

Time-Varying Idiosyncratic Risk and Aggregate Consumption Dynamics

Alisdair McKay

Boston University

April 2015

Abstract

This paper presents an incomplete markets business cycle model in which idiosyncratic risk varies over time in accordance with recent empirical findings. The model's

1 Introduction

Recent empirical studies using large panel datasets on individual earnings portray recessions as times when households face substantially larger downside risks to their earnings prospects. Moreover, these risks appear to have highly persistent effects on household earnings. Davis and von Wachter (2011) show that earnings losses from job-displacement are large, long-lasting, and roughly twice as large when the displacement occurs in a recession as opposed to an expansion. The differential impact of displacement in a recession is evident even twenty years after the event occurred. Similarly, Guvenen et al. (2013) show that the distribution of *ve*-year earnings growth rates displays considerable pro-cyclical skewness meaning severe negative events are more likely in a recession. According to this empirical evidence, recessions are times when workers face considerably more risk to their long-term earnings prospects.

The purpose of this paper is to

on the dynamics of aggregate consumption. Finally, the paper concludes with Section 5.

2 Model

I analyze a general equilibrium model with heterogeneous households and aggregate uncertainty. At the aggregate level, the model is similar to that of Krusell and Smith (1998). At the microeconomic level, I incorporate time-varying idiosyncratic risk with an income process similar to the one estimated by Guvenen et al. (2013).

2.1 Population, preferences and endowments

The economy is populated by a unit mass of households. Households survive from one period to the next with probability $1 - \delta$ and each period a mass δ (0; 1) of households is born leaving the population size unchanged. At date 0, a household seeks to maximize preferences given by

$$E_0 \sum_{t=0}^{\infty} \beta^t (1 - \delta)^t \frac{C_t^{1-\sigma}}{1-\sigma};$$

where C_t is the household's consumption in period t . I allow for different rates of time preference across households in order to generate additional heterogeneity in wealth holdings.

Households can be either employed ($n = 1$) or unemployed ($n = 0$) and transition between these two states exogenously. Let λ and ρ be the job-finding and -separation rates, respectively. Let $u \in [0; 1]$ be the unemployment rate.

If employed, a household exogenously supplies e efficiency units of labor, where y is the household's individual efficiency. The cross-sectional dispersion in efficiency units could be due to differences in wage or due to differences in hours. For lack of a better term, I will refer to y as "skill." This skill evolves according to

$$y_t = \alpha + \epsilon_t; \\ \epsilon_t = \rho + \eta_t,$$

where ϵ_t is a transitory shock distributed $N(0; \sigma^2)$. I choose the constant parameters of the distribution for ϵ_t such that $E[\epsilon_t] = 0$. η_t is a permanent shock to the individual's skill.

The model assumes that agents learn in period t what the distribution of shocks will be between t and $t + 1$. This is an important point because it allows the households time to react to this news about risk.

As the three idiosyncratic labor income shocks ϵ , η , and n are independent, using a law of large numbers the aggregate labor input is

$$L = E[e^{\epsilon} e^{\eta} e^{\eta} (1 - u)] = 1 - u:$$

It would be natural to assume that there is a correlation between shocks to skill and shocks to employment. I have experimented with including such a correlation and found that it has little impact on the results. Intuitively, if households are well self-insured against unemployment risks then the existence of this risk is not important to their consumption behavior and therefore the correlation of this risk with other risks is not important.

It is important that the model includes mortality risk, which allows for a finite cross-sectional variance of skills despite the fact that innovations to skills are permanent. When a household dies, it is replaced by a newborn household with no assets and skill, normalized to one. The unemployment rate among newborn households is the same as prevails in the surviving population at that date. A household's rate of time preference is fixed throughout its life and drawn initially from a stable two-point distribution.

2.2 Technology, markets, and government

A composite good is produced out of capital and labor according to

$$Y = e^{\zeta} K^{\alpha} L^{1-\alpha} \tag{4}$$

where ζ is an exogenous total factor productivity (TFP) and aggregate quantities are denoted with a bar. Capital depreciates at rate δ and evolves according to

$$C + K^0 = Y + (1 - \delta)K:$$

The factors of production are rented from the households each period at prices that satisfy the representative firm's static profit maximization problem

$$W = (1 - \tau) e^z K^\alpha L^{1-\alpha} \quad (5)$$

$$R = e^z K^{\alpha-1} L^{1-\alpha} + 1 - \delta \quad (6)$$

Here R is the return on capital and W is the wage paid per efficiency unit. Households save in the form of annuities and the return to surviving households is $R = (1 - \delta)$. I assume that savings must be non-negative due to borrowing constraints. Given the income process, in which the shocks to log-income are unbounded, the zero borrowing limit is the natural borrowing limit.

The data reported by Guvenen et al. (2013) refer to pre-tax earnings. As taxes and transfers provide insurance against idiosyncratic risks it is important to incorporate this insurance into the model. Let the net tax payment of an employed individual with earnings $W e^y$ be $W e^y (1 - \tau) e^{(1 - b^y)y}$. The parameters τ and b^y control the level and progressivity of the tax, respectively. For incomes less than $(1 - \tau)^{1/b^y}$ the average tax rate is negative and the household receives a transfer from the government. Heathcote et al. (2014) discuss the properties of this type of tax system in detail.

