Rochester Center for
Economic Research

Income and Wealth Heterogeneity in the Macroeconomy

Krusell, Per and Anthony A. Smith, Jr.

Working Paper No. 399
March 1995

University of
Rochester

Income and Wealth Heterogeneity in the Macroeconomy

Per Krusell and Anthony A. Smith Jr.

Rochester Center for Economic Research
Working Paper No. 399

March 1995

Income and Wealth Heterogeneity in the Macroeconomy

Per Krusell*
University of Rochester
University of Pennsylvania

Anthony A. Smith, Jr.*
Carnegie Mellon University

March 1995
First version: September 1993

Abstract

How do movements in the distribution of income and wealth affect the macroeconomy? We analyze this question theoretically, using numerical methods, in the context of a calibrated version of the stochastic growth model with partially uninsurable idiosyncratic risk and movements in aggregate productivity. Our main finding is that, in the stationary stochastic equilibrium, the behavior of all relevant aggregates can be almost perfectly described using only the mean of the wealth distribution. This result is robust to substantial changes in both parameter values and model specification. A separate finding is that the equilibria of our baseline economy and of some extensions of it have aggregate time-series statistics which are very similar to those of the corresponding representative-agent economies (i.e. the complete-markets economies with no wealth differences across agents). However, when agents also differ in preferences the economies with complete and incomplete markets begin to look more dissimilar.

*We thank conference participants at: Macroeconomics with "Frictions", held at the Federal Reserve Bank of Minneapolis, the NBER Summer Workshop on the Aggregate Implications of Microeconomic Consumption Behavior, the Northwestern Summer Workshop on Applied General Equilibrium Analysis, the Annual Meetings of the Society for Economic Dynamics and Control, and the North American Summer Meetings of the Econometric Society; and seminar participants at Carnegie Mellon, Columbia, Institute for International Economic Studies, Pompeu Fabra, SMU, Texas at Austin, Wharton, Yale, and York. We also benefitted particularly from comments by Wouter den Haan, Mark Huggett, Robert Lucas, Víctor Ríos, Tom Sargent, and José Scheinkman. Financial support from the Bank of Sweden Tercentenary Foundation, the National Science Foundation, and the Social Sciences and Humanities Research Council of Canada is gratefully acknowledged.

1 Introduction

Modern macroeconomic theory, which we take to mean dynamic general equilibrium theory, draws heavily on the representative-agent abstraction: a dominant assumption is that the economy behaves "as if" it is inhabited by a single consumer, or by a single type of consumer. At the same time, modern macroeconomics attempts to take seriously the microfoundations of macroeconomic models, selecting parameter values on the basis of our existing empirical and theoretical knowledge and then using the models to generate quantitative statements. At first glance, the representative-agent assumption appears to be inconsistent with a serious treatment of microfoundations. There are two circumstances, however, under which the representative-agent construct seems reasonable. First, it is possible that the theoretical assumptions needed to justify the use of an aggregate consumer are roughly met in the data. This, however, seems hard to argue; one problem is that it is difficult to justify the assumption that there are complete insurance markets for consumers' idiosyncratic risks. A second possibility is that the aggregate variables in theoretical models with a more realistic description of the microeconomic environment actually behave like those in the representative-agent models. In this paper we begin the exploration of this second possibility.

The goal of the present paper is to extend the standard macroeconomic model to include substantial heterogeneity in income and wealth. More precisely, we consider a calibrated version of the stochastic growth model where there is a large number of consumers. In addition to uncertain aggregate productivity, agents face idiosyncratic income shocks—interpreted as employment shocks—which are only partially insurable: capital is available as a buffer stock, but there is no other insurance. The absence of full insurance implies that the distributions of income and wealth are nontrivially determined.[1] We characterize stationary stochastic equilibria of this model numerically, and we then compare the aggregate properties of these equilibria with those implied by the corresponding representative-agent model. The characterization of the stochastic behavior of the income and wealth distributions is central to our task, since aggregate variables depend on these distributions. An important component of our analysis involves dealing with the main computational difficulty of dynamic heterogeneous-agent models: in order to predict prices, consumers need to keep track of the evolution of the wealth distribution. With a continuum of agents, this distribution is an infinite-dimensional object.

We obtain two key results. First, we are able to compute an equilibrium which satisfies our numerical accuracy requirements. In this equilibrium, all aggregate variables—consumption, the capital stock, and relative prices—can be almost perfectly described as a function of two simple statistics: the mean of the wealth distribution, and the aggregate productivity shock. In other words, the consumers in our equilibrium face maneageable prediction problems, since the distri-

[1] The complete-markets version of our baseline model allows exact aggregation for the class of preferences that we consider.

bution of aggregate wealth is almost completely irrelevant for how the aggregates behave in the equilibrium. The finding that only the mean of capital matters for describing the evolution of total wealth in this economy is not an exact result, since there is no aggregation theorem for this environment. However, the finding is remarkably robust to changes in both parameter values and the specification of the model. The alternative specifications that we examine include models with valued leisure and with nontrivial heterogeneity in preferences among agents. Our findings thus show that it is feasible using current computational technology to analyze an interesting class of dynamic equilibrium models with substantial agent heterogeneity.

Second, we find that for many but not all of the variations of the baseline model that we consider, the aggregate variables behave almost as they do in the corresponding representative-agent model. For example, the "real-business-cycle" model with heterogeneity in income and wealth is very difficult to distinguish from its representative-agent counterpart (i.e. the version where idiosyncratic risks are fully insurable) by looking at aggregate variables. In particular, the second moments of these data are very similar across the two setups. The main difference between the two setups is that, due to precautionary savings, the capital stock is somewhat higher in the heterogeneous-agent version of the model.

In contrast, the models with heterogeneity in preferences which we examine behave differently than their representative-agents counterparts along some dimensions (here, there would be one representative agent for each kind of preferences). In particular, the differences are large for some correlations between macroeconomic aggregates, such as the consumption-output correlation. The models with heterogeneity in preferences are also interesting because they are capable of generating wealth distributions which do a better job of matching certain features of observed wealth distributions—such as large fractions of agents with low levels of wealth and thick right tails—than do models without heterogeneity in preferences.[2]

The intuitive explanation for our main result—the collapse of the state space—to a large part derives from the properties of optimal consumption behavior in stationary equilibria. Although marginal propensities to save do differ across agents with different income and wealth levels, they are *almost* the same for all but the agents with the very lowest levels of asset holdings. Moreover, since low asset holdings lead to low levels of utility, and since agents in the economy use capital actively to self-insure against this kind of outcome, it follows that in equilibrium almost all agents have almost the same marginal propensities to save. Nonetheless, although the differences in marginal propensities across agents are small, they are large enough that there exist wealth redistributions which would lead to significant changes in total savings. It turns out, however, that in a stationary stochastic equilibrium such wealth redistributions occur very infrequently: in effect, by focusing on stationary equilibria, we are restricting the set of wealth distributions. In sum, we find that changes

[2]For a discussion of the key features of wealth distributions in the United States, see Huggett (1994).

in the way the total capital stock is distributed across agents with different marginal propensities to consume account for almost no movements at all in total savings.

Some recent studies use dynamic equilibrium theory to analyze the importance of consumer heterogeneity for asset prices (see den Haan (1994a, 1994b), Heaton and Lucas (1994a,1994b), Lucas (1994), Marcet and Singleton (1991), and Telmer (1993)). Some of these models display a significant dependence of equilibrium asset prices on higher-order moments of the asset distribution. A key distinguishing feature between our model and these other frameworks is the existence of capital. When capital accumulation is permitted, the economy as a whole can work its way toward a stationary region where the incomplete markets and borrowing constraints are not constraining for most agents. When capital accumulation is not permitted, one agent's gain is another agent's loss: any increase in a given agent's wealth must be balanced by decreases in the wealth of other agents since the sum of all the asset holdings is constant. Thus, changes in the distribution of assets inevitably force agents toward asset holdings where the borrowing constraints begin to influence behavior significantly, unless the total sum of assets is a large amount and the borrowing constraint is generous. Finally, it should also be said that our findings are related to the theoretical results in Bewley (1977), who studies a decision-theoretic setup which under certain conditions leads to permanent-income type behavior similar to that which we observe in our equilibrium model. We discuss the connection with Bewley's results in Section 3 below.

We examine the robustness of our results in the following ways: we vary the parameters of the baseline model within a quantitatively reasonable range, and we look at models with valued leisure, models with two types of agents whose discount factors or degrees of risk aversion differ, and models with fixed costs associated with capital accumulation. Although some of the model parameters are important for the results—we find, for example, that stronger discounting makes our results weaker—these parameters must be changed substantially from empirically reasonable values in order to alter our main conclusions.

Perhaps the most interesting extensions we look at are the ones with heterogeneity in preferences, since this heterogeneity provides one way of generating highly skewed wealth distributions. For example, in a stationary equilibrium in which half of the agents have a low and the other half a high discount factor, the distribution of wealth is quite skewed, with the group of patient agents owning almost all the capital. Between the two groups and, particularly, within the group of impatient agents, there are more pronounced differences in propensities to save than in the baseline economy. However, since almost all the aggregate capital is held by a group of agents among whom propensities to save differ very little—these agents are not constrained by asset market frictions—the evolution of aggregate capital can be very well understood without looking at its distribution. Very similar results also apply for a model where there are two types of agents whose preferences have different degrees of risk aversion.

The permanent-income hypothesis is at the core of our infinitely-lived-agent model economy, and one important question is whether life-cycle-savings motives would produce different results. Assuming, say, that old agents have significantly lower propensities to save than young agents, it seems clear that the demographic structure will be important for understanding the propagation of aggregate shocks. The extent to which the wealth distribution influences aggregate variables in such a setup once one controls for the demographic state variables is an open question. Relatedly, Ríos-Rull (1992) looks at a yearly life-cycle economy with a demographic structure which has been calibrated to conform with real-world life spans, life-expectancy rates, and age-earnings profiles. He finds that the fluctuations in that framework are very similar to those coming out of a representative-agent model.

We also briefly describe our computational algorithm, since it is essential for understanding how our approximate equilibrium differs from an exact equilibrium. The main computational task is to calculate the law of motion for the distribution of capital over individuals. Our approach is to calculate equilibria where, by assumption, agents have a limited ability to calculate this law of motion, and then to show that the bound on ability almost does not constrain the agents at all. More precisely, we compute approximate equilibria by postulating that the law of motion can be described by a stochastic process for a finite-dimensional vector of moments of the capital distribution. For any given vector of moments m, an approximate equilibrium is calculated as a fixed point in a class $\mathcal{S}$ of (possibly nonlinear) first-order Markov processes. A given Markov process $S \in \mathcal{S}$ is a fixed point if: (1) S is, in a statistical sense, the best approximation in the class $\mathcal{S}$ to the dynamic behavior of m when the moments comprising m are calculated from simulations using the aggregated decision rules of agents; (2) agents' decision rules derive from dynamic maximization problems in which the behavior of the aggregate state is described by S.

In other words, the calculated object satisfies all the standard equilibrium conditions except the agents' ability to make perfect forecasts. One way to assess how much this ability is constrained is to measure how well individual agents can forecast future prices using their law of motion S. We use a measure of forecasting accuracy in the algorithm to decide whether or not to increase the agents' ability to make forecasts—increase the dimension of m or expand the class $\mathcal{S}$—or to stop. In these terms, our main finding is that, when $\mathcal{S}$ is the class of linear first-order Markov processes and m consists only of the mean of the distribution of capital, we obtain extremely high forecasting accuracy. Indeed, the accuracy is so high that we find it very hard to argue on the basis of the "irrationality" of the agents in our model that our approximate equilibrium is a less satisfactory economic model than an exact equilibrium.

The organization of the paper is straightforward: in Section 2 we describe the baseline model and the computational strategy, and in Section 3 we describe the computed heterogeneous-agent equilibria. Section 3 includes analyses of the baseline model (Section 3.1), the economy with preference heterogeneity (Section 3.2), and the real-business-cycle model (Section 3.3). Section 4 then compares the heterogeneous-agent models analyzed with their representative-agent counterparts.

Section 5 concludes with some remarks.

