

Income and Wealth Heterogeneity in the Macroeconomy

Per Krusell

University of Rochester, Centre for Economic Policy Research, and Institute for International Economic Studies

Anthony A. Smith, Jr.

Carnegie Mellon University

How do movements in the distribution of income and wealth affect the macroeconomy? We analyze this question using 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 the macroeconomic 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. Our benchmark model, whose only difference from the representative-agent framework is the existence of uninsurable idiosyncratic risk, displays far less cross-sectional dispersion

We thank conference participants at Macroeconomics with “Frictions,” held at the Federal Reserve Bank of Minneapolis; the National Bureau of Economic Research 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; the North American summer meetings of the Econometric Society; and the Stanford Institute for Theoretical Economics Summer Workshop and seminar participants at Carnegie Mellon, Columbia, Institute for International Economic Studies, Pompeu Fabra, Southern Methodist University, Texas at Austin, Wharton, Yale, and York. We also benefited particularly from comments by Wouter den Haan, Mark Huggett, Robert Lucas, Víctor Ríos-Rull, Tom Sargent, José Scheinkman, Chris Telmer, Stan Zin, and two anonymous referees. 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.

[*Journal of Political Economy*, 1998, vol. 106, no. 5]

© 1998 by The University of Chicago. All rights reserved. 0022-3808/98/0605-0003\$02.50

and skewness in wealth than U.S. data. However, an extension that relies on a small amount of heterogeneity in thrift does succeed in replicating the key features of the wealth data. Furthermore, this extension features aggregate time series that depart significantly from permanent income behavior.

I. Introduction

Most of dynamic general equilibrium macroeconomic theory relies heavily on the representative-agent abstraction: it is assumed that the economy behaves “as if” it is inhabited by a single (type of) consumer. At the same time, this macroeconomic theorizing attempts to take microfoundations seriously; the parameters of the models are typically selected on the basis of existing empirical and theoretical knowledge, and the models are then used 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 would be a reasonable modeling strategy. First, it is possible that the theoretical assumptions needed to justify the use of a representative consumer are roughly met in the data. This view, however, is hard to defend; 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 analysis 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 in which there is a large number of consumers who, in addition to uncertain aggregate productivity, face idiosyncratic income (employment) shocks. Following Bewley (1977), Scheinkman and Weiss (1986), and others, we assume that consumers cannot insure directly against these shocks, but that they can buy and sell an asset subject to an exogenous lower bound on the asset holding. In our macroeconomic framework, as in Aiyagari (1994), this asset is aggregate capital. Savings can thus be precautionary and allow partial insurance against the idiosyncratic shocks. Because of the lack of full insurance, this model generates an endogenous distribution of wealth across consumers. An important problem, therefore, is how to characterize the interaction between the distribution of wealth and the macroeconomic aggregates.

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. One of the contributions of our work is to show how equilibria can be approximated numerically, despite the fact that the state of the economy at any point in time is an infinite-dimensional object (we assume a continuum of agents). This methodological contribution opens the possibility of characterizing a large class of new macroeconomic models in which heterogeneity in income and wealth plays a key role. For example, this class of models allows a much richer analysis of the interactions between business cycles and inequality than existing frameworks do.

Our main insight is that the macroeconomic model with heterogeneity features *approximate* aggregation. By approximate aggregation, we mean that, in 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. Therefore, the consumers in our equilibrium face manageable prediction problems since the distribution of aggregate wealth is almost completely irrelevant for how the aggregates behave in the equilibrium. Furthermore, this finding is remarkably robust to changes in both parameter values and the specification of the model.

When the representative-agent model is altered only by adding idiosyncratic, uninsurable risk, the resulting stationary wealth distribution is quite unrealistic: there are too few very poor agents, and much too little concentration of wealth among the very richest. For this reason, we consider a version of the model with preference heterogeneity: agents have random discount factors, whose values have a symmetric distribution with a small variance and whose transition probabilities are such that the average duration, or life length, of a discount factor equals that of a generation. In this fashion, we incorporate genetic differences in the population that are passed on imperfectly from parents to children. We show that this model does succeed quite well in matching the key features of the wealth distribution.

The model with preference heterogeneity also gives rise to interesting aggregate time-series behavior. In this model, although aggre-

gate wealth is mainly in the hands of the rich, poor agents have a large influence on aggregate consumption. Since these agents are also impatient on average, they can be characterized as “hand-to-mouth” consumers. Thus, in the aggregate, we observe a significant departure from permanent income behavior, in contrast to standard representative-agent models.

The explanation for our main result—the collapse of the state space—builds on the properties of optimal savings behavior in our class of models. The key insight is related to earlier findings from similar models that utility costs from fluctuations in consumption are quite small and that self-insurance with only one asset is quite effective.¹ Self-insurance in our model is not very effective in terms of smoothing individual relative to aggregate consumption; for example, the unconditional standard deviation of individual consumption is about four times that of aggregate consumption, and the unconditional correlation of the consumption of any two agents is very close to zero. However, in utility terms, agents in our stationary equilibria are insured well enough that the marginal propensity to save out of current wealth is almost completely independent of the levels of wealth and labor income, except at the very lowest levels of wealth. Furthermore, although some very poor agents have substantially different marginal savings propensities at any point in time, the fraction of total wealth held by these agents is always very small (this is particularly true in the model with a realistic wealth distribution). Because it is so small, higher-order moments of the wealth distribution simply do not affect the accumulation pattern of total capital, even though these moments do move significantly over time.

Our computational algorithm 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 in which, by assumption, agents have a limited ability to predict the evolution of this distribution. We then show that this bound on ability almost does not constrain the agents at all. More precisely, we compute approximate equilibria by postulating that the law of motion perceived by agents can be described by a stochastic process for a finite-dimensional vector of moments of the wealth distribution. For any given vector of moments \mathbf{m} , a candidate approximate equilibrium is a fixed point in a class \mathcal{S} of (possibly non-

¹ See, e.g., Robert Lucas (1987), Cochrane (1989), and Krusell and Smith (1996*b*) and the incomplete-markets asset pricing literature as represented by Marcet and Singleton (1991), Telmer (1993), Deborah Lucas (1994), Heaton and Lucas (1995, 1996), den Haan (1996*a*, 1996*b*), and Krusell and Smith (1997).

linear) first-order Markov processes. A given Markov process $S \in \mathcal{S}$ is a fixed point if (1) agents' decision rules derive from dynamic maximization problems in which the behavior of the aggregate state is described by S and (2) S is the best approximation in the class \mathcal{S} to the dynamic behavior of \mathbf{m} implied by the aggregated decision rules of agents. 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 S . We use a measure of forecasting accuracy in the algorithm to decide whether to increase the agents' ability to make forecasts—increase the dimension of \mathbf{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 \mathbf{m} consists only of the mean of the distribution of capital, we obtain extremely high forecasting accuracy. 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.