Unemployed households receive taxable unemployment insurance payments with a replacement rate b^u . The post-government income of a household with employment status $s \in \{0, 1\}$ and skill e^y is therefore

$$(1 - \tau) W e^{(1 - b^y)y} [n + b^u (1 - n)] \quad (7)$$

I assume the level of the tax system, τ , is adjusted to balance the budget of the tax and transfer system period by period, which requires

$$1 - \tau = \frac{1 - u}{Q(1 - u + b^u u)} \quad (8)$$

where $Q = E e^{y(1 - b^y)}$ reflects the fact that a progressive income tax raises more revenue when

incomes are more dispersed. As explained in Appendix A, evolves according to

$$Q^0 = (1 - \beta)Q + \beta Q^0 \quad (9)$$

where $Q = E[e^{(1-\beta)y}]$ and $Q^0 = E[e^{(1-\beta)y^0}]$.

2.3 Aggregate shock processes

I assume the following processes for aggregate shocks. TFP evolves according to

$$z^0 = \rho_z z + \epsilon_z^0 \quad (10)$$

For the labor market, I assume that aggregate shocks occur at the start of a period and labor market outcomes in period t reflect the shocks realized at date t . I assume that the unemployment rate and job-finding rate follow AR(1) processes with correlated innovations. Specifically,

$$\hat{u}^0 = (1 - \rho_u)\hat{u} + \epsilon_u^0 \quad (11)$$

$$\hat{\lambda}^0 = (1 - \rho_\lambda)\hat{\lambda} + \epsilon_\lambda^0; \quad (12)$$

where \hat{u} is the inverse-logistic transformation⁶ of the unemployment rate and $\hat{\lambda}$ is similarly defined. \hat{u} and $\hat{\lambda}$ are constant parameters that determine the mean unemployment and job-finding rates, respectively. The job-separation rate, ρ_λ , is determined implicitly by the law of motion

$$u^0 = (1 - \rho_\lambda)u + \rho_\lambda(1 - u); \quad (13)$$

The process for skill risk, x , follows

$$x^0 = \rho_x x + \epsilon_x^0; \quad (14)$$

where the innovations, ϵ_x , are correlated with ϵ_u and ϵ_λ .

⁶That is, u and \hat{u} are related according to $u = 1/(1 + e^{-\hat{u}})$.

2.4 The household's decision problem

The individual state variables of the household's decision problem are its cash on hand, call it A , its permanent skill, θ , and its employment status, n . Households also differ in their rates of time preference although these are not state variables as they are fixed within a household's lifetime. The aggregate states are $\{z; u; x; g\}$, where z is the distribution of households over the state space from which one can calculate aggregate capital, Q , and the unemployment rate, u .

Symbol	Description	Value
	Risk aversion	2
	Capital share	0.36
	Depreciation rate	0.02
z	Persistence of TFP	0.96
z	St. dev. of TFP innovation	0.0081
δ	Mortality rate	0.005
b^u	Unemployment insurance replacement rate	0.30
b^y	Tax-and-transfer progressivity	0.151
low	Discount factor	0.96645
high	Discount factor	0.98865
μ_2	Mean of right tail of distribution	0.355
μ_3	Mean of left tail of distribution	-0.298
σ_1	St. dev. of center of distribution	0.0143
σ_2	St. dev. of right tail of distribution	0.1041
σ_3	St. dev. of left tail of distribution	0.1041
	St. dev. of transitory income shock	0.1580
p_1	Weight of center of distribution	0.8948
p_2	Weight of right tail of distribution	0.0526
p_3	Weight of left tail of distribution	0.0526

Table 1: Calibrated parameter values.

V , and policy rule, F , and pricing functions W and R . In an equilibrium, V and F are optimal for the household's problem, $R = \mathbb{R} = (1 - \delta)$ and W satisfy (5)-(6), and H is induced by F and the idiosyncratic income process.

3 Parameters and computation

I begin by describing the calibration of the income process before turning to the other parameters of the model and finally the computational methods.

3.1 The idiosyncratic income process

Calibrating the model requires an empirical counterpart to the variable α_t in the model, which changes the distribution of idiosyncratic risk and I construct this using a simulated method of moments procedure. The empirical moments describe the year-by-year distribution of one-year, three-year, and five-year earnings changes reported by Guvenen et al. (2013). While the

Guvenen et al. data is available at an annual frequency, business cycles are typically analyzed at the quarterly frequency. Therefore I use the Guvenen et al. data to construct a quarterly time series for x_t . The assumption underlying my approach is that developments in the labor market drive both x_t and observable indicators of labor market conditions that are available at a quarterly frequency. I use four such indicators: the ratio of short-term unemployed (fewer than 15 weeks) to the labor force, the same ratio for long-term unemployed (15 or more weeks), an index of average weekly hours, and the labor force participation rate. Note that the employment-population ratio can be expressed as a function of these variables. I then posit that x_t is a linear combination of these four series with factor loadings to be determined. After these factor loadings are determined I use them to construct a quarterly sequence x_t from the quarterly labor market indicators. In addition to the factor loadings, I simultaneously search for values for $\rho_2, \rho_3, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3$, and γ while imposing the restrictions $\rho_3 = \rho_2$ and $\alpha_2 = \alpha_3$.

For each candidate parameter vector, I simulate the income process for a panel of households including employment and mortality shocks and form an objective function that penalizes the distance between the model-implied moments and the empirical moments. The moments I seek to match are the year-by-year values for the median, 10th percentile and 90th percentile of the one-year, three-year and five-year earnings growth distributions. The Guvenen et al. data range from 1978 to 2011 and in total there are 279 moments.