2 Model Framework

In this section we describe our model economy. The key source of heterogeneity is an assumption that idiosyncratic income shocks are partially uninsurable. In our baseline setup there is only one type of consumer, i.e. all consumers have the same preferences, and the setup is one which in all other ways is like the standard stochastic growth model. We also consider a number of perturbations on the baseline setup, including diversity of preferences, valued leisure, and fixed costs associated with the accumulation of assets.

2.1 The Environment

We consider a version of the stochastic growth model with a large (measure 1) population of infinitely-lived consumers. There is only one consumption good per period and in our baseline setup we assume that all agents have the same preferences over streams of consumption:

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t)$$

with

$$u(c) = \lim_{\nu \to \sigma} \frac{c^{1-\nu} - 1}{1 - \nu}.$$

Each agent is endowed with one unit of time. This unit of time has stochastic productivity as labor input, ϵ: it can take on either the value 0 or 1. When $\epsilon = 1$, we think of the agent as employed, and when $\epsilon = 0$ we think of him as unemployed. There is also a stochastic shock to aggregate productivity, which we denote z. As for the idiosyncratic uncertainty, we consider two possible states: either the state is good, and $z = z_g$, or it is bad, and $z = z_b$.

The aggregate shock follows a first-order Markov structure given by the transition matrix

$$\begin{pmatrix} \pi_{gg} & \pi_{bg} \\ \pi_{gb} & \pi_{bb} \end{pmatrix},$$

where $\pi_{ss'}$ is the probability that the aggregate shock next period is $z_{s'}$ given that it is z_s this period.

The individual and aggregate shocks are correlated, and the individuals' shocks are assumed to satisfy a law of large numbers. More specifically, the number of agents who are unemployed always equals u_g in good times and u_b in bad times. In other words, controlling for z, individual shocks are uncorrelated. Because of the correlation between aggregate and individual shocks, we have to describe their joint distribution; hence, we assume the following first-order Markov process for (z, ϵ):

$$\begin{pmatrix} \pi_{gg00} & \pi_{bg00} & \pi_{gg10} & \pi_{bg10} \\ \pi_{gb00} & \pi_{bb00} & \pi_{gb10} & \pi_{bb10} \\ \pi_{gg01} & \pi_{bg01} & \pi_{gg11} & \pi_{bg11} \\ \pi_{gb01} & \pi_{bb01} & \pi_{gb11} & \pi_{bb11} \end{pmatrix},$$

where $\pi_{ss'\epsilon\epsilon'}$ denotes the probability of transition from state (z_s, ϵ) today to $(z_{s'}, \epsilon')$ tomorrow. This transition matrix has to satisfy the restrictions

$$\pi_{ss'00} + \pi_{ss'01} = \pi_{ss'10} + \pi_{ss'11} = \pi_{ss'}$$

and

$$u_s \frac{\pi_{ss'00}}{\pi_{ss'}} + (1 - u_s)\frac{\pi_{ss'10}}{\pi_{ss'}} = u_{s'}$$

for all four possible values of (s, s').

The technology is restricted to give output, y, as a Cobb-Douglas function of capital input, k, and labor input l:

$$y = zk^{\alpha} l^{1-\alpha},$$

with $\alpha \in [0, 1]$. Output can be transformed into future capital, k', and current consumption according to

$$c + k' - (1 - \delta)k = y,$$

where $\delta \in [0, 1]$ is the rate of depreciation.

2.2 The Market Arrangement

For the above economic environment, the assumption of complete markets gives an aggregation theorem: it is possible to determine full contingent plans for total capital accumulation, and to solve for all state-contingent prices, by considering one consumer only.[3] Here, however, we assume that there are incomplete markets. In particular, we assume that there is only one asset—call it capital—which plays the twin roles of being a store of value for the individual agent and a means of self-insurance against the income shocks. In this respect, our approach parallels those adopted in a number of recent papers, including İmrohoroğlu (1989), İmrohoroğlu (1992), Huggett (1993), Aiyagari and Gertler (1991), Aiyagari (1994), and Díaz-Giménez et al. (1992). Thus, let k denote the holdings of capital. We restrict capital holdings to satisfy $k \in \mathcal{K} \equiv [0, \infty)$, reflecting the need for a ruling-out of Ponzi schemes and a restriction that loans are paid back.[4] To maintain consistency with related literature, we refer to the lower bound on capital as the borrowing constraint (even when this lower bound is nonnegative).

[3] See Chatterjee (1994) or Krusell and Ríos-Rull (1994) for a discussion.

[4] One could also consider a negative lower bound on holdings of capital. As Aiyagari (1994) shows, in a context of no aggregate productivity shocks an always-pay-back constraint implies that a lower bound for capital which is less than zero is equivalent to one where it is zero. In this sense, our lower bound is generous. Later, we also consider positive lower bounds.

Consumers collect income from working and from the services of their capital. If the total amount of capital in the economy is denoted $\bar{k}$ and the total amount of labor supplied is denoted $\bar{l}$, our constant-returns-to-scale production function implies that the relevant prices are $w(\bar{k}, \bar{l}, z) = (1-\alpha)z\left(\bar{k}/\bar{l}\right)^{\alpha}$ and $r(\bar{k}, \bar{l}, z) = \alpha z\left(\bar{k}/\bar{l}\right)^{-\alpha}$, respectively.

We consider a recursive equilibrium definition, which includes, then, as a key element, a law of motion of the aggregate state of the economy. The aggregate state is (Γ, z), where Γ denotes the current measure (distribution) of consumers over holdings of capital and employment status. The part of the law of motion which concerns z is exogenous; it can be described by z's transition matrix. The part which concerns updating Γ is denoted H; in other words, $\Gamma' = H(\Gamma, z, z')$.

For the individual agent, the relevant state variable is his holdings of capital, his employment status, and the aggregate state: $(k, \epsilon; \Gamma, z)$. The role of the aggregate state is to allow the consumer to predict future prices. His optimization problem can therefore be expressed as follows:

$$v(k, \epsilon; \Gamma, z) = \max_{c,\, k'} \{u(c) + \beta E[v(k', \epsilon'; \Gamma', z')|z, \epsilon]\}$$

subject to

$$c + k' = r(\bar{k}, \bar{l}, z)k + w(\bar{k}, \bar{l}, z)\,\tilde{l}\epsilon + (1-\delta)k$$

$$\Gamma' = H(\Gamma, z, z')$$

$$k' \geq 0$$

and the stochastic laws of motion for z and ϵ. If the agent is employed ($\epsilon = 1$), he supplies an exogenous amount of labor $\tilde{l}$. The decision rule for the updating of capital coming out of this problem is denoted by the function f: $k' = f(k, \epsilon; \Gamma, z)$.

A **recursive competitive equilibrium** is then a law of motion H, a pair of individual functions v and f, and pricing functions (r, w) such that (i) (v, f) solves the consumer's problem, (ii) r and w are competitive (i.e. given by marginal productivities as expressed above), and (iii) H is generated by f, i.e. the appropriate summing up of agents' optimal choices of capital given their current status in terms of wealth and employment.

Our computational methods focus on finding stationary stochastic equilibria only, and we therefore need to define what we mean by a stationary stochastic equilibrium for our economy. We mean by a stationary stochastic equilibrium a recursive equilibrium described by an ergodic set $\mathcal{D}$ of distributions (i.e. a set such that once there, the economy never leaves the set) and an invariant probability measure $\mathcal{P}$ over this set. We therefore seek a good approximation to the function H over the set $\mathcal{D}$.

2.3 Computational Strategy

We now outline our algorithm for computing equilibria numerically.[5] This description is nontechnical and is included in the main body of the paper because the procedures are intimately connected with the economic mechanisms we wish to emphasize. The state variable of the economy which is endogenous, Γ, is a high-dimensional object, and in the numerical analysis it is necessary to restrict its representation to a finite vector. This could be done in many ways. Our strategy is to approximate the distribution with a finite set of the first I moments: $m \equiv (m_1, m_2, \ldots, m_I)$.[6]

Next, to approximate the law of motion for the distribution, we use a class $\mathcal{S}$ of functions H_I giving m', i.e. the vector of I moments in the next period, as a function of the I current moments: $m' = H_I(m, z, z')$. Our final approximation of an equilibrium builds on the following. When agents take as given the law of motion H_I of the vector m representing the distribution, the resulting accumulation of capital can be represented by a decision rule f_I, and this decision rule can be simulated for a large number of agents to yield time series for the I moments. Our final approximation of an equilibrium is a function H_I which, when taken as given by the agents, (i) yields the best fit within the class $\mathcal{S}$ to the resulting time series; and (ii) yields a fit which is close to perfect (H_I tracks the resulting simulated behavior almost exactly, i.e. with very small errors). In a computed, "approximate" equilibrium, thus, agents do not take into account all the moments of the distribution, but the errors in forecasting prices that result from this omission are very small.

More precisely, our algorithm amounts to the following iterative procedure: (1) Select I. (2) Guess on H_I in the form of some given parameterized functional form, and guess on parameter values. (3) Solve the consumer's problem given H_I. This step, which builds on a nonlinear approximation of the value function, is described in more detail in the Appendix. (4) Use the obtained decision rules to simulate the behavior of N agents (with N a large number) over a large number, T, of time periods. (5) Use the stationary region of the simulated data to estimate a set of parameters for the functional forms assumed above.[7] At this stage, we obtain a goodness-of-fit. (6) If the estimation gives parameter values that are very close to those guessed initially, and the goodness-of-fit is satisfactory, stop. If the parameter values have converged but the goodness-of-fit is not satisfactory, increase I, or as an intermediate step, try a different functional form for H_I.[8]

As an illustration, consider the following example. Assume that $I = 1$ and that H_I is linear:

[5]Our algorithm bears some similarities to one proposed in Díaz-Giménez and Ríos-Rull (1992).

[6]By m_1 we are referring to a 2×1 vector consisting of the first moments of Γ; m_2 is a 2×2 matrix consisting of the second moments of Γ, and so on.

[7]To define this region, we discard an initial part of the time series and check that the behavior of the moments of interest in the remaining part of the series appears to be stationary.

[8]In our implementation, we use a more flexible procedure in which the state vector can consist not only of moments of the distribution but also of other statistics describing the distribution, such as tail probabilities, which are themselves nonlinear functions of the distribution's moments.

$$z = z_g : \quad \bar{k}' = a_0 + a_1 \bar{k}$$

$$z = z_b : \quad \bar{k}' = b_0 + b_1 \bar{k}.$$

The agent then solves the following problem:

$$v(k, \epsilon; \bar{k}, z) = \max_{c,k'} \{u(c) + \beta E[v(k', \epsilon'; \bar{k}', z')|z, \epsilon]\}$$

subject to:

$$c + k' = r(\bar{k}, \bar{l}, z)k + w(\bar{k}, \bar{l}, z)\tilde{l}\epsilon + (1 - \delta)k$$

$$\bar{k}' = a_0 + a_1 \bar{k} \text{ if } z = z_g$$

$$\bar{k}' = b_0 + b_1 \bar{k} \text{ if } z = z_b$$

$$k' \geq 0$$

and the law of motion for (z, ϵ). We thus obtain a (nonlinear) decision rule

$$k' = f_I(k, \epsilon; \bar{k}, z)$$

which, when simulated, allows us to compare the aggregate behavior of the moments with their description H_I.

3 Results

We now describe the results, beginning with those of the baseline setup in Section 3.1. In Section 3.2 we look at a model with two types of agents who differ in patience or risk aversion, and in Section 3.3 we consider valued leisure: the real-business-cycle model. Finally, we briefly comment on a few other extensions of the model in Section 3.4.

3.1 The Baseline Setup

We first list the parameter values we use in our calculations:

A. Model parameters

We use $\beta = 0.99$ and $\delta = 0.025$, reflecting a period of one quarter, a relative-risk-aversion parameter σ of 1, and a capital share α of 0.36. We let the value of the technology shock be 1.01 in good times and 0.99 in bad times. The unemployment rate is chosen to equal 0.04 in good times and 0.1 in bad times. These values for the technology shock and the unemployment rate lead to fluctuations in the macroeconomic aggregates which have roughly the same magnitude as the fluctuations in observed postwar U.S. time series. We select the values of the stochastic process for (z, ϵ) so that the average duration of both good and bad times is 8 periods and so that the average duration of an unemployment spell is 1.5 in good times and 2.5 in bad times. The exogenous amount of labor supplied by an employed agent ($\tilde{l}$) is set equal to 0.3271.[9] We also impose $\pi_{gb00} = 1.25\,\pi_{bb00}$ and $\pi_{bg00} = 0.75\,\pi_{gg00}$.