In Sections II and III we describe the benchmark model, the computational strategy, and the main result. Section IV discusses the wealth distribution data and presents the model with preference heterogeneity. That section also describes the aggregate time-series statistics from the various models and makes comparisons with representative-agent models. Section V concludes with some remarks.

II. 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 benchmark setup there is only one type of consumer (i.e., all consumers have the same preferences), and the setup is one that in all other ways is like the standard stochastic growth model. Later, in Section IV, we extend the benchmark model to include preference heterogeneity.

A. *The Environment*

We consider a version of the stochastic growth model with a large (measure one) population of infinitely lived consumers. There is only one good per period, and we assume that the preferences over streams of consumption of each agent are given by

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

with

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

Production of the good, y , is a Cobb-Douglas function of capital input, k , and labor input l : $y = zk^\alpha l^{1-\alpha}$, with $\alpha \in [0, 1]$. The 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.

Each agent is endowed with one unit of time, which gives rise to $\epsilon \tilde{l}$ units of labor input, where ϵ is stochastic and can take on the value zero or one. When $\epsilon = 1$, we think of the agent as employed and supplying \tilde{l} units of labor input; when $\epsilon = 0$, we think of him as unemployed. There is also a stochastic shock to aggregate productivity, which we denote z . There are two possible aggregate 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 probabilities $\pi_{ss'}$: 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. By virtue of the law of large numbers, the only exogenous source of aggregate uncertainty in the economy is the aggregate productivity shock. 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, when one controls for z , individual shocks are uncorrelated. We use $\pi_{ss'\epsilon\epsilon'}$ to denote the probability of transition from state (z_s, ϵ) today to $(z_{s'}, \epsilon')$ tomorrow. The transition probabilities have 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') .

B. *The Market Arrangement*

For the economic environment described above, 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 without knowing how wealth is distributed across consumers.² Here, however, we assume that there are incomplete markets: there is only one asset—capital. This asset plays the twin roles of being a store of value for the individual agent and a means of self-insurance against the income shocks.³ Thus let k denote the holdings of capital. In order to rule out Ponzi schemes and to guarantee that loans are paid back, we restrict capital holdings to satisfy $k \in \mathcal{K} \equiv [0, \infty)$.⁴ We refer to the lower bound on capital as the borrowing constraint.

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 $w(\bar{k}, \bar{l}, z) = (1 - \alpha)z(\bar{k}/\bar{l})^\alpha$ and $r(\bar{k}, \bar{l}, z) = \alpha z(\bar{k}/\bar{l})^{\alpha-1}$, 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 that concerns z is exogenous; it can be described by z 's transition matrix. The part that 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

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

² See Chatterjee (1994) or Krusell and Ríos-Rull (1997) for a discussion.

³ In this respect, our approach parallels those adopted in a number of recent papers, including İmrohoroğlu (1989, 1992), Aiyagari and Gertler (1991), Díaz-Giménez et al. (1992), Huggett (1993), and Aiyagari (1994).

⁴ One could also consider a negative lower bound on holdings of capital. As Aiyagari (1994) shows in a model without aggregate productivity shocks, if the lowest individual income realization is zero, then a constraint to always pay back implies that a lower bound for capital that is less than zero is equivalent to one that is zero. In this sense, our lower bound is generous. Later, we also consider positive lower bounds.

subject to

$$\begin{aligned} c + k' &= r(\bar{k}, \bar{l}, z)k + w(\bar{k}, \bar{l}, z)\bar{l}\epsilon + (1 - \delta)k, \\ \Gamma' &= H(\Gamma, z, z'), \\ k' &\geq 0 \end{aligned}$$

and the stochastic laws of motion for z and ϵ . The decision rule for the updating of capital implied by 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 , that is, the appropriate summing up of agents' optimal choices of capital given their current status in terms of wealth and employment.

C. 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 endogenous state variable of the economy, Γ , is a high-dimensional object. It is well known that numerical solution of dynamic programming problems becomes increasingly difficult as the size of the state space increases. Our way of dealing with this problem is to assume that agents are boundedly rational in their perceptions of how Γ evolves over time and to increase the sophistication of these perceptions until the errors that agents make because they are not fully rational become negligible.

Therefore, suppose that agents do not perceive current or future prices as depending on anything more than the first I moments of Γ (in addition to z); denote these moments as $\mathbf{m} \equiv (\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_I)$.⁶ Since current prices depend only on the total amount of capital and not on its distribution, limiting agents to a finite set of moments is restrictive only as far as future prices are concerned. In particular, to know future prices, it is necessary to know how the total capital stock evolves. Since savings decisions do not aggregate,

⁵ Our algorithm bears some similarities to one proposed in Díaz-Giménez and Ríos-Rull (1991).

⁶ By \mathbf{m}_1 we are referring to a 2×1 vector consisting of the first moments of Γ ; \mathbf{m}_2 is a 2×2 matrix consisting of the second moments of Γ , and so on.

the total capital stock in the future is a nontrivial function of all the moments of the current distribution.

Agents thus perceive the law of motion for \mathbf{m} to be given by a function H_t that belongs to a class \mathcal{S} . Each of these functions expresses \mathbf{m}' , that is, the vector of I moments in the next period, as a function of the I current moments: $\mathbf{m}' = H_t(\mathbf{m}, z, z')$. Given the law of motion H_t , each agent's optimal savings decision can then be represented by a decision rule f_t . Given such a decision rule f_t for individuals and an initial wealth and labor shock distribution, it is possible to derive the implied aggregate behavior—a time-series path of the distribution of income and wealth—by simulating the behavior of a large number of consumers. The resulting distributions, therefore, are restricted only by initial conditions, shocks, and the decision rules of agents. Moreover, they can be used to compare the simulated evolution of the specific vector of moments \mathbf{m} to the perceived law of motion for \mathbf{m} on which agents base their behavior. Our approximate equilibrium is a function H_t that, when taken as given by the agents, (i) yields the best fit within the class \mathcal{S} to the behavior of \mathbf{m} in the simulated data and (ii) yields a fit that is close to perfect in the sense that H_t tracks the behavior of \mathbf{m} in the simulated data almost exactly, that is, 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.

