Models of Business Cycles, R Lucas

# I

Research in my field of specialization – macroeconomics, or monetary and business cycle theory – has undergone rapid change in the past 15 years. One way of describing some of these changes is in terms of ideological contests between rival schools of thought: the 'Keynesian revolution', the 'monetarist counter-revolution', and so on. There is no doubt something to be learned by tracing the main ideological currents in macroeconomic research, but I myself find most of this discussion of crises, revolutions and so on, unintelligible, and almost wholly unconnected with the most interesting current research. Recent macroeconomic controversy seems to me much more easily interpreted as a surface manifestation of much deeper and more important developments in economic theory, developments that long predate current macroeconomic controversies and that will be of importance long after these controversies have gone the way of the liquidity trap, operation twist, the loanable funds doctrine, and all the other macroeconomic issues that seemed so important as they were occurring and are so hard to remember now.

Dynamic economic theory – I mean theory in the sense of models that one can write down and *do* something with, not in the sense of 'opinion' or 'belief' – has simply been reinvented in the last 40 years. It is now entirely routine to analyze economic decision-makers as operating through time in a complex, probabilistic environment, trading in a rich array of contingent-claim securities, and to study agents situated in economies with a wide variety of possible technologies, information structures, and stochastic disturbances. While Keynes and the other founders of what we now call macroeconomics were obliged to rely on Marshallian ingenuity to tease some useful dynamics out of purely static theory, the modern theorist is much better equipped to state *exactly* the problem he wants to study and then to study it.

This new ability to incorporate dynamic and probabilistic elements into economic theory, at the same level of rigor on which we study the problem of a single decision-maker making a one-time choice at given prices, has already had a deep, permanent influence on virtually every branch of applied economics. What people refer to as the 'rational expectations revolution' in macroeconomics is mainly the manifestation, in one field of application, of a development that is affecting *all* fields of application. To try to understand and





explain these events as though they were primarily a reaction to Keynes and Keynesianism is futile. One may as well try to understand parallel developments in industrial organization as a reaction to Viner or Chamberlain or in public finance as a reaction to Pigou or Musgrave.

In macroeconomics, certainly, this process of dynamicization is very far from completed, and I do not believe it has yet taken us to a satisfactory theoretical account of the events we group together under the term 'business cycle'. Yet even at this early stage, I think it is evident that certain habits of thought that are central to the way we have thought about macroeconomics and macroeconomic policy since the Second World War need to be replaced with a point of view that is in some ways radically different.

I would like to use these lectures as an opportunity to describe some of these developments, and the changes in viewpoint they seem to me to necessitate. I will begin, in section II, by discussing some general considerations that seem to me to be important in deciding which kinds of economic models might permit us to determine the effects that economic policies will have on people's consumption streams and welfare. Section III confronts a hypothetical consumer, in an experimental spirit, with some different consumption paths, and traces out his reactions to them, in an attempt to get a quantitative idea of the importance of stabilization policy relative to other economic questions.

Section IV moves to an exposition and discussion of a particular model of aggregate activity due to Finn Kydland and Edward Prescott. This model seems to me to offer a useful definition of the current frontier in business cycle research and I will use it as a point of departure in discussing outstanding questions.

Sections V and VI deal with two issues from which Kydland and Prescott abstract, but which many believe to be at the center of the problem of business cycles. Section V surveys certain aspects of the theory of unemployment. Section VI introduces monetary complications.

Section VII considers the prospects for a model of business cycles centered on the role of monetary disturbances. Though this is the subject, of all those considered in the lectures, to which I have devoted the most thought – and with good judgement, I continue to believe – it is also that on which the least is known and I will rely heavily on conjecture. I will sum up these matters in the last section.

When I began work on these lectures, I had in mind something like a broad and balanced survey of recent developments. Now that I am finished, it is easy enough to see the ways in which I have been guided by research prejudices that cannot (at least as yet!) be said to have been verified by hard results and also the ways in which I have been limited by the narrowness of my own technical range. In particular, I will not treat in any detail the





econometric developments that I think may well be the most lasting contribution of John Muth's idea of rational expectations. A good introduction to these ideas is contained in a volume Thomas Sargent and I edited,[1] and to which Sargent is the leading contributor, but I will not be reviewing these developments here.