B. Solution and simulation parameters

We solve the consumer's problem by computing an approximation to the value function on a grid of points in the state space. We use cubic spline and polynomial interpolation to compute the value function at points not on the grid. See the Appendix for a detailed description of the numerical algorithm used to solve the consumer's problem. In our simulations we include 5,000 agents and 11,000 periods, where we discard the first 1,000 time periods. Typically, the initial wealth distribution in the simulations is one in which all agents hold the same level of assets. We find that our results are not sensitive to changes in the initial wealth distribution.

3.1.1 Equilibrium properties: only the mean matters

The main finding is that the fit is very good even when only the first moment is used. We therefore first present our approximate equilibrium in a version where the endogenous state variable, as in our example above, is the population mean of capital only. For this case we characterize the equilibrium and document the good fit in a few different ways. We then describe analyses where we allow more moments to play a role in equilibrium.

A. Only one moment

With a log-linear functional form and only the population mean of the capital stock as a state variable, we obtain the following approximate equilibrium:[10]

$$\log \bar{k}' = 0.095 + 0.962 \log \bar{k}$$

$$R^2 = 0.999998 \quad \hat{\sigma} = 0.0028\%$$

in good times and

$$\log \bar{k}' = 0.085 + 0.965 \log \bar{k}$$

$$R^2 = 0.999998 \quad \hat{\sigma} = 0.0036\%$$

in bad times. There are two measures of fit: R^2 and the standard deviation (in per cent) of the regression error, $\hat{\sigma}$. Using our simulated sample (consisting of 10,000 observations) we plot tomorrow's aggregate capital against today's aggregate capital. This graph (see Figure 1) is a clear illustration of the high R^2 and the low $\hat{\sigma}$.[11]

[9] We choose $\tilde{l} = 0.3271$ so that the supply of labor in the setup with exogenous labor supply is the same on average as the equilibrium supply of labor in the setup with valued leisure (see Section 3.3). This choice allows a direct comparison of the average capital stocks in the two economies.

[10] We also used a nonlinear flexible functional form for the law of motion with virtually identical results.

[11] The top line in the graph yields the law of motion for aggregate capital in good times, while the bottom yields the law of motion in bad times. The middle line is the 45-degree line.

Insert Figure 1 here.

In terms of R^2's and $\hat{\sigma}$'s, these fits are extremely good. However, could it be that other measures of fit are less impressive?[12]

In addition, we report percentage standard errors of the regression as well as the forecast errors agents make by using one moment only in predicting future means. The following table summarizes the accuracy of the 1-quarter and 25-year forecasts of the mean capital stock and the prices of capital and labor services. (In the table below, *corr*$(x, \hat{x})$ denotes the correlation between variable x and its forecast $\hat{x}$, where the forecasts assume that the sequence of future aggregate shocks z is known with certainty. The column labeled *max % error* lists the maximum percentage forecast error across the 10,000 simulated time periods.)

Forecasting accuracy				
	1 quarter ahead		*25 years ahead*	
Variable	*corr*$(x,\hat{x})$	*max % error*	*corr*$(x,\hat{x})$	*max % error*
Capital	0.999999	0.0143	0.99961	0.237
Rental rate	1.000000	0.0091	0.99989	0.152
Wage rate	0.999999	0.0051	0.99958	0.086

We see that forecasts are very good, even on as long a horizon as 25 years. Neither the fits nor the forecasts, however, are perfect, so it is relevant to ask whether, in our computed equilibrium, agents could do better by using more than just the current mean to predict future means of the capital stock. If we let an agent in our equilibrium include three more moments to explain the movements in the mean in the simulated data, the following regression results emerge (the numbers in parentheses are t-statistics):

$$\log \bar{k}' = 0.092 + 0.963 \log \bar{k} + \underset{(69.9)}{0.00087 \log s_2} - \underset{(-18.0)}{0.00018 \log s_3} + \underset{(11.9)}{0.00011 \log s_4}$$

$$\begin{array}{ll} R^2 = 0.999999 & \text{(before: } 0.999998) \\ \hat{\sigma} = 0.0018\% & \text{(before: } 0.0028\%) \end{array}$$

in good times and

$$\log \bar{k}' = 0.081 + 0.965 \log \bar{k} + \underset{(57.5)}{0.0012 \log s_2} - \underset{(-23.3)}{0.00029 \log s_3} + \underset{(15.3)}{0.00019 \log s_4}$$

$$\begin{array}{ll} R^2 = 0.999999 & \text{(before: } 0.999998) \\ \hat{\sigma} = 0.0024\% & \text{(before: } 0.0036\%) \end{array}$$

[12] For an AR(1) process with autoregressive coefficient given by ρ, the R^2 associated the regression of current on lagged values equals ρ^2 in a large sample, regardless of the standard deviation of the innovation. Here, we have a somewhat more complicated process for which the large-sample R^2's in the two regressions depend not only on the two autoregressive coefficients but also on the sizes of the errors. In particular, as the errors become arbitrarily small, one can show that the two R^2's go to 1. As the errors become arbitrarily large, on the other hand, we observe a phenomenon similar to that for an AR(1) process: the R^2's in the two equations have lower bounds which depend on the values of the autoregressive coefficients (if the autoregressive coefficients are both equal to ρ, then the lower bounds are both ρ^2). Notice, though, that the coefficients on $\log \bar{k}$ in our regressions are endogenous, so in this sense the R^2's are not *a priori* bound to be high.

in bad times, where s_2, s_3, and s_4 are, respectively, the standard deviation, skewness, and kurtosis of the (marginal) capital stock distribution.[13] Two facts are worth pointing out about these regressions. First, the other moments enter significantly (in some cases highly so), and second, the coefficients on the other moments are very small and hence the quantitative significance in terms of improved forecasting ability is very small. The fact that the other moments are significant is not a surprise: if indeed there is no aggregation theorem for the economy with incomplete insurance, so that the other moments do matter, then t-values go to infinity as the sample size is increased even if the coefficients are very close to zero. Our sample size is large and the significance levels simply reflect this fact.

The conclusion we draw from these results is that although the other moments do play a role in our computed equilibrium, their quantitative importance is extremely small. In this sense, then, agents are very close to optimal behavior, which is what equilibrium dictates. It should also be said that this statement can be made in utility terms, or in consumption equivalents: in our equilibrium consumers could do better by using all the moments for predicting future prices, but this would give them no more than a vanishingly small increase in utility expressed in terms of an average percentage increase in consumption across all dates and states.[14]

It should be pointed out that our computational procedure does not provide bounds on how far our approximate equilibrium deviates from an exact equilibrium. This shortcoming, however, is true for almost all numerical procedures used in economics, and in Section 5 below we make a few more comments on these important issues. We cannot, for example, rule out that the exact equilibrium is such that higher moments matter in a more significant way than what the results above suggest. To gain further insight about what happens if agents use higher moments in equilibrium, we also present results for the case in which a second moment is added to the state vector.

B. More moments

In this section we add another aggregate state variable to our aggregate law of motion. We first use the natural next candidate given our algorithm—the standard deviation of the marginal distribution of capital. These results are reported on in detail below. We also comment on what happens when we use a hybrid of higher moments as an additional element in the aggregate state vector: the fraction of agents at the lowest levels of asset holdings.[15]

Letting s_2 denote the standard deviation of the (marginal) capital stock distribution, the fixed

[13] In the simulated data, s_2 varies in the range 4.8 to 5.6, s_3 varies in the range 0.65 to 1.22, and s_4 varies in the range 3.4 to 5.7.

[14] As documented in our earlier work (Krusell and Smith (1994)), the calibrated versions of this kind of economy give objective functions for agents in equilibrium which are quite flat around the optimum.

[15] Specifically, we incorporate as an additional state variable the fraction of agents with wealth holdings less than or equal to 6. In the simulated data for the baseline model, this fraction ranges from 6% to 18%, with an average value of 11.5%.

point we obtain is as follows:

$$\log \bar{k}' = 0.094 + 0.963 \log \bar{k} + 0.00016 \log s_2$$

$$R^2 = 0.999999 \quad \hat{\sigma} = 0.0019\%,$$

$$\log s_2' = 0.048 - 0.019 \log \bar{k} + 0.999 \log s_2$$

$$R^2 = 0.9998 \quad \hat{\sigma} = 0.043\%$$

in good times and

$$\log \bar{k}' = 0.084 + 0.965 \log \bar{k} + 0.00030 \log s_2$$

$$R^2 = 0.999999 \quad \hat{\sigma} = 0.0026\%,$$

$$\log s_2' = 0.057 - 0.019 \log \bar{k} + 0.994 \log s_2$$

$$R^2 = 0.9995 \quad \hat{\sigma} = 0.073\%$$

in bad times. We see that the R^2's associated with the law of motion for aggregate capital go up (and the $\hat{\sigma}$'s go down) relative to the baseline case with one moment, suggesting that the addition of the second moment does improve the fit. However, the results also suggest quite clearly that the original findings—that only the mean matters—are not sensitive to the addition of further moments. The law of motion for the mean changes very little; the coefficient on current capital is almost exactly the same as before, and the coefficient on the second moment is very small. Relatedly, although we do not provide the plots here, inspection of the simulated time series in the cases with one and two moments reveals that it is virtually impossible to tell the difference between the two sets of series. Finally, the R^2 is high for both equations, although the fit it slightly worse for the updating of the second moment.

The results using the second higher-moment measure—the fraction of agents at the lowest levels of asset holdings—are similar to those using the standard deviation. The fit of the aggregate law of motion is improved relative to the one-moment case, but the corresponding change in the law of motion is very slight. In addition, the fit using the standard deviation as the additional higher-moment measure is slightly better than the fit using the fraction of agents with low levels of asset holdings. Once again, it is difficult to distinguish simulated time series for the one-moment case from simulated time series for the two-moment case.

An important issue, about which we are unable to draw any theoretical conclusions, is that of uniqueness of our stationary stochastic equilibrium. It is for example conceivable that the simplicity of the law of motion perceived by agents is "self-fulfilling". However, based on the fact that we do not find signs of multiple equilibria in our numerical computations, we conjecture that equilibria are unique for the environments we look at. The result in this subsection that adding a second moment to the state vector does not change the economy's aggregate implications also supports this conjecture.

3.1.2 Why is only the mean important?

One way of understanding our main findings is to account for savings throughout the population. In the graph below (see Figure 2), we show the decision rules of the private agent in the two employment cases.[16]

Insert Figure 2 here.

It is clear from the graph that the slopes are very similar both across employment states and levels of capital. This fact is an important part of the explanation for our results: if the marginal propensity to save out of current capital is the same for all agents, redistributions of the total stock of capital have no effect on aggregate savings. The slopes of the decision rules depicted in Figure 2 vary between 0.993 and 1.001 for employed agents and between 0.985 and 1.001 for unemployed agents for most of the range of capital holdings, with the exception being the very lowest range (capital holdings less than 5), where slopes decline significantly as capital holdings approach zero.

To account for how these differences in slopes aggregate, it is crucial to look at how the distribution of capital behaves in our equilibrium. Clearly, if most of the capital is redistributed to the agents in the very lowest range, there would be large changes in aggregate savings. Here, the fact that we are studying a stationary equilibrium is important, because it limits the set of distributions we are considering. It turns out, because low levels of capital are associated with relatively large risks of large drops in utility, that very few agents have low levels of capital—for example, the fraction of agents with a significantly different marginal propensity to save (i.e. with capital holdings less than 5) varies between 3% and 12%, with an average value of 6.5%. We illustrate this point in the diagram below (see Figure 3), where we display a typical equilibrium distribution of the capital stock.

Insert Figure 3 here.

Since most agents are far away from their borrowing constraint (i.e. the lower bound on capital holdings), they behave as if markets were complete most of the time. Furthermore, in our baseline environment the agents' Engel curves (which express consumption as a function of individual wealth) are linear if markets are complete, i.e. there is aggregation. Thus, the ability of capital to shield agents from the market frictions is clearly an important element behind our findings.

