Frontiers of Business Cycle Research

Thomas F. Cooley, Editor

Princeton University Press Princeton, New Jersey

Chapter 4
Models with Heterogeneous Agents

José-Víctor Ríos-Rull

1. Introduction

Some questions addressed in the business cycle literature require environments with multiple agents. One obvious question we need to ask is, How robust are the findings of this literature to modifications of the model that incorporate observed heterogeneity of agents in characteristics such as age and skill level, as well as imperfect insurance markets? Other, more specific questions involve the cyclical behavior of the distribution of income and wealth, or the role played by liquidity constraints, and of asymmetric information. Moreover, multiagent models are being used to address certain features of the data that are difficult to explain within the representative-agent framework. These include the cyclical behavior of the factor shares of national income, and the risk premium puzzle, which is a related issue in financial economics.

Computable models that are capable of dealing with multiagent environments have been developed to address many of these questions. Computation of equilibria in these models is usually substantially more difficult than in standard representative-agent models, as equilibrium laws of motion become functions not only of aggregate variables, but also of the distribution of these variables across different types of agents. Solving for the laws of motion of such distributions is a nontrivial task.

We review these models by grouping them according to the criteria that give rise to the heterogeneity. Section 2 deals with economies where agents face some idiosyncratic shocks to productivity, and where there are no markets to insure such risks. In Section 3, overlapping generations economies are reviewed. In these two sections, there are a large number of agent types. Section 4 is about a variety of two-agent models, each addressing a specific question, while Section 5 gives some indications regarding the course of future research.

2. Models with Uninsurable Individual Risks

Potential Uses

The existence of a full set of Arrow-Debreu markets makes the distribution of income irrelevant, as agents only face one budget constraint in the initial period. A property of this class of complete-markets economies is that all individual risks get perfectly insured at no cost. The existence of an insurance industry in the United States that accounts for a nontrivial portion of output reminds one that whatever insurance is available is not free. If instead of complete markets, trading arrangements only include asset holdings with returns that cannot be made contingent upon the realization of individual shocks, then agents trade every period, and they will insure themselves against adverse realizations by holding assets that will be depleted during bad times. These arrangements represent an alternative, and perhaps more accurate, description of the trading opportunities in actual economies. In contrast with complete-markets arrangements they give an important role to the distribution and intertemporal movements of income and wealth.

One interesting use of such models has been to evaluate the cost of business cycles when the model accounts for some of the observed heterogeneity of people in the data. Another use has been to assess the welfare properties of various monetary policies when money is modeled as a store of value. Finally, they have been used to assess the properties of a variety of government policies, such as unemployment insurance, and social security, that operate as a substitute for perfect insurance.

General Description

In these economies, agents differ on how they are affected by some idiosyncratic risk, as well as in the assets they are able to hold. The key feature that makes these economies different from representative-agent models is that the set of possible trades available for agents is restricted.[1] Typically, agents cannot write contracts contingent on their individual shocks, and in some cases, they cannot hold negative quantities of any asset. This prevents various aggregation results from holding[2] (see, e.g., Deaton 1992), and thus computing the equilibrium requires keeping track of the distribution of agents. At any point in time, the state of the economy is characterized by how agents are positioned across levels of asset holdings, individual shocks, and, perhaps, an aggregate exogenous shock to the economy. In general, equilibrium prices and quantities depend on both the distribution of agents and the aggregate shock. The individual agents, when solving their maximization problem, have to be able to predict future prices, and thus they have to use the distribution of agents and its law of motion as an input to their decisionmaking process. Note that the aggregate level of asset holdings would not be enough to characterize the individual state, since in order to predict next period's level, we would also have to know today's distribution, as different agents accumulate different amounts, depending on their individual states.

These considerations imply that computing the competitive equilibrium amounts to finding a fixed point in the space of functions from a set of measures into itself, an unmanageable computational problem. The key computational difficulties arise when solving the problem of the agent because a measure becomes a state variable. A measure is not a standard state variable, such as, say, aggregate physical capital; it is a far more complex mathematical object. Its characterization, even if it is approximated, demands that a lot of information be stored. This means that the individual state space is a large set, and that decision rules have to be solved for a large variety of circumstances. Furthermore, inherent features of the problem (for example, the restrictions imposed in the set of feasible trades) prevent the use of the cheapest computational method, linear-quadratic approximations, because the optimal solution sometimes hits a corner. The other standard methods available, discretization of the state space, parameterized expectations, backward solving, and so on, have severe shortcomings in environments with a large set of state variables.

Note also that solving the problem of the agent for a given law of motion of the distribution is not enough. The correct law of motion of the distribution has to be found. This requires iterating at this level too, which introduces another layer of computational complexity into the process.

For all these reasons, it is necessary to reduce the dimensionality of the problem of the agent. Typically, the procedure followed is to prevent the distribution of agents from affecting relative prices, in particular, the relative wage and the rate of return of capital. This drastically simplifies the problem of the agent and completely avoids iterations on the law of motion of the distribution, which can then be computed residually. In this fashion, İmrohoroğlu (1989) utilizes a storage technology that pins down exogenously the rate of return of savings. In Díaz-Giménez and Prescott (1989) and Díaz-Giménez (1990), the government commits itself to a certain inflation rate policy that does not depend on the asset distribution.

The approach that these models use to compute the equilibrium is to create a grid in the set of possible asset holdings, and to have a finite number of possible realizations of both individual and aggregate exogenous variables. In this case, the maximization problem of the agent, if written in the form of a value function, becomes a finite problem, as the value function and the distribution of agents become vectors.

Next, we develop a general structure that encompasses as special cases most of the papers in the area.

A General Heterogeneous-Agent Model with Liquidity Constraints

The economy consists of a continuum of agents that is taken to be of measure one. They have standard preferences over streams of consumption and leisure, which can be represented by $E_0\{\sum_{t=0}^{\infty} u(c_t, n_t)\}$, where for each period t, the pair (c_t, n_t) is restricted to belong to C, the per-period consumption possibility set of the agent, which might specify nonnegative consumptions and restrict the choice of possible hours worked by the agent. There is a limited set of markets to which agents have access. Every period, they engage in production activities, and accumulate assets, which are restricted to be in a set $A \subset \mathbf{R}^{n_a}$. We can think of $a \in A$ as including real assets, which we generically label as a_1, and which affect production possibilities, and a variety of financial assets, such as money, a_2, and bonds, a_3, which differ in their denomination sizes and rates of return. For convenience, we measure these assets in real terms at their acquisition values for the previous period; for instance, real capital is measured in goods of the period it was acquired. All these assets are not identical, and therefore they might command different rates of return. We thus think of r_1 as the real rate of return on capital, while $\frac{1}{r_2} - 1$ is the inflation rate.

The after-tax production capability of the agents depends partly on an exogenous idiosyncratic shock, $s \in S$, and on an economywide shock, $z \in Z$. We assume that both s and z can only take a finite number of values N_S and N_Z. These shocks are assumed to be Markov processes, with transition matrix $\pi(s', z' \mid s, z)$. The problem of the agents is to maximize expected utility subject to a sequence of budget constraints, such as $c_t + e^T \cdot a_{t+1} = a_t^T \cdot r_t + n_t w_t + tr_t$, where $a_t \in A$ is the vector of asset holdings; r_t is the vector of gross after-tax rates of return on those assets (or one plus the rate of return); e is just a vector of ones used to aggregate (and the superindex T indicates that the vector is transposed, that is, that it is a row vector); w_t is the after-tax real wage; and tr_t are government transfers, which might depend on the agents' actions (unemployment insurance is the obvious example). Given processes for w_t, r_t, and tr_t, the agents are characterized by (z_t, s_t, a_t), but in order to be able to solve their maximization problem, those processes have to be specified. For the reasons noted above, the processes r_t, w_t, and tr_t have to be independent of the distribution of agents, but they can depend on the particular state of each agent, and, even on the agents' labor decisions.

