

Time-Varying Idiosyncratic Risk and Aggregate Consumption Dynamics

Alisdair McKay*

Boston University

December 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 income process is calibrated to match the cyclicalities of earnings risks documented by Guvenen et al. (2014). Market incompleteness raises the volatility of aggregate consumption growth by 43 percent relative to a complete markets benchmark. Half of this increase is due to the time-varying precautionary saving motive that results from changes in idiosyncratic risks over the cycle. Idiosyncratic risk spiked during the Great Recession and this contributed 2.0 percentage points to the decline in aggregate consumption.

Keywords: consumption, incomplete markets, business cycle.

*Contact: amckay@bu.edu. I would like to thank Susanto Basu, Emi Nakamura, Jón Steinsson, Ricardo Reis and seminar participants at Columbia, UCL, the Philadelphia Workshop on Macroeconomics, Federal Reserve Bank of St. Louis, NYU, the Green Line Macro meeting, and the Federal Reserve Bank of Minneapolis for helpful comments and suggestions.

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. (2014) show that the distribution of five-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 explore how the cyclical dynamics of these risks alter the precautionary savings motive and the dynamics of aggregate consumption over the business cycle. To do so, I develop a general equilibrium business cycle model with uninsurable idiosyncratic shocks to earnings. One of the aggregate shocks that drives the model is a risk shock that alters the distribution of risks that households face. This shock results in a time-varying precautionary savings motive and serves as an additional source of aggregate consumption fluctuations.

I find that time-varying idiosyncratic risk has a quantitatively substantial effect on the dynamics of aggregate consumption. In particular, the standard deviation of quarter-to-quarter aggregate consumption growth is 43 percent larger than it is in a complete markets version of the model. 21 percentage points of this difference are due to the cyclicity of idiosyncratic earnings risk. Most of the remaining difference is due to the greater sensitivity of consumption to income among low-wealth households.

In calibrating the model, I construct a quarterly time series for risk shocks that best fits the observed distribution of earnings growth rates reported by Guvenen et al. (2014). According to my constructed series for risk shocks, the Great Recession brought about a large spike in risk to household earnings. I use the model to simulate the consumption response to the deterioration of labor market conditions in the Great Recession including the changes in idiosyncratic risk. The model predicts a 3.7 percentage point drop in aggregate consumption between the NBER

peak and the the first quarter of 2009, which is close to the 3.6 percent drop in non-durable and services consumption in the data. I find that 2.0 percentage points of this drop are due to the increase in idiosyncratic earnings risks. While the model predicts that idiosyncratic risk lead to large swings in aggregate consumption during the Great Recession, it also predicts that the cross sectional dispersion of consumption remained almost completely unchanged. This prediction is in line with the empirical findings of Perri and Steinberg (2012). In this way I demonstrate that evidence of a constant cross-sectional distribution does not preclude the possibility that idiosyncratic risk and distributional shocks played an important role in aggregate dynamics.

Krusell and Smith (1998) show the business cycle dynamics of a heterogeneous agent version of the neoclassical growth model are generally close to those of the representative agent economy, however, when the model matches the distribution of household net worth aggregate consumption shows an increased correlation with aggregate income. This result is the reflection of the greater sensitivity of consumption to income among low wealth households.

Market incompleteness can affect aggregate consumption through an alternative channel besides hand-to-mouth behavior. In particular, if the uninsurable risk that households face varies over time, there will be a time-varying precautionary savings motive that can generate additional fluctuations in aggregate consumption that are disconnected from aggregate income. The existing literature that studies the contribution of uninsurable idiosyncratic income risk to the business cycle has focussed on fluctuations in the unemployment rate as a source of time-varying uninsurable risk (Krusell and Smith, 1998; Challe and Ragot, 2013; Ravn and Sterk, 2013; Challe et al., 2015; Krueger et al., 2015). However, as unemployment is generally a short-lived shock to earnings it is easily smoothed through self-insurance and only households with very low levels of savings will alter their consumption behavior in a meaningful way when the unemployment risk changes. One response is to make unemployment more painful by limiting the self insurance that households have. Challe and Ragot (2013), Ravn and Sterk (2013), and Challe et al. (2015) pursue this approach by calibrating their models such that the majority of households have little or no wealth. This approach could be justifiable on the grounds that much of household wealth takes the form of illiquid assets that cannot easily be used for consumption smoothing.

The contribution of this paper is to introduce a rich income process that incorporates cyclical variation in the risks to long-term earnings prospects as documented by Guvenen et al. (2014). As these earnings shocks are persistent they are more difficult to self-insure and even households with large amounts of savings will respond to changes in the distribution of earnings risks. Therefore, this model is able to match the distribution of net worth and still predict substantial changes in the dynamics of aggregate consumption due to changes in the precautionary savings motive.

The nature and source of cyclical changes in the earnings process are still somewhat poorly understood.¹ This paper gives a particular interpretation to the facts on the distribution of earnings changes in expansions and recessions—these changes in income are uninsured and unforeseen risks—and then goes on to consider the implications for aggregate consumption dynamics. Other implications of this type of risk have also been studied. For example, Constantinides and Duffie (1996), Storesletten et al. (2007), Constantinides and Ghosh (2014) and Schmidt (2015) investigate the asset pricing implications of this type of risk. The welfare cost of business cycles in the presence of these risks have been analyzed by Storesletten et al. (2001), Krebs (2003, 2007), and De Santis (2007). However, it does not appear that the consequences for the business cycle dynamics of aggregate quantities have been studied in the literature.

This paper is also related to the recent literature that investigates the role of uncertainty in business cycle fluctuations. In particular, Basu and Bundick (2012), Leduc and Liu (2012), and Fernández-Villaverde et al. (2013) emphasize the precautionary savings effect that follows an increase in uncertainty surrounding aggregate conditions such as preferences, technology or taxes. In contrast to those studies, the focus here is on cyclical variation in microeconomic uncertainty faced by heterogeneous households. Other studies analyze the impact of cyclical microeconomic uncertainty faced by firms. This work is motivated by evidence of countercyclical dispersion in firm-level productivity, sales growth rates, and other measures of business conditions.² Bloom (2009), Bloom et al. (2012) and Bachmann and Bayer (2013)

¹Davis and von Wachter (2011) show that structural models of the labor market have difficulty explaining the size and cyclicity of present-value earnings losses after job displacement. Huckfeldt (2014) presents a model that performs better but still struggles to explain the strong cyclicity. Jarosch (2014) shows that job insecurity at the bottom of the job ladder can explain long-term earnings losses for displaced workers, but he does not address the cyclicity of these earnings losses.

²See Bloom (2014) for a review of the evidence.

study the interaction of this microeconomic uncertainty with non-convex adjustment costs for investment and hiring. Arellano et al. (2010) and Gilchrist et al. (2014) explore the interaction of firm-level risks and financial frictions. This paper contributes to this literature by studying the importance of variations in the microeconomic uncertainty surrounding household incomes for aggregate consumption.

The paper is organized as follows: Section 2 presents the model. Section 3 discusses the choice of parameters and the construction of the time series for idiosyncratic risk. Section 4 presents the results on the impact of household heterogeneity and time-varying earnings risks 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, the model incorporates time-varying idiosyncratic risk with an income process similar to the one estimated by Guvenen et al. (2014).

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 $1 - \omega$ and each period a mass $\omega \in (0, 1)$ of households is born leaving the population size unchanged. At date 0, a household seeks to maximize preferences given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t (1 - \omega)^t \frac{C_t^{1-\gamma}}{1 - \gamma},$$

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 ζ be the job-finding and -separation rates, respectively. Let $u \in [0, 1]$ be the unemployment rate.

If employed, a household exogenously supplies e^y efficiency units of labor, where y is the

household’s individual efficiency. The cross-sectional dispersion in earnings 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

$$\begin{aligned} y &= \theta + \xi, \\ \theta' &= \theta + \eta', \end{aligned}$$

where ξ is a transitory shock distributed $N(\mu_\xi, \sigma_\xi)$. I choose the constant parameters of the distribution for ξ such that $\mathbb{E}[e^\xi] = 1$. η is a permanent shock to the individual’s skill. Assuming that this shock is permanent as opposed to persistent has the advantage that it allows one state variable to be eliminated from the household’s decision problem as described in Appendix C.³ This type of income process is known to fit longitudinal earnings data well as shown by MaCurdy (1982) and Abowd and Card (1989).

The permanent shock, η , is drawn from a time-varying distribution the tails of which vary over the business cycle in such a way to generate pro-cyclical skewness as documented in the data by Guvenen et al. (2014). I assume η is drawn from a mixture of three normals.⁴ Specifically

$$\eta' \sim \begin{cases} N(\mu_{1,t-1}, \sigma_{\eta,1}) & \text{with prob. } p_1 \\ N(\mu_{2,t}, \sigma_{\eta,2}) & \text{with prob. } p_2 \\ N(\mu_{3,t}, \sigma_{\eta,3}) & \text{with prob. } p_3, \end{cases}$$

where $\sum_{j=1}^3 p_j = 1$. The time-varying parameters $\mu_{1,t}$, $\mu_{2,t}$, and $\mu_{3,t}$ will, respectively, control the center, right tail, and left tail of the distribution. I assume that these distributional

³See also Carroll et al. (2013).

