Queen's Economics Department Working Paper No. 1270

Precautionary Savings and Wealth Inequality: a Global Sensitivity Analysis

Marco Cozzi
QueenŠs University

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

6-2011

Precautionary Savings and Wealth Inequality: a Global Sensitivity Analysis

Marco Cozzi, Queen's University

June 2011

Abstract

This paper applies Canova JAE 1994 methodology to perform a thorough sensitivity analysis for the Aiyagari QJE 1994 economy. This is a calibrated GE model with incomplete markets and uninsurable income risk, designed to quantify the size of precautionary savings and the degree of wealth inequality. The results of this global robustness analysis are broadly consistent with Aiyagari's findings. Even when considering priors for the parameters uncertainty which are highly dispersed, the size of the precautionary savings is modest: at most, they account for an 11% increase in the saving rate. However, the results show that the parameter representing the exogenous borrowing limit seems to lead to relatively large changes in measures of wealth inequality. The Gini index increases by 15 points when considering values of the borrowing limits that lead to empirically plausible shares of households with a negative net worth. The parameters that quantitatively have the largest effects on determining the wealth Gini index are the capital share, the borrowing limit, and the depreciation rate. The parameters affecting most significantly precautionary savings are the risk aversion and the standard deviation of the income shocks.

JEL Classification Codes: E21, D52, D58.

Keywords: Precautionary Savings, Calibration, Heterogeneous Agents, Incomplete Markets, Computable General Equilibrium, Monte Carlo.

Aknowledgements: *I am grateful to Giovanni Gallipoli, Makoto Nakajima and Josep Pijoan-Mas for several comments and constructive criticism. Corresponding author's e-mail: mcozzi@econ.queensu.ca. Updated versions of the paper can be found at:*

http://qed.econ.queensu.ca/pub/faculty/cozzi/Webpage/

1 Introduction

What is the size of precautionary savings? What accounts for the degree of wealth inequality observed in the data? These questions are hard to address relying solely on data regarding income, asset levels and consumption. A quantitative theory is a useful approach to understanding how uncertainty and limited insurance opportunities affect aggregate savings and the shape of the wealth distribution.

Several contributions after the seminal paper by Aiyagari (1994) have tried to measure the size of precautionary savings, typically finding that it is small and that the standard incomplete markets model cannot account for the high degree of wealth concentration found in the data.[1]

A possible source of criticism on the results obtained in this literature is the limited scope of the robustness checks that are usually carried out in quantitative macroeconomics analyses. In order to tackle these kind of objections, this paper applies Canova (1994) methodology to perform a thorough sensitivity analysis for the Aiyagari (1994) economy and makes two contributions. On the one hand it confirms Aiyagari (1994) findings, which suggest that the size of precautionary savings is small and that uninsurable income risk is not enough to explain wealth inequality. On the other hand it shows that with today's computational power, global sensitivity analysis of the type proposed by Canova (1994) in the simple Real Business Cycle framework are also feasible for richer Heterogenous-Agent (HA) economies.

This type of exercise is important for two reasons.

First, there is often no natural empirical counterpart to the model's parameters, hence some relevant calibration targets (e.g. the capital/output ratio, or the labor share) have to be chosen for the model to match. In the absence of aggregate uncertainty, the researcher is confronted with two possible choices. The first one focuses the attention on a single year, making the measurements for all variables consistent with the steady-state of the model, whereas the second choice takes long-run averages of the statistics of interest instead. Moreover, selecting the moments relevant for the analysis together with their weighting matrix can be a delicate step. Either way, sampling variability and parameter uncertainty for the parameters that are taken from related empirical studies (such as the ones characterizing the exogenous stochastic processes for labor income, or the degree of risk aversion) can cast doubts on how general the quantitative findings are.

Second, a carefully executed global robustness analysis can deliver bounds and distributions for the measurement of the economic outcomes under study, showing how likely these outcomes are. A more widespread implementation of this methodology can give the calibration approach to empirical research sounder evidence for its findings.

[1] See, for example, the surveys by Cagetti and De Nardi (2008), Heathcote, Storesletten and Violante (2009) and Quadrini and Rios-Rull (1997).

The results of our global robustness analysis are consistent with Aiyagari's findings. Even when considering priors for the parameter uncertainty which are very different from each other (Uniform Vs. Symmetric and Asymmetric Beta distributions) and highly dispersed, the size of the precautionary savings is modest and the degree of wealth concentration falls short of the empirical observations. However, the results show that the parameter representing the exogenous borrowing limit seems to lead to relatively large changes in measures of wealth inequality. The (median) wealth Gini index increases by more than 10 points when considering values of the borrowing limits that lead to empirically plausible shares of households with a negative net worth, moving from 0.45 to up to 0.61. This number, however, is still well below its empirical counterpart of 0.8, as found for the U.S. by Budria Rodriguez, Diaz-Gimenez, Quadrini, and Rios-Rull (2002).

Finally, the simulated data make it possible to study which parameters affect the outcomes most significantly. As for the wealth Gini index, the parameters that quantitatively have the largest effects are the capital share, the borrowing limit, and the depreciation rate. As for the precautionary savings, the parameters affecting them the most are the risk aversion and the standard deviation of the income shocks.

The rest of the paper is organized as follows. Section 2 presents the theoretical model. Section 3 is devoted to the description of the Monte Carlo experiments. Section 4 presents the main results, while Section 5 concludes. A series of appendices provide more details on the methodology together with some additional results.

2 The Aiyagari (1994) HA Economy

This is a GE model with incomplete markets and uninsurable income risk, designed to quantify the size of precautionary savings and the degree of wealth inequality.[2] Time is discrete. The economy is populated by a measure one of infinitely lived ex-ante identical agents.

2.1 Preferences

Agents' preferences are assumed to be represented by a time separable utility function $U(.)$. Agents' utility is defined over stochastic consumption sequences $\{c_t\}_{t=0}^{\infty}$: their aim is to choose how much to consume (c_t), and how much to save in an interest bearing asset (a_{t+1}) in each period of their lives,

[2] In the interest of space, just a sketch of the model is presented. For more details see Appendix A, Aiyagari (1994) and Rios-Rull (1999).

in order to maximize their objective function. The agents' problem can be defined as:

$$\max_{\{c_t,a_{t+1}\}_{t=0}^{\infty}} E_0 U(c_0, c_1, ...) = \max_{\{c_t,a_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t u(c_t)$$

where E_0 represents the expectation operator over the stochastic efficiency units of labor ε_t. $\beta \in (0,1)$ is the subjective discount factor. We assume that $u(c_t) = \frac{c_t^{1-\sigma}-1}{1-\sigma}$: the per period utility function is strictly increasing, strictly concave, satisfies the Inada conditions, and has a CRRA$=\sigma$.

2.2 Endowments

Agents are all born with the same asset endowment a_0. There is a stochastic process for the effective units of labor ε a worker is going to supply in the labor market. This process is assumed to be an exogenous continuous first order Markov process, specified as an AR(1): $\varepsilon_{t+1} = \rho_y \varepsilon_t + \eta_{t+1}, \eta_t \sim iid\ N(0, \sigma_y^2)$.

2.3 Production

The production side of the model is represented by a constant returns to scale technology of the Cobb-Douglas form, which relies on aggregate capital K_t and labor L_t to produce the final output Y_t. $Y_t = F(K_t, L_t) = K_t^{\alpha} L_t^{1-\alpha}$. The labor input L_t is the sum of the workers' efficiency units $L_t = \int \varepsilon_t d\mu_t(\varepsilon)$, where $\mu_t(\varepsilon)$ is the distribution over the labor endowments implied by the markov process.

2.4 Other market arrangements

All markets are competitive. Capital is supplied by rental firms that borrow from workers at the risk-free rate r and invest in physical capital, which depreciates at rate δ.

There are no state-contingent markets to insure against income risk, but workers can self-insure by saving into the risk-free asset. The agents also face a borrowing limit, denoted as $b \geq 0$.

3 Aiyagari QJE 94 meet Canova JAE 94

Can we complement standard calibration methods by undertaking more robust (i.e. global) sensitivity analysis in equilibrium HA models? The answer is yes.

This paper applies the methodology proposed by Canova (1994) in order to parametrize the model economy.

[Table 1 about here]

Table 1 provides a list of the parameters in the Aiyagari (1994) economy, together with their upper and lower values that are going to be used in the simulations. The economy under study has seven independent parameters: $\alpha, \beta, \delta, \sigma, \rho_y, \sigma_y$, and b. Canova (1994) provided a very simple procedure to implement global robustness checks for fully parametrized quantitative macroeconomic models, which has not been fully exploited by the applied general equilibrium calibration community. The essential feature of Canova's approach is to acknowledge parameter uncertainty. This uncertainty can come from uncertainty on the moments to be matched, or from the intrinsic uncertainty of the estimation studies whose results are used to pin down exogenously some parameters. Rather than computing the quantitative implications of a theory for a unique set of calibrated parameters, the methodology assumes prior distributions for them, and solves repeatedly the model for many, many different calibrations obtained by drawing each parameter vector from the prior distributions.

Notice that this type of exercise is not meant to be a substitute for more traditional calibration studies. Once a quantitative theory is proposed and a set of empirically relevant values for the parameters are available, one can provide more robust evidence on the original quantitative findings, possibly suggesting in which cases and, more importantly, for which economic channels some results are more likely to be observed.

3.1 Monte Carlo Calibration as a Global Sensitivity Analysis

The calibration procedure consists of performing a series of Monte Carlo experiments, which rely on two steps.

In the first step, the prior distributions from which the parameters are going to be drawn are postulated. In the second step 2,500 economies are simulated and solved. These economies differ only in the parameters' vector that is drawn at each iteration m. For every economy, two equilibria are computed: the equilibrium of the incomplete markets economy (IM) and the equilibrium of the corresponding complete markets one (CM). Notice that the computation of the equilibrium in the IM economy requires an iterative procedure on the interest rate until the asset market clearing is achieved, while the equilibrium in the CM economy is easily found by exploiting the condition $1 = (1 + r)\,\beta^m$. In the CM economy the equilibrium interest rate is pinned down by the discount factor, hence it will depend on the actual realization of this parameter in each Monte Carlo replication: every iteration m will deal with a different β^m, hence it will involve a different equilibrium interest rate $r^m = \frac{1}{\beta^m} - 1$, and different allocations. Substituting the equilibrium r^m into the firms' first order conditions gets the aggregate capital in the CM economy for each simulated combination of parameters (which in turn is going to be affected by the draws of α^m and δ^m, that determine the marginal productivity and the user cost of capital, and by the draws of ρ_y^m and σ_y^m, that determine the total labor input L^m).

Finally, the size of precautionary savings is quantified as the difference between the aggregate savings in the IM economy and the aggregate savings in the CM one.

These experiments are repeated several times: with different prior distributions for the parameters, and, as an additional robustness check, with different assumptions on the degree of market incompleteness (i.e. how tight the borrowing constraint is).[3]

[Table 2 about here]

Table 2 reports the list of prior distributions that are going to be used. More in detail, independent Uniform and Beta priors are specified.

The choice settled on these specific distributions for two reasons: 1) with the priors centered around the typical calibration values, it is possible to assess if there is a tendency for the model to "compress" the outcomes around a set of results; 2) by comparing the results of the Uniform and Beta calibrations, it is possible to appreciate the effect of relying on symmetric vs. asymmetric priors, which, at the same time, differ in their dispersion.

The first set of experiments, denoted by *U1-U4*, assume that all parameters are uniformly distributed. The average of such distributions corresponds to the values that are used in the typical calibrations of these models: $\alpha^{avg} = 0.35, \beta^{avg} = 0.96, \delta^{avg} = 0.075, \sigma^{avg} = 2, \rho_y^{avg} = 0.9$ and $\sigma_y^{avg} = 0.145$.[4] These values are chosen to match some macroeconomic facts and to be consistent with the available empirical evidence. The first three parameters match the capital share of income, the capital/output ratio, and the investment/output ratio, while the remaining three are borrowed from related applied studies.[5] In order to check if the findings are robust to relatively large changes in the parameters, and if they are well behaved with respect to them, the parameters' bounds are chosen to allow for a wide range of possible calibrations. At the same time, an attempt is made to prevent the

[3] In order to minimize the effect of sampling variability, for the same values of ρ_y^m and σ_y^m, all simulations share the same sequence of income shocks. The solution algorithm, some computational aspects and simulation details are explained in Appendices B and C.

[4] The length of the model period is set to one year. Appendix E reports the results for another specification, with $\underline{\rho_y} = 0.5$. The results related to the precautionary savings are very similar, while the ones related to inequality are slightly different. However, the empirical studies on income risk tend to provide evidence dismissing values for the persistence parameter below 0.8.

[5] See, among others, Attanasio, Banks, Meghir, and Weber (1999), Guvenen (2009) and the papers cited therein.

model to grossly miss some quantitative implications. This would happen in several replications, if the parameters' ranges were to be too wide.

The four uniform experiments differ in only one aspect, which is related to the borrowing constraint b. The case *U1* is going to assume for the Monte Carlo simulations that agents cannot borrow, namely an economy for which the parameter b is fixed at its most stringent value, $b = 0$. This is the case that is most often considered in the literature, including Aiyagari (1994). The other three Uniform prior cases, *U2-U4*, relax this aspect, allowing for borrowing constraints that vary from one experiment to the other (while keeping the sequence of realizations for all the remaining parameters the same in all replications). These additional cases rely on a uniform prior, whose upper bound is the no borrowing case (i.e. $\overline{b} = 0$), while the lower bound is less and less stringent, being equal to $\underline{b} = -1.0, -2.0, -2.5$. These three values were chosen because they imply agents with a negative net worth on average equal to $8\%, 11\%$ and 15%, respectively. The first and last values are consistent with the data reported in Cagetti and De Nardi (2008) and Wolff (1998).

Table 2 reports the second set of prior distributions as well. In these experiments, denoted by *B1-B7*, more flexibility is allowed for. It is assumed that all the parameters are distributed according to a Beta distribution. This distribution has four parameters: two shape parameters together with two location parameters. The location parameters allow to define the support of the distribution in an interval different from the $[0, 1]$ one, and are kept constant across all experiments. For comparability with the Uniform cases, these coincide with the upper and lower bounds reported in the last two columns of Table 1. Notice that we kept exactly the same bounds for all parameters as in the uniform case that was matching the average share of households in debt across simulations found by Wolff (1998), that is case *U4*, with $\underline{b} = -2.5$. As for the other two (shape) parameters that are needed to fully characterize the Beta distribution, two different set-ups are considered. Four cases, *B1-B4*, rely on the priors $Beta(2, 2), Beta(10, 10), Beta(5, 2), Beta(2, 5)$ for all the parameters. The first two specifications deal with a Beta distribution which is symmetric around its mean, because the two shape parameters are equal. The pair $(10, 10)$ implies a distribution which is substantially more concentrated around its mean, when compared to the $(2, 2)$ one. The other two cases deal with asymmetric priors. The $Beta(5, 2)$ is a left skewed distribution, while the $Beta(2, 5)$ is a right skewed one. Figure (1) shows the shape of the four Beta priors over the $[0, 1]$ interval.

[Figure (1) about here]

The last three cases, *B5-B7*, introduce an empirically motivated mix of Beta priors. All parameters, but ρ_y and σ_y, rely on the $Beta(10, 10)$ prior, i.e. the most concentrated case. As it will be discussed below, this assumption was found to minimize the number of implausible quantitative implications of the model. However, for the two parameters describing the earnings process, we rely on eitehr a $Beta(5,2)$ or a $Beta(10,10)$ prior. The reason is simple: the bulk of the available econometric evidence, discussed for example in Guvenen (2009), tends to find highly persistent processes, with many studies estimating $\rho_y > 0.95$. However, depending on the specification of the transitory component of the income shocks and whether income profile heterogeneity is included as well in the econometric model, this estimate can be as low as 0.8. These results suggest what is a reasonable range for the earnings persistence parameter, and that one should try to analize several priors that give different importance to some subset of the parameter's support. Similar comments apply to the standard deviation of the innovation σ_y.

4 Results

This Section presents the main results. First we show how allowing for very flexible calibrations does not alter considerably typical measures of precautionary savings. Second we discuss how measures of wealth inequality behave in this global robustness exercise.

[Table 3 about here]

Table 3 reports a set of statistics for the two outcomes that are the focus of our analysis. The minimum, maximum, mean, median and standard deviation are listed for all priors.

Moreover, in order to provide a more intuitive summary of such a rich analysis, several figures are included. All figures show eleven panels. Each panel represents the non-parametric kernel density estimate of a variable for a particular prior. The first four panels deal with the Uniform experiments, while the remaining seven deal with the Beta ones.

[Figures (2), (3) and (4) about here]

4.1 Precautionary Savings

Figure (2) displays a measure of precautionary savings, the percentage increase in the aggregate saving rate of the IM economy compared to the CM one. The results are striking. When considering a set of eleven Monte Carlo calibrations of size 2,500, each coming from a different prior, the percentage increase in the aggregate saving rate is always modest. There are virtually no cases in which the increase is as high as 6%, with the bulk of the replications being included in the $0.5\% - 3\%$ increase range.

In Table 3 we see that the maximum increase in the saving rate is 11%, which is, not surprisingly, found in experiment *U1* (the no borrowing case), while the minimum increase is always very close to zero. The median increase is less than a percentage point in eight out of the eleven priors. According to this quantitative theory, precautionary savings are small indeed.