Our algorithm amounts to the following iterative procedure: (1) Select I . (2) Guess on a parameterized functional form for H_t and on the parameters of this function. (3) Solve the consumer's problem given H_t . This step, which builds on a nonlinear approximation of the value function, is described in more detail in the Appendix. (4) Use consumers' 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 measure of 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_t .⁸

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

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

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

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

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

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

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

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

$$k' \geq 0$$

and the law of motion for (z, ϵ) . We thus obtain a (nonlinear) decision rule $k' = f_I(k, \epsilon; \bar{k}, z)$ that, when simulated, allows us to compare the aggregate behavior of the moments with their description H_I . The idea is thus to find a fixed point for H_I in the form of a vector $(a_0^*, a_1^*, b_0^*, b_1^*)$. If this H_I is satisfactorily reproduced in simulations (i.e., if the goodness of fit is high), stop. Otherwise, consider a more flexible functional form for H_I or add another moment.

III. Results

We shall now show our approximate aggregation result for the benchmark model. Our approach is to select parameter values that are in line with those used in similar studies (which in turn are based on microeconomic data or long-run model considerations) and then to examine whether the results are robust to changes in the parameter values. We comment on the robustness analysis at the end of this section.

A. Model Parameters for the Benchmark Setup

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 set the shock values to $z_g = 1.01$ and $z_b = 0.99$ and the unemployment rates to $u_g = 0.04$ and $u_b = 0.1$, implying that the

describing the distribution, such as tail probabilities, which are themselves nonlinear functions of the distribution's moments.

fluctuations in the macroeconomic aggregates have roughly the same magnitude as the fluctuations in observed postwar U.S. time series. The process for (z, ϵ) is chosen so that the average duration of both good and bad times is eight quarters and so that the average duration of an unemployment spell is 1.5 quarters in good times and 2.5 quarters in bad times. We also impose $\pi_{gb00}\pi_{gb}^{-1} = 1.25\pi_{bb00}\pi_{bb}^{-1}$ and $\pi_{bg00}\pi_{bg}^{-1} = 0.75\pi_{gg00}\pi_{gg}^{-1}$.⁹

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; 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.

C. Equilibrium Properties: Only the Mean Matters

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

$$\log \bar{k}' = 0.095 + 0.962 \log \bar{k}; \quad R^2 = .999998, \quad \hat{\sigma} = 0.0028\%,$$

in good times and

$$\log \bar{k}' = 0.085 + 0.965 \log \bar{k}; \quad R^2 = .999998, \quad \hat{\sigma} = 0.0036\%$$

in bad times.¹⁰ There are two measures of fit: R^2 and the standard deviation (percent) of the regression error, $\hat{\sigma}$. Using our simulated sample (consisting of 10,000 observations), we plot tomorrow's ag-

⁹ Our labor income process is similar to that in İmrohoroğlu (1989).

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




FIG. 1.—Tomorrow's vs. today's aggregate capital (benchmark model)

gregate capital against today's aggregate capital. This graph (see fig. 1) is a clear illustration of the high R^2 and the low $\hat{\sigma}$.¹¹

In terms of R^2 's and $\hat{\sigma}$'s, these fits are extremely good. Thus an agent perceiving these simple laws of motion for aggregate capital makes extremely small mistakes compared to using the correct law of motion. In terms of errors in forecasting future prices due to the omission of higher moments, we calculate that forecasts of prices 25 years ahead (given the sequence of future aggregate productivity shocks) have *maximum* errors of less than 0.1 percent. In other words, although the inclusion of more moments in these predictions must significantly improve forecasts in a statistical sense (since strict aggregation does not obtain), these improvements are minuscule in quantitative terms.¹² Moreover, from a welfare perspective, superior prediction techniques would only lead to vanishingly small increases in the agent's utility (expressed, say, as a uniform percentage increase in consumption across all dates and states). In this sense, then, agents are very close to optimal behavior, which is what equilibrium dictates.

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

¹² For details on these experiments and others, see the working paper version of the present paper (Krusell and Smith 1996a).




FIG. 2.—An individual agent's decision rules (benchmark model, aggregate capital = 11.7, good aggregate state).

Like most numerical procedures, the present one does not provide bounds on how far the approximate equilibrium deviates from an exact equilibrium. In particular, one might imagine that there are self-fulfilling approximate equilibria: because agents perceive a simple law of motion, they behave accordingly. However, there is nothing in the theoretical link between these perceptions and the aggregate savings behavior that is suggestive of self-fulfilling equilibria. Nevertheless, we did check whether more sophisticated perceptions tend to change significantly the equilibrium properties. When we include a higher moment—we tried various dispersion measures for capital—in the perceived law of motion, our approximate equilibrium fixed point remains virtually identical to the simpler case in terms of both the aggregate processes and individual behavior.

D. *Why Is Only the Mean Important?*

In this economy, aggregation would obtain if all agents had the same propensity to save out of wealth, because then changes in the distribution of the total stock of capital would have no aggregate effects. Figure 2 depicts the individual decision rule for savings—the amount of capital carried into the next period as a function of to-

day's capital, given values for today's individual and aggregate shocks and for aggregate capital.¹³

The figure reveals that the marginal propensities to consume are almost identical for agents with different employment states and levels of capital.¹⁴ As the agent's wealth increases, the slope increases toward one; a slope of one would amount to exact permanent income behavior, as discussed in Bewley (1977).

In the stationary state, the distribution of capital does move around significantly: for example, the standard deviation, skewness, and kurtosis of the capital stock distribution all display substantial variation in the simulated aggregate time series. Most of the capital, however, is held by agents with essentially the same savings propensity. Very few agents—the very poorest ones—have a much lower propensity, and the capital that they hold is negligible. For these reasons, aggregation is almost perfect. The insight that different savings propensities are associated with very low levels of wealth is an important one. As we shall see below, especially in Section IV, it is possible to construct models in which a much larger number of agents have low propensities. However, since these agents are also very poor in relative terms, they are not important in the aggregate.

Why, then, are marginal savings propensities almost independent of wealth in the benchmark model? In the class of models that we consider here, the utility costs from accepting fluctuations in consumption are very small, even when these fluctuations are several times larger than for aggregate consumption.¹⁵ Moreover, the access to one aggregate asset is sufficient for providing the agent with very good insurance in utility terms. These findings are consistent with the results in Robert Lucas (1987), Cochrane (1989), and Krusell and Smith (1996*b*) and in the incomplete-markets asset pricing literature; see, for example, Singleton (1991), Marcet and Telmer (1993), Deborah Lucas (1994), Heaton and Lucas (1995, 1996), den Haan (1996*a*, 1996*b*), and Krusell and Smith (1997). As a consequence, in the stationary state, most agents have enough capital that

¹³ The graph assumes a good aggregate state and a typical value of aggregate capital. The top line is the decision rule of an employed agent, and the bottom line is the decision rule of an unemployed agent. The middle line is the 45-degree line.

¹⁴ Conditional on a given amount of individual capital, there is a notable difference in the savings of employed and unemployed agents. This difference reflects the insurance role of capital. If an agent becomes unemployed, current labor income declines; and in order for consumption to be independent of the unemployment shock, the amount of capital carried into the next period would need to drop by the amount of the wage income loss. Here, since complete insurance cannot be obtained using capital alone, the actual drop in savings is smaller.