Not only is the presence of capital important for understanding the behavior of this model, but it also helps to explain why propensities to consume may differ more substantially in endowment

[16] The graph shows the decision rules for a typical agent assuming a good aggregate state and a value of aggregate capital equal to 11.7. The top line is the decision rule of an employed agent, while the bottom line is the decision rule of an unemployed agent. The middle line is the 45-degree line.

economies (see den Haan (1994a,1994b), Heaton and Lucas (1994a,1994b), Lucas (1994), Marcet and Singleton (1991), and Telmer (1993)). If the aggregate endowment is not very large in the latter economies, there are by necessity more agents close to those endowment levels where propensities to consume are significantly dispersed: the fact that one agent has been lucky and reached a large asset holding means that the opposite must have happened to some other agent.

The similarities in propensities to consume are not the whole story, however: even in the range of asset holdings where most agents are located, the differences in marginal propensities are large enough that it is possible to construct redistributions which would significantly change total savings. However, the movements in the distributions in our stationary equilibrium are simply not large enough to give rise to these kinds of redistributions: for a given capital mean, the set of distributions which appear in equilibrium does not display enough diversity to allow these differences in marginal propensities to produce significant changes in aggregate savings.

It may be informative also to provide complementary evidence that the distributions which appear in our computed equilibrium are not of a one-parameter family. In the graph below (see Figure 4) we use our simulated sample to graph the standard deviation of the (marginal) capital stock distribution against its mean. It seems clear from this graph that there is no one-to-one relation between these two moments.[17]

Insert Figure 4 here.

3.1.3 Are the results special to our parameter values?

To gauge the sensitivity of our main finding, and to analyze in more detail which features of the economic environment are critical in producing it, we look at a number of perturbations of the parameter values of the baseline setup. First, we increase both the frequency and duration of unemployment spells so that agents find it harder to stay away from low levels of capital, and we examine the extent to which marginal propensities to consume display more diversity in the population in equilibrium. Second, we tighten the borrowing constraint.[18] Third, we change the discount factor and the depreciation rate. In connection with this experiment we also discuss the effects of changing the length of calendar time corresponding to one model period. Fourth, we study the effects of varying the risk aversion parameter.

Later (in Section 3.2) we also add heterogeneity in preferences in the population, partly as a way of making our setup diverge more from the permanent-income tradition—with agents with different degrees of patience in the population, there are presumably larger differences in savings

[17]Note that this finding does not imply that omitting the second moment significantly reduces forecasting performance. In this context, note from above that, in our approximate equilibrium with two moments, the first and second moments come close to following univariate laws of motion. Hence the relation between the first and second moments, which in the graph looks quite weak, is not so important for forecasting.

[18]Recall that we refer to the lower bound on capital holdings as the borrowing constraint, even when this lower bound is nonnegative.

propensities. The question is then how large these differences become in equilibrium, and whether the aggregate stock of capital is still sufficient for making good forecasts of future prices. In Section 3.3 we also look at a setup where leisure is valued—a "real-business-cycle" model—since this economy has the added feature that current prices depend on the distribution of capital in a more complicated way than in a model without value leisure. This nontrivial dependence arises because total labor input is not fixed in this case: since work effort varies with wealth, the rental rate of capital and wages depend not only on the total amount of capital but also on its distribution. Finally, we also comment briefly on two other extensions.

A. More frequent and persistent unemployment spells: making bad shocks worse

When we increase the frequency and duration of unemployment spells, the fraction of agents with low levels of asset holdings (i.e. with asset holdings less than 4) increases slightly relative to the baseline model, implying that marginal propensities to save vary across the population somewhat more than in the baseline model. The fraction of agents with asset holdings close to the borrowing constraint increases only slightly because low levels of asset holdings are potentially more painful when agents face the possibility of longer and more frequent unemployment spells. In this case, agents respond endogenously to this increased frequency and duration by avoiding low levels of asset holdings.

At the same time, the average stock of capital increases in response to the increased risk of getting long runs of bad luck. Thus, again we see that the possibility of capital accumulation serves to diminish the effects of market frictions: when the incompleteness of markets has particularly bad consequences for unlucky agents, the capital stock in the stationary equilibrium is particularly high.

On net, changes in forecasting performance using log-linear laws of motion for aggregate capital are not substantial when compared to the baseline model. In particular, R^2's are a little lower here than in the baseline model, but are still very high, and agents continue to make very close to optimal forecasts. In the experiment, we choose the unemployment rate to be 0.1 in good times and 0.14 in bad times, and the average duration of unemployment spells to be 5 quarters in good times and 7 quarters in bad times.[19] The law of motion for aggregate capital is as follows:

$$\log \bar{k}' = 0.099 + 0.960 \log \bar{k}$$

$$R^2 = 0.999993 \quad \hat{\sigma} = 0.0043\%$$

in good times and

$$\log \bar{k}' = 0.094 + 0.961 \log \bar{k}$$

$$R^2 = 0.999988 \quad \hat{\sigma} = 0.0060\%$$

in bad times.

[19]We also set π_{gb00} equal to $1.1\,\pi_{bb00}$ rather than $1.25\,\pi_{bb00}$.

B. Tighter borrowing constraints

If we make the borrowing constraint tighter (i.e. make the lower bound on asset holdings larger) more agents should be forced toward levels of wealth associated with small marginal propensities to save (i.e. toward levels of wealth close to the borrowing constraint). We use a borrowing limit of 8 rather than 0 as in the baseline economy (on average, 26% of agents have capital holdings less than 8 in the baseline economy). We find that the accuracy is still very good using only the mean: the R^2 for the law of motion for the mean is 0.999993 ($\hat{\sigma} = 0.0046\%$) in good times and 0.999978 ($\hat{\sigma} = 0.0085\%$) in bad times.

One effect of this experiment is to increase the average level of capital holdings by 5% (to 12.2 from 11.6) relative to the baseline economy, so there is again a sense in which the economy endogenously works its way out of situations in which market frictions hurt. In contrast, endowment economies that have a low total endowment, tight borrowing constraints, or very persistent idiosyncratic income processes do not offer this flexibility.

C. The discount factor and the rate of depreciation

The most important parameter from the point of view of the sensitivity of our results is the discount factor. Intuitively, this can be understood as follows. A key feature of our results is that agents' behavior has a permanent-income flavor: decision rules for capital are close to linear in current capital with a coefficient near unity. As the discount factor decreases, the permanent-income flavor is bound to become weaker, and we should expect decreases in the propensities to save out of current capital. When we depart from permanent-income behavior, there is also no reason to continue to expect that propensities to save are very similar across agents.

In our experiments, we also let the depreciation rate decrease toward zero as we increase the discount factor toward one, and vice versa. We choose to report the simultaneous change in β and δ because both these parameters are calibrated with reference to the length of a period: as the length of a period changes, both these parameters need to be changed appropriately. Our results are as follows:

Goodness-of-fit measures				
Period in quarters	β	δ	R^2	$\hat{\sigma}$
1/12	0.9992	0.0019	1.000000	0.000014%
1/3	0.997	0.0083	0.999999	0.0011%
1	0.990	0.025	0.999998	0.0029%
4	0.961	0.1	0.99997	0.013%
8	0.923	0.2	0.9998	0.021%
20	0.818	0.5	0.9986	0.027%
40	0.669	1.0	0.9948	0.032%

Clearly, as β decreases (and δ increases) the R^2 do fall (the reported measures of fit are for the law of motion in good times). Substantial changes are necessary, however, to make a big difference for the fit. Similarly, when β is increased, the fit improves.

It is informative to compare these experiments to the theoretical results in Bewley (1977).[20] Bewley considers a decision-theoretic model where the agent can save at a safe return of unity, and where the income stream is stochastic. His analysis shows that as the discount factor goes to one, and the agent is given a larger and larger initial asset holding, the savings function becomes linear, with a coefficient on current capital of unity. In our equilibrium framework, letting (β, δ) approach $(1, 0)$ seems to have the same effect: heuristically, the average level of capital in a stationary stochastic equilibrium goes to infinity, since the marginal productivity of capital is forced to go to zero (cf. the agent's first-order condition). Hence, the return to savings becomes more and more safe and close to unity, while the initial holdings of assets increase without bound in the corresponding stationary equilibrium, with the problem of the agent thus corresponding to the case Bewley looked at.[21]

D. The degree of risk aversion

When the degree of risk aversion, σ, is changed none of the main findings change. We computed approximate equilibria for values of σ ranging from 1 to 5. The forecasting accuracy improves slightly as the degree of risk aversion increases. Some quantitative aspects of the model, of course, do vary with σ. Higher risk aversion leads to smaller variability in consumption, to a higher average capital stock, and to a tighter distribution of capital. The increased need to smooth consumption thus is accomodated in two ways: first, by increasing the average amount of capital held; and second, given the average level of capital, by forcing capital to stay further away from zero so that consumption levels close to zero become very rare.

3.2 Preference Heterogeneity and the Dispersion of Aggregate Wealth

Suppose there are other differences among agents than wealth differences. In particular, suppose there is preference heterogeneity. Then, we might expect larger differences in savings behavior across agents, implying a more important role for higher moments of the asset distribution in determining prices and aggregate variables. We look at two economies with preference heterogeneity: one where there are two types of agents with differing discount factors, and one where the two types differ in their degrees of risk aversion.

The heterogeneity in discount factors brings in some features which are inherent in models with life-cycle behavior: the patient (young) agent would be expected to have a larger savings propensity

[20] We thank José Scheinkman for drawing our attention to the connection between our setup and that studied by Bewley.

[21] Notice also how Bewley's experiment relates to letting the length of the time period decrease toward zero, i.e. to moving toward a setting with continuous time. However, in his and in our exercise, the employment transition probabilities are not changed, nor is the shape of the utility function, as the discount factor is moved toward one. Reasonable calibration would in contrast demand changes in these parameters as well. In this case, for example, the average unemployment duration (measured in terms of the number of model time periods) would increase as the time period's length is shortened, thereby ensuring that unemployment spells continue being similar in length to those in the data. When the employment transition probabilities are not changed, the limit case essentially amounts to no idiosyncratic uncertainty: over any positive time interval, the accumulated labor income becomes deterministic.

than the impatient (old) agent. Consequently, at least if the fractions of wealth held by the two types of agents vary over time, it is possible that our results would weaken considerably.

Notice that the incomplete-markets economy with differences in patience among agents gives a much less trivial outcome for the stationary asset distribution than does the corresponding complete-markets economy. In the latter, stationary equilibria have the feature that the impatient agents have no wealth and are virtually absent—impatient agents borrow at the beginning of time and spend the rest of their lives paying back debt with zero consumption in the limit. In contrast, the incomplete-markets economy allows a nontrivial stationary distribution of wealth where the impatient agents hold a positive fraction of total wealth; this is because there are limits to how much they can borrow, and because good luck keeps coming back to increase their level of wealth. In particular, our experiment suggests differences in discount factors as one avenue for explaining why the holdings of capital are so concentrated in the U.S. data.[22]

We solve for an equilibrium with two types of agents, 50% of each type, whose discount factors are 0.99 and 0.984. The findings are as follows: the aggregate law of motion for capital with only one moment used in H is

$$\log \bar{k}' = 0.093 + 0.963 \log \bar{k}$$

$$R^2 = 0.99997 \quad \hat{\sigma} = 0.011\%$$

in good times and

$$\log \bar{k}' = 0.086 + 0.964 \log \bar{k}$$

$$R^2 = 0.99995 \quad \hat{\sigma} = 0.014\%$$

in bad times. So in fact, although the goodness of fit worsens somewhat as compared to the baseline model, it is still very good.[23]

Note that the set of distributions we observe in our simulations of the two-patience-types economy is very different than the one recorded in the previous experiments. Given that the discount rates of both agents are close to (but higher than) the equilibrium interest rate, small changes in either the discount rate or the interest rate induce large changes in the average asset holdings of the two types of agents. For example, we see in Figures 5 and 6, in which typical capital stock distributions are plotted for each of the two kinds of agents, that impatient agents typically hold much less capital than patient agents. Indeed, the average capital holdings of a patient agent is 20.9, while the average capital holdings of an impatient agent is 2.2. Hence, a small amount of heterogeneity in the degree of patience causes the distribution of asset holdings to be much more dispersed. In particular, note that the cross-sectional distribution of capital for patient agents has a significantly thicker (longer) right tail than does the distribution of capital in the economy with

[22] For attempts to match these data, see Huggett (1994) and Castañeda, Díaz-Giménez, and Ríos-Rull (1994).