There is a government in this world that commits itself to policies $w(z, s, a, n)$, $r(z, s, a, n)$, and $tr(z, s, a, n)$. The government can do this through a variety of fiscal, monetary, and debt policies. With this arrangement, the government is restricted by technology considerations, and by its budget constraint. It operates in asset markets, and it imposes taxes on capital and labor income that guarantee that labor and capital prices, w and r, and transfer payments, tr, are independent of the distribution of agents across individual states. In this fashion, the policies that are implemented by the government have to be functions of the aggregate state of the economy. They do not enter, however, into the problem of the agent, as the after-tax pricing functions that agents use to make their decisions are not functions of the aggregate state. In fact, this is what the government does: it implements policies that prevent the distribution of agents from affecting agents' decisions.

At any point in time, the economy is characterized by the aggregate shock, $z \in Z$, and a distribution of agents according to their asset holdings and individual states. Mathematically, this can be represented by a pair, (z, μ), where μ is a measure defined over $\mathcal{B}$, an appropriate family of subsets of $(S \times A)$, the set of possible states for individual agents. The problem of the agents' is to maximize utility, subject to their budget constraint and to the trading restrictions, and taking as given the after-tax pricing functions, $w(z, s, a, n)$ and $r(z, s, a, n)$, and the transfer function, $tr(z, s, a, n)$.

Production possibilities depend on the total amount of capital and labor input, and the aggregate shock, z. The labor input is constructed by aggregating hours worked over agents that have idiosyncratic shock, s. Let (K, N) be the total amount of factor inputs; then total output is given by a constant-returns-to-scale production function, $f(K, N, z)$.

Equilibrium Defined

We have completed the characterization of the environment, so we are ready to define equilibrium. It consists of a set of government policies; after-tax real wages, $w(z, s, a, n)$, after-tax rates of return for real capital, $r_1(z, s, a, n)$, for money, $r_2(z, s, a, n)$, and for bonds, $r_3(z, s, a, n)$; and for transfers $tr(z, s, a, n)$; tax rates on capital income, $\tau(z, \mu, s, a, n) \in \mathbf{R}^{n_a}$; tax rates on labor income, $\tau_0(z, \mu, s, a, n)$; government expenditures, $G(z, \mu) \geq 0$, bonds issued by the government, $b_g(z, \mu)$, and monetary policy, $m_g(z, \mu)$; decision rules of the agents for consumption, hours worked, and asset holdings, $\{c(z, s, a), n(z, s, a), a'(z, s, a)\}$; the value function of the agents' problem, $v(z, s, a)$; functions for aggregate inputs, $K(z, \mu)$ and $N(z, \mu)$; and a law of motion for the distribution of agents in the economy, $\mu' = g(z, \mu, z')$, such that

1) given $r(z, s, a, n)$, $w(z, s, a, n)$, and $tr(z, s, a, n)$, agents' decision rules, $\{c(z, s, a), n(z, s, a), a'(z, s, a)\}$, solve their problem:

$$v(z, s, a) = \max_{\{a', c, n\}} u(c, n) + \beta E\{v(z', s', a' \mid s, z\}$$

$$\text{s.t.} \quad a' \in A$$

$$(c, n) \in C$$

$$a^T \cdot r(z, s, a, n) + n\, w(z, s, a, n) + tr(z, s, a, n) \geq e^T \cdot a' + c \tag{1}$$

for all $(z, s, a) \in Z \times S \times A$;

2) The goods market clears, and government expenses are positive:

$$\int_{A,S} \left[a_1'(z, s, a) + c(z, s, a)\right] d\mu + G(z, \mu) = f[K(z, \mu), N(z, \mu), z], \tag{2}$$

for all (z, μ).

3) The money market clears:

$$\int_{A,S} a_2'(z, s, a) d\mu - m_g(z, \mu) = 0, \tag{3}$$

for all (z, μ);

4) The bonds markets clear:

$$\int_{A,S} a_3'(z, s, a) d\mu - b_g(z, \mu) = 0, \tag{4}$$

for all (z, μ);

5) Aggregate factor inputs are generated by decision rules of the agents:

$$K(z, \mu) = \int_{A,S} a_1\, d\mu,$$

and

$$N(z, \mu) = \int_{A,S} s\, n(z, s, a) d\mu, \tag{5}$$

for all (z, μ);

6) Pretax factor prices are marginal productivities:

$$w(z, s, a, n) = [1 - \tau_0(z, \mu, s, a, n)]\, s\, f_2[K(z, \mu), N(z, \mu), z],$$

and

$$r(z, s, a, n) = [1 - \tau_1(z, \mu, s, a, n)] f_1[K(z, \mu), N(z, \mu), z]; \tag{6}$$

and

7) Individual and aggregate behavior are consistent:

$$\mu'(S_0, A_0) = g(S_0, A_0)(z, \mu, z')$$

$$= \int_{S_0, A_0} \left\{ \int_{S,A} 1_{a'=a'(z,s,a)} \pi(s' \mid s, z, z') d\mu \right\} da' ds', \tag{7}$$

for all $(S_0, A_0) \in \mathcal{B}$, and all (z, μ, z'), with $1_{a'=a'(z,s,a)}$ being an indicator function that takes the value one if the statement is true and zero otherwise.

An important feature that we have avoided is the budget constraint of the government—nowhere in this definition have we talked about it. This is because Walras's law takes care of it. Typically, equilibrium is defined by requiring that all agents satisfy their budget constraints, and by requiring that all markets except one clear, with aggregation (Walras's law) taking care of the last market. Here, we have required that all markets clear, and that all agents except the government satisfy their budget constraint. Aggregation implies that the government's budget constraint is also satisfied.

An Example without Aggregate Endogenous Variables

"Cost of Business Cycles with Indivisibilities and Liquidity Constraints" (İmrohoroğlu 1989) is perhaps the first published paper where the equilibrium of an artificial economy with heterogeneous agents, calibrated to match some key U.S. observations, is computed. Its purpose is to find whether measurements of the social cost of business cycles that treat every person as perfectly insured against idiosyncratic risks—in particular, the evaluation made by Lucas (1987)—severely understates such cost relative to an economy where liquidity constraints are pervasive. İmrohoroğlu considers three different environments. The first is one with liquidity constraints, where the assets of the agents cannot be negative and where there is no possibility of writing contingent contracts. Another institutional market arrangement allows for the existence of credit markets but not contingent markets for idiosyncratic or aggregate risks. In this economy, holding assets remains the only mechanism of insurance. Finally, she considered an economy where perfect insurance prevails. For all three economies, the allocation with and without aggregate fluctuations is computed. Comparisons between the allocations of the economy with aggregate fluctuations and the economy without them permit the computation of the additional consumption required to make agents indifferent between the two. If this difference is similar across institutional market arrangements, we could conclude that the existence of liquidity constraints in an economy is not an important feature for welfare considerations. It is very easy to specialize our general model described above to her environment. In particular, there is no need to introduce a government since relative prices do not depend on the total amount of assets held in the economy.