To simulate the model, I use the following parameters: $\rho_2 = 0.9701$, $\rho_3 = 0.9701$, $\alpha_2 = 0.6828$, $\alpha_3 = 0.6828$, $\beta_1 = 1.793$, $\beta_2 = 1.793$, $\beta_3 = 1.793$, and $\gamma = 11.95$.

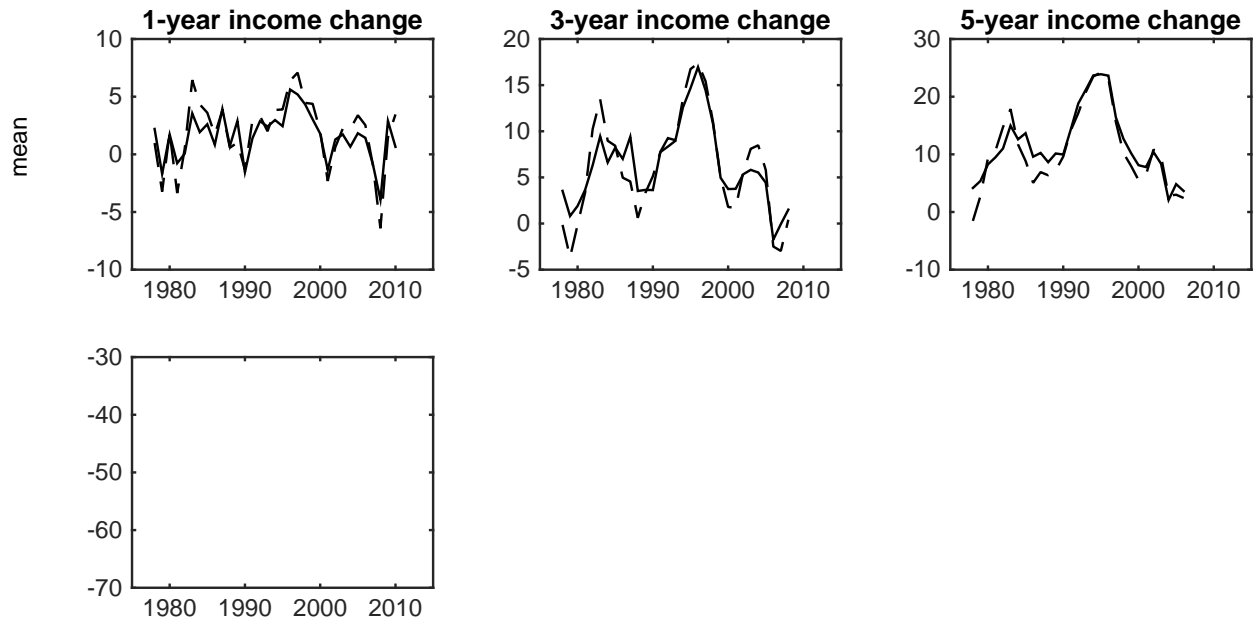


Figure 1: Simulated (dark line) and empirical (light line) moments of the earnings process.




Figure 2: PDF for distribution of h for $x = 0$ and $x = 0:2$.

of matching the moments of the three-year and v -year earnings changes. While the model fails to generate the volatility of the 10th and 90th percentiles for one-year changes, this is not too worrisome as the three-year and v -year earnings changes are a better reflection of long-term earnings risks that are of particular interest here.

The left panel of Figure 2 shows the PDF of h for $x = 0$. There is a large mass near zero and dispersed tails. The right panel of Figure 2 shows the effect of an increase in x to 0.2 on the distribution of h .

X_t

where D is a diagonal matrix with diagonal elements $[0.965, 0.945, 0.801]$ and the decomposed covariance matrix of ϵ is

$$\Sigma_{\epsilon} = \begin{pmatrix} 0 & & & \\ 0.0033 & 0 & 0 & \\ 0.0626 & 0.0563 & 0 & \\ 0.0364 & 0.0263 & 0.0594 & \end{pmatrix}$$

3.2 Other parameters

The coefficient of relative risk aversion is set to 2, the depreciation rate is set to 2 percent per quarter. I set the persistence of the productivity process to 0.96 in line with typical estimates for the US. The labor share is set to 64 percent and the mortality risk is set to 0.5 percent per quarter for an expected working lifetime of 50 years.

I set the unemployment insurance replacement rate β^u , to 0.3, which is in line with replacement rates for the United States reported by Martin (1996). The skill insurance parameter β^s is set to 0.151, which is the progressivity of the tax-and-transfer system estimated by Heathcote et al. (2014) to fit the relationship between pre- and post-government income in PSID data.⁸

I assume that there are two values of θ_i in the population with 80 percent of the population having the lower value and 20 percent having the higher value. I choose these values, and the volatility of the productivity process to match the following moments in an internal calibration: a capital-output ratio of 3.32, the wealth share of the top 20 percent by wealth equal to 83.4 percent of total wealth (see Diaz-Gimenez et al., 2011), and the standard deviation of log output growth equal to 0.0084. The resulting parameter values appear in Table 1.