1 Robert E. Lucas, Jr and Thomas J. Sargent (eds), *Rational Expectations and Econometric Practice* (University of Minnesota Press, Minneapolis, 1981).

# II

Discussions of economic policy, if they are to be productive in any practical sense, necessarily involve quantitative assessments of the way proposed policies are likely to affect resource allocation and individual welfare. This means that useful policy discussions are ultimately based on *models*, not in the sense that policy decisions can be automated once and for all without the need for individual judgements, but in the sense that participants in the discussion must have, explicitly or implicitly, some way of making a quantitative connection between policies and their consequences. I want to spend almost all of my time in these lectures, then, on issues involved in the construction of useful economic models. Although this will lead us into technical issues that seem far removed from current policy debates, I think this is ultimately the most – perhaps the *only* – practical way to proceed. Macroeconomics receives a great deal of attention in the newspapers, but this is not the level at which progress is made or continuity is to be discovered.





In a general way, the problem of macroeconomics – really, of *all* applied economics – is to go from non-experimental observations on the past behavior of the economy to inferences about the future behavior of the economy under alternative assumptions about the way policy is conducted. In terms of models, then, we want a model that fits historical data and that can be simulated to give reliable estimates of the effects of various policies on future behavior. But *what* data? And what do we mean by *fit?* And when can we expect that particular simulations will be *reliable?* These are hard questions, harder and more open than is commonly acknowledged, and nothing is more dangerous than the illusion that they can be answered by 'common sense'. It will be helpful, I think, to consider these questions in a very general way before turning to the construction of specific, operational models.

A useful model, to begin at the very end of the story, is going to take the form of an explicit description of the way the economy evolves through time. We will want to consider stochastically disturbed systems, so let $e_t$ denote independent drawings of an exogenous shock from some fixed distribution $G(e)$, and let the law of motion of $s_t$, a complete description of the 'state of the system' at date $t$, be denoted:

$$s_{t+1} = F(s_t, e_t).$$

Think of $s_t$ as a vector with many components, including the economic variables of primary interest, such as the date $t$ consumption of the agents in the system, underlying intermediate variables like stocks of capital goods and inventories, policy variables and perhaps variables of no intrinsic importance except that they contain information useful in forecasting future values of these other variables. Similarly, think of the shock vector $e_t$ as a vector with many components.

We could think of proceeding empirically by identifying the state vector $s_t$ with observable magnitudes and estimating $F$ and $G$ from economic data. This, very roughly, would give us what an econometrician would call a 'reduced form' description of the system. Such a description might be useful in forecasting future $s_t$ values (this would depend on the stability of the estimated reduced form over time), and perhaps, as Christopher Sims and George Stigler have argued in different ways, this is all we can hope to do anyway.[1] But if, as I

1 I associate with Stigler (for example, George J. Stigler, *The Citizen and the State* (University of Chicago Press, Chicago, 1975)) and Sims (for example, Christopher A. Sims, 'Policy analysis with econometric models', *Brookings Papers on Economic Activity* (1982), pp. 107–64) the view of social science as most usefully standing outside the political as well as the economic system, attempting to forecast and understand the behavior of both without attempting to influence either. The normative, welfare-economic tradition is, in contrast, unabashedly utopian, treating the behavior of the private sector as something to be predicted, and that of the public sector as something to be reformed. Stigler, with some justice, calls this a 'deeply schizophrenic view of the state'.





have been taking for granted and will continue to do, our objective is to determine how *changes* in those aspects of the system we call 'policies' induce *changes* in individual consumption and welfare, we need much more knowledge about $F$ and $G$ than simply the forms that fit certain historical data.

The problem, of course, is that the motion $F$ of an economic system depends, in a very complicated way, not only on actions we would want to call 'policies' but on the actions of all of the agents that make up the system. At an abstract level, we can say that a policy change involves a shift from one law $F$ to another, $F'$ say, but operationally, we have no direct way of knowing how particular policy changes translate into $F$-changes.

We can move toward what an econometrician would call a 'structural' description of the system, by refining the specification of $F$ so as to differentiate between actions taken by private agents and those taken by 'nature' (including in the latter term 'policies'). To set up a notation that reflects this

Sims criticizes the 'rational expectations revolution' for 'destroying or discarding much that was valuable in the name of utopian ideology'. See Thomas J. Sargent, 'Autoregressions, expectations and advice', *American Economic Review* **74** (1984), pp. 408–15, for a useful discussion of Sim's position.