These results could be potentially driven by an odd behavior of the saving rate, or by the model failing to account for some other features of the data related to the saving behavior, such as the capital/output ratio. As figures (3) and (4) display, generally this is not the case. The saving rate falls in reasonable ranges, as the capital/output ratio does. There are indeed some replications that imply values that are grossly over or understated, just like the extreme cases in panels 1 to 4, namely the Uniform priors. In order to tackle this potential issue, we provide some additional robustness checks below. However, when focusing on the Beta cases, on *B2* in particular (which has the least counterfactual quantitative implications), we see that precautionary savings are highly concentrated around the 1% value.[6]

Interestingly, relaxing the borrowing constraint does not have any major impact on precautionary savings. This is shown in the first four panels of figure (2). Even when agents can borrow a substantial amount of resources (up to three times their average income, as displayed in figure 7), allowing them to better smooth consumption, this does not affect much precautionary savings: they do decrease, but not spectacularly so. At the same time, we can appreciate the effect of a highly left skewed prior: for the $Beta(2, 5)$ case reported in panel *B4* precautionary savings become a fraction of 1%. Income shocks that are not very persistent and with a very low variance, together with relatively impatient and not very risk averse individuals, explain the result.

[Figures (5), (6), and (7) about here]

[6] This very set of comments apply when we consider another measure of precautionary savings, the percentage increase in the economy's capital stock when moving from the CM economy to the IM one. See figures (15) and (16) in Appendix D.

4.2 Inequality and Exogenous Borrowing Constraints

It has generally been found that the Aiyagari (1994) model (and many of its variants) cannot generate high levels of wealth inequality, as described by the Gini index.[7] This result is confirmed by our global sensitivity analysis, with some caveats.

As for our model economy, figure (5) shows that wealth tends to be more concentrated than in the calibrations reported by Aiyagari (1994). The first panel shows that, without access to borrowing, the Gini index does not vary much from its average value of 0.45, a value above the calibrations provided by Aiyagari (1994). In Table 3 we see that for the *U1* case the minimum is 0.24, the maximum is 0.55, and the s.d. is 0.03.

However, Aiyagari (1994) did not provide any robustness checks related to the borrowing limit. Panels 1 to 4 in figure (5) present such an analysis. They consider the equilibrium effects of relying on a less strict exogenous borrowing constraint. When moving from the no borrowing case to a case where people can borrow, we can appreciate a first order effect on the Gini index. The average of the index increases, together with its dispersion, being equal to 0.56 in case *U4*. The minimum is now 0.25, the maximum is 0.88, and the s.d. is 0.09.

Related to this, there is an important caveat. When allowing for negative assets, the Gini coefficient need not be between zero and one. In order to deal with a variable that can take negative values while giving a well-behaved Gini index which is comparable to non-negative variables, we apply the adjustment proposed by Chen, Tsaur and Rhai (1982). Without correcting the index this way, the same panels would show an unpleasant feature: for some economies the Gini index would be above one.[8]

Overall, the median wealth Gini index now ranges from 0.49 to 0.64, depending on the prior being considered. The first order effect of relaxing the borrowing constraint notwithstanding, these figures are still far below from what is observed in the data, approximately 0.8. Hence, apart from a handful of calibrations in experiments *U3*, *U4*, *B1* and *B4*, the model cannot match the observed degree of wealth inequality in the U.S.

Figure (6) reports the percentage of households that have a negative net worth. As discussed above, the average (and median) values are consistent with what is observed in the data. Finally, the amount of resources that can be borrowed compared to the average income appear to be plausible numbers, as the graphs in figure (7) show.

[7] We do not comment the results on another inequality measure, the coefficient of variation. The estimated densities had the very same shape as the ones for the corresponding Gini index. See figures (18), (19) and (20) in Appendix D.

[8] See figure (25) in Appendix E, which considers a similar graph without implementing the adjustment for an alternative case.

4.3 (Global) Comparative Dynamics

With the results of the simulations in our hands, we can perform a surrogate of comparative dynamics analysis. Rather than taking the analytical derivative of our outcomes with respect to a parameter, we can run a regression, with either precautionary savings or the Gini index on the left hand side, and the corresponding simulated parameters on the right hand side.[9]

[Table 4 about here]

Table 4 shows the standardized beta coefficients of such regressions. It is interesting to notice that the same parameter has different effects on the two outcome variables. As for the wealth Gini index, the parameters that quantitatively have the largest impact on this inequality measure are the capital share α, the borrowing limit b, and the depreciation rate δ. The intuition goes as follows. The higher α, the higher the marginal product of capital and the higher the interest rate. This price change induces an income effect, which is going to kick in more likely for wealthier agents, hence the reduction in inequality. The depreciation rate has a similar, but opposite, mechanism. The intuition for the borrowing limit is simple: the less resources people can borrow, the higher the precautionary savings and (possibly) the lower the range of wealth. This reduces the degree of wealth concentration. Although their effect is quantitatively less important, a higher discount factor β and a higher risk aversion σ tend to compress inequality, by inducing all agents to save more. Surprisingly, the parameters driving the uncertainty in the economy, ρ_y and σ_y, have a relatively low impact on the Gini index.

As for precautionary savings, the parameters that affect them the most are the relative risk aversion σ and the s.d. of the income shocks σ_y.[10]

[9] This is a well posed exercise, because the parameters are changing randomly across replications. As a result, the source of variation is exogenous. However, when drawing the parameters, the independence of the distributions is not imposed, hence some spurious correlation is detected in our artificial dataset. We ran a 2SLS estimation procedure as well, and the results changed very little. All regressions include a constant term, not reported in the table.

[10] A non-linear relationship is found for most parameters. The square of the parameters is significant in almost every experiment, while including them as additional regressors increases only marginally the explanatory power (usually by less than a percentage point). The point estimates of the square terms are positive only for α and σ in the Gini regressions, and for α, σ, σ_y and b in most of the precautionary savings ones. Higher order polynomials, as usual, introduce multicollinearity for most specifications.

4.4 Robustness[2]

This subsection considers a robustness check for our global robustness analysis. With some abuse of language, it could be said that we are going to perform robustness squared.

A possible source of criticism can come from the relatively naive way in which parameter vectors are selected, and their potentially counterfactual quantitative implications. Unfortunately, before solving the model it is very complicated to guess which replications are going to "fail" in some economic dimension. Restricting ex-ante the parameters' space and the set of plausible priors is somewhat challenging.

However, in order to tackle this conceivable issue, the simulations can be restricted ex-post. After having solved the model for 2,500 times relying on a specific prior, the simulations that imply implausible values for some simulated moments are dropped.

More precisely, the following restrictions on the capital/output ratio, the saving rate, and the percentage of people in debt are imposed: 1) $1.5 \leq \frac{K^{IM}}{Y^{IM}} \leq 3.5$, 2) $0.125 \leq \frac{\delta K^{IM}}{Y^{IM}} = \frac{S^{IM}}{Y^{IM}} \leq 0.275$, and 3) $0.05 \leq \%NNW \leq 0.2$. Only the simulations that satisfy all three criteria at the same time are kept. This means that we are discarding between 595 replications for the case *B2* and 2,066 replications for the case *B3*. It goes without saying that the benchmark case *U1* does not satisfy the third requirement.

[Table 5 about here]

Table 5 displays the new ranges of the parameters once the restrictions are implemented. With the uniform priors, most parameters' bounds are not affected. The changes tend to be small, and affect mostly $\overline{\alpha}$, $\underline{b}$ and $\overline{b}$. Differently, with the Beta experiments, most parameters' bounds do change. Some changes are relatively minor. For example, $\underline{\beta}$ is not affected much, its most stringent value now being 0.946 for case *B3*. Others are larger. For example $\overline{\sigma}$ becomes 2.965 for case *B4*, $\underline{\sigma}$ becomes 1.162 for case *B3*, $\overline{\rho_y}$ becomes 0.957 for case *B2*, and $\overline{\sigma_y}$ becomes 0.190 for case *B4*.

Figures (8), (9) and (10) display the new distributions of the three moments that are now subject to the restrictions above.

[Figures (8), (9), (10), (11) and (12) about here]

Estimating the densities only on the cases that meet the above criteria delivers figures (11) and (12). From the graphs we can conclude that the results presented in the previous sections are not altered substantially by these additional requirements.

Since precautionary savings were already found to be consistently small, table 3 shows the results for the restricted cases only for wealth inequality. These are reported in parenthesis in the first 5 columns. As expected, the extreme values are eliminated by the restrictions: the minima increase by up to 19 points and the maxima decrease by up to 23 points. As a consequence, it is possible to conclude that all the replications that were found to get close to the data were failing in some of the three moments above. Finally, the median values of the wealth Gini index now range from 0.52 to 0.57.

4.5 Discussion

At least four aspects of the analysis call for further discussion.

It goes without saying that also this empirical methodology is not free of potential issues. On the one hand, the priors' choice is somewhat arbitrary: there are no clear indications on how to specify these distributions. This aspect of the procedure can only be problem-dependent. Hence, in general, it is not possible to assess how much the results are going to be contingent on the specific priors that are being used. However, this obstacle can be partially overcome by considering several different cases, with relatively flexible and general specifications. At the same time, the priors should try to minimize the number of replications that entail implausible quantitative implications along some dimension. Finally, relevant empirical studies should be used as an additional guide in this choice. This is the approach attempted in this paper.

Another complication is the computational burden. The computational time that is needed to complete the procedure can become easily intractable, even with powerful computers. Unless a researcher has access to a supercomputer, the computational costs are still high: on a moderately fast desktop computer it takes approximately a week of continuous computer time to complete an experiment.[11] Hence, the solution method must strike a balance between the computational complexity, the implied numerical errors and its reliability. The last point seems to be quite important for non-linear models, such as the one we are considering here. The solution method has to guarantee that convergence problems are not going to riddle too many replications, especially for non-random subsets of the parameters. Otherwise, this could potentially bias the results. However, once a robust solution method is available, the extension for the Monte Carlo calibration methodology is extremely simple: it boils down to an additional outer loop where a vector of random parameters is drawn at each replication.

[11] Incidentally, this consideration made it unfeasible to rely on $10,000$ replications, as in Canova's original contribution. However, for the HA economy under study, when considering only the first $1,000-1,500$ cases the results were very similar to the ones of the full experiment, suggesting that the number of replications is large enough.

A natural question that arises is whether this method can be applied to models that are currently used in the quantitative macroeconomics field. The most recent contributions proposing HA models allow for several different sources of uncertainty and several layers of heterogeneity. It goes without saying that the more complex the model, the less likely the feasibility and the computational success of Canova's methodology. However, recent developments in numerical methods for economics can prove useful. The relatively simple and powerful endogenous grid method can cut the computational time substantially, as found by Barillas and Fernandez-Villaverde (2007) and Carroll (2006). The same objective can be achieved by more efficient simulation methods, as proposed by Nishimura and Stachurski (2010). Trivially, a supercomputer with (say) 256 nodes would make the computational burden tractable for several richer models. However, this does not seem to be a reasonable requirement. Resource constraints and the added complexity of parallel computing in large clusters make this option not viable for many quantitative macroeconomists. A more plausible option for sizeable computational gains could come from GPU computing, as recently discussed in Aldrich, Fernandez-Villaverde, Gallant, and Rubio-Ramirez (2011), but this opportunity is yet to be carefully explored.

Finally, unlike fully Bayesian methods proposed for example by DeJong, Ingram, and Whiteman (1996), Canova's procedure does not allow the model to exploit the information contained in the data, to achieve a better parameterization while reducing the model's parameter uncertainty.

5 Conclusions

This paper contributed to the literature on calibration methods for equilibrium HA macroeconomic models. It showed that global (and more robust) sensitivity analyses are also feasible for such rich economies. Unless one is willing to rely on supercomputers, the computational costs are still high: on a moderately fast desktop computer it takes approximately a week of continuous computer time to complete an experiment. However, once a robust solution method is available, the extension for the simulation methodology is extremely simple: it boils down to an additional outer loop where a vector of random parameters is drawn.

It is worth stressing that this type of exercise is meant to complement more traditional calibration studies (or sophisticated structural estimation ones), not to substitute them. Once a quantitative theory is proposed and a set of empirically relevant values for the parameters are available, one can provide more robust evidence on the original quantitative findings, possibly suggesting in which cases some outcomes are more likely to be observed.

The results of our experiments confirm Aiyagari's findings: precautionary savings are small indeed, even when considering priors for the parameter uncertainty that are very different and sometimes highly dispersed. Morever, the results show that the exogenous borrowing limit leads to relatively large changes in the wealth Gini index. However, even with increases in the median of the index of up to 15 points, the model is still very far from accounting for the one observed in the U.S. economy.

The findings are confirmed even when restricting the calibrations to satisfy some relevant features of the available macroeconomic data.

The parameters that quantitatively have the largest effects on determining the wealth Gini index are the capital share, the borrowing limit, and the depreciation rate. The parameters affecting precautionary savings the most are the risk aversion and the standard deviation of the income shocks.

A similar methodology can be implemented for other HA models, quantifying for example the welfare effects of eliminating the social security system or the ones arising from changing the tax code. A positive aspect of this analysis is that, when applied to the evaluation of policy reforms, it provides distributions and ranges of welfare effects, an approach to empirical research pushed forward by Manski (1995), among others. We leave these extensions and modifications for future work.

Parameter	*Description*	*Min*	*Max*
Model Period	*Yearly*		
α	*Capital share*	$\underline{\alpha} = 0.25$	$\overline{\alpha} = 0.45$
β	*Rate of time preference*	$\underline{\beta} = 0.94$	$\overline{\beta} = 0.98$
δ	*Capital depreciation rate*	$\underline{\delta} = 0.03$	$\overline{\delta} = 0.12$
σ	*CRRA*	$\underline{\sigma} = 0.5$	$\overline{\sigma} = 3.5$
ρ_y	*Persistence of the AR(1) earnings process*	$\underline{\rho_y} = \{0.5, 0.8\}$	$\overline{\rho_y} = 0.99$
σ_y	*S.d. of the innovation in the AR(1) earnings process*	$\underline{\sigma_y} = 0.04$	$\overline{\sigma_y} = 0.25$
b	*Borrowing limit*	$\underline{b} = \{0, -1, -2, -2.5\}$	$\overline{b} = 0$

Table 1: Model parameters and their support

Experiment 1 (U1)	*Experiment 2 (U2)*	*Experiment 3 (U3)*	*Experiment 4 (U4)*
$U\left(\underline{p},\overline{p}\right);\underline{b}=\overline{b}=0$	$U\left(\underline{p},\overline{p}\right);\underline{b}=-1.0$	$U\left(\underline{p},\overline{p}\right);\underline{b}=-2.0$	$U\left(\underline{p},\overline{p}\right);\underline{b}=-2.5$

Experiment 5 (B1)	*Experiment 6 (B2)*	*Experiment 7 (B3)*	*Experiment 8 (B4)*
$Beta(2,2;\underline{p},\overline{p})$	$Beta(10,10;\underline{p},\overline{p})$	$Beta(5,2;\underline{p},\overline{p})$	$Beta(2,5;\underline{p},\overline{p})$

Experiment 9 (B5)	*Experiment 10 (B6)*	*Experiment 11 (B7)*
$Beta(10,10;\underline{p},\overline{p})$	$Beta(10,10;\underline{p},\overline{p})$	$Beta(10,10;\underline{p},\overline{p})$
$Beta(10,10;\underline{\rho_y},\overline{\rho_y})$	$Beta(5,2;\underline{\rho_y},\overline{\rho_y})$	$Beta(5,2;\underline{\rho_y},\overline{\rho_y})$
$Beta(5,2;\underline{\sigma_y},\overline{\sigma_y})$	$Beta(10,10;\underline{\sigma_y},\overline{\sigma_y})$	$Beta(5,2;\underline{\sigma_y},\overline{\sigma_y})$

Table 2: Experiments - Uniform and Beta Prior Distributions over a generic parameter p

	Wealth Gini					*Prec.*	*Savings*:	*%Change*		
	Min	*Max*	*Mean*	*Median*	*S.d.*	*Min*	*Max*	*Mean*	*Median*	*S.d.*
U1	.24	.55	.45	.45	.03	.001	11.0	1.3	.8	.01
U2	.25 (.41)	.76 (.61)	.50 (.52)	.49 (.52)	.06 (.04)	.001	10.7	1.1	.7	.01
U3	.25 (.44)	.85 (.61)	.54 (.53)	.53 (.53)	.09 (.04)	.001	10.7	1.1	.7	.01
U4	.25 (.42)	.88 (.61)	.56 (.54)	.54 (.54)	.09 (.04)	.001	10.7	1.1	.6	.01
B1	.36 (.44)	.82 (.61)	.56 (.54)	.55 (.54)	.07 (.04)	.001	6.8	1.0	.8	.01
B2	.46 (.49)	.71 (.61)	.56 (.56)	.56 (.56)	.03 (.02)	.001	2.9	.9	.9	.01
B3	.39 (.46)	.72 (.60)	.51 (.53)	.50 (.52)	.04 (.03)	.001	5.3	1.5	1.4	.01
B4	.43 (.47)	.84 (.61)	.61 (.56)	.61 (.57)	.06 (.03)	.001	3.2	.4	.3	.01
B5	.46 (.47)	.72 (.61)	.57 (.56)	.57 (.56)	.04 (.02)	.001	4.2	1.7	1.6	.01
B6	.48 (.49)	.69 (.61)	.57 (.56)	.57 (.56)	.03 (.02)	.001	2.5	.7	.7	.01
B7	.47 (.48)	.71 (.60)	.57 (.56)	.57 (.57)	.04 (.02)	.001	4.3	1.3	1.2	.01

