



Queen's Economics Department Working Paper No. 1333

Heterogeneity in Macroeconomics and the Minimal Econometric Interpretation for Model Comparison

Marco Cozzi
Queen's University

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

11-2014

Heterogeneity in Macroeconomics and the Minimal Econometric Interpretation for Model Comparison

Marco Cozzi, Queen's University

November 2014

Abstract

This paper formally compares the fit of various versions of the incomplete markets model with aggregate uncertainty, relying on a simple Bayesian empirical framework. The models differ in the degree of households' heterogeneity, with a focus on the role of preferences. For every specification, empirically motivated priors for the parameters are postulated to obtain the models' predictive distributions, which are interpreted as being the distributions of population moments. These are in turn contrasted with the posterior distributions of the same moments obtained from an atheoretical (Bayesian) econometric model. It is shown that aggregate data on consumption and income contain valuable information to determine which models are more likely to have generated the data. In particular, despite its generality, a model with both risk aversion and discount factor heterogeneity displays a very low marginal likelihood, and should not be employed for the design of macroeconomic policies and welfare analysis. It is also found that the other models display similar posterior odds, with the Bayes factors ranging between 1 and 3. Finally, it is shown that practitioners in the field should carefully calibrate the values of the unemployment rate in booms and expansions, as they heavily affect the autocorrelation of aggregate consumption and the correlation between consumption and income. This finding suggests that the magnitude of welfare effects computations is likely to be influenced considerably by these two parameters.

JEL Classification Codes: C63, C68, E21, E32, D52, D58.

Keywords: Heterogeneous Agents, Incomplete Markets, Unemployment Risk, Business Cycles, Bayesian Methods, Calibration, Minimal Econometric Interpretation, Model Comparison.

Acknowledgements: *I am grateful to SHHRC for providing financial support for this project (IDG Grant # 430-2013-511) and to Kevin Andrew for useful comments and help with the data. Contact details: Department of Economics, Queen's University, 94 University Avenue, Kingston, ON K7L 3N6, Canada. Tel: 1-613-533-2264, Fax: 1-613-533-6668, Email: mcozzi@econ.queensu.ca.*

1 Introduction

Quantitative structural modeling represents a widely popular way of undertaking macroeconomic analysis. Models with household-level heterogeneity, a mix of idiosyncratic and aggregate shocks, and incomplete markets, as in Krusell and Smith (1998), are especially valuable because they allow to perform counterfactual computational experiments, study the welfare implications of different policy regimes, and evaluate the distributional impacts of various policy interventions. A number of contributions assessing the welfare consequences of eliminating business cycles with this class of models have found the welfare benefits of eliminating aggregate risk to be large, approximately an order of magnitude larger than those originally documented by Lucas (1987). Notable studies in this literature are Storesletten, Telmer, and Yaron (2001), Mukoyama and Sahin (2006) and Krusell, Mukoyama, Sahin and Smith (2009). Castaneda, Diaz-Gimenez and Rios-Rull (1998), Heathcote (2005) and Chiu and Molico (2010) propose variants of the baseline framework to quantify the macroeconomic outcomes and distributional effects of cyclical variations in the income distribution, and of fiscal and monetary policy.

Although the use of sophisticated empirical methodologies has a long and well established tradition in models with representative agents, the computational challenges of solving models with incomplete markets and aggregate shocks make basic calibration procedures the most common methodology to consider the quantitative implications of a model.¹ It goes without saying that this approach contributes to shape the quantitative answers in a deep and not always fully understood way. In order to perform reliable welfare analyses and provide a sound guidance for policy design, it is desirable to confront these models with the data in a systematic way. Furthermore, given that the researcher is free to consider several different dimensions of heterogeneity, among other modeling choices, one would like to be able to assess the relative performance of different models with an internally consistent empirical framework.

In particular, this paper considers whether preference heterogeneity helps accounting for the time series behavior of aggregate consumption and its correlation with income. These are two key elements of any dynamic macroeconomic model, because they are intimately linked to the households' consumption smoothing behavior. In models with incomplete markets, the welfare consequences of a policy intervention are affected by both the insurance properties of the policy being evaluated and by the households' response in terms of their consumption/saving decisions, which can trigger interesting general equilibrium effects. Ultimately, the level of trust in the quantitative assessment of the welfare effects arising from, say, a stabilization policy fundamentally rests on a model's ability to adequately

¹See, among others, the surveys by Krusell and Smith (2006), An and Schorfheide (2007), Canova (2009), Fernandez-Villaverde (2010), and Fernandez-Villaverde, Guerron and Rubio-Ramirez (2010).

capture the dynamics of consumption at business cycle frequencies, and its correlation with income. Relying on the Krusell and Smith (1998) set-up as a starting point, more novel variants of the model are proposed, where Panel Study of Income Dynamics (PSID) data are used to provide a parsimonious specification for heterogeneity in risk aversion. The emphasis is on treating this aspect as observed heterogeneity. The different versions of the incomplete markets model with aggregate uncertainty are compared by means of the Minimal Econometric Interpretation (MEI) of the model, a Bayesian empirical framework proposed by Geweke (2010), which in turn represents a generalization of DeJong, Ingram and Whiteman (1996). The first goal is to assess the empirical performance of these models, along the dimensions they are designed to tackle, namely the behavior of both consumption and income in a time series sense. The second goal is to provide evidence on which specification of the model is more likely to have generated the data, in an empirical framework that a) does not assume that one of the competing models is the true one, b) interprets the model as a theoretical tool, providing information on population moments.² The third goal is to understand how the presence of parameter uncertainty affects the quantitative findings. To the best of my knowledge, this paper is the first one to evaluate formally models with heterogeneous agents and aggregate uncertainty.

The results show that all the models possess some empirical validity. However, using the Bayes factor as a formal measure of fit, the most general model (featuring heterogeneity in both the risk aversion and the rate of time preference) is clearly dominated by the other specifications (which shut down at least one of these channels). It is also found that the simpler models display similar posterior odds, with the related Bayes factors ranging between 1 and 3. Finally, it is shown that practitioners in the field should carefully calibrate the values of the unemployment rate in booms and expansions, as they heavily affect the autocorrelation of aggregate consumption and the correlation between consumption and income. This finding suggests that the magnitude of welfare effects computations is likely to be influenced considerably by these two parameters.

The rest of the paper is organized as follows. Section 2 briefly presents the models. Section 3 discusses the choice of prior distributions for the parameters. Section 4 outlines the empirical methodology. Section 5 describes the main results, while Section 6 concludes. A set of appendices discuss in more detail the numerical methods used and present some additional results.

²Perhaps implicitly, this is the interpretation that currently applies to most applications relying on models with heterogeneous agents and aggregate uncertainty. The endogenous variables are typically obtained from arbitrarily long simulations (after an indispensable burn-in), instead of considering the actual data series length. Since these models have heterogeneous agents, the simulations are often performed with large synthetic panel data, with a cross sectional dimension that doesn't match the size of real longitudinal datasets. This comment is all the more true when the model's solution doesn't rely on simulations, and the aggregate law of motion is solved for with projection methods, as described in den Haan (2010).

2 Preference Heterogeneity in Macroeconomic Models with Aggregate Risk

Does preference heterogeneity help account for both the correlation between aggregate consumption and income and the autocorrelation of aggregate consumption? In order to address this question, I consider four versions of the incomplete markets model with heterogeneous agents and aggregate risk, where the specific models differ in the degree of households' heterogeneity. The starting point is the framework proposed by Krusell and Smith (1998), and I am going to consider both their baseline model (denoted as $M1_{(0)}$) and the extension with preference heterogeneity in the discount factors β (denoted as $M2_{(\beta)}$). The other two models introduce another layer of heterogeneity, allowing agents to differ also in their preferences for risk. One version of the model focuses only on heterogeneity in risk aversion γ (denoted as $M3_{(\gamma)}$), while another version assumes that households' preferences differ not only in the relative risk aversion, but also in the degree of patience (denoted as $M4_{(\beta,\gamma)}$).³

[Table 1 about here]

Table 1 lists the four specifications of the models that are going to be compared. The simplest framework coincides with the standard model introduced by Krusell and Smith (1998). The only difference compared to their formulation consists of the income received by the unemployed. Following den Haan, Judd and Juillard (2010), I assume the existence of a budget-balanced Unemployment Insurance (UI) scheme, which raises contributions by taxing employed workers and distributes unemployment benefits to workers without a job.

Time is discrete. The economy is populated by a measure one of infinitely lived agents. The framework is a production economy with both idiosyncratic and aggregate risk. Aggregate productivity shocks hit the economy every period, inducing aggregate fluctuations. Idiosyncratic shocks are correlated with aggregate shocks: agents face different employment histories and self-insure by accumulating a single risky asset. An exogenous borrowing constraint (b) hampers households' ability to smooth consumption. Agents can be prevented from borrowing the desired amount of resources in periods where they obtain a low income.

Technology: The production side of the economy is modeled as a constant returns to scale technology of the Cobb-Douglas form, which relies on aggregate capital K_t and labor L_t to produce final

³In Cozzi (2014b) I show that a model with risk aversion heterogeneity and endogenous sorting into risky jobs accounts for many features of the U.S. wealth distribution. That model, however, abstracts from aggregate uncertainty.

output $Y_t = z_t K_t^\alpha L_t^{1-\alpha}$. The aggregate shock takes only two values: $z_t = \{z_G, z_B\}$, with $z_G > z_B$. The aggregate shock follows a symmetric first-order Markov chain such that booms and recessions last the same number of periods. Capital depreciates at the exogenous rate δ and firms hire capital and labor every period from competitive markets. Total labor services are $L_t = lN_t$, namely they are the product of the share of the time endowment devoted to market activities, l (normalized to 1), and the employment rate N_t . The firm's first order conditions give the expressions for the net real return to capital r_t and the wage rate w_t :

$$r_t = \alpha z_t \left(\frac{L_t}{K_t} \right)^{1-\alpha} - \delta, \quad (1)$$

$$w_t = (1 - \alpha) z_t \left(\frac{K_t}{L_t} \right)^\alpha. \quad (2)$$

Government: The government taxes the employed agents' labor income at rate τ_t to finance a budget-balanced UI scheme. Unemployed agents receive UI benefits equal to a fixed replacement rate ρ of the going labor income. Since labor supply is fixed, and the aggregate unemployment rate can only take two values (u_G when $z_t = z_G$ and u_B when $z_t = z_B$), the equilibrium tax rate is $\tau_t = \rho(1 - N_t)/N_t$, with $N_t = 1 - u_t$.

