Journal of Economic Dynamics and Control 18 (1994) 849–864. North-Holland

Heterogeneous agents in quantitative aggregate economic theory*

Finn E. Kydland

Carnegie Mellon University, Pittsburgh, PA 15213, USA

Received September 1993, final version received November 1993

Modern macroeconomics, with increasing frequency, is devoting attention to the fact that agents are heterogeneous, especially in their role as providers of labor input into production. These differences can take various forms: the skills for market production are different; some workers are and some are not engaged in market activity in any given period; people are at different stages of the life cycle; people are faced with uninsurable idiosyncratic shocks, perhaps combined with liquidity constraints; and so on. In some cases, the equilibrium still can be computed from a stand-in planner's problem. The most exciting recent development, however, is the dramatically expanded ability to compute equilibriums for interesting model economies for which that is not the case. Examples are given.

Key words: Computable equilibrium models; Heterogeneous agents
JEL classification: C10

1. Introduction

These conferences on economic dynamics and control have a long history. The present one is listed as number 15, but these were preceded by similar conferences whose proceedings were published in the *Annals of Economic and Social Measurement*, which ceased publication a year before this journal started. Personally, I was especially pleased to be asked to give this keynote speech, since

Correspondence to: Finn E. Kydland, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA 15213-3890, USA. Fax: (412) 268-7357.

*Keynote address delivered at the 15th Annual Conference of the Society for Economic Dynamics and Control, Nafplio, Greece, 22–25 June, 1993. With minor changes, the speech format has been maintained in this written version.

Editors' Note
SEDC annual meeting in Nafplio, Greece, 1993: The JEDC is pleased to publish Finn Kydland's keynote address and summaries of the invited plenary session talks by Masanao Aoki, Manuel Santos, and Kenneth West. We also include a copy of the program.

0165-1889/94/$07.00

this month represents the 20th anniversary of the first conference at which I had a paper on the program and, moreover, that was also a conference on economic dynamics and control – in Chicago in June 1973. I suppose I had more fun the next time, in Cambridge, Mass., in 1975. There, I presented some research in progress with Ed Prescott. I started the presentation, naturally, with the title of the paper – something about the time inconsistency of optimal government policy. And all hell broke loose, with words covering the whole spectrum from 'misguided' to 'plain wrong'. I guess we eventually managed to convince a few people, so that, when again I attended an economic dynamics and control conference, in Austin a couple of years later, that research was taken seriously enough that, in spite of – or perhaps because of – the somewhat negative perspective, I was invited to appear on a panel to debate the role of optimal control theory for policy evaluation.

I went also to Princeton in 1980, but as it turned out, that was my last SEDC conference until Capri, two years ago. I'm not sure about the reason for the long break. It seemed perhaps that, at least in the U.S., for a while the conferences did not have quite the same visibility. Perhaps they were still as important in Europe – I don't know. What made me take notice again, a few years ago, was a letter from Tom Sargent expressing his commitment to invigorating the society. And was he successful in reviving interest! We are all thankful for his efforts, because the SEDC can play an extremely important role in the development and use of economic dynamics, which by now is used in just about all fields of economics.

One awkward aspect of my first participations at these conferences, especially the second and third, was that my message was interpreted to be quite critical and to apply directly or indirectly to many of the other papers presented at those very same conferences – not a pleasant situation. Today, I'm in a much happier position. There is every reason to be upbeat about a lot of what we're seeing at conferences like this and about the future of quantitative economic theory. So, a main purpose of my talk is to attempt to give a flavor of what one can do, of how it is now feasible to construct computable general equilibrium model economies in which agents are heterogeneous, at least in the sense that their outcomes are different, and to compute stochastic time paths for these economies. This ability has expanded the scope of questions to which we can obtain quantitative answers. The main dimension considered will be the actors' decisions as related to the labor input and the context mainly business cycles.

2. Computational experiments

Given the emphasis placed historically on computation at these conferences, let me first say a few words about an important tool – the computational experiment. (I could call it an econometric tool, but many, for some reason, find that provocative.) A computational experiment is the act of placing a model economy's people in the environment desired by the experimenter and then recording the time paths of their behavior. To carry out such experiments with actual people of course would be prohibitively expensive. Instead, model environments are constructed in the form of computer programs. A model environment is designed to address a particular question or issue. Examples of the types of questions for which the computational experiment has been used are:

(i) What are the welfare consequences of changing policy A to policy B?

(ii) How much of fact(s) X is accounted for by Y?

(iii) Does 'established' theory display feature Z?

(iv) Does the introduction of feature F in standard theory account for part of deviation D from standard theory?