Table 3: Monte Carlo Results - Wealth Gini (Restricted cases in parenthesis) and Precautionary Savings

	U1	*U2*	*U3*	*U4*	*B1*	*B2*	*B3*	*B4*	*B5*	*B6*	*B7*
Wealth Gini											
α	-0.54	-0.59	-0.58	-0.57	-0.57	-0.60	-0.49	-0.60	-0.58	-0.60	-0.59
β	-0.21	-0.20	-0.17	-0.16	-0.16	-0.16	-0.15	-0.18	-0.17	-0.14	-0.14
δ	0.48	0.51	0.49	0.49	0.50	0.52	0.42	0.54	0.50	0.52	0.51
σ	-0.31	-0.20	-0.15	-0.14	-0.13	-0.15	-0.15	-0.13	-0.17	-0.13	-0.15
ρ_y	0.34	0.13	0.06	0.04	0.07	0.12	0.19	0.01	0.11	0.07	0.08
σ_y	0.30	0.22	0.19	0.18	0.15	0.12	0.06	0.27	0.14	0.15	0.15
b	–	-0.43	-0.54	-0.56	-0.55	-0.56	-0.68	-0.42	-0.55	-0.53	-0.53
R^2	0.83	0.90	0.91	0.92	0.94	0.98	0.94	0.97	0.99	0.97	0.97
Prec. Savings											
α	-0.16	-0.13	-0.11	-0.10	-0.09	-0.09	-0.10	-0.07	-0.08	-0.08	-0.05
β	-0.15	-0.14	-0.14	-0.14	-0.12	-0.13	-0.18	-0.07	-0.11	-0.09	-0.08
δ	0.01	-0.02	-0.04	-0.05	-0.05	-0.07	-0.04	-0.04	-0.06	-0.06	-0.06
σ	0.48	0.49	0.50	0.50	0.49	0.54	0.45	0.50	0.52	0.44	0.42
ρ_y	-0.21	-0.21	-0.20	-0.20	-0.15	-0.13	-0.48	-0.01	-0.13	-0.51	-0.50
σ_y	0.65	0.65	0.64	0.64	0.73	0.79	0.63	0.79	0.79	0.65	0.66
b	–	0.06	0.10	0.11	0.09	0.10	0.10	0.06	0.10	0.07	0.07
R^2	0.74	0.74	0.74	0.74	0.83	0.95	0.88	0.86	0.95	0.88	0.89

Table 4: Comparative Dynamics - Linear Regressions, Standardized Beta Coefficients

Parameter	*U2*	*U3*	*U4*	*B1*	*B2*	*B3*	*B4*	*B5*	*B6*	*B7*
$\overline{\alpha}$	0.445	0.445	0.445	0.441	0.408	0.417	0.407	0.408	0.408	0.408
$\underline{\alpha}$	0.250	0.250	0.250	0.260	0.287	0.280	0.254	0.291	0.287	0.287
$\overline{\beta}$	0.980	0.980	0.979	0.979	0.973	0.979	0.969	0.972	0.973	0.971
$\underline{\beta}$	0.940	0.940	0.940	0.941	0.946	0.946	0.940	0.946	0.946	0.946
$\overline{\delta}$	0.120	0.120	0.120	0.118	0.106	0.119	0.100	0.106	0.106	0.106
$\underline{\delta}$	0.039	0.039	0.039	0.038	0.051	0.054	0.037	0.051	0.051	0.051
$\overline{\sigma}$	3.499	3.499	3.499	3.477	2.968	3.478	2.965	2.978	2.968	2.968
$\underline{\sigma}$	0.501	0.501	0.501	0.602	0.992	1.162	0.505	0.992	0.992	0.992
$\overline{\rho_y}$	0.990	0.990	0.990	0.988	0.957	0.989	0.963	0.958	0.989	0.989
$\underline{\rho_y}$	0.800	0.800	0.800	0.801	0.832	0.841	0.800	0.832	0.836	0.836
$\overline{\sigma_y}$	0.250	0.250	0.250	0.247	0.207	0.248	0.190	0.248	0.204	0.248
$\underline{\sigma_y}$	0.040	0.040	0.040	0.040	0.078	0.079	0.042	0.080	0.078	0.080
$\overline{b}$	−0.008	−0.011	−0.014	−0.087	−0.329	−0.044	−0.293	−0.329	−0.329	−0.329
$\underline{b}$	−1.000	−1.997	−2.496	−2.402	−2.057	−1.603	−2.453	−1.954	−2.041	−2.004

Table 5: Model parameters and their support - Restricted simulations

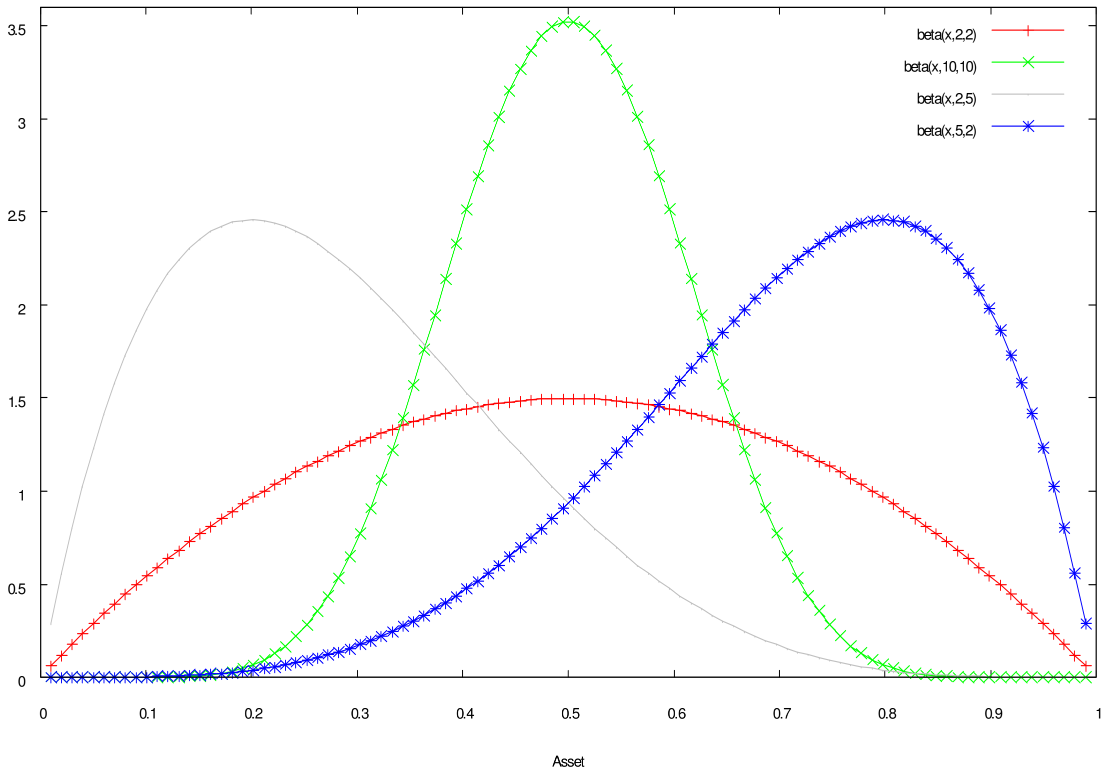

Figure 1: Monte Carlo Experiments: $Beta(p, q)$ priors.

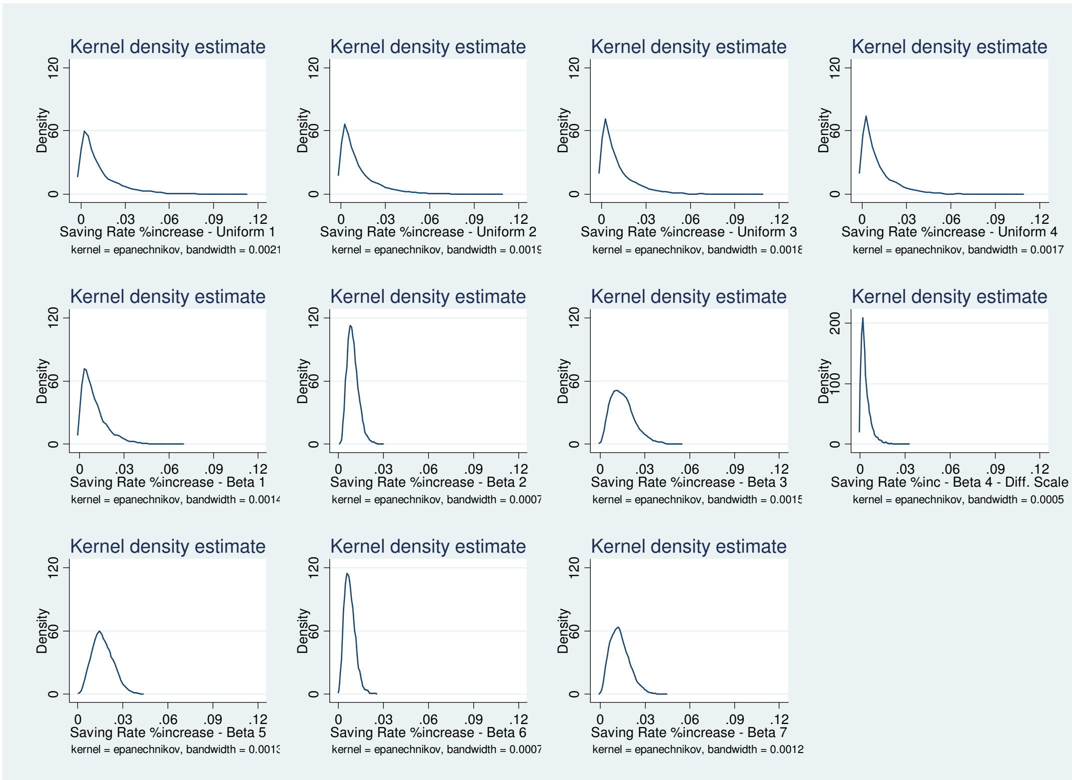

Figure 2: Precautionary Savings - Estimated Densities for the eleven Monte Carlo Experiments

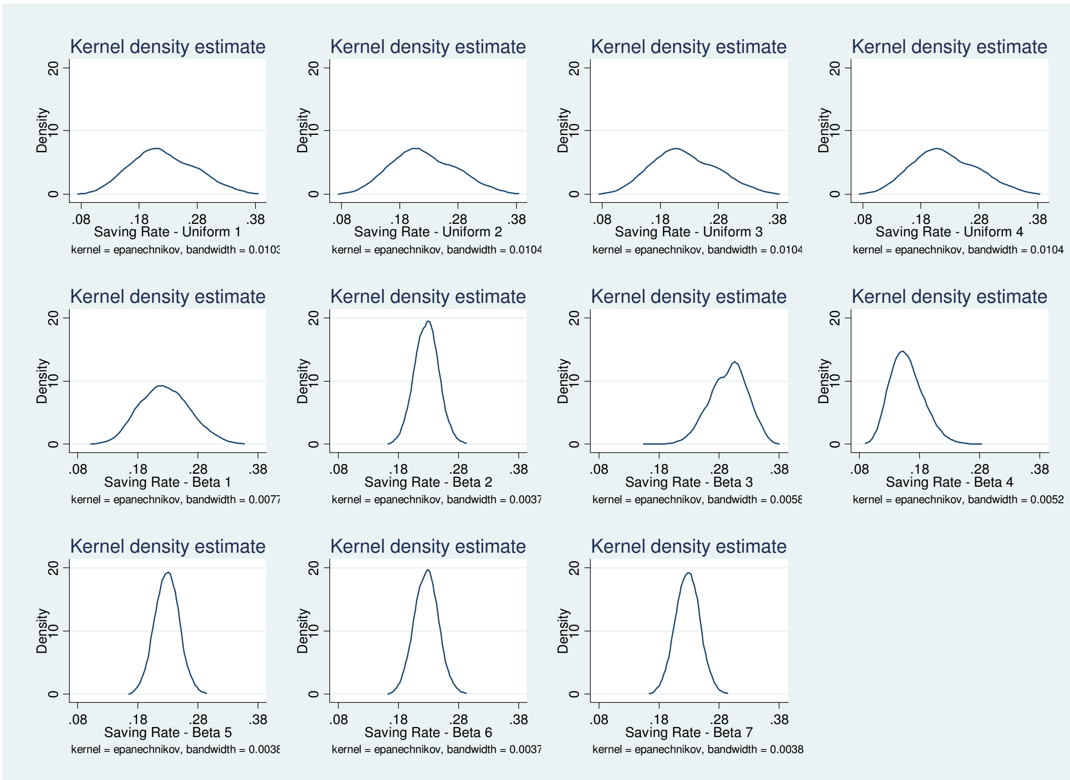

Figure 3: Saving Rates - Estimated Densities for the eleven Monte Carlo Experiments

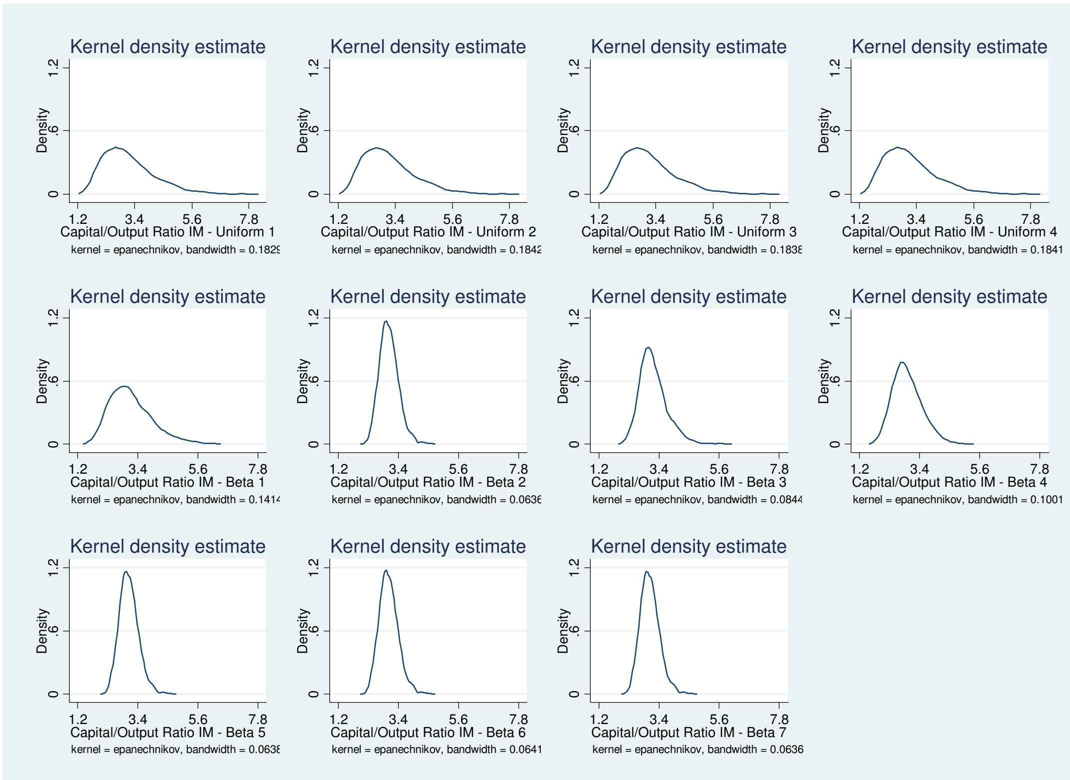

Figure 4: Capital/Output Ratios (IM) - Estimated Densities for the eleven Monte Carlo Experiments

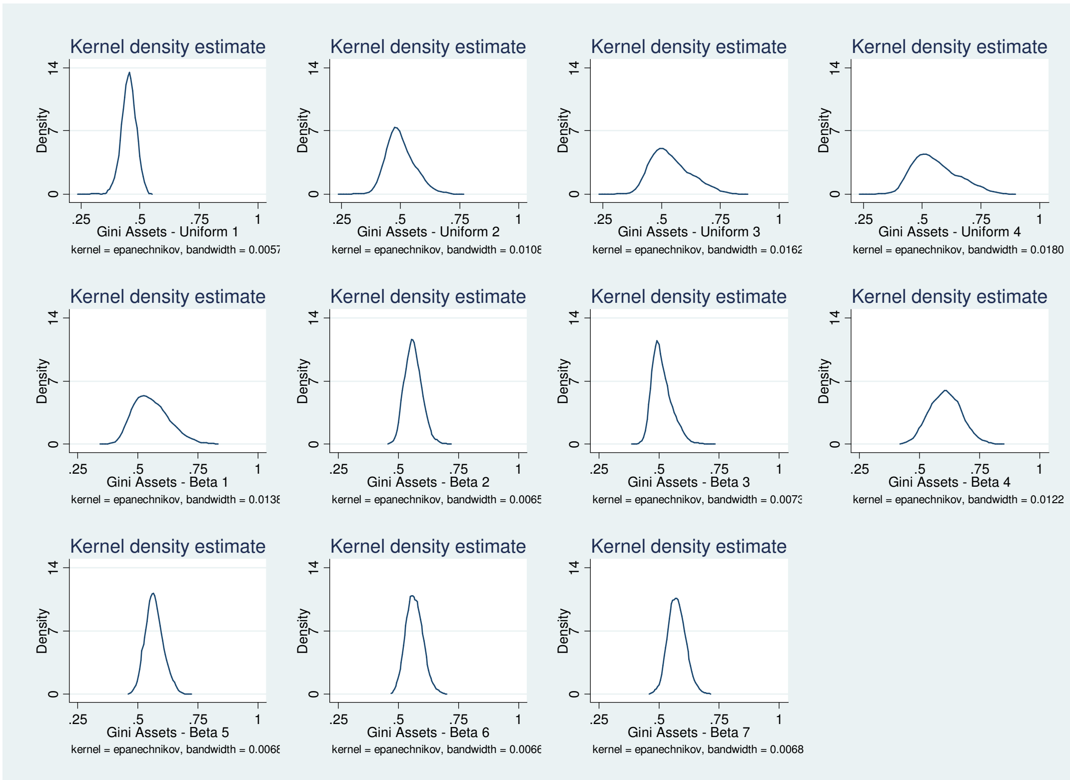

Figure 5: Wealth Gini Index - Estimated Densities for the eleven Monte Carlo Experiments