¹⁵ The standard deviation of individual consumption is about four times that of aggregate consumption in the model.

their savings behavior is guided mainly by intertemporal concerns rather than by insurance motives. The availability of enough capital is automatic here since our model has a neoclassical production function with high marginal returns to capital at low levels of capital and a realistic capital/output ratio.

E. Robustness

We performed a large number of sensitivity checks by changing parameter values within the context of the benchmark specification. The working paper version of this paper (Krusell and Smith 1996*a*) documents these experiments in detail. The basic finding from these experiments is that it is extremely difficult to find exceptions to the approximate aggregation result. For example, if individual shocks are more volatile or more persistent (alternatively, if borrowing is more restricted or if agents are more risk averse), the aggregate economy responds by accumulating just enough extra capital to provide most agents with a large enough buffer that the shocks do not hurt them much in utility terms.

Experiments with the discount rate deserve special mention: we observe that more impatience leads to lower propensities to consume on average and to more dispersion in marginal propensities. There is an explicit connection to Bewley's (1977) work here.¹⁶ Bewley considers a decision-theoretic framework with an agent facing income risk and a sure return to savings of unity. He shows that as the agent's discount factor approaches unity and as the agent's initial wealth grows larger, the agent's savings function becomes linear with a slope of unity—permanent income behavior. In our economy, if we were to let the discount rate and the depreciation rate (β , δ) approach (1, 0), the capital stock would increase to infinity and the gross rate of return on savings would become unity, thus placing the agent in a situation identical to the limit that Bewley considers. Our computations indeed show that higher discount rates strengthen (and lower discount rates weaken) the aggregation result. However, large decreases in β are necessary in order for the goodness of fit to significantly worsen; for example, the percentage standard deviation of the regression error is only about 10 times higher for a β as low as 0.67.

We also considered models with valued leisure, various forms of heterogeneity in preferences (risk aversion and patience), and fixed costs of adjusting capital, and our main finding holds up in these

¹⁶ We thank José Scheinkman for drawing our attention to the connection between our setup and that studied by Bewley.

extensions as well. The valued-leisure (“real business cycle”) extension is especially interesting since it complicates the determination of prices: wages and rental rates no longer depend only on aggregate capital and the aggregate shock, but also on the total work effort.¹⁷ With significant dependence of individual work effort on wealth, thus, aggregation might fail. However, it turns out that even in a formulation with large wealth effects, the relation between wealth and effort is almost linear for most agents. Thus our approximate aggregation result continues to hold in a model with valued leisure. The Appendix contains a more detailed description of this model.

IV. Matching the Wealth Distribution

Is the benchmark model of the last section a reasonable model of income and wealth heterogeneity? The labor income process is very simple, but it is calibrated so as to at least roughly match income variability due to employment variation. Heterogeneity in wealth, however, is entirely nontrivially determined. As it turns out, one problem with the present model is that it does not generate realistic wealth heterogeneity: the data display significantly more skewness for wealth than the model does.¹⁸ More precisely, too few agents hold low levels of wealth, and the concentration of wealth among the richest agents is far too small. On the basis of data in Wolff (1994) and Díaz-Giménez, Quadrini, and Ríos-Rull (1997), for example, the poorest 20 percent of the population have about zero wealth on average, whereas the richest 5 percent of the population hold roughly half of all the wealth. In contrast, the benchmark model predicts that the poorest 20 percent hold (on average) 9 percent of total wealth whereas the richest 5 percent hold (on average) 11 percent: there is significant skewness, but not nearly as much as in the data.

One of the main purposes of our line of research is to extend the standard macroeconomic framework to allow heterogeneity among consumers. Therefore, it seems important to make sure that the heterogeneity in the new framework is quantitatively adequate. In this section we construct one model that roughly matches the observed income and wealth distributions. We show that the approximate aggregation result still obtains in this model, and we go on to make a

¹⁷ In separate work (Krusell and Smith 1997) we consider an extension with another nontrivial market: a market for a riskless bond. Approximate aggregation also obtains in that setup.

¹⁸ This is also true for similar frameworks that have richer processes for labor income than we do (see, e.g., Aiyagari 1994; Huggett 1996; Castañeda, Díaz-Giménez, and Ríos-Rull, in press): for a given, realistic, labor income distribution, there is far too little skewness in wealth.

few remarks about how this model compares to the standard representative-agent framework.

Amending the model to generate a large group of poor agents is straightforward. For example, one can follow Hubbard, Skinner, and Zeldes (1995) in assuming realistic settings for taxes, subsidies, and social insurance at low levels of income, which they show implies that low-income agents have incentives not to save. In addition, equilibrium frameworks with overlapping generations of consumers who are not altruistically linked, as studied in Huggett (1996), tend to generate stationary wealth distributions with more agents close to zero wealth. We follow the former of these routes, although in a stylistic fashion: we assume that unemployed agents receive income too. We set their income to a number that is about 9 percent of the average employed wage. In this case, agents are less afraid of having low asset holdings since their compensation when unemployed now serves as partial insurance. At the same time, we set the lower bound on capital holdings to a negative number rather than zero. That is, we allow agents to borrow, with maximum allowable borrowings being set at about half of average annual earnings.¹⁹

To find assumptions that lead to a long (thick) right tail is harder; it seems necessary either to make rich agents have higher propensities to save or to give them higher returns on saving (or both). Quadrini (1996) and Quadrini and Ríos-Rull (1996) do the latter, whereas we explore a setting with preference heterogeneity. In particular, we assume that agents' preferences are *ex ante* identical but that discount factors are random and follow a Markov process. Of course, there are no direct observations on discount factors, but we think that it is reasonable to assume some heterogeneity across generations within a dynasty (i.e., we do not have to assume that different dynasties face different discount rate processes). We subject the experiment to the requirements (1) that the differences in discount factors are not large and (2) that their distribution is symmetric around its mean. More precisely, we assume that β can take on three values, 0.9858, 0.9894, and 0.9930, and that the transition probabilities are such that (i) the invariant distribution for β 's has 80 percent of the population at the middle β and 10 percent at each of the other β 's, (ii) immediate transitions between the extreme values of β occur with probability zero, and (iii) the average duration of the highest and lowest β 's is 50 years. We choose the latter number to roughly match the length of a generation since we view the model

¹⁹ Specifically, we set the wage of an "unemployed" agent equal to 0.07 and we set the borrowing constraint equal to -2.4 (which is stricter than an always pay back constraint).