The amount of detail included in a model economy depends on the question being addressed as well as on what is computationally tractable. This situation is no different from that in the natural sciences, for example, in which computer simulations have a long history as a scientific tool. In his overview of climate modeling for the *Scientific American* some years ago, Schneider states: 'Although all climate models consist of mathematical representations of physical processes, the precise composition of a model and its complexity depend on the problem it is designed to address.' And later: 'Often it makes sense to attack a problem first with a simple model and then employ the results to guide research at higher resolution.' In the natural sciences as in economics, confidence in a particular framework or approach is gained when its use time after time provides interesting answers as well as clear demonstrations of remaining deviations or anomalies relative to that framework, at least deviations that appear possible to reconcile within that framework through further research.

Several developments over the past 30 or 40 years have contributed to making the computational experiment feasible in economic contexts in which intertemporal behavior under uncertainty is at the heart of the question being addressed. Among these developments is the progress in our understanding of decision theory and of general equilibrium theory. These tools already have been relied upon extensively in contexts in which the focus was not so much on intertemporal movements [see, e.g., the overview in Whalley (1985)]. For explicitly dynamic analysis under uncertainty, the development of recursive methods have been especially useful [see Stokey and Lucas (1989)]. With these methods, the elements being computed are decision or policy rules that describe, for a given environment, the decisions made by rational individuals as functions of the state of the economy. Typical members of the state vector, in addition to sufficient information about current and past shocks, are stocks of various sorts, such as productive capital in the business sector, inventories, consumer durables, and human capital. Once the equilibrium decision rules have been computed, the equilibrium aggregate behavior of the people in the model economy is fully described, and the experimenter easily can generate time series on the computer, time series that are analogous, for this environment, to national income and product accounts series.

The model environments that I have in mind and which have been used so far all share certain characteristics. They are inhabited by a large number of people whose decision problems are described explicitly. In addition to the household sector, the business sector usually plays a central role. For some questions, a government sector or a foreign sector is included as well. The assumption that everyone is alike is reasonable for some purposes but not for others. Some questions, for example those for which the demographic pattern may be important, dictate that abstractions with heterogeneous agents be used. In either case, the equilibrium behavior of these economies can display rich intertemporal patterns of movement that mimic those of actual economies to a degree sufficient to make the answers credible and influence the way economists think about aggregate issues.

An important feature of these economies is that growth and business cycles are integrated. Business cycle models are stochastic versions of the neoclassical growth theory. A dictionary definition of *theory* is 'a formulation of apparent relationships or underlying principles of certain observed phenomena which has been verified to some degree' (Webster's New World Dictionary). Neoclassical growth theory certainly satisfies that criterion. Central to this theory is the description of aggregate production possibilities. Adding to the business sector an explicit description of the household sector, with its focus on the time-allocation decision, makes this an internally consistent framework for addressing business cycle questions, as well as other questions of interest to the macroeconomist.

In order to place the model economy's people in the desired situation for the experiment, the economic environment typically is described in the form of a computer program. Using the neoclassical growth framework, the parameters are those describing preferences, technology, information sets, and institutional arrangements, including policy rules.

As a simple illustration, consider the most basic growth framework with stochastic technology shocks, but extended to include a time-allocation decision. There are no externalities and all households are alike. Under those circumstances, the second welfare theorem implies that the equilibrium can be computed by solving the optimization problem of a stand-in planner whose objective function corresponds to the utility function of the typical household. Thus, the competitive equilibrium can be found by maximizing

$$\mathrm{E}\sum_{t=0}^{\infty} \beta^t u(c_t, l_t)\,,$$

subject to the constraints on goods and time,

$$c_t + x_t \leq z_t f(n_t, k_t), \qquad n_t + l_t \leq 1,$$

and the laws of motion,

$$k_{t+1} = (1 - \delta)k_t + x_t, \qquad z_{t+1} = \rho z_t + \varepsilon_{t+1},$$

where c_t is consumption, n_t and l_t are time spent in market and nonmarket activity, respectively, x_t is investment, k_t is the capital stock at the beginning of period t, z_t is the level of technology, and ε_t is a random disturbance to the technology with positive mean. With explicit parametric forms for the utility and production functions, with values assigned to the parameters, and with a probability distribution for the shocks, one can use this economy to perform computational experiments. The choice of parameter values often is quite straightforward.

The objects that first need to be computed are the aggregate decision functions for c, x, and n. In other words, the decisions are viewed as functions of the list of state variables that provide sufficient information about the state of the economy, here k and z. These three decision functions, the laws of motion for the two state variables, and the probability distribution for the shocks are what we need for the computer to generate equilibrium time series for the model economy. For each t, given k_t and z_t inherited from the previous step, this is done by computing the values of c_t, x_t, and n_t from the decision functions. The computer then generates a random number from the distribution for ε, and uses the laws of motion to update the two state variables for period $t + 1$.