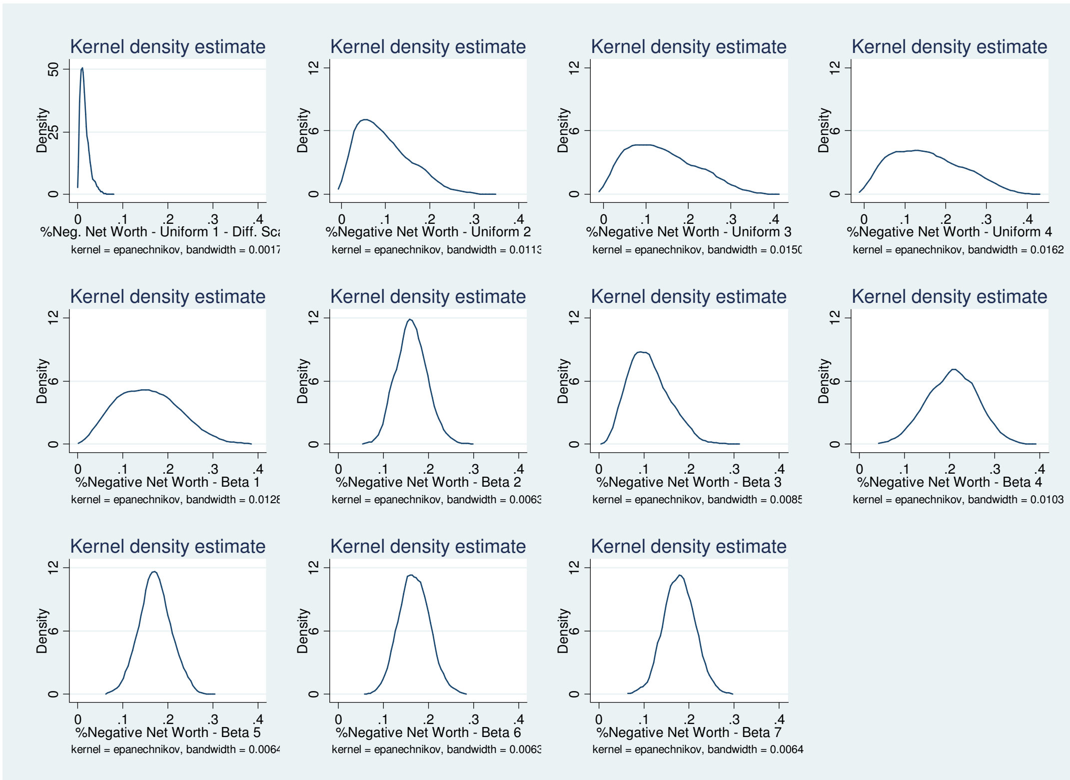

Figure 6: % with Negative Net Worth - Estimated Densities for the eleven Monte Carlo Experiments

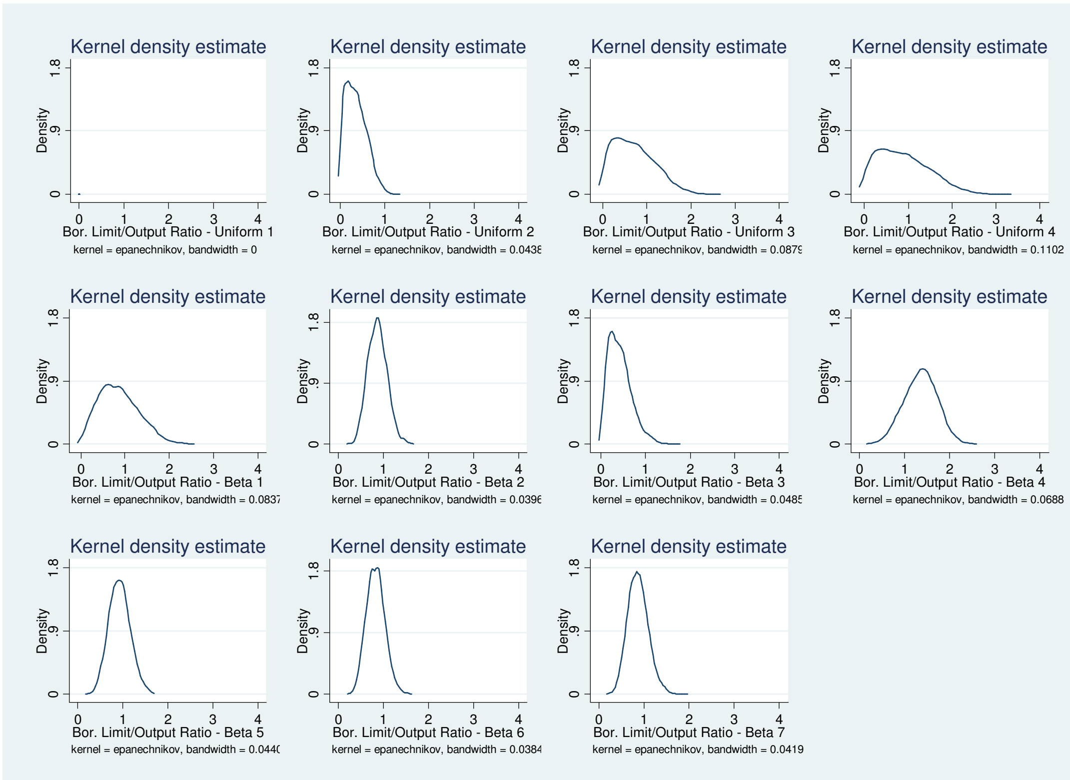

Figure 7: Borrowing Limit/Ouput Ratios - Estimated Densities for the eleven Monte Carlo Experiments

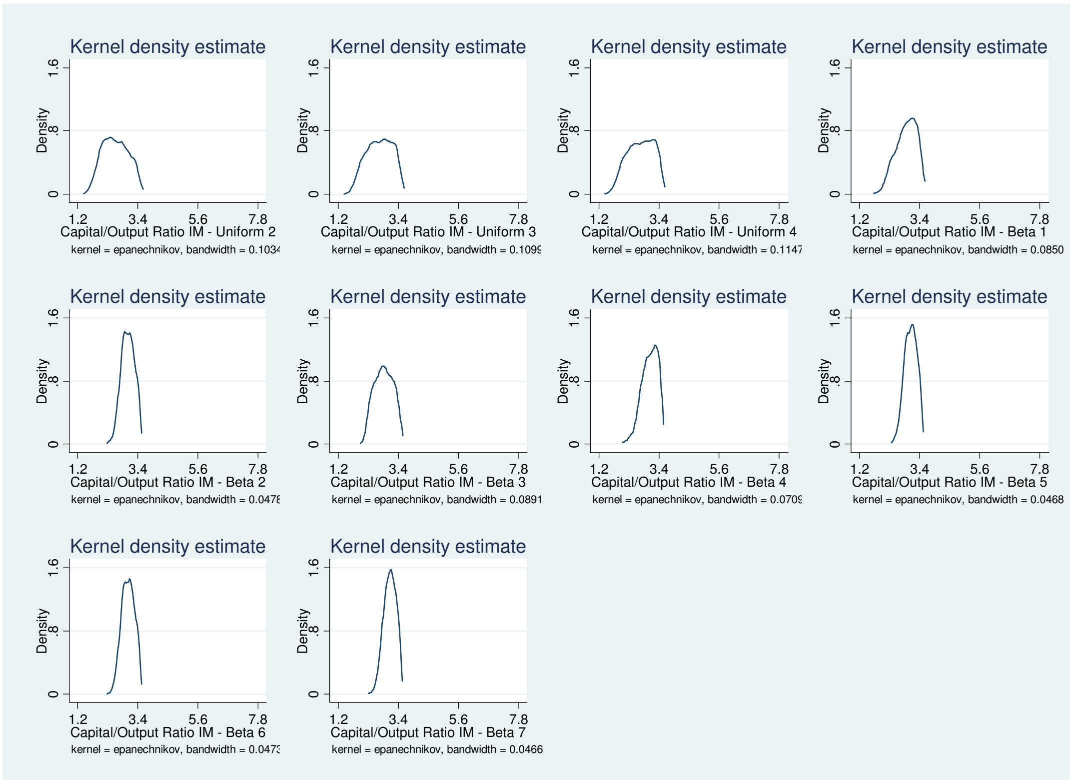

Figure 8: Capital/Output Ratios (IM) - Estimated Densities for the eleven Restricted Monte Carlo Experiments

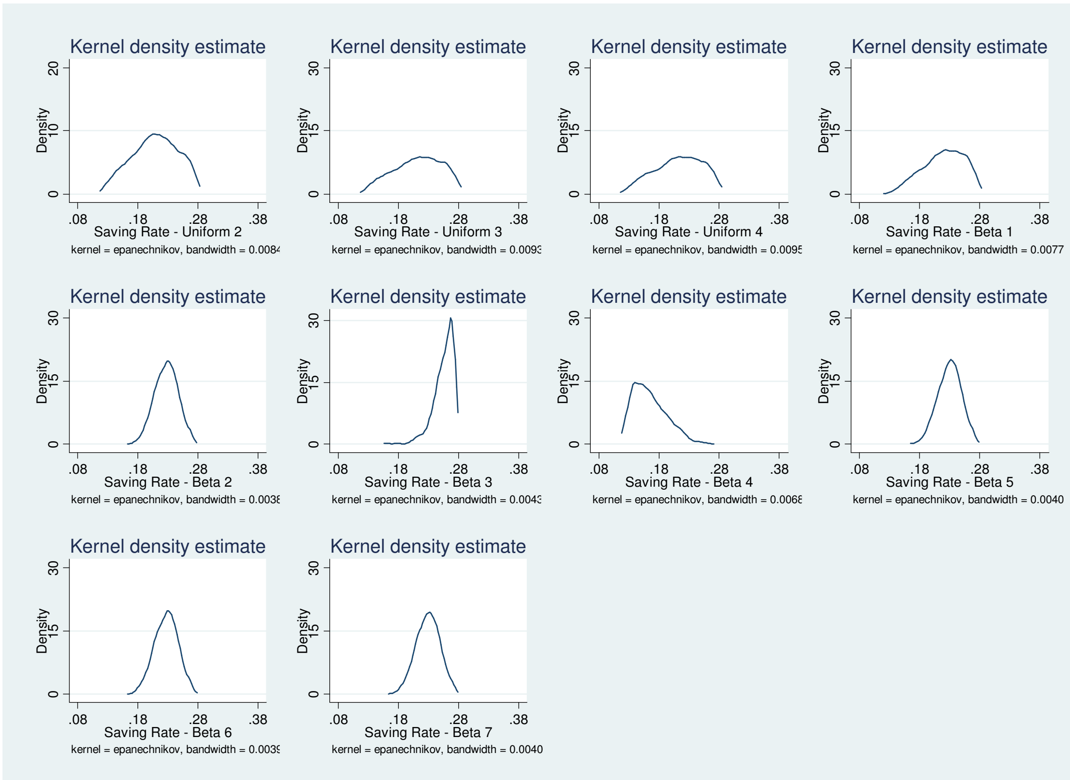

Figure 9: Saving Rates - Estimated Densities for the eleven Restricted Monte Carlo Experiments

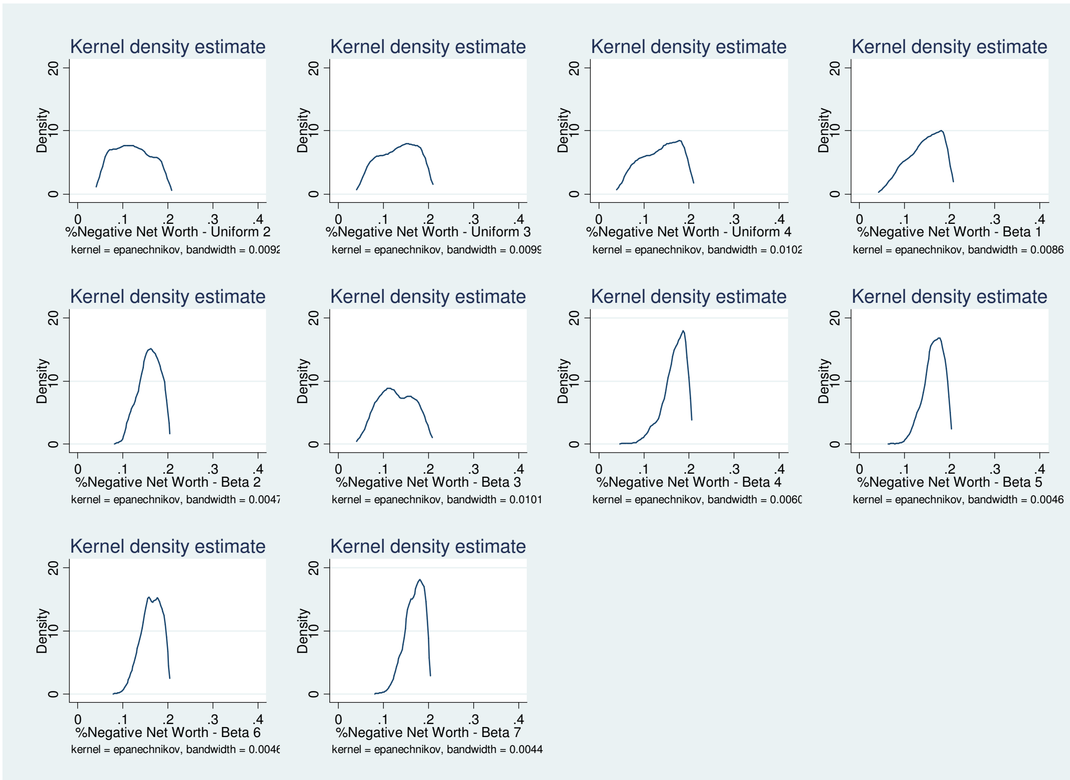

Figure 10: % with Negative Net Worth - Estimated Densities for the eleven Restricted Monte Carlo Experiments

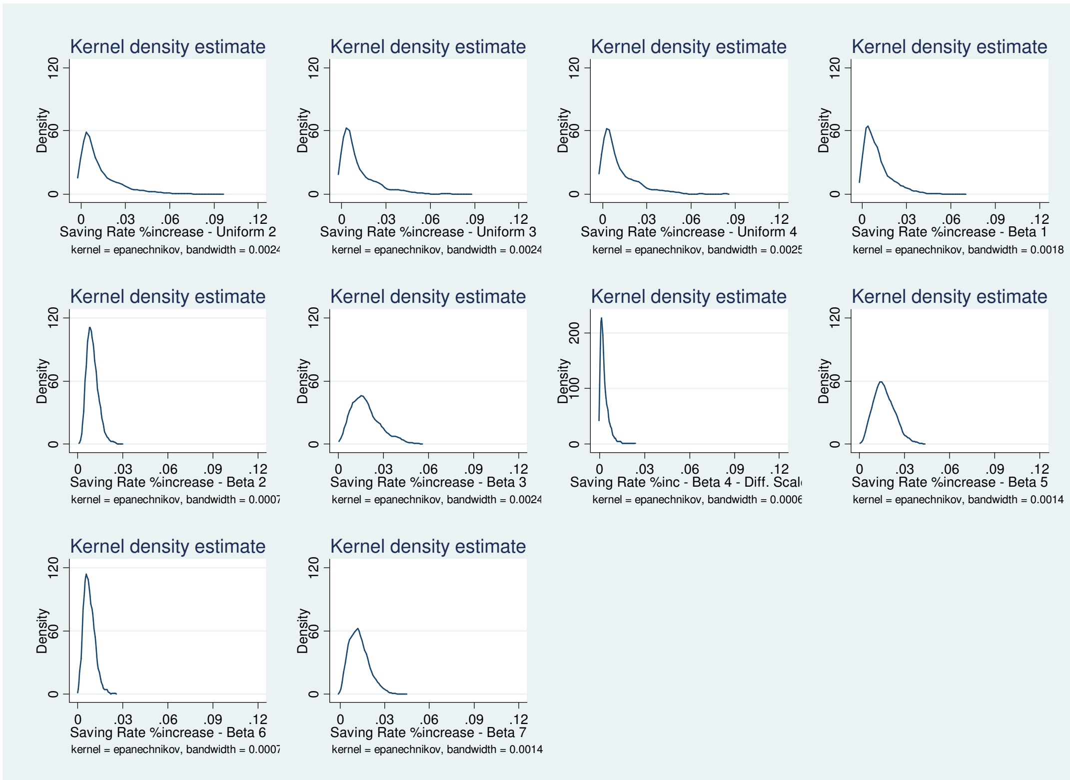

Figure 11: Precautionary Savings - Estimated Densities for the eleven Restricted Monte Carlo Experiments

