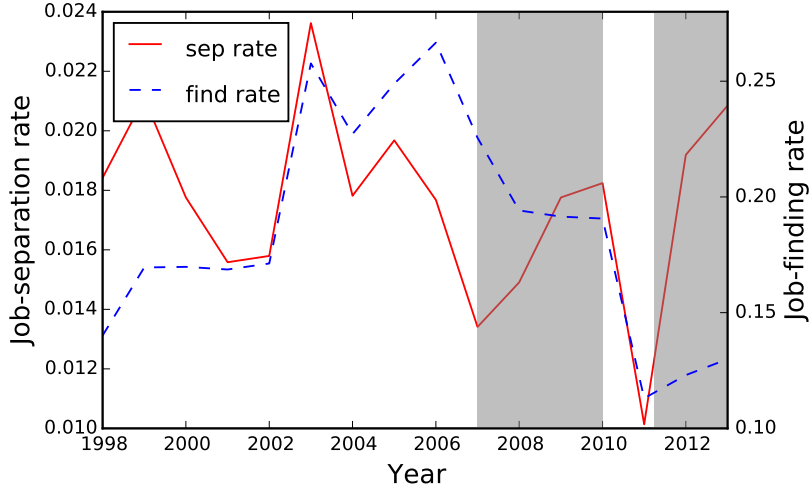
Figure 6: *Estimated quarterly job-separation and job-finding rates, Italy 1998-2013. ECRI recession dates are indicated by the Shaded areas.*

Normalizing the time scale, we observe $u_{t+d}^{<d}, u_{t+d}$ in year $t+d$. We infer u_t by taking the weighted geometric average of u_{t+d} and u_{t+d-4},

$$u_t = u_{t+d}^{d/4} u_{t+d-4}^{(4-d)/4}.$$

The OECD data allow us to observe unemployment rates with duration less than $1, 3, 6$ and 12 months. Accordingly, we estimate the flow rates $f_t^{<1/3}, f_t^{<1}, f_t^{<2}, f_t^{<4}$. Then, we compute the average job finding rate f_t as the simple average of these four variables.[17]

Given f_t and u_t we can infer s_t from the law of motion for the aggregate unemployment rate:

$$\frac{\partial u}{\partial t} = s_t(1 - u_t) - f_t u_t. \tag{31}$$

Assuming that the flows are constant over a year and solving (31) one year forward, we have

$$u_t = \kappa_t u_t^* + (1 - \kappa_t) u_{t-4}, \tag{32}$$

where $\kappa_t = 1 - e^{-4(s_t + f_t)}$ and $u_t^* = \frac{s_t}{s_t + f_t}$. Given f_t, we use (32) to solve for the separation rate s_t.

The resulting quarterly job-finding and job-separation rates are shown in Figure 6. Between the start of the Eurozone crisis in 2011 and the last period of observation in 2013, the separation rate increased from 1.01 percent to 2.08 percent, a relative increase of 106 percent.

---

[17]Elsby et al. (2013) use an optimal weighting scheme based on minimizing the mean squared error of the estimate.

Parameters	Value	Target	Empirical	Model
Discount rate, β	0.97	Mean norm. fin. assets	1.72	X
Nondurable weight, α	0.95	Mean norm. dur. assets	1.24	X
Depreciation rate, δ	0.038			
Adjustment cost, h	0			
Collateral constraint, $\hat{\chi}$	0.95	Share neg. fin. wealth	0.11	0.10

Table 6: *The calibrated parameters of the flexible-adjustment model. The targeted moments are computed from the SHIW sample used in Section 4. For moments that concern the stock and flow of durable goods, we use the stock and flow of motor vehicles in the data.*

There is a small computational error in how we use these estimated transitiion rates. $\bar{\zeta}$ and λ are job-separation and job-finding probabilities. In our computations, those have been calibrated to equal the average job-separation and job-finding rates in the data. This means that we understate job-finding probability in the model by approximately 1.5 percentage points (a relative difference of approximately 10 percent). For the job-separation probability, the difference is negligible. The discrepancy will be corrected in the next revision of the paper, but we do not expect it to significantly affect any of our results.