⁴Guvenen et al. (2014) estimate a parametric income process in which the shocks are drawn from a mixture of two normals with the distribution changing discretely between expansions and recessions. Here I assume that the business cycle is driven by continuous shocks rather than discrete regime switches and in this context I found that a mixture of three normals is better able to generate the cyclical skewness observed in the data than a mixture of two normals.

parameters are driven by a single stochastic process, x_t , according to

$$\mu_{1,t} = \bar{\mu}_t \tag{1}$$

$$\mu_{2,t} = \bar{\mu}_t + \mu_2 - x_t \tag{2}$$

$$\mu_{3,t} = \bar{\mu}_t + \mu_3 - x_t. \tag{3}$$

An increase in x_t moves the tails of the distribution to the left relative to the center of the distribution and will generate negative skewness in the distribution. $\bar{\mu}_t$ is a normalization such that $\mathbb{E}[e^\eta] = 1$ in all periods. This normalization in turn implies that $\mathbb{E}[e^\theta] = 1$ given suitable initial conditions.⁵ Therefore movements in x_t only affect the distribution of income and are not first moment shocks. The details of this normalization appear in Appendix B.

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— θ , ξ and n —are independent, using a law of large numbers the aggregate labor input is

$$\bar{L} \equiv \mathbb{E} [e^{\theta+\xi}n] = \mathbb{E} [e^\theta] E [e^\xi] (1 - u) = 1 - u.$$

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 as this 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 $\theta = 0$. The unemployment rate among newborn households is the same as prevails in the surviving population

⁵To verify this observe that

$$\mathbb{E}[e^{\theta'}] = (1 - \omega)\mathbb{E}[e^{\theta+\eta'}] + \omega = (1 - \omega)\mathbb{E}[e^\theta]\mathbb{E}[e^{\eta'}] + \omega = (1 - \omega)\mathbb{E}[e^\theta] + \omega,$$

which implies that $\mathbb{E}[e^\theta]$ converges to one.

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

$$\bar{Y} = e^z \bar{K}^\alpha \bar{L}^{1-\alpha} \quad (4)$$

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

$$\bar{C} + \bar{K}' = \bar{Y} + (1 - \delta)\bar{K}.$$

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

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

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

Here \tilde{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 $R \equiv \tilde{R}/(1 - \omega)$. 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. (2014) 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 We^y be $We^y - (1 - \tau)We^{(1-b^y)y}$. 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 post-tax replacement rate b^u . The post-government income of a household with employment status $n \in \{0, 1\}$ and skill 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 $Q \equiv \mathbb{E} [e^{y(1-b^y)}]$ reflects the fact that a progressive income tax raises more revenue when incomes are more dispersed. As explained in Appendix A, Q evolves according to

$$Q' = (1 - \omega)Q\tilde{Q}^{\eta'} + \omega\tilde{Q}^{\xi} \quad (9)$$

where $\tilde{Q}_\xi \equiv \mathbb{E} [e^{(1-b^y)\xi}]$ and $\tilde{Q}^\eta \equiv \mathbb{E} [e^{(1-b^y)\eta}]$.

2.3 Aggregate shock processes

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

$$z' = \rho_z z + \epsilon'_z. \quad (10)$$

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,

$$\hat{u}' = (1 - \rho_u)\hat{u}^* + \rho_u\hat{u} + \epsilon'_u \quad (11)$$

$$\hat{\lambda}' = (1 - \rho_\lambda)\hat{\lambda}^* + \rho_\lambda\hat{\lambda} + \epsilon'_\lambda, \quad (12)$$

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, ζ , is determined implicitly by the law of motion

$$u' = (1 - \lambda')u + \zeta'(1 - u). \quad (13)$$

The process for skill risk, x , follows

$$x' = \rho_x x + \epsilon_x, \quad (14)$$

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

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 $S \equiv \{z, \lambda, u_{-1}, x, \Gamma\}$, where Γ is the distribution of households over the state space from which one can calculate aggregate capital, \bar{K} , Q , and the unemployment rate, u . The lagged unemployment rate, u_{-1} is needed in order to calculate the job-separation probability from equation (13). Appendix C describes how the model can be normalized to eliminate some of these state variables. The household's decision variable is end-of-period savings, K' .

The household's decision problem is then

$$V(A, \theta, n, S) = \max_{K' \geq 0} \left\{ \frac{(A - K')^{1-\gamma}}{1-\gamma} + \beta(1 - \omega)\mathbb{E}[V(A', \theta', n', S')] \right\}$$

subject to

$$A' = R(S')K' + (1 - \tau(S'))W(S')e^{(1-b^y)(\theta+\eta'+\xi')}[n' + b^u(1 - n')].$$

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

The prices in the household's problem depend on the aggregate state S through (5) and (6). The law of motion for the aggregate state is given by (10), (11), (12), (14), $u'_{-1} = u$, and a law of motion for the distribution of idiosyncratic states, $\Gamma' = H_\Gamma(S, u', \lambda')$.⁷

2.5 Equilibrium

Let $F(A, \theta, n, S)$ be the optimal decision rule for K' in the household's problem. Aggregate savings are

$$K' = \sum_n \int_A \int_\theta F(A, \theta, n, S) \Gamma(dA, d\theta, n), \quad (15)$$

Given a set of exogenous stochastic processes for z , u , λ , and x , a recursive competitive equilibrium consists of the law of motion for the distribution, H_Γ , household value function, V , and policy rule, F , and pricing functions W and R . In an equilibrium, V and F are optimal for the household's problem, $R = \tilde{R}/(1 - \omega)$ and W satisfy (5)-(6), and H_Γ 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 x_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. (2014). 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

⁷ H_Γ depends on the labor market shocks in the next period as Γ' is the distribution of households after labor market transitions have occurred.

Symbol	Description	Value
γ	Risk aversion	2
α	Capital share	0.36
δ	Depreciation rate	0.02
ρ_z	Persistence of TFP	0.96
σ_z	St. dev. of TFP innovation	0.0081
ω	Mortality rate	0.005
b^u	Unemployment insurance replacement rate	0.30
b^y	Tax-and-transfer progressivity	0.151
β_{low}	Discount factor	0.9703
β_{high}	Discount factor	0.9892
μ_2	Mean of right tail of η distribution	0.3071
μ_3	Mean of left tail of η distribution	-0.2508
$\sigma_{1,\eta}$	St. dev. of center of η distribution	0.0143
$\sigma_{2,\eta}$	St. dev. of right tail of η distribution	0.1041
$\sigma_{3,\eta}$	St. dev. of left tail of η distribution	0.1041
σ_ξ	St. dev. of transitory income shock	0.1580
p_1	Weight of center of η distribution	0.8948
p_2	Weight of right tail of η distribution	0.0526
p_3	Weight of left tail of η distribution	0.0526

Table 1: Calibrated parameter values.

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. For a given set of factor loadings I can construct a quarterly sequence for x_t from the quarterly labor market indicators. In addition to the factor loadings, I simultaneously search for values for $p_2, p_3, \mu_2, \mu_3, \sigma_{1,\eta}, \sigma_{2,\eta}, \sigma_{3,\eta}$, and σ_ξ while imposing the restrictions $p_3 = p_2$ and $\sigma_{2,\eta} = \sigma_{3,\eta}$.

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 need estimates of λ_t and ζ_t . I estimate these from the relationships

$$u_t^s = \zeta_t(1 - u_{t-1}) \tag{16}$$

$$u_t - u_t^s = (1 - \lambda_t)u_{t-1}, \tag{17}$$

where u_t is the unemployment rate and u_t^s is the short-term unemployment rate measured as those with durations less than 15 weeks. I simulate quarterly data and then aggregate to annual observations to conform to the Guvenen et al. data. Appendix B contains further discussion of the implementation of this method.