The model generated distribution of wealth appears in Table 2. The baseline model does an excellent job of matching the data all along the Lorenz curve including the holdings of the very rich. That the model can generate extremely wealthy households is partially due to

⁸Those authors discuss the fact that the tax-and-transfer system became more progressive during the Great Recession. Whether or not this time-varying insurance is important depends on how constrained households are. If households are unconstrained, the precautionary savings motive is driven by changes in the households entire future earnings path. As the shocks to earnings that arise during the recession have long lasting effects, what is particularly relevant is the degree of insurance over the household's consumption of log

	Share of wealth by quintile					and held by richest			Gini
	1st	2nd	3rd	4th	5th	10%	5%	1%	
Baseline	0.01	0.02	0.04	0.09	0.84	0.70	0.56	0.30	0.79
Common-	0.03	0.07	0.11	0.18	0.60	0.45	0.34	0.17	0.56
Data	0.00	0.01	0.05	0.11	0.83	0.71	0.60	0.34	0.82

Table 2: Distribution of wealth. Data refer to net worth from the 2007 Survey of Consumer Finances as reported by Diaz-Gimenez et al. (2011).

preference heterogeneity as shown by the comparison with the second row of the table in which all households have the same rate of time preference. Even without preference heterogeneity, however, some households accumulate large wealth positions by virtue of good luck in their income draws coupled with a strong precautionary motive. In this regard the model has some similarity to that of Castaneda et al. (2003) where large wealth positions result from large income shocks. The model implies a distribution of earnings that is somewhat more dispersed than found in the data. The Gini index for earnings is 0.69 as compared to 0.64 in the Survey of Consumer Finances.

3.3 Computation

The model presents two computational challenges. First, the aggregate state of the model includes the endogenous distribution of households over individual states. I use the Krusell-Smith algorithm and replace this distribution with the first moment for capital holdings, K , the unemployment rate, u , and the measure of income inequality, Q . The aggregate state is then $S_t = \{z; u; K; \beta; \gamma; x; Q\}$, which is seven continuous variables. The second computational challenge is the curse of dimensionality as the model includes seven aggregate states, three individual states and four aggregate shocks. To compute solutions to the household's problem efficiently, I make use of the algorithm introduced by Judd et al. (2012) to construct a grid on the part of the aggregate state space that the system actually visits. This approach reduces the computational cost of having many state variables while still allowing for accurate solutions

⁹There are three individual states as opposed to four because the household's problem is homogeneous in $\exp(\beta)g$ so it is sufficient to normalize cash on hand by this value and eliminate one state. See Appendix C.

by avoiding computing the solution for combinations of states that are very unlikely to arise

raises the volatility of consumption growth and reduces the correlation of consumption and income growth. The increase in consumption volatility is 28 percent of the complete markets volatility. Moreover, time-varying idiosyncratic risk greatly reduces the correlation of output growth and consumption growth. Both the increase in consumption volatility and the decrease in the correlation of output and consumption growth reflect the fact that time-varying risk is an additional source of consumption volatility that is imperfectly related to changes in aggregate income.

While consumption growth is more volatile when idiosyncratic risk is time-varying, the level of consumption is slightly more stable. As risk tends to be counter-cyclical, it raises savings and investment in recessions, which actually stabilizes output and the level of consumption. These outcomes are direct implications of the aggregate resource constraint: if the resources are not consumed they are invested¹¹.

The difference between the aggregate consumption dynamics generated by the baseline and complete markets models is less pronounced if one compares levels as opposed to growth rates. For example, the standard deviation of the level of consumption is only two percent smaller than in the complete markets model. The level of consumption reflects low-frequency developments more strongly than growth rates do. In particular, as the extent of idiosyncratic risk appears to spike in recessions and quickly recede to more normal levels¹² as shown in Figure 3 its effects on the level of consumption are shorto yalWhileeo ya2l of consumption ns

	Y	C		Y	C	Y; C	Y;C
	Relative						
(i) Baseline	0.834	0.439	1.493	3.965	3.110	0.760	0.947
(ii) Constant risk	0.834	0.357	1.212	3.966	3.112	0.983	0.955
(iii) Common- , constant risk	0.833	0.309	1.052	4.171	3.146	0.983	0.947
(iv) Complete markets	0.837	0.294	1.000	4.384	3.162	0.987	0.948
(v) Data	0.845	0.520	1.770	4.353	2.982	0.540	0.919

Table 3: Standard deviations () and correlations () of aggregate output (Y) and consumption (C) growth rates (denoted with) and log-levels. Standard deviations are scaled by 100. Empirical moments for log-levels refer to real GDP and consumption of non-durables and services linearly detrended.

the complete markets economy. If consumption already responds strongly to income, adding constrained agents and hand-to-mouth behavior will not make such a large difference to the overall dynamics of aggregate consumption. Comparing rows (iii) and (iv) shows that in the absence of preference heterogeneity and time-varying risk, the dynamics of the incomplete markets model are 2845 (74e1.2845 (74e7es1407v4(2845 38(incomp3er28457-1000(incy49(and)]TJ 0