For most applications, a useful approach when making a comparison with actual data over a period of a particular length, say T quarters, is to determine from the model economy the sampling distribution of T-period samples. That is, one produces repeated samples of length T from the model economy. One can then compute statistics that indicate the degree of consistency across samples of key characteristics of the model economy. These statistics can be of help in judging the reliability of the quantitative answer obtained from a particular set of experiments, and the presence of any remaining anomalies relative to the data.

Subject to being computable, the selection of model environment depends on the question being addressed. Model-economy selection should not depend on the answer provided. Moreover, searching within some parametric class of economies for the one that best fits aggregate time series according to some statistical criterion makes little sense. It is hard to think of interesting questions for which such practice would provide an answer. In some cases, it would be downright silly. For example, if the question is of the type: How much of fact(s) X is accounted for by Y, then it does not make sense to choose the parameter values with the objective of making the amount accounted for as large as possible according to some metric. A model economy of course is an abstraction and by definition false. With enough data, hypothesis testing surely would reject it along some dimension. The model is useful insofar as it provides a quantitative answer to an interesting issue. A given model may be appropriate for some question (or class of questions) but not for others.

Calibration

A few words about calibration are in order. Originally, in the physical sciences, calibration referred to the graduation of measuring instruments. For example, a thermometer is calibrated so that it registers zero when immersed in water that contains ice and one hundred when immersed in boiling water. Theory is used: mercury expands approximately linearly within this range of temperatures. Computational experiments are, like thermometers, in a sense measuring devices. In physics they are artificial physical systems or models that are used to estimate quantitatively what will happen under different contingencies. Generally, there are some questions to which we know the answer and to which the model should give an approximately correct answer if we are to have any confidence in the answer the model economy provides to the original question posed. The model systems are calibrated so that this is indeed the case. This task may entail varying some of the parameters of the model until the model system mimics reality on the key dimensions. In the physical sciences, this activity has come to be called calibration.

Economic systems are, of course, different from physical models. Economic models are inhabited by people who anticipate and make decisions that are in their *ex ante* best interest, given that other model people behave likewise. In spite of the difference between physical and economic models, however, the same principle applies: we have more confidence in the answer to the question posed if the model gives correct answers to questions to which we already know the answer.

The task of calibration often involves merely computing a few averages. For example, if the standard Cobb–Douglas production function is used in the simple growth model above, that is, we let $f(n, k) = n^{\theta}k^{1-\theta}$, then a numerical value for the parameter θ can be obtained by computing the average over a period of years for the labor share of aggregate output. Several other growth relations map more or less directly into parameter values for typical models within the neoclassical growth framework, at least if they have been formulated with the ability to calibrate in mind. As a consequence, the computational experiments reproduce the key growth or long-term relations among the model aggregates.

Most growth relations have remained essentially constant, in the sense of not changing much on the average from one cycle to the next, for decades.

There are notable exceptions, however. The inventory stock as a fraction of GDP has declined slowly. Real purchases of durables as a fraction of total consumption has risen, although the nominal expenditure share has remained essentially flat. Depending on the associated pattern in the corresponding relative price, such information often enables the researcher to obtain a quite precise estimate of some elasticity of substitution. At the same time, it may be entirely acceptable to abstract from the difference in growth rates of the same two quantities if that feature does not play a role for the answer to the question posed.

A good example is the fact that per household hours of work have remained nearly constant for several decades in spite of a large rise in the real wage rate over the same period. Modelers have taken this fact to indicate an elasticity of substitution between consumption and nonmarket time near one. Still, many business cycle models abstract from the long-run productivity growth that would give rise to this sort of wage growth. The reason is that the answer to the questions addressed in those studies would have been the same. For example, Hansen (1989) compares otherwise identical model economies and permits growth in one version and not in the other. The model without growth needs a slight adjustment in the capital depreciation rate in order to be calibrated to the investment share of output and the observed capital/output ratio. With this adjustment, both models produce the same estimate of the role of Solow residuals for cyclical fluctuations.

Since these model economies are populated by people, another source of calibration is averages across large numbers of the relevant people in the actual economy. For example, some model environments employ a utility function in consumption and leisure which, like the production function above, has a share parameter. The approximate empirical counterpart is the average fraction of time spent in market activity. This fraction, in principle, can be obtained from panel data covering large samples of individuals. An example of a useful and careful measurement study is Ghez and Becker (1975). To carry out such a study, one needs to make choices about a variety of issues. What should be the upper and lower age limits for people to be included? What is a reasonable definition of the total time endowment? The choice by Ghez and Becker to exclude sleeping time and time for personal care is a reasonable one.