The resulting parameters appear in Table 1. Figure 1 shows the model's fit to the earnings growth distribution at one-year, three-year and five-year horizons. The model does a good job of matching the moments of the three-year and five-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 five-year earnings changes are a better reflection of

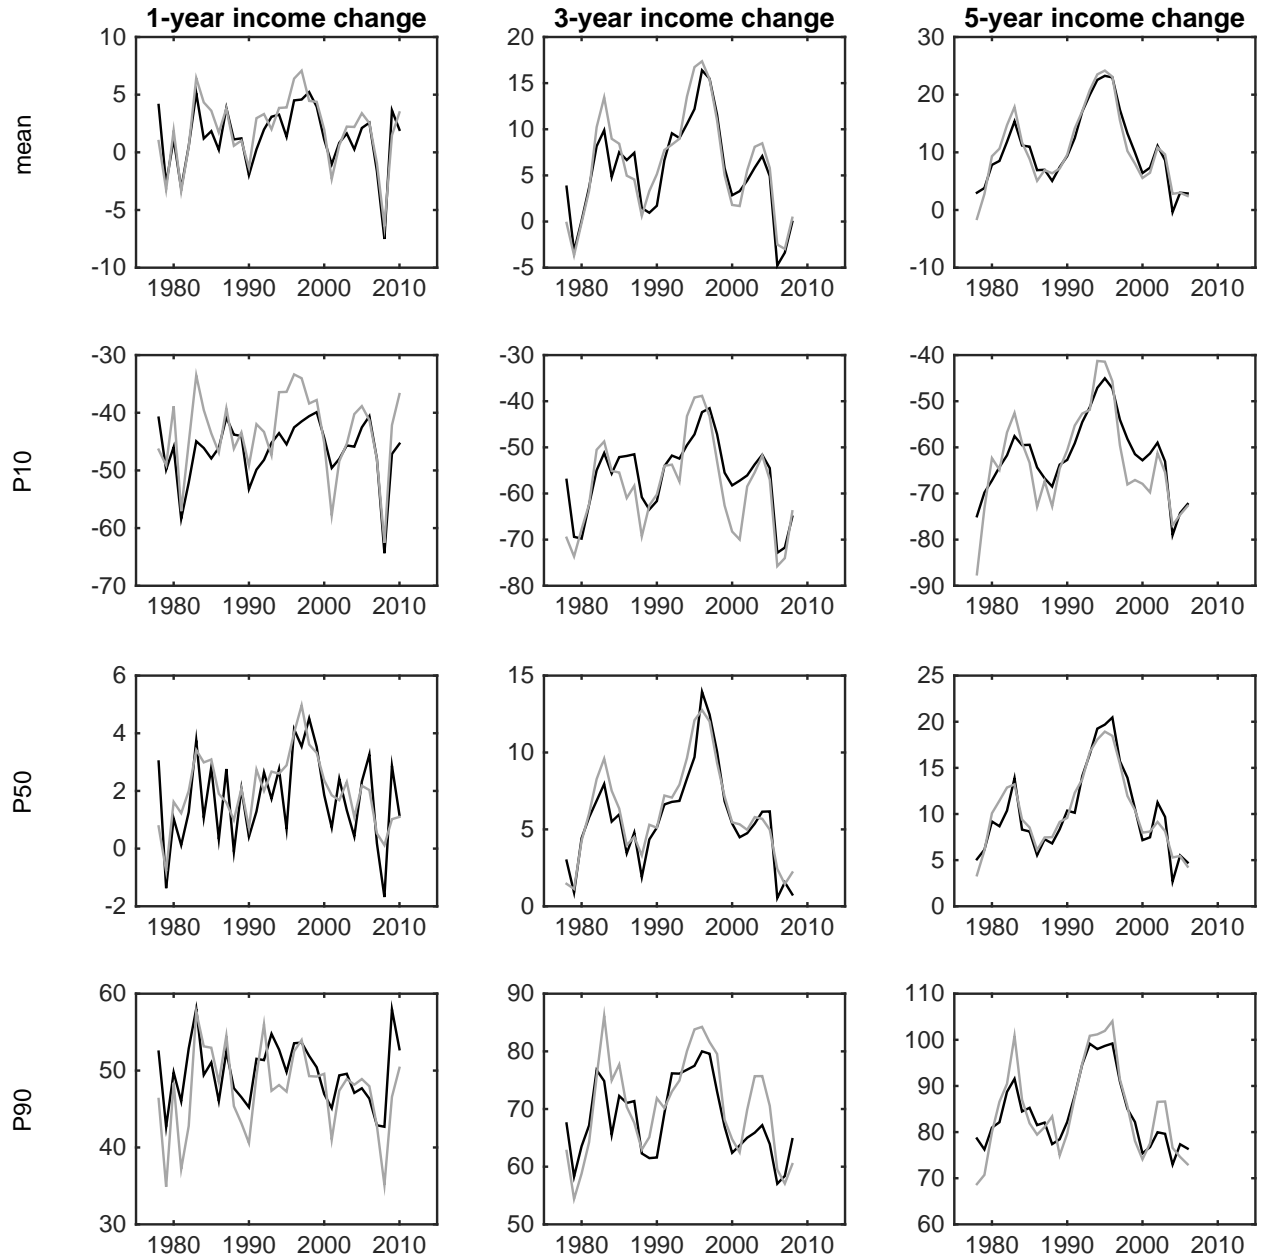


Figure 1: Simulated (dark line) and empirical (light line) moments of the earnings process.

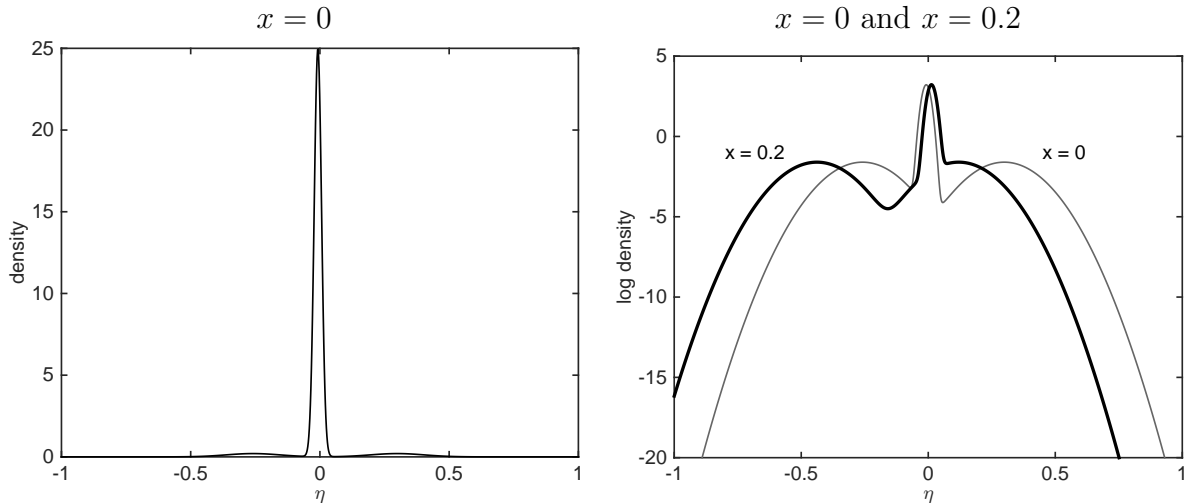


Figure 2: PDF for distribution of η .

long-term earnings risks that are of particular interest here.

The left panel of Figure 2 shows the PDF of η 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 η now in terms of the log density to emphasize the tails of the distribution. The negative skewness caused by $x = 0.2$ is evident as the left tail now shifts away from the central mass and the right tail compresses towards it.

The left panel of Figure 3 shows the time series for x_t that is generated by this procedure. One can see that there are sharp spikes in this measure of idiosyncratic risk during recessions with an especially large spike in the Great Recession. The correlation of this series with short-term unemployment is 0.8. As short-term unemployment is high when workers are flowing into unemployment during recessions one interpretation is that labor market events that lead to flows into unemployment are also associated with negatively skewed innovations in permanent skill. This interpretation is consistent with the findings of Davis and von Wachter (2011) who present evidence that job layoffs are associated with large and long-lasting reductions in earnings and that long-term earnings losses are roughly twice as large for layoffs that occur in recessions.

My finding that idiosyncratic risk is closely related to short-term unemployment is supported by the work of Schmidt (2015). Schmidt (2015) also creates a quarterly-time series of idiosyncratic risk based on the Guvenen et al. (2014) data. He uses annual observations of

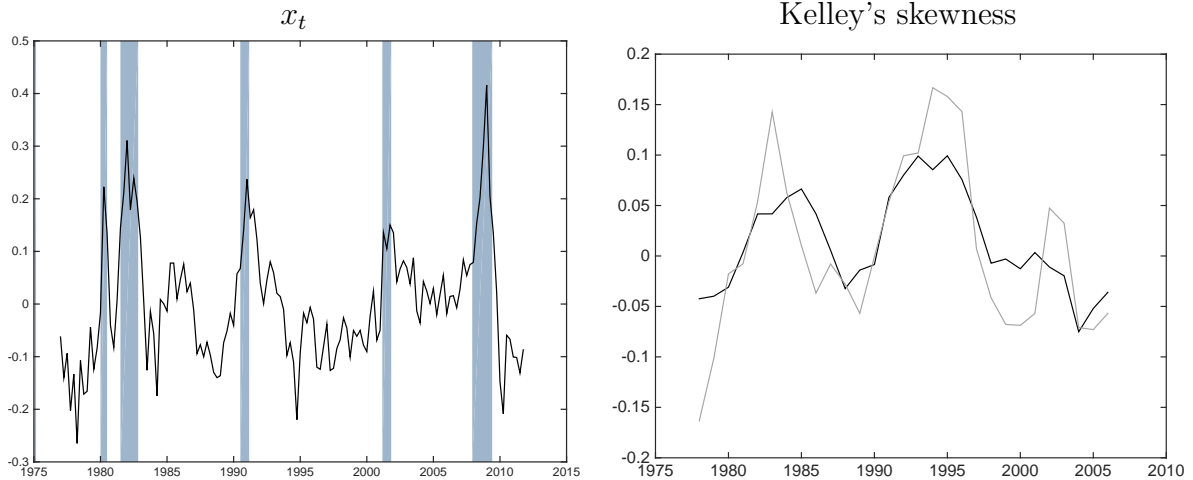


Figure 3: Empirical measure of x_t and Kelley's skewness of five-year earnings changes for model (dark line) and data (light line).

the skewness of earnings growth rates and then interpolates this skewness index to a quarterly time series for the skewness of idiosyncratic shocks using 109 macroeconomic time series. He finds that the resulting skewness index is closely related to initial claims for unemployment insurance, which is similar to my finding that idiosyncratic risk is highly-correlated with the number of short-term unemployed. Schmidt (2015) also finds that idiosyncratic risk reached unprecedented levels in the Great Recession.