TABLE 1
DISTRIBUTION OF WEALTH: MODELS AND DATA

MODEL	PERCENTAGE OF WEALTH HELD BY TOP					FRACTION WITH WEALTH < 0	GINI COEFFICIENT
	1%	5%	10%	20%	30%		
Benchmark model	3	11	19	35	46	0	.25
Stochastic- β model	24	55	73	88	92	11	.82
Data	30	51	64	79	88	11	.79

as capturing some elements of an explicit overlapping generations structure with altruism (i.e., parents care about the utility of their children) and less than perfect correlation in genes between parents and children (i.e., there is “regression to the mean” in the rate at which the current generation discounts the utility of future generations).

A. *The Shape of the Wealth Distribution*

Table 1 and figure 3 summarize the average shape of the wealth distribution in the benchmark model, in the stochastic- β model, and in the observed data.²⁰

Table 1 shows that the stochastic- β model, unlike the benchmark model, replicates some of the key features of the observed data on the distribution of wealth. In particular, the large fraction of agents with negative wealth in the stochastic- β model matches the data, and the Gini coefficient is also quite close to that in the data. The table also shows that the stochastic- β model predicts somewhat too low a concentration in the extreme upper tail of the wealth distribution and too high a concentration in the middle; these shortcomings reflect the restrictions of symmetry and so forth that we impose on the discount factor distribution.²¹ Figure 3 graphs the Lorenz curves for the data, the stochastic- β model, and the benchmark model.

In the stochastic- β model, poor agents are poor because they (their dynasties) have chosen to be poor. This choice, in turn, is based purely on (exogenous) genetics: the degree of patience turns

²⁰ When we simulate the behavior of the stochastic- β economy, we use 30,000 agents, so that the invariant distribution across discount factors has 3,000 agents at each of the extreme values of β and 24,000 agents at the middle value of β .

²¹ The data are based on the studies by Wolff (1994) and Díaz-Giménez et al. (1997). Since measures of wealth vary depending on the definitions of wealth and household unit, we simply used unweighted averages across the different numbers in these studies.




FIG. 3.—Lorenz curves for wealth holdings (+ refers to the data, □ to the benchmark model, and ○ to the stochastic-β model).

out to be crucial for the accumulation of wealth of an agent, and the equilibrium interaction between agents with different degrees of patience forces very large differences in wealth from seemingly small differences in patience.

B. Approximate Aggregation Again

Clearly, there is now substantially more wealth heterogeneity in the model. For a large fraction of the consumers there is a significant difference between their current discount rate and the equilibrium interest rate. This difference, in turn, leads to larger differences in propensities to save. In particular, the middle- and (especially) the low-β consumers are more hand-to-mouth consumers than in a model with no discount factor heterogeneity. The important thing to notice, however, is that the accuracy of the aggregate law of motion for capital is still excellent:

$$\log \bar{k}' = 0.100 + 0.961 \log \bar{k}; \quad R^2 = .999991, \hat{\sigma} = 0.0056\%$$

in good times and

$$\log \bar{k}' = 0.095 + 0.961 \log \bar{k}; \quad R^2 = .999985, \hat{\sigma} = 0.0077\%$$

in bad times. That is, in comparison to the benchmark model, the percentage standard deviations of the regression errors do go up by

TABLE 2
AGGREGATE TIME SERIES

Model	Mean(k_t)	Corr(c_t, y_t)	Standard Deviation (i_t)	Corr(y_t, y_{t-1})
Benchmark:				
Complete markets	11.54	.691	.031	.486
Incomplete markets	11.61	.701	.030	.481
$\sigma = 5$:				
Complete markets	11.55	.725	.034	.551
Incomplete markets	12.32	.741	.033	.524
Real business cycle:				
Complete markets	11.56	.639	.027	.342
Incomplete markets	11.58	.669	.027	.339
Stochastic- β :				
Incomplete markets	11.78	.825	.027	.459

about a factor of two, but nonetheless the fit is still remarkably good. The explanation for this result follows the arguments made in Section III D: although marginal propensities differ more across consumers in the stochastic- β model than in the benchmark model, almost all the wealth is held by well-insured consumers. Therefore, only the rich agents matter for determining the aggregates. These rich agents, in turn, behave like the agents from the benchmark model, where almost everyone has sufficient wealth to be well insured in utility terms. The rich thus exhibit permanent income savings behavior with marginal propensities to consume that (almost) do not depend on current wealth or income.

C. Aggregate Time-Series Properties

As may already be clear, the fact that a large fraction of consumers in the stochastic- β model have a significant wedge between their subjective rate of discount and the market rate of return makes the “average consumer” look more impatient. This fact will influence some of the aggregate time-series statistics of the model. We now examine these time-series statistics in some detail since a motivating question for our analysis is to examine to what extent the heterogeneous-agent model of the macroeconomy differs from the standard representative-agent model.

One way to make the comparison is to contrast three separate frameworks: (i) the complete-markets, representative-agent model; (ii) the incomplete-markets model without preference heterogeneity (and an unrealistic wealth distribution); and (iii) the incomplete-markets model with heterogeneity in thrift (and a reasonable wealth distribution). Table 2 shows some selected aggregate statistics for

these models.²² The table also covers two additional extensions of the benchmark model: one with a high value of risk aversion ($\sigma = 5$) and one with valued leisure (the real business cycle model).

Two main kinds of observations can be made using table 2. First, the table allows an evaluation of the effects of market incompleteness. Second, the table makes it possible to gauge the effects of introducing preference heterogeneity and, thus, a realistic wealth distribution on the aggregate time series.

Table 2 shows that the lack of full insurance always raises the steady-state aggregate capital stock since capital has the additional value of insuring against risk. The amount of precautionary savings for the benchmark model is about 0.6 percent (calculated as the percentage increase in the capital stock when insurance markets are closed), which is quite small. However, with more risk-averse agents, the amount of precautionary savings rises significantly: with a degree of relative risk aversion of five, the amount of precautionary savings is 6.7 percent. The reason for the increase is apparent: the agents in the economy increase their total buffer since being poorly insured hurts more when agents are more risk-averse. In contrast, the real business cycle model gives somewhat lower precautionary savings since varying leisure allows a complementary way of adjusting consumption in response to shocks.

The incomplete-markets economies have second-moment properties that are different from their representative-agent counterparts, but not by large amounts. The difference is especially small for the benchmark economy, which shows very small effects of market incompleteness overall. In fact, this observation extends beyond the particular moments reported in table 2: simulated realizations show that the two market structures lead to virtually indistinguishable time series, except for the difference in means. In the setups with higher risk aversion and valued leisure, the differences are somewhat larger. Overall, consumption tends to be a little more correlated with income when markets are incomplete, investment is slightly less volatile, and income is somewhat less serially correlated. In conclusion, in the models we examine that do not have preference heterogeneity, the market structure matters only marginally for aggregate time-series behavior.