Technology and preferences are identical across economies. Labor is inelastic, so $C = \mathbf{R}_+ \times 1$, and preferences are $E_0\{\sum_{t=0}^{\infty} \beta^t u(c_t)\}$. Agents face two possible individual states, $s \in S = \{e, u\}$; in state e there is an employment opportunity, and in state u there is none. An agent with an individual state s produces $w(s)$ units of output, where $w(e) > w(u)$. One can think of $s = e$ as being employed and $s = u$ as being unemployed, but being able to work at home. The transition of this state depends on the state of the economy as a whole, $z \in Z = \{g, b\}$, which can be in either good or bad times. The matrix $\pi(s', z' \mid s, z)$ is the transition matrix of the aggregate state of the economy and the individual state. The difference between economies that have business cycles and those that do not is determined by the specification of π. In economies with business cycles, the transition matrix depends on z, while in economies without business cycles it does not. In these economies, as in the data, recessions are characterized by agents' having both a higher probability of losing a job and a higher probability of not finding a job than in booms. Of course, this difference shows up not

only in individual transition probabilities, but also in the aggregate level, with unemployment typically building up during recessions, and falling during booms. There is a publicly available storage technology that transforms a units of the time t good into $a\,r$ units in $t+1$. These considerations imply that we can write the production function as $\int_{A,S}[ra_1 + w(s)]d\mu$, since the only asset available is the storage good.

Liquidity Constraints

The first market arrangement considered has liquidity constraints and no contracts contingent on either individual states or the aggregate state. In this economy, the only means that agents have to smooth consumption over time is to store the good as a precaution against future bad times. Since there is no money ($a_2 = 0$), or bonds ($a_3 = 0$), the set of possible asset holdings becomes $A = \mathbf{R}_+ \times \{0\} \times \{0\}$. The return of the unique asset is that of the storage technology, r. The problem of the agent can be written in the form of a value function as

$$v(z, s, a_1) = \max_{\{a_1', c\} \geq 0} u(c) + \beta E\{v(z', s', a_1') \mid s, z\}$$

$$\text{s.t.} \quad a_1 r + w(s) = a_1' + c. \tag{8}$$

An equilibrium for this economy consists of decision rules $c(z, s, a_1)$ for consumption and $a_1'(z, s, a_1)$ for inventory holdings, together with the value function, v, and a law of motion of the population, $g(z, \mu)$, such that

1) $c(z, s, a_1)$ and $a_1'(z, s, a_1)$ solve problem (8) given $w(s)$; and

2) The goods market clears:

$$\int_{A,S} [a_1'(z, s, a_1) + c(z, s, a)]\, d\mu = \int_{A,S} [a_1(z, s, a_1) + w(s)]\, d\mu, \tag{9}$$

Notice that the first condition guarantees that the second condition is satisfied. Condition 1 requires that the budget constraint be satisfied for every agent, which is just another way of requiring it to be satisfied for every pair (s_1, a_1). Then it will also be satisfied for the integral with respect to the measure that that describes the state of the economy. This is another way of saying that storage and the budget constraint guarantee that the goods market will clear.

Conditions 3–6 in the definition of equilibrium for the general problem are irrelevant, as there are no bonds, no labor, and no marginal productivities.

A version of condition 7 in the definition of equilibrium for the general problem has to be satisfied. All that is required is that the law of motion, g, of the distribution of agents across asset holdings, a_1, and individual shocks, s, evolve according to agents' choices and the stochastic properties of s and z. We write this condition as

$$\mu'(S_0, A_0) = g(S_0, A_0)(z, \mu, z')$$
$$= \int_{S_0, A_0} \left\{ \int_{S,A} 1_{a_1'=a_1'(z,s,a_1)} \pi(s' \mid s, z, z') d\mu \right\} da_1' ds', \quad (10)$$

Pure Credit

The second market arrangement allows for noncontingent loan contracts. In the model, the rates at which the agent can borrow and lend are different, since intermediation uses real resources. The problem of the agent becomes slightly different than in the first market arrangement, since the agent's choice set is enlarged by allowing for negative assets. We model this arrangement by letting borrowing simply be negative bond holdings ($a_3 \leq 0$), with gross rate of return bigger than one. Here, $A = \mathbf{R}_+ \times \{0\} \times \mathbf{R}_-$, where the first component is the amount stored or lent, with gross return r, and the third component is the amount borrowed with interest rate i. Note that both rate r and rate i can be considered parameters since both are technologically determined, r by storage and i by the sum of both storage and intermediation costs, since lending and storing have to be perfect substitutes in equilibrium. The problem of the agent is

$$v(z, s, a_1, a_3) = \max_{\{a_1', -a_3', c\} \geq 0} u(c) + \beta E\{v(z_1', s_1') \mid z, s)$$
$$\text{s.t.} \quad a_1 r + a_3(1 + i) + w(s) = a_1' + a_3' + c. \quad (11)$$

The definition of equilibrium is similar to the definition of the first market arrangement, except that in order to guarantee feasibility, total asset holdings cannot be negative, $\int_{A,S}(a_1 + a_3)\, d\mu \geq 0$. This implies that if the interest rate is too low, and the rate differential too small, equilibrium might not exist because aggregate savings cannot be negative. Savings can be positive because of the storage.

Complete Markets

The last market arrangement considered allows for perfect insurance, or complete markets. The allocation of this economy can be found as the solution to a planner's problem, as the welfare theorems hold for this economy. The planner's problem chosen is the one that treats everyone equally. The allocation found this way corresponds to the equilibrium when the initial condition is equal initial wealth distribution. It is found by solving the following maximization problem:

$$v(z, N, a_1) = \max_{a_1' \geq 0} u[ra + Nw(e) + (1 - N)w(u) - a_1']$$
$$+ \beta E\{v(z', N', a_1') \mid z\}$$

$$\text{s.t.} \quad N'(N, z') = N\pi(e, z' \mid e, z) + (1 - N)\pi(e, z' \mid u, z), \quad (12)$$

where N is aggregate employment, and its law of motion is given in the constraint.

Calibration of this economy is standard except for the transition, π, and the ratio between unemployment and employment income, $w(u)/w(e)$. The values of π are chosen so that the aggregate state is symmetric and the average length of a cycle 4 years, and so that the average length of a spell of unemployment matches what we see in the United States in both good and bad times. With these considerations, π is completely pinned down. To choose a value for the relative income between the individual states, observations on the size of unemployment insurance and on how widespread it is are used.

The equilibria for these artificial economies are computed, and their statistical properties are empirically determined. It is found that among the properties of these models are that the cost of the business cycle in economies with liquidity constraints is at least three times bigger than in economies with perfect insurance.[5]

Computing the Equilibrium

The most popular approach to solving this type of model is to discretize the state space. A grid on $\{Z \times S \times A\}$ is constructed. Agents' choices are restricted to belong to a finite set. This property determines that the value function of agents can be represented as a vector of size, $(N_Z \times N_S \times N_A)$, where N_A is the total number of grid points in A, an element of a finite-dimensional Euclidean space. Decision rules specify for each of the elements in $\{Z \times S \times A\}$ elements in A, and C, and this can be represented not only by means of a real-valued vector, but by an integer-valued vector, as there are only a finite number of possible choices. All this implies that the Bellman equation is one with finite states, whose solution can be found using numerical methods.

The computations required to obtain the value function and the decision rules are as follows:

Step 1. Initialize value functions $v_0 \in \mathbf{R}^{N_z \times N_s \times N_a}$ to arbitrary initial values. Note that decision rules are also vectors, as there are only a finite number of states. Moreover, they can be represented as integers, as there are only N_a possible saving levels and typically only a decision of whether to work or not to work on the part of the agent.