There is no point in arguing over which of these positions is correct; the answer must be 'neither' or 'both'. Throughout these lectures, I will be consistently taking the schizophrenic, utopian point of view.

distinction, let 'nature' take action $z_t = z(s_t)$ when the state of the economy is $s_t$, and let agent $i$ take action $a_{it} = a_i(s_t)$, with $a_t = a(s_t)$ denoting the vector of *all* actions taken by all agents. Then if $s_{t+1} = H(z_t, a_t, s_t, e_t)$ when the shock is $e_t$ and nature and private agents take the actions $(z_t, a_t)$, we can express $F$ as

$$F(s_t, e_t) = H(z(s_t), a(s_t), s_t, e_t).$$

Is this translation progress? This depends on whether the changes in policy we want to analyze are easily translated into changes in the way we specify the function $z(\cdot)$ (describing 'nature's' actions) *and* if we have reason to believe that the functions $H(\cdot)$ and $a(\cdot)$ do not vary in responses to these changes in policy $z(\cdot)$. This is exactly the case for 'structural' estimation (estimation of $H$, $a$ and $z$ separately) as opposed to 'reduced form' estimation that Marschak made in his classic Cowles paper: if policy shifts the supply function in a known way, and if it does not shift the demand function, then knowing demand and supply parameters separately permits us to predict the price and quantity effects of the policy, while knowing only the reduced form does not.[2]

2 Jacob Marschak, 'Economic measurements for policy and prediction', in William C. Hood and Tjalling G. Koopmans (eds), *Studies in Econometric Method*, Cowles Commission Monograph 14 (Wiley, New York, 1953), pp. 1–26.





There are interesting problems in economics for which this is as far as one needs to go, but the issues in economic dynamics that arise in macroeconomics are not among them. The problem, as is now widely recognized, is that there is no reason to expect that the function $a(\cdot)$ describing agents' actions will remain invariant under changes in the function $z(\cdot)$ describing nature's actions. If agents are interested in maximizing an objective function, say the expected utility from the stochastic consumption streams they receive, they will choose actions $a_{it}$ in order to try to achieve this objective. I have not as yet spelled out the decision problem faced by these agents in a way that would permit us to see exactly what function or 'decision rule' will achieve this, but it is clear enough simply from the facts that the agents are interested in their future consumption and are operating in a changing environment that one would expect their decision rules $a_i(\cdot)$ to adapt in response to changes in policy $z(\cdot)$. To imagine otherwise is to assume that the solution to a maximum problem does not vary with changes in the function being maximized! Given this, the identification of the 'structure' in the sense of the functions $H$, $a$ and $z$ gives us no more ability to determine the consequences of policy changes than does knowledge of the 'reduced form' $F$. We need a deeper idea of what we mean by 'structure', not because 'depth' is desirable in itself – a key to success in applied science, I think, it to operate on as shallow a level as one can get by with – but because a model has to be able to isolate those aspects of behavior that remain invariant to policy shifts from those that do not if it is to be of any use in assessing the consequences of the shift.[3]

It will clarify what I mean by a 'useful structure' to set out some additional formalism, though I must say in advance that the formalism I will write down is both too general to do anything with, analytically, and also too 'special' in the sense that it is easy to think of interesting dynamic models that it does not subsume as special cases. Let us imagine that at a given date, when the system is in state $s$, nature selects an action $z$ and each agent $i$ selects an action $a_i$ from an opportunity set $\Omega_i(a_{-i}, s, z)$ that is determined by the state $s$, nature's action $z$, and the actions $a_{-i}$ taken by all other agents. Denote the immediate utility or 'pay-off' to agent $i$, given *all* of these actions, by $R_i(a, s, z)$. Assume that agent $i$ seeks to maximize the expected, subjectively discounted sum of these pay-offs, or

$$E\left\{\sum_{t=0}^{\infty} \beta^t R_i(a_t, s_t, z_t)\right\}. \tag{1}$$

3 This argument is elaborated and illustrated with examples in Robert E. Lucas, Jr, 'Econometric policy evaluation: a critique', in *The Phillips Curve and Labor Markets*, Vol. 1 of the Carnegie–Rochester Series on Public Policy, Karl Brunner and Allan H. Meltzer (eds) (North-Holland, Amsterdam, 1976), pp. 19–46.