The right panel of Figure 3 shows a measure of skewness in the five-year earnings changes for the model and the data. Kelley's skewness is Guvenen et al.'s preferred measure of skewness because it is less sensitive to extreme observations. It is calculated from the 10th, 50th and 90th percentiles of the distribution as $((P90 - P50) - (P50 - P10)) / (P90 - P10)$. The model slightly understates the volatility in this measure of risk.

To parameterize the aggregate shock processes, I estimate AR(1) processes for the three series \hat{u}_t , $\hat{\lambda}_t$, and x_t . Given these estimates, I calculate the covariance matrix of the residuals and perform a Cholesky decomposition of the covariance matrix yielding the following system

$$\left[\hat{u}' + 2.7926, \hat{\lambda}' - 0.7679, x' \right]^T = D \left[\hat{u} + 2.7926, \hat{\lambda} - 0.7679, x \right]^T + \epsilon',$$

where D is a diagonal matrix with diagonal elements $[0.9650, 0.9457, 0.7887]$ and the decom-

posed covariance matrix of ϵ is

$$\begin{pmatrix} 0.0033 & 0 & 0 \\ -0.0626 & 0.0563 & 0 \\ 0.0321 & 0.0259 & 0.0576 \end{pmatrix}.$$

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, b^u , to 0.3, which is in line with replacement rates for the United States reported by Martin (1996). The skill insurance parameter b^y 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 β 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 preference heterogeneity as shown by the comparison with the second row of the table in which

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 remaining lifetime as opposed to the progressivity of the system at a point in time. However, if a substantial portion of households are constrained, the degree of insurance at a point in time could be important in that transfers have a strong effect on current consumption. I assume a constant tax-and-transfer system for simplicity.

	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	0.55	0.28	0.78
Common- β	0.03	0.08	0.12	0.19	0.57	0.42	0.30	0.14	0.52
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).

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, \bar{K} , the unemployment rate, u , and the measure of income inequality, Q . The aggregate state is then $S_t = \{z, u, \bar{K}, \lambda, u_{-1}, x, Q\}$, 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 by avoiding computing the solution for combinations of states that are very unlikely to arise in practice. For individual cash on hand, I use an endogenous grid point method and place 100 grid points on K' . Appendix D provides further discussion of the methods and presents several accuracy checks.

4 Results

I now assess the extent to which household heterogeneity and uninsurable idiosyncratic risk alters the dynamics of aggregate consumption. To do so, I compare four economies: (i) the baseline model described above; (ii) a version of the model with a distribution of idiosyncratic risk that is stable over time (i.e. $x_t = 0$ for all t); (iii) a version with a stable distribution of risk and a single rate of time preference so there is less wealth heterogeneity; and (iv) a complete markets version of the model. In the complete markets model, households have a common discount rate and all shocks are insurable including mortality risk, which leads to the standard Euler equation for aggregate consumption

$$\bar{C}^{-\gamma} = \beta \mathbb{E}_t \left[\bar{C}'^{-\gamma} \tilde{R}' \right]$$

as shown in Appendix E. For both the model without time preference heterogeneity and the complete markets model I recalibrate the discount rate to match the capital-output ratio from the baseline model.

4.1 Unconditional second moments

Table 3 displays standard deviations and correlations of output and consumption both in log-levels and in growth rates.⁹ The standard deviation of consumption growth is 43 percent larger in the baseline model than in the complete markets version of the model. This difference reflects all aspects of household heterogeneity. To isolate the effects of time-varying idiosyncratic risk, one can compare the first and second rows, which show time-varying risk raises the volatility of consumption growth and reduces the correlation of consumption and income growth. The increase in consumption volatility is 21 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

⁹Simulated consumption growth is especially sensitive to sampling variation. For a fixed set of aggregate shocks, I continue increasing the number of households in the simulation until the standard deviation of consumption growth stabilizes. For the baseline model this required simulating a panel of 7.2 million households.

aggregate income.

While consumption growth is more volatile when idiosyncratic risk is time-varying, the volatility of the level of consumption is hardly affected. Output is somewhat more stable in the baseline model than in the complete markets model although this finding is not specific to the model with time-varying risk. To understand the effects on the volatilities of consumption and output in levels, consider the following: if market incompleteness pushes consumption down in a recession this creates more investment and raises output in the next period. In the next period the effects of market incompleteness are ambiguous as precautionary savings motives and borrowing constraints may be working to reduce consumption while the increase in income is supporting consumption. While the effects on consumption are ambiguous, output falls by less than the complete markets case as a result of the increase in the capital stock. In this simple model these outcomes are direct implications of the aggregate resource constraint. One way of addressing this co-movement problem is to introduce nominal rigidities and constraints on monetary policy as in Basu and Bundick (2012) although I do not pursue that here.

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 and is dominated by movements in the capital stock and TFP. As the extent of idiosyncratic risk appears to spike in recessions and quickly recede to more normal levels—as shown in Figure 3—its effects on the level of consumption are short lived. These high-frequency movements in aggregate consumption therefore are much more visible in the standard deviation of consumption growth than they are in the volatility of the level of consumption.

Rows (ii) to (iv) of Table 3 are closely related to the three models presented in Krusell and Smith (1998). Row (ii) is akin to their stochastic- β economy. Like Krusell and Smith, I too find that consumption and income are more strongly correlated when the model has a realistic degree of wealth inequality (comparing rows (ii) and (iii)). This follows from the hand-to-mouth behavior of low-wealth households. However, the change in correlation is quite small here being only a difference of 0.956 versus 0.947. Meanwhile, Krusell and Smith find a much larger difference of 0.825 versus 0.701 between their stochastic- β model and their

		$\sigma_{\Delta\bar{Y}}$	$\sigma_{\Delta\bar{C}}$		$\sigma_{\bar{Y}}$	$\sigma_{\bar{C}}$	$\rho_{\Delta\bar{Y},\Delta\bar{C}}$	$\rho_{\bar{Y},\bar{C}}$
			Relative					
(i)	Baseline	0.834	0.420	1.427	3.987	3.112	0.797	0.948
(ii)	Constant risk	0.834	0.359	1.221	3.983	3.117	0.982	0.956
(iii)	Common- β , constant risk	0.833	0.307	1.045	4.165	3.099	0.984	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 (σ) and correlations (ρ) of aggregate output (\bar{Y}) and consumption (\bar{C}) growth rates (denoted with Δ) 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.

baseline model. An important difference in my analysis is that the aggregate shocks are much more persistent here and therefore consumption responds more strongly to income even in 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 similar to the complete markets model. Therefore the conclusions from the baseline economy in Krusell and Smith (1998) carry over to this model when income risks are not time-varying and wealth inequality is modest.

The last row of Table 3 shows the empirical moments for comparison. The two principal effects of time-varying idiosyncratic risk are to make aggregate consumption growth more volatile and less correlated with aggregate income growth. In both of these dimensions, the baseline model is closer to the data than any of the three benchmarks.