Step 2. For all $(z, s, a) \in Z \times S \times A$, solve the following problem:

$$\max_{\{1 \leq i \leq N_a, 0 \leq j \leq 1\}} u[a^T \cdot r(z, s, a, n) + n_j\, w(z, s, a, n) + tr(z, s, a, n)$$
$$- e^T \cdot a_i', n_j] + \beta E\{v_0(z', s', a_i') \mid z, s\} \quad (13)$$

Step 3. For all $(z, s, a) \in Z \times S \times A$, update the value function:

$$v_1(z, s, a) = u[a^T \cdot r(z, s, a, n) + n_{j^*} w(z, s, a, n) + tr(z, s, a, n) - e^T \cdot a'_{i^*}, n_{j^*}] + \beta E\{v_0(z', s', a'_{i^*} \mid z, s)\} \quad (14)$$

where the stars refer to the solutions obtained in Step 2. Obviously, these solutions are functions of the state (z, s, a).

Step 4. Check for convergence. If $\max_{\{z,s,a\}} |v_1(z, s, a) - v_0(z, s, a)| > \epsilon$, then $v_0 = v_1$, and go back to Step 2. Otherwise we will assume convergence, and therefore that the solution of the agent's problem has been found.

Once the problem of the agent has been solved, we still have to compute the tax policies, $\tau(z, \mu, S, a, n)$, that equate pretax relative prices to marginal productivity of factor inputs. The required government expenditures needed to implement government policies are computed and checked to be nonnegative.

The remaining object of the equilibrium to be computed is the law of motion of the distribution, $\mu' = g(z, \mu)$. There is no need to compute the function g as such. Given a distribution, μ, and aggregate shock, z, in order to compute next period's distribution, all that is needed is the transition matrix, π, and the decision rule, $a'(z, s, a)$. First note that with a finite state space, a distribution is a nonnegative vector of $(N_S \times N_A)$, elements that sum up to one. In order to compute $\mu'(s', a')(z, \mu)$, the measure of agents that have assets a' and individual shock s' tomorrow given that today the distribution is μ and the aggregate shock z, we look over all pairs $(\bar{s}, \bar{a})$ such that $a' = a'(\bar{s}, \bar{a})$, and we add $\pi(s' \mid \bar{s})\mu(\bar{s}, \bar{a})$.

It is hard to say what is in general a sufficient number of points. One possible rule is to keep increasing them until the findings are no longer sensitive to finer grids. In İmrohoroğlu (1989), for example, the grid considered for the set A has 301 points (for economies 1 and 3), and the one for economy 2 has 602, although they are all equally spaced. Computing the value function is solving a finite-state Bellman equation, which always has a solution. Given the two individual states, the two aggregate states, and the size of the grid in İmrohoroğlu (1989), the value function becomes a vector with 1,204 entries.

An Example with Aggregate Endogenous Variables

In the previous example, relative prices (rates of return of assets) were pinned down by technological factors. This is typically not the case, as the interactions among agents typically play a crucial role in determining prices. For example, in "The Risk Free Rate in Heterogeneous-Agents, Incomplete Markets Economies," Huggett (1993) asks the question, What is the interest rate of an economy where there are no insurance possibilities, and no capital? Equivalent versions of his economy with perfect insurance have an interest rate of $i = 1/\beta - 1$. He argues that a puzzle in the U.S. data is that the risk free interest rate is too low in the post–World War II period (in fact, this rate is about half a percentage point, real, per year). Huggett tries to assess the importance of the role played by the lack of insurance. His paper does not have aggregate uncertainty, but to solve for the equilibrium of his economy a nontrivial market-clearing condition has to be satisfied. Another reason to look at this model is the fact that Huggett developed a different computational method, which proves very useful when market-clearing conditions are important. In his environment, agents can lend and borrow up to certain limits at a rate that is endogenously determined by market-clearing conditions. Agents are subject to idiosyncratic labor market shocks of the same type that İmrohoroğlu's agents' experience. In this case, all assets can be thought of as bonds, a_3, which can be either positive or negative, and the problem of the agent can be written as

$$v(s, a_3; i) = \max_{\{a'_3, c\} \geq 0} u(c) + \beta E\{v(s', a'_3; i) \mid s\}$$

$$\text{s.t.} \quad a_3(1 + i) + w(s) = a'_3 + c \quad (15)$$

Note the explicit dependence of the value function on the interest rate i. A key element will be, of course, determination of the interest rate that clears the bond market.

In this economy, we can define stationary equilibria, as there is no aggregate uncertainty. This definition would be a pair of decision rules for bonds holdings, $a'_3(s, a_3; i^*)$, and consumption, $c(s, a_3; i^*)$, together with a value function, $v(s, a_3; i^*)$, that solve the problem of the agent, an interest rate, i^*, and a stationary distribution, μ^*, that clears the bonds market:

$$\int_{S,A} a'_3(s, a_3; i^*)\, d\mu^* = 0, \quad (16)$$

Let us consider what a nonstationary equilibrium would look like. The first thing to note is that we would need a law of motion for the distribution of agents: $\mu' = g(\mu)$. Also, the agents would need to know the relevant interest rate, i, every period, and this has to be a function of μ, $i(\mu)$. However, in this case, the distribution of agents would enter the problem of the agent, making it a computational nightmare. The nice property about a stationary equilibrium is that it is a fixed point of the function $g(.)$, and we can find it without computing the whole function.

To compute a stationary equilibrium for this type of economy, we define an aggregate excess demand for bonds function, $\varphi(i)$, in the following steps:

Step 1. Fix i.

Step 2. Solve the problem of the agent and obtain $a'_3(s, a_3; i)$.

to the U.S. business cycle. The government follows a monetary policy that sets the inflation rate equal to $r_2(z)^{-1}$ (recall that in our notation r_2 is the real rate of return of a_2, money). This inflation policy either takes the form of a "Phillips curve" (procyclical investment with a 4 percent average and a 2 percent premium in good times) or is at a constant 4 percent rate, it also imposes a labor income tax of $\tau_0 = .25$. The equilibrium of this economy can now be computed, as well as the process for government expenditures. In order to make meaningful comparisons with a perfect insurance environment, the two economies must share the government policies regarding expenditures and labor income taxes. For this, the optimal arrangement that agents can achieve is the maximum utility of the average agent, taking as given the tax sequences and the necessary transfers that the government has to make to balance its budget. This becomes a static problem and it is easily solved once the first-order maximization conditions are specified. Then the results of the simulations of the two environments that share economic policy are processed to obtain meaningful statistics that can be compared.

The Size of Precautionary Savings

The issue of precautionary savings, in particular, those motivated by self-insurance against idiosyncratic risk, has a long tradition in economics. There have been very few attempts to measure the size of this factor. One such attempt is Aiyagari (1993), which sets forth a model similar to that of Huggett (1993). Obviously, this question can be addressed only in an environment where aggregate savings can exist (note that Huggett's is an endowment economy, meaning that there are no possibilities for the economy as a whole to save). There is also a sense in which the level of aggregate savings should affect society's abilities to provide for goods. Aiyagari incorporates these two features by using the technological framework of the standard neoclassical growth model. In this economy, the market-clearing condition is similar to (16)). We write it as

$$\int_{S,A} a_1'(s, a; R(K))\, d\mu^* = K, \tag{18}$$

where $R(K)$ is the real rate of return associated with a level of aggregate capital K. This condition can be stated by saying that in equilibrium, such a level of aggregate capital generates a real rate of return that induces agents' decisions that in turn generate such a level of aggregate capital.

Aiyagari found that the effect of precautionary savings is small: uninsured idiosyncratic risk accounts for only a 3 percent increase in the aggregate savings rate, at least for moderate and empirically plausible parameter values.

Related Non-Business Cycle Research

The type of models described in this paper can also be used to investigate other issues that involve abstraction from aggregate fluctuations. For example, they

have been used to study the real cost of inflation when money plays a well-defined precautionary role and to study the welfare properties of such insurance schemes as unemployment compensation in the presence of moral hazard considerations that arise from the fact that the shock, s, is not observed; only the employment choice, n, is observed.