Households: Agents' preferences are assumed to be represented by a time-separable utility function $U(\cdot)$. Every household $i \in [0, 1]$ chooses consumption $(c_{i,t})$ and future asset holdings $(a_{i,t+1})$, and the households objective function is:

$$\max_{\{c_{i,t}, a_{i,t+1}\}_{t=0}^{\infty}} \mathbf{E}_0 \sum_{t=0}^{\infty} \beta_{i,t}^t \frac{c_{i,t}^{1-\gamma_{i,t}} - 1}{1 - \gamma_{i,t}}$$

where \mathbf{E}_0 is the expectation operator. In the simplest set-up, preferences are homogeneous, and all agents share the same discount factor β and the same risk aversion γ . In general, however, these parameters will differ across agents and they are indexed by i to highlight this possibility. In particular, I consider all possible combinations, namely a model without preference heterogeneity, a model with heterogeneity in the discount factors, a model with heterogeneity in the risk aversion, and a model with heterogeneity in both parameters. Furthermore, the parameters representing the agents' preferences are indexed by t to indicate that they can evolve over time.

$\beta_{i,t} \in (0, 1)$ is the agents' discount factor, and can take up to three different values, $\beta_{i,t} \in \{\beta_l, \beta_m, \beta_h\}$, with $\beta_l < \beta_m < \beta_h$. Each agent's discount factor can vary over time according to a three-state Markov chain. Similarly, the risk aversion $\gamma_{i,t} > 0$ can take up to three different values

$\gamma_{i,t} \in \{\gamma_l, \gamma_m, \gamma_h\}$, with $\gamma_l < \gamma_m < \gamma_h$. The risk aversion parameter can also vary over time according to a three-state Markov chain.

Agents can be employed, $s = e$, or unemployed, $s = u$. The employment probabilities follow a first-order Markov process, and depend on both the idiosyncratic employment status s and on the aggregate state of the economy z .

I use recursive methods to solve the model, and the value function associated with this problem is denoted with $V(a, \beta, \gamma, s, z, K)$. This represents the expected lifetime utility of an agent whose current asset holdings are equal to a , whose current discount factor is β , whose current risk aversion is γ , whose current employment status is s , facing the aggregate shock z and in an economy with K units of aggregate capital. The Bellman equation is:

$$V(a, \beta, \gamma, s, z, K) = \max_{c, a'} \left\{ \frac{c^{1-\gamma} - 1}{1-\gamma} + \beta E_{\beta', \gamma', s', z' | \beta, \gamma, s, z} V(a', \beta', \gamma', s', z', K') \right\}$$

s.t.

$$c + a' = (1+r)a + (1-\tau)wl, \text{ if } s = e$$

$$c + a' = (1+r)a + \rho wl, \text{ if } s = u$$

$$c \geq 0, \quad a' \geq b$$

The appropriate Markov chains for β, γ, s and z

$$\ln K' = \theta_{0,j} + \theta_{1,j} \ln K, \text{ if } z = z_j, j = \{G, B\} \quad (3)$$

Agents have to optimally set their consumption/savings plans. They enjoy utility from consumption and face several uncertain events in the future. Notice that, according to the algorithm used to solve this model, the relevant state variable in the agents' problem is just aggregate capital K , rather than the whole current endogenous distribution over idiosyncratic states. Hence, the agents forecast future prices relying on the (equilibrium) evolution of the aggregate capital stock, the Aggregate Law of Motion (ALM) being specified as the pair of equations (3). Moreover, every version of the model will have the laws of motion (i.e., the Markov chains) for the evolution of the exogenous stochastic state variables (β, γ, s and z) that apply to each specific case.

3 Parameterizing the Models

The issue of parameter uncertainty is ubiquitous in quantitative macroeconomics. The models' level of abstraction frequently leads to mismatches between model variables and their empirical counterparts (e.g., think of capital: housing represents a large component of households' wealth, and it is usually included in the value of the capital stock, but to what extent can we consider it as a factor of production?). Similarly, the estimates for prices and quantities of interest (and their variability) often depend on the time period chosen for the analysis (e.g., should we include the great recession data?). Parameter uncertainty, however, can be easily accommodated by specifying prior distributions for the parameters, embracing one of the building blocks of Bayesian empirical analysis.

[Table 2 about here]

Table 2 reports the list of prior distributions that are going to be used. This Table focuses on the parameters that are common across all four models. Specifically, independent uniform priors are specified, reflecting the idea that a-priori information on the model is not very informative, apart from providing bounds for the parameters.

The bounds for the capital share α are chosen on the basis of the labor share values found in the Penn World Tables 8.0. In the period 1950 – 2011 the labor share has fluctuated between 62% and 68%. Since the downward trend in the labor share seems to be very persistent, and still on-going, I use an upper bound for the capital share of 40%. Moreover, since most studies in the RBC literature use a capital share of 36%, it seems appropriate to center the prior distribution at this value.

The model considers only one type of capital, hence its depreciation rate δ is the weighted average of the depreciation rates of many different capital goods. A range for δ between 8% and 10% on an annual basis is broadly consistent with the estimates available in the literature.

The borrowing limit b is set between the most extreme case of market incompleteness, $b = 0$ such that borrowing is not allowed, and a value corresponding to (minus) the average quarterly income, $b = -1$.

As for the prior related to the exogenous labor supply l , Juster and Stafford (1991) documented that the share of time devoted to market activities, averaged between males and females, is 31.9%. Accordingly, the prior's average is set to this value, $l = 0.319$. Since the consensus is that measurement error is a pervasive issue in answers to the time allocation survey questions, I consider a range of $\pm 10\%$ around this estimate.

For comparability with Krusell and Smith (1998), I work with their parameterization for the (average) values of the unemployment rate during booms (u_G) and recessions (u_B): $u_B = 10\%$ and

$u_G = 4\%$. For both parameters, I consider a range which is ± 1 percentage point around these averages. In order to understand whether these ranges are plausible, I split the Bureau of Labor Statistics data on the quarterly unemployment rates for the 1948Q1 – 2007Q4 period into two groups, one group with the observations above the average unemployment rate of 5.59%, and the other group with the observations below it. The 3% – 5% range corresponds to the 5th – 65th percentiles of the unemployment rates below the average, while the 9% – 11% range corresponds to the top 5% of the unemployment rate distribution for the unemployment rates above the average (with 10.7% being the highest recorded unemployment rate). This statistic, hence the choice of the range for the unemployment rate in recessions, might seem extreme. However, it should be interpreted as including the marginally attached workers, an enlarged concept of the unemployment rate for which consistent measurements are available only since 1994. In the last twenty years the special quarterly unemployment rate, as it is labeled in the FRED database, has been 1.35 percentage points above the standard unemployment rate. Considering this adjustment, the 7.65% – 9.65% unemployment rate range corresponds to the 85th – 97th percentiles, which appears to be somewhat less extreme.

The priors for the remaining parameters are model dependent and are presented in Table 3.

[Table 3 about here]

The parameters whose priors still need to be discussed are the discount rate and the risk aversion.⁴ Because of current data limitations, discount rate heterogeneity must be considered as unobserved heterogeneity. Differently, if one is willing to make a structural distributional assumption, (relative) risk aversion heterogeneity can be obtained from the data, exploiting questions on attitudes towards risk currently included in a number of large datasets, such as the PSID.

In the simpler version of the model $M1_{(0)}$ these two parameters are the same for all the agents. In this case, the range for γ is [1.0, 3.0], which spans most of the available estimates for this parameter obtained with a variety of methods, surveyed for example by Attanasio and Weber (2010). Since there is not much information on the subjective discount factor, choosing a prior is more complicated. I work with a range for β equal to [0.990, 0.992], because the equilibrium annual interest rate in the corresponding complete markets economy is between 3% and 4%, and it is well known that in the benchmark Krusell and Smith (1998) model precautionary savings have a limited quantitative impact.

In the model $M2_{(\beta)}$ the range for γ is still [1.0, 3.0], while there are three time-varying discount factor types. In this case, the range for β_l is [0.9848, 0.9868], the range for β_m is [0.9884, 0.9904], and

⁴Given the fairly limited number of replications that are computationally feasible, and the non trivial way of how to impose that their probabilities add up to one, the probabilities in the various Markov chains are always kept fixed across replications.

the range for β_h is $[0.992, 0.994]$. The three ranges are chosen for the discount factors never to overlap. Furthermore, since in Cozzi (2014a) I found that this economy solved at the priors' averages matches a wealth Gini index of 0.80 and attains an average annual interest rate of 4%, I consider fairly tight priors around that calibration.

