

Idiosyncratic risk, insurance, and aggregate consumption dynamics: a likelihood perspective

Alisdair McKay*

July 2013

Abstract

This paper shows how the empirical implications of incomplete markets models can be assessed using the same full-information methods that are commonly used for representative agent models. It then asks what features of the microeconomic insurance arrangement are important for understanding the dynamics of aggregate consumption as it relates to aggregate labor income and employment conditions. A model with a low level of insurance against unemployment risk and an intermediate level of insurance against individual skill shocks provides the best fit of the aggregate data. A model that matches the strong consumption responses to fiscal stimulus payments does not improve the overall fit to the aggregate data.

*Contact: amckay@bu.edu. I would like to thank Simon Gilchrist, Zhongjun Qu, and Ricardo Reis for helpful discussions and Carlos Ramos for excellent research assistance.

1 Introduction

A range of microeconomic evidence is difficult to reconcile with the complete-markets, rational-expectations model of consumption that is now standard in most models of the business cycle. For instance: consumption does respond to idiosyncratic income changes (Cochrane, 1991; Attanasio and Davis, 1996) and consumption does respond to anticipated income changes (Souleles, 1999; Parker, 1999; Johnson et al., 2006). This evidence has led to a large literature that explores models with incomplete markets and heterogeneous households (Deaton, 1991; Carroll, 1992; Aiyagari, 1994; Krusell and Smith, 1998). Models in this class allow for very limited insurance against idiosyncratic income shocks as households only have access to self-insurance through savings. Some aspects of the microeconomic data point to households having more insurance than just self-insurance (Blundell et al., 2008; Heathcote et al., 2012), while in other respects the data suggest less (Kaplan and Violante, 2011).

This paper asks what features of the microeconomic insurance arrangement are important for understanding the dynamics of aggregate consumption as it relates to aggregate labor income and employment conditions. There are several reasons to think that the level of insurance against idiosyncratic risk could be important for aggregate consumption dynamics. For example, a deterioration of labor market conditions that increases the idiosyncratic risk of a prolonged unemployment spell may lead to a large drop in consumption when this risk is uninsurable. Additionally, the presence of idiosyncratic risk and binding borrowing constraints can lead aggregate consumption to be more sensitive to changes in income.

Methodologically, the analysis partially integrates two distinct literatures. Empirical macroeconomists have developed a set of techniques for formally comparing structural models to time series data. These techniques use a variety of structural shocks so that the model generates a rich covariance structure for observed data. One then conducts inference using the full range of empirical implications of the model (An and Schorfheide, 2007). However, the models used in these analyses rely on a representative agent abstraction so they are not useful for exploring the question at hand, which centers around relaxing the complete

markets assumption underlying the representative agent.

Incomplete markets models have so far not been incorporated into the standard toolbox of empirical macroeconomics in part due to a lack of formal methods.¹ Existing work with incomplete-markets business-cycle models has followed a calibration approach that compares the model’s implications to selected empirical moments. The methodological contribution of this paper is to show how these incomplete markets models can be analyzed using the full-information approach that has previously been used to study representative agent models. The distinction is not between estimation and calibration—the procedure I use is somewhere between the two—but the basis upon which inference is conducted.

I use these methods to study an economy populated by households that face two idiosyncratic shocks and several aggregate shocks. The idiosyncratic shocks are to the household’s individual skill and employment status. At the aggregate level, there are shocks to the aggregate wage level and shocks to the two transition probabilities that move households into and out of unemployment. The households collectively participate in an insurance scheme that partially insures against skill and unemployment shocks with the extent of insurance against the two shocks depending on two parameters. With parameters corresponding to low levels of insurance, the model is close to the standard incomplete markets model in which households rely exclusively on self-insurance. With parameters corresponding to full insurance, the model becomes a representative agent model. By varying the insurance parameters it is also possible to analyze cases between these two extremes.

I first compare the low-insurance and full-insurance economies. The results show that the low-insurance economy is much closer to the data primarily because it is able to generate considerably more volatility in aggregate consumption growth, which brings the model closer to the data. In addition, removing the insurance against idiosyncratic risk brings the model

¹There are a handful of papers that estimate versions of the incomplete markets model using microeconomic data, but these models are not suitable for aggregate time series data because they are either models of a stationary aggregate environment (Gourinchas and Parker, 2002; Guvenen and Smith, 2010) or include variation in the dispersion of household productivity levels but not in the overall level of productivity or wages (Heathcote et al., 2012).

closer to the data by making consumption growth serially correlated and correlated with the number of households entering unemployment. Using a summary measure of fit proposed by Watson (1993), I show that the variance of residuals needed to bridge the gap between model and data is roughly 50% larger for the full-insurance economy than it is for the low-insurance economy. Information from the likelihood function provides a similar comparison of the two insurance arrangements. The results in this section are related to work by Krusell and Smith (1998) and Challe and Ragot (2013) who compare the incomplete markets models and representative agent models by comparing model and data moments.

Second, I explore whether the aggregate data support an intermediate level of insurance. Using microeconomic data, Blundell et al. (2008) and Heathcote et al. (2012) find evidence in favor of a partial insurance arrangement that is between full-insurance and self-insurance alone. I find that allowing for some degree of partial insurance against skill shocks brings the model closer to the data than the low-insurance economy. With regard to unemployment risks, the model fits the data best with a low level of insurance, but the presence or absence of insurance against these risks affects the model's overall fit much less than the insurability of skill shocks.

Finally, Kaplan and Violante (2011) have recently observed that the standard incomplete markets model is inconsistent with the microeconomic evidence on the way household consumption responds to fiscal stimulus payments as estimated by Johnson et al. (2006) and others. This evidence suggests that households have less insurance than self-insurance as consumption responds more strongly to transitory income changes than the standard incomplete markets model implies. Kaplan and Violante show that illiquid assets can explain why households behave as if they are liquidity constrained even if they have considerable net worth. One interpretation of this explanation is that illiquid assets are less useful for self-insurance than are liquid assets. Motivated by Kaplan and Violante's work, I extend the model to include a cost of adjusting household asset holdings and I calibrate this adjustment cost to match the sensitivity of household consumption to transitory income changes as measured by the response to fiscal stimulus payments. The extended model has very

different predictions for how aggregate consumption responds to transitory shocks, but its predictions are nearly identical to the baseline model for the response to persistent shocks. As my estimates imply that most of the variance in aggregate consumption results from persistent shocks, the overall dynamics for consumption and the model’s ability to match the data are little affected.

Elements of this paper are related to work by Berger and Vavra (2012) who show that certain features of the business cycle dynamics of aggregate durable consumption expenditures can be well accounted for by a model in which heterogeneous households face fixed costs of adjusting their stocks of durables. My work differs in that it analyzes consumption of non-durables and services and uses a different set of empirical methods.

The paper is organized as follows: section 2 presents the model, section 3 discusses the empirical strategy and methods used to solve the model and compare its implications to the data, and section 4 discusses the data that the model is asked to explain as well as the calibration and estimation of model parameters. I present the results in three sections: section 5 compares low- and full-insurance, section 6 investigates partial insurance, and section 7 presents the extension with illiquid assets. Section 8 explores whether richer complete markets settings provide better descriptions of consumption dynamics. Finally, section 9 concludes.

2 Model

The model is populated by a unit mass of households who differ in their employment status and skill. Households maximize the expected discounted sum of period utilities given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{c_{i,t}^{1-\chi}}{1-\chi}. \quad (1)$$

The income of an individual household varies over time with shocks to its employment status, its skill level, and the aggregate after-tax wage in the economy, w_t . I assume that the log of the aggregate after-tax wage is the sum of a persistent AR(1) process and a transitory disturbance. The transitory disturbance is able to capture short-lived changes

such as temporary changes in taxes. In addition to these stochastic components, the wage grows in line with a deterministic trend. In summary,

$$\log w_t = z_t + A_t + \varepsilon_t^T, \quad (2)$$

$$z_t = z_{t-1} + g,$$

$$A_t = \rho^A A_{t-1} + \varepsilon_t^A, \quad (3)$$

where g is the trend growth rate. I assume that ε_t^T and ε_t^A are normally distributed with mean zero and standard deviations σ^T and σ^A , respectively.

Households face two types of idiosyncratic risk: skill shocks and employment shocks. A household with skill $s \in S$ will receive after-tax labor income $w_t s$ when employed and no labor income when unemployed. I assume that all households in the economy share unemployment and skill risk by pooling their incomes and then allocating after-transfer incomes such that a household's income is a function of its skill and employment status given by

$$y_t(e, s) = \frac{[e + b^u(1 - e)]s^{1-b^s}}{\int [e_{j,t} + b^u(1 - e_{j,t})]s_{j,t}^{1-b^s} dj} \int e_{j,t} w_t s_{j,t} dj, \quad (4)$$

where $e = 1$ if the household is employed and $e = 0$ otherwise. The parameters b^s and b^u control the degree of insurance against skill and unemployment shocks, respectively. If both are set to zero, equation (4) simplifies to $y_t(e, s) = esw_t$, which implies no insurance. If $b^s = b^u = 1$ all households have an equal share of aggregate income. In reality, households can receive insurance from many sources including financial markets, informal arrangements, and the government. I hope to capture all of these forces in the model, but the spirit of the income pooling is more in line with an informal risk-sharing arrangement than a formal market or government program.

Households transition between skill levels and between employment statuses according to a Markov chain with transition matrix T_t , which describes transitions from $(e_{i,t}, s_{i,t})$ to $(e_{i,t+1}, s_{i,t+1})$. Transition probabilities across skill groups are constant across time and independent of employment status. I assume the economy has reached the ergodic distribution over skills so the skill composition of the labor force is constant. Let $T^s(s, s')$ be the probability of moving from s to s' .

Employment risk differs across skill groups and across time. It is potentially important to allow unemployment risk to depend on skill because low-skill workers are more likely to be liquidity constrained and therefore changes in the unemployment rate may have a bigger effect on aggregate consumption if unemployment spells are concentrated on this group. If we take education as a measure of skill, it is well-known that unemployment rates fall with education. Using data from the Current Population Survey, Elsby et al. (2010) show that job-finding rates are similar across education groups, but job-separation rates are decreasing with education. Moreover, their Figure 8 shows that the differences in job-separation rates are fairly constant over the business cycle. Based on this evidence, I assume that the job-finding probability for any unemployed worker is λ_t and the job-separation probability for an employed worker with skill s is $\zeta_t + \zeta^s$. ζ^s is a fixed adjustment for skill-group and I normalize it to zero for the highest skill group. The driving processes of the labor market are the stochastic processes

$$\zeta_t = (1 - \rho^\zeta)\bar{\zeta} + \rho^\zeta\zeta_{t-1} + \varepsilon_t^\zeta \quad (5)$$

$$\lambda_t = (1 - \rho^\lambda)\bar{\lambda} + \rho^\lambda\lambda_{t-1} + \varepsilon_t^\lambda, \quad (6)$$

with ε_t^ζ and ε_t^λ normally distributed with mean zero and standard deviations σ^ζ and σ^λ . The overall transition matrix, T_t , is constructed by combining these employment transition probabilities with the skill transition matrix to describe the evolution of idiosyncratic states from date t to $t + 1$. I assume that ζ_t and λ_t are known at t , which means the aggregate unemployment rate for date $t + 1$ becomes known at date t , but individuals do not learn their individual employment status for date $t + 1$ until that date arrives.