In the model with heterogeneity in thrift, what stands out most is

²² In table 2, k_t denotes aggregate capital at time t , y_t aggregate output, c_t aggregate consumption, and i_t aggregate investment. The statistics in the table are computed using a simulated sample consisting of 11,000 observations, with the first 1,000 observations discarded in order to diminish dependence on initial conditions. In order to facilitate comparisons across models, we use the same sequence of aggregate shocks for all model economies when generating the simulated time series.

the consumption-output correlation. The large correlation reflects the hand-to-mouth behavior of the least patient (low- β) consumers. Although these impatient consumers have very little capital and, thus, matter little for what happens to the evolution of economywide wealth, their consumption behavior is a more significant part of the total; the most patient (high- β) consumers are much richer but consume more only by the amount of the interest payments on their assets. In this model, the interest rate is slightly below the discount rate of the most patient agents (the difference arises since assets to some extent are used for precautionary savings). The difference between the interest rate and the discount rate of the least patient agents, however, is larger. This larger wedge leads to a stronger departure from permanent income behavior. In fact, the correlation between the aggregate consumption of the least patient agents and aggregate output is (on average) .90 in this model; for the most patient agents this correlation is .61.²³

In a complete-markets setting in which all consumers have the same discount rate, consumers exhibit permanent income behavior regardless of the degree of their impatience (for reasonably calibrated income processes): if the discount rate is lowered, the stationary equilibrium interest rate adjusts so that the two remain very close. In contrast, in a model in which agents with different discount rates coexist, departures from permanent income behavior arise naturally. An important related point is that the more impatient consumers in the incomplete-markets model do not smooth consumption and thus also may give the impression of being up against a borrowing constraint. Although the borrowing constraint does indeed limit consumption possibilities, the chief reason for impatient agents' rush to consume is precisely that they are more impatient than others. Empirically, it is not obvious how to distinguish a binding borrowing constraint from a lower than average degree of patience.

In sum, we find that the heterogeneous- and representative-agent versions of some of the models generate aggregate behavior with important similarities. We also find, however, that these similarities

²³ We did not solve the complete-markets version of the multiple- β model (it does not aggregate). However, we did compare economies with two groups of agents, one with a permanently high β and one with a permanently low β , since the complete-markets version of such an economy does aggregate: in the stationary state, only the high- β agents are economically active. The effects of incomplete markets on second moments in this case are very large. For example, the consumption-output correlation is .691 in the complete-markets version and .865 in the incomplete-markets version. Since the model with permanently different β 's for different agents is quite similar to the stochastic- β model, in which β 's are highly persistent, the effects of changing the market structure are likely similar for the stochastic- β model.

do not extend to all properties of the models nor to all the models. In addition, heterogeneity appears to be a key factor in generating some of the most prominent differences, such as the tendency of the aggregate, or average, consumer to depart from permanent income behavior. Although this departure is more marked in some of the models than in others, it occurs in all the heterogeneous-agent models that we study.

V. Concluding Remarks

In the Introduction to this paper we set out to investigate whether representative-agent models of the macroeconomy might be justified by showing that models with consumer heterogeneity give rise to aggregate time series that are in fact close to those of the representative-agent models. For this purpose, we studied a fairly rich class of neoclassical production economies with incomplete insurance markets for idiosyncratic risk and heterogeneity in income, wealth, and preferences. Our first and main finding 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, despite substantial heterogeneity in the population with respect to wealth as well as to some preference parameters. Hence, one need track only the evolution of the “aggregate budget” in order to analyze the dynamic behavior of the macroeconomic aggregates.

A natural question is thus whether it is possible to describe preferences of an “aggregate consumer” such that when this agent’s utility is maximized subject to the aggregate budget the outcome is a set of time series that matches those calculated here. Our second finding is that this is sometimes, but not always, an easy task. In our benchmark economy and some extensions to it, the incomplete-markets (heterogeneous-agent) economy behaves almost identically to its complete-markets (representative-agent) counterpart, except for a difference in levels: in the stationary stochastic equilibrium, capital is slightly higher in the incomplete-markets version. In some extensions to the benchmark setup, however, especially those with preference heterogeneity, some second moments of the aggregate time series are substantially different. Although we have not explicitly tried to formulate preferences of a fictitious representative consumer whose behavior might match the aggregate consumption behavior generated by the heterogeneous-agent models, it seems necessary to move outside the class of models with a single infinitely lived agent with time-additive preferences. In particular, in representative-agent models with time-additive preferences, it seems dif-

difficult to obtain the departures from permanent income behavior that we observe in the incomplete-markets models with heterogeneity in preferences. In other words, the interaction of consumers with heterogeneous preferences in an incomplete-markets setting leads to new insights. Among the models we study, those that come the closest to matching real-world wealth distributions are precisely models with heterogeneous preferences and incomplete markets.

An important task is to investigate further the robustness of our findings. First, it would be interesting to know whether life cycle savings motives might lead to different results. Ríos-Rull (1996) shows that calibrated life cycle models without altruism but with complete markets generate aggregate behavior that is very similar to that of the benchmark (or real business cycle) model studied in this paper. These results suggest that life cycle considerations alone do not lead to significant departures from permanent income behavior on the part of individual consumers. An open question is whether introducing incomplete markets into a life cycle model would be sufficient to generate significant departures from permanent income behavior on the part of individual consumers.²⁴ Second, models with endogenous borrowing constraints, such as in Kiyotaki and Moore (1997), might ascribe a more important role to the distribution of wealth. Third, one simple extension of the standard growth model that may give rise to stronger effects from the distribution of capital is to force consumers to use their own production technologies (in contrast to our setup in which the return to savings is equal for all agents). Finally, as Banerjee and Newman (1993) have shown, fixed costs in capital accumulation may also lead to a more fundamental dependence of the aggregates on the distribution of resources across agents.

Although not in main focus in this paper, how the state of the macroeconomy and macroeconomic policy shape the distribution of consumption and wealth is an important question, at least judging from contemporary public debate. The methodological findings in this paper suggest that such issues can be feasibly studied: it now seems within our ability to begin using equilibrium models to analyze the interrelation between business cycles, inequality, and economic policy. In this context, an important task is to analyze further the determinants of the wealth distribution. As we have shown in this paper, introducing preference heterogeneity into the standard model allows a closer match between model and data. It is necessary, however, to consider carefully other hypotheses as well. The use of

²⁴ Storesletten, Telmer, and Yaron (1996) investigate a life cycle model with incomplete markets and persistent idiosyncratic shocks.

household data should prove fruitful for comparing competing hypotheses.

Appendix

A. Numerical Solution Method for the Agent's Problem

This Appendix describes 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 benchmark 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 accommodate a different model specification or additional moments.²⁵

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 each of these 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 we do not restrict choices for capital to points on the grid.