3. Towards heterogeneity

The simple framework just described has permitted me to discuss the specification of preferences, of productive opportunities, and of the laws of motion for the state variables. In addition to representing a natural economic language for aggregate issues, perhaps its simplicity was also a factor in making it attractive. It turned out to be easy to introduce a few bells and whistles, for example to turn it into a business cycle model, while maintaining tractability, at least as long as one was willing to stick to one key assumption – homogeneous agents. In many environments, indeed in a lot of those used in the past ten years, the equilibrium can then be computed by solving a stand-in planner's problem – a great saving in computational detail. Because of its tractability, it was natural that economists first would try to push this abstraction as far as they could. In many cases, the question was one for which it was *a priori* quite clear that heterogeneity was not at the heart of the issue. And even if one suspects heterogeneity may potentially affect the answer, a preliminary answer still can be valuable in guiding future research. In Schneider's words, it's the simple model that precedes research at higher resolution.

When Prescott and I wrote the paper published in *Econometrica* in 1982, we did not consider departures from the homogeneous-agent assumption. I suppose the finding was sufficiently striking – that more than half of U.S. postwar business-cycle volatility could be accounted for by technology shocks, also called Solow residuals – that it was unlikely to depend too much on this assumption. The modified version of a neoclassical growth model, augmented by features that we thought important in accounting for fluctuations observed in quarterly data, obviously still abstracted from many things. In spite of that, it replicated many of the key cyclical facts well enough, we thought, that the credibility of the estimate of the role of this source of impulse had to be taken seriously. Relative volatilities of consumption and investment were quite similar to those in the data, as was the behavior of the capital input. The one property that, initially at least, seemed like an anomaly was the volatility of the labor input, measured empirically by aggregate hours of work. In the data, it was much greater than in the model.

Now, in my opinion, the fact that this economy was calibrated in such a way that there was little room for playing with parameter values made this deviation stand out all the more clearly. It made me think about differences across workers in skill levels as a possible source of additional labor-input volatility. Thus, when I was invited to write a paper on the cyclical behavior of the labor market for the 10th-anniversary Carnegie–Rochester Conference in 1983, the role of labor-force heterogeneity to me was the obvious interesting topic. Differences in cyclical behavior depending on human capital seemed particularly important.

So the first task was to see if there was any empirical basis for this casual impression. For that purpose, I took from University of Michigan's Panel Study of Income Dynamics (PSID) data on prime-age males, the group of the labor force that presumably is cyclically the least volatile. These were divided into five different categories depending on years of schooling. It turned out that the cyclical volatility of annual hours of work was substantially higher for the lower-education workers than for those with more years of schooling.

A simple abstraction thought to capture the just-mentioned feature was one with the population divided into two equal-sized groups, one much more skilled for market production than the other. This difference in human capital was taken to be exogenous. I then considered equilibriums obtained by solving a Pareto problem with weights on each type's utility functions. The weights were calibrated so that average hours per period would match observations, that is, the high-skilled workers in the steady state worked more hours to an extent corresponding to the data.

The finding was that, especially if the technology is such that the equilibrium wage of the skilled relative to that of the unskilled varies slightly and is countercyclical, an assumption consistent with empirical findings by Reder (1962), then the volatility of aggregate hours for a given volatility of Solow residuals increases substantially. This is obviously a simple way of introducing heterogeneity while still having the equilibrium solve a stand-in planner's problem, but it encourages one to think about what the production function should look like with different types of workers, about sources of differences in trade-offs between market and nonmarket production, and so on. This is an area of business-cycle research that I think should be emphasized more. At the firm level, there is a wealth of data on different classifications of workers (blue collar vs white collar and so on), data that have been used by Berndt and Christensen (1974) and others to estimate parameters characterizing interrelations between different types of labor inputs.

An alternative to this modeling strategy of incorporating multiple skills is to improve the empirical hours series to make it more appropriate as a measure of labor input. In particular, one needs to weight the hours of different workers by their human capital. Recently, Prescott and I (1993) did so for a sample from the PSID consisting of all demographic groups, with the finding that for the particular sample period in question, late 60s to early 80s, the constructed labor-input series fluctuated substantially less than the corresponding aggregate hours series. The focus of this talk, however, is more on what can be done from a modeling standpoint, so I'll concentrate on that.