Models $M\mathcal{B}_{(\gamma)}$ and $M\mathcal{A}_{(\beta,\gamma)}$ introduce heterogeneity in the agents' preferences for risk. Unlike discount rate's heterogeneity, thanks to hypothetical lottery questions included in the PSID in 1996, this can be dealt with as observed heterogeneity. In particular, I am going to apply the same procedure proposed in Kimball, Sahn, and Shapiro (2008). With appropriate methods, I find that γ is log-normally distributed in the population, according to $LN(\mu_\gamma = 1.07, \sigma_\gamma^2 = 0.76)$.⁵ Although the true underlying distribution is a log-normal distribution, it is feasible to consider only three risk aversion types. First, there is a data limitation issue that will be discussed below. Second, this model is computationally costly, and I need to simplify the problem in order to solve it many times (I draw 2,000 combinations of parameters from their priors). In particular, I partition the PSID data in three groups with mass 20%, 30% and 50%. Then, exploiting the CDF of the true distribution, I consider the conditional averages of γ for each of these subgroups. This provides me with the three desired values for the risk aversion types $\gamma_l = 0.92$, $\gamma_m = 2.12$, and $\gamma_h = 7.55$. There are some additional facts about the risk aversion distribution that are going to guide the next modeling choices. The first fact is documented by Kimball, Sahn, and Shapiro (2008), using Health and Retirement Study (HRS) data. In the HRS the lottery questions have been asked in several years to the same respondent, leading to the conclusion that measurement error is large. Hence I consider fairly wide ranges for each parameter. The range for γ_l is $[0.62, 1.22]$, the range for γ_m is $[1.82, 2.42]$, and the range for γ_h is $[7.25, 7.85]$. Another finding is that the distribution of answers to the lottery questions is virtually stationary over time. It is therefore appropriate to assume that the evolution of the risk aversion is governed by a time invariant Markov chain. More in detail, I obtain an estimate for the Markov chain probabilities by using a moment-matching procedure. I start by imposing appropriate restrictions on the structure of the Markov Chain, so that it is parameterized by only three probabilities: $p_{\gamma_l}, p_{\gamma_m}$, and p_{γ_h} .⁶ Then I can use the observed cross sectional distribution of risk preferences to identify two of these probabilities, while the third one is backed out to replicate the intergenerational correlation between the risk aversion of parents and their children documented by Kimball, Sahn, and Shapiro

⁵For more details, see Cozzi (2014b). In that model I work with the assumption that preference types are permanent, so it is worthwhile considering the implications of time varying risk aversions.

⁶Note that other parameterizations of the CRRA Markov chain suffer from identification issues. For example, it is easy to show that a chain that imposes the additional restriction $p_{\gamma_h} = p_{\gamma_l}$ cannot be identified. Similarly, more general chains with four or more unknown probabilities would require more pieces of information from the data and cannot be recovered.

(2009). The Markov chain (Π) that I have to recover from the data is:

$$\Pi \equiv \pi(\gamma_i, \gamma'_j) = \begin{bmatrix} p_{\gamma_l} & 1 - p_{\gamma_l} & 0 \\ p_{\gamma_m} & 1 - 2p_{\gamma_m} & p_{\gamma_m} \\ 0 & 1 - p_{\gamma_h} & p_{\gamma_h} \end{bmatrix}, \text{ with } i, j \in \{l, m, h\}$$

The stationarity of the risk aversion distribution allows me to consider the system of linear equations $\mu_\gamma^* = \Pi' \mu_\gamma^*$, where $\mu_\gamma^* = [\mu_{\gamma_l}^*, \mu_{\gamma_m}^*, 1 - \mu_{\gamma_l}^* - \mu_{\gamma_m}^*]$ denotes the stationary distribution over the vector of risk aversions $[\gamma_l, \gamma_m, \gamma_h]$. Since the entries in μ_γ^* must add up to one, I can use only two fractions from the data to assign values to the Markov chain probabilities: $\widehat{\mu}_{\gamma_l}^* = 0.201$ and $\widehat{\mu}_{\gamma_m}^* = 0.303$.

Exploiting the stationarity condition, it is possible to write the following system of equations:

$$\left\{ \begin{array}{l} \mu_{\gamma_l}^* = p_{\gamma_l} \mu_{\gamma_l}^* + p_{\gamma_m} \mu_{\gamma_m}^* \\ \mu_{\gamma_m}^* = (1 - p_{\gamma_l}) \mu_{\gamma_l}^* + (1 - 2p_{\gamma_m}) \mu_{\gamma_m}^* + (1 - p_{\gamma_h}) \mu_{\gamma_h}^* \\ \mu_{\gamma_h}^* = p_{\gamma_m} \mu_{\gamma_m}^* + p_{\gamma_h} \mu_{\gamma_h}^* \end{array} \right.$$

which, for given \widehat{p}_{γ_m} , $\widehat{\mu}_{\gamma_l}^*$ and $\widehat{\mu}_{\gamma_m}^*$, leads to the following expressions for the two remaining probabilities:

$$\left\{ \begin{array}{l} \widehat{p}_{\gamma_l} = 1 - \widehat{p}_{\gamma_m} \frac{\widehat{\mu}_{\gamma_m}^*}{\widehat{\mu}_{\gamma_l}^*} \\ \widehat{p}_{\gamma_h} = 1 - \widehat{p}_{\gamma_m} \frac{\widehat{\mu}_{\gamma_m}^*}{\widehat{\mu}_{\gamma_h}^*} \end{array} \right. \quad (4)$$

The estimation procedure then relies on a simulated minimum distance estimator. The following steps are considered: a) construct a fine grid for p_{γ_m} , b) consider one grid point at a time, whose value is denoted with \widetilde{p}_{γ_m} , c) get \widetilde{p}_{γ_l} and \widetilde{p}_{γ_h} from the system (4), d) simulate the implied Markov chain for 200 periods (50 years \times 4 quarters per year) with 100,000 agents (i.e., risk aversion types), d) compute the intergenerational correlation ρ_γ^Π , e) select \widehat{p}_{γ_m} as $\widehat{p}_{\gamma_m} = \underset{\widetilde{p}_{\gamma_m}}{Arg \min} [\rho_\gamma^\Pi(\widetilde{p}_{\gamma_m}) - \rho_\gamma^{PSID}]^2$.

The procedure leads to the following estimates for the Markov chain probabilities: $p_{\gamma_h} = 0.9948$, $p_{\gamma_m} = 0.0085$ and $p_{\gamma_l} = 0.9872$.⁷

Finally, choosing a prior for the CRRA Markov chain probabilities boils down to choosing a range for p_{γ_m} . This is obtained by relying on the empirical distribution of 2,000 minimum distance estimates

⁷As shown in Figure 5 in the Appendix, given two of the three probabilities, the remaining one is uniquely identified.

obtained by simulating the chain with 5,662 artificial agents, which corresponds to the cross sectional dimension of the household heads in the PSID sample that provided an answer to the lottery questions.

A shortcoming of model $M3_{(\gamma)}$ is that, at the average priors, its wealth Gini index of 0.42 is well below the data counterpart. Hence I consider model $M4_{(\beta,\gamma)}$, where the heterogeneity in the discount factor is chosen for the model to display a plausible wealth Gini index. It was found that using the same specification as in $M2_{(\beta)}$, namely three discount factor types that evolve over time through a Markov chain, was not increasing the wealth Gini index considerably. I then opted for a specification with only two permanent discount factor types, β_l and β_h , with μ_β denoting the share of low discount factor types. In this case, the range for β_l is $[0.9785, 0.9805]$, the range for β_h is $[0.9885, 0.9905]$, and the range for μ_β is $[0.83, 0.87]$. This economy, when solved at the priors' average, achieves an average wealth Gini index of 0.81. Finally, the priors for the risk aversion types are kept the same as in model $M3_{(\gamma)}$.

4 The Empirical Framework

At this stage, the analysis does not have explicit empirical content, because the models provide quantitative assessments about population moments. Moreover, for a given parameterization, each model delivers a point for every moment of interest. By assuming priors for the parameters one obtains the predictive distributions, namely the distributions of selected moments of the endogenous variables generated by the parameter uncertainty. The framework needs another element to be able to link the models' implications to the information contained in the data, as stressed by DeJong, Ingram and Whiteman (1996). This link takes the form of a reduced form Bayesian econometric model, which allows us to relate the observables to the population moments. The interaction between these two elements of the MEI equips the researcher with an empirical framework where formal model comparison can be undertaken. Perhaps, one additional appealing feature of this approach is that it treats symmetrically the uncertainty in the moments of both the theoretical model and the data. Under the MEI, prior distributions are specified for both the theoretical and the empirical model, which induce probability distributions over functions of interest. The degree of overlap between these distributions is then used to assess the model's fit.

Since the main focus of heterogeneous-agent models is to carefully microfound the determination of consumption/saving decisions, I am going to assess their performance in two dimensions, both related to the behavior of aggregate consumption.

Data: The time period is 1948Q1 – 2007Q4, and the number of observations is 240.⁸ Following the basic principles of quantitative analysis, I make sure that model and measurement are consistent with each other. As customary in the RBC literature, I consider as aggregate consumption C_t the sum of the expenditures on Services and Non-Durables.⁹ Since all model economies are closed, and there is no government, aggregate output Y_t is defined as the sum of Services, Non Durables and Investment. Both series are HP filtered, with a standard smoothing parameter of 1,600.

A Bayesian VAR(1): In this framework there is the need to link the distributions of population moments to the observables. In order to study the empirical characteristics of the autocorrelation of consumption and its correlation with income, I specify a Bayesian VAR(1) process on detrended aggregate income Y_t and detrended aggregate consumption C_t . Vague (flat) priors are assumed for all the parameters.

$$\begin{pmatrix} Y_t \\ C_t \end{pmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{pmatrix} Y_{t-1} \\ C_{t-1} \end{pmatrix} + \begin{pmatrix} \eta_t^Y \\ \eta_t^C \end{pmatrix}$$

$$\begin{pmatrix} \eta_t^Y \\ \eta_t^C \end{pmatrix} \stackrel{iid}{\sim} N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma = \begin{bmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{bmatrix}\right)$$

With a more compact notation, the VAR can be represented as $\mathcal{Y}_t = A\mathcal{Y}_{t-1} + \mathcal{E}_t$, where \mathcal{Y}_t is a $T \times 2$ matrix, whose columns are Y_t and C_t , A is the 2×2 matrix of parameters to be estimated, and \mathcal{E}_t is the $T \times 2$ matrix of innovations. The error terms η_t^Y and η_t^C are assumed to be normally distributed, with a zero mean and a variance/covariance matrix denoted by Σ .

It is well known, see for example Koop and Korobilis (2009), that in VAR models with normally distributed shocks the flat priors assumption allows one to get analytical formulas for the posteriors' marginal distributions, which are multivariate t distributions centered around the OLS estimates. However, since the objects of interest are not the posteriors of the parameters, but functions of them (i.e., some moments), I use appropriate simulation methods to obtain draws from their posteriors. All seven parameters ($a_{11}, a_{12}, a_{21}, a_{22}, \sigma_{11}, \sigma_{12}, \sigma_{22}$) are then estimated with a posterior simulator.¹⁰

⁸The simple structure of the models' aggregate shocks makes it impossible to generate a recession as severe as the great recession. This suggests to drop the financial crisis years. Conducting the analysis with the whole sample, i.e. until 2014Q1, doesn't alter the main results.