D   Calibration of Flexible-Adjustment Model

For the calibration of the flexible-adjustment model used in Section 3, we set $h = 0$ and take the calibrated depreciation rate of the baseline model. We then calibrate β, α and $\hat{\chi}$ to match the mean level of normalized financial assets, the mean level of normalized durable assets and the share of households with negative wealth in the data – the same targets that were used to calibrate the baseline model. The recalibrated parameters are shown in Table 6.

E   Model Moments When Varying the Discount Factor

To construct Figure 4 in Section 3, we computed the responses to the job-separation shock in the baseline and the flexible-adjustment model when varying the discount factor β within the range $[0.920, 0.973]$. In Tables 7 and 8, we show how the moments used for the calibration change when varying β. As seen, varying β has a large effect on the mean normalized financial asset holdings and the share of household with negative financial asset holdings, but it does not significantly affect the other calibration moments.

F   Data Description

In this appendix, we describe the SHIW data in more detail and give precise definitions of all variables used.

β	Fin. Assets	Dur. Assets	Purch. Freq.	Purchase Size	Share Neg. Fin. Assets
0.920	-0.48	0.95	0.15	0.83	0.96
0.928	-0.46	0.99	0.15	0.87	0.93
0.937	-0.42	1.04	0.15	0.91	0.88
0.945	-0.33	1.08	0.15	0.96	0.79
0.945	-0.33	1.08	0.15	0.96	0.79
0.946	-0.31	1.09	0.15	0.97	0.77
0.948	-0.28	1.10	0.15	0.98	0.74
0.949	-0.26	1.11	0.15	0.99	0.71
0.951	-0.22	1.12	0.15	1.00	0.68
0.952	-0.19	1.13	0.15	1.02	0.65
0.954	-0.14	1.14	0.15	1.03	0.61
0.955	-0.09	1.15	0.15	1.04	0.57
0.956	-0.03	1.16	0.15	1.05	0.53
0.958	0.05	1.18	0.15	1.07	0.48
0.959	0.13	1.19	0.14	1.09	0.43
0.961	0.24	1.20	0.14	1.10	0.37
0.962	0.36	1.22	0.14	1.12	0.32
0.964	0.51	1.23	0.14	1.14	0.27
0.965	0.69	1.25	0.14	1.16	0.22
0.967	0.91	1.26	0.14	1.18	0.17
0.968	1.19	1.28	0.14	1.20	0.13
0.969	1.54	1.30	0.14	1.23	0.10
0.971	2.00	1.32	0.14	1.25	0.08
0.972	2.59	1.34	0.14	1.28	0.06

Table 7: *Moments of the baseline model when varying the discount factor* β. *For the definition of the moments, see Section 2.*

β	Fin. Assets	Dur. Assets	Share Neg. Fin. Assets
0.921	-0.87	1.16	1.00
0.929	-0.83	1.16	1.00
0.938	-0.75	1.17	1.00
0.946	-0.63	1.17	0.98
0.946	-0.63	1.17	0.98
0.947	-0.60	1.18	0.97
0.949	-0.57	1.18	0.95
0.950	-0.53	1.18	0.93
0.952	-0.48	1.18	0.91
0.953	-0.43	1.18	0.88
0.955	-0.37	1.18	0.84
0.956	-0.31	1.19	0.79
0.958	-0.23	1.19	0.72
0.959	-0.13	1.19	0.64
0.960	-0.02	1.20	0.55
0.962	0.11	1.20	0.46
0.963	0.29	1.20	0.35
0.965	0.48	1.21	0.27
0.966	0.69	1.21	0.21
0.968	0.96	1.22	0.16
0.969	1.31	1.23	0.12
0.971	1.77	1.24	0.09
0.972	2.35	1.25	0.06
0.973	3.09	1.26	0.05

Table 8: *Moments of the flexible-adjustment model when varying the discount factor* β*. For the definition of the moments, see Section* 2*.*

F.1 Sampling structure

Since 1998, the survey includes approximately 8000 households in each wave, of which about half has been interviewed in previous surveys, with gradual replacement over time. Table 9 shows the structure of the panel.