In the model, there are no markets for insurance against idiosyncratic risk except as reflected in the income pooling mechanism described above. Instead, households have access to a single asset that they can use to smooth their consumption. This asset pays a risk-less return of r per period. Households are unable to borrow.

Normalization. In the absence of aggregate shocks, household income and assets will grow smoothly at the rate g . To render the economy stationary, I normalize these variables by e^{zt} .

The balanced growth path of the economy then corresponds to the stationary equilibrium of the transformed economy in which the distribution of normalized assets and income is constant as in Aiyagari and McGrattan (1998). For a generic variable x_t , let $\tilde{x}_t \equiv x_t/e^{zt}$.

Decision problem. The states of an individual's decision problem at date t includes their assets, skill and employment status, the aggregate stochastic processes and also the distribution of households over skill and employment states, which I summarize with the unemployment rates for the three skill levels denoted $U = \{u_s : s \in S\}$. The latter becomes a state variable because the income-pooling insurance scheme makes the take-home income of one household depend on the pre-insurance incomes of other households. Let $\Omega = \{A, \tilde{w}, \lambda, \zeta, U\}$ denote the aggregate state variables.

After normalization, the preference ordering becomes

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \tilde{\beta}^t \frac{\tilde{c}_{i,t}^{1-\chi}}{1-\chi}, \quad (7)$$

where $\tilde{\beta} = \beta(1+g)^{1-\chi}$. A household with assets \tilde{a} , skill s , and employment status $e \in \{0, 1\}$ solves

$$V(\tilde{a}, e, s; \Omega) = \max_{\tilde{c}, \tilde{a}'} \left\{ \frac{\tilde{c}^{1-\chi}}{1-\chi} + \tilde{\beta} \mathbb{E} [V(\tilde{a}', e', s'; \Omega')] \right\}$$

subject to

$$(1+g)\tilde{a}' + \tilde{c} = \tilde{y}(e, s) + (1+r)\tilde{a}$$

$$\tilde{a}' \geq 0.$$

The aggregate state evolves according to the laws of motion for the exogenous processes in Eq.s (2), (3), (5), and (6). The skill-specific unemployment rates evolve following

$$P_{s'} u_{s', t+1} = \sum_s P_s T^s(s, s') [(1 - \lambda_{s,t}) u_{s,t} + (\zeta_t + \zeta^s)(1 - u_{s,t})], \quad (8)$$

where P_s is the mass of households with skill level s , which is constant.

The Euler equation for this consumption decision is

$$\tilde{c}^{-\chi} (1+g) \geq \tilde{\beta}(1+r) \mathbb{E} [\tilde{c}'^{-\chi}], \quad (9)$$

where the expectation operator in the Euler equation integrates over both aggregate and idiosyncratic uncertainty. The Euler equation will hold with equality for all households except those that are constrained.

3 Empirical strategy and methods

The central question is whether the joint dynamics of aggregate consumption, income and labor market conditions generated by the model are closer to the data when the model is parameterized with some insurance parameters as compared to others. The model takes wages and employment conditions to be exogenous so the focus of the analysis is on the behavior of consumption and how it co-moves with the other variables. I consider two metrics that measure the distance between model and data. The first is the measure of fit proposed by Watson (1993), which is a lower bound on the variance of measurement error needed to reconcile the model with the data. The second is to evaluate the likelihood of the data under different combinations of parameters.

Throughout, I adopt an approach of exploring the distance between model and data at selected insurance arrangements as opposed to full-fledged estimation of the insurance parameters. A main reason for this choice is computational feasibility as solving the model for one vector of parameters takes too long to use the usual MCMC or numerical maximization algorithms. Nevertheless, I believe that the information I report here gives a good sense of which parameters values are supported by the data and which are not.

3.1 Watson’s measure of fit

Watson (1993) adopts the view that the economic model is an approximation to the stochastic process generating the data and asks how much “measurement” error would have to be attached to each variable to reconcile the autocovariances of the model (plus errors) with the data. The measurement error represents the abstraction of the model and is not truly measurement error in the typical sense. There are many measurement error processes that

could rationalize the observed data and Watson proposes to select the one with the lowest variance so the contribution of measurement error is kept to a minimum.² Watson's measure of fit can be summarized by the ratio of the measurement error variance to the data variance and this can be interpreted as $1 - R^2$ from a regression as it is akin to the residual sum of squares over the total sum of squares. Watson's measurement error process is computed in the frequency domain for each frequency separately and the measure of fit can be computed separately for different frequencies or integrated across ranges of frequencies. I will consider the model's fit overall as well as at high, business-cycle, and low frequencies.³

3.2 The value of the likelihood function

The likelihood of the data as a function of model parameters is the central component of Bayesian estimation and of maximum likelihood estimation (MLE). In the case of Bayesian estimation, the shape of likelihood will, in most cases, determine the regions of highest posterior density and therefore inferences about parameters. In the case of MLE, the shape of the likelihood clearly determines the estimated parameter values. Therefore, understanding the shape of the likelihood function is at the heart of both forms of statistical analysis.

3.3 Model solution and state-space representation

Methodologically, the main challenge to confront is to develop methods to evaluate these measures of fit for a model that includes a distribution of heterogeneous agents. While the model I analyze is a partial equilibrium model, the methods described here can equally well be applied to richer general equilibrium models.

To begin, it is useful to consider how one proceeds with a representative agent model. The most common approach is to derive a linear approximation to the equilibrium conditions and then apply a method for solving linear rational expectations models to obtain a linear

²Specifically, Watson minimizes the trace of the measurement error covariance matrix.

³Watson allows for the variables to be weighted when the measurement error variance is minimized. I use equal weights on all variables.

state-space representation of the model dynamics. Assuming Gaussian shocks, one can then apply the Kalman filter to evaluate the likelihood function. For the Watson measure of fit, one needs to find the model's spectral density matrices, which are easily computed from the state space representation (see Tkachenko and Qu, 2012).

My approach for the heterogeneous agent model is very similar to that for the representative agent model. It begins by solving the model using the algorithm developed by Reiter (2009) for models with endogenously evolving distributions in the style of Krusell and Smith (1998). This algorithm replaces the continuous distribution of wealth with a histogram with a large number of bins and it replaces the household savings decision rules with splines with a large number of knots. In this way, the economy is summarized by a finite-dimensional vector, X_t , that describes the histogram, the savings rules and aggregate variables. The equilibrium conditions of the economy are then a set of forward-looking, non-linear difference equations that X_t must satisfy, which can be written as

$$F(X_t, X_{t+1}, \eta_{t+1}, \varepsilon_{t+1}) = 0,$$

where η_{t+1} are forecast errors and ε_{t+1} is a vector of aggregate shocks. In addition to aggregate relationships, these equations describe how the distribution of wealth evolves and require that the Euler equation holds exactly for a large number of idiosyncratic states. Appendix B describes the specific equations used to solve the model. The Reiter algorithm involves linearizing these equations with respect to aggregate states around the stationary economy in which there are idiosyncratic shocks, but no aggregate shocks. The dynamics of X_t can then be found by solving the resulting linear system using standard techniques such as Sims (2002). The resulting solution is linear in aggregate states, but non-linear in idiosyncratic states. This solution algorithm is useful for this application because a) it allows for a large number of aggregate states, which allows me to incorporate a variety of persistent shocks into the model and b) it results in a linear representation of the aggregate economy, which is useful for statistical analysis.

The Reiter solution algorithm delivers a linear system for X_t of the form

$$X_{t+1} = \Psi_X X_t + \Psi_\varepsilon \varepsilon_{t+1}.$$

X_t includes the aggregate variables that we wish to observe and an observation matrix can be used to select them.

3.4 Model reduction

At this stage, the model solution has been expressed in state space form, but before proceeding with the analysis I undertake a model reduction step. In typical applications, the dimension of X_t will range from several thousand to more than ten thousand, which makes direct analysis of the system computationally demanding. Linear systems theory provides tools to construct smaller linear systems that provide a similar mapping from any history of shocks to observable variables. This is accomplished by identifying dimensions of the state space that are either unlikely to be reached or that have a limited impact on observable variables (are difficult to observe). The state space of the system can be transformed so that dimensions of the state space that are difficult to reach are also difficult to observe. These dimensions can then be discarded with little loss of accuracy in the full system.⁴ This model reduction step reduces the dimension of X_t from more than 3,600 to fewer than 50 yet results in a negligible loss of accuracy. Appendix B describes the methods and their implementation in greater detail.

3.5 Discussion of methods

The approach used here has some advantages and disadvantages relative to alternative methods. The most natural alternative is to solve the model fully non-linearly and then compare the properties of simulated data to the observed data. The first advantage concerns the

⁴The distance between two linear systems can be measured in terms of the discrepancy in their impulse responses or in other terms, the impulse response of the error system. For the methods that I use, results from linear systems theory place bounds on the maximal amplitude on the Laplace transform of these error system impulse responses. See Antoulas (2009) for a textbook treatment of model reduction techniques. Reiter (2010) has used these techniques to reduce the state space before applying the linear rational expectations solver, which is more involved than the approach taken here where the solver is applied before the model reduction step.

number of aggregate state variables. The Reiter solution method easily accommodates large numbers of aggregate state variables just as perturbation-based solution methods do in other contexts. This is important because models used for empirical analysis of time series data commonly incorporate a large number of structural shocks to reduce the colinearity of the model generated data. For example, Smets and Wouters (2007) include seven persistent structural shocks each of which becomes an aggregate state variable. Standard non-linear solution methods cannot easily be applied to such high-dimensional problems due to the curse of dimensionality. The second advantage, is that the resulting state-space representation of the economy facilitates the application of statical methods. A third advantage is that the Reiter method is easily applied to models with rich aggregate features such as nominal rigidities (see McKay and Reis, 2013).

The chief disadvantage relative to a fully non-linear solution is that the solution may be less accurate especially if there are large shocks that drive the economy far from the stationary equilibrium. To investigate the loss of accuracy, I conduct an experiment in which I solve a version of the model using the Reiter method and using a non-linear method. I then compare the properties of simulated data from the two models. In order to apply a standard non-linear solution method for this comparison, I reduce the number of aggregate shocks to two. I find that while there are some differences between the fully non-linear solution and the partially linear solution generated by Reiter's method, these discrepancies are small relative to the difference between the low-insurance and full-insurance economies or the difference between the model and the data. I also explore the effect of the model reduction step and find that it results in no appreciable loss of accuracy. Appendix C has the details.