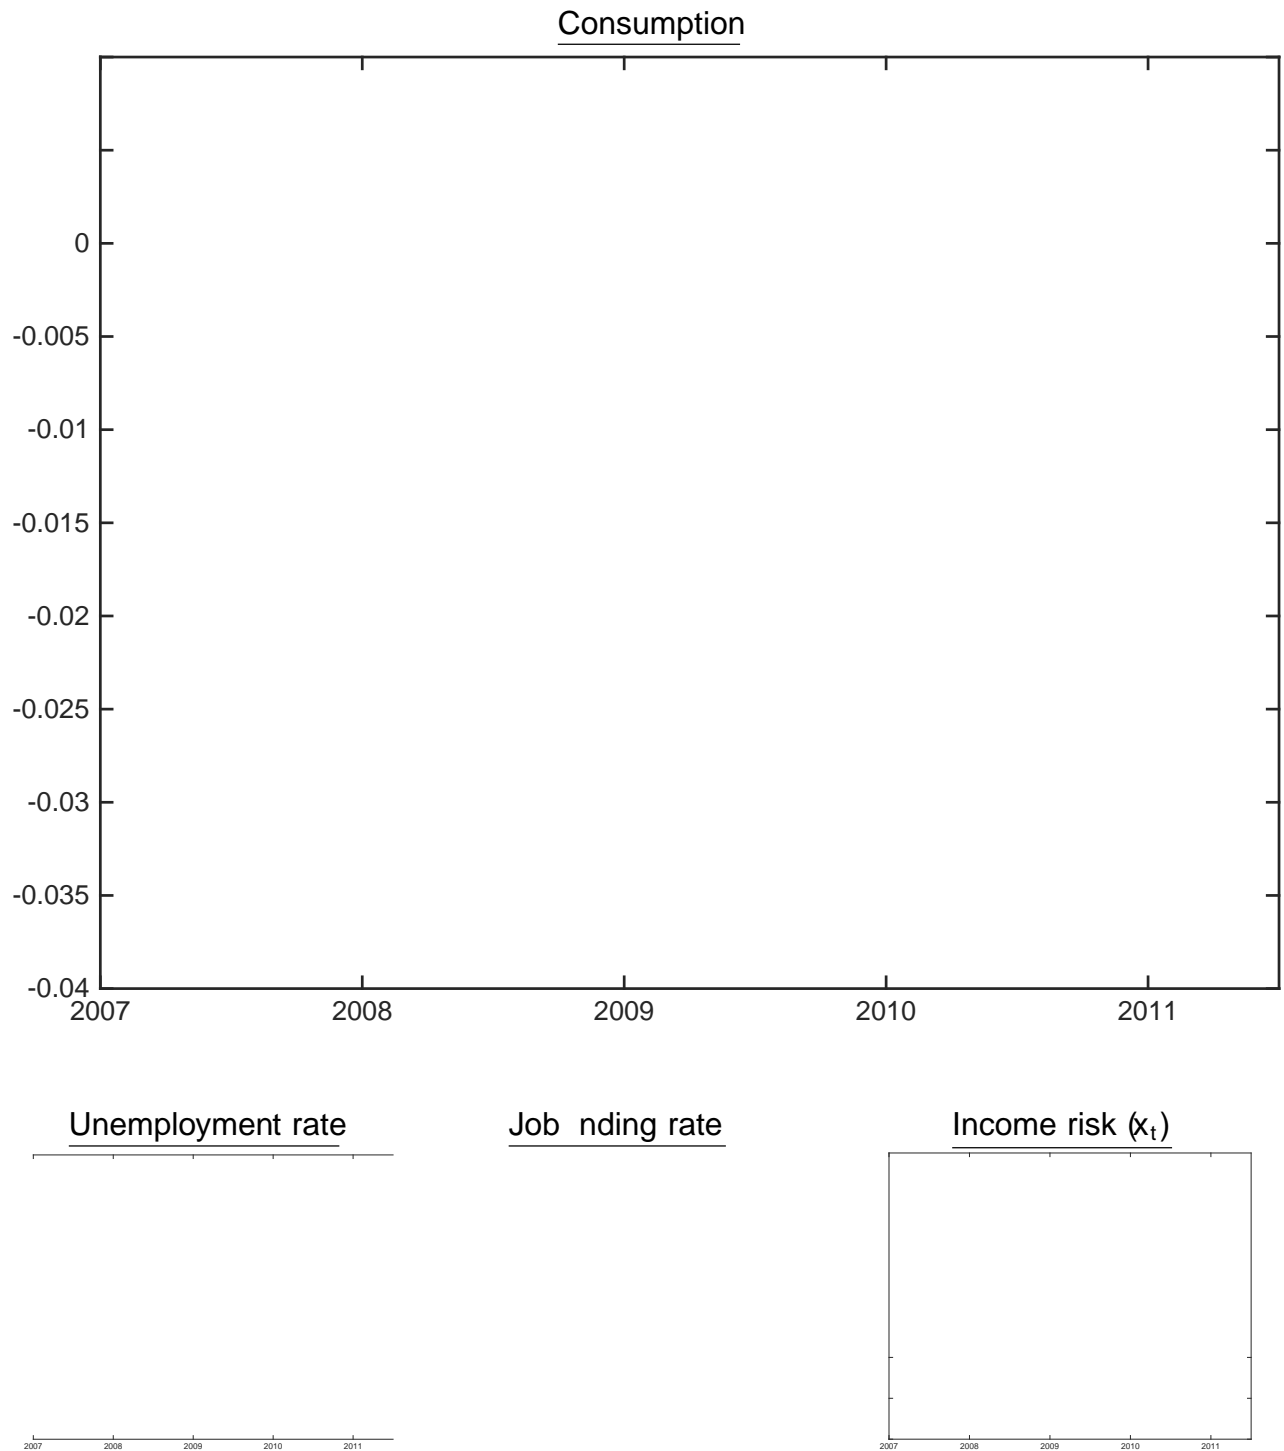


Figure 4: Dynamics of aggregate consumption implied by labor market shocks in the Great Recession. Data refer to per capita consumption of non-durables and services de ated with the GDP de ator and detrended with the HP lter with smoothing parameter 1600.

assumption about the initial condition in 2007:I, I then use equations (11), (12), and (14) to solve for sequences of u_t , c_t , and x_t . I feed these shocks into the model and report the path for aggregate consumption. I also perform the same experiment with the three benchmark models considered in Table 3.