A simplifying assumption in the two-skill model was the exogenous endowment of human capital. Recently, Víctor Ríos-Rull (1993a) has studied an overlapping generations model in which human capital is endogenous. Part of the motivation is the fact, mentioned already, that over the past few decades, real compensation has risen dramatically while hours per household devoted to market work has remained about constant. Cross-sectionally, however, there is a clear positive correlation between market work and the real wage. Moreover, the volatility of hours of work per year is much lower for high-wage than for low-wage earners. The key features of his economy are the existence of a home-produced good with poor market substitutes, and the choice in the first period of whether to acquire skills through schooling. Consequently, there is meaningful heterogeneity in each generation even though everyone is born alike. The model accounts well for some of the striking movements both cross-sectionally and secularly. Cyclically, however, a remaining discrepancy is that, unlike the U.S. data, the model's volatility of market hours for unskilled workers does not exceed that for the skilled.

In the economies discussed so far, a literal interpretation is that everyone works, and that all the labor-input volatility is in hours per worker. This assumption has been criticized by many. [See Heckman (1984), for example.] In the U.S., in fact, there is substantially more volatility in employment than in hours per worker. Presumably, much of the employment volatility is related to fixed costs, more or less, of moving in and out of the market sector – a source of nonconvexity. Building on a theoretical insight by Rogerson (1988), Gary Hansen (1985), in an influential paper, incorporates employment fluctuation in a business-cycle model with indivisible labor. Indeed, he goes to the opposite extreme and assumes that *all* the labor-input volatility is of this form. With two possible levels of hours per worker – zero or some positive number – the key is to regard the decision variable as a probability, namely that of working in that period. That is, the relation between workers and firms can be such that everyone understands that there is a probability of being out of work, but everyone still gets compensated. That probability then varies cyclically. The introduction of this lottery convexifies the problem and makes the model tractable computationally.

Hansen finds that when labor input takes only this form, the theoretical labor-input volatility increases substantially. Also, an important message in the context of his economy is that the aggregate elasticity of intertemporal hours substitutability could be very large even though individuals have standard utility functions that, at the individual level, imply elasticities that are low, as labor economists tend to find for panel data using their empirical methods.

Some have criticized the Hansen economy. There is too much insurance. Consumption is the same whether you work or not – seemingly a steal for the nonworking, as they also get to enjoy the full time endowment as leisure. This property is far from general, however, and merely reflects Hansen's use of a logarithmic utility function. By adding curvature in the form of a larger risk-aversion parameter, consumption of those outside the market sector will be lower than for those who are in. Whether the degree of insurance would be hard to mimic, within the household, through government institutions, by holdings of precautionary liquid assets, or in other ways, will have to be investigated. And it may not make all that much difference to the quantitative answer. The main point is very clear and transparent due to the simplicity of Hansen's abstraction.

Observations indicate that both margins are important. For example, if we remove trends using the Hodrick–Prescott filter, the cyclical standard deviation of quarterly total hours over the period 1954–91 according to the household survey is 1.49 percent. For hours per worker, it is only 0.54, while it is 1.09 for employment. Both components of total hours are highly procyclical, with employment lagging output by about one quarter.

In a paper published in 1991, Prescott and I introduced both margins in a business-cycle model. An interesting issue is what production function to use. The standard one presumably would treat total hours – the product of employment n and hours per worker h – as the labor input, that is, output equals $z(hn)^{\theta}k^{1-\theta}$. This would mean that the marginal product of labor in a given period is affected in the same way for a given hours increase regardless of what form it takes. But machines are typically designed for a particular number of workers. Thus, with the period's capital fixed, adding hours by hiring more workers on the average should increase output less than adding the same amount of hours by having the existing work force work longer hours in that period. Thus, we regarded it as a better approximation to let output be linear in hours per worker: $zhn^{\theta}k^{1-\theta}$.

Recently, some economists have attempted to interpret h as effort rather than hours per worker. The idea is that it reflects labor hoarding. Given the choice of interpreting h as hours per worker, for which we have good measurements, or effort, for which I'm not aware of any measurements, I know where I would put my money. There are of course implications for the measurement of Solow residuals for such economies. In the hours vs employment economy, a way to deal with them is to input into the model Solow residuals with properties according to statistics based on standard measurements, compute model time series, and then compute from the model the Solow residuals according to the standard method. The magnitude of the bias can then be corrected for. [For more detail, see Kydland and Prescott (1991).]

There are other margins as well. One could consider more than one shift, for example. The business-cycle model that goes furthest in this direction is that by Cooley, Hansen, and Prescott (1993). There is good basis in measurements for adding such detail to the technology.

4. Overlapping generations economies

The infinitely-lived agent model is one of the two basic frameworks that have been used. The other is the overlapping generations model which introduces heterogeneity in a more generic way, one might say. It allows for life-cycle behavior, which may be essential for issues such as the role of demographic changes for savings behavior. These economies tend to be computationally quite intensive, but that is becoming less and less of a constraint.