4.2 The Great Recession

In order to illustrate the implications of time-varying risk, I now assess its contribution to the path of aggregate consumption during the Great Recession. I assume that the economy is in its risky steady state¹⁰ in 2007:I. Beginning from this starting point, I simulate the economy

¹⁰Coeurdacier et al. (2013) define the risky steady state as the point to which the economy will converge if the realization of aggregate shocks is zero for all periods. This concept differs from the deterministic steady

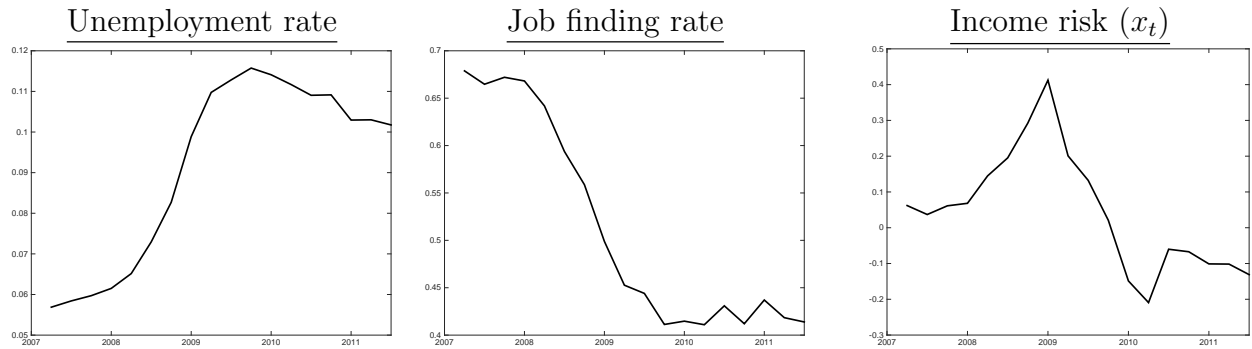
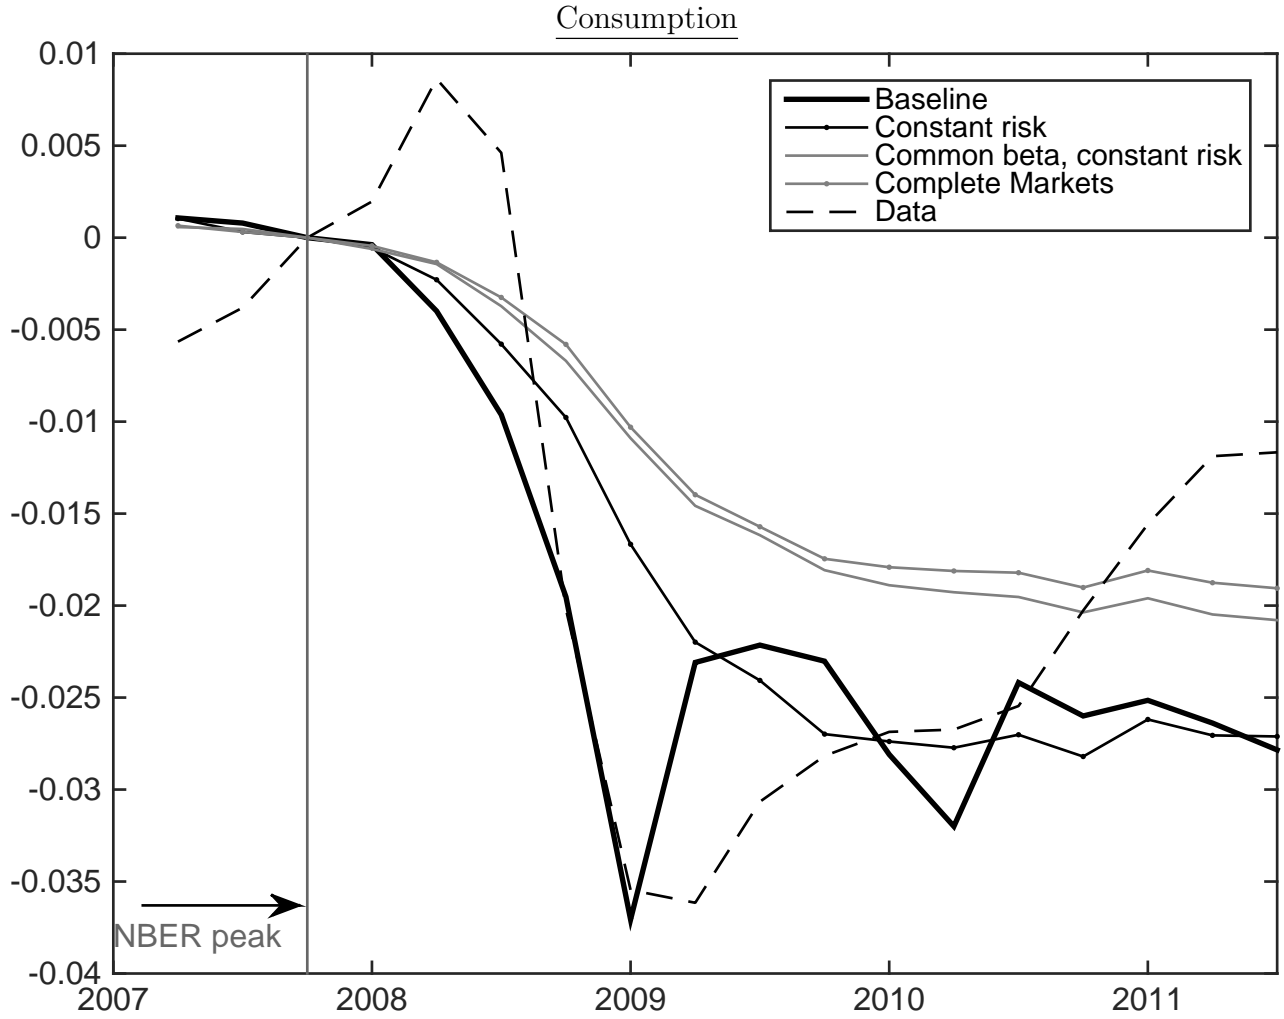


Figure 4: Dynamics of aggregate consumption implied by labor market shocks in the Great Recession. Data refer to per capita consumption of non-durables and services deflated with the GDP deflator and detrended with the HP filter with smoothing parameter 1600.

using the shocks taken from the data for the unemployment rate, u , the skewness of the income shock process, x , and the job-finding rate, λ . The construction of the series for x and λ is described in Section 3. Given the assumption about the initial condition in 2007:I, I then use equations (11), (12), and (14) to solve for sequences of $\epsilon_{u,t}$, $\epsilon_{\lambda,t}$, and $\epsilon_{x,t}$. 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 3.7 percent decline. Had idiosyncratic risks remained stable the decline at this date would have only been 1.7 percent so time-varying risk reduced aggregate consumption by 2.0 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.

Perri and Steinberg (2012) document that there was little change in the cross-sectional distribution of disposable income and consumption during the Great Recession. One might ask whether the increase in income risk in my analysis creates a counterfactual increase in the dispersion of consumption in the cross section. Figure 5 shows that the model is consistent with the Perri-Steinberg observations by plotting the ratios of 90th and 50th percentiles and state in that the agents believe that aggregate shocks can occur and this contributes to their precautionary savings motive.

of the 50th and 20th percentiles of consumption and of disposable income. The model predicts that these ratios are nearly constant as in the data.

While the invariance of the consumption distribution can be partly attributed to insurance, either self insurance or social insurance, this cannot be the whole explanation as the increase in risk leads to a substantial drop in the average level of consumption. There are two ingredients that are needed to understand how a change in risk can lead to a drop in aggregate consumption without affecting the cross-sectional distribution of consumption. First, the cross sectional distributions of income and consumption reflect the accumulation of past shocks and these distributions are only marginally affected by one quarter's innovations. As the aggregate shock to risk is short-lived the changes in the distribution of risks do not accumulate to substantial differences in the cross sectional distribution of levels. Second, the increase in risk leads all households reduce their consumption more or less in unison as all households face more risk and self insurance is not very effective in smoothing the persistent income shocks so high-wealth households respond similarly to low-wealth households. As a result, the increase in idiosyncratic risk leads to a substantial drop in aggregate consumption with little effect on the cross sectional distribution of consumption.

5 Conclusion

The deterioration of labor market conditions in the Great Recession has renewed interest in the effects of idiosyncratic risk on the business cycle. Changes in idiosyncratic risk will only have strong effects on consumption if households are not self-insured against these risks. This paper focusses on changes in the distribution of shocks to the persistent component of earnings. As 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 can explain the sharp 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

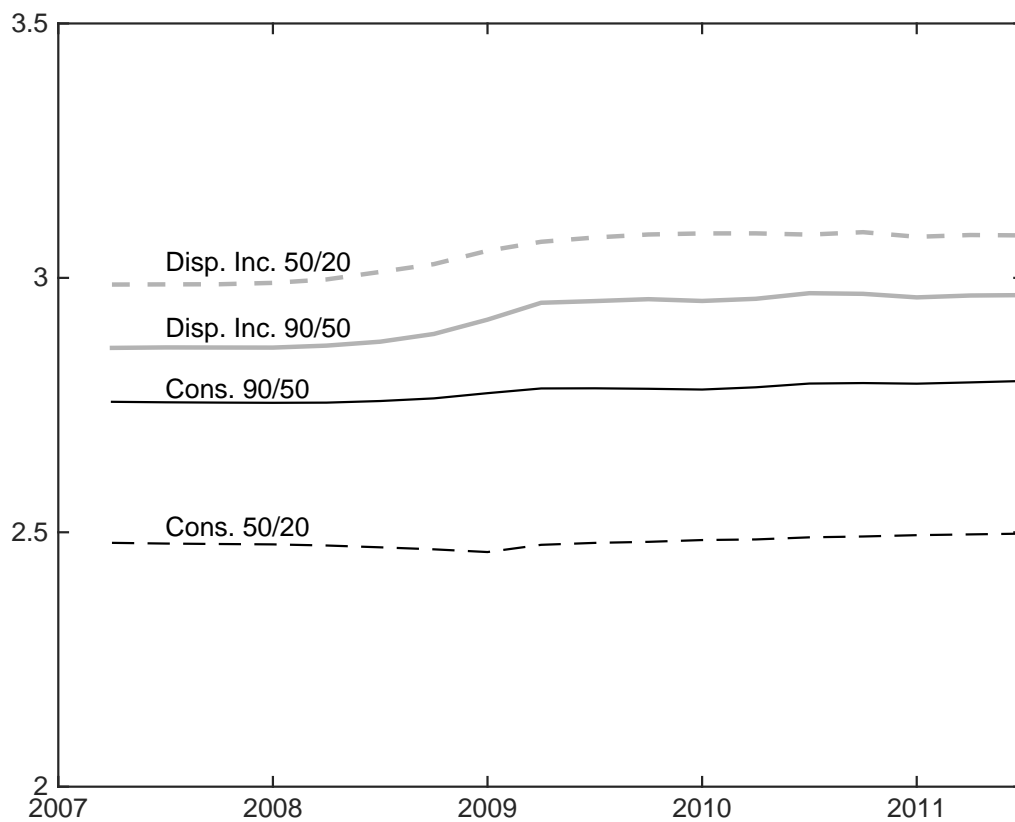


Figure 5: Cross-sectional distribution of levels of consumption and disposable income. Ratios of 90th and 50th percentiles and 50th and 20th percentiles.