Year	1998	2000	2002	2004	2006	2008	2010	2012	2014
1998	7147	3873	2591	1855	1476	1284	1088	915	634
2000		8001	3605	2522	1951	1682	1418	1171	804
2002			8011	3604	2623	2207	1834	1511	1025
2004				8012	3957	3202	2620	2142	1420
2006					7768	4345	3476	2790	1834
2008						7977	4621	3596	2315
2010							7951	4611	2894
2012								8151	4459
2014									8156

Table 9: *The sampling structure of the SHIW, 1998-2014. Each row corresponds to one survey wave. In each row, the columns show the number of households that have also been surveyed in the year indicated by the column.*

F.2 Definition of all variables used

In this subsection, we provide definitions of all variables used that are non self-explanatory.

Employment status. Category variable that indiciates the employment status held for most part of that year.

Marital status. Category variable that takes 4 values: Married, Single, Separated or Widow/er.

Education level. Category variable that takes 8 values: None, Primary school certificate, Lower secondary school certificate, Vocational secondary school diploma, Upper secondary school diploma, 3-year university degree, 5-year university degree and Postgraduate qualification.

2-digit region indicator. Category variable with 20 values, one for each administrative region of Italy.

1-digit superregion indicator. Category variable with 3 values: North, Centre and South-Islands.

Town size. Category variable with 4 values: 0-20,000, 20,000-40,000, 40,000-500,000 and 500,000+.

Occupation. Category variable that take 5 values conditioned on being employed by second party. Refers to the situation of the person for most of the last 12 months.

Industry. Category variable that take 21 values and indicates where the person currently works.

Household income. Defined as total net disposable income over the last 12 months from summing labor income, pensions and transfers, income from self-employment and income from financial assets and property.

Labor income. Defined as total net payroll income over the last 12 months, including fringe benefits.

Binding liquidity constraint. Indicator variable that takes the value of 1 if the household has reported 1) that a member of the household has applied for a loan and been partly or fully refused and/or 2) a member of the household considered applying for a loan but later changed his/her mind in anticipation that the loan would be refused.

Furniture stock. The self-estimated value of all household belongings of furniture, furnishings, household appliances and sundry equipment.

Vehicle stock. The self-estimated value of all household belongings of cars and other means of transport.

Durable stock. The sum of furniture and vehicle stock.

Net financial assets. The sum of all financial assets, e.g. deposit accounts, savings accounts, stocks, bonds, funds, shares in partnerships etc., net of debt owed to other households and debt for purchases of consumption goods, thus not including mortgage debt.

Net total assets. The sum of real assets (property, jewellry, business equity), financial assets and liabilities.

Nondurable expenditures. Self-estimated total spending less of expenditures on furniture, motor vehicles and jewellery. Does not include actual or imputed rents. Does include fringe benefits.

Vehicle expenditures. Net expenditures on cars and other means of transport.

Vehicle purchase. A dummy variable that takes value 1 if vehicle expenditures in year t exceed 1/104 times the cross-sectional mean of household income in year t, that is, it exceeds one half of the average household income for one week.

	Value/Mean	Std. Deviation	Minimum	Maximum
Share males	0.81	.	.	.
Age	45.86	7.46	25.0	60.0
Share with durable purchase	0.41	.	.	.
Share with vehicle purchase	0.14	.	.	.
Share with furniture purchase	0.32	.	.	.
Nondurable consumption/income	0.57	0.23	0.07	4.26
Net total assets/income	4.52	4.29	-3.51	40.22
Net financial assets/income	0.43	0.89	-3.16	34.12
Unemployment probability	0.03	0.04	0.0	0.36
Unemployment probability growth	0.0	0.02	-0.2	0.22
Share denied/discouraged credit applicants	0.04	.	.	.