[23] Similarly, the forecasting accuracy, which we do not report, worsens somewhat but not enough to change the overall picture.

one type of agent.

Insert Figures 5 and 6 here.

Why, then, is the goodness of fit still so good? In our stationary stochastic equilibrium, a given agent's asset holdings fluctuate randomly around the average level of asset holdings associated with the agent's type, in effect making the equilibrium look like two subeconomies, one with a high level of average asset holdings and one with a low level of average asset holdings. Both types of agents continue to behave according to a permanent-income story, but, as expected, the marginal propensities to save out of current assets are different across the types (on average 0.937 for the low and 0.999 for the high discount factor types). Although the differences in slopes across agents are much larger here than in the baseline economy, it turns out not to be important in equilibrium. First, the fraction of total wealth held by the impatient agents is quite small (less than 10%). Since the subeconomy of patient agents accounts for most of the movements in aggregate wealth, and since the distribution of wealth within this subeconomy does not matter for the same reasons as in the previous sections, the distribution has little importance for the aggregate in this case as well. Second, the interaction between the two subeconomies is not very strong, as witnessed by relatively small movements in the respective fractions of wealth held by the two groups.

Turning to the economy with risk-aversion heterogeneity, we focused on a case where half of the population has $\sigma = 1$ (logarithmic preferences) and the other half $\sigma = 3$ (which corresponds to a higher degree of risk aversion). As for the case with multiple discount factors, the R^2's and $\hat{\sigma}$'s remain impressive: 0.999996 and 0.0049% in good times and 0.999996 and 0.0057% in bad times, respectively. Average equilibrium wealth of agents with $\sigma = 3$ is significantly higher than for agents with $\sigma = 1$: the average asset holdings for these groups are 15.7 and 7.7, respectively. In other words, risk-aversion heterogeneity does not change the total amount of capital in this economy much but it creates a much larger dispersion of wealth.

That higher risk aversion implies a higher average asset holding for precautionary purposes should be clear; that the difference between the two groups is so high is perhaps less clear. The finding can be understood by first looking at the nature of the agent's problem. The agent's optimal average (long-run) asset holding in this type of setup is very sensitive to the interest rate in the region where the interest rate is close to the discount rate.[24] In particular, as we have seen above, decision rules for asset holdings in this environment have slopes which are very close to 1. Hence a small change in the average interest rate, while inducing only a small change in the overall shape of decision rules, leads at the same time to large movements in the point of intersection between the 45-degree line and the decision rule for employed agents, thereby leading in turn to a large change in the long-run average holdings of capital. Similarly, for a given interest rate process, small changes in preference parameters give rise to large movements in the average capital holdings for the agent. Second, in the stationary equilibrium, it turns out that the the interest rate process adjusts so that

[24]This observation is also made in Aiyagari (1994).

the population-wide average of capital is at a level implying a return to saving slightly less than the discount rate. With only one type of agent in the economy, whether this agent has a σ of 1 or 3, the total stock of capital is therefore forced to a value of around 11.6, and hence the per-capita holding of capital for all agents is also pinned to that number. With two types of agents, the total stock of capital is close to the same number, but because of the sensitivity of decision rules to changes in preferences discussed above, the total capital stock must be split up very unevenly across the two types of agents.

On net, the added variation in marginal propensities across agents due to differences in patience and risk aversion has little impact on our basic finding that only the mean of asset holdings seems necessary in order to understand the evolution of the macroeconomic aggregates.

3.3 Valued Leisure: The Real-Business-Cycle Model Revisited

In this section we look at a real-business-cycle model—one which is similar in spirit to those developed in Kydland and Prescott (1982) and Long and Plosser (1983)—from the perspective that there is wealth and income heterogeneity due to idiosyncratic risk and incomplete insurance opportunities. We ask whether it is true in this economy as in our baseline setup that the mean of capital holdings is (nearly) sufficient to characterize aggregate time series properties. We will, subject to minor qualifications, answer this question affirmatively. Later (in Section 4), we ask whether the time series properties of the incomplete-markets economy are similar to those of the corresponding representative-agent model. As we will see, the answer to this question is affirmative as well, thus lending some theoretical support to the existing work in the representative-agent real-business-cycle literature.

Turning to a version of the model where leisure is valued, assume that time spent off work, $1-l$, where l is the amount of labor supplied, is valued according to

$$u(c,l) = \lim_{\nu \to \sigma} \frac{(c^{\theta}(1-l)^{1-\theta})^{1-\nu} - 1}{1-\nu}.$$

In our recursive equilibrium definition, there is now an additional element to consider: the way leisure is supplied at each point in time. Let the aggregate amount of hours worked be given by the function L:

$$\bar{l} = L(\Gamma, z).$$

This function is needed as an input in each agent's decision; to know prices, $\bar{l}$ needs to be determined, and this is what L delivers. Optimal decisions of the agent thus lead to the decision rule

$$l = g(k, \epsilon; \Gamma, z)$$

specifying how much to work at each value of the state. The equilibrium condition for L thus states that at any given state (Γ, z), $L(\Gamma, z)$ equals the total labor supply when integrated over the population using individual supplies given by the g's.

In a model where the leisure parameter θ satisfies $\theta/(1-\theta) = 1.9$ and the risk aversion parameter $\sigma = 1$, and where the labor supply function is approximated as a log-linear function of the total stock of capital, our results are as follows:

$$\log \bar{k}' = 0.123 + 0.951 \log \bar{k}$$

$$R^2 = 0.999994 \quad \hat{\sigma} = 0.0040\%$$

$$\log \bar{l} = -0.544 - 0.252 \log \bar{k}$$

$$R^2 = 0.992 \quad \hat{\sigma} = 0.039\%$$

in good times and

$$\log \bar{k}' = 0.114 + 0.953 \log \bar{k}$$

$$R^2 = 0.999993 \quad \hat{\sigma} = 0.0049\%$$

$$\log \bar{l} = -0.592 - 0.255 \log \bar{k}$$

$$R^2 = 0.988 \quad \hat{\sigma} = 0.054\%$$

in bad times.[25]

We first comment on the accuracy of the approximation. We see that the fit is still very good for the law of motion for capital (see also Figure 7, in which we use our simulated sample to plot tomorrow's aggregate capital against today's). For the aggregate labor function, the fit is good, but not as good as for the law of motion of capital—the R^2's are lower, and there are more significant clouds in the plot of aggregate labor supply against aggregate capital (see Figure 8).

Insert Figures 7 and 8 here.

To understand this finding, it is possible to note that the nonlinearities in an individual agent's decision to supply labor as a function of his own capital holdings are stronger and extend over a wider range than for the savings decision (the marginal propensity to take leisure is much higher for poor agents). This fact together with the fact that there are significant movements in the distribution of capital account for the clouds observed in the graph for aggregate labor. In reproducing the tables with forecasting errors and checking the forecasting ability of the other moments, it is possible to detect a decrease in accuracy as compared to the baseline model, but the changes are not large (because these results are quite similar to the those for the baseline model, we omit a full

[25] We also computed an approximate equilibrium using two moments of the (marginal) capital stock distribution: its mean and its standard deviation. In this case, the aggregate labor supply function depends on both of these moments. Including an additional moment leads to significantly better fits for the aggregate labor functions: in good times, $R^2 = 0.998$ ($\hat{\sigma} = 0.021\%$), while in bad times, $R^2 = 0.993$ ($\hat{\sigma} = 0.039\%$). In addition, the fit of the law of motion for aggregate capital improves slightly. Aggregate time series, however, are virtually unchanged.

report). The overall summary is that the results are close to those of the baseline setup.

3.4 Other extensions

Finally, let us comment briefly on two other extensions. First, consider an economy where agents must pay a fixed cost for choosing a level of capital different from some target value k^*. In this case, optimal savings behavior is characterized by a savings propensity of zero for a fraction of the population, whereas the remaining agents have a propensity closer to one. The reason for studying this setup is that it gives a simple way of forcing large differences in savings propensities, which in turn allows us to study more directly the importance for our results of movements in the distribution of capital across groups with different savings propensities. In this economy, we find that although the fit worsens, it is still surprisingly high: the R^2's are 0.9998 ($\hat{\sigma} = 0.017\%$) and 0.9999 ($\hat{\sigma} = 0.013\%$) for good and bad times, respectively. To understand this result, it is key to study the movements in the fraction of capital held by agents with a zero savings propensity. This fraction moves substantially over the cycle, but it turns out that for a given level of aggregate capital, the range of fractions which one observes in a stationary equilibrium is fairly small. In other words, given a level of aggregate capital, it is possible to make a reasonably accurate prediction about the fraction of agents with a zero savings propensity.

Second, we also wish to refer briefly to experiments which we did not perform ourselves: Castañeda et al. (1994) introduce heterogeneity in the processes for labor income, and the goodness of fit of that setup is of the same order of magnitude as in the current model.

4 Comparisons with the one-agent model

Given that only the mean level of capital holdings seems to matter for the equilibrium behavior of the baseline economy and its various extensions, can we conclude that the properties of aggregate time series in these economies are very close to those produced by a corresponding representative-agent model? Although we find that for many of the economies studied above there is indeed a close correspondence between the time series in the two types of setups, this correspondence is not a logical implication from our main results of the previous section.

To make the question more precise, suppose markets in our baseline economy are complete. In this case, there is an aggregation theorem which states that aggregate time series properties depend only the mean, and not the distribution, of capital across consumers. It is a straightforward computational problem to characterize this equilibrium, and we do so in order to compare it to the incomplete-markets economy. The following tables summarize the findings.[26]

[26] In the tables in this section, "h-a" refers to heterogeneous-agent models, while "r-a" refers to representative-agent models. Also, k denotes aggregate capital, y denotes aggregate output, c denotes aggregate consumption, i denotes aggegate investment, and l denotes aggregate hours. The numbers reported are sample means computed using a simulated sample consisting of 11,000 observations, with the first 1,000 observations being discarded in order to diminish dependence on initial conditions. In order to facilitate comparisons across models, we use the same sequence

Aggregate time series properties			
Variable, Model	*mean*	*std. dev.*	*corr. w/ y*
k, h-a	11.61	0.292	0.527
k, r-a	11.54	0.303	0.529
y, h-a	1.128	0.039	1
y, r-a	1.126	0.039	1
c, h-a	0.838	0.015	0.701
c, r-a	0.837	0.015	0.691
i, h-a	0.290	0.030	0.935
i, r-a	0.289	0.031	0.937

Correlation of output with its l-period lag						
Model	l = 1	l = 2	l = 3	l = 4	l = 5	l = 6
h-a	0.812	0.667	0.561	0.481	0.415	0.366
r-a	0.814	0.670	0.565	0.486	0.421	0.372

Although there is a slight difference in the average stock of capital across the two economies—the incomplete-markets economy employs more capital since it is used as a means of self-insurance—the aggregate time series are very similar in the two economies.[27] It is clear that even though the aggregation theorem fails in the incomplete markets-economy, macroeconomists would make almost no errors at all by treating the world as one in which all agents have the same income and wealth.

Although we do not report in detail how each of the economies studied above compares to its representative-agent counterpart, we do show results for some of them: the baseline model with a higher degree of risk aversion ($\sigma = 5$), the model with discount factor heterogeneity, and the real-business-cycle economy.[28] The results are as follows:

of aggregate shocks for all model economies when generating the simulated time series.

[27]It is interesting to note that precautionary savings seem to have little aggregate importance unless risk aversion and/or the severity of the idiosyncratic shocks is increased substantially (for the economy with more severe unemployment shocks described in Section 3.1.3, the mean level of aggregate capital is 11.55, as compared to 10.93 in the corresponding complete-markets economy). These findings confirm what was first pointed out in Aiyagari (1994) in the context of an economy without aggregate shocks.

[28]The representative-agent economy corresponding to the model with heterogeneity in discount factors has only one type of economically active agent—the patient agent—since the asymptotic behavior of the impatient agents is characterized by zero consumption when markets are complete. Note that the real-business cycle model with complete markets permits aggregation for the class of preferences that we use.