Some of the earliest computational experiments with mortal-consumer economies are in a macroeconomics textbook – by Miller and Upton (1974). People live for four periods. They earn labor income only in periods 2 and 3, but prefer a smooth consumption stream. Miller and Upton use this economy to illustrate numerically several things, such as the time path toward steady state of a growing economy, the role of government debt with and without bequest motive, the real effects of unexpected changes in monetary policy when debt contracts between generations are nominal, and so on. It's amazing that, after 20 years, macro textbooks still have not caught up with Miller and Upton in this respect.

A real breakthrough in the application of OLG models was Auerbach and Kotlikoff (1987), in which they compute time paths of economies with people who live for 55 periods. In other words, the period length is a year, which is important for the ability to calibrate the economy. They studied the quantitative role of various aspects of fiscal policy with life-cycle behavior and showed that such economies were computable, at least in the absence of uncertainty and with the focus restricted to movement from one steady state to another.

Since then, the profession has made great strides in its ability to deal with such models. The one who has pushed them furthest is Víctor Ríos-Rull. What he can do computationally is impressive, especially since his model economies have a strong theoretical and factual basis. Moreover, they allow for uncertainty, for example about individuals' lifetimes and about aggregate productivity, whose effect on a generation's earnings depends on the stage of the life cycle.

Let me start with a robustness kind of question. The prototype real business-cycle economy is inhabited by immortal consumers, that is, no life-cycle behavior is assumed. A question is: If we addressed the same sort of question, say, about the role of technology shocks in an economy with mortal consumers, would the findings be different? Victor's answer is no. Sure, there are some minor differences in the properties of some of the model statistics, but none that would lead one to question the basic findings from the immortal-consumer framework.

Another interesting question is: What quantitative role would be played by the lack of insurance markets to hedge against certain shocks? The comparison is with environments in which people can trade contingent securities with other generations. For example, with a hump-shaped earnings pattern over the life cycle, technology shocks affect individuals differently depending on their age. Also, people face uncertain lifetimes.

Let's try to get some of the flavor of how such an economy may be modeled. I shall focus on the budget and resource constraints and emphasize two important features that can play a role in such an economy, in addition to the simple fact that there are different age cohorts. [For a fuller description, see Ríos-Rull (1993b).]

Assume that people live for up to I periods. The population grows at the rate of λ, and the size of each cohort i is μ_i. If we normalize so that μ_1 is one, then $\mu_i = (1 + \lambda)^{-i}$. The main source of uncertainty is aggregate productivity. That is, aggregate output is given by $zf(K, N)$, where $z \in Z = \{z_1, z_2, \ldots, z_{n_z}\}$, and where f is a standard production function displaying constant returns to scale. To insure against these shocks, people trade in securities which deliver one unit of capital goods next period. The delivery of these capital goods is dependent on elements of a partition of Z: $\mathscr{F} = \{\phi_1, \phi_2, \ldots, \phi_{n_\phi}\}$, where $n_\phi \leq n_z$. Let b_i^ϕ be the amount of the state-contingent capital good purchased by generation i and q^ϕ its price. If a_i denotes the net wealth accumulated by age-i agents, then their budget constraints are as follows (dropping time subscripts):

$$a_1 = 0\,,$$

$$Ra_i + W(1 - l_i)\varepsilon_i = \sum_{\phi \in \mathscr{F}} q^{\phi} b_i^{\phi} + c_i\,,$$

$$a'_{i+1}(\phi) = b_i^{\phi}\,,$$

$$b_I^{\phi} \geq 0, \quad \phi \in \mathscr{F}, \quad i = 1, \ldots, I\,,$$

where R is the return on asset holdings and W the price per unit of labor input, both determined in equilibrium from the respective marginal products. Each individual has time, normalized to one, that can be used for nonmarket activity, l, or market activity, $1 - l$. Moreover, a unit of market time translates into different quantities of labor input depending on age, as determined by the exogenous sequence of factors $\{\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_I\}$. Using observations for thousands of individuals, one can measure average age-specific ε's. Ríos-Rull finds that this life-cycle pattern is hump-shaped, with the peak about three-quarters higher than ε_1, and with a noticable decline approaching retirement.