Table 10: *Descriptive statistics of the sample used for the final regressions.*

Furniture expenditures. Net expenditures on furniture, furnishings, household appliances and sundry equipment.

Furniture purchase. A dummy variable that takes value 1 if furniture expenditures in year t exceed 1/104 times the cross-sectional mean of household income in year t, that is, it exceeds one half of the average household income for one week.

Durable expenditures. The sum of vehicle and furniture expenditures.

Durable purchase. A dummy variable that takes value 1 if either or both of the vehicle or furniture purchase variables take value 1.

F.3 Descriptive statistics

In Table 10, we provide descriptive statistics of the main variables that we use in the analysis, including the generated variable Uprob_{it}.

F.4 Data quality

To gauge the quality of the SHIW data, we compare aggregate statistics from the survey with corresponding variables from the national accounts. The results are shown in Figure 7. As seen, disposable income and durable expenditures are well aligned. Total and nondurable expenditures show a significant increase between the years 2002 and 2004 in the SHIW, which is not apparent in the national accounts data. Otherwise, these series also seem well-aligned.
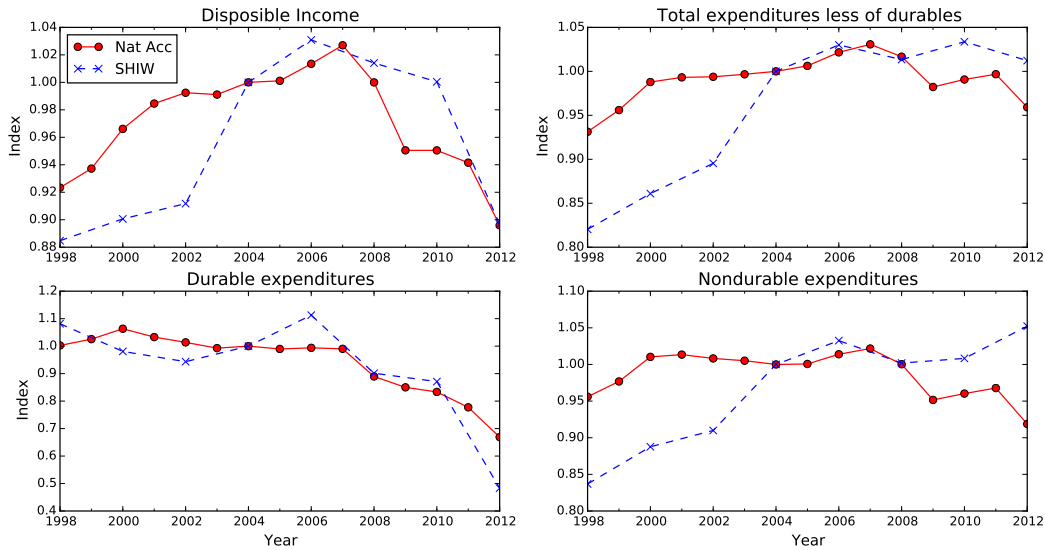
Figure 7: *Comparison of aggregate SHIW data to the national accounts (source: OECD). Total expenditures less of durables is defined as total household expenditure on all items, including housing maintenance cost, actual and imputed rents, less of durable goods. Nondurable expenditures is defined as expenditure on all items, excluding housing maintenance costs, actual rents, imputed rents and durable goods. All variables are deflated by the CPI (source: OECD). National accounts variables have been divided by population size, while SHIW variables are calculated as means using sample weights and dividing by household size.*

G   Robustness Exercises to the Empirical Analysis

In this section, we perform various robustness checks to the empirical exercise in Section 4.