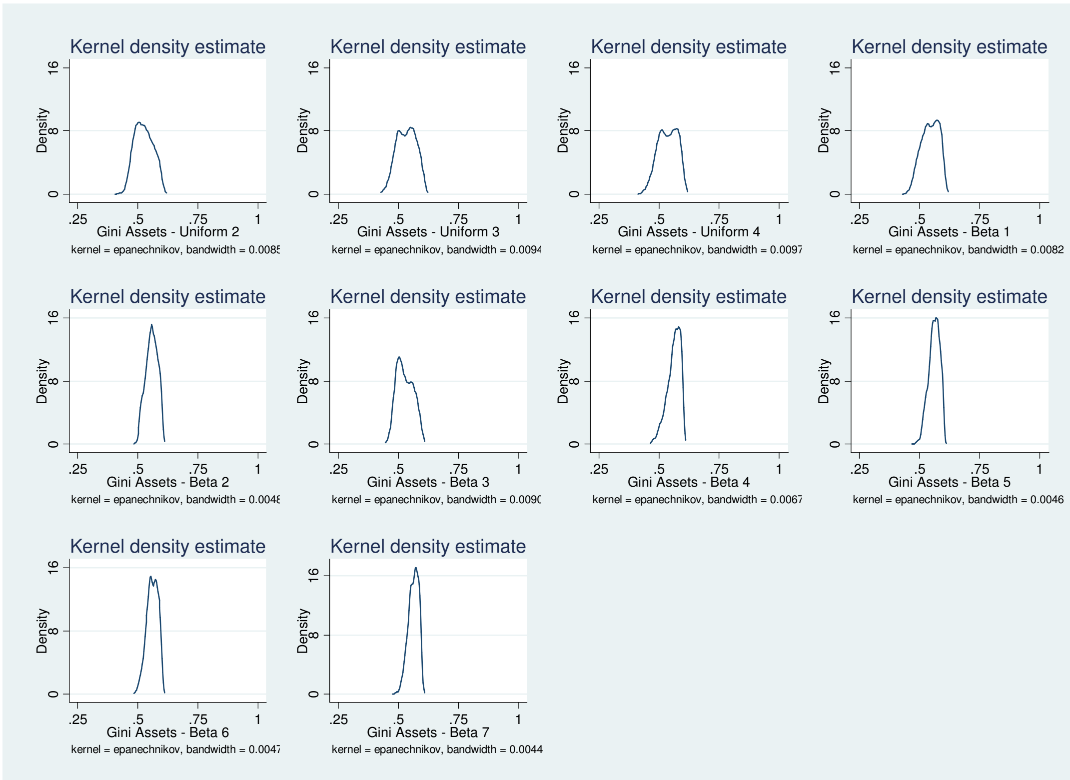

Figure 12: Wealth Gini Index - Estimated Densities for the eleven Restricted Monte Carlo Experiments

References

Aiyagari, R. (1994). "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics,* Vol. 109, 659-684.

Aldrich, E., Fernandez-Villaverde, J., Gallant, R., and Rubio-Ramirez, J. (2011). "Tapping the supercomputer under your desk: Solving dynamic equilibrium models with graphics processors," *Journal of Economic Dynamics and Control*, Vol. 35, 386-393.

Attanasio, O., Banks, J., Meghir, C., and Weber, G. (1999). "Humps and Bumps in Lifetime Consumption," *Journal of Business and Economic Statistics*, Vol. 17, 22-35.

Barillas, F., and Fernandez-Villaverde, J. (2007). "A generalization of the endogenous grid method," *Journal of Economic Dynamics and Control*, Vol. 31, 2698-2712.

Budria Rodriguez, S., Diaz-Gimenez, J., Quadrini, V., and Rios-Rull, V. (2002). "Updated Facts on the U.S. Distributions of Earnings, Income, and Wealth," *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol. 26, 2-35.

Cagetti, M., and De Nardi, M. (2008). "Wealth Inequality: Data and Models," *Macroeconomic Dynamics*, Vol. 12, 285-313.

Canova, F. (1994). "Statistical Inference in Calibrated Models," *Journal of Applied Econometrics*, Vol. 9, S123-S144.

Carroll, C. (2006). "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics Letters*, Vol. 91, 312-320.

Chen, C., Tsaur, T., and Rhai, T. (1982). "The Gini Coefficient and Negative Income," *Oxford Economic Papers*, Vol. 34, 473-478.

DeJong, D., Ingram, B. and Whiteman, C. (1996). "A Bayesian Approach to Calibration," *Journal of Business and Economic Statistics*, Vol. 14, 1-9.

Diaz-Gimenez, J., Quadrini, V., and Rios-Rull, J.V. (1997). "Dimensions of Inequality: Facts on the U.S. Distributions of Earnings, Income, and Wealth," *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol. 21, 3-21.

Guvenen, F. (2009). "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, Vol. 12, 58-79.

Heathcote, J, Storesletten, K., and Violante, G. (2009). "Quantitative Macroeconomics with Heterogeneous Households," *Annual Review of Economics*, Vol. 1, 319-354.

Manski, C. (1995). *Identification Problems in the Social Sciences*, Harvard University Press, Cambridge (MA).

Nishimura, K., and Stachurski, J. (2010). "Perfect Simulation of Stationary Equilibria," *Journal of Economic Dynamics and Control*, Vol. 34, 577-584.

Rios-Rull, V. (1999). "Computation of Equilibria in Heterogenous Agent Models," in *Computational Methods for the Study of Dynamic Economies: An Introduction*, (Ramon Marimon and Andrew Scott, Eds.), Oxford University Press.

Quadrini, V., and Rios-Rull, J.V. (1997). "Understanding the U.S. Distribution of Wealth," *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol. 21, 22-36.

Wolff, E. (1998). "Recent Trends in the Size Distribution of Household Wealth," *Journal of Economic Perspectives,* Vol.12, 131-150.

Appendix A - The Model and its Recursive Representation

6 Stationary Equilibrium

In this Section we first define the problem of the agents in their recursive representation, then we define the problem of the firm, finally we provide a formal definition of the equilibrium concept used in this model, the recursive competitive equilibrium.

The individual state variables are the labor endowment $\varepsilon \in \mathcal{E} = [0, \overline{\varepsilon}]$, and asset holdings $a \in \mathcal{A} = [-b, \overline{a}]$. The stationary distribution is denoted by $\mu(\varepsilon, a)$.

6.1 Problem of the workers

The value function of an agent whose current asset holdings are equal to a, and whose current labor endowment is ε is denoted with $V(\varepsilon, a)$. The problem of these agents can be represented as follows:

$$V(\varepsilon, a) = \max_{c,a'} \left\{ u(c) + \beta E_{\varepsilon'|\varepsilon} V(\varepsilon', a') \right\} \tag{1}$$

$$s.t.$$

$$c + a' = (1+r)\, a + w\varepsilon$$
$$\varepsilon' = \rho_y \varepsilon + \eta', \eta \sim iid\ N(0, \sigma_y^2)$$

$$a_0\ given, \quad c \geq 0, \quad a' > -b$$

Agents have to set optimally their consumption/savings plans. They enjoy utility from consumption, and face some uncertain events in the future. In the next period they will still have the same risk aversion parameter, but their labor income can go up or down, depending on the future realizations of the earnings shock η. Finally, they are subject to an exogenous borrowing constraint, $b \geq 0$.

6.2 Problem of the firm

The production side of the model is represented by a constant returns to scale technology of the Cobb-Douglas form, which relies on aggregate capital K and labor L to produce the final output Y.

$$Y = F(K, L) = K^{\alpha} L^{1-\alpha}.$$

Capital depreciates at the exogenous rate δ and firms hire capital and labor every period from competitive markets. From the first order conditions of the firm we obtain the expression for the net real return to capital r and the wage rate per efficiency unit w:

$$r = \alpha \left(\frac{L}{K} \right)^{1-\alpha} - \delta, \tag{2}$$

$$w = (1-\alpha) \left(\frac{K}{L} \right)^{\alpha}. \tag{3}$$

Notice that the marginal productivity of labor is always positive, hence firms will rely on the total sum of the efficiency units of labor. It follows that in the steady-state:

$$L = \int \varepsilon d\mu \left(\varepsilon \right)$$

where $\mu \left(\varepsilon \right)$ is the stationary distribution over the labor endowments implied by the markov process.

6.3 Recursive Stationary Equilibrium

Definition 1 *A recursive stationary equilibrium is a set of decision rules* $\left\{ c(\varepsilon, a), a'(\varepsilon, a), k = \frac{K}{L} \right\}$*, value functions* $V(\varepsilon, a)$*, prices* $\{r, w\}$ *and stationary distributions* $\mu(\varepsilon, a)$ *such that:*

- Given relative prices $\{r, w\}$, the individual policy functions $\{c(\varepsilon, a), a'(\varepsilon, a)\}$ solve the household problem (1) and $V(\varepsilon, a)$ are the associated value functions.
- Given relative prices $\{r, w\}$, k solves the firm's problem (2)-(3).
- The asset market clears
$$K = \int a d\mu(\varepsilon, a)$$
- The goods market clears
$$F(K, L) = \int c(\varepsilon, a; \sigma) d\mu(\varepsilon, a) + \delta K$$
- The stationary distributions $\mu(\varepsilon, a)$ satisfy

$$\mu(\varepsilon', a') = \sum_{\varepsilon} \int_{a: a'(\varepsilon, a) = a'} \pi \left(\varepsilon', \varepsilon \right) d\mu(\varepsilon, a) \tag{4}$$

In equilibrium the measure of agents in each state is time invariant and consistent with individual decisions, as given by the equation (4) above.[12]

[12] Notice that the equation already exploits the Markov Chain representation of the continuous process for ε.

Appendix B - 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 Q6600 Quad Core 2.4 Ghz processor.

- On average, each replication takes 5 minutes. The actual computing time depends essentially on the discount factor β that was drawn, and on how far the initial guess on the interest rate is from the equilibrium one (from 1 to 13 iterations on the interest rate are needed to find each equilibrium). This means that, for any prior distribution, the whole Monte Carlo procedure takes between 8 and 13 days to complete. However, the Quad Core processor allows to run at least four cases simultaneously, with a 100% load, but without losing any computational speed. Obviously, more recent Quad and Hexa Core CPU's could handle even more simultaneous simulations, that would be completed in less time.

- In the actual solution of the model we need to discretize the two continuous state variables ε, a. As for ε, we rely on Tauchen's method, which approximates the $AR(1)$ process for the efficiency units with a Markov chain. We use a seven-state approximation: this is a common number of points, as it strikes a good balance between approximation error and speed. As for a, we rely on an unevenly spaced grid, with the distance between two consecutive points increasing geometrically. In order to keep the computational burden manageable, we use 301 grid points on the asset space, the lowest value being the borrowing constraint and the highest one being a value $a_{\max} > \overline{a}$ high enough for the saving functions to cut the 45 degree line ($a_{\max} = 50$). This is done to allow for a high precision of the policy rules at low values of a, that is where the change in curvature is more pronounced.

- A collocation method is implemented, that is we 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 we need to solve a non linear equation. A time iteration scheme is applied to get the policy functions, i.e. we compute the first order conditions with respect to a' and through the envelope condition we obtain a set of euler equations, whose unknowns are the policy functions, $a'(\varepsilon, a)$. We start from a set of guesses, $a'(\varepsilon, a)_0$, and keep on iterating until a fixed point is reached, i.e. until two successive iterations satisfy:

$$\underset{a}{Sup}\,|a'(\varepsilon,a)_{n+1}-a'(\varepsilon,a)_n|<10^{-8},\forall\varepsilon$$

Between 115 and 1,500 iterations are needed to reach a fixed point.

- The stationary distributions $\mu(\varepsilon, a)$ are computed by simulating a large sample of 10,000 individuals for 3,000 periods, which ensure that the statistics of interest are stationary processes, and that they do not vary substantially when considering more individuals. For more details, see Rios-Rull (1999). This stage is particularly time consuming (approximately 30 seconds per iteration on the interest rate). Notice that in order to minimize the effect of sampling variability affecting our results, we do not rely on the simulation method used by Aiyagari (1994) (the simulation of only one long history). For each statistic of the simulated sample, we consider the time average of their cross sectional values for the last 1,000 periods, rather than the value obtained from one long simulated history. As for the approximation method, we rely on a linear approximation scheme for the saving and consumption functions, for values of a falling outside the grid.

- If the numerical procedure fails to converge in some of his steps, the related results are disregarded. This happens when the time iteration procedure gets stuck in a cycle, rather than converging to the true policy functions.

Appendix C - Monte Carlo Algorithm

The computational procedure used to solve the Monte Carlo experiments can be represented by the following algorithm:

1. Draw 2,500 combinations of parameters from their prior distributions and store them.
2. Calibrate the model by reading the first vector of simulated parameters and begin the model solution.
3. Generate a discrete grid over the asset space $[-b, ..., a_{\max}]$.
4. Generate a discrete grid over the efficiency units space $[\varepsilon_{\min}, ..., \varepsilon_{\max}]$.
5. Get the aggregate labor supply L.
6. Guess on the interest rate r_0.
7. Get the capital demand k.
8. Get the wage rate per efficiency units w.
9. Get the saving functions $a'(\varepsilon, a)$.
10. Simulate the stationary distributions $\mu(\varepsilon, a)$.
11. Get the aggregate capital supply.
12. Check asset market clearing; Get r_1.
13. Update $r_0' = \varpi r_0 + (1 - \varpi) r_1$ (with bisection).
14. Iterate until market clearing.
15. Get the consumption functions $c'(\varepsilon, a)$.
16. Check final good market clearing.
17. Save the output and repeat from step 2 for all the 2,500 combinations of simulated parameters.

Appendix D - Additional Figures: Baseline and Restricted Experiments

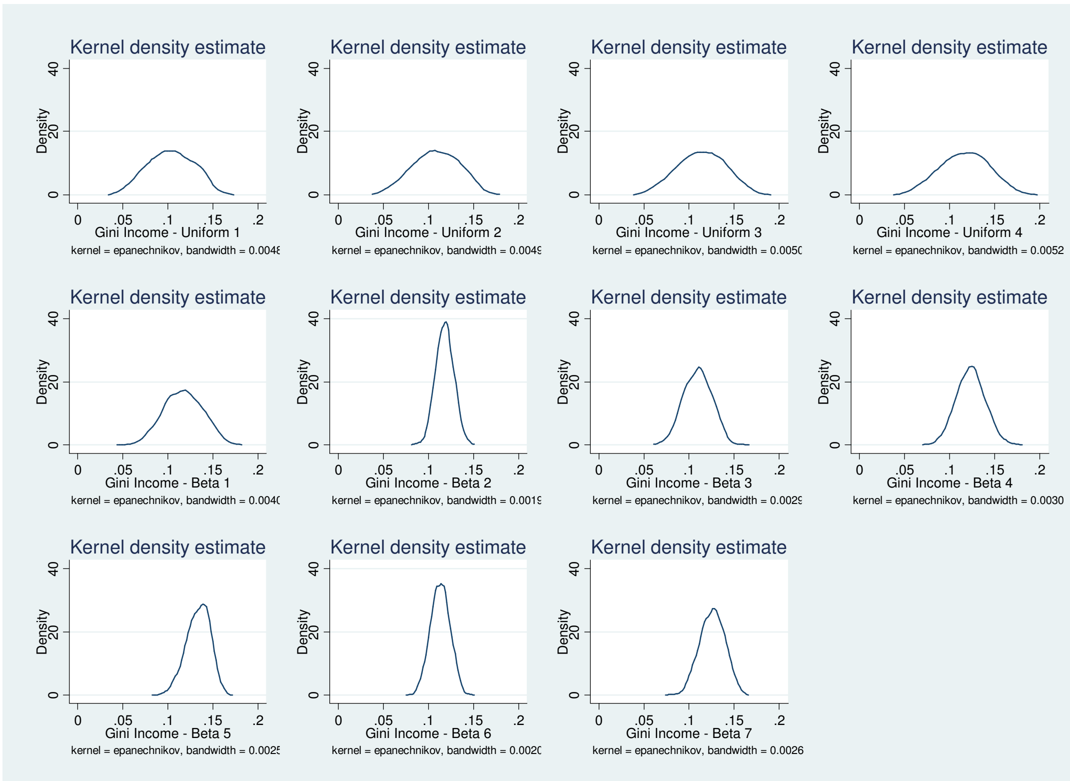

Figure 13: Income Gini Index - Estimated Densities for the eleven Monte Carlo Experiments