For (1) to be meaningful, the content of the operator $E\{\cdot\}$ must be specified. I will assume that expectations are *rational* in John Muth's sense, or that agents *know* the functions $a(\cdot)$ and $z(\cdot)$, and the distribution defined by $F$ and $G$ of the state vectors $s_t$, and that they take the expectation $E\{\cdot\}$ correctly.[4] For this discussion, I will assume that all agents have the common information $s_t$ at each date $t$, so that $E\{\ \}$ means an expectation conditioned on the initial information $s_o$.

A central construct in discussing agent $i$'s decision problem is his *value function* $v_i(s)$, interpreted as the value of his objective function (1) when the system begins in state $s$, and he chooses the optimal action $a_i$ from the set $\Omega_i$. This value function must satisfy a 'Bellman equation' (after the late mathematician Richard Bellman), which we can derive by thinking through the consequences of a particular $a_i$ choice. The *immediate* payoff of such a choice, the first term in the sum (1), is just $R_i(a, s, z(s))$. The longer-term consequences, the remaining terms in (1), are

4 John F. Muth, 'Rational expectations and the theory of price movements', *Econometrica* **29** (1961), pp. 315–35. The term 'rational expectations', as Muth used it, refers to a consistency axiom for economic models, so it can be given precise meaning only in the context of specific models. I think this is why attempts to define rational expectations in a model-free way tend to come out either vacuous ('People do the best they can with the information they have') or silly ('People know the true structure of the world they live in').

summarized by $\beta E\{v_i(s')\}$: the value of the agent's objective as of the beginning of *next* period, when the state is $s'$ (an expectation, since $s'$ is not perfectly predictable), discounted to the present by the discount factor $\beta$. Now $s'$ is determined in part by the agent's own action $a_i$, in part by other agents' actions, $a_{-i}$, and in part by nature: $s' = H(z(s), a, s, e)$. Optimal behavior means maximizing the sum of immediate and long-term pay-offs, and it yields the value $v_i(s)$. The Bellman equation is thus:

$$v_i(s) = \max_{a_i \in \Omega_i(a_{-i}, s, z)} \{R_i(a, s, z) + \beta \int v_i(H(z(s), a, s, e))dG(e)\}. \qquad (2)$$

The system (2) for $i = 1, \ldots, n$ is in *equilibrium* when each agent $i$ chooses the action $a_i$ that attains the right side of (2), given the actions $(a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n) = a_{-i}$ chosen by all other agents. Clearly the nature of such an equilibrium $a(s) = (a_1(s), \ldots, a_n(s))$ will depend on everything that determines the nature of this game: the return functions $R_i$, the opportunity sets $\Omega_i$, the function generating policies, $z(s)$, and the law of motion $H, G$ of the system as a whole. Any *change* in policy, in the sense of a change in the function $z(\cdot)$ that describes the way policy variables react to the state of the system, will induce changes in the solution functions $a(\cdot)$, and hence in the reduced form behavior $F(\cdot)$ of the system as a whole.





For a model of the form (2) to be useful in evaluating policies, it is essential in the first place that we be able to calculate the solution functions $v_1, \ldots, v_n$. There is certainly no generally useful algorithm for doing this, and as we shall see, successful applications to date involve imposing stringent restrictions on the formulation to give tractability. This is a mathematical frontier on which there is much to be done. Even if (2) can be solved, the system will be successful in simulating policy effects only if the preferences $R_i$ and the technology and rules of the game $\Omega_i$ are invariant under changes in policy $z$. This property is not one that can be guaranteed logically: if it could, we would not need to *test* economic models but could simply build them up from impeccable axioms. Of course, in practice all axioms for models we can actually solve will be crude approximations at best, and determining which axioms produce reliable models will involve judgement, testing and luck.