4 Data and model parameters

4.1 Data

I judge the model specifications on how well they explain the joint dynamics of consumption of non-durable goods and services, labor income net of taxes and government transfers,⁵ and two constructed variables that relate to short-term and long-term unemployment. Consumption and income are deflated with the GDP deflator, expressed per capita and transformed by $100 \times \Delta \log(\cdot)$. I use quarterly data from 1966:I to 2012:III.

In the context of imperfect insurance, the choice of data on labor market conditions is more complicated than it is under complete markets because the distribution of labor income is now relevant to the dynamics of aggregate consumption. I ask the model to match the contribution of aggregate hours to fluctuations in aggregate labor income, but to avoid complicating the model I abstract from labor force participation and fluctuations in hours per worker. Instead, I assume that all fluctuations in hours are driven by unemployment. While this assumption overstates the amount and variability of unemployment risk, I can use the parameter b^u to smooth out this risk although this approach does complicate the interpretation of that parameter as a replacement rate.

I use $1 - h_t / \bar{h}$ as the empirical observation of the unemployment rate, where h_t is aggregate hours per capita and \bar{h} is aggregate hours at full employment. To define a full-employment level of aggregate hours per capita I look to 1999QIV as the quarter since 1960 with the largest value of aggregate hours per capita. I then define full employment hours per capita as the labor force participation rate in 1999QIV times the index of average weekly hours in that quarter.

The unemployment rate alone does not convey information about the duration of unemployment spells. I therefore use data on the composition of the unemployed pool by splitting

⁵These data come from NIPA Table 2.1. I use the sum of compensation of employees (line 2) and personal current transfer receipts (line 16) less contributions for social insurance and personal current taxes (lines 25 and 26). This is equivalent to disposable personal income less proprietor's income, rental income and asset income.




Figure 1: Constructed short-term and long-term unemployment data.

the pool into those with fewer than 15 weeks and at least 15 weeks of unemployment. As the model period is one quarter, I take less than 15 weeks to be individuals who are unemployed, but were employed in the previous quarter. These data define the shares of short-term and long-term unemployment in total unemployment. I use these shares to split the constructed unemployment rate into short-term and long-term unemployment pools. Figure 1 shows the constructed unemployment rate and its decomposition by duration. From this figure, one can see there are some low-frequency movements in the unemployment rates, which are partly attributable to demographic factors. To remove these low-frequency trends, I detrend the unemployment series using an HP filter with smoothing parameter 100,000 following Shimer (2005).

Relation between model variables and data. In the model, the mass of households in the first period of unemployment is given by

$$u_{t+1}^{\text{short}} = \sum_s P_s(\zeta_t + \zeta^s)(1 - u_t^s) \quad (10)$$

and the mass of households with unemployment durations greater than one period satisfies

$$u_{t+1}^{\text{long}} = \sum_s P_s(1 - \lambda_t)u_t^s. \quad (11)$$

Similarly, aggregate labor income is

$$Y_t = w_t \sum_s s(1 - u_{s,t})P_s = w_t \int e_{i,t}s_{i,t}di. \quad (12)$$

Finally, aggregate consumption is simply

$$C_t = \int c_{i,t}di. \quad (13)$$

4.2 Stochastic singularity and measurement error

As written, the model is not obviously stochastically singular—a condition that typically arises when there are fewer structural shocks than observable variables. While the model has only three driving processes and four observable series, the wage process is driven by two shocks—persistent and transitory—so there are an equal number of shocks and observable variables. In principle, independent movements in consumption growth could be explained in terms of offsetting persistent and transitory wage shocks. However, relying on the two wage shocks so heavily seems rather unrealistic and in practice the covariance matrices may be close to singular. Therefore, at least one more source of uncertainty is needed to explain movements in consumption growth that are not related to the other data series.

The literature has pursued several ways of breaking stochastic singularity. One approach is to augment the model with additional structural disturbances.⁶ A second approach, which is the one that I adopt, is to augment the model with measurement errors.⁷ There are two

⁶Early examples of this approach are Leeper and Sims (1994) and Ingram et al. (1994).

⁷Early examples of this approach are Sargent (1989) and McGrattan (1994).

reasons to include measurement errors beyond the need to break stochastic singularity. First, the measurement errors can be interpreted as bridging the gap between the abstraction of the model and reality, which is to say they need not be considered to be measurement errors in the typical sense but are some measure of model misspecification (Watson, 1993). Second, the data really are measured with error as demonstrated by the disagreement between different measures of aggregate employment or different measures of inflation (Boivin and Giannoni, 2006).

One must make some kind of assumption about the stochastic process that the measurement errors follow. Here I investigate two such assumptions. First, in using the Watson measure of fit, the measurement error process is chosen to be the one with the lowest variance while reconciling the autocovariances of the model and data. In this case, measurement errors are assumed to enter all four data series. Second, when I construct the likelihood function, I assume that the data on aggregate consumption growth contain an i.i.d. measurement error. In this case, the other three data series are assumed to be error-free.

4.3 Calibration and driving stochastic processes

The parameters of the model are divided into four groups (see Table 1). The first group is the insurance parameters that are the primary objects of interest. The empirical strategy is to explore how the model's fit changes as these parameters are varied. The next group consists of a single parameter, the discount factor. Changes in the degree of insurance in the economy have important consequences for the precautionary savings motive and so for the asset-income ratio. Therefore, I increase β as I increase the insurance in the economy so that the model's balanced growth path matches the observed mean asset-income ratio, which is measured as the ratio of household net worth from the flow of funds to disposable personal income.

Panel C of Table 1 lists those parameters of the model that are calibrated on the balanced growth path. To calibrate the skill process, I follow Domeij and Heathcote (2004) and Heathcote (2005) to construct a three-point Markov chain that is consistent with estimates of

Symbol	Parameter	Value	Target/Prior
<i>Panel A. Objects of interest</i>			
b^u	Unemployment insurance	0.3	
b^s	Skill insurance	0	
<i>Panel B. Calibrated for each specification</i>			
β	Discount factor	0.958	Aggregate assets $5\times$ annual income.
<i>Panel C. Calibrated on balanced growth path</i>			
χ	Risk aversion	2	
r	Interest rate	0.0075	3% annual interest rate.
$\bar{\lambda}$	Avg. job finding rate	0.679	Mean long-term unemployment.
$\bar{\zeta}$	Avg. high-skill job separation rate	0.037	Mean short-term unemployment.
ζ^s	Differences in separation rate by skill		See Appendix A.
T^s	Skill transition matrix		See Appendix A.
<i>Panel D. Estimated driving processes</i>			
g	Trend income growth	0.004	Uniform[0,1].
ρ^A	Autoregressive coefficient of A	0.951	Beta: mn. = 0.5, var. = 0.04
ρ^λ	Autoregressive coefficient of λ	0.920	Beta: mn. = 0.5, var. = 0.04
ρ^ζ	Autoregressive coefficient of ζ	0.924	Beta: mn. = 0.5, var. = 0.04
σ^A	Standard deviation of ε^A	1.040	Inverse Gamma: mn. = 1, var. = 4
σ^λ	Standard deviation of ε^λ	2.591	Inverse Gamma: mn. = 1, var. = 4
σ^ζ	Standard deviation of ε^ζ	0.432	Inverse Gamma: mn. = 1, var. = 4
σ^T	Standard deviation of ε^T	0.290	Inverse Gamma: mn. = 1, var. = 4

Table 1: Parameter values, targets and priors for the low-insurance economy.

the dynamics of wages and such that the model delivers a reasonable degree of heterogeneity in wealth by matching the Gini coefficient and the Lorenz curve at the 40th percentile. The mapping from parameters to the distribution of wealth depends on the insurance parameters and I calibrate the skill process at the benchmark low insurance values of $b^u = 0.3$ and $b^s = 0$. I choose the skill process to match the dispersion in wages before taxes and transfers and in this respect it is comparable to $w_t s_{i,t}$ in the model as opposed to $y_t(e, s)$ which is inclusive of insurance transfers. The fixed differences in job-separation risk by skill level are calibrated to match data on the differences in unemployment rates across education groups. The low-skill and high-skill groups are small in the calibrated model so it is reasonable to interpret them as high-school dropouts and college graduates, respectively. I then set the job-separation rates for these skill groups to reflect the relative unemployment rates of high-school dropouts and college graduates, respectively. Appendix A has additional details of the calibration procedure for skills and the job-separation rates.

Finally, panel D shows the parameters of the driving stochastic processes, which are estimated from the data on labor income and short-term and long-term unemployment. I employ a Bayesian procedure and set the parameters by finding the posterior mode. To construct the likelihood, I linearize the relevant model equations and apply the Kalman filter as described above. Prior distributions for the parameters are listed in the table although these do not exert a strong influence on the results. A feature of the estimated parameter values that is particularly important for some of the results below is the decomposition of the aggregate wage volatility between the persistent and transitory shocks. I find that the persistent shock explains the vast majority of the variance in the (log) aggregate wage accounting for 99% of the unconditional variance. This finding is consistent with those of Lettau and Ludvigson (2004) who perform a permanent-transitory decomposition of after-tax labor income and find that permanent shocks account for 97% to 100% of the variance depending on the specification and forecast horizon.

5 Comparing low-insurance and full-insurance

I now turn to the first set of results, which contrast the low-insurance and full-insurance economies. In this section, I compare the models and their relation to the data using a number of different tools in order to make clear what features of the models are closer or further from the data. I conclude the section with the two measures of fit described above, which summarize all of this information.

5.1 Impulse response functions

Figure 2 shows the impulse responses of consumption to one-standard deviation shocks to the driving processes for the low-insurance economy ($b^u = 0.3$, $b^s = 0$) and the full-insurance economy ($b^u = 1$, $b^s = 1$). The impulse responses are shown relative to the trend growth rate. Under full insurance, the permanent income hypothesis implies that the impulse responses are entirely flat as the representative agent adjusts consumption once and for all when the shock occurs. With imperfect insurance however, there are extra dynamics coming from various sources. First, some households are borrowing constrained with the result that the dynamics of consumption are the same as the dynamics of income. This is evident in the bottom-right panel, which shows the response to a transitory wage shock as there is a clear spike in consumption on impact. The existence of constrained households also helps to explain why there are hump-shaped dynamics in response to job-finding and -separation shocks as the unemployment rate falls for several periods after the shocks. Second, the endogenously evolving distribution of wealth and precautionary savings generate additional dynamics for aggregate consumption.

Finally, with incomplete markets households are relatively impatient in the sense that the discount factor is less than the inverse of the gross interest rate. This impatience is a reflection of the precautionary savings motive, which allows the model to generate a realistic wealth-to-income ratio despite their impatience. As a result of this impatience, consumption becomes more sensitive to current income.