in investment and an increase in output in future periods. Moreover, there is no endogenous feedback between the level of consumption and the extent of risk. Ravn and Sterk (2013) and Challe et al. (2015) explore an amplification mechanism that runs from unemployment risk to precautionary savings to reductions in aggregate demand and back to unemployment risk. This same chain of events could be triggered even more powerfully by the type of time-varying risk studied here. More generally, future work might incorporate cyclicity in the distribution of persistent earnings shocks as an important source of fluctuations in aggregate consumption in richer models of the business cycle.

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.
- Castaneda, A., Díaz-Giménez, J., and Ríos-Rull, J.-V. (2003). Accounting for the us earnings and wealth inequality. *Journal of Political Economy*, 111(4):818–857.
- Challe, E., Matheron, J., Ragot, X., and Rubio-Ramirez, J. F. (2015). Precautionary saving and aggregate demand.
- Challe, E. and Ragot, X. (2013). Precautionary saving over the business cycle. *Paris School of Economics manuscript*.
- Coeurdacier, N., Rey, H., and Winant, P. (2013). Financial integration and growth in a risky world.

- Constantinides, G. M. and Duffie, D. (1996). Asset pricing with heterogeneous consumers. *Journal of Political economy*, 104(2):219.
- Constantinides, G. M. and Ghosh, A. (2014). Asset pricing with countercyclical household consumption risk. Technical report, National Bureau of Economic Research.
- Davis, S. J. and von Wachter, T. (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity*, (2).
- 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.
- Diaz-Gimenez, J., Glover, A., and Rios-Rull, J.-V. (2011). Facts on the distributions of earnings, income, and wealth in the united states: 2007 update. *Quarterly Review*.
- Fernández-Villaverde, J., Guerrón-Quintana, P. A., Kuester, K., and Rubio-Ramírez, J. (2013). Fiscal volatility shocks and economic activity. Technical report, University of Pennsylvania.
- Gilchrist, S., Sim, J. W., and Zakrajšek, E. (2014). Uncertainty, financial frictions, and investment dynamics. Working Paper 20038, National Bureau of Economic Research.
- Guvenen, F., Ozkan, S., and Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2014). Optimal tax progressivity: An analytical framework. Working Paper 19899, National Bureau of Economic Research.
- Huckfeldt, C. (2014). The scarring effect of recessions: A quantitative analysis.
- Jarosch, G. (2014). Searching for job security and the consequences of job loss. University of Chicago Working Paper.
- Judd, K. L. (1992). Projection methods for solving aggregate growth models. *Journal of Economic Theory*, 58(2):410–452.
- Judd, K. L., Maliar, L., and Maliar, S. (2012). Merging simulation and projection approaches

- to solve high-dimensional problems. Working Paper 18501, National Bureau of Economic Research.
- Krebs, T. (2003). Growth and welfare effects of business cycles in economies with idiosyncratic human capital risk. *Review of Economic Dynamics*, 6(4):846–868.
- Krebs, T. (2007). Job displacement risk and the cost of business cycles. *The American economic review*, pages 664–686.
- Krueger, D., Mitman, K., and Perri, F. (2015). Macroeconomics and heterogeneity, including inequality. University of Pennsylvania working paper.
- 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.
- Perri, F. and Steinberg, J. (2012). Inequality and redistribution during the great recession. Economic Policy Paper 12-1, Federal Reserve Bank of Minneapolis.
- Ravn, M. and Sterk, V. (2013). Job uncertainty and deep recessions. UCL Working Paper.
- Schmidt, L. (2015). Climbing and falling off 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

$$\begin{aligned}\tilde{Q}^\theta &= \mathbb{E} [e^{\theta(1-b^y)}] \\ \tilde{Q}^\eta &= \mathbb{E} [e^{\eta(1-b^y)}] \\ \tilde{Q}^\xi &= \mathbb{E} [e^{\xi(1-b^y)}],\end{aligned}$$

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

$$Q = \tilde{Q}^\theta \tilde{Q}^\xi.$$

\tilde{Q}^θ evolves according to

$$\begin{aligned}\tilde{Q}^{\theta'} &= (1 - \omega) \mathbb{E} [e^{(\theta+\eta')(1-b^y)}] + \omega \\ \tilde{Q}^{\theta'} &= (1 - \omega) \tilde{Q}^\theta \tilde{Q}^{\eta'} + \omega.\end{aligned}$$

And as \tilde{Q}^ξ is constant one can then write

$$\begin{aligned}\tilde{Q}^{\theta'} \tilde{Q}^\xi &= (1 - \omega) \tilde{Q}^\theta \tilde{Q}^{\eta'} \tilde{Q}^\xi + \omega \tilde{Q}^\xi \\ Q' &= (1 - \omega) Q \tilde{Q}^{\eta'} + \omega \tilde{Q}^\xi.\end{aligned}$$

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. (2014).

Step 1. Calculate λ_t and ζ_t implied by the data. To do so, use the data on short-term unemployment described in Section 3 and solve for λ_t and ζ_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-term unemployed (15 or more weeks), an index of average weekly hours, and the labor force participation rate.¹¹ Note that the employment-population ratio can be expressed as a function of these variables. I transform these four series to have mean zero and unit standard deviation and then express the resulting series in terms of their principal components. Orthogonalizing the series into principal components should not affect the results in theory, but it is helpful for the numerical analysis. These quarterly data cover 1977:I to 2011:IV. Store these in a matrix X .

Step 3. Guess a vector of parameters

$$\Theta \equiv [\phi_1, \dots, \phi_4, \sigma_\xi, \mu_2, \mu_3, \sigma_{\eta,1}, \sigma_{\eta,2}, \sigma_{\eta,3}, p_2],$$

and ϕ_j is a loading on the j th labor market indicator. Also guess a sequence $\{m_t\}_{t=1978}^{2011}$. m_t is the quarterly growth rate of average income in year t , which shifts the entire distribution from which η is drawn. While I simulate quarterly data, I assume the mean growth rate is constant in each year as the observed data are at an annual frequency.

Step 4. Calculate $\mu_{1,t}$, $\mu_{2,t}$ and $\mu_{3,t}$ from Equations (1) - (3) with $x = X\phi$. The normalization $\bar{\mu}$ is chosen to satisfy $\mathbb{E}[e^\eta] = 1$ and this requires

$$\bar{\mu} = -\log(p_1 \exp(\sigma_{\eta,1}^2/2) + p_2 \exp(\mu_2 - x + \sigma_{\eta,2}^2/2) + p_3 \exp(\mu_3 - x + \sigma_{\eta,3}^2/2)). \quad (\text{A1})$$

Step 5. Simulate employment, skill, and mortality shocks for a panel of households. The employment transition probabilities are the values for λ_t and ζ_t computed in step 1. I simulate 10,000 individuals from 1977 through 2011. The results are not sensitive to the way the

¹¹These data series constructed from the series with the following codes in the Federal Reserve Bank of St. Louis FRED database: CLF16OV, UNEMPLOY, UEMP15OV, PRS85006023, and CIVPART.

distribution of θ is initialized because the objects of interest are related to earnings changes as opposed to levels. I initialize 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. (2014). 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. (2014).