I have deliberately set out this formalism without saying anything at all specific about the nature of the 'game' agents are assumed to 'play' at each date $t$. I have described the actions $a_i$ simultaneously chosen by agents as a (Nash) *equilibrium*, but the term equilibrium in this (now entirely standard) context obviously does not refer to a system 'at rest', nor does it necessarily mean 'competitive' equilibrium in the sense of price taking agents, nor does it have in general any connection with social optimality properties of any kind. All it *does* mean is that, in the model, the objectives of each agent and the situation he faces are made explicit, that each agent is doing the best he can in light of the actions taken by others, and that these actions taken together are technologically feasible.

It is not going to be possible to say anything very specific about business cycles or about stabilization policy at this level of generality, so I will soon consider much more specific (and controversial) models. But the main criticisms of Keynesian models and their use in formulating policies that one associates with the idea of 'rational expectations' are all straightforward consequences of the acceptance of the general formalism of dynamic games that I am using here. Let me review what these are.

Keynesian models – at least as this term was used in the 1960s and 1970s – consist of a set of equations most of which are the decision rules of agents, the functions $a_i(s)$ in my notation. That is, their structural equations describe the levels of consumption, investment, money demand, and so on, that agents choose as functions of variables describing their situation. Some equations also describe the decision rules of fictional agents, like the auctioneer who is assumed to move prices or wages in the 'Phillips curve'. This is the nature of the large-scale econometric models based on Keynesian ideas. It also describes textbook models of the 'IS-LM' variety, and their many theoretical





refinements. Now over a period of time in which policies are generated (or can be viewed as generated) by a fixed function $z(s)$ one would expect the 'structure' $a(s)$ of such models to remain fixed as well, and hence also $F(s, e) = H(z(s), a(s), s, e)$. Then forecasts based on this reduced form $F$ should be accurate (up to unavoidable errors due to unpredictable $e_t$ shocks). On the other hand, an attempt to simulate the effects of policy *changes* – changes in the *function* $z(s)$ – carried out by holding $a(s)$ fixed, will not yield accurate answers as long as equation (2) holds. I think the general logic of this criticism is widely understood. Its practical significance was dramatically illustrated by the failure of models containing 'stable' Phillips curves to deal with the effects of 1970s inflation, but the point is obviously not specific to this particular relationship.

*If* one had an accurate model of the form (2), simulations of a change in $z(s)$ to $z'(s)$ could be carried out accurately by substituting $z'(s)$ for $z(s)$ in (2), re-solving the system for new equilibrium decision rules $a'(s)$, and using this new *pair* $(z', a')$ to describe the new motion $s_{t+1} = H(z'(s_t), a'(s_t), s_t, e_t)$ of the system. That is, one would assess policies by working out the equilibrium of the new 'game' they imply. I do not see any way short of this to estimate the likely effects of policy changes.

These remarks refer to correct and incorrect ways of simulating the effects of changes in some *function* $z(s)$ that is assumed to generate individual policy decisions $z_t$, so the thought-experiment involves a once-and-for-all change from one fixed function $z$ to another, $z'$. Examples of such a policy shift would be a shift from a monetary policy directed at stabilizing interest rates to one directed at achieving a money growth target, or from a fiscal policy directed at budget balancing to one directed at some kind of 'leaning against the wind'. Yet in general, at least since the Second World War, macroeconomic policy is not discussed in terms of functions or rules at all, but in terms of selecting current policy *numbers:* this year's deficit, this year's money growth, and so on. The idea – formalized in the USA in the Employment Act of 1946 – is that policy ought not to be pre-set in *any* sense, but rather that government should have broad discretion to deal with each year's unique economic situation as it arises.

The framework I have just described does not equip the economic expert to participate in this kind of policy discussion. According to (2), agents will *have* to form an opinion as to how *future* policy is to be made in order to decide how to react to *current* policies. If an expert is asked to predict how people will react to a particular choice of $z_t$ independent of the way future $z_t$-values are to be chosen, he will simply have to say that he doesn't know, or else reformulate the question to the point where it has an answer.

This is what I meant in my introduction when I said that modern theoretical developments lead to





radical changes in the way we think about policy. If the viewpoint of dynamic games, in any form, is a useful way of thinking about policy then the use of Keynesian models to set policy on a year-to-year basis is not. Trying to reconcile these two points of view by interpreting ideas based on rational expectations as an alternative way of solving the year-to-year management problems to which the Keynesian framework addresses itself will just lead to confusion and misunderstanding. I will come back to this at the conclusion of the lectures, when I will summarize what I think economic theory *does* have to say about economic policy.