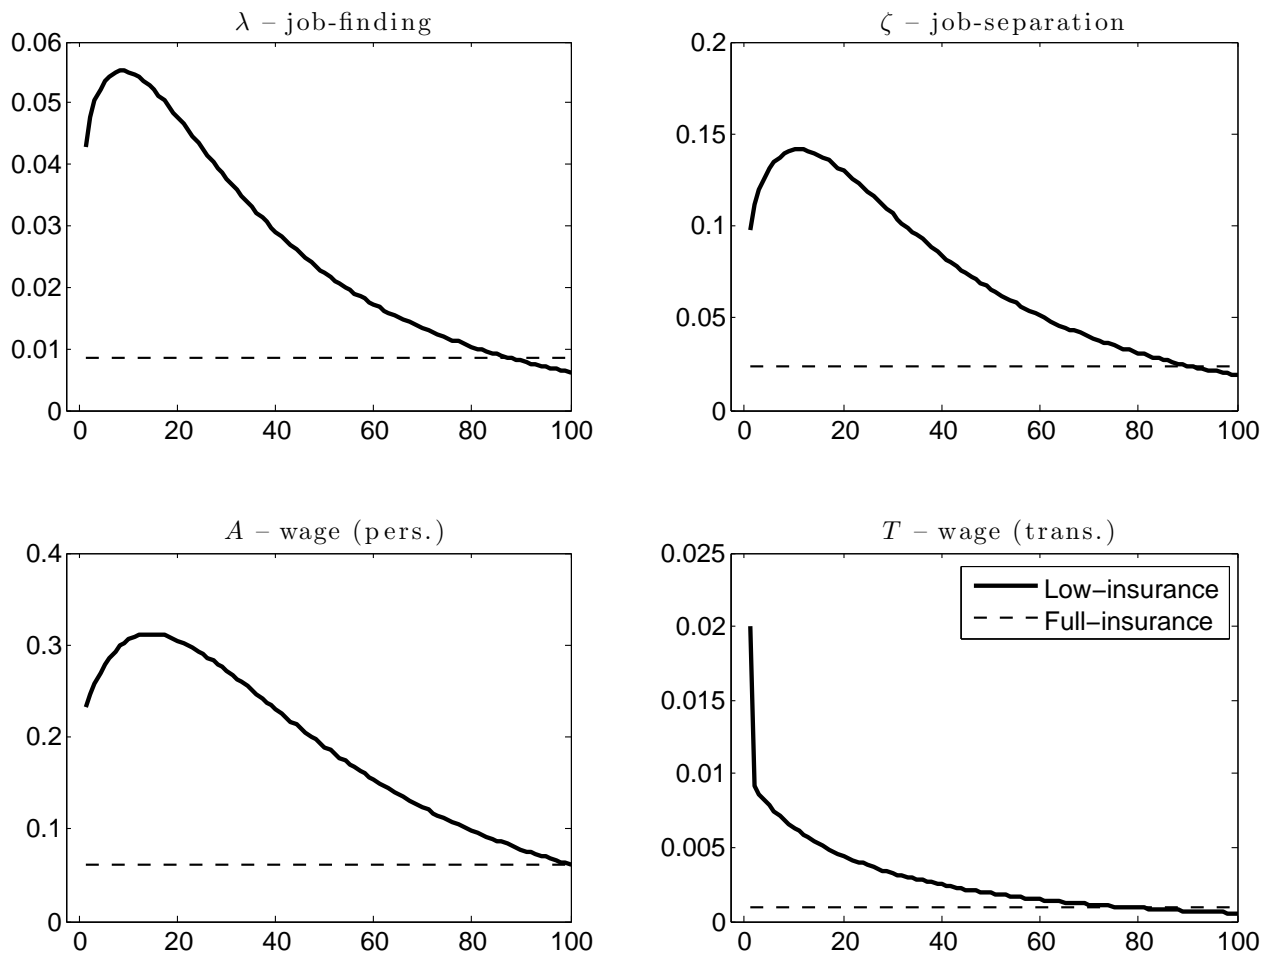


Figure 2: Impulse response of consumption to the four shocks. The plots show $100 \times \log$ change in response to one standard deviation shock. The plot for ζ shows a negative shock to ζ .

5.2 Moments

Table 2 shows selected empirical and model-generated moments. The standard deviation of consumption growth is very low for the full-insurance economy and substantially higher for the low-insurance economy although still only half that in the data. Turning to other data series, one can see that the income process that is fed into the model is somewhat too volatile relative to the data.

Panel B of the table shows that consumption growth data are positively correlated with income growth and negatively correlated with short-term unemployment. This is also true of the low-insurance economy although the correlation with income growth is too high and the correlation with short-term unemployment is rather low. The full-insurance economy differs in that the correlation with the unemployment rate is almost exactly zero.

Finally, panel C of the table shows that consumption growth is positively autocorrelated in the data and to a lesser extent in the low-insurance economy. In the full-insurance economy, however, consumption growth is not autocorrelated due to the random walk behavior of consumption.

5.3 Spectra

A useful way of comparing the full set of second moments of models and data is to plot model and data spectral densities as in King and Watson (1996). The spectral density matrices are helpful in diagnosing the features of the model that drive the likelihood results in section 5.5 because the time-domain likelihood is closely approximated by the distance between the model and data spectral density matrices (see Tkachenko and Qu, 2012).

Figure 3 shows the spectral density of consumption growth that I have estimated from the data along with a 95% confidence interval.⁸ In addition, the figure shows the model-

⁸To estimate the spectral density non-parametrically, I apply a kernel smoother to the periodogram and construct confidence intervals using the methods described in Brockwell and Davis (2006) and implemented by Tkachenko and Qu (2012). The 95% confidence interval is based on the asymptotic distribution of the smoothed periodogram.

<u>A. Standard deviation</u>				
	ΔC_t	ΔY_t	u_t^{short}	u_t^{long}
Data	0.535	1.029	0.921	1.143
Low-insurance	0.261	1.231	0.939	0.948
Full-insurance	0.066	1.231	0.939	0.948

<u>B. Correlation of ΔC_t with</u>			
	ΔY_t	u^{short}	u^{long}
Data	0.271	-0.339	0.064
Low-insurance	0.800	-0.039	-0.013
Full-insurance	0.789	0.001	-0.001

<u>C. Autocorrelation of ΔC_t</u>				
Lags	1	2	3	4
Data	0.407	0.199	0.130	0.062
Low-insurance	0.099	0.085	0.074	0.068
Full-insurance	-0.001	0.000	0.001	0.005

Table 2: Moments from model and data. Model moments are calculated from 10,000 quarters of simulated data.

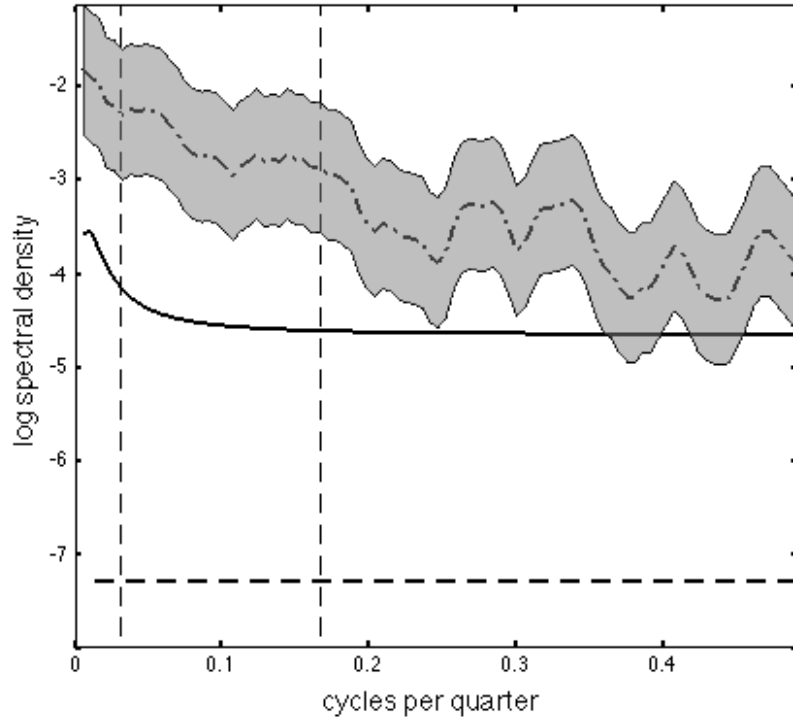


Figure 3: Estimated and model-implied spectral densities. Dash-dot line refers to data with 95% confidence interval (shaded area), solid line refers to low-insurance (incomplete markets) economy and dashed line refers to full-insurance (representative agent) economy. Vertical lines show the range of business cycle frequencies: $1/32$ cycles per quarter and $1/6$ cycles per quarter.

implied spectral density of consumption growth from the low-insurance and full-insurance economies without measurement error. The area under the spectral density of a time series is equal to the unconditional variance of the series. The spectral density therefore shows how fluctuations at different frequencies contribute to the overall variance of the series.

The data show that consumption growth has more power at low frequencies than at high frequencies. Low-frequency movements in consumption growth are well-known in the contexts of a declining personal saving rate and an increasing share of consumption in GDP. For the full-insurance economy, the random walk behavior of consumption results in con-

sumption growth that is white noise and so the spectral density is flat across frequencies. In addition, the spectral density is far too low at all frequencies indicating that the model generates consumption that is too smooth at all frequencies. For the low-insurance economy, the spectral density has some of the downward sloping shape that is seen in the data and it has much more power at all frequencies than does the low-insurance model, but it is still too low compared to the data except possibly at high frequencies. The downward-sloping spectral density in the low-insurance economy is a reflection of the distribution of wealth contributing slow-moving state variables to the dynamics of the system.

Plots (not shown here) of the spectral densities of income growth and the unemployment rates show that the driving processes do a fairly good job of matching the volatility of these series across different sets of frequencies.

5.4 Watson's measure of fit

Watson's measure of fit is one way of measuring the distance between model and data spectral densities or equivalently autocovariances. Table 3 shows the ratio of the measurement error variance to the data variance for each of the four observable variables. Larger values of this ratio indicate that larger measurement errors are needed to reconcile the model's autocovariances with those of the data. Watson's procedure selects the measurement error process by minimizing its variance, this ratio is a lower bound on what is needed to bridge the gap between the model and data.

In panel A., the measurement errors for consumption growth have a variance equal to 89% of the data under full insurance while for the low-insurance economy it is much lower at 59%. Turning to the other variables, the model fits rather well across all frequencies. While the processes for ΔY , u^{short} , and u^{long} are the same in the two models, the measurement error processes are not identical because the calculations allow for movements in consumption to be partly explained by measurement errors in other series. However, the results across models for these series are generally very similar.

The other panels of the table show the models' fit at different frequency bands. The

performance of both models in fitting consumption growth at low frequencies is similar to their performance across all frequencies. At business cycle frequencies, the fit of the low-insurance economy deteriorates somewhat. Both models fit consumption growth best at high frequencies as compared to other frequencies and the ratio drops to 50% for the low-insurance economy.