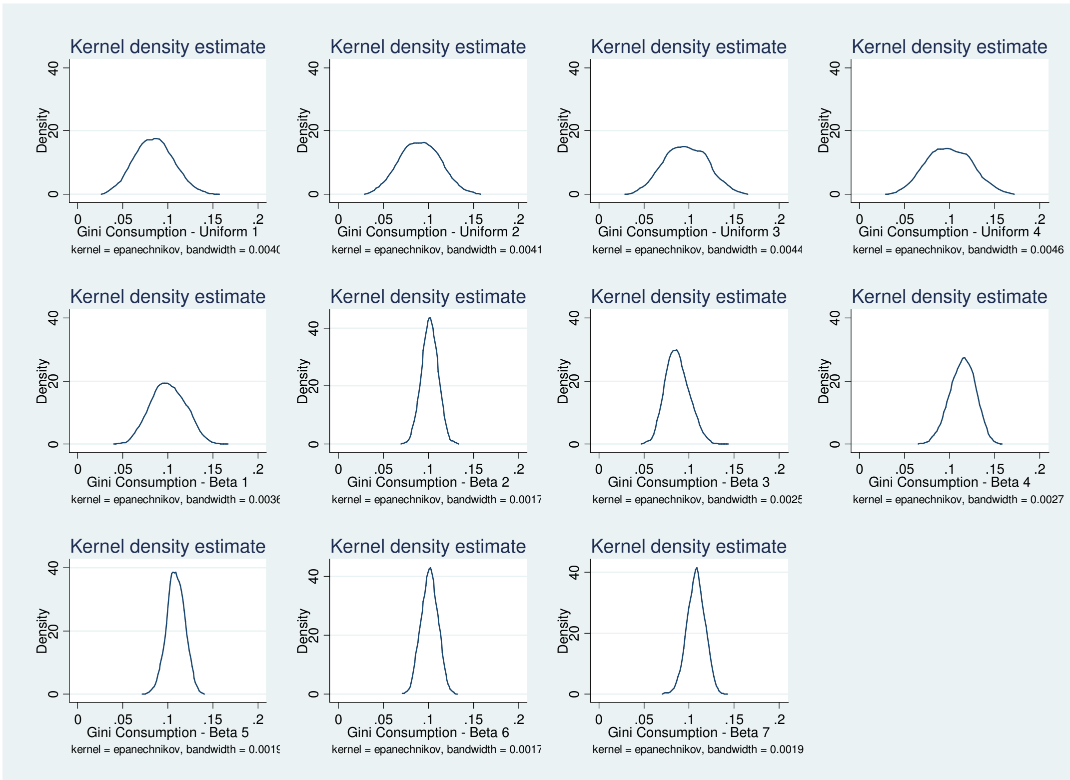

Figure 14: Consumption Gini Index - Estimated Densities for the eleven Monte Carlo Experiments

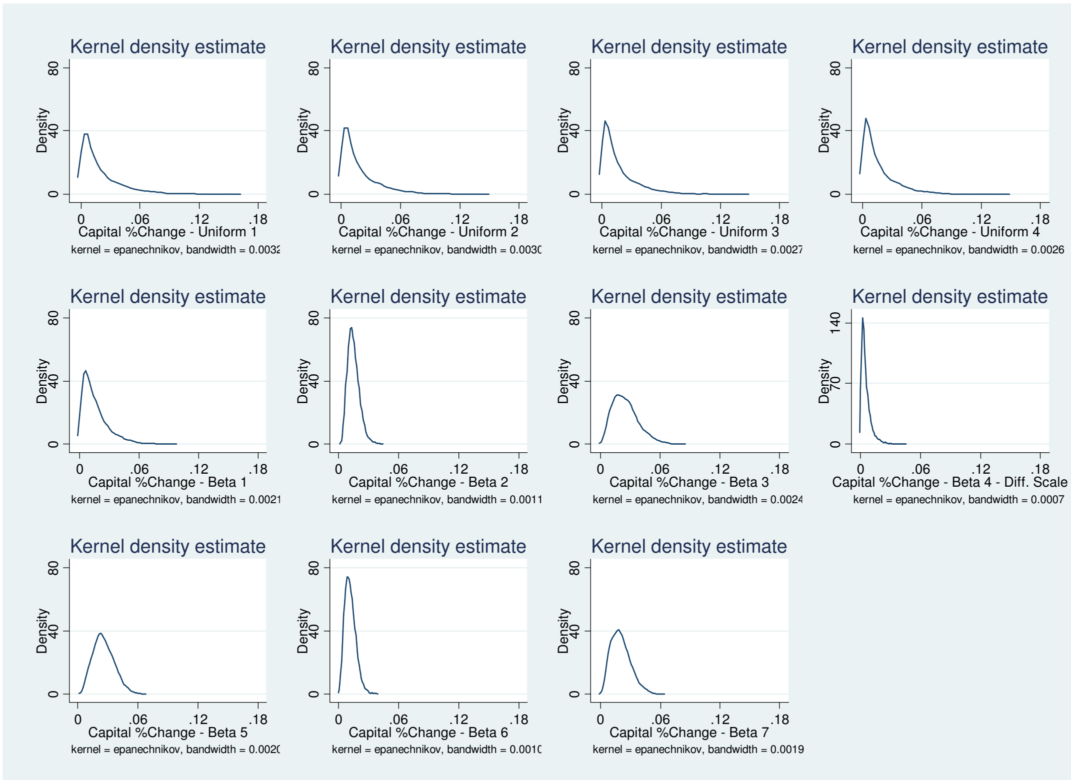

Figure 15: Precautionary Savings - Estimated Densities for the eleven Monte Carlo Experiments

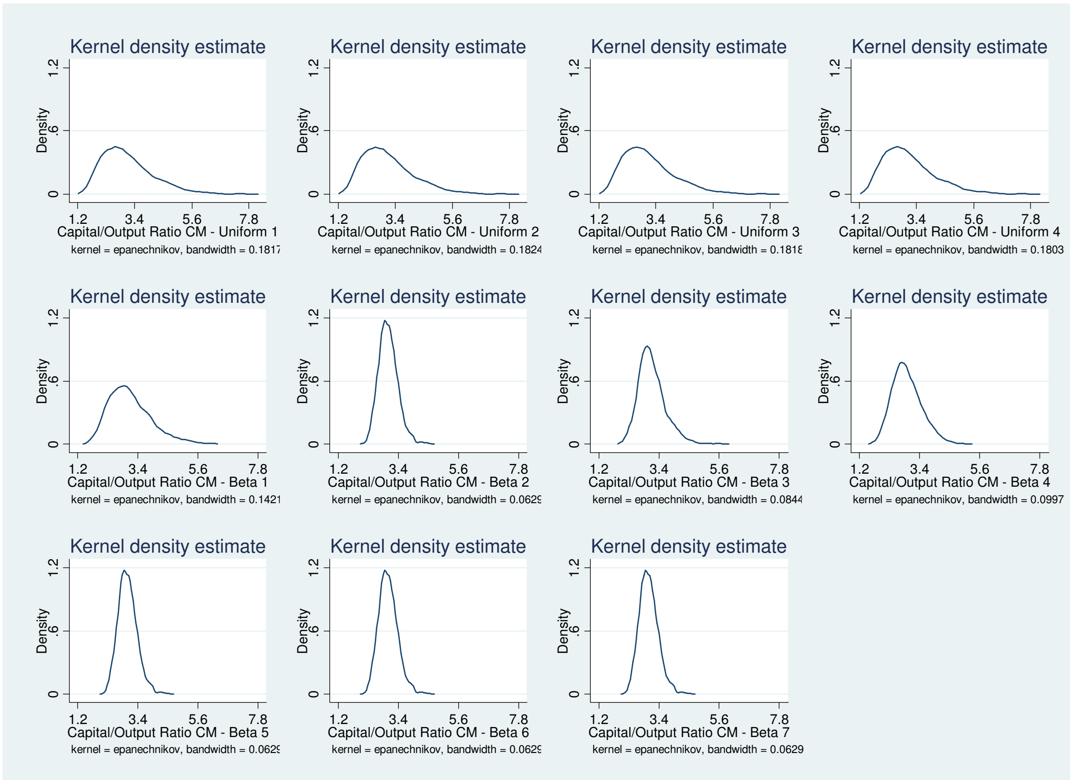

Figure 16: Capital/Output Ratios (CM) - Estimated Densities for the eleven Monte Carlo Experiments

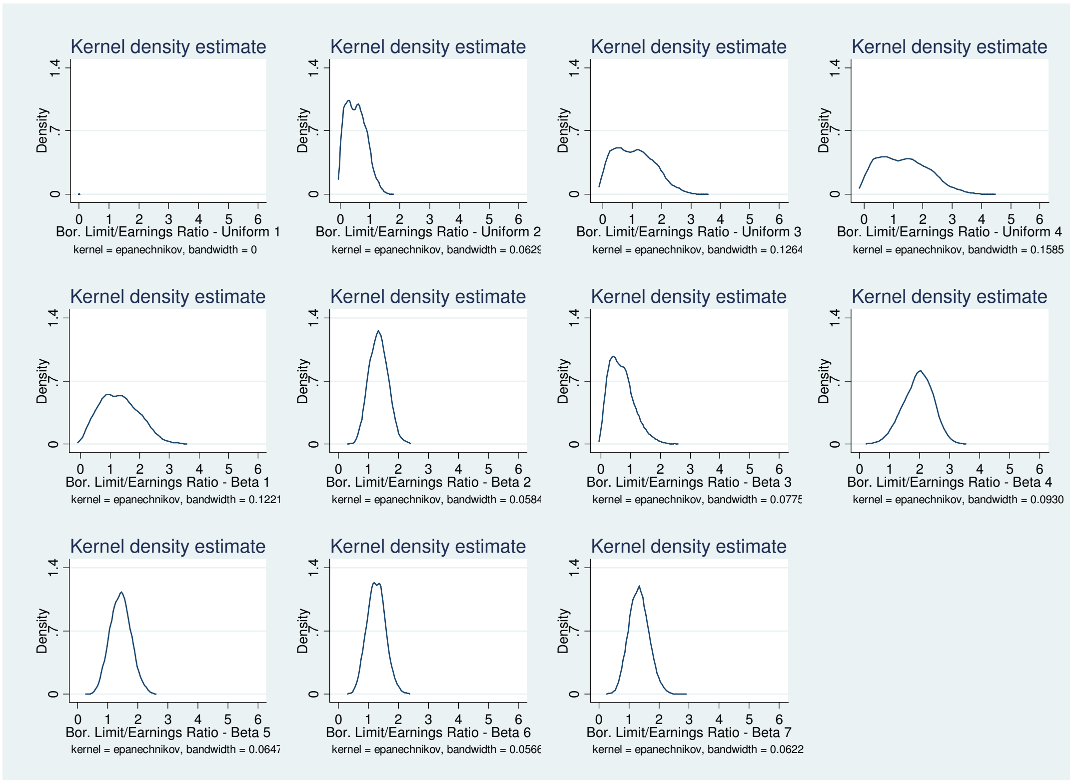

Figure 17: Borrowing Limit/Labor Earnings Ratios - Estimated Densities for the eleven Monte Carlo Experiments

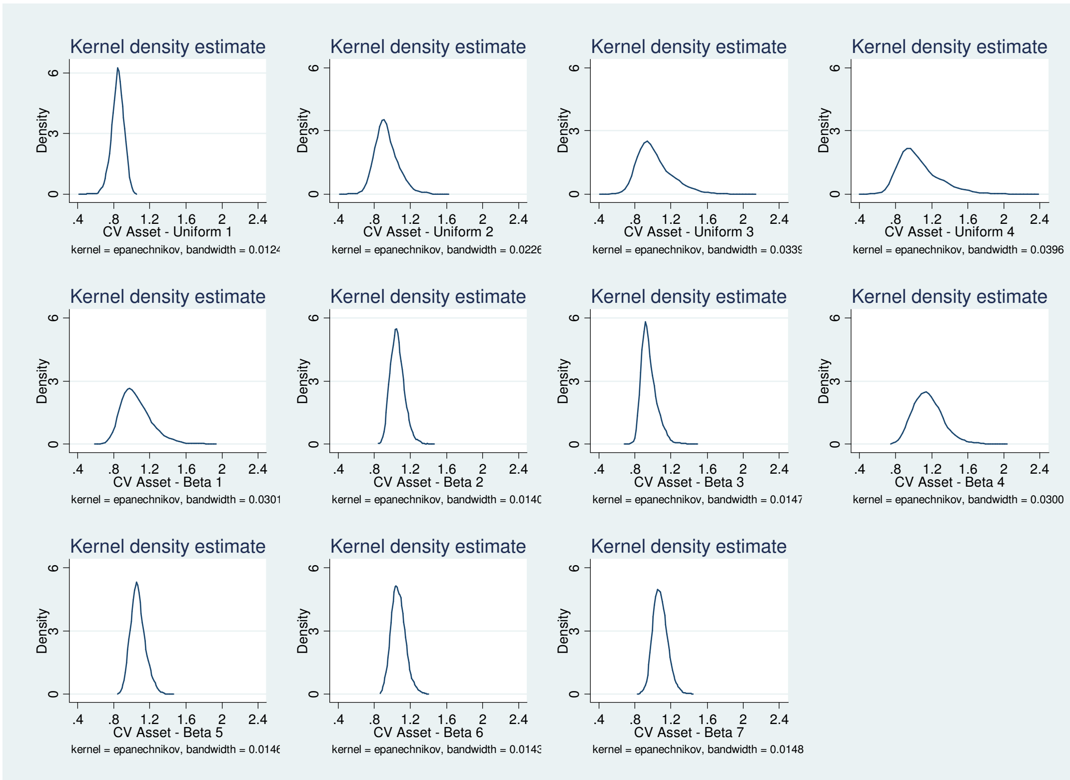

Figure 18: Wealth Coefficient of Variation - Estimated Densities for the eleven Monte Carlo Experiments

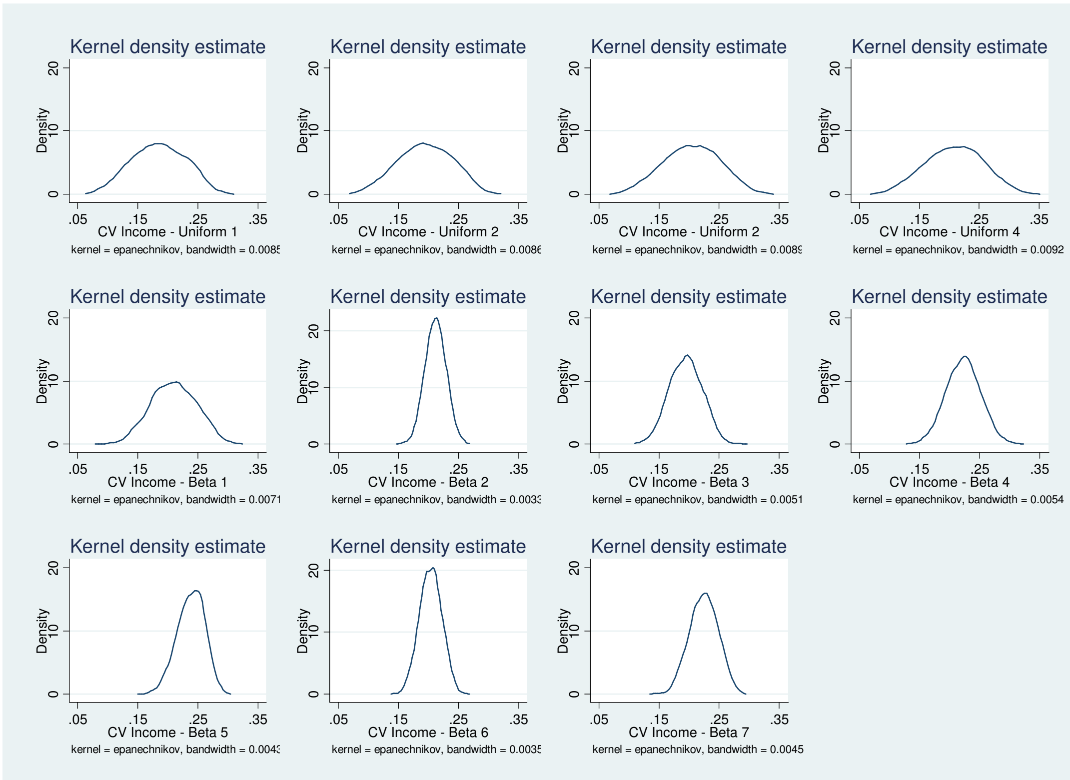

Figure 19: Income Coefficient of Variation - Estimated Densities for the eleven Monte Carlo Experiments