# III

If we are to move beyond the general formalism I have been discussing to models that account for specific features of business cycles, it will be necessary to be more specific about the agents in the system, about the technology at their disposal, and about the way they interact. In this section, I will consider consumer preferences only, saying nothing about the other, much more difficult, aspects of the problem. It is remarkable how much one can say about the importance of macroeconomic questions on the basis of preferences alone.

Any economic model is going to have at its center a collection of hypothetical consumers whose decisions, together with the technology and market structure, determine the operating characteristics of the system and whose welfare is the explicit subject of normative analysis. A typical household will consume a collection $c_t$ of goods at date $t$, possibly contingent on probabilistically-determined events between dates zero and $t$, and will evaluate an entire sequence, or process, $\{c_t\}$ of





consumption according to a utility function such as (1), specialized here to:

$$E\left\{\sum_{t=0}^{\infty}\beta^t U(c_t)\right\} \tag{3}$$

If we are to think about economic policy starting from this viewpoint, what we mean, in the first place, is that we want the ability to be able to determine how different policies will induce different consumption sequences $\{c_{it}\}$ for each agent $i$ in this economy. In the second place, we mean to evaluate policies normatively according to their effects on agents' welfare as measured by (3).

We can move to some sharper, quantitative conclusions for macroeconomic problems if we abstract from issues involving the *mix* of consumption goods at date $t$ and limit discussion to policies affecting consumption of goods-in-general. At the simplest level, let us identify $c_t$ with real consumption at date $t$, and specialize preferences to:

$$E\left\{\sum_{t=0}^{\infty}\beta^t \frac{1}{1-\sigma}(c_t^{1-\sigma}-1)\right\} \tag{4}$$

where $\beta \in (0, 1)$ is a constant discount factor and $\sigma > 0$ is the constant coefficient of relative risk-aversion. This two-parameter preference family will not be adequate for every problem I will want to address, but it will serve to get the discussion going in a concrete way.

Before situating this consumer (4) in a model of the sort outlined in Section II, let us examine his attitudes toward some purely hypothetical consumption streams by simply *asking* him about them. Since I am particularly interested in his attitudes toward growth and fluctuations, it will be useful to work with a class of consumption streams with 'trend' and 'cycle' components, such as:

$$c_t = (1+\lambda)(1+\mu)^t e^{-\frac{1}{2}\sigma_z^2} z_t,\ t = 0, 1, \ldots, \tag{5}$$

where $\{z_t\}$ is a stationary stochastic process with a stationary distribution given by:

$$\ln(z_t) \sim N(O, \sigma_z^2).$$

Then $E(e^{-\frac{1}{2}\sigma_z^2} z_t) = 1$, so that mean consumption under these assumptions is $(1+\lambda)(1+\mu)^t$.[1] Setting

1 The assumption that mean consumption follows a deterministic trend, and hence is perfectly predictable, is not innocuous in this context, and it would be desirable to work through the calculations below under alternative assumptions. Charles R. Nelson and Charles I. Plosser, 'Trends and random walks in macroeconomic time series, some evidence and implications', *Journal of Monetary Economics* **10** (1981), pp. 139–62, have recently argued that most year-to-year variability in US real GNP can be attributed to a random walk, or stochastic trend, component. It is likely that similar methods applied to consumption would lead to a similar conclusion. I agree with John H. Cochrane, 'How big is the random walk in GNP?' (University of Chicago Work ing Paper, 1986) that Nelson and Plosser's methods have considerably overstated the importance of the random walk component, but even so it seems clear that something inter-





the parameter $\lambda$ is just a matter of units: I will use it later on to provide 'compensation' for variations in the parameters $\mu$ and $\sigma_z^2$. For the USA the annual growth rate in total consumption is about 3 per cent, so $\mu = 0.03$ can serve as a benchmark value. For the post-war period in the USA the standard deviation of the log of consumption about trend is about 0.013 (so that average consumption that is as much as 2.5 per cent above or below trend is a 'two-sigma' event), so that $\sigma_z^2 = (0.013)^2$ is a benchmark value.[2] Thus we may take (5) with $(\lambda, \mu, \sigma_z^2) = (0, 0.03, (0.013)^2)$ as a rough description of the consumption behavior the average American family is used to, and examine its attitudes toward changes.