5.5 Likelihood

In order to compute the likelihood of the data, I need to make specific assumptions about the measurement error process. As I demonstrate below, measurement errors can complicate the identification of the insurance parameters from the likelihood function. To see this, notice that the results so far show that the low-insurance economy comes closer than the full-insurance economy to matching the data on consumption growth across several dimensions with the most striking difference between the models being the standard deviation of consumption growth or equivalently the height of the spectral density in Figure 3. Adding i.i.d. measurement error to consumption growth will raise the spectral density uniformly across all frequencies. The statistical models that result from adding i.i.d. measurement error to the low-insurance and full-insurance economies will therefore perform relatively similarly in this dimension, which was previously what allowed us to distinguish the two economies.

One way to proceed is to restrict the magnitude of measurement error that one is willing to consider. Table 4 shows the likelihood of the joint dynamics of the four series for different values of the standard deviation of measurement error for both the low-insurance and full-insurance economies. If the variance of the measurement error is low, then the likelihood under full-insurance is substantially below that under low-insurance. However, at the variances that maximize the likelihood function, the likelihood values are very similar as was predicted. Allowing for autocorrelated measurement errors does not change this conclusion.

One might ask how Table 4 can be reconciled with the Watson measure of fit? The two measures of fit take very different views of measurement error: in the case of maximizing the likelihood function, measurement errors are a completely legitimate source of variation

<u>A. All frequencies</u>				
	ΔC_t	ΔY_t	u_t^{short}	u_t^{long}
Low-insurance	0.592	0.127	0.225	0.266
Full-insurance	0.882	0.118	0.218	0.264
<u>B. Low frequencies</u>				
	ΔC_t	ΔY_t	u_t^{short}	u_t^{long}
Low-insurance	0.578	0.659	0.221	0.195
Full-insurance	0.928	0.597	0.222	0.194
<u>C. Business-cycle frequencies</u>				
	ΔC_t	ΔY_t	u_t^{short}	u_t^{long}
Low-insurance	0.678	0.512	0.203	0.304
Full-insurance	0.913	0.485	0.203	0.304
<u>D. High frequencies</u>				
	ΔC_t	ΔY_t	u_t^{short}	u_t^{long}
Low-insurance	0.502	0.029	0.511	1.049
Full-insurance	0.828	0.026	0.351	0.945

Table 3: Watson’s measure of fit. Each entry shows the ratio of measurement error variance relative to data variance for the listed variable and frequency band.

σ^v	Low-insurance	Full-insurance	Difference
0.1	-2328	-2683	356
0.2	-876	-900	24
0.3	-613	-615	3
0.4	-543	-543	0
0.5	-527	-526	-1
0.5269	-526		
0.5280		-526	

Table 4: Log likelihood as a function of measurement error standard deviation (σ^v). The bottom two rows show the respective maximum likelihood estimates of σ^v .

in the data and explaining the data through measurement errors makes the model just as “successful” as explaining the data through structural innovations. By contrast, Watson’s measure treats measurement errors as residuals whose contribution should be minimized. Watson’s measure is the extent to which measurement errors cannot be avoided and in this sense it penalizes them heavily. In the context of the likelihood function, penalizing measurement errors can be interpreted as a prior that the measurement error variance should be small. Given the shape of the likelihood shown in Table 4, if one applies a strong prior that the measurement error variance should be small one will arrive at a marginal likelihood that favors the low-insurance economy.

The results in Table 4 do not favor the low-insurance economy as much as the Watson measure of fit in part because it assumes the income and labor market data are measured without error. The low-insurance model benefits more from including measurement errors in those series because errors in those series have more explanatory power for aggregate consumption. The difference in explanatory power is a reflection of the impulse response functions: measurement errors in income are better able to reconcile the consumption growth data if a given change in the income data due to measurement error leads to a larger change in consumption growth.

Another way of proceeding is to use the model with measurement errors to generate smoothed estimates of the structural shocks and then form a predicted consumption growth series implied by these shocks in the absence of measurement errors. Fernández-Villaverde and Rubio-Ramírez (2007) perform a similar exercise to summarize what fraction of the variation in the data is accounted for by structural shocks rather than measurement errors. For the low-insurance economy I find that the predicted consumption growth series has a standard deviation that is 46% of that for the data and has correlation with the actual consumption growth series equal to 0.29. For the full-insurance economy, the ratio of standard deviations is 15% and the correlation is 0.21. Therefore it appears the low-insurance economy does not just generate a higher standard deviation, but also is better able to explain the joint dynamics of consumption growth and the other series.

My interest is in how the structural model is able to explain consumption growth dynamics, which is best captured by the likelihood function evaluated with a low standard deviation of measurement error. Therefore, I set $\sigma^v = 0.1$ in most of what follows.

6 Partial insurance

The previous section compared two extreme insurance arrangements, which are just two points in a range of possible insurance arrangements. The parameters b^s and b^u allow me to explore other levels of insurance against skill shocks and unemployment shocks, respectively. In investigating the possibility of partial insurance, there is a similarity to the work of Blundell et al. (2008) and Heathcote et al. (2012) who have investigated the possibility of partial insurance using panel data on household income and consumption and have found support for partial insurance against permanent income shocks. Here, however, I ask what can be learned about partial insurance from aggregate time series data.

To investigate partial insurance, I explore how the model fit changes with b^s and b^u . For each set of insurance parameters, I recalibrate the discount factor, β , to match the wealth-income ratio as before while other parameter values are left unchanged. Table 5 shows Watson’s measure of fit and the log likelihood of the data at selected insurance arrangements.

			Watson's measure of fit				Std. dev.	log \mathcal{L}
	b^s	b^u	ΔC_t	ΔY_t	u_t^{short}	u_t^{long}	ΔC_t	($\sigma^v = 0.1$)
<i>Low-insurance</i>	0.0	0.3	0.592	0.127	0.225	0.266	0.261	-2328
	0.3	0.3	0.561	0.128	0.226	0.266	0.293	-2242
<i>Best fit</i>	0.5	0.3	0.561	0.128	0.227	0.265	0.296	-2211
	0.8	0.3	0.681	0.123	0.223	0.265	0.187	-2359
	0.0	0.6	0.607	0.127	0.224	0.266	0.257	-2337
	0.3	0.6	0.578	0.128	0.225	0.265	0.289	-2262
	0.5	0.6	0.580	0.128	0.226	0.265	0.290	-2242
	0.8	0.6	0.704	0.122	0.222	0.265	0.178	-2402
	0.0	1	0.619	0.127	0.222	0.265	0.254	-2335
	0.3	1	0.589	0.128	0.224	0.265	0.286	-2266
	0.5	1	0.591	0.128	0.225	0.265	0.287	-2251
<i>Full-insurance</i>	1	1	0.882	0.118	0.218	0.264	0.066	-2683

Table 5: Watson's measure of fit and standard deviation of consumption growth for alternative insurance arrangements. All frequencies are included in the calculations.

As one would expect from above, only the fit to ΔC is sensitive to the insurance arrangement. Variation in the unemployment insurance parameter b^u appears to have little influence on the model fit, however, variation in the skill-insurance parameter, b^s , has a strong influence. The fit is not monotonically deteriorating as we move from low values of b^s to high values. Instead, $b^s = 0.5$ provides a fit that is superior to $b^s = 0$. The non-monotonicity in fit is related to a non-monotonicity in the standard deviation of consumption growth. As the level of insurance changes, there is a trade-off between increasing the risk that households face, which tends to raise their marginal propensity to consume, and reducing the resources of low-income households. If low-income households have few resources in all periods, then their consumption must be a small part of the aggregate and their choices cannot have a large impact on the dynamics of aggregate consumption.

7 Matching the response to tax rebates

The standard incomplete markets model has recently been criticized by Kaplan and Violante (2011) for failing to match empirical evidence on consumption responses to transitory income fluctuations without resorting to a counterfactual distribution of household net worth. The data show a greater sensitivity of consumption to current income than the model is able to generate. The model could generate more constrained households whose consumption will move one-for-one with income, but to do so requires a counterfactually large number of households with low net worth. There are several empirical measures of the sensitivity of consumption to transitory income. One such measure comes from the staggered timing of fiscal stimulus payments during recent recessions in the United States. The timing of these payments of several hundred dollars was effectively randomized leading to comparable groups that only differ in the timing of the payment. Johnson et al. (2006) assess the sensitivity of the change in consumption to the receipt of the payment by regressing the change in consumption of household i at time t on the transfer received by that household at that date. The resulting coefficient on the tax rebate is estimated to be 20% to 40%. Further research using new methods and new data has produced estimates of this rebate coefficient

in the neighborhood of 20% (Misra and Surico, 2011; Parker et al., 2013).

Kaplan and Violante (2011) show that the standard incomplete markets model is inconsistent with this evidence on fiscal stimulus payments and generates a response of consumption that is just 1.8% of the transfer in their model. Conducting the same experiment in the low-insurance economy presented above, I find an even lower response of just 0.1%.⁹

The rebate coefficient differs from the MPC because it compares the change in consumption in a period in which a transfer is received to the change in consumption in a period in which a transfer has already been received or is yet to be received. The MPC, by contrast, would compare consumption when the transfer is received to a counterfactual without a transfer in any period. The MPC out of unanticipated transitory income is 5.6% in the low-insurance economy. This MPC is also below empirical estimates of MPCs, which are generally in the range of 20% to 50% with many estimates near 20%.¹⁰

Given the inconsistency between the model and microeconomic evidence on the response of consumption to income changes, one might ask whether a model that performed better in these dimensions would give a more accurate description of the dynamics of aggregate consumption. To explore this, I conduct an experiment that is motivated by Kaplan and Violante's illiquid asset model. Kaplan and Violante show that a model that includes illiquid assets, which make net worth less effective for smoothing consumption, is able to generate more constrained consumption behavior and responses to transfers as high as 21%. I incorporate the essence of the Kaplan-Violante model by introducing an adjustment cost on household assets that reduces their effectiveness for consumption smoothing. I then calibrate this adjustment cost to match a 25% response to the receipt of a fiscal stimulus payment.

The Kaplan-Violante model involves two state variables to summarize a household's

⁹To calculate the rebate coefficient I follow Kaplan and Violante (2011) and assume the economy is in steady state when one group of households receives a transfer and another group of households learns that they will receive the transfer in the following period. The two groups are identical except for the timing of the transfer and initially represent a random sample from the model's steady state distribution over assets and income. I simulate the consumption behavior of these two groups and then calculate the rebate coefficient using the same regression as in Johnson et al. (2006) described above.

¹⁰Examples include Hall and Mishkin (1982); McCarthy (1995); Lusardi (1996); Parker (1999).