Aggregate time series properties									
Variable,	*Model with $\sigma = 5$*			*Model with two β's*			*RBC model*		
Model	*mean*	*std. dev.*	*corr. w/ y*	*mean*	*std. dev.*	*corr. w/ y*	*mean*	*std. dev.*	*corr. w/ y*
k, h-a	12.32	0.459	0.544	11.57	0.243	0.494	11.58	0.228	0.452
k, r-a	11.55	0.500	0.574	11.54	0.303	0.529	11.56	0.233	0.454
y, h-a	1.152	0.041	1	1.127	0.038	1	1.127	0.033	1
y, r-a	1.126	0.041	1	1.126	0.039	1	1.128	0.032	1
c, h-a	0.844	0.012	0.741	0.837	0.016	0.865	0.838	0.011	0.669
c, r-a	0.837	0.013	0.725	0.837	0.015	0.691	0.839	0.011	0.639
i, h-a	0.308	0.033	0.967	0.289	0.025	0.946	0.289	0.027	0.951
i, r-a	0.289	0.034	0.966	0.289	0.031	0.937	0.289	0.027	0.952
l, h-a	—	—	—	—	—	—	0.304	0.008	0.954
l, r-a	—	—	—	—	—	—	0.305	0.008	0.950

Correlation of output with its l-period lag						
h-a, $\sigma = 5$	0.826	0.693	0.596	0.524	0.465	0.421
r-a, $\sigma = 5$	0.836	0.710	0.619	0.551	0.495	0.453
h-a, two β's	0.803	0.650	0.539	0.456	0.387	0.336
r-a, two β's	0.814	0.670	0.565	0.486	0.421	0.372
h-a, RBC	0.787	0.622	0.502	0.412	0.339	0.285
r-a, RBC	0.786	0.622	0.503	0.414	0.342	0.289

We see that the differences between the economies with and without wealth heterogeneity are very slight both for the $\sigma = 5$ and the RBC frameworks. More precisely, the aggregate level of capital is somewhat higher in the incomplete-markets economies due to the insurance role of capital, but all the other first and second moments are almost identical. In particular, we thus note that the aggregate time series properties of the standard real-business-cycle model are robust to introducing wealth heterogeneity of the kind we consider in this paper.

The third model reported in the last set of tables—the one with patient and impatient agents—also shows similarities across its representative-agent and heterogeneous-agent versions. However, here there are some clear dissimilarities as well. In particular, population-wide means and standard deviations are quite close, whereas correlations across variables differ substantially. For example, the correlation between consumption and output is much higher in the heterogeneous-agent than in the representative-agent model. This increase reflects the difference in the response of patient and impatient agents to shocks, with impatient agents displaying a higher marginal propensity to consume out of added income. Consequently, the correlation between aggregate output and the aggregate consumption of patient agents is 0.63, whereas the correlation between aggregate output and the aggregate consumption of impatient agents is 0.95.[29]

The final finding is interesting since it demonstrates the point that the "only the mean matters for prices" property (which is satisfied in all our models) is *not* equivalent to saying that adding

[29]We find fewer differences between the complete-markets and incomplete-markets versions of the economy with heterogeneity in risk-aversion parameters. Here, the comparison is less straightforward since the representative-agents (there are now two representative agents) version of the model has a continuum of stationary distributions indexed by the relative wealth of each of the two agents, whereas the incomplete-markets version gives a unique process for the wealth distribution. As relative wealth is varied in the version with (two) representative agents, the aggregate time-series properties vary continuously and in a monotone way between the two extreme cases.

incomplete markets and wealth heterogeneity into a representative-agent model does not change the model predictions. The equivalence seems to be roughly met for some frameworks but clearly not for the economy with heterogeneity in discount factors. Thus, our methodological findings are that models with wealth heterogeneity are feasible to study, and our substantantive findings are that some such models give new aggregate predictions, whereas others do not.

5 Concluding Remarks

Our main finding in this paper is that a low-dimensional object—the total capital stock and the value of the aggregate productivity shock—seems to be sufficient for characterizing the stochastic behavior of all the macroeconomic aggregates in the stochastic growth model with wealth heterogeneity and incomplete markets. We also show that many but not all models with substantial heterogeneity in wealth and income display aggregate behavior which is very similar to that displayed by their representative-agent counterparts. To explore the latter issue quantitatively we would need to pay more attention to the determination of the shape and movements of the wealth and income distributions. The analyses in Huggett (1994) and in Castañeda et al. (1994) show that straightforward extensions of standard macroeconomic models to allow a nontrivial wealth distribution give rise to substantially less skewness than what is observed in the data. Our experiments with preference heterogeneity indicate one route for generating more skewness in the wealth distribution, and we intend to explore thir route further in future work.

At this point one might ask in what kinds of economic environments the distribution of income and wealth and heterogeneity would play a more important role in affecting the aggregate economy. We suspect that one set of such environments would have heterogeneity/incomplete markets on the side of the firms, so that the allocation of production inputs is prevented from equalizing marginal rates of transformation across technologies/plants. Some existing studies indeed display this feature, e.g. the papers by Banerjee and Newman (1993) and Kiyotaki and Moore (1995). In the same spirit, one simple extension of the standard growth model which may give rise to stronger effects from the distribution of capital would have consumers who are isolated and hence forced to use their own production technologies (this is in contrast to our setup where the direct return to savings is equal for all agents). However, in order to make quantitative statements about the role of heterogeneity one would also want to restrict the extent and nature of the incompleteness of markets in an empirically reasonable way. At present, this discipline seems harder to work with on the firm side than on the side of consumers.

Finally, let us comment on the nature of our computational approximations. As we point out above, we are not able to provide bounds on how far our approximated equilibrium is from a true equilibrium. This is of course disturbing, but we do not find it more disturbing here than in any other numerical computation. Our computational algorithm builds on depriving the agents of access to more than I moments in their perceived rule for how the aggregate capital stock evolves, so for a given I this economy has an interpretation involving boundedly rational agents. This is not, however, how we wish to interpret our results; we wish to interpret our equilibrium as one where agents are fully rational, and the fact that agents could improve their utility ever so slightly in this equilibrium is no different than what it is in any other numerically calculated equilibrium. In any equilibrium of this kind agents are, strictly speaking, boundedly rational, but their deviations from full rationality are very, very small. At least to us, to argue that numerically calculated equilibria are unacceptable on the grounds that agents are only 99.999% rational as opposed to 100% rational seems extremely hard to justify on scientific grounds.

References

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

Aiyagari, S.R. & M. Gertler (1991), Asset Returns with Transactions Costs and Uninsured Individual Risk. *Journal of Monetary Economics* 27, 311–331.

Banerjee, A. and A. Newman (1993), Occupational Choice and the Process of Development. *Journal of Political Economy* 101, 274–298.

Bewley, T. (1977), The Permanent Income Hypothesis: A Theoretical Formulation. *Journal of Economic Theory* 16, 252–292.

Castañeda, A., J. Díaz-Giménez, and V. Ríos-Rull (1994), Unemployment Dynamics, Aggregate Fluctuations, and Income Distribution. Manuscript.

Chatterjee, S. (1994), Transitional Dynamics and the Distribution of Wealth in a Neoclassical Growth Model. *Journal of Public Economics*, 54, 97-119.

de Boor, C. (1978), *A Practical Guide to Splines*. Springer-Verlag, New York.

den Haan, W.J. (1994a), Heterogeneity, Aggregate Uncertainty, and the Short-Term Interest Rate. Manuscript.

den Haan, W.J. (1994b), Credit Conditions, Cross-Sectional Dispersion, and ARCH Effects in a Dynamic Equilibrium Model. Manuscript.

Díaz-Giménez, J., E.C. Prescott, T. Fitzgerald, and F. Alvarez (1992), Banking in Computable General Equilibrium Economies. *Journal of Economic Dynamics and Control* 16, 533–559.

Díaz-Giménez, J. and J.-V. Ríos-Rull (1991), Heterogeneous Agents, Capital Accumulation, and Business Cycles. Manuscript.

Heaton, J. and D.J. Lucas (1994a), Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing. Manuscript.

Heaton, J. and D.J. Lucas (1994b), The Importance of Investor Heterogeneity and Financial Market Imperfections for the Behavior of Asset Prices. Manuscript.

Huggett, M. (1993), The Risk-Free Rate in Heterogeneous-Agents, Incomplete Markets Economies. *Journal of Economic Dynamics and Control* 17, 953–969.

Huggett, M. (1994), Wealth Distribution in Life Cycle Economies. Manuscript.

İmrohoroğlu, A. (1989), The Cost of Business Cycles with Indivisibilities and Liquidity Constraints. *Journal of Political Economy* 97, 1364–83.

İmrohoroğlu, A. (1992), The Welfare Cost of Inflation under Imperfect Insurance. *Journal of*

Economic Dynamics and Control 16, 79–91.

Johnson S.A., J.R. Stedinger, C.A. Shoemaker, Y. Li, and J.A. Tejada-Guibert (1993), Numerical Solution of Continuous-State Dynamic Programs Using Linear and Spline Interpolation. *Operations Research* 41, 484–500.

Kiyotaki, N. and J.H. Moore (1995), Credit Cycles. Manuscript.

Krusell, P. & J.-V. Ríos-Rull (1994), What Constitutions Promote Capital Accumulation? A Political-Economy Approach. Manuscript.

Krusell, P. & A.A. Smith, Jr. (1994), Rules of Thumb in Macroeconomic Equilibrium: A Quantitative Analysis. *Journal of Economic Dynamics and Control* (forthcoming).

Kydland, F.E. and E.C. Prescott (1982), Time to Build and Aggregate Fluctuations. *Econometrica* 50, 1345–70.

Long, J.B. & C. Plosser (1983), Real Business Cycles. *Journal of Political Economy* 91, 1345–1370.

Lucas, D.J. (1994), Asset Pricing with Undiversifiable Income Risk and Short Sales Constraints: Deepening the Equity Premium Puzzle. *Journal of Monetary Economics* (forthcoming).

Marcet, A. and Singleton, K. (1991), Equilibrium Asset Prices and Savings of Heterogeneous Agents in the Presence of Incomplete Markets and Portfolio Constraints. Manuscript.

Press, W.H., B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling (1993), *Numerical Recipes*. Cambridge University Press, Cambridge.

Ríos-Rull, J.-V. (1992), Life-Cycle Economies and Aggregate Fluctuations. Manuscript.

Telmer, C.I. (1993), Asset Pricing Puzzles and Incomplete Markets. *Journal of Finance* 48, 1803–1832.

APPENDIX

Here we give a description of the numerical techniques used to solve the consumer's dynamic programming problem. The algorithm is similar to one used in Johnson et al. (1993). The description here assumes that we are solving the baseline economy. In addition, it assumes that only one moment of the capital distribution (i.e. $\bar{k}$) is included in the law of motion for aggregate capital. It is straightforward to modify the algorithm so as to accomodate a different model specification and/or additional moments.

The objective of the numerical algorithm is to approximate the four functions $v(k,1;\bar{k},z_g)$, $v(k,1;\bar{k},z_b)$, $v(k,0;\bar{k},z_g)$, and $v(k,0;\bar{k},z_b)$. We accomplish this task by approximating the values of the functions on a coarse grid of points in the $(k,\bar{k})$ plane and then using cubic spline and polynomial interpolation to calculate the values of these functions at points not on the grid. The numerical algorithm is in many ways analogous to value function iteration, except in that we do not restrict choices for capital to points on the grid.

The following steps describe the numerical procedure. First, choose a grid of points in the $(k,\bar{k})$ plane (we give some details below about how we choose these points). Second, choose initial values for each of the four functions at each of the grid points. (It is generally feasible to use the zero function as the initial condition for each of the functions.) Third, for each of the four (z,ϵ) pairs, maximize the right-hand side of Bellman's equation at each point in the grid. In this maximization, we allow the agent to select any value for capital. We use various interpolation schemes to compute the value function at points not on the grid (we describe the interpolation schemes in greater detail below). For large values of k (i.e. values for which the borrowing constraint does not bind), we use a Newton-Raphson procedure for finding the optimal choice of capital. For small values of k (i.e. for values close to the borrowing constraint), we use a bisection procedure to map out the objective function. This procedure allows for the possibility that the borrowing constraint binds (in which case the optimal value of capital is at a corner, so that the first derivative of the objective function is not zero at the optimum). Fourth, compare the new optimal values generated by the third step to the original values. If the new values are close to the old values, then stop; otherwise, repeat step three until the new and old values are sufficiently close. (In practice, value functions typically converge more slowly than the decision rules associated with these value functions. Thus it is generally more efficient to stop the iterations when the optimal decisions at each of the grid points stop changing, even if the value functions have not yet fully converged.)