Given any choice of $(\lambda, \mu, \sigma_z^2)$ we could simply calculate the value of (4) under the consumption

mediate to Nelson and Plosser's model and (5) would provide a better description of consumption behavior than (5) does.

2 This figure is from Finn E. Kydland and Edward C. Prescott, 'Time to build and aggregate fluctuations', *Econometrica* **50** (1982), pp. 1345–70, table IV, p. 1365. Since 0.013 is a quarterly figure, it overstates the standard deviation of annual consumption, but not by much, since consumption is highly serially correlated. This number and others used in this section are intended to give a rough idea of the relative importance of certain issues. The reader will agree, I think, upon reaching the end of the section that its conclusions do not hinge on delicate questions of measurement.

behavior (5) and call the indirect utility function so defined $U(\lambda, \mu, \sigma_z^2)$. But we will obtain a measure that is easier to think about if we use *compensating* variations in $\lambda$ to evaluate various $\mu$ and $\sigma_z^2$ changes. To evaluate changes in the growth rate $\mu$, for example, let us define $f(\mu, \mu_0)$ by:

$$U(f(\mu, \mu_0), \mu, \sigma_z^2) = U(0, \mu_0, \sigma_z^2), \tag{6}$$

so that $f(\mu, \mu_0)$ is the percentage change in consumption, uniform across all dates and values of the shocks, required to leave the consumer indifferent between the growth rates $\mu$ and $\mu_0$. (In general, $\sigma_z^2$ would appear as an argument of $f$, too, but under these 'constant relative risk-aversion' preferences it drops out.) A direct calculation gives:

$$f(\mu, \mu_0) = \left(\frac{1+\mu_0}{1+\mu}\right)^{\beta/(1-\beta)} - 1.$$

Here is a table of this function $f$, which I will call simply the *cost* of reduced growth, for $\beta = 0.95$ and a base growth rate of $\mu_0 = 0.03$.

At the parameters used in table 1, then, consumers would require a 20 per cent across-the-board consumption increase to accept voluntarily a reduction in the consumption growth rate from 0.03 to 0.02, and would surrender 42 per cent across the board to obtain an increase in the growth rate from 0.03 to 0.06.

I hasten to add that I have said nothing about the *feasibility* of these growth rate changes; I am





TABLE 1 Cost of reducing growth from $\mu_0 = 0.03$ when $\beta = 0.95$

| $\mu$ | $f(\mu, \mu_0)$ |
|---|---|
| 0.01 | 0.45 |
| 0.02 | 0.20 |
| 0.03 | 0.00 |
| 0.04 | −0.17 |
| 0.05 | −0.31 |
| 0.06 | −0.42 |

simply running experiments on this fictional consumer. Indeed, under a standard, neoclassical technology, policies that affect growth usually do so only over a transient period only, not permanently as in table 1. But the range of growth rates in table 1 is not large relative to what we observe across countries, and the welfare consequences of 'small' changes are *enormous*, relative to anything we will see in what follows. I shall return to this later on.

The costs of economic instability can be measured in a way that is identical conceptually to this way of measuring the costs of reduced growth. To this end, define the function $g(\sigma_z^2)$ by

$$U(g(\sigma_z^2), \mu, \sigma_z^2) = U(0, \mu, 0). \tag{7}$$

That is, $g(\sigma_z^2)$ is the percentage increase in consumption, uniform across all dates and values of the shocks, required to leave the consumer indifferent between consumption instability of $\sigma_z^2$ and a *perfectly* smooth consumption path. I will call $g(\sigma_z^2)$ the cost of consumption instability.

By direct calculation, and using the approximation $\ln(1+\lambda)=\lambda$ (which is entirely safe in this context), $g$ is given by:

$$g(\sigma_z^2) \cong \tfrac{1}{2}\sigma \cdot \sigma_z^2. \tag{8}$$

Table 2 shows the function $g$, for various $\sigma$ and $\sigma_z^2$ values. The coefficient $\sigma$ of risk-aversion can be estimated from a variety of different samples: estimates vary widely. A value of unity means logarithmic preferences; people appear to be more risk-averse than this. No available estimates are as large as 20, but some do exceed 10.