⁹Notice that I exclude durables from the analysis. A possible alternative could be to include durables in investment. Although the empirical results are fairly similar, this procedure would be inconsistent with the model's assumption that total investment contributes to the accumulation of capital.

¹⁰For more details on how the computations are actually implemented, see the Appendix. DeJong and Ripoll (2007)

The incomplete markets models are evaluated by comparing the draws from the joint posteriors of the autocorrelation of consumption (ρ_C) and the correlation between consumption and income (ρ_{CY}) obtained from the Bayesian VAR with the draws for the same (population) moments predicted by the theoretical models.

5 Results

In this section I present the results. I begin by discussing the model comparison with a simple Bayesian approach. Then I report which parameters are the most important ones quantitatively in affecting the cyclical behavior of consumption and its correlation with income. Finally, I examine model performance along a number of dimensions that are not considered in their empirical evaluation.

5.1 Model Comparison

The model comparison is ultimately performed on the basis of the Bayes factors, and the models' marginal likelihoods have to be computed as an intermediate step. These are reported in Table 4.

[Table 4 and Figure 1 about here]

Perhaps surprisingly, the model with the most general set-up, $M_4(\beta, \gamma)$, is also the model with the lowest log-marginal likelihood, which is -11.29 . The other three models show log-marginal likelihoods that are fairly similar, ranging from -3.70 for models $M_2(\beta)$ and $M_3(\gamma)$ to -2.66 for model $M_1(0)$. The model without preference heterogeneity displays the highest log-marginal likelihood, suggesting that preference heterogeneity does not provide a quantitatively important improvement in accounting for the consumption dynamics at business cycle frequencies.

Keeping the number of dimensions along which the models are confronted with the data to two has the advantage of allowing to show a number of interesting plots. Figure 1 presents the scatter plots for the minimal econometric interpretation. Every panel plots the distributions of the two autocorrelations of interest: ρ_C and ρ_{CY} . The top and bottom rightmost panels represent the posterior distribution obtained from the econometric model, plotted twice for convenience, while the other four panels represent the predictive distributions of the four incomplete markets models.

A first comment relates to the models' evaluation. A simple visual comparison of the scatter plots reveals that all the models' predictive distributions do overlap with the moments' posterior

use a similar methodology, relying on the predictive distributions to perform model comparison. However, their models' relatively simple set-up allows them to perform a full Bayesian estimation.

distribution. The models possess some empirical validity, whose degree of success will be assessed more formally below.

Another interesting result is that the models impose a non-trivial restriction on the data. As the Figure shows, there is a systematic relationship between the autocorrelation of consumption and the correlation between consumption and income. The posterior draws obtained from the VAR model show a positive relationship between these two correlations, while the models predictive distributions display the opposite outcome for a large portion of their supports.

[Figure 2 about here]

Figure 2 makes the same point by plotting the bivariate Kernel densities of the predictive distributions for each of the four incomplete markets models. Model $M1_{(0)}$ (the top left plot), model $M2_{(\beta)}$ (the top right one), and model $M3_{(\gamma)}$ (the bottom left one) display rather similar joint predictive distributions. In contrast, it is evident that $M4_{(\beta,\gamma)}$ (portrayed in the bottom right panel) delivers the tightest predictions, with its predictive density being highly concentrated around values of ρ_{CY} close to 1 and values of ρ_C close to 0.65.

The pairwise comparisons between all models are based on the Bayes factor, which is the ratio of two models' posterior odds. The computation of the posterior odds ratios, namely the ratios of the models' marginal likelihoods, is going to be achieved with bivariate kernel density approximations.

The Bayes factors can be computed for any pair of models M_i and M_j ($i, j = 1, \dots, 4$), and they are the ratio of two marginal likelihoods, $P(M_i|data, E)$ and $P(M_j|data, E)$, whose expressions are as follows:

$$\begin{aligned}
& \frac{P(M_i|data, E)}{P(M_j|data, E)} = \frac{P(M_i|E)P(data|M_i, E)}{P(M_j|E)P(data|M_j, E)} \\
& = \frac{P(M_i|E) \int \int P(\rho_C, \rho_{CY}|M_i)P(data|\rho_C, \rho_{CY}, E)d\rho_C d\rho_{CY}}{P(M_j|E) \int \int P(\rho_C, \rho_{CY}|M_j)P(data|\rho_C, \rho_{CY}, E)d\rho_C d\rho_{CY}} \\
& \propto \frac{\int \int P(\rho_C, \rho_{CY}|M_i)P(\rho_C, \rho_{CY}|data, E^*)d\rho_C d\rho_{CY}}{\int \int P(\rho_C, \rho_{CY}|M_j)P(\rho_C, \rho_{CY}|data, E^*)d\rho_C d\rho_{CY}} \\
& \approx \frac{\frac{1}{N_{M_i}N_{E^*}} \sum_{u=1}^{N_{M_i}} \sum_{v=1}^{N_{E^*}} K\left(\rho_{C,u}^{M_i}, \rho_{CY,u}^{M_i}; \rho_{C,v}^{E^*}, \rho_{CY,v}^{E^*}\right)}{\frac{1}{N_{M_j}N_{E^*}} \sum_{u=1}^{N_{M_j}} \sum_{v=1}^{N_{E^*}} K\left(\rho_{C,u}^{M_j}, \rho_{CY,u}^{M_j}; \rho_{C,v}^{E^*}, \rho_{CY,v}^{E^*}\right)}
\end{aligned}$$

Posterior odds ratios, denoted as $P(M_i|data, E)/P(M_j|data, E)$, are a common tool in Bayesian empirical work. Notice how these depend on both the data and an incomplete econometric model denoted by E . This model is needed to provide the framework with empirical content. In particular, E specifies a conditional distribution of observables $P(Y, C|\rho_C, \rho_{CY}, \theta, E)$ together with a prior for the parameters θ in the econometric model. This econometric model is incomplete, because it does not provide a joint prior distribution for the moments ρ_C and ρ_{CY} . The prior is instead provided by each theoretical model M_i in turn. Under appropriate conditions, see Geweke (2010), the density $P(data|\rho_C, \rho_{CY}, E)$ that appears in the double integral can be replaced by the density $P(\rho_C, \rho_{CY}|data, E^*)$. The latter density is easy to work with, because it corresponds to the posterior density for the moments of interest implied by an auxiliary econometric model E^* , which in this application corresponds to the VAR model introduced above.

In the formula approximating the integrals, I rely on an independent bivariate normal kernel $K(., .)$.¹¹ The numbers N_{M_i} and N_{E^*} stand for the number of draws for the theoretical models (2,000 each) and for the Bayesian VAR (10,000).¹²

Intuitively, the marginal likelihood of each model M_i is a measure of the overlap between the predictive distributions of a heterogeneous-agent model and the posterior distributions of the VAR model. This overlap is easy to portray when focusing on one correlation at a time. In this case, the marginal distributions are plotted for both the predictive density and the posterior.

[Figures 3 and 4 about here]

Figure 3 displays the four plots for ρ_{CY} , while Figure 4 for ρ_C . These plots confirm visually that the models $M1_{(0)}$, $M2_{(\beta)}$ and $M3_{(\gamma)}$ have similar marginal likelihoods, with large overlaps between the predictive and posterior marginals. Differently, $M4_{(\beta, \gamma)}$ has a much poorer fit, with the predictive marginal densities being more concentrated for values of the correlations that are in the tails of the posterior marginal densities.

[Table 5 about here]

¹¹For all density approximations, I use Silverman’s “optimal” bandwidth. Using a fixed bandwidth equal to 0.01 doesn’t alter the results considerably.

¹²It is worth mentioning that in both models with risk aversion heterogeneity the ALM doesn’t converge and shows cycles in a non-trivial number of replications, which are discarded. This virtually never happens in the more traditional versions of the model, and in Cozzi (2014a) I have shown with a perturbation approach that self-fulfilling equilibria are not likely to arise in the incomplete markets model with aggregate shocks and discount factor heterogeneity.

To conclude with, the six Bayes factors are reported in Table 5 and they are as follows. For the comparison $M1_{(0)}$ Vs. $M2_{(\beta)}$ the Bayes factor is $2.83 : 1$, for $M1_{(0)}$ Vs. $M3_{(\gamma)}$ it is $2.85 : 1$, while for $M1_{(0)}$ Vs. $M4_{(\beta,\gamma)}$ it is a stunning $5572.56 : 1$. In the comparison $M2_{(\beta)}$ Vs. $M3_{(\gamma)}$ the Bayes factor is $1.01 : 1$, while in the $M2_{(\beta)}$ Vs. $M4_{(\beta,\gamma)}$ comparison it is $1966.93 : 1$, and for $M3_{(\gamma)}$ Vs. $M4_{(\beta,\gamma)}$ it is $1954.29 : 1$.

An unambiguous result is that all pairwise comparisons between model $M4_{(\beta,\gamma)}$ and any other model provide decisive evidence in favor of the latter. Mostly because of the time-constant heterogeneity in the discount factor, the model with the highest degree of heterogeneity is also the least reliable one when confronted with the data, and dominated by all the other alternatives. The time-varying risk aversion notwithstanding, in model $M4_{(\beta,\gamma)}$ the economy essentially consists of two subgroups of agents. The patient/risk averse households, that are the wealthy ones, and the vast majority of agents, that are relatively impatient and wealth poor. The larger fraction of the latter group is needed to match a high concentration of wealth. At the same time, this induces the noticeably high correlation between aggregate consumption and income, with consumption almost tracking income for most model parameterizations, due to the limited precautionary savings made by these agents. In this model only few household types save extensively, allowing them to have smooth consumption profiles over time and across states of the world. In the aggregate, these agents have a limited quantitative importance, with the economy also displaying a very low autocorrelation of consumption.

Finally, the pairwise comparisons among the other three models do not provide strong evidence favoring one of them.