	Watson's measure of fit				Std. dev.	log \mathcal{L} ($\sigma^v = 0.1$)
	ΔC_t	ΔY_t	u_t^{short}	u_t^{long}		
Baseline low-insurance economy	0.592	0.127	0.225	0.266	0.261	-2328
With asset adjustment cost	0.595	0.129	0.227	0.266	0.303	-2478

Table 6: Watson's measure of fit, standard deviation of consumption growth, and log likelihood with and without adjustment costs. All frequencies are included in the calculations.

portfolio. The methods I have described might in principle be applied to a model with two continuously-distributed state variables, but there are a number of challenges in doing so. Rather than include a second state variable for household assets, I assume a household that enters the period with assets in amount a and wishes to save a' for the next period must pay a cost $\Gamma(a, a'; e^z) \geq 0$, where e^z is the trend level of income. I assume that Γ is homogeneous of degree one in all three arguments, which allows for a balanced growth path. By including e^z as an argument of the adjustment cost function, I allow for non-homogeneity in the cross-section while still preserving the balanced growth path in the aggregate. In particular, the portfolios of low-wealth households are dominated by liquid assets and so it is plausible that adjustment costs are less relevant for these households (Campbell, 2006). Moreover, there is a difficulty with specifying an adjustment cost function that is linearly homogeneous in a and a' alone as households can have assets near zero. To see the difficulty, consider the commonly used quadratic adjustment cost function $\frac{\Phi}{2} \frac{(a'-a)^2}{a}$. The cost of increasing savings by, say, one unit goes to infinity as assets go to zero. Therefore I use the modified adjustment cost function $\Gamma(a, a'; e^z) = \frac{\Phi_1}{2} \frac{(a'-a)^2}{a + \Phi_2 e^z}$, where Φ_2 is a parameter that reduces the importance of adjustment costs at low wealth levels.

Φ_1 controls the strength of the asset adjustment cost function. I choose this parameter so that the model generates a rebate coefficient of 25%. At this calibration, the MPC out of unanticipated transitory income is 20%. To calibrate Φ_2 , it is useful to rewrite the adjustment cost function as $\Gamma(a, a'; e^z) = \frac{\Phi_1}{2} \frac{(a'-a)^2}{a} \frac{a}{a + \Phi_2 e^z}$, where the first two terms are the linearly homogeneous adjustment cost for illiquid assets and the third term, which

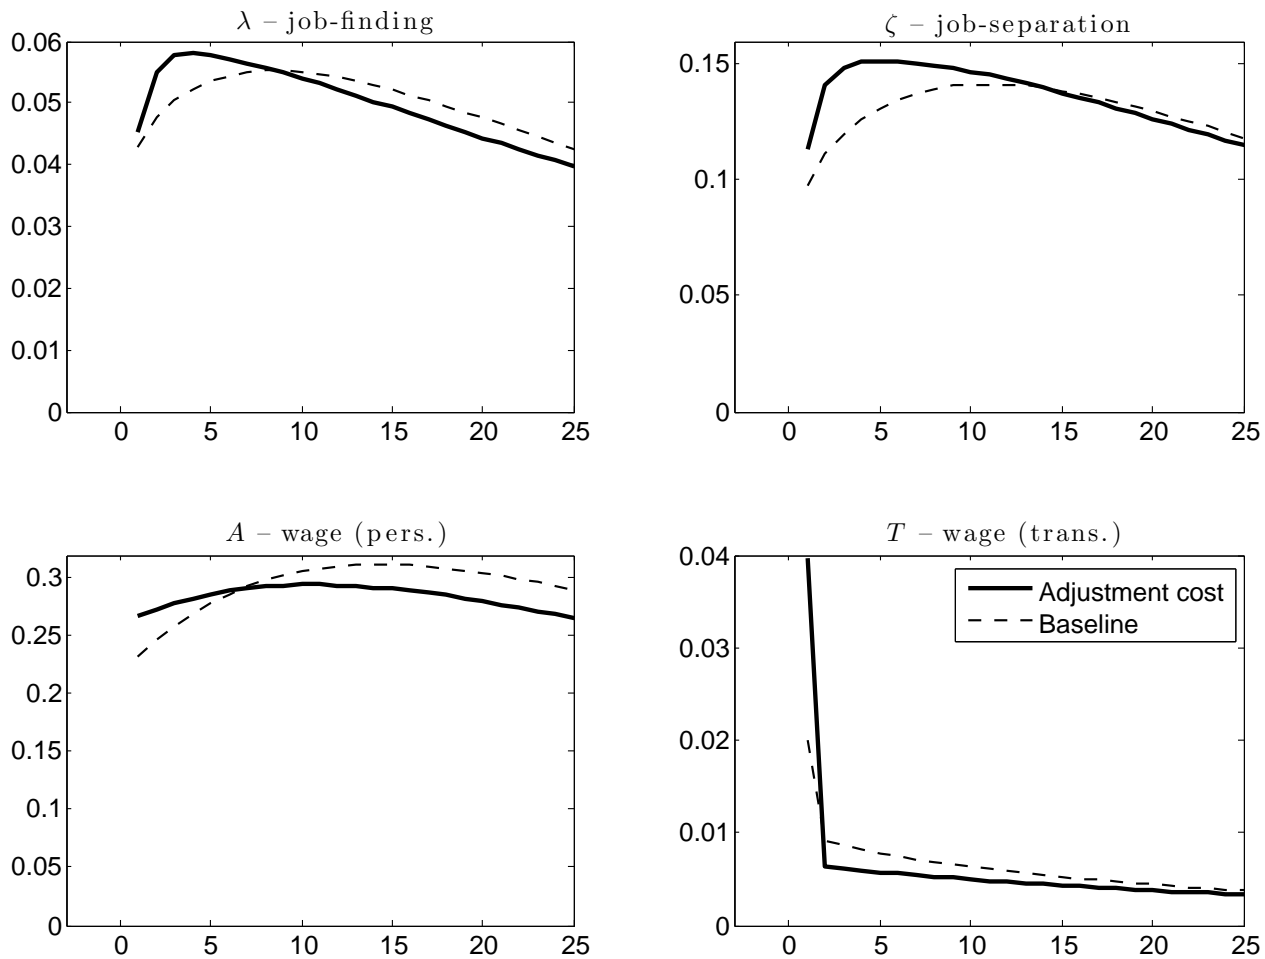


Figure 4: Impulse response of consumption to the four shocks. The plots show $100 \times$ log change in response to one standard deviation shock.

takes values in $[0, 1)$, is the extent to which these adjustment costs apply to a household with wealth a . Figure 3 in Campbell (2006) shows that transactions accounts are only an important part of total household assets for very low wealth households. Therefore I set $\Phi_2 = 1.5$ or approximately the average labor income per quarter of an employed household. The interpretation is that a household with assets equal to one quarter's income will face half the "full" transaction cost. In addition to these two new parameters, I recalibrate the low-insurance economy to match the same targets as in Table 1.¹¹

Table 6 compares the model with the adjustment costs to the low-insurance economy without the adjustment costs and a 5.6% marginal propensity to consume. It appears that the adjustment cost either has little effect on the dynamics of consumption or results in a worse fit to the data depending on which metric one uses. To understand why this is, it is helpful to look at the impulse response functions in Figure 4. Notice that the impact of a transitory wage shock doubles when the adjustment cost is included and it is this response to the transitory shock that the fiscal stimulus experiment is capturing. However, the responses to other shocks, which are more persistent, are not much affected. Suppose that income shocks are permanent and households follow the permanent income hypothesis consumption rule. In that case, consumption adjusts immediately in response to changes in income and assets do not change at all rendering the adjustment costs irrelevant. This logic provides the intuition for why the adjustment costs affect the response to the transitory shock more than the response to the persistent shock although the case being considered is not this extreme. Finally, notice that the transitory wage shock accounts for a small share of the variance in total income so the change in this impulse response function has a limited influence on the overall dynamics of consumption. This last point follows from the low estimate of the variance of the transitory shock as reflected by the overall scale of the impulse responses in Figure 4.

Table 6 is an example of the usefulness of the full-information analysis. The model with

¹¹The only appreciable change in parameter values is that the discount factor rises to 0.962. Solving the model with the adjustment costs involves the same steps as used to solve the original model with the difference being that there are new terms that enter the household's Euler equations and budget constraints.

the adjustment cost generates more volatility in consumption growth and would seem to bring the model closer to the data in this regard. But it appears that this extra volatility in consumption growth does not fit well with the dynamics of the other series and the model with the adjustment cost results in a substantially lower likelihood.

These findings show that the evidence on the response of consumption to transitory income fluctuations does not necessarily invalidate the standard incomplete markets model for the purpose of understanding the time series behavior of aggregate consumption. However, this evidence might be of first-order importance if one is specifically interested in the response to a transitory shock, such as those generated by fiscal stimulus payments.

8 Alternative explanations

The full-insurance version of the model is rather simplistic so it is worth asking how the improvements in model fit for the low-insurance economy compare to those that can be achieved by enriching the environment with other factors while maintaining the complete markets assumption. In doing so, I show that certain modifications of the model can greatly improve the performance of the full-insurance economy. Nevertheless, there is still a considerable discrepancy between the model and the data. These findings suggest that incorporating market incompleteness into richer models of the business cycle may be a useful way forward.

8.1 The role of interest rates

One of the main objections that could be raised to the analysis so far is that it assumes a constant interest rate. Changes in expected real interest rates have direct implications for the consumption-savings problem. Moreover, the interest rate is the key channel through which general equilibrium effects might alter household consumption choices. An important question is therefore whether the consumption dynamics would be more realistic if the model were to include the observed dynamics of interest rates? Doing so will reveal whether the low volatility of consumption growth documented above can be attributed to the constant

	Watson's measure of fit					Std. dev.	log \mathcal{L}
	ΔC_t	ΔY_t	u_t^{short}	u_t^{long}	r_t	ΔC_t	($\sigma^v = 0.2$)
Baseline low-insurance	0.592	0.127	0.225	0.266		0.261	-876
Baseline full-insurance	0.882	0.118	0.218	0.264		0.066	-900
Stoch. interest rates	3.349	0.160	0.239	0.279	0.399	1.344	-6200
Labor complementarity	0.599	0.118	0.226	0.271		0.193	-779

Table 7: Watson's measure of fit, standard deviation of consumption growth and log-likelihood for alternative complete-markets environments. All frequencies are included in the calculations.

interest rate and it will also incorporate the real-world (as opposed to model-generated) general equilibrium effects to the extent that they act through interest rates.

To answer this question, I expand the data to be explained to include expected real interest rates, which I construct by regressing ex post real interest rates on lagged information as described by Mishkin (1981).¹² I then fit an AR(1) process for expected real interest rates in the same manner that I constructed the driving processes above.¹³

The observed fluctuations in interest rates are not useful in explaining consumption growth dynamics. Table 7 shows results for the full-insurance economy with a constant interest rate and with the stochastic interest rate. One can see that the fit of consumption growth deteriorates substantially with the stochastic interest rates. Watson's measure of fit is actually much above 1, which is possible if the model's predictions are negatively correlated with the data so the measurement errors need to explain the entirety of the data and then some. This is in fact what is happening. As in section 5.5, I use smoothed estimates of the shocks to the driving processes (now including interest rates) and then use these to

¹²For the ex post real interest rate I use the difference between the 3-month Treasury bill rate from date t and the growth of the GDP deflator from date t to $t + 1$. I then use predicted values from regressing this ex post rate on lags of nominal interest rates and inflation rates dated t and earlier.