The aggregate inputs into production, then, are

$$K = \sum_{i=1}^{I} \mu_i a_i \quad \text{and} \quad N = \sum_{i=1}^{I} \mu_i (1 - l_i)\varepsilon_i\,,$$

and the aggregate feasibility constraints are

$$\sum_{i=1}^{I} \mu_i (b_i^{\phi} + c_i) \leq z f(K, N) + (1 - \delta)K \quad \text{for all} \quad \phi \in \mathscr{F} \quad \text{and} \quad z \in Z.$$

An important aspect of this economy is the sequence of ε_i's over the life cycle. Thus, when there is an aggregate shock to the technology level z, the effect will differ depending on the stage of the life cycle. Consequently, life-cycle behavior is different when contingent securities are present than when they are not. One can then ask what is the quantitative magnitude of this difference, say, as reflected in standard cyclical statistics, when quantities are aggregated across all people. With a standard utility function added to the model, such questions can be addressed.

There are some slight differences in model behavior in the very early stage of the life cycle, and somewhat larger differences towards the end of the life cycle. More surprising is the finding that the presence of contingent securities plays essentially no quantitative role for the aggregate behavior across generations.

An extension of this model economy is to permit laws of motion for the population movement, as reflected in the μ_i's, rather than simply assuming constant population growth. This is done in Ríos-Rull (1993c) for the purpose of predicting savings behavior over the next couple of decades under alternative birth-rate scenarios. Both in the U.S. and in many other countries, there is still a baby-boom generation making itself felt in different ways as it ages. A lot of information exists on the basis of which one can calibrate the law of motion for the population. The biggest problem is what to use as the initial condition. One possibility is to assume that the latest year available characterizes the steady-state distribution. This is not a likely scenario. An alternative, used by Ríos-Rull, is to simulate until a population distribution, with associated asset distribution, resembling the current one is obtained, and then use it as the initial condition for the subsequent computational experiment. This is hard to do, especially for the U.S., in part because immigration may be an important factor, more so than for many other countries. For that reason, Ríos-Rull considers Spain instead. An interesting finding is that future fertility has implications for future equilibrium real interest rates that are quantitatively sizable. To work out the joint implications of birth-rate *and* immigration scenarios remains an interesting question for future research.

Another way in which aggregate shocks can affect people, due in part to the age–earnings profile, is through a change in the progressiveness of the income-tax schedule. In Altig and Carlstrom (1991), such changes take place through variation in the inflation rate. The question they address is: How much does inflation volatility, with a nonindexed tax system, contribute to the volatilities of real aggregates? In addition to being progressive, a property of their tax environment is that capital income for tax purposes is overstated in inflationary periods. In their model, as in the work of Ríos-Rull, the life-cycle pattern of productivity is important. A finding is that inflation volatility has the biggest effect on labor-input decisions.

Two tentative conclusions can be drawn from the discussion so far:

(i) Heterogeneity associated with age profiles is probably not a big deal for most business-cycle questions.

(ii) Skill differences that arise for different reasons than simply the age profile play an important role.

5. Computational issues

So far, I have described two types of economies that incorporate heterogeneity of labor input. One is such that equilibria still can be computed by solving for the Pareto optimum. An example is in Kydland (1984). Compared to a corresponding homogeneous-worker economy, this involved only adding two decision variables and one state variable. Using as an objective function the quadratic approximation around the steady state, the decision rules describing aggregate behavior can be computed in less than a minute even on a 386.

The computational effort required to solve models like those of Ríos-Rull, however, may seem prohibitively high even though he uses quadratic approximations. Because of the mortal-consumer generational structure, his models have numbers of variables running into the hundreds. Even on a fast work station, this may take a day or two of computing. Perhaps this is not so bad, since it provides an incentive to do the calibration thoroughly so that only a few computational experiments are needed to give a reliable answer to the question addressed.

The class of models for which the computation of equilibria is the toughest is that in which agents face idiosyncratic shocks. [See the survey in Ríos-Rull (1993d).] The environment is that agents cannot write contracts on their individual shocks and, in some cases, cannot hold negative quantities of any asset. This environment prevents standard aggregation results from holding. The basic problem is that computing equilibria requires keeping track of the distribution of agents. The reason is that, in making decisions, agents have to forecast future prices, and these prices (or their laws of motion) depend on the changing distribution of agents. Iterating over the usual stuff is often hard enough. Iterating over distributions is extremely difficult and time-consuming. Moreover, because constraints often are binding, the usual quadratic-approximation methods cannot be used. Instead, the space over which the variables may range typically is discretized, which further increases the computational burden. But then again, presumably it's only a matter of time before computers get large enough to deal even with such issues.

In the meantime, the models in this class that have been studied make assumptions preventing the distribution across agents from affecting future prices. Imrohoroglu (1989) introduces a storage technology which pins down the relevant price. In several examples, for example Díaz-Giménez and Prescott (1989) and Díaz-Giménez (1990), the key is the assumed behavior of the government. Indeed, as Ríos-Rull (1993d) says, in these heterogeneous-agent models with liquidity constraints, 'this is what the government does, implements policies that prevent the distribution of agents from affecting agents' decisions.' Finally, we've found a useful role for the government!