5.2 Comparative dynamics

A quantitative framework that relies on prior distributions for the parameters allows us to address the question of which parameters matter the most in shaping the model's outcomes in a fairly straightforward way. Rather than taking the analytical derivative of one of the outcomes with respect to a parameter, one can run a regression, with either the autocorrelation of consumption or the correlation between consumption and income on the left hand side, and the corresponding combination of parameters that generated that moment on the right hand side. This represents a surrogate of comparative dynamics analysis, performed around the mean of the moments under scrutiny.

It goes without saying that there is a deterministic mapping between a model's parameters and the related outcomes. However, this mapping can be extremely complex. A way to approximate this relationship is to consider a simple linear regression. Even though the parameters are drawn randomly from their priors, this regression does not have any meaningful interpretation in a statistical sense. However, it can provide useful information on which parameters matter the most in affecting the

results. Table 6 reports the OLS regression with the correlation between consumption and income as the dependent variable, while Table 7 the OLS regression with the autocorrelation of consumption as the dependent variable.¹³

[Tables 6 and 7 about here]

A first comment is that most regressions have a high R^2 , meaning that an approximation of the first order adequately accounts for the response of the endogenous variables to changes in the parameters. However, the R^2 of models $M3_{(\gamma)}$ and $M4_{(\beta, \gamma)}$ is somewhat lower.

A second comment is that in most cases the same regression coefficient has opposite signs in the two regressions. This is intuitive, as both have implications for the consumption smoothing behavior. A (say) negative effect on ρ_{CY} implies a lower correlation between consumption and income, hence stronger consumption smoothing. This is consistent with a positive effect on ρ_C , where the higher autocorrelation is induced by the relatively flat dynamics of consumption in presence of a stronger consumption smoothing. In a few cases, a regression coefficient has the same sign in the two regressions. I interpret this outcome as being due to the mis-specification of the regression, especially for ρ_{CY} as suggested by its lower R^2 . In this case, the coefficients are likely to be picking up also the effect of higher order and interaction terms that are not included in the regression.

A third comment relates to the magnitude of the regression coefficients. Since the coefficients reported in both Tables refer to standardized regressions (i.e., both the regressors -the models' parameters- and the dependent variable are divided by their standard deviations), they provide direct information on which parameters have the largest quantitative impact.

By inspecting both Tables, it is apparent that in models $M1_{(0)}$, $M2_{(\beta)}$ and $M3_{(\gamma)}$ most parameters have a small effect on the two correlations. There are only two parameters that consistently have a very large effect: the unemployment rates in good and bad times. Since the range of their priors was fairly small, this result tells us that even minor differences in the value of these parameters have a sizable impact on the model's outcomes. Although we have extremely good data on the unemployment rate and its dynamics, the simple stochastic structure assumed in these models is too parsimonious to capture in an unambiguous way how to parameterize the Markov chains. Furthermore, the model is also too abstract with respect to its characterization of the labor market, making it hard to argue that there is a dominant approach to calibrate it. Given the results, it follows that in these models an extensive robustness analysis should always be performed along this margin. Interestingly, the two unemployment rates have opposite effects on the two correlations, and the size of their coefficients

¹³All regressions include a constant term, not reported in the Tables.

is quite similar, even though it is always larger for u_B . The effect of u_B is easy to rationalize, as a higher unemployment rate in a recession is the worst outcome an agent could find himself in. This increases the precautionary savings motive, and consequently the degree of consumption smoothing, which is consistent with a negative effect on ρ_{CY} and a positive one on ρ_C . The effect of u_G has a similar interpretation. During a boom agents accumulate assets to be able to smooth consumption when a recession will eventually hit the economy. However, a larger unemployment rate during a boom reduces the share of individuals that are in a position to accumulate assets for future rainy days. This explains the positive effect on ρ_{CY} and the negative one on ρ_C .

A different set of comments applies for $M4_{(\beta,\gamma)}$. In this case, there are more parameters that have an important role in shaping the two correlations, such as the capital share α , the capital depreciation δ and the labor supply l . However, the relatively low R^2 suggests to take these results with caution.

5.3 The Models' Performance in other Dimensions

It goes without saying that the model has implications for other variables as well, and it is interesting to consider them in some detail. Table 8 reports a number of statistics related to the predictive distributions of the correlation between investment and income, the autocorrelation of investment, and the correlation between consumption and investment. The corresponding sample statistics are $\rho_{IY} = 0.975$, $\rho_I = 0.821$ and $\rho_{CI} = 0.652$, respectively.

[Table 8 about here]

Also for these moments, all variants of the incomplete markets model seem to put evident restrictions on the data.

The correlation between investment and income is very high and extremely concentrated around its mean. The correlation's range is really small, its largest range across models being [.988, .998], and there is no overlap between the sample statistic and the models' predictions. Nevertheless, this does not represent conclusive evidence that the model is inconsistent with the data, because: a) the discrepancy is quantitatively small, b) this comparison is far from being ideal, as argued by Eichenbaum (1991), DeJong, Ingram and Whiteman (1996) and Geweke (2010).

A more troublesome aspect is related to the models' implications for the time series behavior of investment. Each variant of the model fails in getting close to the empirical autocorrelation. In all models the largest autocorrelation of investment is always below .60, while the data counterpart is remarkably higher. A way to address this shortcoming that has been used in several DSGE models would be to consider a more general structure for the shocks. Adding investment or depreciation

shocks could help the models getting closer to the data. However, this avenue would open the issue of the sources of identification for those additional shocks.

Finally, for the correlation between investment and income all models seem to perform well. Also in this case, model $M4_{(\beta,\gamma)}$ has a mean (.88) which is noticeably different from the other models (.60 for $M1_{(0)}$, .69 for $M2_{(\beta)}$, and .70 for $M3_{(\gamma)}$), and far from the sample statistic. Differently, for the other models the sample statistic is well within the support of their predictive distributions.

5.4 Calibration Vs. Estimation

The debate on whether DSGE models should be calibrated, and compared to the data with informal procedures, or estimated, and tested with formal statistical methods, is a long-lived one. Calibrators argue that DSGE models are too simple and abstract hence necessarily false, while econometricians respond that without a coherent statistical apparatus it is impossible to assess whether a model is suitable for policy analysis.

Analysing a workhorse DSGE model, Rios-Rull, Schorfheide, Fuentes-Albero, Kryshko and Santaaulalia-Llopis (2012) stressed that the findings in quantitative work in macroeconomics are mostly affected by the identification of key parameters, rather than the quantitative methodology chosen to select parameter values per se (i.e., calibration Vs. Bayesian estimation).

Currently, for models with heterogeneous agents and aggregate shocks, this antagonism is vacuous. Given their computational requirements, the Bayesian estimation of these models is still unlikely to be feasible, even with supercomputers. In order to be confident that the sequence of draws in the Monte Carlo-Markov Chain procedure accurately approximates the posterior distributions, hundred of thousands of parameter draws (and model solutions) are necessary. This, together with the demanding step of computing the likelihood function through some non-linear filters, make the Bayesian estimation of these models a very challenging endeavour. Given the computational requirements of the models considered in this paper, my results suggest that even a parsimonious model would likely take several months worth of computer time to be estimated.¹⁴

On the basis of these considerations, the MEI framework achieves a reasonable and viable compromise in this dispute. It is feasible, even for fairly sophisticated models, it allows for a formal comparison among models that can all be false, and it embeds a global sensitivity analysis by default, as explained in Canova (1994). However, the main drawback is that, by avoiding the computation of the likelihood function, posterior estimates for the model's parameters are not part of the empirical framework, giving more importance to the prior distributions than what is done in Bayesian

¹⁴Parallelized over 8 cores on a top-of-the-line workstation, the sequence of 2,000 parameterizations of each model takes more than two weeks to complete.

estimation.

Another related issue deserves some further discussion. Currently, in the literature on heterogeneous-agent models without aggregate shocks there is a widespread belief that researchers are using a GMM procedure to obtain parameters estimates. Although there are several careful applications, the reality of the matter is that in a number of instances this is not the case.

The source of the confusion, I believe, arises from what can be considered as a moment in this class of models. These models lack aggregate uncertainty, hence the economy is deterministic at the aggregate level. As a consequence, estimation can be performed only on variables that have cross-sectional variation, because only in this case is a moment well defined. Perhaps a simple example will clarify the problem. A popular and, without any doubt, desirable property is for the model to match a plausible interest rate. However, taken from the perspective of a stationary model, this is not a moment condition. The interest rate is constant over time, and for most models it does not vary in the cross sectional dimension, because there is only one asset and a common interest rate that applies to all economic agents. Differently, the counterpart in the data is typically chosen to be the rate of return between safe and risky assets averaged over time. The model delivers a number, while the data consist of two time series. Clearly, there is a mismatch between the model and the empirical dimension it is called to replicate.

Choosing parameters to match a number of empirical targets is not necessarily an application of a method of moments estimator, hence it is not necessarily an estimation exercise. In some applications this moment-matching methodology is not sound from a statistical perspective, because standard errors cannot be computed. Moreover, often the model is implicitly used to quantify population moments. Regardless, the model's moments are directly compared to their sample counterparts, and also this step is not free of potential pitfalls.

To conclude with, the models' predictive distributions show that even for the model with the lowest marginal likelihood there are calibrations that are consistent with the empirical correlations found in the data. Searching in the parameter space to match these moments would lead to a successful calibration exercise. However, this is where both prior information about the model's parameters and a well-posed criterion to confront these models with the data play a crucial role. With a very popular and quite sophisticated moment/target matching procedure, a researcher would be able to hit the empirical targets of interest. Abstracting from accumulated (prior) knowledge on some key parameters shows that one could end up working with a fairly general model, which is shown to be very fragile when evaluated empirically, such as the model with both risk aversion and discount factor heterogeneity. These considerations provide an insight on the advantages of considering Bayesian elements into the analysis.

6 Conclusions

DSGE models are tools that sharpen our understanding of several macroeconomic phenomena. In order to be used with confidence to inform policy design, the models have to be confronted with the data in a systematic way. Consumption smoothing behavior is at the core of any dynamic macroeconomic model, and models with agent heterogeneity in household wealth, employment status and other characteristics have been developed to allow for a better representation of this phenomenon. Although several solution methods for models with heterogeneous agents and aggregate shocks have been proposed in the literature, surveyed by den Haan, Judd and Juillard (2010), the bulk of macroeconomic analysis is still performed with models that abstract from heterogeneity, relying on the representative consumer paradigm. This assumption narrows the scope of the welfare analysis that can be carried out, limits our understanding of the distributional consequences of a policy reform, and neglects a number of trends that have been documented for several countries, such as the spectacular increase in residual wage inequality discussed in Krueger, Perri, Pistaferri and Violante (2010) and the other contributions included in that volume.