¹³The resulting process has an autoregressive coefficient of 0.941 and a standard deviation of 0.194. I continue to assume a mean of 0.0075.

generate a predicted consumption growth series. The correlation of this series with the actual consumption growth data is -0.21 . The same experiment with a constant interest rate produces a correlation of 0.21 . So while interest rate volatility can generate substantial volatility in consumption growth, it is unrelated to the consumption growth data. This disconnect between consumption growth and interest rates is not a new finding as it is what drives the low estimates of the elasticity of intertemporal substitution in the literature (Hall, 1988; Campbell and Mankiw, 1989).

8.2 Non-separable preferences

In the model considered here, preferences are separable across time and across goods and leisure. Would relaxing these assumptions help the model fit the consumption growth data? The most commonly used form of intertemporal non-separability is to include external habit formation in the preferences. Typically, the benefit of habit formation is to smooth out consumption and generate hump-shaped dynamics. Notice that the incomplete markets version of the model here is able to generate some amount hump-shape dynamics in consumption even without habits. Habit formation, however, has the effect of reducing consumption volatility. I have experimented with adding habit formation to the model with stochastic interest rates and find that it can deliver the desired level of consumption volatility, but the predicted consumption growth series is still negatively correlated with the data.

On the other hand, a non-separability between consumption of goods and leisure has the potential to generate pro-cyclical consumption volatility if households prefer to consume more when they are working longer hours in an expansion—i.e. consumption and labor supply are complements.

To explore the role of non-separable preferences in generating consumption volatility, I consider the following specification of instantaneous utility

$$\frac{c_t^{1-\chi}}{1-\chi} \exp(\gamma L_t),$$

where L_t is hours worked. In modifying the preferences I am only interested in the impact of labor-consumption non-separabilities on consumption volatility. In particular, I continue to

assume that labor supply is exogenous. The parameter γ now controls the extent to which labor supply and consumption are complementary. This parameter is difficult to estimate based on microeconomic data and so I instead estimate by maximum likelihood from the aggregate time series data and find a point estimate for γ equal to 0.766.

Table 7 shows that the consumption-labor complementarity is very useful in explaining consumption growth and greatly increases the level of consumption volatility in the full insurance economy. According to the Watson measure the economy with the complementarity performs similarly to the low-insurance economy and according to the likelihood function it performs better.

From these results, it would seem that labor-consumption complementarities are a useful component of the model, however, some caution is required in interpreting these results as the estimated complementarity may reflect factors besides preferences. In particular, a deterioration of labor market conditions that increases idiosyncratic risk may lead to a drop in consumption through the precautionary savings channel. If this channel is not captured by the model, one might attribute the drop in consumption to preferences. One check on whether this explains the results is to compare the predicted consumption series from the low-insurance economy to that from the economy with non-separable preferences. Each series is generated by using the smoothed shocks without measurement error. The predicted consumption series have a correlation of just 0.09, which suggests that success of the two models is driven by distinct factors.

9 Conclusion

This paper shows how newly-developed methods for analyzing incomplete markets models can be combined with many of the tools of empirical macroeconomics and uses these tools to analyze the role of idiosyncratic risk in the fluctuations in aggregate consumption. Three main substantive conclusions emerge: first, the incomplete markets version of the model provides a much more realistic description of the data in that smaller residuals or “measurement errors” are needed to reconcile the model with the data. This finding stems primarily

from the larger response of consumption to changes in income leading to a higher variance of consumption growth. Second, a model with partial insurance against idiosyncratic skill shocks provides a better fit to the data than a version of the model without skill-insurance that is close to the standard incomplete markets model. This finding from aggregate data mirrors results in the literature from microeconomic data. The third main conclusion is that the model's failure to match empirical estimates of the response to transitory income changes does not necessarily invalidate it as a model of consumption dynamics in general. I show that the MPC out of transitory income changes can be altered without much affecting the way consumption responds to persistent shocks. As persistent income shocks account for most of the dynamics of consumption according to my estimates, the overall dynamics of consumption are not very different in two versions of the model with quite different MPCs out of transitory income. Of course, accurately predicting the response of consumption to transitory income changes is crucial to some applications such as predicting the effects of fiscal stimulus payments.

References

- Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–84.
- Aiyagari, S. R. and McGrattan, E. R. (1998). The optimum quantity of debt. *Journal of Monetary Economics*, 42(3):447–469.
- An, S. and Schorfheide, F. (2007). Bayesian analysis of dsge models. *Econometric Reviews*, 26(2-4):113 – 172.
- Antoulas, A. C. (2009). *Approximation of large-scale dynamical systems*, volume 6. Society for Industrial and Applied Mathematics.
- Attanasio, O. and Davis, S. J. (1996). Relative wage movements and the distribution of consumption. *Journal of Political Economy*, 104(6):pp. 1227–1262.
- Berger, D. and Vavra, J. (2012). Consumption dynamics during recessions. Technical report, Working Paper.
- Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921.
- Boivin, J. and Giannoni, M. (2006). Dsge models in a data-rich environment. Technical report, National Bureau of Economic Research.
- Brockwell, P. J. and Davis, R. A. (2006). *Time series: theory and methods*. Springer, New York, second edition.
- Campbell, J. Y. (2006). Household finance. *Journal of Finance*, 61:1553–1604.
- Campbell, J. Y. and Mankiw, N. G. (1989). Consumption, income and interest rates: Reinterpreting the time series evidence. In *NBER Macroeconomics Annual 1989, Volume 4*, pages 185–246. MIT Press.

- Carroll, C. D. (1992). The buffer-stock theory of saving: Some macroeconomic evidence. *Brookings papers on economic activity*, 1992(2):61–156.
- Challe, E. and Ragot, X. (2013). Precautionary saving over the business cycle. *Paris School of Economics manuscript*.
- Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of political economy*, pages 957–976.
- Deaton, A. (1991). Saving and liquidity constraints. *Econometrica*, 59(5):1221–1248.
- Domeij, D. and Heathcote, J. (2004). On the distributional effects of reducing capital taxes. *International economic review*, 45(2):523–554.
- Elsby, M. W., Hobijn, B., and Sahin, A. (2010). The labor market in the great recession. *Brookings Papers on Economic Activity*, pages 1–48.
- Fernández-Villaverde, J. and Rubio-Ramírez, J. F. (2007). Estimating macroeconomic models: A likelihood approach. *The Review of Economic Studies*, 74(4):1059–1087.
- Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1):pp. 47–89.
- Guvenen, F. and Smith, Jr., A. A. (2010). Inferring labor income risk from economic choices: An indirect inference approach. Technical report, Yale University.
- Hall, R. E. (1988). Intertemporal substitution in consumption. *Journal of Political Economy*, pages 339–357.
- Hall, R. E. and Mishkin, F. S. (1982). The sensitivity of consumption to transitory income: Estimates from panel data on households. *Econometrica*, 50(2):461–481.
- Heathcote, J. (2005). Fiscal policy with heterogeneous agents and incomplete markets. *Review of Economic Studies*, 72(1):161–188.

- Heathcote, J., Perri, F., and Violante, G. L. (2010). Unequal we stand: An empirical analysis of economic inequality in the united states, 1967-2006. *Review of Economic Dynamics*, 13(1):15 – 51.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2012). Consumption and labor supply with partial insurance: An analytical framework. Research Department Staff Report 432, Federal Reserve Bank of Minneapolis.
- Ingram, B. F., Kocherlakota, N. R., and Savin, N. E. (1994). Explaining business cycles: A multiple-shock approach. *Journal of Monetary Economics*, 34(3):415–428.
- Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5):1589–1610.
- Jung, P. and Kuester, K. (2011). The (un)importance of unemployment fluctuations for the welfare cost of business cycles. *Journal of Economic Dynamics and Control*, 35(10):1744 – 1768.
- Kaplan, G. and Violante, G. L. (2011). A model of the consumption response to fiscal stimulus payments. Working Paper 17338, National Bureau of Economic Research.
- King, R. G. and Watson, M. W. (1996). Money, prices, interest rates and the business cycle. *The Review of Economics and Statistics*, 78(1):pp. 35–53.
- Krusell, P. and Smith, Jr., A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896.
- Leeper, E. M. and Sims, C. A. (1994). Toward a modern macroeconomic model usable for policy analysis. In *NBER Macroeconomics Annual 1994, Volume 9*, pages 81–140. MIT Press.
- Lettau, M. and Ludvigson, S. C. (2004). Understanding trend and cycle in asset values: Reevaluating the wealth effect on consumption. *American Economic Review*, 94(1):276–299.

- Lusardi, A. (1996). Permanent income, current income, and consumption: Evidence from two panel data sets. *Journal of Business & Economic Statistics*, 14(1):81–90.
- McCarthy, J. (1995). Imperfect insurance and differing propensities to consume across households. *Journal of Monetary Economics*, 36(2):301–327.
- McGrattan, E. R. (1994). The macroeconomic effects of distortionary taxation. *Journal of Monetary Economics*, 33(3):573–601.
- McKay, A. and Reis, R. (2013). The role of automatic stabilizers in the u.s. business cycle. Working Paper 19000, National Bureau of Economic Research.
- Mishkin, F. S. (1981). The real interest rate: An empirical investigation. *Carnegie-Rochester Conference Series on Public Policy*, 15:151 – 200.
- Misra, K. and Surico, P. (2011). Heterogeneous responses and aggregate impact of the 2001 income tax rebates. Technical report, CEPR.
- Parker, J. A. (1999). The reaction of household consumption to predictable changes in social security taxes. *American Economic Review*, 89(4):959–973.
- Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*. Forthcoming.
- Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, 33(3):649–665.
- Reiter, M. (2010). Approximate and almost-exact aggregation in dynamic stochastic heterogeneous-agent models. Institute for Advanced Studies, Economics Series 258.
- Rouwenhorst, K. G. (1995). *Asset Pricing Implications of Equilibrium Business Cycle Models.*, chapter 10, pages 294 – 330. Princeton:.
- Sargent, T. J. (1989). Two models of measurements and the investment accelerator. *Journal of Political Economy*, pages 251–287.

- Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *The American Economic Review*, 95(1):pp. 25–49.
- Sims, C. A. (2002). Solving linear rational expectations models. *Computational Economics*, 20(1-2):1 – 20.
- Smets, F. and Wouters, R. (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review*, 97(3):586 – 606.
- Souleles, N. S. (1999). The response of household consumption to income tax refunds. *The American Economic Review*, 89(4):947–958.
- Tkachenko, D. and Qu, Z. (2012). Frequency domain analysis of medium scale dsge models with application to smets and wouters (2007). In *Advances in Econometrics: DSGE Models in Macroeconomics-Estimation, Evaluation, and New Development*, volume 28, pages 319–385.
- Watson, M. W. (1993). Measures of fit for calibrated models. *Journal of Political Economy*, 101(61).