In the type of economies that we have been reviewing, there is a natural role for public policies that provide some kind of insurance that is infeasible for private agents. For example, an unemployment insurance scheme is Pareto improving in such an environment. However, one of the main criticisms of public insurance mechanisms is that if moral hazard is widespread, private agents will not have the incentives to follow the unrestricted social optimum. Hansen and İmrohoroğlu (1992) attempt to make a quantitative assessment of how important these moral hazard considerations are for the welfare implications of unemployment insurance. In their model, in every period agents get a realization of an idiosyncratic variable that is interpreted as receiving or not receiving a job offer. Moral hazard is modeled as a situation in which the government audits only a fraction of those that claim the insurance. They find that the optimal level of insurance (the fraction of the wage to be paid when unemployed) is very low in the presence of even very small amounts of moral hazard. In İmrohoroğlu the welfare cost of inflation is investigated when money is held for precautionary purposes.

3. Overlapping-Generations Models

Overlapping-generations models provide a natural partition of the population according to features that can be readily observed. Standard data-collecting procedures provide a variety of descriptions of individual behavior by age group. For instance, the cyclical behavior of hours worked varies by age group. It is natural to ask whether these are features that arise naturally in equilibrium models of the business cycle type.

In the absence of measures of how the population is distributed with respect to its attitude toward risk, overlapping-generations models also provide an environment in which to assess the properties of risk sharing in an economy. In these economies it is usually assumed that agents live a large number of periods in order to relate the length of the periods for which the data are collected to the length of the life of people.

A Simple Overlapping-Generations Model

The demographics of this model will be kept at the most simple level. There is a maximum number of periods that an agent can live, I. Compared to the frequencies that interest business cycle researchers, the age structure of the population moves very slowly; therefore populations are taken to have a fixed age distribution. We normalize the size of the first cohort, μ_1, to one. The size of age i cohort, μ_i, becomes $(1+\lambda_\mu)^{-i}$, where λ_μ is the rate of growth of the population.[11] Agents are endowed with one unit of time per period, which can be enjoyed as leisure or can be used as an input to produce, jointly with capital, a consumption good. One unit of time of an age i agent can be transformed into ϵ_i units of labor input, where $\epsilon = \{\epsilon_1, \ldots, \epsilon_I\}$ is a vector of exogenously given parameters. Preferences are standard, and are represented by the expected discounted sum of a strictly concave current utility function of leisure and a consumption good. For a newborn agent this can be written as

$$E\left\{\sum_{i=1}^{I} \beta^{i-1} u(c_i, l_i)\right\}, \tag{19}$$

while an agent of age i has remaining utility of $E\left\{\sum_{j=i}^{I} \beta^{j-1} u(c_j, l_j)\right\}$.[12]

There is a neoclassical production function, $f(K, N)$, that is affected by a multiplicative shock, $z \in Z = \{z_1, z_2, \ldots, z_{N_z}\}$. This shock follows a Markov process with transition matrix π and is observed at the beginning of the period. Output can be used either for consumption in the same period that production takes place or for increasing the capital stock next period. Capital depreciates at rate δ, and for notational simplicity, we will assume that undepreciated capital can also be used for consumption purposes. If we denote next period's variables by primes, all this can be written as

$$\sum_{i=1}^{I} c_i\, \mu_i + K' = z\, f(K, N) + (1-\delta)K. \tag{20}$$

These models are used, among other things, to explore the quantitative relevance of different market arrangements. This is done by comparing the equilibrium allocations of these alternative market institutions. In particular, these models are very useful in exploring the issue of whether the existence of possibilities of insurance against aggregate fluctuations is quantitatively important or not. A useful way of formulating the market structure is to allow trade in securities that deliver one unit of the capital good next period. At the beginning of the period, the shock is observed, then the contracts for the securities are honored, and finally, production takes place. This choice of timing of delivery is somehow nonstandard, as the securities do not deliver units of consumption, but units of the capital good before production takes place. As we will see, this timing implies simple restrictions on security prices. A general way of setting up different degrees of market completeness is to allow for trade on securities contingent on the elements of $\mathcal{F} = \{\varphi_1, \varphi_2, \ldots, \varphi_{n_\varphi}\}$, a partition[13] of Z, with $n_\varphi \le N_z$, as for all $n \in \{1, \ldots, N_z\}$, there exists $m \in \{1, \ldots, n_\varphi\}$, such that $z_n \in \varphi_m$. When $\mathcal{F} = \{Z\}$, savings can only be made uncontingent, as there are no possibilities of signing contracts that depend in any ways on the value of next's period's shock.

When $\mathcal{F} = \{z_1, z_2, \ldots, z_{N_z}\}$, the full set of securities can be traded. This implies that agents face the following budget constraints:

$$a_1 = 0, \tag{21}$$

$$a_i R + W(1-l_i)\epsilon_i = \sum_{\varphi \in \mathcal{F}} b_i^\varphi q^\varphi + c_i, \tag{22}$$

$$a'_{i+1}(z') = b_i^{\varphi(z')}, \quad \text{for } \varphi \in \mathcal{F},\ i = 1, \ldots, I, \tag{23}$$

and

$$b_I^\varphi \ge 0, \tag{24}$$

where a_i is the net wealth accumulated so far by an age i agent, b_i^φ constitutes the amount of state-contingent capital goods bought by agents of age i, q^φ is the price of the state-contingent assets, R is the return on asset holdings, and W is the price of one unit of labor input in terms of the consumption good. Condition (23) specifies that net wealth tomorrow depends on the realization of the shock tomorrow.

Note that aggregate gross savings (including undepreciated capital) are $zf(K, N) + (1-\delta)K - \sum_i \mu_i c_i$, and they become next period's aggregate capital, K'. The decisions that lead to the determination of K' are all made this period; therefore K' cannot depend in tomorrow's value of the shock, z'. On the other hand, aggregate savings can also be obtained by adding up contingent savings of all individuals: $(\sum_i \mu_i b_i^\varphi)$. Note that this is true for every φ. This implies that aggregate contingent savings has to be the same for all $\varphi \in \mathcal{F}$, as their size is determined today by the total amount of nonconsumed real resources. In other words, next period's aggregate capital stock cannot be made contingent on tomorrow's shock, although its distribution can. This is just another way of stating the market-clearing condition for the securities. Hence, the aggregate feasibility constraint becomes:

$$\sum_i \mu_i (b_i^\varphi + c_i) = zf(K, N) + (1-\delta)K, \quad \text{for all } \varphi \in \mathcal{F},\ \text{all } z \in Z. \tag{25}$$

In the feasibility constraint, aggregate inputs are, respectively,

$$N = \sum_i \mu_i\, (1-l_i)\epsilon_i \quad \text{and} \quad K = \sum_i \mu_i\, a_i. \tag{26}$$

Note that the price of a unit of capital for sure for next period can be acquired by the purchase of one unit of b^φ for all $\varphi \in \mathcal{F}$, at a price of $\sum_\varphi q^\varphi$. Note also that this capital good can be obtained by staying outside the market and saving one unit of the good. Obviously, this restricts the set of possible prices of the securities. If $\sum_\varphi q^\varphi > 1$, an agent could sell arbitrarily large amounts of these securities by buying the good, and delivering it next period while making arbitrarily large profits. If $\sum_\varphi q^\varphi < 1$, an agent could do the opposite operation and also obtain arbitrarily big consumption possibilities. These two properties together imply that in equilibrium it must be the case that $\sum_{\varphi} q^{\varphi} = 1$, since unbounded consumption possibilities cannot be feasible.