Number of factors used for estimation of unemployment probabilities.   In the estimation of unemployment probabilities, we used 6 factors based on maximizing the adjusted predictive power of our regressions. Here, we report the the estimates from between 4 and 12 factors in the estimation. We report the results in Table 11. Except when using 4 and 12 factors, all point estimates for all durable goods categories stays within one standard deviation of the baseline estimate. When using 4 factors, the estimate for vehicle purchases is significantly larger. When using 12 factors, the same estimate is significantly smaller. When using 8 factors, we also find a significant negative effect of unemployment probability growth on nondurable expenditure growth.

Varying the income cutoff in the sampe selection.   The sample only includes households where the labor income of the household head is at least 40 percent of total household income. We make this restriction to ensure that an unemployment spell is a meaningful income loss for the household. The exact

	(1)	(2)	(3)	(4)
	Nondurables	Durables	Furniture	Vehicles
ΔUprob$_t$ (4 factors)	0.11	-2.72***	-2.27*	-4.71**
	(0.30)	(0.96)	(1.19)	(2.12)
ΔUprob$_t$ (6 factors)	-0.30	-2.42***	-2.36**	-2.83
	(0.26)	(0.82)	(1.02)	(1.84)
ΔUprob$_t$ (8 factors)	-0.47**	-1.85**	-1.94**	-2.96*
	(0.24)	(0.75)	(0.93)	(1.66)
ΔUprob$_t$ (10 factors)	-0.22	-1.65**	-1.73**	-2.25
	(0.21)	(0.68)	(0.83)	(1.55)
ΔUprob$_t$ (12 factors)	-0.05	-1.22**	-1.57**	-0.74
	(0.19)	(0.60)	(0.74)	(1.39)
Char. variables	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Financial variables	Yes	Yes	Yes	Yes

Table 11: *Regression results from estimating* (12) *and* (13)*, when using between 4 and 12 factors for the estimation of unemployment probabilities. For columns (2)-(4), the left-hand side variable is an indicator of whether a durable good has been purchased in the last 12 months. The coefficients show the average marginal effect normalized by the unconditional purchase probability, estimated with a probit model. In column (1), the left-hand side variable is the log growth of expenditures on nondurable goods, estimated by OLS. For explanation of the control variables, we refer to the text in Section 4.2. *, **, *** indicate that the coefficients are significant at the 10%, 5% and 1% level, respectively.*

cutoff, however, is somewhat arbitrary. We rerun the regression when setting this cutoff to 30 and 50 percent, respectively. The results are shown in Table 12. The estimated coefficients of unemployment probability growth on durable purchases are smaller in both cases, but still sizable and substantially larger than the coefficient for nondurables. When increasing the cutoff to 50 percent, we lose almost 30 percent of the sample size, and all coefficients becomes insignificant.

Endogeneity of unemployment probability to future income growth If a firm responds to an adverse demand shock by cutting down on the number of workers, it might also inhibit future wage growth among the remaining workers. Therefore, it is possible that our estimates of the unemployment probability is correlated with expected future wage growth. In this case, our estimated coefficient of unemployment probability growth on expenditures could be biased by the response of expenditures to expected wage growth. To overcome the omitted variable bias, we add income growth between periods t and t+2 as a control variable to the expenditures regressions. The results are shown in Table 13. As seen, the inclusion of this variable does not affect our results.