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. (2014). Figure 6 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 θ_i has income proportional to $e^{(1-b^y)\theta_i}$. Given the formulation of the unemployment insurance scheme, this proportionality holds even for unemployed households. This scaling along with homothetic preferences and permanent shocks to θ can be exploited to eliminate one state variable. Specifically, use lower case letters

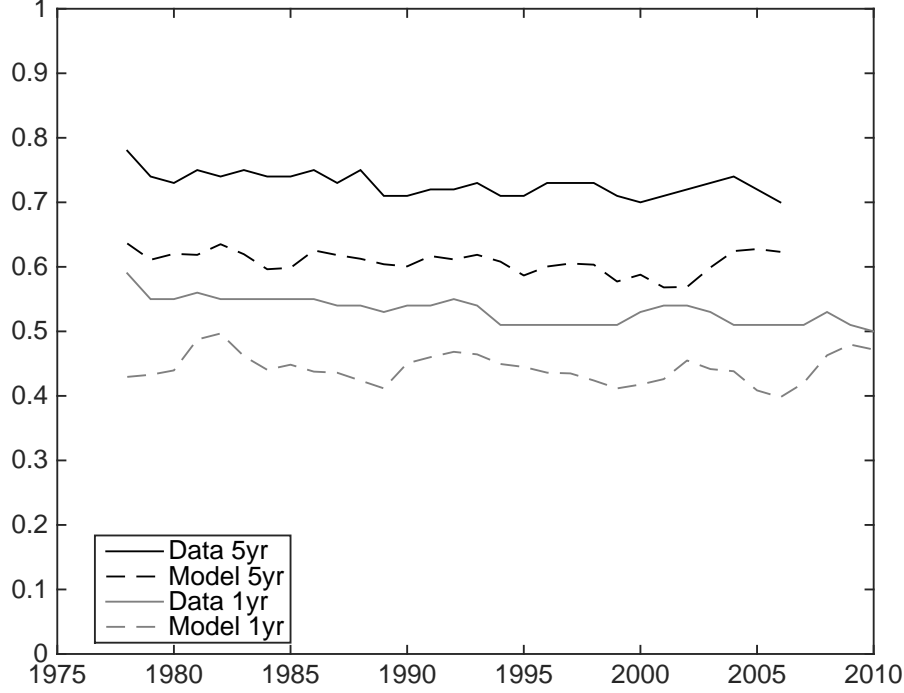


Figure 6: Simulated and empirical standard deviations of the income process.

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

$$c_i = \frac{C_i}{e^{(1-b^y)\theta_i}}, \quad a_i = \frac{A_i}{e^{(1-b^y)\theta_i}}, \quad k'_i = \frac{K'_i}{e^{(1-b^y)\theta_i}}.$$

The household's Euler equation and budget constraint are

$$C_{i,t}^{-\gamma} \geq \beta_i(1-\omega)\mathbb{E}_t [R_{t+1}C_{i,t+1}^{-\gamma}]$$

$$C_{i,t} + K_{i,t} = RK_{i,t-1} + (1-\tau)W_t e^{(1-b^y)y_{i,t}} [n_{i,t} + b^u(1-n_{i,t})].$$

and in terms of normalized variables these are

$$c_{i,t}^{-\gamma} \geq \beta_i(1-\omega)\mathbb{E}_t [e^{-\gamma(1-b^y)\eta_{i,t+1}} R_{t+1} c_{i,t+1}^{-\gamma}] \quad (\text{A2})$$

$$c_{i,t} + k_{i,t} = Rk_{i,t-1} e^{-(1-b^y)\eta_{i,t}} + (1-\tau)W_t e^{(1-b^y)\xi_{i,t}} [n_{i,t} + b^u(1-n_{i,t})]. \quad (\text{A3})$$

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 $\mu_{1,t}$, $\mu_{2,t}$, $\mu_{3,t}$, $\bar{\mu}_t$, z , u , λ , x , ζ , Q , τ , W , R , and \bar{K} . Closing the model requires determining the aggregate

capital stock \bar{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

$$\bar{K}' = h(z, \hat{\lambda}, \hat{u}, \hat{u}_{-1}, x, \bar{K}, Q). \quad (\text{A4})$$

I assume that $h(\cdot)$ is a complete second-order polynomial. In simulating the model, \bar{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 (A2) and (A3) are independent of the household's permanent income, $e^{(1-b^y)\theta}$, and so the decision for K' will be proportional to this permanent income. In particular the household's savings policy rule can be expressed as

$$F(A, \theta, n, S) = e^{(1-b^y)\theta} f(a, n, S).$$

In normalized terms we then have

$$\bar{K}' = \sum_n \int_A \int_\theta e^{(1-b^y)\theta} f(Ae^{-(1-b^y)\theta}, n, S) \Gamma(A, d\theta, n). \quad (\text{A5})$$

The household's value function can also be expressed in normalized terms. Define v as the value function in normalized terms

$$v(a, n, S) \equiv V(A, \theta, n, S)e^{-(1-\gamma)(1-b^y)\theta}.$$

The Bellman equation is then

$$e^{(1-\gamma)(1-b^y)\theta} v(a, n, S) = e^{(1-\gamma)(1-b^y)\theta} \max_{k' \geq 0} \left\{ \frac{(a - k')^{1-\gamma}}{1-\gamma} + \beta(1-\omega) \mathbb{E} \left[e^{(1-\gamma)(1-b^y)\theta'} v(a', n', S') \right] \right\}$$

so $v(a, n, S)$ satisfies

$$v(a, n, S) = \max_{k' \geq 0} \left\{ \frac{(a - k')^{1-\gamma}}{1-\gamma} + \beta(1-\omega) \mathbb{E} \left[e^{(1-\gamma)(1-b^y)\theta'} v(a', n', S') \right] \right\}. \quad (\text{A6})$$

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 below, I approximate functions of aggregate states with complete second-order polynomials.¹² As the dimension of S is seven there are 36 terms in these polynomials that will be determined by the value of the function on this grid.

Solving the household's problem I solve the household's problem using the endogenous grid point method (Carroll, 2006). A household's decision rule can be written in terms of cash on hand relative to permanent income

$$F(A, \theta, n, \beta, S) = e^{(1-b^y)\theta} f(a, n, \beta, S),$$

where $a = Ae^{-(1-b^y)\theta}$. Even though it is not a state, I have included β among the household's states in order to be explicit about the different types of households whose decision rules must be solved for. For given values of $n \in \{0, 1\}$, β , and aggregate state S , I approximate the household's savings function with a piece-wise linear function of 100 knots with more knots placed at low levels of savings and $k'_{[1]} = 0$. I fix a grid on end-of-period savings k' such that

$$k'_{[j]} = e^{(1-b^y)\theta} \hat{f}(a_{[j]}(n, \beta, S), n, \beta, S),$$

where $[j]$ indexes grid points and $a_{[j]}(n, \beta, S)$ is the value of normalized cash on hand (relative to $e^{(1-b^y)\theta}$) for which a household with states (n, β, S) saves $k'_{[j]}$. n and β both take two discrete values and there are 100 values of j so the algorithm must find 400 functions that map S to particular values of $a_{[j]}(n, \beta, S)$. I approximate each of these functions as a complete second-order polynomial in S . As there are more grid points than terms in the polynomials

¹²The use of second-order polynomials should not be confused with a second-order perturbation approximation method. The projection method used here minimizes the residual in the model equations across a grid over the state space as opposed to a perturbation method which uses information from the derivatives of the model equations at a single point in the state space.

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 η 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 x . For the transitory shock, ξ , 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 9.

Accuracy of the law of motion for capital To assess the accuracy of the law of motion for the capital stock, Figure 7 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.¹³ 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 0.0048 and the mean absolute log difference is 0.0022. Another commonly-reported accuracy check is the R^2 of the one-step

¹³As 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.

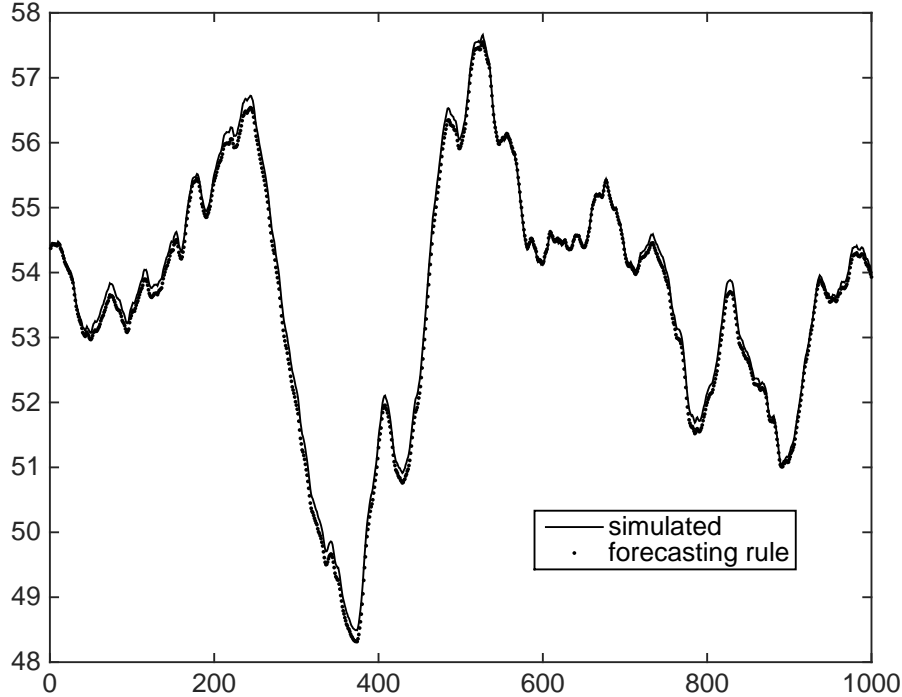


Figure 7: Simulated aggregate capital stock with implied values from $\log K' = h(S)$.

ahead forecast, which is 0.999993.