Equilibrium Defined

Equilibrium is defined recursively. The state of the economy is characterized by the economywide shock, $z \in Z$, and by the distribution of asset holdings by agents in each age group, $k \in \mathcal{A} \subset \mathbf{R}^I$. The state variables for any given agent are the agent's own asset holdings, $a \in \mathcal{A}_i \subset \mathbf{R}$, and the economywide state, (z, k). Note that since all agents of the same age share the same strictly concave utility function and convex choice set, they do the same thing in equilibrium, that is, $a_i = k_i$. Therefore, the aggregate laws of motion for the economy depend only on the decision rules of the agents and the process for the shock, and, since the latter follows a Markov process, its current value is all the information needed to know tomorrow's distribution. With the type of securities chosen, next period's state, k', depends on next period's shock, z'. Then $k'(z') = b^{\varphi}$, where $\varphi(z')$ is the element of $\mathcal{F}$ that contains z'. The value functions involved in the definition are indexed by the age of the agent, and their value is expected remaining utility. A recursive competitive equilibrium is a set of decision rules, $b_i^{\varphi}(z, k, a), c_i(z, k, a)$, $l_i(z, k, a)$, for all $i = 1, \dots, I$, $\varphi \in \mathcal{F}$, $z \in Z$, $k \in \mathcal{A}$, $a \in \mathcal{A}_i$; a set of pricing functions $W(z, k)$, $R(z, k)$, $q^{\varphi}(z, k)$, for all $\varphi \in \mathcal{F}$, $z \in Z$, $k \in \mathcal{A}$; a set of value functions, $v_i(z, k, a)$, $i = 1, \dots, I$, $z \in Z$, $k \in \mathcal{A}$, $a \in \mathcal{A}_i$; a law of motion for the capital stocks, $k'_{i+1}(z') = g_i^{\varphi(z')}(z, k)$, $i = 1, \dots, I-1$, $z \in Z$, $k \in \mathcal{A}$, $z' \in Z$; and a pair of functions for aggregate variables. $K(z, k)$, $N(z, k)$, $z \in Z$, $k \in \mathcal{A}$, such that the following conditions are satisfied:

1) The allocation is feasible, that is, for all z, k, and all φ,

$$\sum_i \left[b_i^{\varphi}(z, k, k_i) + c_i(z, k, k_i) \right] \mu_i = zf[K(z, k), N(z, k)] + (1 - \delta) \sum_i k_i \mu_i. \quad (27)$$

2) Factor prices are competitive, that is, they are the marginal productivities:

$$W(z, k) = zf_2[K(z, k), N(z, k)], \quad (28)$$

and

$$R(z, k) = 1 - \delta + zf_1[K(z, k), N(z, k)]. \quad (29)$$

3) Given the law of motion for capital stocks, the price functions, and the transition for z, decision rules for age i agents solve their maximization

problem:

$$\left\{ b_i^{\varphi}(z, k, a), c_i(z, k, a), l_i(z, k, a) \right\} \in$$

$$\arg\max_{b^{\varphi}, c, l} u(c, l) + \beta E \left\{ v_{i+1}\left(z', g^{\varphi(z')}(z, k), b^{\varphi(z')}\right) \mid z \right\}$$

$$\text{s.t.} \quad aR(z, k) + (1 - l)\epsilon_i W(z, k) = c + \sum_{\varphi \in \mathcal{F}} b^{\varphi} q^{\varphi}(z, k). \quad (30)$$

4) The value functions are generated by the policy functions and $v_{I+1} = 0$, and the budget constraint is satisfied:

$$v_i(z, k, a) = U\left[c_i(z, k, a), l_i(z, k, a)\right] + \beta E \left\{ v_{i+1}\left[z', g^{\varphi}(z, k), b_i^{\varphi}(z, k, a)\right] \mid z \right\}, \quad (31)$$

and

$$aR(z, k) + [1 - l_i(z, k, a)]W(z, k)\epsilon_i = c_i(z, k, a) + \sum_{\varphi \in \mathcal{F}} b_i^{\varphi}(z, k, a) q^{\varphi}(z, k). \quad (32)$$

5) The law of motion of the capital stocks is generated by the decision rules of the agents:

$$b_i^{\varphi}(z, k, k_i) = g_{i+1}^{\varphi}(z, k), \quad \text{for all } z' \in Z, \text{ all } i = 1, \dots, I-1. \quad (33)$$

6) Aggregate functions K and N are generated by aggregation and the decision rules of the agents:

$$K(z, k) = \sum_i \mu_i k_i, \quad (34)$$

and

$$N(z, k) = \sum_i \mu_i [1 - l_i(z, k, k_i)] \epsilon_i. \quad (35)$$

If the agent problem (30) of condition 3 has been explicit about the fact that agents can arbitrage by storing the capital good themselves, then a requirement on prices that prevents arbitrage opportunities does not have to be imposed as it is implicit in the fact that there is a well-defined solution to (30), which can only happen if $\sum_{\varphi \in \mathcal{F}} q^{\varphi}(z, k) = 1$. If we are not explicit about that fact we can impose such a requirement as an equilibrium condition. It is also a good idea to impose such a condition for computational reasons. We write this as equilibrium condition 7:

7) There are no arbitrage opportunities:

$$\sum_{\varphi \in \mathcal{F}} q^{\varphi}(z, k) = 1, \text{ or } q(z, k) \in \Delta^{n_{\varphi}}, \quad \text{for all } z \in Z,\ k \in \mathcal{A}. \tag{36}$$

It is also important to explain the role played by market-clearing conditions. The above definition includes only those for the factors of production, but we also have markets for securities, so where are their market clearing conditions? The answer is, of course, that it is implicit in condition 1, the feasibility requirement. Note that as it is written, it has to hold for all φ. Thus every period, we have not only one feasibility requirement but n_{φ} of them. Therefore, if we have a set of functions, $b^{\varphi}(z, k, a)$ that satisfy condition 1, then they will also satisfy

$$\sum_i \mu_i b_i^{\varphi}(z, k, k_i) = \sum_i \mu_i b_i^{\varphi'}(z, k, k_i), \quad \text{for all } \varphi, \varphi' \in \mathcal{F}. \tag{37}$$

Computation of the Equilibrium

The typical procedure used to compute the equilibrium law of motion for these economies consists in calculating a linear-quadratic approximation of a set of reduced-form utility functions around the steady state. This requires us first to find the steady state. As Auerbach and Kotlikoff (1987) show, this involves finding the solution of a nonlinear equation of one variable. The easiest way to do it is to define an equation that given a capital-labor ratio, returns the capital-labor ratio that is generated by the economy.[14]

Once this is done, we know the levels of asset holdings, consumption, and leisure at every age. We then substitute factor prices for the expressions of marginal productivities in terms of the shock, z, the asset distribution, k, and aggregate employment, N. We set contingent prices, q^{φ}, equal to the conditional probabilities of the events the associated securities are contingent on (which will depend on the current shock, z), and contingent security holdings, b_i^{φ}, equal to the steady state levels of asset holdings at each age. All this allows us to use the budget constraint to substitute for current consumption, and we obtain a current utility function $R_i(z, k, a, N, q^{\varphi}, b^{\varphi}, l)$. We approximate these functions by quadratic functions. Next, iterations are performed on the set of value functions. In each iteration, expressions for $(N, q^{\varphi}, k'^{\varphi})$ as functions of the aggregate state (z, k) have to be found. This requires the inversion of a large matrix. Chapter 2 gives a generic description of how to perform these procedures. However, for a detailed explanation of how to compute the equilibria in overlapping-generation economies with a variety of market structures, see Ríos-Rull (1993a). For details on how to incorporate sophisticated demographic structures, see Ríos-Rull (1992b).