This paper showed that it is finally possible to undertake formal model evaluation also in models with heterogeneous agents and aggregate uncertainty, relying on a simple Bayesian framework developed by DeJong, Ingram and Whiteman (1996) and generalized by Geweke (2010).

I have shown how to incorporate observable heterogeneity in risk preferences, even though this additional layer of heterogeneity ultimately does not improve the model's fit. I have applied the MEI showing that macroeconomic data can help discriminate between different specifications of the model. In particular, when evaluating the model's fit with respect to the correlation between consumption and income, and the autocorrelation of consumption, it was found that the model with the lowest marginal likelihood is also the model with the most general specification. The Bayes factors for the other variants of the model are very similar, preventing us from favoring one specification over the other two.

Finally, I showed which parameters have the largest effects on the outcomes of interest. It was found that the unemployment rates in booms and recessions have by far the most important quantitative impact. Typically, robustness checks along this dimension are not performed, which can potentially lead to an unsound welfare analysis.

Given the infeasibility of fully Bayesian methods, currently the MEI represents a viable and coherent approach for model comparison for fairly sophisticated versions of the incomplete markets model with aggregate risk.

<i>Label</i>	<i>Model Set-up</i>
$M1_{(0)}$	<i>No Preference Heterogeneity</i>
$M2_{(\beta)}$	<i>β Heterogeneity (No γ Heterogeneity)</i>
$M3_{(\gamma)}$	<i>γ Heterogeneity (No β Heterogeneity)</i>
$M4_{(\beta,\gamma)}$	<i>β and γ Heterogeneity</i>

Table 1: Models Description.

<i>Parameter</i>	<i>Description</i>	<i>Min</i>	<i>Max</i>
α	<i>Capital share</i>	$\underline{\alpha} = 0.32$	$\bar{\alpha} = 0.40$
δ	<i>Capital depreciation rate</i>	$\underline{\delta} = 0.020$	$\bar{\delta} = 0.025$
b	<i>Borrowing limit</i>	$\underline{b} = -1.0$	$\bar{b} = 0$
l	<i>Hours worked (share of time endowment)</i>	$\underline{l} = 0.287$	$\bar{l} = 0.351$
u_G	<i>Unemployment rate in expansions</i>	$\underline{u}_G = 0.03$	$\bar{u}_G = 0.05$
u_B	<i>Unemployment rate in recessions</i>	$\underline{u}_B = 0.09$	$\bar{u}_B = 0.11$

Table 2: Model parameters and their prior's support. The parameters in the table are common to the four models and the priors are all uniformly distributed. The model period is a quarter.

<i>Model</i>	<i>Parameter</i>	<i>Description</i>	<i>Min</i>	<i>Max</i>
$M1_{(0)}$	β	<i>Discount Factor</i>	$\underline{\beta} = 0.990$	$\overline{\beta} = 0.992$
	γ	<i>Relative Risk Aversion</i>	$\underline{\gamma} = 1.0$	$\overline{\gamma} = 3.0$
$M2_{(\beta)}$	β_h	<i>Discount Factor (type-specific)</i>	$\underline{\beta}_h = 0.992$	$\overline{\beta}_h = 0.994$
	β_m	"	$\underline{\beta}_m = 0.9884$	$\overline{\beta}_m = 0.9904$
	β_l	"	$\underline{\beta}_l = 0.9848$	$\overline{\beta}_l = 0.9868$
	γ	<i>Relative Risk Aversion</i>	$\underline{\gamma} = 1.0$	$\overline{\gamma} = 3.0$
$M3_{(\gamma)}$	β	<i>Discount Factor</i>	$\underline{\beta} = 0.983$	$\overline{\beta} = 0.993$
	γ_h	<i>Relative Risk Aversion (type-specific)</i>	$\underline{\gamma}_h = 7.247$	$\overline{\gamma}_h = 7.847$
	γ_m	"	$\underline{\gamma}_m = 1.818$	$\overline{\gamma}_m = 2.418$
	γ_l	"	$\underline{\gamma}_l = 0.623$	$\overline{\gamma}_l = 1.223$
	p_{γ_j}	<i>CRRA Markov Chain Probabilities</i>	<i>See text</i>	
$M4_{(\beta,\gamma)}$	β_h	<i>Discount Factor (type-specific)</i>	$\underline{\beta}_h = 0.9885$	$\overline{\beta}_h = 0.9905$
	β_l	"	$\underline{\beta}_l = 0.9785$	$\overline{\beta}_l = 0.9805$
	μ_β	<i>Share with low discount factor β_l</i>	$\underline{\mu}_\beta = 0.83$	$\overline{\mu}_\beta = 0.87$
	γ_h	<i>Relative Risk Aversion (type-specific)</i>	$\underline{\gamma}_h = 7.247$	$\overline{\gamma}_h = 7.847$
	γ_m	"	$\underline{\gamma}_m = 1.818$	$\overline{\gamma}_m = 2.418$
	γ_l	"	$\underline{\gamma}_l = 0.623$	$\overline{\gamma}_l = 1.223$
	p_{γ_j}	<i>CRRA Markov Chain Probabilities</i>	<i>See text</i>	

Table 3: Model specific parameters and their prior's support. The prior for the CRRA Markov Chain probability is the empirical distribution of 2000 minimum distance estimates obtained by simulating the chain with 5662 artificial agents (which corresponds to the cross sectional dimension of the household heads in the PSID sample). All other priors are uniformly distributed.

<i>Model</i>	<i>Log Marginal Likelihood</i>
$M1_{(0)}$ - <i>No Preference Heterogeneity</i>	-2.661
$M2_{(\beta)}$ - β <i>Heterogeneity</i>	-3.702
$M3_{(\gamma)}$ - γ <i>Heterogeneity</i>	-3.709
$M4_{(\beta,\gamma)}$ - β and γ <i>Heterogeneity</i>	-11.287

Table 4: Log Marginal Likelihoods under the Minimal Econometric Interpretation. Silverman's "optimal" bandwidth is used for all density approximations.

<i>Comparison</i>	<i>Bayes Factor</i>
$P(M1_{(0)} \cdot)/P(M2_{(\beta)} \cdot)$	2.833
$P(M1_{(0)} \cdot)/P(M3_{(\gamma)} \cdot)$	2.851
$P(M1_{(0)} \cdot)/P(M4_{(\beta,\gamma)} \cdot)$	5572.561
$P(M2_{(\beta)} \cdot)/P(M3_{(\gamma)} \cdot)$	1.006
$P(M2_{(\beta)} \cdot)/P(M4_{(\beta,\gamma)} \cdot)$	1966.930
$P(M3_{(\gamma)} \cdot)/P(M4_{(\beta,\gamma)} \cdot)$	1954.286

Table 5: Model Comparison: Posterior Odds Ratios.

$Corr(C_t, Y_t)$	$M1_{(0)}$	$M2_{(\beta)}$	$M3_{(\gamma)}$	$M4_{(\beta, \gamma)}$
α	-.055	-.089	.144	.632
δ	.025	.026	-.005	-.218
b	.003	.014	-.013	-.049
l	-.009	-.002	.021	.201
u_B	-.678	-.664	-.602	-.378
u_G	.634	.624	.561	.316
β	-.036	—	-.090	—
β_h	—	.079	—	.044
β_m	—	-.064	—	—
β_l	—	-.017	—	-.087
μ_β	—	—	—	.031
γ	.076	-.076	—	—
γ_h	—	—	-.004	.015
γ_m	—	—	-.009	-.009
γ_l	—	—	.013	-.007
p_{γ_m}	—	—	.032	—
R^2	.88	.86	.72	.67

Table 6: Comparative Dynamics - Linear Regressions, Standardized Beta Coefficients

$Corr(C_t, C_{t-1})$	$M1_{(0)}$	$M2_{(\beta)}$	$M3_{(\gamma)}$	$M4_{(\beta, \gamma)}$
α	.008	.032	-.212	-.678
δ	.019	.012	.055	.263
b	.001	-.010	.020	.054
l	-.002	-.002	-.023	-.196
u_B	.707	.712	.630	.421
u_G	-.650	-.666	-.588	-.351
β	-.007	-	-.020	-.064
β_h	-	-.054	-	-.015
β_m	-	.035	-	-
β_l	-	.011	-	.064
μ_β	-	-	-	-.031
γ	-.185	-.092	-	-
γ_h	-	-	.006	-.012
γ_m	-	-	.005	.006
γ_l	-	-	-.005	.007
p_{γ_m}	-	-	-.024	-
R^2	.96	.97	.79	.78

Table 7: Comparative Dynamics - Linear Regressions, Standardized Beta Coefficients

	<i>Min</i>	<i>Max</i>	<i>Mean</i>	<i>Med</i>	<i>S.d.</i>
<i>M1</i> ₍₀₎					
<i>Corr</i> (C_t, Y_t)	-.229	.951	.658	.727	.235
<i>Corr</i> (C_t, C_{t-1})	.654	.933	.803	.799	.074
<i>Corr</i> (I_t, Y_t)	.991	.998	.997	.997	.001
<i>Corr</i> (I_t, I_{t-1})	.595	.598	.597	.598	.001
<i>Corr</i> (C_t, I_t)	-.302	.932	.599	.668	.250
<i>M2</i> _(β)					
<i>Corr</i> (C_t, Y_t)	-.003	.955	.757	.796	.151
<i>Corr</i> (C_t, C_{t-1})	.656	.924	.784	.777	.062
<i>Corr</i> (I_t, Y_t)	.988	.998	.995	.996	.002
<i>Corr</i> (I_t, I_{t-1})	.582	.597	.594	.595	.002
<i>Corr</i> (C_t, I_t)	-.085	.933	.693	.735	.165
<i>M3</i> _(γ)					
<i>Corr</i> (C_t, Y_t)	-.227	.973	.757	.805	.175
<i>Corr</i> (C_t, C_{t-1})	.638	.924	.769	.761	.070
<i>Corr</i> (I_t, Y_t)	.994	.998	.996	.996	.001
<i>Corr</i> (I_t, I_{t-1})	.592	.598	.597	.597	.001
<i>Corr</i> (C_t, I_t)	-.297	.953	.704	.752	.189
<i>M4</i> _(β, γ)					
<i>Corr</i> (C_t, Y_t)	.506	.992	.929	.947	.058
<i>Corr</i> (C_t, C_{t-1})	.629	.895	.689	.678	.040
<i>Corr</i> (I_t, Y_t)	.991	.996	.994	.993	.001
<i>Corr</i> (I_t, I_{t-1})	.577	.596	.592	.593	.003
<i>Corr</i> (C_t, I_t)	.413	.971	.885	.905	.069

Table 8: Statistics of the Predictive Distributions for the four Models.