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 errors.¹⁴ For a given state of the economy, S , the distribution of wealth, the capital stock, and exogenous variables are predetermined.

$$\text{Pre-determined and exogenous: } K, z, u, u_{-1}, \lambda, x, Q, \Gamma.$$

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

¹⁴See 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

$$\text{Approx. solutions: } a_{[j]}(n, \beta, S) \quad \forall j, n, \beta.$$

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

$$\frac{\beta(1 - \omega)\mathbb{E} [e^{-\gamma(1-b^y)\eta'} R(S')c(a', n', \beta, S')^{-\gamma}]^{-1/\gamma}}{c(a, n, \beta, S)} - 1$$

where $a' = R(S')k'e^{-(1-b^y)\eta'} + (1 - \tau(S'))W(S')e^{(1-b^y)\xi'}[n' + b^u(1 - n')]$ and the consumption functions satisfy $c + k' = a$. Here \mathbb{E} represents an expectation over aggregate and idiosyncratic shocks. For aggregate shocks, I use Gaussian quadrature with seven points in each dimension. For idiosyncratic shocks I use the same Gaussian quadrature methods as used to solve the model. For households who are borrowing constrained the Euler equation should not hold. For these households consumption is determined from the borrowing constraint and there is no Euler equation error.

Using these steps, I can compute the Euler equation error for a household with a particular set of states (aggregate and idiosyncratic). To choose a set of aggregate states at which to evaluate the errors, I simulate the economy for 1000 starting from the risky steady state and repeat this 100 times for different sets of random shocks. This produces 100 points that can be considered as draws from the model's ergodic distribution over the aggregate state space. For idiosyncratic states, I construct a fine grid on normalized cash on hand. Specifically, I use 1000 equally spaced points from 1/1000 to 1000.

Figure 8 summarizes the errors across points in the state space. Each of the panels corresponds to a set of discrete states for a household with the top row showing less patient households and the bottom row more patient and the left column unemployed and the right column employed. Each panel plots the mean and maximum absolute errors across the 100 aggregate states that were tested.

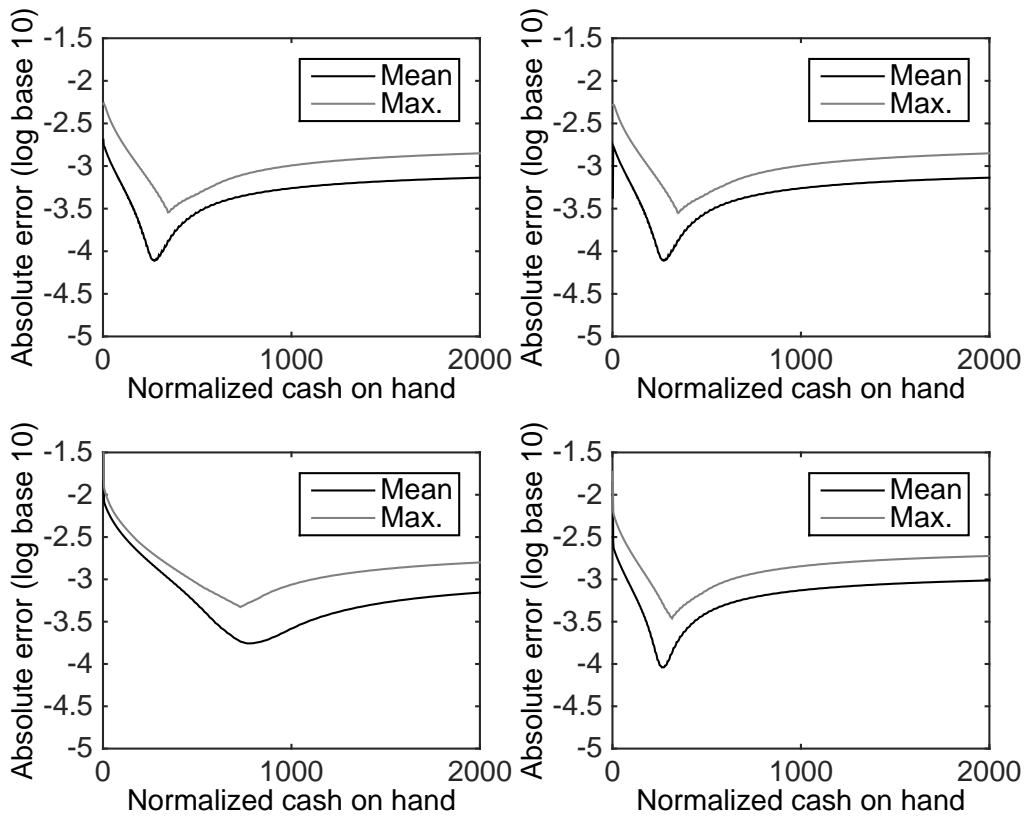


Figure 8: Euler equation errors. Left column: unemployed; right column: employed; top row: less patient; bottom row: more patient. Maximum and mean across 100 aggregate states.

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 $I_{i,t}$ 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

$$\sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi_t(s^t) \frac{C_{i,t}^{1-\gamma}}{1-\gamma} I_{i,t}(s^t).$$

Notice that a household only values consumption when it is living. In order to prevent them from choosing negative values of consumption when not living I impose $C_{i,t}(s^t) \geq 0$ for all i , t , and s^t . The household's present-value budget constraint is

$$\sum_{t=0}^{\infty} \sum_{s^t} p_t(s^t) \left[C_{i,t}(s^t) + K_{i,t+1}(s^t) - W_t(s^t) \ell_{i,t}(s^t) - \tilde{R}_t(s^t) K_{i,t}(s^{t-1}) \right],$$

where $\ell_{i,t}$ is the household's endowment of efficiency units of labor. I assume that all households are identical when trade occurs and each is endowed with an equal share of the initial capital stock.

The household's Lagrangian is

$$\begin{aligned} \mathcal{L} = & \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi_t(s^t) \frac{C_{i,t}(s^t)^{1-\gamma}}{1-\gamma} I_{i,t}(s^t) \\ & - \Xi_i \sum_{t=0}^{\infty} \sum_{s^t} p_t(s^t) \left[C_{i,t}(s^t) + K_{i,t+1}(s^t) - W_t(s^t) \ell_{i,t}(s^t) - \tilde{R}_t(s^t) K_{i,t}(s^{t-1}) \right] \\ & + \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi_t(s^t) \psi_{i,t}(s^t) C_{i,t}(s^t), \end{aligned}$$

where Ξ_i and $\psi_{i,t}(s^t)$ are Lagrange multipliers. The first order condition with respect to consumption is

$$\beta^t \pi_t(s^t) \left[C_{i,t}(s^t)^{-\gamma} I_{i,t}(s^t) + \psi_{i,t}(s^t) \right] = p_t(s^t) \Xi_i.$$

The complementary slackness condition is $\psi_{i,t}(s^t) C_{i,t}(s^t) = 0$. By symmetry, the Lagrange multiplier Ξ_i is common across households. It follows that $C_{i,t}(s^t)^{-\gamma} I_{i,t}(s^t) + \psi_{i,t}(s^t)$ must be common across households at a particular date and history. Suppose the household is living, then consumption is positive and common across living households and $\psi_{i,t}(s^t) = 0$. If the household is not living then $\psi_{i,t}(s^t)$ takes the value of the marginal utility of consumption for living households and consumption is zero. This establishes that all living households consume the same amount regardless of their labor income history or age. Define $\bar{C}_t(s^t)$ as the common level of consumption for living households and note $\bar{C}_t(s^t)^{-\gamma} = C_{i,t}(s^t)^{-\gamma} I_{i,t}(s^t) + \psi_{i,t}(s^t)$.

The first order condition with respect to $K_{i,t+1}(s^t)$ is

$$\Xi_i p_t(s^t) = \sum_{s^{t+1}|s^t} \Xi_i p_{t+1}(s^{t+1}) \tilde{R}_{t+1}(s^{t+1}).$$

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

$$\begin{aligned}
& C_{i,t}(s^t)^{-\gamma} I_{i,t}(s^t) + \psi_{i,t}(s^t) \\
&= \beta \sum_{s^{t+1}|s^t} \frac{\pi_t(s^{t+1})}{\pi_t(s^t)} [C_{i,t+1}(s^{t+1})^{-\gamma} I_{i,t+1}(s^{t+1}) + \psi_{i,t+1}(s^{t+1})] \tilde{R}_{t+1}(s^{t+1}) \\
\bar{C}_t(s^t)^{-\gamma} &= \beta \sum_{s^{t+1}|s^t} e^{q_{t+1}(s^{t+1})} \frac{\pi_t(s^{t+1})}{\pi_t(s^t)} \bar{C}_{t+1}(s^{t+1})^{-\gamma} \tilde{R}_{t+1}(s^{t+1}) \\
\bar{C}^{-\gamma} &= \beta \mathbb{E}_t [\bar{C}'^{-\gamma} \tilde{R}'].
\end{aligned}$$