Altig and Carlstrom (1991) used a different procedure, which can be described not as computing the exact (up to computer accuracy) equilibria of an approximated economy, but as an approximation to the equilibria of the original economy. Their method is based on the algorithm developed by Auerbach and Kotlikoff (1987) to compute transition paths between steady states. The process can be described in the following way. Starting from the steady state, an innovation in the exogenous shock occurs. Assume that from that period on the economy will behave deterministically, with the value of future shocks set at their conditional expectation. Then compute the transitional path to the steady state. From this transition, obtain the value of prices and quantities for the first period. Next, a new value for the exogenous shock is obtained, and the process is repeated. This method presents two problems. The first is that it applies the notion of certainty equivalence in a context where it does not apply. The other problem is that it is very computer intensive, with the computations proportional to the length of the sample that is produced.

Special Calibration Issues

In these economies some special calibration issues arise. The production side of the economy is the same as that of representative-agent economies, so the same calibration considerations apply. The important difference is that the demographic structure has to be specified. For this, the procedure is to construct a stable population that shares the U.S. birth rates and age-specific mortality. (To explicitly include mortality requires certain adjustments in the model, as we should face the issue of the possible role played by markets for annuities; see Ríos-Rull [1992a] for details). The demographic structure has strong implications for the preference parameters. The discount rate, β, no longer should be calibrated to match historical average rates of return. Its value must be chosen, based on estimation from individual data. In this respect, Hurd (1989) estimates individual preference parameters, taking into account demographic features that make his estimates very attractive for our purposes. The values for the per-period utility function also must be chosen from microeconomic studies, but these are not typically detailed enough to be able to make the parameters age dependent.

A very important set of parameters is that given by ϵ, the vector of efficiency units of labor by age group. Its value can be obtained from different sources: Ríos-Rull (1992a) uses Hansen's (1993) Consumer Population Survey (CPS) data to construct it, while Auerbach and Kotlikoff (1987) and Altig and Carlstrom (1991) use Welch's (1979) regression coefficients. It is also important to note that under this assumption of how individual hours aggregate into the labor input, the standard measure of the Solow residual is no longer valid, and its process has to be reevaluated. Details of these and related issues are discussed in Ríos-Rull (1992a).

Review of the Literature

Life Cycle Economies and Aggregate Fluctuations

It is a well-known fact that the volatility of hours worked is not constant across age groups, but is highest for the youngest and the oldest. Ríos-Rull (1992a) studies the general business cycle properties of these economies and, in particular, the relative volatility of hours across age groups. His findings were that the business cycle properties of these economies are very similar to those of the representative-agent models, which are discussed elsewhere in the book. With respect to the behavior of hours worked, he found that aggregate hours have small volatility, a finding that is similar to other business cycle models with no nonconvexities in the labor choice. The relative volatility of hours worked across age groups shows some differences between the data and the model. While volatility of the hours worked by young agents in the model is greater than that of prime-age agents, as they are in the data, agents between 45 and 64 years of age show a much higher volatility of hours than that of the 25–44 age group, contradicting the behavior observed in the data.

Inflation, Personal Taxes, and Real Output: A Dynamic Analysis

The objective of Altig and Carlstrom (1991) is to answer the question, What consequences do interactions between inflation and the nominal taxation of capital income have for the cyclical behavior of the economy? In their model, there is no role for money. Inflation simply introduces a distortion in the measurement of individual capital income. In particular, an inflation rate of π overstates capital income in an amount equal to $\pi/(1+\pi)$ per unit. They impose a progressive tax structure that mimics the 1965 U.S. tax code. Given the role that inflation has in this model, it is sufficient to calibrate its process with a univariate representation. They estimate an autoregressive process of order two, with an associated average inflation of 4 percent. They find that with progressive taxes, the volatility of inflation, on top of the exogenous variation in the Solow residual, drastically increases the volatility of hours worked, compared to a constant inflation process (about 80 percent more), and moderately increases the volatility of consumption, while leaving the volatility of investment, output, and capital roughly unchanged.

The inclusion of the inflation effects transforms equation (21) into

$$[a_i R + W(1-l_i)\epsilon_i]\{1 - \tau[a_i R \frac{1+2\pi}{1+\pi} + W(1-l_i)\epsilon_i]\}$$

$$- \tau[a_i R \frac{1+2\pi}{1+\pi} + W(1-l_i)\epsilon_i]a_i R \frac{\pi}{1+\pi} + tr = \sum_{\varphi \in \mathcal{F}} b_i^{\varphi} q^{\varphi} + c_i, \quad (38)$$

where the first term is the after-tax income, taking into account the progressive nature of the tax, the next term reflects the distortion that inflation induces on the tax base through the capital income and the last term of the right-hand side is the transfer required to balance the government budget given that there are no public expenditures.

The state space has to be enlarged to include current and lagged inflation since inflation is supposed to follow a second-order autoregressive process. The market structure considered does not include contingent markets; therefore, $\mathcal{F} = \{Z\}$. The definition of the equilibrium also requires a transfers function, $tr(z, k, \pi, \pi_{-1})$, that the agents use to solve their problem, and that has to balance the government budget.

On the Quantitative Importance of Market Completeness

In Ríos-Rull (1993a), the question explored is, How different are the equilibrium allocations across economies that differ only in whether there are Arrow securities for economywide productivity shocks in the absence of idiosyncratic shocks of any kind? The economies studied vary only on the specification of $\mathcal{F}$, ranging from $\mathcal{F} = \{Z\}$, where savings can only be made noncontingent, to $\mathcal{F} = \{z_1, z_2, \ldots, z_{N_z}\}$, where markets are dynamically complete. The answer found is that the aggregate series obtained for the various market specifications are very similar to one another, to the point that it would be almost impossible to identify the type of economy from this type of data alone. If data on individual allocations are observed, slight differences across economies arise as agents actively engage in risk sharing trades. These differ, of course, depending on the trade possibilities allowed in each economy. In any case, differences are very small. There are differences across individual asset holding data, that can be readily observable. This seems to indicate that different market structures offer different mechanisms that allow an economy to reach almost the same objective: an allocation as close as possible to a Pareto optimum.

Related Non–Business Cycle Research

The ability to compute nonstationary equilibria in overlapping-generations economies where agents live a large number of periods makes it possible to address some other interesting questions. For example, we can study the role that demographic changes play in the economy. Of special interest is their influence on capital accumulation and social insurances. The methods developed can address these because they take into account the effect that the induced changes in prices have on the behavior of agents. Ríos-Rull (1992b) assesses the magnitude of the changes in savings rates that can be expected the result from the current aging of the Spanish population. Similar topics have been studied by Auerbach et al. (1989) and by Danthine and Surchat (1991), where certain shortcuts have been assumed to simplify calculations. Ríos-Rull (1993b) uses a simple computable overlapping-generations model with two-period–lived agents to study some joint facts about the relation between wages and hours worked from four different points of view: cross-section, long-term, cyclical, and age-specific. This model gives some insights into how heterogeneity arises in environments where all agents start equal, but it is very different from the models in this book, as there is no physical capital and periods are very long compared with the quarterly or yearly frequency typical of business cycle models.

4. Other Types of Heterogeneity

The Cyclical Behavior of Factor Shares

Economies with agents differing in preferences have been used to address some additional business cycles issues. Gomme and Greenwood (1993) use such an economy to study the cyclical behavior of factor shares. The long-term behavior of the factor shares is basically constant, as is the capital output ratio, and the rate of return on capital. This restricts the elasticity of substitution between capital and labor to be one, and leaves Cobb-Douglas as the primary candidate for the production function. But this production function implies that factor shares are constant at every frequency, while the data show that labor share moves countercyclically, while profits are procyclical. To address this issue, Gomme and Greenwood introduce labor contracts as risk-sharing arrangements between two types of agents: workers and entrepreneurs.