The top panel of Figure 4 plots the path for consumption starting in 2007:II and normalized to one in 2007:IV, which was the peak of the expansion as defined by the NBER. In addition to the four versions of the model, the figure also plots the data on aggregate consumption of services and non-durable goods detrended with the HP filter.

In the data, consumption falls by 3.6 percent by 2009:I while the baseline model predicts a 3.9 percent decline. Had idiosyncratic risks remained stable over, the decline at this date would have only been 1.6 percent so time-varying risk reduced aggregate consumption by 2.3 percent in this quarter. The deterioration in the distribution of risks had similar albeit smaller effects during the latter part of 2008. From 2009:II onwards, the worst part of the recession had passed in terms of idiosyncratic risk and time-varying risk played a smaller role.

There is also a notable difference between the predictions of the constant risk model and the model that has both constant risk and a single rate of time preference. These differences reflect the stronger relationship between consumption and current income in the model with time-preference heterogeneity. In particular, with preference heterogeneity, the path for aggregate consumption more strongly reflects the path for the unemployment rate, which rises steadily throughout the recession and remains elevated in 2010 and 2011.

Overall, the changes in the distribution of idiosyncratic risks appears to have contributed substantially to the decline in aggregate consumption at the start of the Great Recession when risk was elevated.

these shocks are highly-persistent they are difficult to self-insure and even wealthy households are sensitive to changes in these risks. The results show that time-varying idiosyncratic risks substantially raises the volatility of aggregate consumption growth and played a major part in generating the decline in aggregate consumption during the Great Recession.

This paper has focussed on the dynamics of aggregate consumption. At the aggregate level, the model is a version of the flexible-price real business cycle model with exogenous labor supply and as a result an increase in household savings necessarily leads to an increase in investment and an increase in output in future periods. Moreover, there is no endogenous

References

Abowd, J. M. and Card, D. (1989). On the covariance structure of earnings and hours changes. *Econometrica: Journal of the Econometric Society* pages 411{445.

Arellano, C., Bai, Y., and Kehoe, P. (2010). Financial markets and fluctuations in uncertainty. Federal Reserve Bank of Minneapolis Working Paper

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

Basu, S. and Bundick, B. (2012). Uncertainty shocks in a model of effective demand. Working Paper 18420, National Bureau of Economic Research.

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

Bloom, N. (2014). Fluctuations in uncertainty. *The Journal of Economic Perspectives* 28(2):153{175.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. (2012). Really uncertain business cycles. Working Paper 18245, National Bureau of Economic Research.

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

Carroll, C. D., Slacalek, J., and Tokuoka, K. (2013). Buffer-stock saving in a krusell-smith world. ECB Working Paper No. 1517. <https://www.ecb.int/press/pr/2013/09/130917en.pdf>

De Santis, M. (2007). Individual consumption risk and the welfare cost of business cycles. *American Economic Review* 97(4):1488 { 1506.

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

- Krusell, P. and Smith, Jr., A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867{896.
- Leduc, S. and Liu, Z. (2012). Uncertainty shocks are aggregate demand shocks. Federal Reserve Bank of San Francisco Working Paper 10.
- MaCurdy, T. E. (1982). The use of time series processes to model the error structure of earnings in a longitudinal data analysis. *Journal of econometrics* 18(1):83{114.
- Martin, J. P. (1996). Measures of replacement rates for the purpose of international comparisons: a note. *OECD Economic Studies* 26(1):99{115.
- Ravn, M. and Sterk, V. (2013). Job uncertainty and deep recessions. UCL Working Paper.
- Schmidt, L. (2015). Climbing and falling on the ladder: Asset pricing implications of labor market event risk. UCSD Working Paper.
- Shell, K. (1971). Notes on the economics of infinity. *The Journal of Political Economy*, 79(5):1002.
- Storesletten, K., Telmer, C. I., and Yaron, A. (2001). The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk. *European Economic Review* 45(7):1311{1339.
- Storesletten, K., Telmer, C. I., and Yaron, A. (2007). Asset pricing with idiosyncratic risk and overlapping generations. *Review of Economic Dynamics* 10(4):519{548.

Appendix

A Dynamics of Q

To calculate the dynamics of the tax adjustment Q , in equation (8) define

$$Q = E e^{(1 - b^y)}$$

$$Q = E e^{(1 - b^y)}$$

$$Q = E e^{(1 - b^y)} ;$$

where expectations are taken across agents. By the independence of the shocks one can write

$$Q = Q Q :$$

Q evolves according to

$$Q^0 = (1 - \beta) E e^{(1 - \beta)(1 - b^y) + \beta}$$

$$Q^0 = (1 - \beta) Q Q^0 + \beta :$$

And as Q is constant one can then write

$$Q^0 Q = (1 - \beta) Q Q^0 Q + \beta Q$$

$$Q^0 = (1 - \beta) Q Q^0 + \beta :$$

B Calibrating the idiosyncratic income process

This appendix provides additional information on the simulated method of moments procedure used to select the parameters of the idiosyncratic income process, which is a variant of the procedure used by Guvenen et al. (2013).

Step 1. Calculate τ_t and τ_t implied by the data. To do so, use the data on short-term unemployment described in Section 3 and solve for τ_t and τ_t from equations (16) and (17).

Step 2. Construct the four labor market indicators. I use four such indicators: the ratio of short-term unemployed (fewer than 15 weeks) to the labor force, the same ratio for long-

distribution of ϵ is initialized because the objects of interest are related to earnings changes as opposed to levels. I initialize ϵ to a 7.5 percent unemployment rate, which is the value reported by the BLS for January 1977.