These economies already have been the basis for many interesting insights. They are particularly useful for studying the welfare implications of certain kinds of government policies. Imrohoroglu, for instance, found welfare costs associated with inflation that far exceeded estimates in models without the precautionary motive for holding assets. Such economies are well suited to evaluate the welfare implications of various institutional arrangements, such as unemployment insurance [Hansen and Imrohoroglu (1992)].

The development of the methods and models that allow us to deal, within general equilibrium economies, with the forms of heterogeneity that we think are important is obviously in its infancy. With the talent present today, I'm sure that many of you will be among those who present the further developments at these conferences in the coming years.

References

Altig, D. and C.T. Carlstrom, 1991, Inflation, personal taxes, and real output: A dynamic analysis, Journal of Money, Credit and Banking 23, 547–571.

Auerbach, A.J. and L.J. Kotlikoff, 1987, Dynamic fiscal policy (Cambridge University Press, Cambridge).

Berndt, E.R. and L.R. Christensen, 1974, Testing for the existence of a consistent aggregate index of labor inputs, American Economic Review 64, 391–404.

Cooley, T.F., G.D. Hansen, and E.C. Prescott, 1993, Idle resources and the cyclical behavior of capital utilization, Working paper (University of Pennsylvania, Philadelphia, PA).

Díaz-Giménez, J., 1990, Business cycle fluctuations and the cost of insurance in computable general equilibrium heterogeneous agent economies, Working paper (Universidad Carlos III, Madrid) reproduced.

Díaz-Giménez, J. and E.C. Prescott, 1989, Asset returns in computable general equilibrium heterogeneous agent economies, Working paper (Federal Reserve Bank, Minneapolis, MN).

Ghez, G. and G.S. Becker, 1975, The allocation of time and goods over the life cycle (Columbia University Press, New York, NY).

Hansen, G.D., 1985, Indivisible labor and the business cycle, Journal of Monetary Economics 16, 309–327.

Hansen, G.D., 1989, Technical progress and aggregate fluctuations, Working paper 546 (University of California, Los Angeles, CA).

Hansen, G.D. and A. Imrohoroglu, 1992, The role of unemployment insurance in an economy with liquidity constraints and moral hazard, Journal of Political Economy 100, 118–142.

Heckman, J., 1984, Comments on the Ashenfelter and Kydland papers, Carnegie–Rochester Conference Series on Public Policy 21, 209–224.

Imrohoroglu, A., 1989, The costs of business cycles with indivisibilities and liquidity constraints, Journal of Political Economy 97, 1364–1383.

Kydland, F.E., 1984, Labor force heterogeneity and the business cycle, Carnegie–Rochester Conference Series on Public Policy 21, 173–208.

Kydland, F.E. and E.C. Prescott, 1982, Time to build and aggregate fluctuations, Econometrica 50, 1345–1370.

Kydland, F.E. and E.C. Prescott, 1991, Hours and employment variation in business cycle theory, Economic Theory 1, 63–81.

Kydland, F.E. and E.C. Prescott, 1993, Cyclical movements of the labor input and its implicit real wage, Federal Reserve Bank of Cleveland Economic Review 29: 2, 12–23.

Miller, M.H. and C.W. Upton, 1974, Macroeconomics: A neoclassical introduction (Richard D. Irwin, Homewood, IL).

Reder, M.W., 1962, Wage differentials: Theory and measurement, in: Aspects of labor economics (National Bureau of Economic Research, New York, NY).

Ríos-Rull, J.V., 1993a, Working in the market, working at home and the acquisition of skills: A general equilibrium approach, American Economic Review 83, 893–907.

Ríos-Rull, J.V., 1993b, On the quantitative importance of market completeness, Working paper (University of Pennsylvania, Philadelphia, PA).

Ríos-Rull, J.V., 1993c, Population changes and capital accumulation: The aging of the baby boom, Working paper (Carnegie Mellon University, Pittsburgh, PA).

Ríos-Rull, J.V., 1993d, Models with heterogeneous agents, in: T.F. Cooley, ed., Frontiers of business cycle research, forthcoming.

Rogerson, R., 1988, Indivisible labor, lotteries and equilibrium, Journal of Monetary Economics 21, 3–16.

Schneider, S.H., 1987, Climate modeling, Scientific American 256: 5, 72–80.

Stokey, N. and R.E. Lucas with E.C. Prescott, 1989, Recursive methods in economic dynamics (Harvard University Press, Cambridge, MA).

Whalley, J., 1985, Trade liberalization among major trading areas (MIT Press, Cambridge, MA).