The following steps describe the numerical procedure. (1) Choose a grid of points in the (k, \bar{k}) plane (we give some details below about how we choose these points). (2) 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.) (3) 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). (4) Compare the new optimal values generated by step 3 to the original values. If the new values are close to the old values, then stop; otherwise, repeat step 3 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

²⁵ There are many available methods for solving the class of decision problems that we consider in this paper. We have chosen a method that we find to be robust to changes in model specification and that allows us to achieve high accuracy.

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 n points in the \bar{k} direction, then polynomial interpolation fits a polynomial of order $n - 1$ to the function values at these points (so that the polynomial fits the values exactly). This polynomial is then used to compute the value function in between grid points. We compute the value of the interpolating polynomial using Neville's algorithm, as described in Press et al. (1989, chap. 3). This algorithm avoids the numerical instabilities associated with computing the coefficients of the interpolating polynomial. We generally use four to six 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. 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. Cubic splines can be computed efficiently by solving a set of tridiagonal linear equations (see, e.g., the description in Press et al. [1989, chap. 3] or de Boor [1978, chap. 4]). Computing cubic splines requires the imposition of two side conditions: we impose that the second derivative of the value function at the first grid point for k is slightly smaller than the second derivative at the second grid point and that the second derivative at the last grid point is slightly larger than at the next to last grid point. We generally use 70–130 grid points in the k direction, with many grid points near zero (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} direction.

To combine these two interpolation schemes, we therefore proceed as follows, where m is the number of grid points in the k direction. (1) For each of the m values of k , use polynomial interpolation to compute the value function at the desired value of \bar{k}' . This set of interpolations yields m values of the value function, one for each value of k in the grid. (2) Use cubic spline interpolation using the m interpolated values to calculate the value function for values of k that are not on the grid. Since the values of \bar{k}' at which interpolated values must be computed are known at the beginning of each of the iterations on the value function, the required cubic splines need to be computed only once for each iteration.

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 150–600 equally spaced points in the k direction and 25–100 equally spaced points in the \bar{k} direction. Optimal decisions at points not on the grid are then computed using bilinear interpolation (see Press et al. 1989, chap. 3). To conserve on computation time, the coefficients determining the bilinear interpolation need to be computed only once, prior to simulating the behavior of the economy. Given these coefficients and an efficient method for finding the appropriate location in the two-dimensional (k, \bar{k}) grid, it is quick and easy to compute individual savings decisions.

As a final point, it is important to ensure that simulations of the economy’s behavior impose the law of large numbers (or at least its first-moment implications). In particular, it is important to make sure that the fraction of unemployed agents is exactly u_g in good times and u_b in bad times. To accomplish this task in the simulated data, we first update the employment status of each agent according to the appropriate conditional probabilities. We then check to see whether the fraction of unemployed agents matches the desired number. If, for example, there are too many employed agents, we choose an employed agent at random, switch his employment status to unemployed, and then continue this process until the fraction of employed agents matches the desired number.

B. The Model with Valued Leisure

Assume that time spent off work, $1 - l$, where l is the amount of labor supplied, is valued according to

$$U(c, l) = \lim_{v \rightarrow \sigma} \frac{[c^\theta (1 - l)^{1-\theta}]^{1-v} - 1}{1 - v},$$

where we set σ to one and θ to 1/2.9. In our recursive equilibrium definition, there is now an additional element to consider: the way in which 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 g .

When we approximate the aggregate labor supply function as a log-linear function of the total stock of capital, our results are as follows:

$$\begin{aligned}\log \bar{k}' &= 0.123 + 0.951 \log \bar{k}; \quad R^2 = .999994, \hat{\sigma} = 0.0040\%, \\ \log \bar{l} &= -0.544 - 0.252 \log k; \quad R^2 = .992, \hat{\sigma} = 0.039\%\end{aligned}$$

in good times and

$$\begin{aligned}\log \bar{k}' &= 0.114 + 0.953 \log \bar{k}; \quad R^2 = .999993, \hat{\sigma} = 0.0049\%, \\ \log \bar{l} &= -0.592 - 0.255 \log k; \quad R^2 = .988, \hat{\sigma} = 0.054\%\end{aligned}$$

in bad times.²⁶ The fit is still very good for the law of motion for capital. 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 in contrast to figure 1, there are noticeable "clouds" in a graph of aggregate labor supply against aggregate capital (the working paper version of this paper [Krusell and Smith 1996a] displays this graph).

Just as for the savings decision, the nonlinearities in the agent's decision rule for labor are stronger at low levels of capital (the marginal propensity to take leisure is higher for poor agents). These nonlinearities, together with the fact that poor agents supply a disproportionately large amount of aggregate labor, account for the "clouds" in a graph of aggregate labor against aggregate capital. The fit, however, is very good: although it is possible to detect a decrease in forecasting accuracy as compared to the benchmark model, the changes are minor. Our assumption that utility is Cobb-Douglas in consumption and leisure clearly works against aggregation; since poor agents work harder than rich agents in the model, the nonlinearities in the decision rule for leisure receive high weight. A preference assumption with no wealth effects—a nesting of the kind $c + (1 - l)^{\gamma}$ —would improve the fit significantly.²⁷ Similarly, if employed agents are allowed to have different labor productivities, those with higher productivity would work more and be richer, leading to significantly smaller "clouds" and improved fits. Further, extensions to include human capital accumulation would tend to lead to a positive correlation between financial and human capital. This extension would give poor agents a low weight in the total supply of effective labor units and, thus, diminish the effects of the nonlinearities. In fact, these alternative setups seem more consistent with observed patterns for hours worked and relative wages. We chose our specification mainly to illustrate that, despite assumptions that work against aggregation, the results are close to those of the benchmark setup.

²⁶ We also computed an approximate equilibrium using two moments of the capital stock distribution: its mean and its standard deviation. In this case, both the law of motion for aggregate capital and the aggregate labor supply function depend on both of these moments. Including an additional moment leads to significantly better fits for the aggregate labor functions: in good times, $R^2 = .998$ ($\hat{\sigma} = 0.021$ percent), and in bad times, $R^2 = .993$ ($\hat{\sigma} = 0.039$ percent). In addition, the fit of the law of motion for aggregate capital improves slightly. The aggregate time series, however, are virtually unchanged.

²⁷ This kind of formulation has been used in Greenwood, Hercowitz, and Huffman (1988).