In their model, agents are born being of one of the two types. Entrepreneurs operate a constant-returns-to-scale stochastic technology, which requires their own labor input as well as workers' labor input and real capital. Preferences are of the type described by Epstein (1983) as stationary cardinal utility, which allows deterministic economies to have a unique invariant distribution of wealth across agents. The key property of these preferences is that future discount factors become a decreasing function of current utility. Preferences of the two types are calibrated so that the entrepreneurs constitute the rich agents, owning 25 percent of the wealth, and constitute 1 percent of the population. A competitive equilibrium with a full set of Arrow securities is calculated. We can think of this equilibrium as having spot markets for labor and capital and Arrow securities, or as the Arrow securities' being embodied in the labor contract, which now is contracted one period in advance, or as the entrepreneurs issuing bonds, which are held by the workers, to acquire capital. Any of these market structures implements the complete-markets allocations, although they imply very different behavior for wages, and consequently for the labor share. The first implementation gives a constant labor share, while the second and the third result in a countercyclical behavior.

Heterogeneous Agents and the Risk Premium

One of the great puzzles in aggregate economics is the size of the premium that risky assets (namely, stocks) have over relatively riskless assets, such as government bonds (see Mehra and Prescott 1985). Standard models (that is, representative-agent endowment economies with time-separable preferences and a process for consumption that shares the stochastic properties of aggregate consumption in the post World War II United States) systematically generate premiums that are far too small. In trying to account for this anomaly, the standard representative-agent model has been modified along several dimensions. The one direction that is relevant to us tries to explore the issue of whether the size of the premium can be

accounted for by lack of markets in multiagent worlds where their endowments are not perfectly correlated. In independent research. Lucas (1990), Ketterer and Marcet (1989), Marcet and Singleton (1990), and Telmer (1992) have constructed models with two agents where individual endowment processes are not perfectly correlated. They explore the nature of the equilibrium allocations and prices under various market structures that include trading on some assets and find that they are very similar to the ones obtained in complete market settings. This result can be explained by noting that in economies that are calibrated to quarterly periods, the average value of their income stream (which is of the order of magnitude of the returns of the assets that agents can use to smooth consumption) is around twenty times the value of average quarterly consumption. This is a very big cushion for smoothing consumption in the presence of reasonable perturbations on the endowments stream. A related paper of Danthine, Donaldson, and Mehra (1992), studies the implications of the fact that stockholders bear a disproportionate share of output uncertainty in a model with non-walrasian features.

Asymmetric Information and Limited Commitment

Marcet and Marimon (1992) have studied the set of implementable allocations in environments where the actions of an infinitely-lived borrower are unobservable and there is no commitment technology to repay the loans. They do this in the context of a small country where the allocation of national output between consumption and investment and the technological shock are both unobservable. They are able to characterize and compute the optimal contract between risk–neutral agents and the small country. They find that when the rationale for borrowing is to increase the capital stock, there is no improvement over autarky. However, when loans are used to smooth consumption, the optimal contract scheme improves significantly over autarky.

5. Current Research Agenda

So far, all models with idiosyncratic shocks and an absence of complete insurance possibilities are either silent on the distribution of wealth (in the sense that it is imposed ex ante) or have no capital. As distributional issues become more interesting (among other features, they have tremendous relevance for public finance and the nature of the tax system), the key innovation will be to develop methods to compute equilibria in economies with capital where the distribution of wealth and income is endogenous. With such a tool, the cyclical behavior of the distributions of income and wealth can be addressed. Also, the redistributional effect of a variety of policies can be quantitatively evaluated, as such methods allow one to compute not only the optimal reaction of agents to the policies but also the ultimate outcome that results from their interaction. To this end, there is some very preliminary work that tries to approximate the distribution of wealth in different ways that might be able to keep both computational costs and the accuracy at reasonable levels.

The radical technological change that affects the computer industry provides economists with increasingly cheaper and more powerful machines. At the same time, computational literacy of economists provides strong externalities in terms of an accumulated body of procedures that become available for all researchers. This is one of the main reasons for the rapid development of this class of models over the last two or three years. As there are no indications of these processes stopping or slowing down, the outlook for future success seems very bright.

Notes

Part of this material has been used to teach a Ph.D. course in Carnegie Mellon University. I am grateful to its students. Thanks also to the comments of Paul Gomme, Jeremy Greenwood, Mark Huggett, Pat Kehoe, Finn Kydland, Lee Ohanian, and the editor of this book, Tom Cooley.

1. This literature takes the fact that some markets are missing as given and studies the properties of such economies; it does not attempt to explain why they are missing.

2. Constantinides and Duffie (1991) and Zin (1992) have examples of heterogeneous-agent models with incomplete markets that can be represented as homogeneous-agent ones. They require the shocks to be such that even in the presence of markets, agents do not gain anything by trading. In other words, in these markets, there is no consumption smoothing, as marginal utilities are restricted to remain constant.

3. As of early 1994, there is preliminary work suggesting that using a small set of moments to approximate the distribution is a promising avenue. In particular, Krusell and Smith (1994) suggest that an affine law of motion for aggregate wealth works very well. Castañeda, Díaz-Giménez, and Ríos-Rull (1994) are using these ideas to study the cyclical behavior of the income distribution. See also den Haan (1993).

4. The reason why μ' is also a function of z' is because π, the joint law of motion of (s, z), might have the property that the distribution of s' can be jointly conditional on $\{s, z, z'\}$ so that the measure of agents with each shock might vary with z'. Of course, the marginal distribution over asset holdings $\mu'(A_0, S)$ for $A_0 \subset A$ will not depend on z'.

5. İmrohoroğlu (1989) actually does not compute the equilibrium of the economy with perfect insurance, but of an economy that does not have storage technology, so her calculations actually provide a lower bound for the difference of the cost of business cycles, as storage can only reduce the costs in perfect-insurance economies.

6. It is possible to pose this type of problem as a linear maximization problem, whose solution is found with the simplex algorithm.

7. Note that by Walras's Law, clearing of the bonds market implies clearing of the goods market.

8. Chapter 3 in this book discusses related computational methods based on approximations to the Euler equation.

9. Actually, this condition is not exactly true. If for all a_3 above the lower bound, equation (17), is negative, then the a_3' that is chosen is precisely the lower bound, and equation (17) is not set to 0.

10. Note that since with perfect insurance, there is no inflation tax and total output is typically different than in the liquidity-constrained environment, there is a need for these transfers to balance the budget.

11. For simplicity, we abstract from early death of the agents. See Ríos-Rull (1992b) for a discussion of the demographics in this type of economy.

12. We could think of more general preference structures by indexing the current return utility function, u, by age, obtaining $u_i, i \in \{1, \cdots, I\}$, by noting that combinations of consumption and leisure are not regarded equally across different ages. Also, we could substitute a common discount factor, such as β, for a system of weights that treat each age's utility in a different manner.

13. Given a set Z, a partition $\mathcal{F}$ of Z is a family of subsets, φ_j, $j \in J$, of Z, such that their union is the set Z, $(\cup_j \varphi_j = Z)$, and they have pairwise empty intersection $(\varphi_j \cap \varphi_{j'} = \emptyset$, for all $j, j' \in J)$.

14. Given K/N, we obtain relative prices W and R. If we assume them constant forever, it is straightforward to find the amounts of asset holdings and labor input that agents want at each age. Aggregating factors of production, and obtaining their ratio yields a new capital-labor ratio.