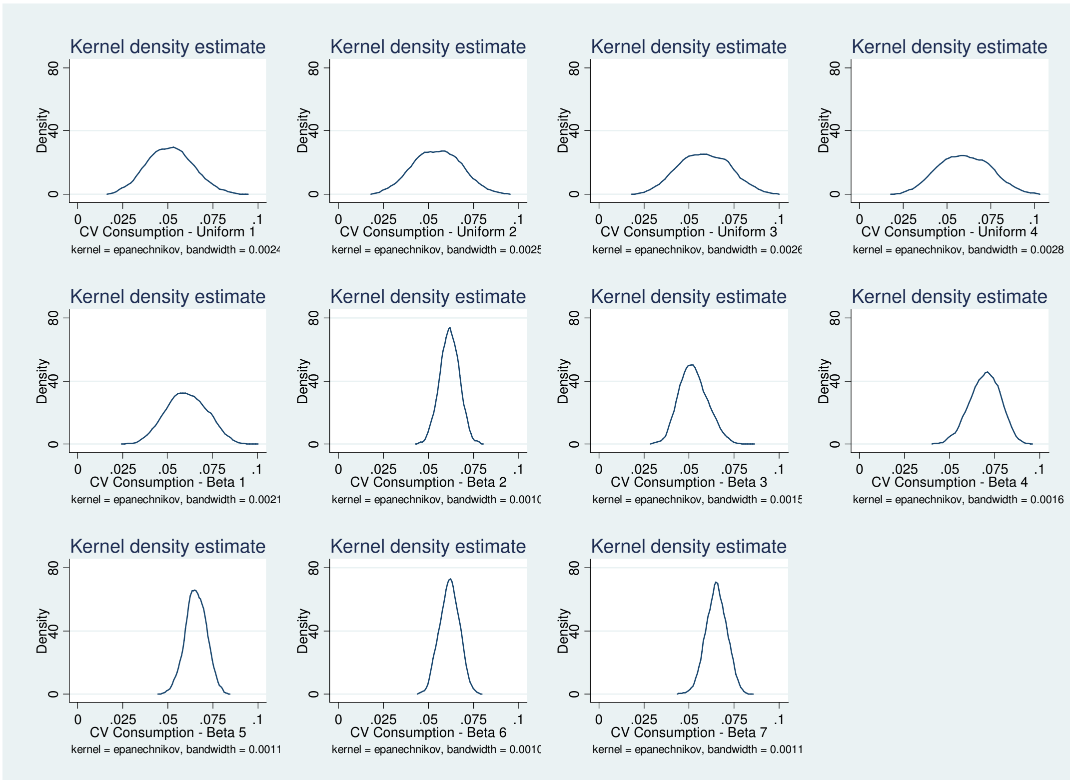

Figure 20: Consumption Coefficient of Variation - Estimated Densities for the eleven Monte Carlo Experiments

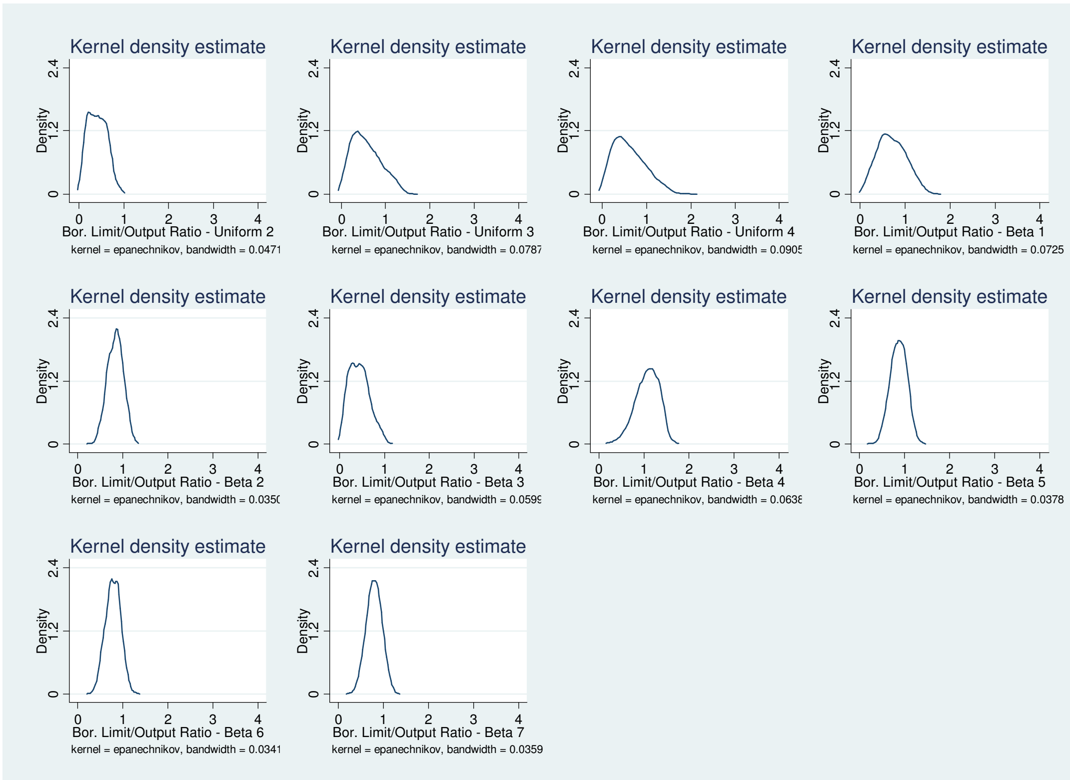

Figure 21: Borrowing Limit/Ouput Ratios - Estimated Densities for the eleven Restricted Monte Carlo Experiments

Appendix E - Additional Figures: $\underline{\rho_y} = 0.5$ Case

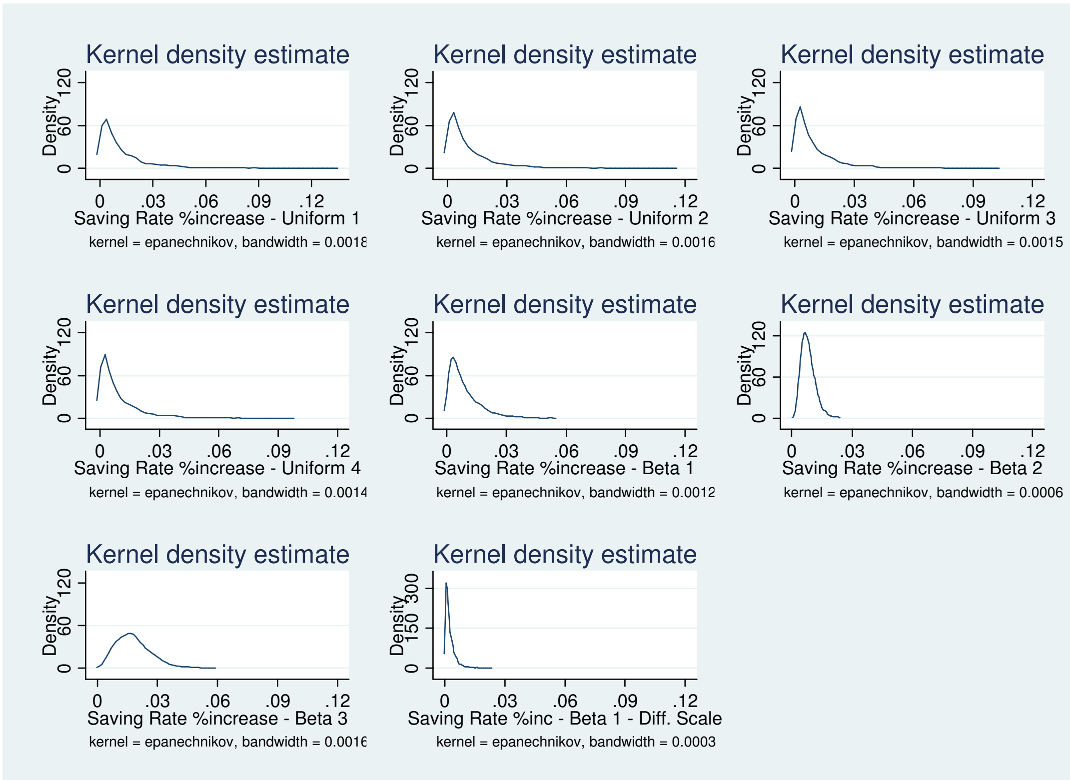

Figure 22: Precautionary Savings - Estimated Densities for the eight Monte Carlo Experiments

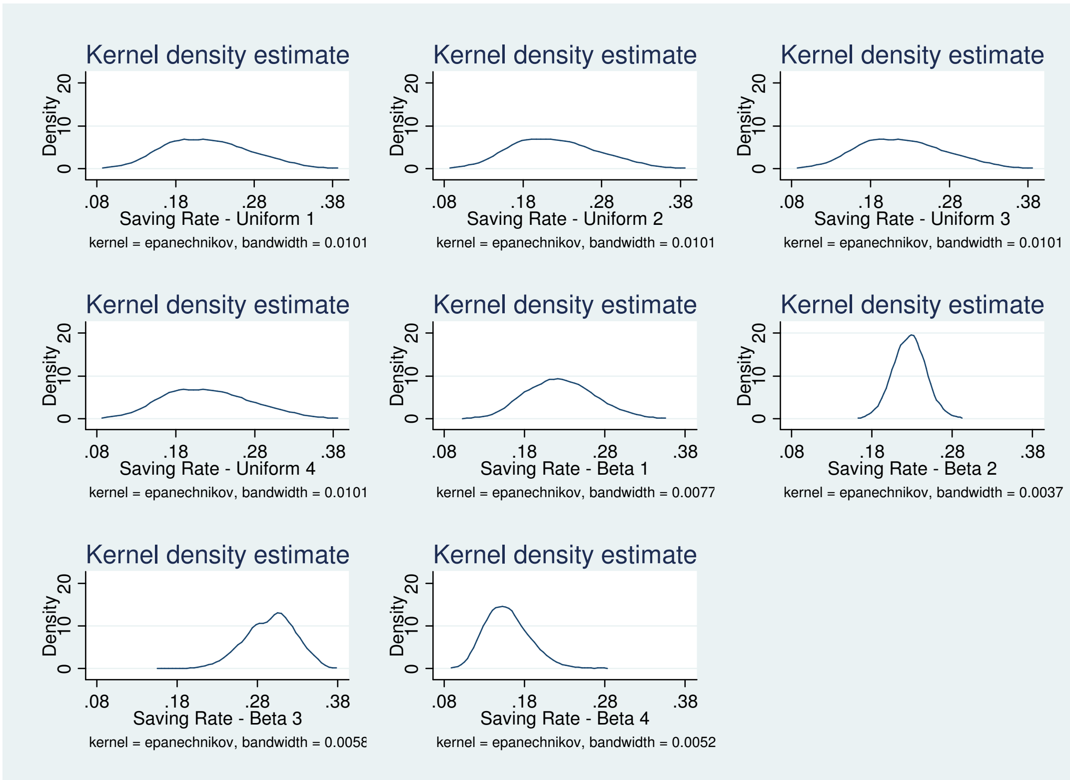

Figure 23: Saving Rates - Estimated Densities for the eight Monte Carlo Experiments

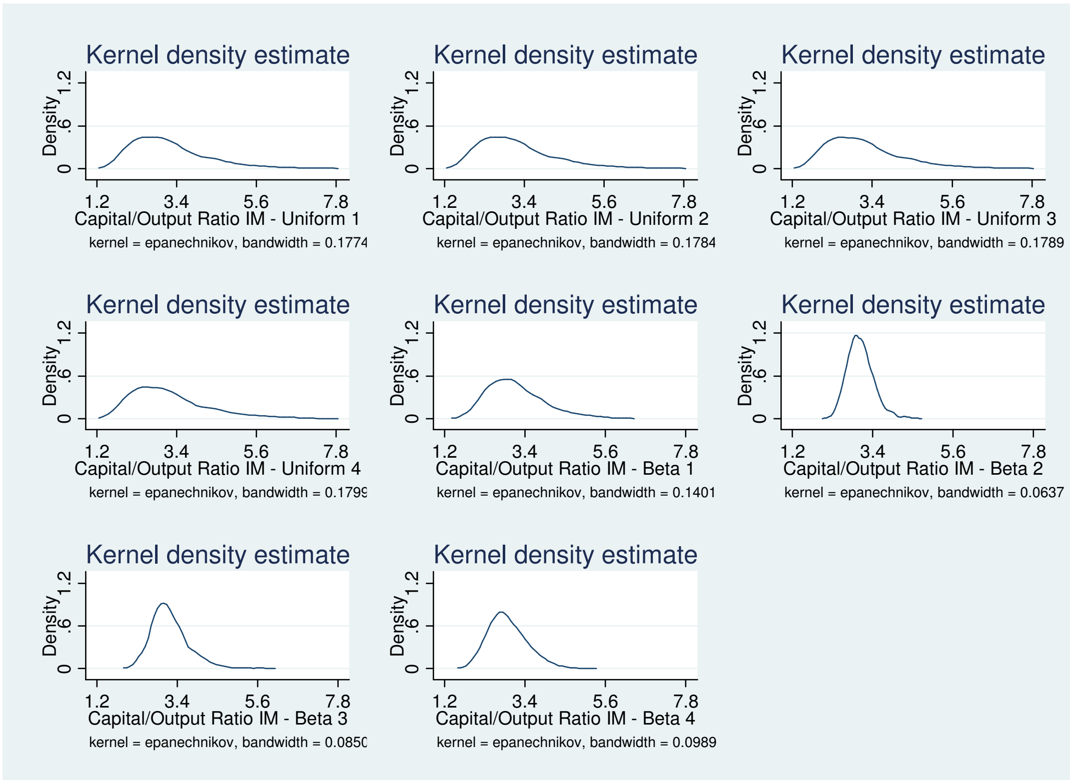

Figure 24: Capital/Output Ratios (IM) - Estimated Densities for the eight Monte Carlo Experiments

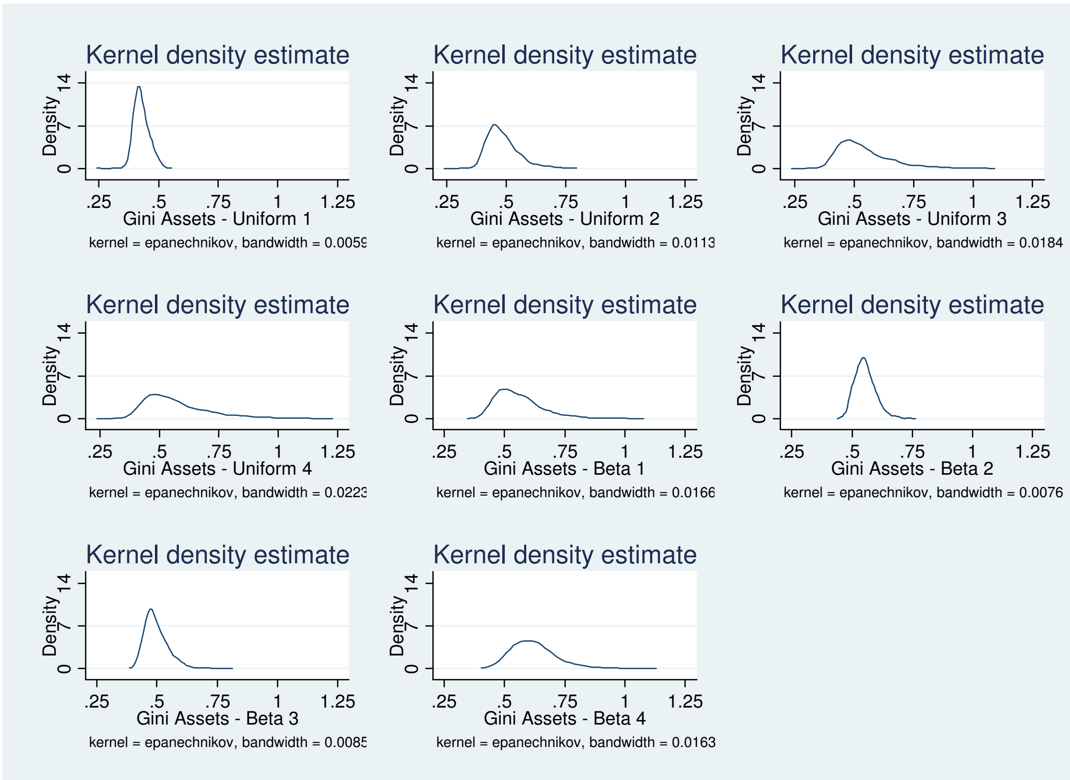

Figure 25: Wealth Gini Index - Estimated Densities for the eight Monte Carlo Experiments

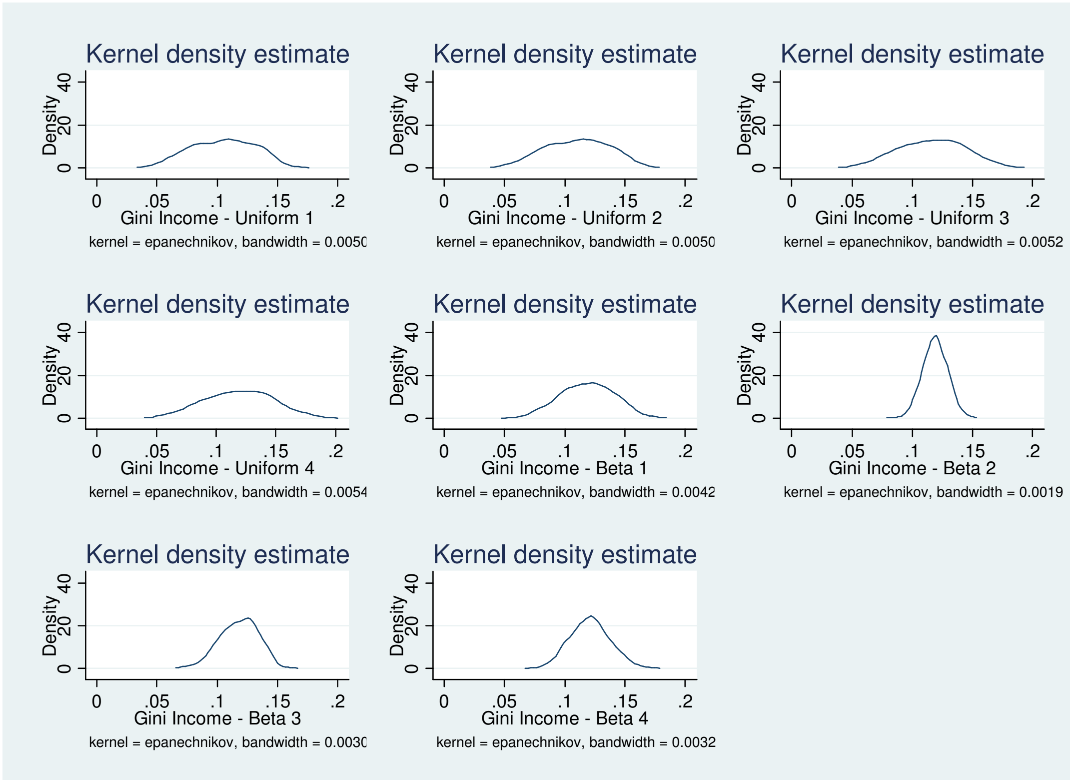

Figure 26: Income Gini Index - Estimated Densities for the eight Monte Carlo Experiments