References

- Aiyagari, S. Rao. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Q.J.E.* 109 (August 1994): 659–84.
- Aiyagari, S. Rao, and Gertler, Mark. "Asset Returns with Transactions Costs and Uninsured Individual Risk." *J. Monetary Econ.* 27 (June 1991): 311–31.
- Banerjee, Abhijit V., and Newman, Andrew F. "Occupational Choice and the Process of Development." *J.P.E.* 101 (April 1993): 274–98.
- Bewley, Truman F. "The Permanent Income Hypothesis: A Theoretical Formulation." *J. Econ. Theory* 16 (December 1977): 252–92.
- Castañeda, A.; Díaz-Giménez, Javier; and Ríos-Rull, José-Victor. "Unemployment Spells, Cyclically Moving Factor Shares, and Income Distribution Dynamics." *J. Monetary Econ.* (in press).
- Chatterjee, Satyajit. "Transitional Dynamics and the Distribution of Wealth in a Neoclassical Growth Model." *J. Public Econ.* 54 (May 1994): 97–119.
- Cochrane, John H. "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives." *A.E.R.* 79 (June 1989): 319–37.
- de Boor, Carl. *A Practical Guide to Splines*. New York: Springer-Verlag, 1978.
- den Haan, Wouter J. "Heterogeneity, Aggregate Uncertainty, and the Short-Term Interest Rate." *J. Bus. and Econ. Statist.* 14 (October 1996): 399–411. (a)
- . "Understanding Equilibrium Models with a Small and a Large Number of Agents." Discussion Paper no. 96-34. La Jolla: Univ. California, San Diego, 1996. (b)
- Díaz-Giménez, Javier; Prescott, Edward C.; Fitzgerald, T.; and Alvarez, Fernando. "Banking in Computable General Equilibrium Economies." *J. Econ. Dynamics and Control* 16 (July–October 1992): 533–59.
- Díaz-Giménez, Javier; Quadrini, Vincenzo; and Ríos-Rull, José-Victor. "Dimensions of Inequality: Facts on the U.S. Distributions of Earnings, Income, and Wealth." *Fed. Reserve Bank Minneapolis Q. Rev.* 21 (Spring 1997): 3–21.
- Díaz-Giménez, Javier, and Ríos-Rull, José-Victor. "Heterogeneous Agents, Capital Accumulation, and Business Cycles." Manuscript. Philadelphia: Univ. Pennsylvania, 1991.
- Greenwood, Jeremy; Hercowitz, Zvi; and Huffman, Gregory. "Investment, Capacity Utilization, and the Real Business Cycle." *A.E.R.* 78 (June 1988): 402–17.
- Heaton, John, and Lucas, Deborah J. "The Importance of Investor Heterogeneity and Financial Market Imperfections for the Behavior of Asset Prices." *Carnegie-Rochester Conf. Ser. Public Policy* 42 (June 1995): 1–32.
- . "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing." *J.P.E.* 104 (June 1996): 443–87.
- Hubbard, R. Glenn; Skinner, Jonathan; and Zeldes, Stephen P. "Precautionary Saving and Social Insurance." *J.P.E.* 103 (April 1995): 360–99.
- Huggett, Mark. "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies." *J. Econ. Dynamics and Control* 17 (September–November 1993): 953–69.
- . "Wealth Distribution in Life-Cycle Economies." *J. Monetary Econ.* 38 (December 1996): 469–94.

- İmrohoroğlu, Ayşe. "Cost of Business Cycles with Indivisibilities and Liquidity Constraints." *J.P.E.* 97 (December 1989): 1364–83.
- . "The Welfare Cost of Inflation under Imperfect Insurance." *J. Econ. Dynamics and Control* 16 (January 1992): 79–91.
- Johnson, Sharon A.; Stedinger, Jerry R.; Shoemaker, Christine A.; Li, Ying; and Tejada-Guibert, José Alberto. "Numerical Solution of Continuous-State Dynamic Programs Using Linear and Spline Interpolation." *Operations Res.* 41 (May–June 1993): 484–500.
- Kiyotaki, Nobuhiro, and Moore, John. "Credit Cycles." *J.P.E.* 105 (April 1997): 211–48.
- Krusell, Per, and Ríos-Rull, José-Victor. "On the Size of U.S. Government: Political Economy in the Neoclassical Growth Model." Manuscript. Rochester, N.Y.: Univ. Rochester, 1997.
- Krusell, Per, and Smith, Anthony A., Jr. "Income and Wealth Heterogeneity in the Macroeconomy." Working Paper no. 399. Rochester, N.Y.: Univ. Rochester, Center Econ. Res.; Working Paper no. 1997-37. Pittsburgh: Carnegie Mellon Univ., 1996. (a)
- . "Rules of Thumb in Macroeconomic Equilibrium: A Quantitative Analysis." *J. Econ. Dynamics and Control* 20 (April 1996): 527–58. (b)
- . "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns." *Macroeconomic Dynamics* 1, no. 2 (1997): 387–422.
- Lucas, Deborah J. "Asset Pricing with Undiversifiable Income Risk and Short Sales Constraints: Deepening the Equity Premium Puzzle." *J. Monetary Econ.* 34 (December 1994): 325–41.
- Lucas, Robert E., Jr. *Models of Business Cycles*. New York: Blackwell, 1987.
- Marcet, Albert, and Singleton, Kenneth J. "Equilibrium Asset Prices and Savings of Heterogeneous Agents in the Presence of Incomplete Markets and Portfolio Constraints." Manuscript. Barcelona: Univ. Pompeu Fabra, 1991.
- Press, William H.; Flannery, Brian P.; Teukolsky, Saul A.; and Vetterling, William T. *Numerical Recipes: The Art of Scientific Computing*. Cambridge: Cambridge Univ. Press, 1989.
- Quadrini, Vincenzo. "Entrepreneurship, Saving, and Social Mobility." Ph.D. dissertation, Univ. Pennsylvania, 1996.
- Quadrini, Vincenzo, and Ríos-Rull, José-Victor. "Models of the Distribution of Income, Wealth and Earnings." Staff report. Minneapolis: Fed. Reserve Bank, Res. Dept., 1996.
- Ríos-Rull, José-Victor. "Life-Cycle Economies and Aggregate Fluctuations." *Rev. Econ. Studies* 63 (July 1996): 465–89.
- Scheinkman, José A., and Weiss, Laurence. "Borrowing Constraints and Aggregate Economic Activity." *Econometrica* 54 (January 1986): 23–45.
- Storesletten, Kjetil; Telmer, Chris I.; and Yaron, Amir. "Persistent Idiosyncratic Shocks and Incomplete Markets." Manuscript. Pittsburgh: Carnegie Mellon Univ., 1996.
- Telmer, Chris I. "Asset-Pricing Puzzles and Incomplete Markets." *J. Finance* 48 (December 1993): 1803–32.
- Wolff, Edward N. "Trends in Household Wealth in the United States, 1962–83 and 1983–89." *Rev. Income and Wealth* 40 (June 1994): 143–74.