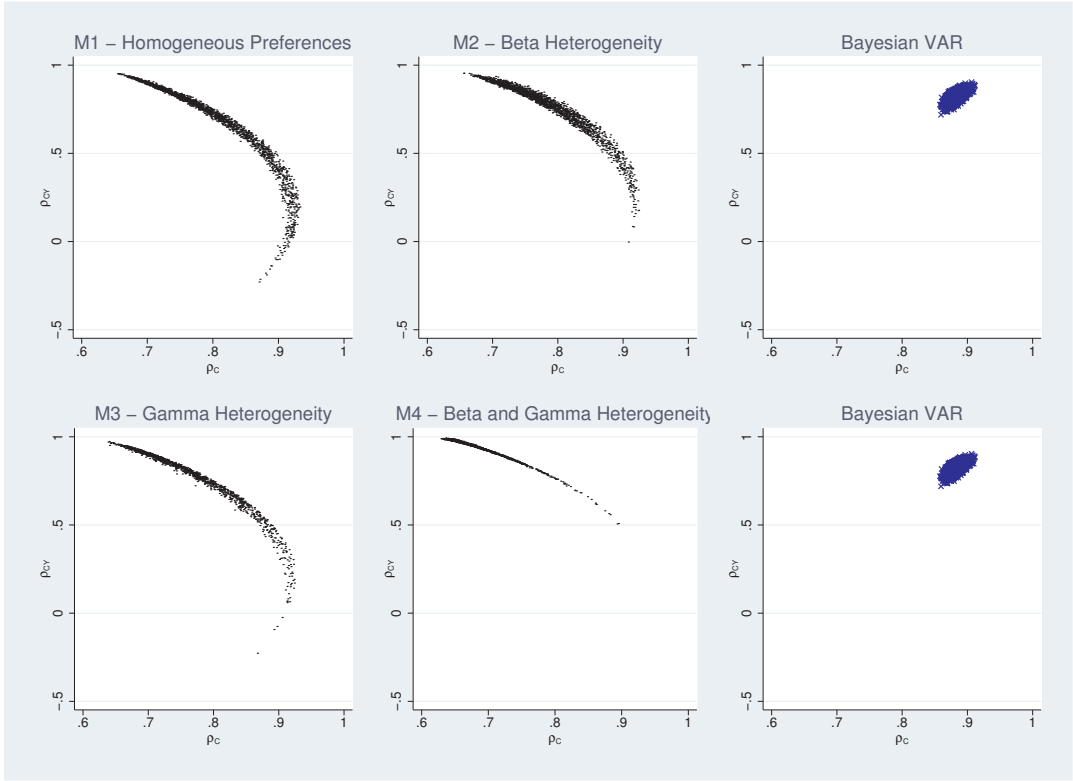


Figure 1: The Minimal Econometric Interpretation. The top and bottom rightmost panels represent the posterior distribution in the econometric model (blue crosses). The other four panels represent the predictive distributions of the four incomplete markets models (black dots).

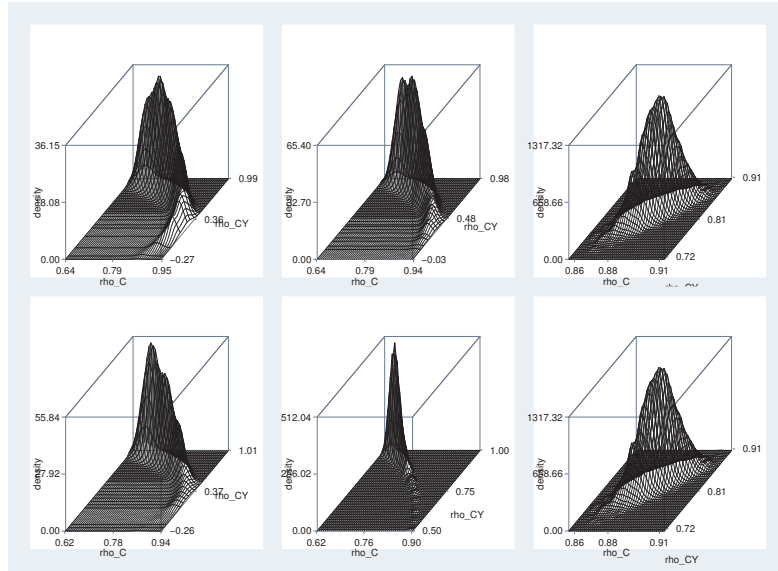


Figure 2: Bivariate Kernel Density estimates of the predictive distributions for each of the four incomplete markets models. The top left plot refers to model $M1$, the top center one to $M2$, the top and bottom right ones to the Bayesian VAR, the bottom left one to $M3$ and the center right one to $M4$.




Figure 3: Models' fit. Predictive distributions (dashed line) Vs. Posterior marginal (solid line) for the correlation between aggregate income and aggregate consumption (ρ_{CY}). Both the Predictive distributions and the posterior marginal are approximated with a Gaussian kernel.

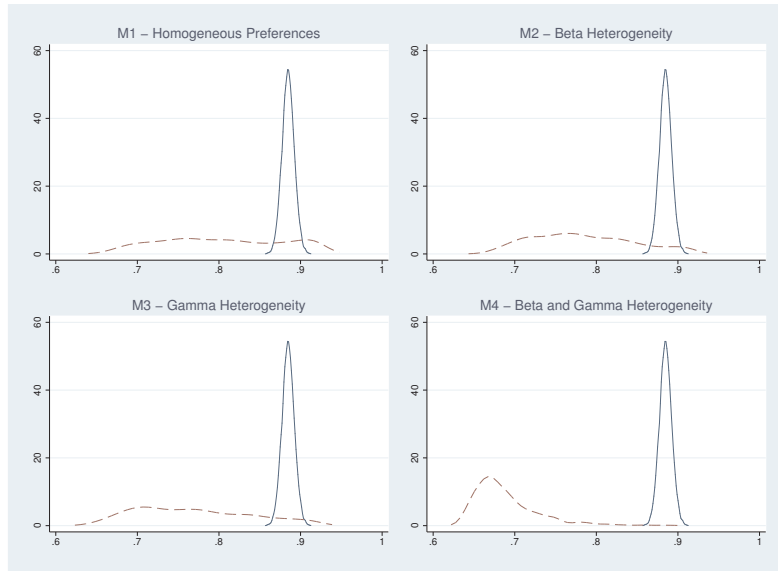


Figure 4: Models' fit. Predictive distributions (dashed lines) Vs. Posterior Marginal (solid line) for the autocorrelation of consumption (ρ_C). Both the Predictive distributions and the posterior marginal are approximated with a Gaussian kernel.

References

- An, S., and Schorfheide, F. (2007). "Bayesian Analysis of DSGE Models," *Econometric Reviews*, Vol. 26, 113-172.
- Attanasio, O., and Weber, G. (2010). "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy," *Journal of Economic Literature* Vol. 48, 693-751
- Canova, F. (2009). "How Much Structure in Empirical Models?" in *Palgrave Handbook of Applied Econometrics*, edited by T. Mills and K. Patterson, Vol. 2, 68-97.
- Canova, F. (1994). "Statistical Inference in Calibrated Models," *Journal of Applied Econometrics*, Vol. 9, S123-S144.
- Castaneda, A., Diaz-Gimenez, J., and Rios-Rull, J.V. (1998). "Exploring the income distribution business cycle dynamics," *Journal of Monetary Economics*, Vol. 42, 93-130.
- Cozzi, M. (2014a). "The Krusell-Smith Algorithm: Are Self-fulfilling Equilibria Likely?" Forthcoming in *Computational Economics*.
- Cozzi, M. (2014b). "Risk Aversion Heterogeneity, Risky Jobs and Wealth Inequality," QED Working Paper n. 1286.
- Chiu, J., and Molico, M. (2010). "Liquidity, Redistribution, and the Welfare Cost of Inflation," *Journal of Monetary Economics*, Vol. 57, 428-38.
- DeJong, D., Ingram, B., and Whiteman, C. (1996). "A Bayesian Approach to Calibration," *Journal of Business and Economic Statistics*, Vol. 14, 1-9.
- DeJong, D., and Ripoll, M. (2007). "Do self-control preferences help explain the puzzling behavior of asset prices?" *Journal of Monetary Economics*, Vol. 54, 1035-1050.
- Eichenbaum, M. (1991). "Real business-cycle theory. Wisdom or whimsy?" *Journal of Economic Dynamics and Control*, Vol. 15, 607-626.
- Fernandez-Villaverde, J., Guerron, P., and Rubio-Ramirez, J. (2010). "The New Macroeconometrics: a Bayesian Approach," *Handbook of Applied Bayesian Analysis*, edited by O'Hagan and West, Oxford University Press.
- Fernandez-Villaverde, J. (2010). "The econometrics of DSGE models," in *SERIEs, Journal of the Spanish Economic Association*, Vol. 1, 3-49.