A Further details of calibration procedure

A.1 Skills

I calibrate the Markov chain for skills using the procedure described in Domeij and Heathcote (2004). The skill distribution has three points and the mean log skill is normalized to zero. The skill transition matrix is restricted to have the form

$$T^s = \begin{bmatrix} T_{11}^s & 1 - T_{11}^s & 0 \\ \frac{1 - T_{22}^s}{2} & T_{22}^s & \frac{1 - T_{22}^s}{2} \\ 0 & 1 - T_{11}^s & T_{11}^s \end{bmatrix},$$

which implies that households do not transit directly from low skill to high skill, the number of workers in the low and high skill states is the same, and the probabilities of moving into (out of) low and high skill is the same. After these restrictions there are four parameters that need to be calibrated: two of the three skill levels, T_{11}^s , and T_{22}^s . The four targets are: the autocorrelation and cross-sectional dispersion of log wages, Gini coefficient for wealth, and the share of wealth held by the poorest 40% of the population.

The parameters needed to match these moments differ from those chosen by Domeij and Heathcote (2004) because the other features of the model affect the distribution of wealth and because here a model period corresponds to one quarter rather than one year. Domeij and Heathcote calibrate their model to generate a variance of log labor productivity of 0.26 and an annual autocorrelation of 0.9, which are inline with commonly-used estimates from PSID data. To adapt this strategy to a quarterly model period, I simulate an AR(1) process at a quarterly frequency and then aggregate to annual averages. I choose the parameters of the quarterly process so that the aggregated data matches an annual autocorrelation of 0.9, but I increase the variance of the process to reflect wages exclusive of taxes and transfers. I set the cross-sectional variance to 0.4, which is inline with estimates for male wages in the CPS (Heathcote et al., 2010). The quarterly process has an autocorrelation of 0.96 and a standard deviation of the innovation of 0.17. This procedure results in $T_{11}^s = 0.9618$ and $T_{22}^s = 0.9932$. The three skill levels are 0.025, 0.960, and 63.0.

	< High school	High school	Some college	College
Labor force	0.10	0.31	0.27	0.31
Unemployment	0.20	0.36	0.26	0.18

Table 8: Shares of labor force and unemployment pool by education 1992 - 2012. Data are for civilians ages 25 and over and are calculated from data on the number of persons in the labor force and unemployment as reported by the Bureau of Labor Statistics.

A.2 Unemployment risk

The model allows the job-separation rate to vary by skill level using data on unemployment by education as a guide to the cross-sectional variation in unemployment risk. Table A.2 shows that high school dropouts account for 20% of the unemployment pool, but only 10% of the labor force implying that their unemployment rate is twice that of the economy as a whole. The first step in the calibration of the skill-specific job-separation rates is to construct average unemployment rates for the three skill groups. The overall unemployment rate is chosen to match the average in the data, which is 9.3%. I then set the low-skill unemployment rate to twice this level and the unemployment rate of the high-skill group to $0.18/0.31 = 0.57$ times this level reflecting the unemployment rate among college graduates. These low- and high-skill groups are small, about 6% of the population, so it is reasonable to interpret them as high school dropouts and college graduates. The unemployment rate of the middle skill group is set to match the overall unemployment rate. With these skill-specific unemployment rates, I then solve the following set of equations to find $\bar{\lambda}, \hat{\zeta}^1, \hat{\zeta}^2, \hat{\zeta}^3$

$$u^{\text{long}} = (1 - \bar{\lambda}) \sum_{s=1}^3 P_s u_s$$

$$P_s u_s = \sum_{\hat{s}=1}^3 P_{\hat{s}} T^s(s|\hat{s}) \left[(1 - \bar{\lambda}) u_{\hat{s}} + \hat{\zeta}^{\hat{s}} (1 - u_{\hat{s}}) \right] \quad \forall s = 1, 2, 3.$$

In these equations u^{long} refers to the average long-term unemployment rate, which is constructed as described in the text and u_s are the skill-specific unemployment rates described

above. After solving these equations, I adopt the normalization $\bar{\zeta} = \hat{\zeta}^3$ and $\zeta^s = \hat{\zeta}^s - \bar{\zeta}$. This leads to $\bar{\lambda} = 0.679$, $\bar{\zeta} = 0.037$, $\zeta^1 = 0.119$, $\zeta^2 = 0.029$, and $\zeta^3 = 0$.

B Implementation of methods

The first step in the Reiter algorithm is to discretize the decision rules and the distribution of wealth. There are six discrete idiosyncratic income states corresponding to the three skill-levels and two employment statuses. For each of these, I represent the savings policy rule using a linear spline with 99 knots, which are concentrated at low asset levels where there is more curvature in the savings policy rule. Following Reiter, I parameterize the level of assets at which the borrowing constraint binds and only approximate the decision rule above that point. Therefore there are a total of 100 parameters for this approximation. To approximate the distribution of wealth for a given idiosyncratic income state, I use a histogram with 500 bins. These too are unevenly placed and more concentrated at low asset levels.

The equations used to solve the model are as follows: first, the Euler equation must hold with equality at the 600 points at which the savings policy rule is approximated. Second, the evolution of the distribution of wealth is described by 3000 equations that specify how the mass of households in one bin updates. These equations are based on Reiter (2010) and depend on the savings rule and distribution of wealth in the previous period. If a household wishes to choose a level of savings between two grid points, their assets are attributed between the nearest two grid points so as to preserve the aggregate assets. Third, I use equations (2), (3), (5), and (6), which describe the exogenous aggregate processes. Fourth, equation (8) describes the evolution of the unemployment rates for each of the three skill levels. Finally, there are the observables in equations (10)-(13). Equations (12) and (13) are used in logs and I introduce a lag of $\log C$ and $\log Y$ so that I can calculate $\Delta \log Y$ and $\Delta \log C$ when I write the model in state space form. The integral in equation (13) is calculated using the discrete approximation to the distribution of wealth. These equations describe the evolution of A_t , \tilde{w}_t , λ_t , ζ_t , $\{u_{s,t} : \forall s \in S\}$, u_t^{short} , u_t^{long} , Y_t , C_t , as well as the savings policy rules and the distribution of wealth.

The model equations are differentiated by automatic differentiation as described in Reiter (2010) and then solved using the algorithm of Sims (2002). To reduce the size of the resulting linear system, I compute the Hankel singular values of the system, which reveal the contribution of each state to the input-output behavior of the overall system. Let n be the number of Hankel singular values greater than 10^{-8} . I then reduce the model to order n using a balanced truncation. Calculating Hankel singular values and balanced truncation are easily performed using the functions ‘`hsvd`’ and ‘`balred`’ from the Matlab Control Systems Toolbox.

C Accuracy of the solution method

To assess the accuracy of the model solution method, I compare the solution that is linear in aggregate states to a fully non-linear solution. The non-linear solution is not the true solution as all calculations of this kind contain some amount of approximation and numerical error. So the spirit of the exercise is to compare the solution generated by the methods used in the paper to a solution generated by a more commonly used method that is well understood.

I conduct this check on a simplified version of the model that has fewer aggregate state variables because the fully non-linear method faces the curse of dimensionality. In what follows, I ignore the transitory wage shock and assume that job-finding and job-separation rates are perfectly negatively correlated so in the end there are two aggregate shock processes. In addition to these changes, I need to eliminate the income-pooling insurance system because this makes the pattern of unemployment across skill levels an aggregate state variable. Instead, I assume that when a household is unemployed they simply receive 30% of what they would have if they were employed. Those who are employed do not make any contribution towards this unemployment insurance. There is no insurance against skill shocks.

I approximate the aggregate shocks using two 11-state Markov chains generated using the Rouwenhorst (1995) algorithm.¹⁴ To solve the model non-linearly, I use the endogenous

¹⁴Above, I estimated the autoregressive coefficient of λ to be 0.920 and that of ζ to be 0.924. For the common labor market shock, I set this coefficient to 0.922. For a given persistence, the Rouwenhorst

grid point algorithm. With the solution in hand, I simulate the solutions using the same sequence of shocks, which are discrete following the Rouwenhorst approximation.

I compare the fully non-linear solution to the result of the Reiter method with and without the model reduction step. Table 9 shows selected moments for the three solution methods. The results for both versions of the Reiter method are identical because the model reduction step leads to very small losses of accuracy. This is not surprising because the model reduction techniques have explicit bounds on the size of errors that can result and these errors can be reduced by including more states in the reduced model. Turning, to the non-linear model the results show some differences from the Reiter method, but these are generally fairly small relative to either the distance between the low-insurance and full-insurance economies or the distance between the model and the data shown in Table 2. Much of the discussion in section 5 centered around the standard deviation of consumption growth. Here it seems that the Reiter method slightly exaggerates this volatility predicting 0.26 as compared to 0.23 for the non-linear solution. Compare this discrepancy to results from Table 2, which shows a value for the full-insurance economy of 0.07 while the empirical moment is 0.54. Other interesting points of comparison show the first-order autocorrelation of consumption growth is higher in the non-linear solution, but again this discrepancy is fairly small relative to the distance to the data and the full-insurance economy. The same applies to the correlation of consumption growth and unemployment.

In addition to the properties of consumption growth, Table 9 also shows some moments of the processes driving income and unemployment. Jung and Kuester (2011) have emphasized the non-linearities inherent in the dynamics of unemployment. From panels A., B., and D., it would appear that these non-linearities result in little loss of accuracy in this application.

algorithm generates a transition matrix that is independent of the mean and variance of the process and generates a grid for the values of the process that shifts with the mean of the process and scales by the unconditional standard deviation of the process. Therefore, I can apply the algorithm to approximate λ_t and then calculate the grid values for ζ_t as $\zeta = \bar{\zeta} - \sigma_\zeta(\lambda - \bar{\lambda})/\sigma_\lambda$. The results would be identical if I were to approximate ζ and then calculate λ from that. It is as though I apply the Rouwenhorst algorithm separately for the two processes and then assume that the transitions across states are perfectly correlated.

<u>A. Mean relative to trend ($\times 100$)</u>				
	ΔC_t	ΔY_t	u_t^{short}	u_t^{long}
Non-linear	0.000	0.000	6.294	3.190
Reiter	0.000	0.000	6.327	3.006
Reiter-reduced	0.000	0.000	6.327	3.006

<u>B. Standard deviation ($\times 100$)</u>				
	ΔC_t	ΔY_t	u_t^{short}	u_t^{long}
Non-linear	0.231	1.237	0.875	1.352
Reiter	0.261	1.229	0.880	1.290
Reiter-reduced	0.261	1.229	0.880	1.290

<u>C. Correlation of ΔC_t with</u>			
	ΔY_t	u_t^{short}	u_t^{long}
Non-linear	0.768	-0.030	-0.019
Reiter	0.776	-0.009	0.005
Reiter-reduced	0.776	-0.009	0.005

<u>D. First-order autocorrelation</u>				
	ΔC_t	ΔY_t	u_t^{short}	u_t^{long}
Non-linear	0.149	0.065	0.894	0.970
Reiter	0.111	0.059	0.895	0.970
Reiter-reduced	0.111	0.059	0.895	0.970

Table 9: Moments from alternative solution methods. Model moments are calculated from 10,000 quarters of simulated data.