Step 6. Compute the moments: aggregate the quarterly earnings observations to annual observations, take 1-year, 3-year, and 5-year changes in log earnings. I use the following moments for each year and for each of the 1-year, 3-year and 5-year changes: the median, and the 10th and 90th percentiles. I express the 10th and 90th percentiles relative to the median (i.e. $50 - 10$ and $90 - 50$). Doing so implies that any differences between the simulated and empirical medians do not change the targets for the widths of the upper and lower tails.

Step 7. Compute the objective function: I take the difference between the simulated moment and the empirical moment from Table A13 in Guvenen et al. (2013). The differences are expressed as squared percentage differences except for the difference in medians, which is expressed relative to the 90th percentile as in Guvenen et al. (2013).

Step 8. Adjust the guess in step 3 and repeat to minimize the objective function from step 7.

As an additional check on the calibrated income process, I compute the standard deviations of the income changes and compared those to the results in Guvenen et al. (2013). Figure 5 shows that the simulated standard deviations are only slightly cyclical while those in the data are more or less acyclical. The simulated standard deviations are somewhat below the observed values.

C Equilibrium conditions

Due to the progressive tax system, a household with skill i has income proportional to $\omega_i^{(1-b^y)}$.




Figure 5: Simulated and empirical standard deviations of the income process.

to denote household variables relative to $e^{(1-b^y)_i}$:

$$c_i = \frac{C_i}{e^{(1-b^y)_i}}; \quad a_i = \frac{A_i}{e^{(1-b^y)_i}}; \quad k_i^0 = \frac{K_i^0}{e^{(1-b^y)_i}}$$

The household's Euler equation and budget constraint are

$$c_{i,t} = \beta E_t [R_{t+1} c_{i,t+1}]$$

$$c_{i,t} + k_{i,t} = R k_{i,t-1} + (1-\delta) W_t e^{(1-b^y)y_{i,t}} [n_{i,t} + b^u(1-n_{i,t})]$$

and in terms of normalized variables these are

$$c_{i,t} = \beta E_t [e^{-(1-b^y)c_{i,t+1}} R_{t+1} c_{i,t+1}] \tag{A2}$$

$$c_{i,t} + k_{i,t} = R k_{i,t-1} e^{-(1-b^y)c_{i,t}} + (1-\delta) W_t e^{(1-b^y)c_{i,t}} [n_{i,t} + b^u(1-n_{i,t})] \tag{A3}$$

The remaining equations needed to solve the model are: (1), (2), (3), (5), (6), (8), (9), (10), (11), (12), (13), (14), and (A1). These are 13 equations in the 14 variables $s_{i,t}$, $z_{i,t}$, $u_{i,t}$, $x_{i,t}$, $Q_{i,t}$, W_t , R_t , and K_t . Closing the model requires determining the aggregate

capital stock K . Following Krusell and Smith (1998) this is done in two ways. In solving the household's decision problem, I make use of a forecasting rule

$$K^0 = h(z; \hat{\alpha}; \hat{\alpha}_1; x; K; Q): \tag{A4}$$

I assume that $h(\cdot)$ is a complete second-order polynomial. In simulating the model, K is determined according to the household decision rules and the dynamics of the distribution of wealth in line with equation (15). To express (15) in normalized terms, note that equations

D Numerical methods

D.1 Method for Section 4

Overview I solve the model using the Krusell-Smith algorithm, which involves solving the household's problem for a given law of motion for the capital stock and updating this law of motion through simulation and least squares curve fitting. For a given law of motion, I solve the household's problem using a projection method on a grid that is constructed from simulated data generated by a guess of the model solution in the manner described by Judd et al. (2012). This requires alternating between solving the decision problem given a grid and simulating the solution and updating the grid. The steps of the algorithm are as follows:

1. Guess household decision rules and a forecasting rule for the aggregate capital stock.
2. Simulate the economy and record aggregate states.
3. Use simulated data to construct a grid for the aggregate state space.
4. Solve the household's decision problem on the grid.
5. Simulate the economy and record aggregate states.
6. Use simulated data to construct a grid for the aggregate state space.
7. If the grid has converged then continue, otherwise return to step 4.
8. Update the forecasting rule with least-squares regression.
9. If the forecasting rule has converged stop, otherwise return to step 4.

Initial guesses A good initial guess is important to the success of this algorithm because a poor guess will lead to a situation in step 5 where the economy is being simulated far from the grid on which the problem was solved. In most cases I have found it sufficient to use the linearized solution for the representative agent model as a starting point. The representative agent's policy rule can be simulated to provide the data for the initial grid and this policy can also serve as a decent guess for the forecasting rule. The success of this guess is premised on the difference between the representative agent and incomplete markets economies being

limited. This is not the case for the baseline economy and this guess is not sufficient for this case. Instead, I found it necessary to gradually build up an initial guess based on versions of the model that are more similar to the representative agent model. I gradually lowered the rate of time-preference of the less patient group to generate this guess.

Constructing the grid See Judd et al. (2012). I target a grid with 45 points. As I explain

approximating these functions, I update the coefficients of the polynomials by least-squares projection.

To compute expectations with respect to aggregate shocks, I use the monomial rule with $2N$ nodes described by Judd et al. (2012). To compute expectations over idiosyncratic shocks I use Gaussian quadrature. Of particular interest is the shock because this has a time-varying distribution. I use Gaussian quadrature with five points in each tail and three points for the central mixture component. As it is only the means of the distributions that are moving with x and not the variance of the mixture components, I construct fixed quadrature grids for each component and shift their locations according to α . For the transitory shock, ϵ , I use Gaussian quadrature with three points.