We now comment on the choice of a grid in the $(k,\bar{k})$ plane and on the interpolation schemes that we use. Since there is generally not much curvature in the value function in the $\bar{k}$-direction, we use a small number of grid points in this direction and we use polynomial interpolation to compute the value function for values of $\bar{k}$ not on the grid. If there are m points in the $\bar{k}$-direction, then polynomial interpolation fits a polynomial of order $m-1$ to the function values at these points (so that the polynomial fits the values exactly), and then uses this polynomial to compute the value function in between grid points. We compute the value of the interpolating polynomial using Neville's algorithm, as described in Chapter 3 of Press et al. (1989). This algorithm avoids the numerical instabilities associated with computing the coefficients of the interpolating polynomial.

We generally use 4 to 7 equally spaced points in the $\bar{k}$ direction.

In the k direction, there is generally a fair amount of curvature in the value function, especially for values of k near the borrowing constraint and for high degrees of risk aversion. In this direction, therefore, we use cubic spline interpolation, which fits a piecewise cubic function through the given function values, with one piece for each interval defined by the grid. This piecewise cubic function satisfies the following restrictions: (1) it matches the function values exactly at the grid points and (2) its first and second derivatives are continuous at the grid points. This smoothness facilitates the use of fast Newton-Raphson algorithms for computing the optimal decision at each point in the grid. If there are m points in the $\bar{k}$ direction, then we need to calculate $8m$ cubic splines in the k direction (4 functions times m values of $\bar{k}$ times 2 values of z). Cubic splines can be computed efficiently by solving a set of tridiagonal linear equations (see, for example, the description in Chapter 3 of Press et al. (1989) or Chapter 4 of de Boor (1978)). Note that the required cubic splines need only be computed once for each iteration of the algorithm; once computed, it is easy to use these splines to calculate interpolated values. We generally use 70 to 130 grid points in the k direction, with many grid points near 0 (where there is a lot of curvature) and fewer grid points for larger values of k (where there is less curvature). We find that our results are not sensitive to increasing the number of grid points in either the k or $\bar{k}$ directions.

In order to simulate the behavior of agents, we need to approximate the decision rules associated with the approximate value function as computed above. (Since these decision rules in general need to be evaluated at many different values of $\bar{k}$ in the course of simulating the dynamic behavior of the economy, it is not efficient to use the interpolation scheme described above to compute optimal decisions at points not on the grid.) We approximate the decision rules by first computing optimal decisions on a fine grid of points in the $(k, \bar{k})$ plane for each value of (z, ϵ). When computing these optimal decisions, we use the approximate value function as computed above. For the purpose of approximating decision rules, we generally use 150 to 600 equally spaced points in the k direction and 25 to 100 equally spaced points in the $\bar{k}$ direction. Optimal decisions at points not on the grid are then computed using bilinear interpolation (see Chapter 3 of Press et al. (1989)).

Figure 1: Tomorrow's vs. Today's Aggregate Capital
(Baseline model)

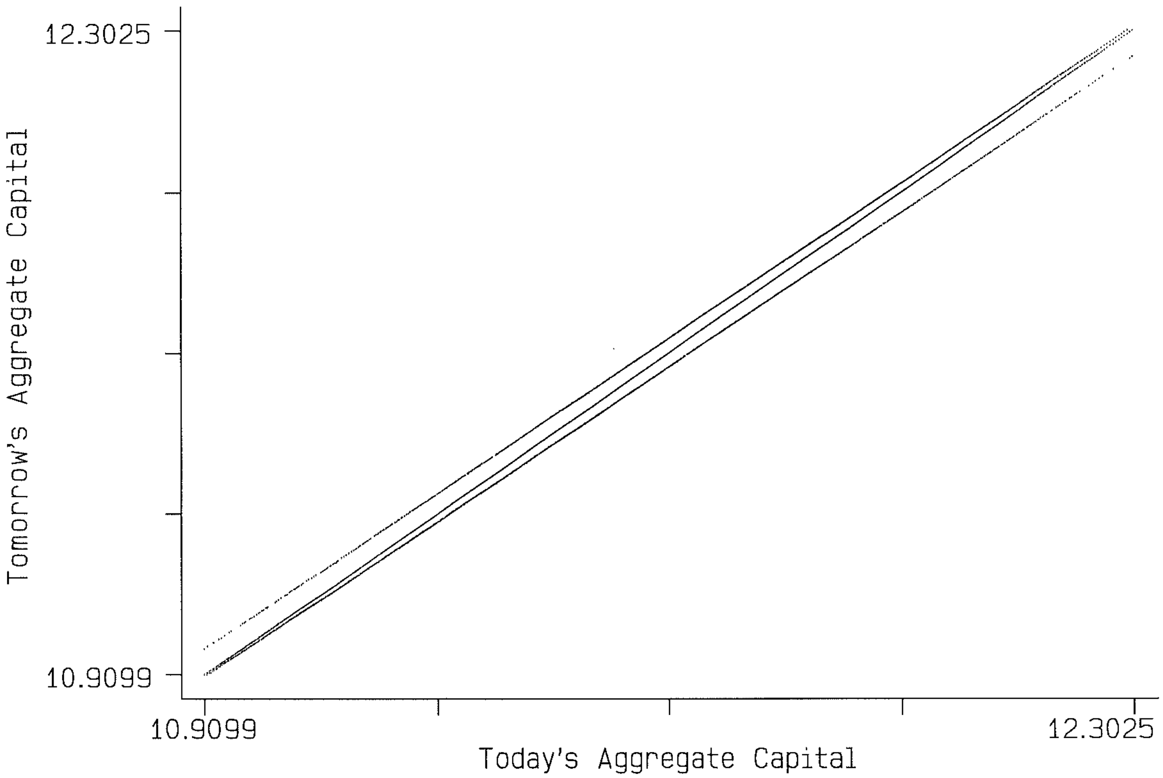

Figure 2: An Individual Agent's Decision Rules
(Baseline model, aggregate capital=11.7, good aggregate state)

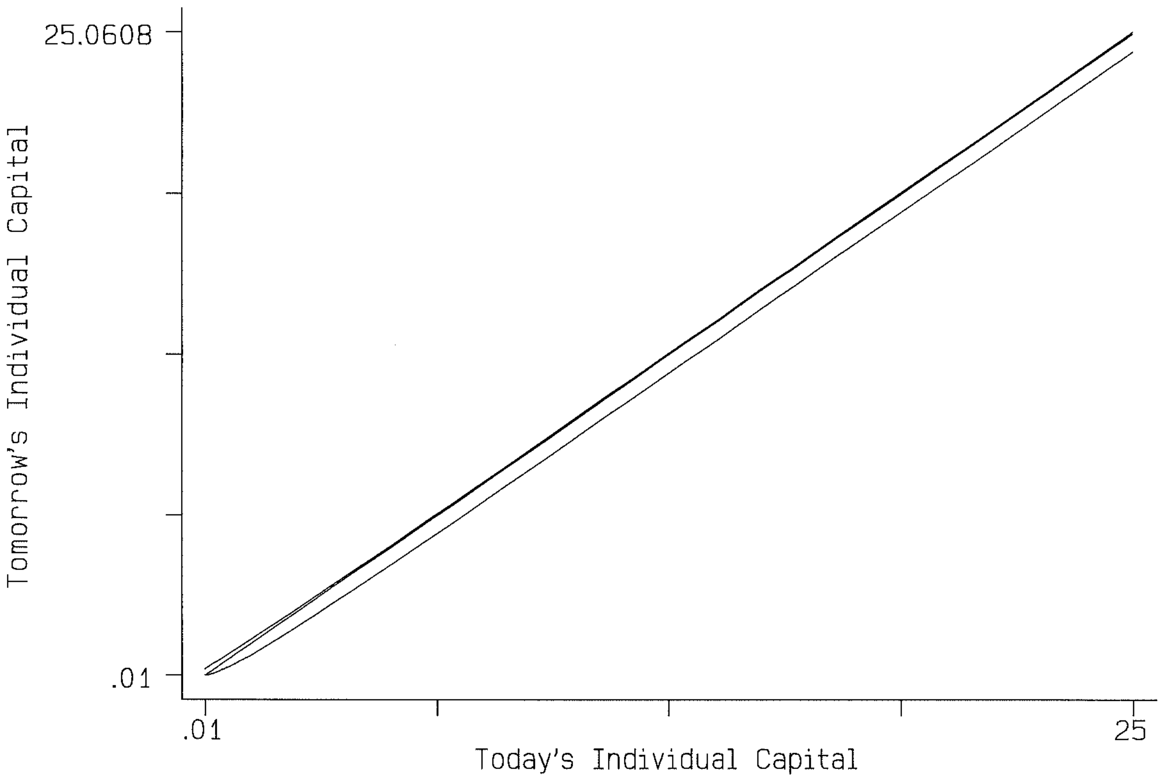

Figure 3: Capital Stock Distribution
(Baseline model, aggregate capital=11.61, good aggregate state)

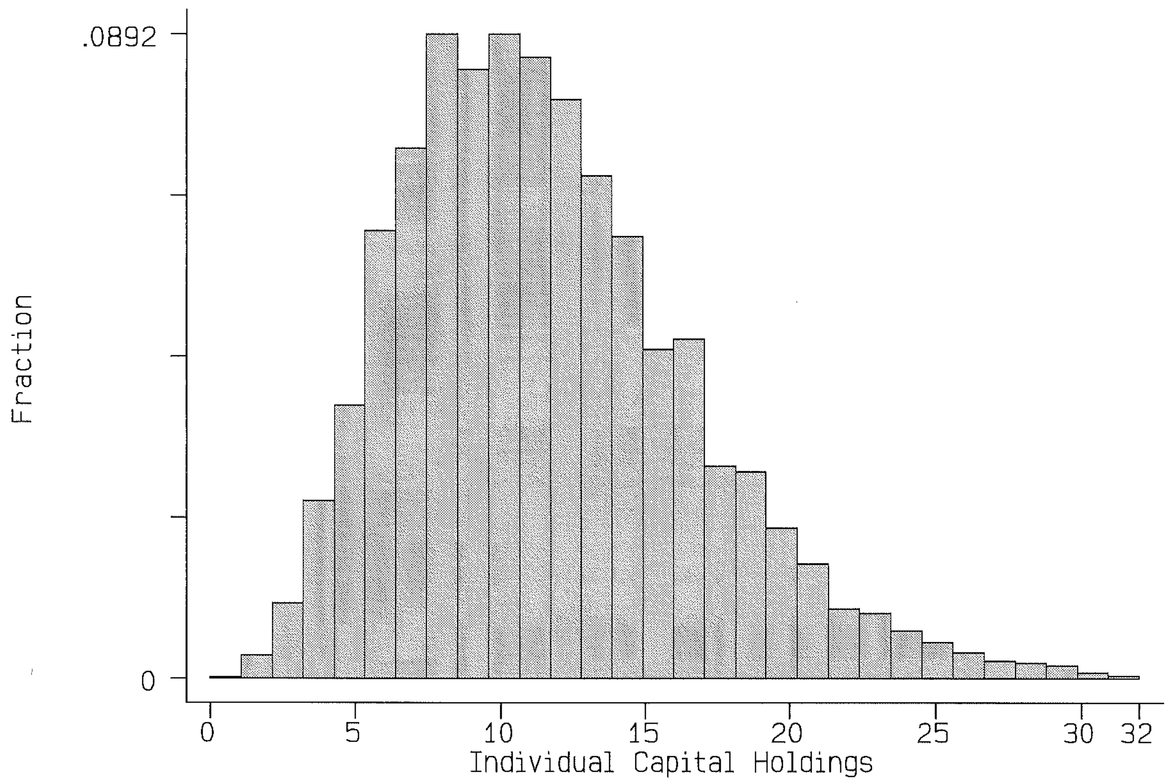

Figure 4: Std. Dev. vs. Mean of Capital Stock Distribution
(Baseline model)