- Geweke, J. (2010). *Complete and Incomplete Econometric Models*. Princeton University Press, Princeton (NJ).
- den Haan, W. (2010). "Comparison of solutions to the incomplete markets model with aggregate uncertainty," *Journal of Economic Dynamics & Control*, Vol. 34, 4-27.
- den Haan, W., Judd, K., and Juillard, M. (2010). "Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty," *Journal of Economic Dynamics & Control*, Vol. 34, 1-3.
- Heathcote, J. (2005). "Fiscal Policy with Heterogeneous Agents and Incomplete Markets." *Review of Economic Studies*, Vol. 72, 161-188.
- Juster, T., and Stafford, F. (1991). "The Allocation of Time: Empirical Findings, Behavioral Models, and Problems of Measurement," *Journal of Economic Literature*, Vol. 29, 471-522.
- Kimball, M., Sahm, C., and Shapiro, M. (2009). "Risk Preferences in the PSID: Individual Imputations and Family Covariation," *American Economic Review: Papers & Proceedings*, Vol. 99, 363-368.
- Kimball, M., Sahm, C., and Shapiro, M. (2008). "Imputing Risk Tolerance From Survey Responses," *Journal of the American Statistical Association*, Vol. 103, 1028-1038.
- Koop, G., and Korobilis, D. (2009). "Bayesian Multivariate Time Series Methods for Empirical Macroeconomics," *Foundations and Trends in Econometrics*, Vol. 3, 267-358.
- Krueger, D., Perri, F., Pistaferri, L., and Violante, G. (2010). "Cross Sectional Facts for Macroeconomists," *Review of Economic Dynamics*, Vol. 13, 1-14.
- Krusell, P., Mukoyama, T., Sahin, A., and Smith, A. (2009). "Revisiting the welfare effects of eliminating business cycles," *Review of Economic Dynamics*, Vol. 12, 393-404.
- Krusell, P., and Smith, A. (2006). "Quantitative Macroeconomic Models with Heterogeneous Agents," in: Blundell, R., Newey, W., and Persson, T. (Eds), *Advances in Economics and Econometrics: Theory and Applications*. Econometric Society Monographs 41, Cambridge University Press, 298-340.
- Krusell, P., and Smith, A. (1998). "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, Vol. 106, 867-896.
- Lucas, R. (1987). *Models of Business Cycles*. Basil Blackwell, Oxford.
- Mukoyama, T., and Sahin, A. (2006). "Costs of business cycles for unskilled workers," *Journal of Monetary Economics*, Vol. 53, 2179-2193.

Rios-Rull, J.V., Schorfheide, F., Fuentes-Albero, C., Kryshko, M., and Santaaulalia-Llopis, R. (2012). “Methods versus substance: Measuring the effects of technology shocks,” *Journal of Monetary Economics*, Vol. 59, 826-846.

Storesletten, K., Telmer, C., and Yaron, A. (2001). “The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk,” *European Economic Review*, Vol. 45, 1311-1339.

Appendix A - Computation

- All codes solving the economies were written in the FORTRAN 95 language, relying on the Intel Fortran Compiler, build 11.1.048 (with the IMSL library). They were compiled selecting the O3 option (maximize speed), and without automatic parallelization. They were run on a 64-bit PC platform, running Windows 7 Professional Edition, with an Intel Xeon *E5-2687Wv2* Octo-core processor clocked at 3.4 Ghz.
- The 2,000 replications are run in parallel across cores and take more than 15 days to complete. Notice that typically from 8 to 15 iterations on the ALM are needed to find each equilibrium.
- In the actual solution of the model I need to discretize the continuous state variables a and K (the discount factor heterogeneity β , the risk aversion heterogeneity γ , the employment status s , and the aggregate productivity shock z are already discrete). For the household assets a I rely on an unevenly spaced grid, with the distance between two consecutive points increasing geometrically. This is done to allow for a high precision of the policy rules at low values of a , where the change in curvature is more pronounced. In order to keep the computational burden manageable, I use 101 grid points on the household assets space (75 grid points for model $M4$), the lowest value being the borrowing constraint b and the highest one being a value a_{\max} high enough not to be binding in the simulations ($a_{\max} = 500$). For the aggregate capital K I use a fairly dense grid. Over the $[3, 27]$ interval I use 25 points, which are far more than the typical 4-6. However, given the large number of parameterizations, in the iterative process on the ALM parameters the simulations do visit regions for aggregate capital that are very far from the support of the ergodic equilibrium distribution, causing convergence issues when using a coarse grid.
- As for the solution method for the household problem, I use a time iteration procedure on the set of Euler equations, guessing the future saving functions, and with linear interpolation in the (a, K) dimensions. This method proved to be more stable than the relatively common value function iteration scheme with cubic spline interpolation.
- A collocation method is implemented, that is I look for the policy functions such that the residuals of the Euler equations are (close to) zero at the collocation points (which correspond to the asset grid). It follows that for all possible combinations of state variables I need to solve a non-linear equation. To get the optimal policy functions, I compute the first order conditions with respect to a' and through the envelope condition I obtain a set of Euler equations, whose unknowns are the policy functions, $a'(a, \beta, \gamma, s, z, K)$. I start from a set of guesses,

$a'(a, \beta, \gamma, s, z, K)_0$, and keep on iterating until a fixed point is reached, i.e. until two successive iterations satisfy:

$$\mathop{Sup}_a |a'(a, \beta, \gamma, s, z, K)_{n+1} - a'(a, \beta, \gamma, s, z, K)_n| < 10^{-6}, \forall \beta, \forall \gamma, \forall s, \forall z, \forall K.$$

- The aggregate dynamics are computed by simulating a large panel of individuals for 5,000 periods, with the first 1,000 periods being discarded as a burn-in. The panel size is 50,000 agents for the economies with relative risk aversion heterogeneity, and 30,000 agents for the other two. As for the approximation method, I rely on a bi-linear approximation scheme for the saving functions, for values of a and K falling outside the grid.
- If the numerical procedure fails to converge in some of its steps, the related results are discarded.

Appendix B - Algorithm for the Models Solution

The computational procedure used to solve the model economies can be represented by the following algorithm:

1. Solve each version of the model at the average of the uniform prior distributions and store the equilibrium ALM parameters Θ^* .
2. Consider model M_i , $i = 1, \dots, 4$.
3. Draw 2,000 combinations of parameters from their prior distributions and store them.
4. Begin the replications and set $j = 1$.
5. Calibrate the model by reading the first vector of simulated parameters and begin the model solution.
6. Generate a discrete grid over the aggregate capital space $[K_{\min}, \dots, K_{\max}]$.
7. Generate a discrete grid over the individual asset space $[-b, \dots, a_{\max}]$.
8. Guess a vector of parameters Θ^g representing the ALM, using in the first iteration Θ^* , the converged parameters for the model solved at the average of the priors.
9. Get the saving functions $a'(a, \beta, \gamma, s, z, K)$.
10. Simulate the model under the guessed ALM, and compute an update $\Theta^{g'}$ as the parameter estimates of OLS regressions on the simulated data.
11. Repeat steps 8 – 10 until the four parameters in Θ converge.
12. Compute the HP-filtered series for consumption and income, and get ρ_{CY} and ρ_C .
13. Save the output, set $j = j + 1$ and repeat steps 5 – 12 N_{M_i} times.
14. Move to the next model.

Appendix C - Algorithm for the Bayesian VAR and the Moments' Posteriors

The computational procedure used to obtain the posterior distributions for the moments of interest can be represented by the following algorithm:

1. Get the OLS estimates \widehat{A}_{OLS} for the VAR parameters and stack the estimates in the vector $\widehat{a}_{OLS} = \text{vec}(\widehat{A}_{OLS})$.
2. Compute the residuals $\mathcal{Y}_t - \widehat{A}_{OLS}\mathcal{Y}_{t-1}$ and the associated matrix $S = \left(\mathcal{Y}_t - \widehat{A}_{OLS}\mathcal{Y}_{t-1}\right)' \left(\mathcal{Y}_t - \widehat{A}_{OLS}\mathcal{Y}_{t-1}\right)$.
3. Get the inverse S^{-1} and $V = \left(\mathcal{Y}'_{t-1}\mathcal{Y}_{t-1}\right)^{-1}$.
4. Begin the replications and set $j = 1$.
5. Since the posterior for the matrix Σ is a Wishart distribution $\Sigma_{POST}|data \sim W(S, v)$, where v denotes the degrees of freedom, draw from the inverse Wishart distribution $\Sigma_{POST}^{-1} \sim W(S^{-1}, v)$.
6. To obtain Wishart draws, get the Cholesky decomposition of S^{-1} , draw v times from a correlated bivariate normal (there are only two variables in the VAR), collect the draws in the matrix D , and set $\Sigma_{POST}^{-1} = D'D$.
7. Get the inverse Σ_{POST} and use it to draw the VAR parameters from their posterior, which is a correlated normal $\widehat{a}_{POST}|\Sigma, data \sim N(\widehat{a}_{OLS}, \Sigma \otimes V)$.
8. Set up a system of three linear equations, whose unknowns are the population second order moments of Y and C . The entries in the vector of constants are the three distinct elements of the current draw Σ_{POST} , while the entries in the matrix of coefficients are simple functions of the current draws of the VAR parameters \widehat{a}_{POST} .
9. Solve the system for the current posterior draws, compute the correlations of interest and store them.
10. Set $j = j + 1$ and repeat steps 5 – 9 N_{E^*} times.

Appendix D - Data

The time series were obtained from the Federal Reserve Bank of St. Louis FRED II data base. The data are quarterly and the range is 1948:I - 2008:I (Source: <http://research.stlouisfed.org/fred2/>).

Aggregate output is defined as the sum of Services, Non Durables and Investment, and the corresponding series are:

- Real personal consumption expenditures per capita: Nondurable goods (series ID: A796RX0Q048SBEA).
- Real personal consumption expenditures per capita: Services (series ID: A797RX0Q048SBEA).
- Real Gross Private Domestic Investment (series ID: GPDIC1).

Since the Investment series represents total investment in the economy, I divide it by the Civilian Non-institutional Population (series ID: CNP16OV).

In the empirical analysis, all time series are HP filtered, and so are the simulated series obtained from the model.

Appendix E - Additional Plots




Figure 5: Identification of the CRRA Markov chain probability p_{γ_m} .




Figure 6: The Models Macroeconomic Performance along all Dimensions.