Simulation and updating the law of motion In solving for the law of motion for the aggregate capital stock, I simulate a panel of 100,000 households for 5,500 quarters and discard the first 500 quarters. When drawing the idiosyncratic shocks I reduce the sampling error by, at each period, requiring the cross-sectional average of idiosyncratic productivities to equal the theoretical value of 1 within both the employed and unemployed groups. Using the simulated aggregate capital stock, I update the law of motion with a least squares regression using the same functional form as for the household decision rules (a complete second-order polynomial in the aggregate state). For computing the moments in Table 3, I simulate a panel of 7.2 million households as described in Footnote 10.

Accuracy of the law of motion for capital To assess the accuracy of the law of motion for the capital stock, Figure 6 shows a plot of the capital stock generated from simulating the model and the approximate capital stock generated by repeatedly applying the approximate law of motion for capital.¹⁴ This is one sample path of shocks for 1000 quarters and the discrepancy between the two lines is the forecast error that the agents are making at different horizons. One can see that the discrepancy is small even at forecast horizons of 1000 quarters. The maximum absolute log difference between the two series is 0.0054153 and the mean absolute log difference is 0.0029717. Another commonly-reported accuracy check is the R^2 of

¹⁴As den Haan (2010) suggests, the sequence of shocks used to simulate the model for the accuracy check differ from those used to calculate the approximate law of motion.




Figure 6: Simulated aggregate capital stock with implied values from $\log^0 = h(S)$.

the one-step ahead forecast, which is $11 \cdot 10^5$.

Accuracy of the policy rules There are several sources of error in the approximate solution. First, there is the error introduced by the discrepancy between the forecasting rule and the actual dynamics of the aggregate capital stock. Second, there are errors associated with the projection method that arise between grid points when the function being approximated is not of the same form as the approximating function.

To assess the accuracy of the solution, I calculate unit-free Euler equation errors¹⁵For a given state of the economy S , the distribution of wealth, the capital stock, and exogenous variables are predetermined.

Pre-determined and exogenous: $K; z; u; u_{-1}; x; Q; \dots$

is generated by simulating a panel of households. I then use the computed solution to

¹⁵See Judd (1992) for an explanation of this accuracy check and the interpretation of the errors in terms of bounded rationality.

determine the household decision rules

$$\text{Approx. solutions: } a_{ij}(n; \cdot; S) = \delta_j; n; \cdot$$

Using these policy rules, one can compute the savings of each household and then aggregate to find K^0 by integrating against \cdot . For a given set of aggregate shocks one can then compute S^0 from (9), (10), (11), (12), and (14). Given S^0 compute $a_{ij}(n; \cdot; S^0)$. The Euler equation error is then

$$(1 - \beta)E e^{-(1-\beta'v)}$$

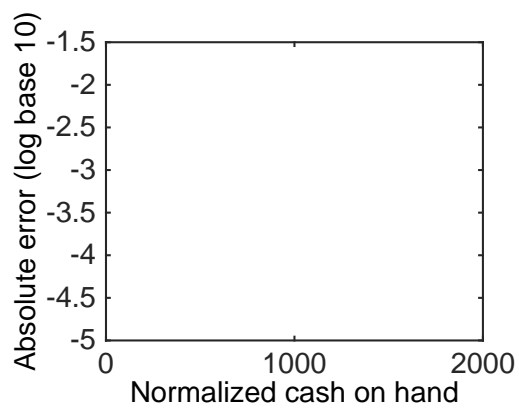


Figure 7: Euler equation errors. Left column: unemployed; right column: employed; top row: less patient; bottom row: more patient. Maximum and mean across 100 aggregate states.

Solving for the policy rule under complete markets For the complete markets model I use the algorithm described in Judd (1992) that iterates on the Euler equation. I again use a complete second-order polynomial for the savings policy rule.

E Complete markets model

This appendix derives the representative agent Euler equation from the environment presented in Section 2 augmented with a complete set of contingent securities. I assume that trade takes place at an initial period prior to date 0 before any uncertainty has been resolved. I also assume that all households have the same rate of time-preference. Like Shell (1971), I assume that all current and future generations meet and trade in this initial period. Let $l_{i,t}$ take the value 1 if household i is alive in period t and zero if it is not. I will treat birth and death as random events against which the household can insure. Specifically, let s^t be a history of stochastic events up to date t the probability of which is $\pi_t(s^t)$. These stochastic events dictate the evolution of all idiosyncratic as well as aggregate developments. Let $p_t(s^t)$ be the date-0 price of a unit of the final good at date t and history s^t . The household's utility function is

$$\sum_{t=0}^{\infty} \beta^t \sum_{s^t} l_{i,t} u(c_{i,t}(s^t))$$

The household's Lagrangian is

$$L = \sum_{t=0}^{\infty} \beta^t \lambda_t (s^t - C_{i,t})$$

Substituting for $p_t(s^t)$ from above yields

$$\begin{aligned}
 C_{i,t}(s^t) &= I_{i,t}(s^t) + \sum_{s^{t+1} | s^t} \frac{p_t(s^{t+1})}{p_t(s^t)} C_{i,t+1}(s^{t+1}) - I_{i,t+1}(s^{t+1}) + \sum_{s^{t+1} | s^t} p_{t+1}(s^{t+1}) R_{t+1}(s^{t+1}) \\
 C_t(s^t) &= \sum_{s^{t+1} | s^t} e^{\alpha_{t+1}(s^{t+1})}
 \end{aligned}$$