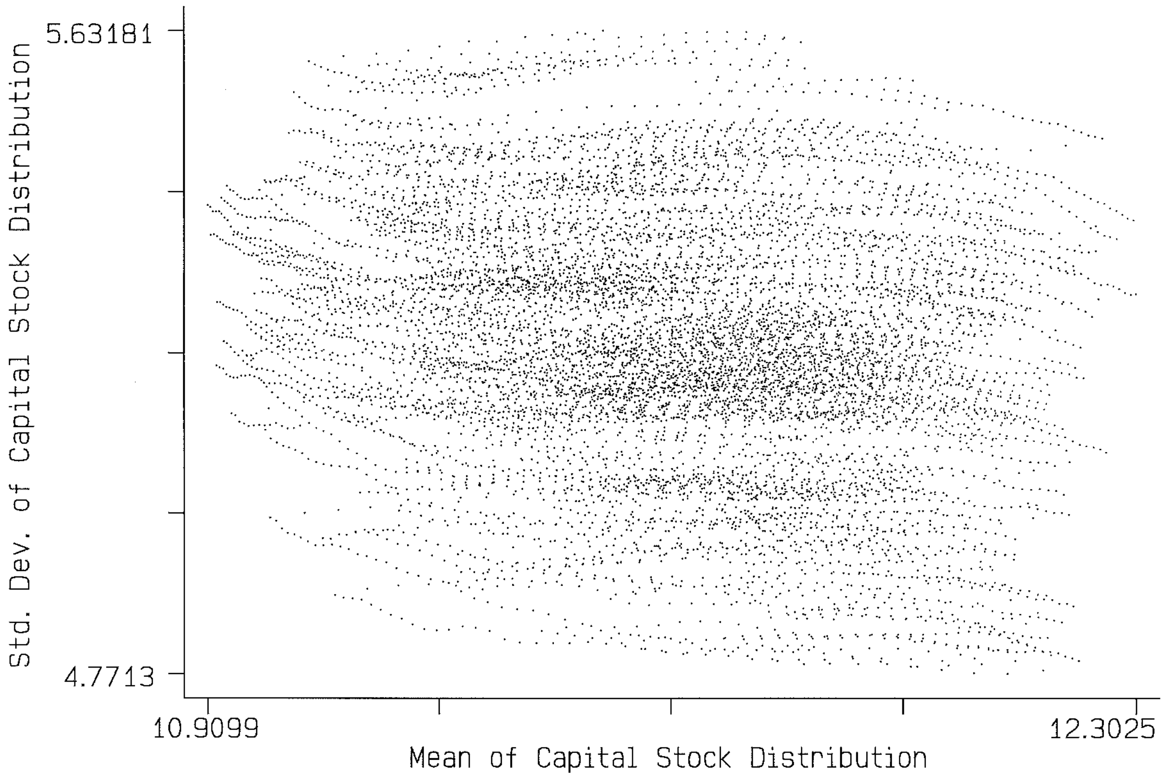

Figure 5: Capital Stock Distribution (Patient Agents)
(Two-agent model, aggregate capital=11.57, good aggregate state)

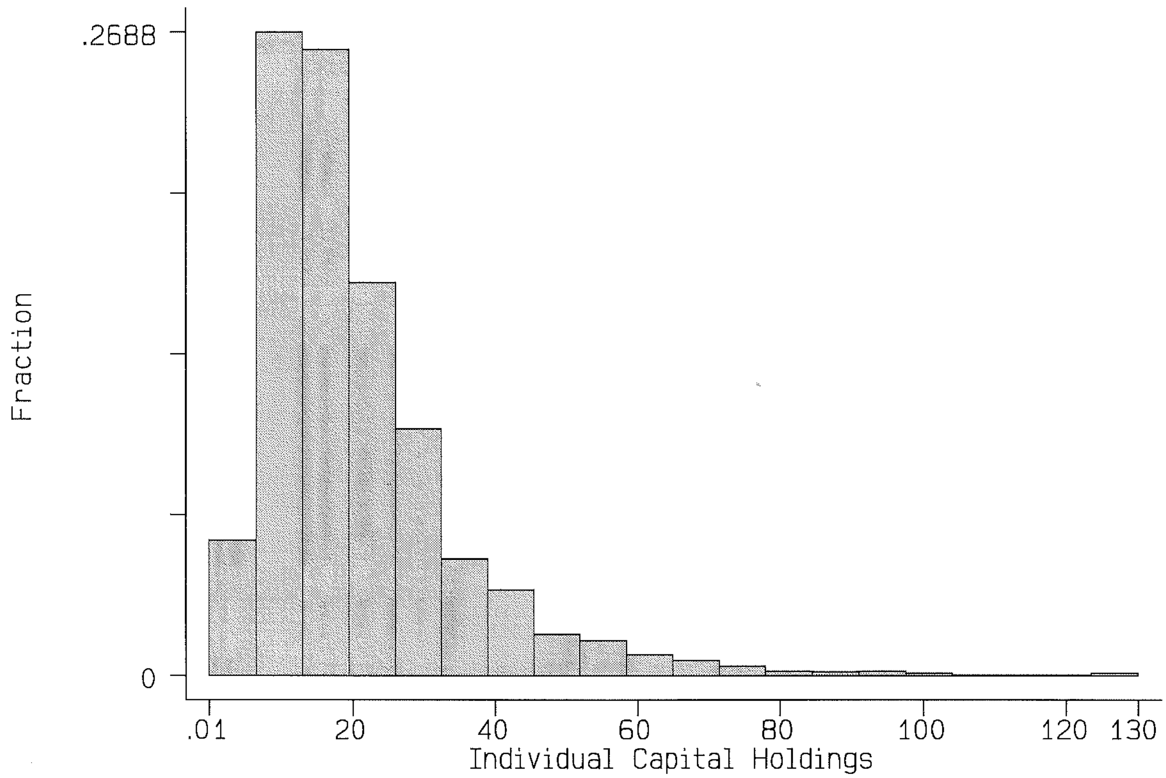

Figure 6: Capital Stock Distribution (Impatient Agents)
(Two-agent model, aggregate capital=11.57, good aggregate state)

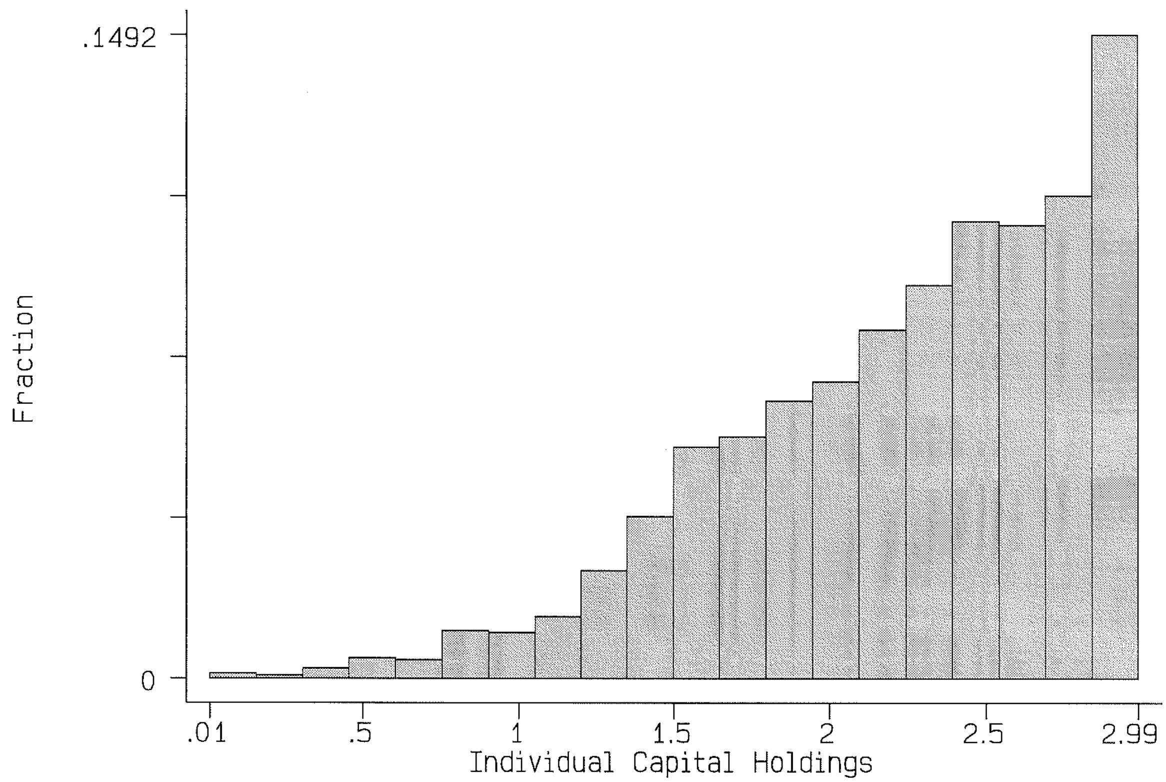

Figure 7: Tomorrow's vs. Today's Aggregate Capital
(Model with valued leisure)

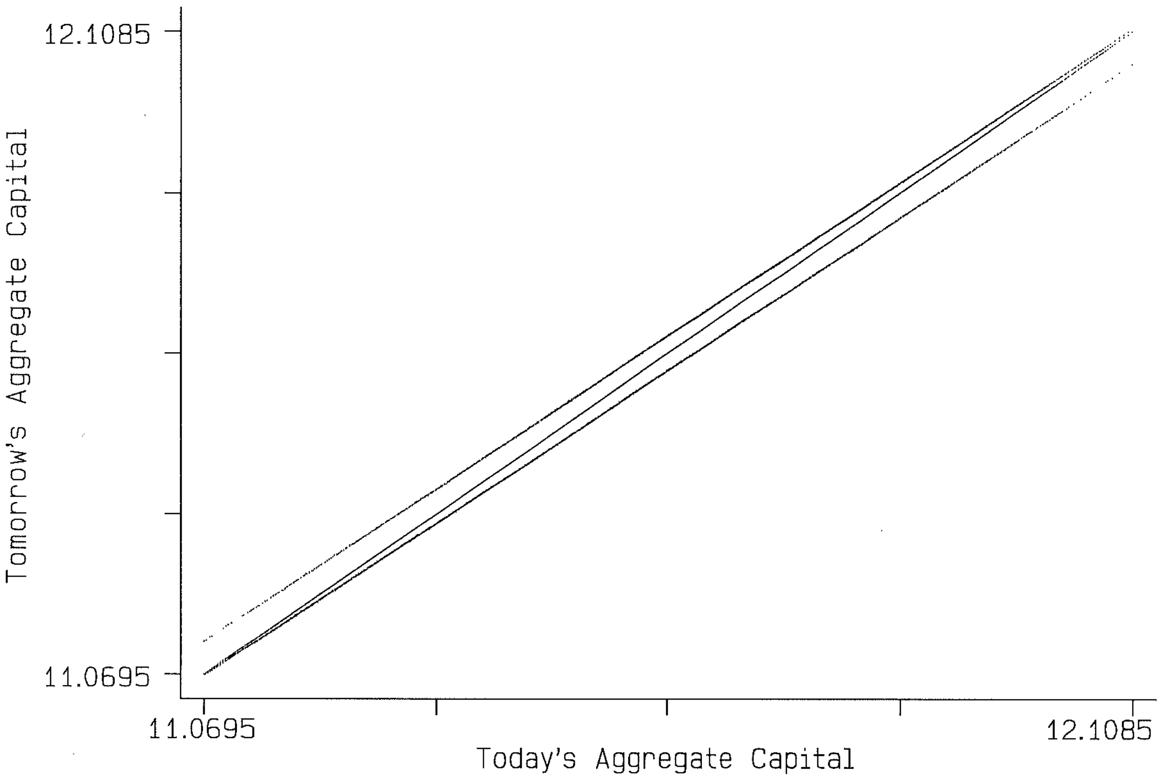

Figure 8: Aggregate Labor vs. Aggregate Capital
(Model with valued leisure)