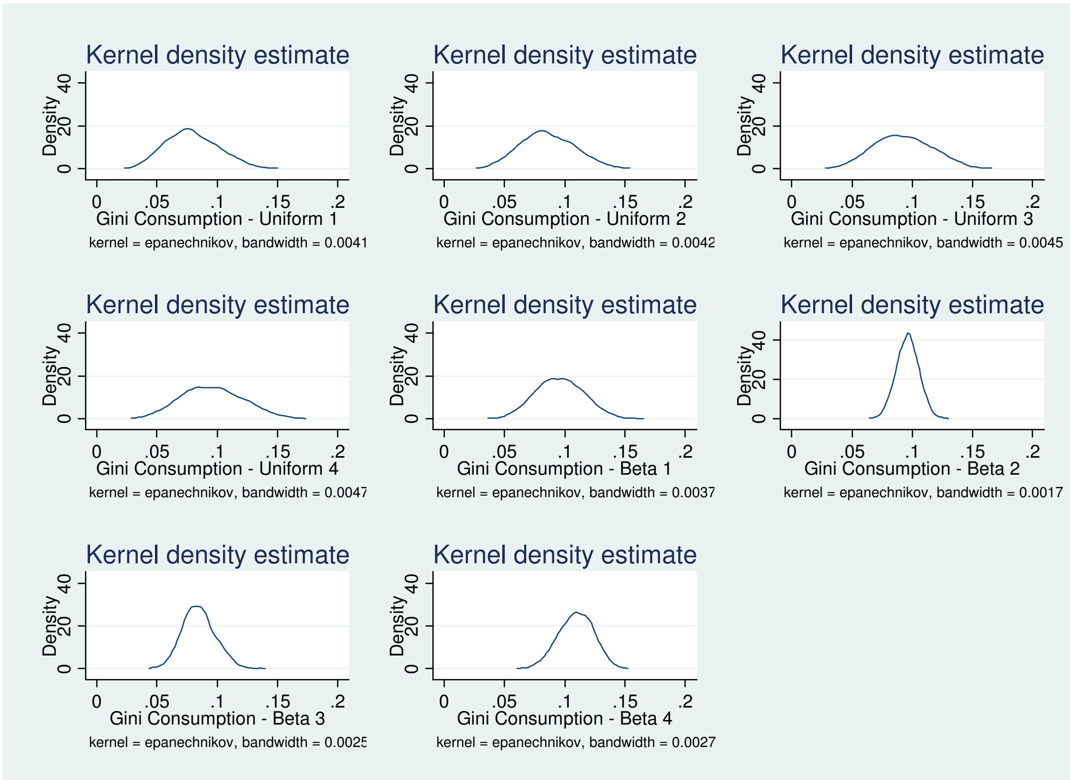

Figure 27: Consumption Gini Index - Estimated Densities for the eight Monte Carlo Experiments

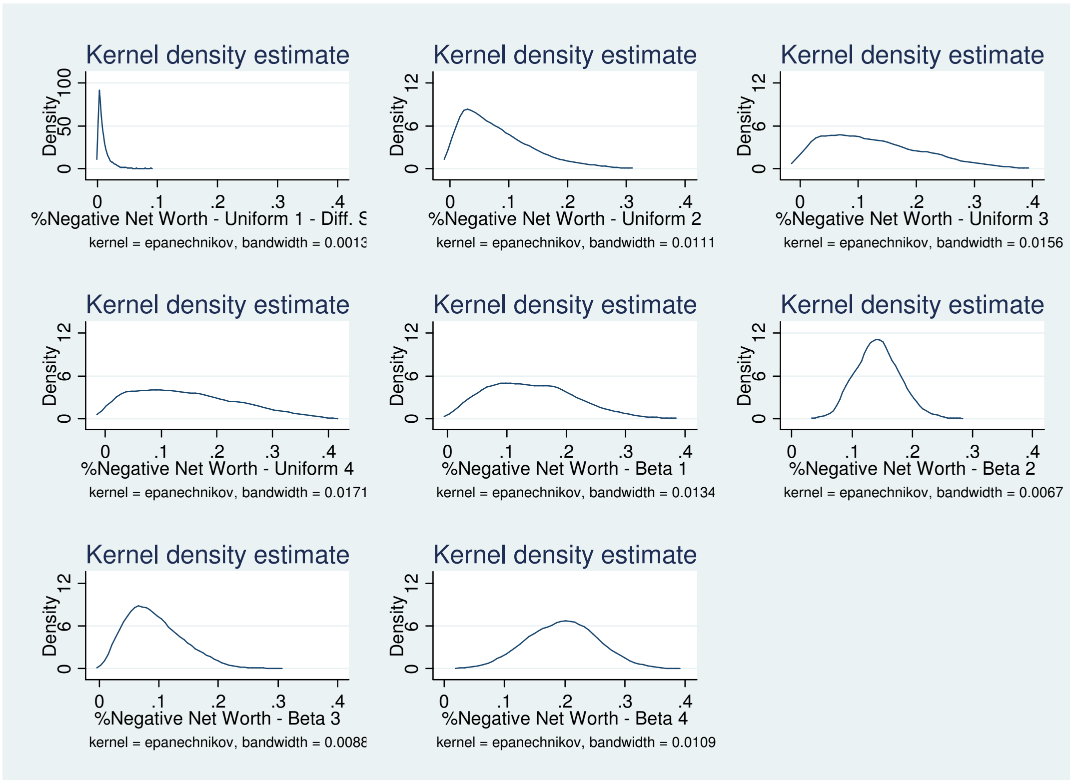

Figure 28: % with Negative Net Worth - Estimated Densities for the eight Monte Carlo Experiments

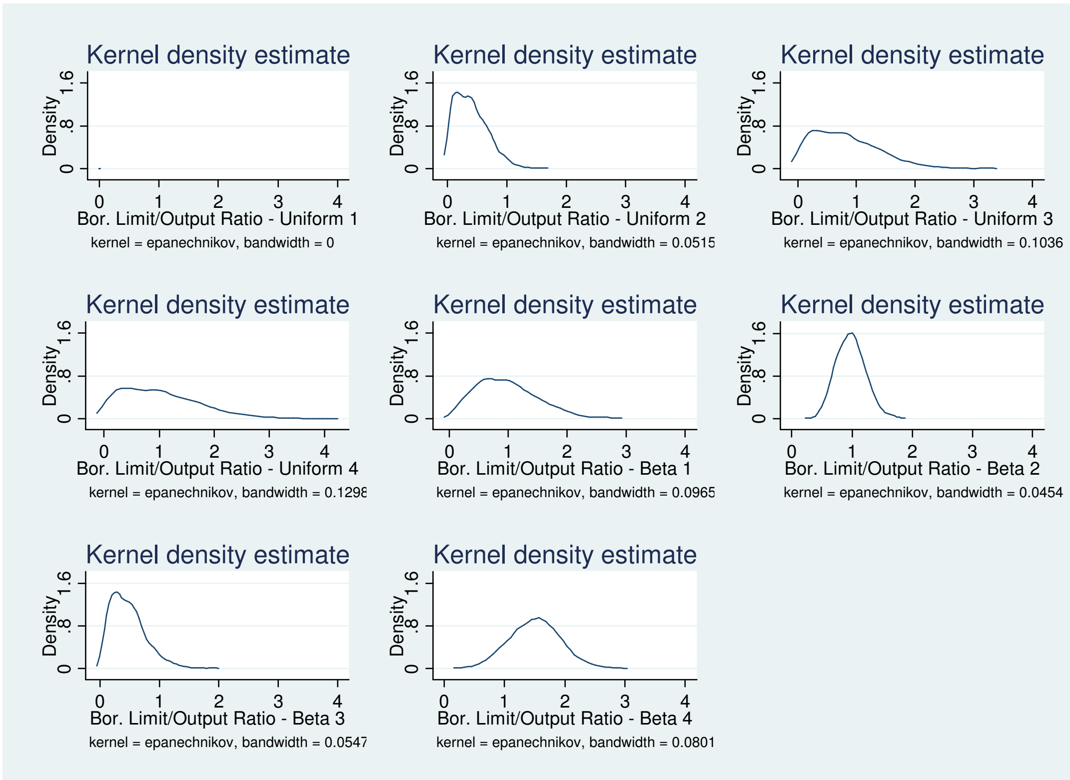

Figure 29: Borrowing Limit/Ouput Ratios - Estimated Densities for the eight Monte Carlo Experiments

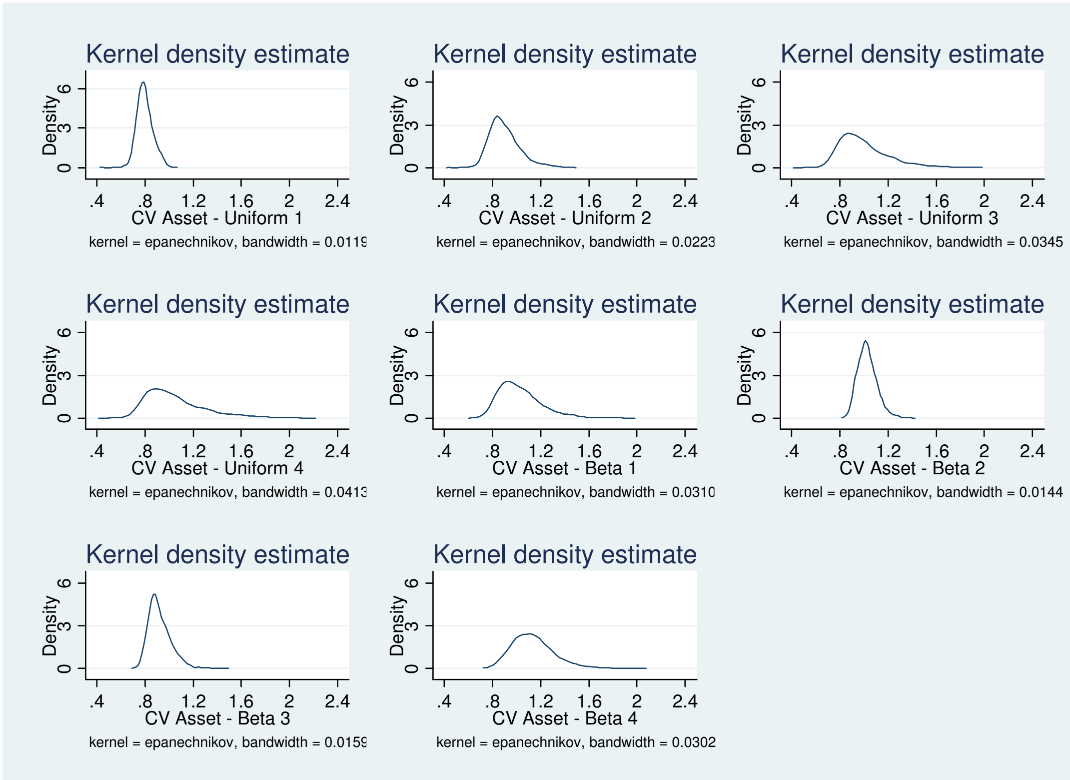

Figure 30: Wealth Coefficient of Variation - Estimated Densities for the eight Monte Carlo Experiments

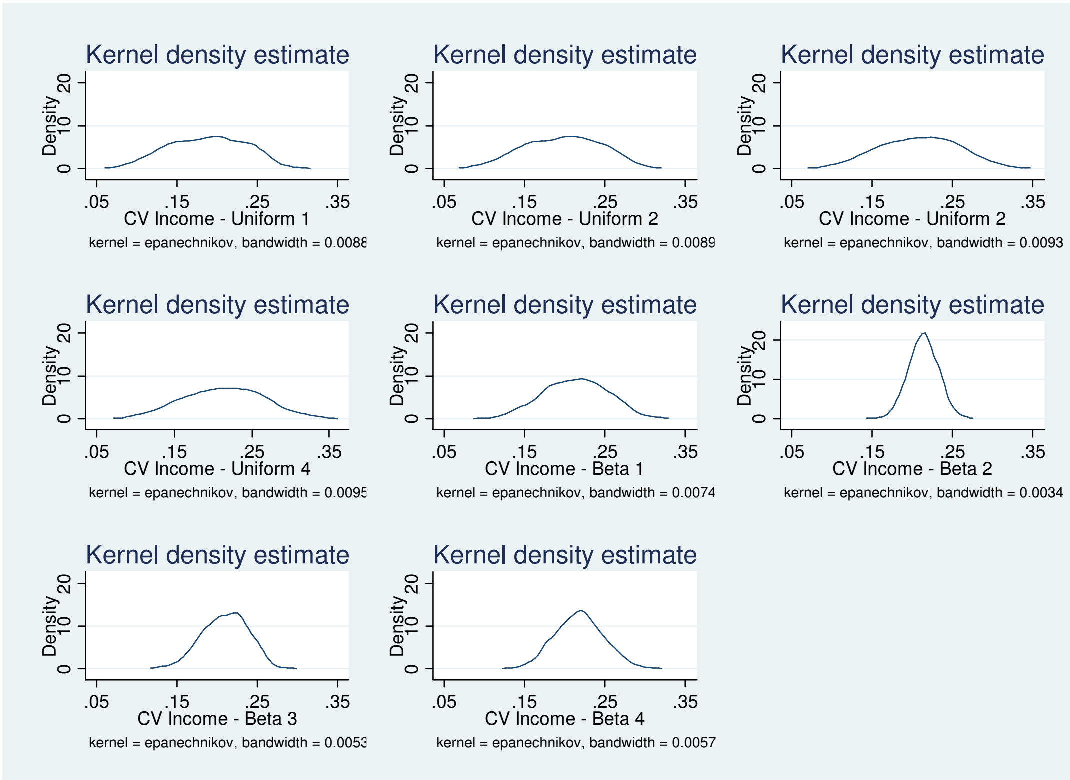

Figure 31: Income Coefficient of Variation - Estimated Densities for the eight Monte Carlo Experiments

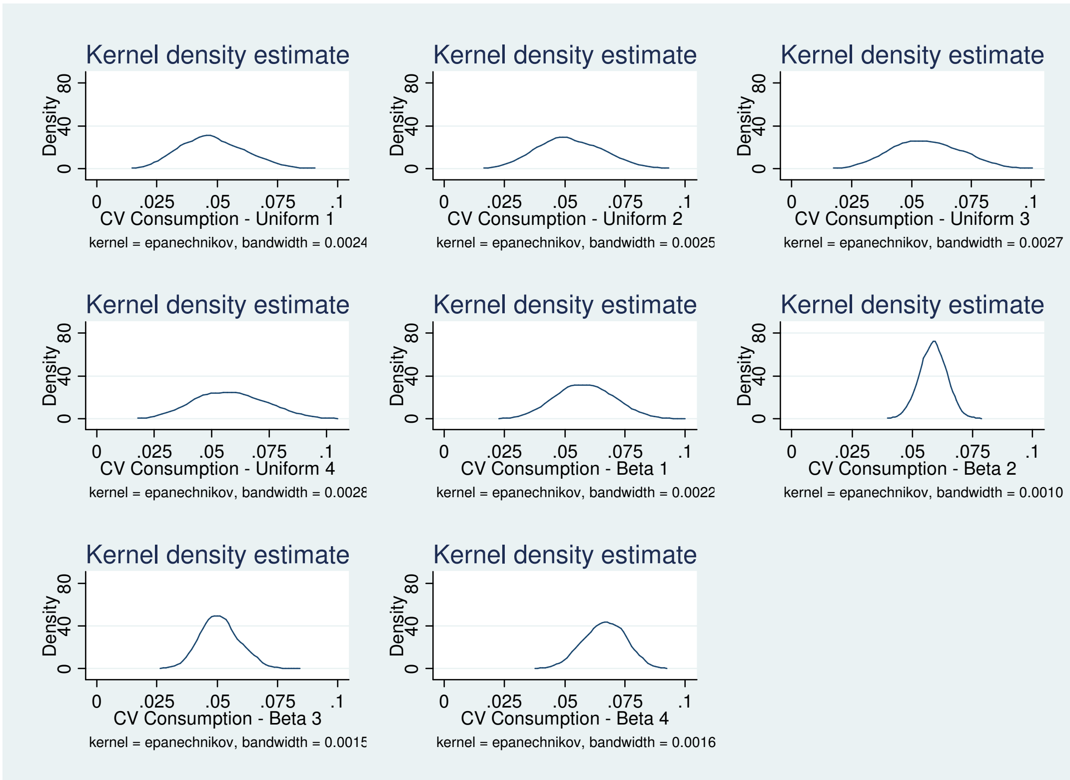

Figure 32: Consumption Coefficient of Variation - Estimated Densities for the eight Monte Carlo Experiments