Endogeneity of unemployment probability to private information We infer the probability of becoming unemployed from households that actually become unemployed. However, it could be the case that households that become unemployed know well in advance whether they will keep their employment or not

	(1) Nondurables	(2) Nondurables	(3) Nondurables	(4) Durables	(5) Durables	(6) Durables
ΔUprob_t	-0.13	-0.30	-0.31	-1.46**	-2.42***	-0.81
	(0.23)	(0.26)	(0.25)	(0.68)	(0.82)	(0.86)
Char. variables	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Financial variables	Yes	Yes	Yes	Yes	Yes	Yes
Income cutoff	0.3	0.4	0.5	0.3	0.4	0.5
R^2	0.32	0.33	0.33	0.07	0.07	0.09
N	5526	4574	3270	5526	4574	3270

	(7) Furniture	(8) Furniture	(9) Furniture	(10) Vehicles	(11) Vehicles	(12) Vehicles
ΔUprob_t	-1.45*	-2.36**	-0.62	-1.50	-2.83	-0.87
	(0.84)	(1.02)	(1.06)	(1.56)	(1.84)	(1.99)
Char. variables	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Financial variables	Yes	Yes	Yes	Yes	Yes	Yes
Income cutoff	0.3	0.4	0.5	0.3	0.4	0.5
R^2	0.07	0.07	0.09	0.09	0.09	0.10
N	5526	4574	3270	5526	4574	3270

Table 12: *Regression results from estimating* (12) *and* (13)*, when varying the income cutoff in the sample selection. For columns (4)-(12), the left-hand side variable is an indicator of whether a durable good has been purchased in the last 12 months. The coefficients show the average marginal effect normalized by the unconditional purchase probability, estimated with a probit model. In columns (1)-(3), the left-hand side variable is the log growth of expenditures on nondurable goods, estimated by OLS. For explanation of the control variables, we refer to the text in Section 4.2. *, **, *** indicate that the coefficients are significant at the 10%, 5% and 1% level, respectively.*

and hence, that the employment status in period $t + 2$ does not have a meaningful stochastic component, but only masks individual private information. If this is the case, the estimated coefficient of unemployment probability growth on expenditures should be driven by those households that actually do become unemployed in period $t + 2$. Therefore, we rerun the regressions excluding households that became unemployed in period $t + 2$ from our sample. The results are shown in Table 13. As seen, this adjustment does not affect the results.

	(1)	(2)	(3)	(4)	(5)	(6)
	Durables	Durables	Durables	Nondurables	Nondurables	Nondurables
$\Delta \text{Uprob}_\text{t}$	-2.42***	-2.45***	-2.64***	-0.30	-0.32	-0.28
	(0.82)	(0.82)	(0.84)	(0.26)	(0.26)	(0.26)
Char. variables	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Financial variables	Yes	Yes	Yes	Yes	Yes	Yes
Future inc. growth	No	Yes	No	No	Yes	No
Excluding unemployed	No	No	Yes	No	No	Yes
R^2	0.07	0.07	0.07	0.33	0.33	0.33
N	4574	4571	4461	4574	4571	4461

	(7)	(8)	(9)	(10)	(11)	(12)
	Furniture	Furniture	Furniture	Vehicles	Vehicles	Vehicles
$\Delta \text{Uprob}_\text{t}$	-2.36**	-2.39**	-2.60**	-2.83	-2.91	-2.92
	(1.02)	(1.02)	(1.04)	(1.84)	(1.84)	(1.87)
Char. variables	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Financial variables	Yes	Yes	Yes	Yes	Yes	Yes
Future inc. growth	No	Yes	No	No	Yes	No
Excluding unemployed	No	No	Yes	No	No	Yes
R^2	0.07	0.07	0.07	0.09	0.09	0.09
N	4574	4571	4461	4574	4571	4461

Table 13: *Regression results from estimating* (12) *and* (13), *when adding future income growth as a control and excluding households that do become unemployed in year* $\text{t} + 2$. *In columns (1)-(3) and (7)-(12), the left-hand side variable is an indicator of whether a durable good has been purchased in the last 12 months. The coefficients show the average marginal effect normalized by the unconditional purchase probability, estimated with a probit model. In columns (4)-(6), the left-hand side variable is the log growth of expenditures on nondurable goods, estimated by OLS. For explanation of the control variables, we refer to the text in Section 4.2. *, **, *** indicate that the coefficients are significant at the 10%, 5% and 1% level, respectively.*