The value 0.013 is the standard deviation of the log of US real quarterly consumption, expressed as a deviation from fitted trend, over the period following the Second World War. Eliminating

TABLE 2 Cost of consumption instability: $g(\sigma_z^2)$

| $\sigma$ | $\sigma_z$ | | |
|---|---|---|---|
| | *0.013* | *0.039* | *0.120* |
| 1 | 0.00008 | 0.00072 | 0.0065 |
| 5 | 0.00042 | 0.0038 | 0.034 |
| 10 | 0.00084 | 0.0076 | 0.068 |
| 20 | 0.0017 | 0.015 | 0.136 |





aggregate consumption variability of this magnitude entirely, would, from table 2, be the equivalent in utility terms of an increase in *average* consumption of something less than one tenth of a percentage point. (Total US consumption in 1983 was $2 trillion, so one-tenth of 1 per cent is $2 billion, which sounds like a sizeable free lunch. But there were 234 million people to feed, so lunch will have to run about $8.50 per person.) I want to propose taking these numbers seriously as giving the order-of-magnitude of the potential marginal social product of additional advances in business cycle theory – or more accurately, as a loose upper bound, since there is no reason to think that eliminating *all* consumption variability is either a feasible or a desirable objective of policy. But I imagine that even one-tenth of a percentage point will seem to many to be an extremely low estimate of the costs of economic instability – at least, it did to me – so it will be useful to digress to discuss some aspects of this estimate.[3]

3 The estimates in table 2 appear somewhat less surprisingly low when compared to estimates of the welfare gains from other (also purely hypothetical) policy changes. For example, Arnold C. Harberger, 'Monopoly and resource allocation', *American Economic Review* **44** (1954), pp. 77–87, found one-tenth of 1 per cent of income to be an upper bound on the welfare gain from the (costless) elimination of *all* product market monopoly in the US economy. Harberger, too, was led to 'confess that I was amazed at this result' (p. 86), but it has not been revised upward by any

The last two columns of table 2 set up what seem to me the two most important qualifications or elaborations of this cost estimate. First, in the period prior to the Second World War, and extending as far back in time as we have usable data, the standard deviation (logarithmic deviations from trend) of consumption was about three times its post-war level.[4] Since this number is squared in the

quantitative research I have seen in the 30 years since his study was published. Perhaps we should be open to the possibility that the intrinsic importance of substantive economic questions is not accurately reflected by the number of journal pages devoted to them.

It is not quite accurate to identify instability in goods consumption with economic instability in general, since consumption of leisure also fluctuates. But since hours worked and goods consumption are positively correlated cyclically, I would guess that taking leisure fluctuations into account more carefully would *reduce* the estimates in the text still further.

4 In a recent paper, Christina Romer, 'Spurious volatility in historical unemployment data', *Journal of Political Economy* **94** (1986, 1–37), has argued that pre-First (not second) World War variability in unemployment rates was, correctly measured, no larger than post-Second World War variability. It appears from related work of hers that the amplitudes of other series (probably including real consumption) were also badly overstated in pre-First World War data. The statement in the text may then rest much more heavily on the experience of the 1930s than I would previously have thought. It would be hard to overstate the importance of the questions on which these findings bear, but I have not attempted to incorporate them into my illustrative calculations.





formula (8), the implied cost estimates are multiplied by nine, becoming something like one-half of 1 per cent of total consumption. As deadweight losses go, *this* is a large number. Second, fluctuations in total consumption do not affect all households equiproportionally, so that variability of total consumption does not capture anything like all of the consumption risk faced by the typical household. Perhaps correcting for this effect would lead to another tripling of the relevant standard deviation, and hence another multiplication by nine in the cost estimates. This is the last column of table 2.

But as a measure of the possible gains from improvements in aggregative policy, this last column is way too high. In so far as the absence of income-risk pooling reflects 'imperfections' in capital markets, and I think it does, the cost of *individual* income variability measures the potential or actual gain from social insurance, not from stabilization policy. Aggregate income variability is but one source of individual income risk, and reduction of aggregate variability – which is all that stabilization policies can accomplish – cannot be expected to eliminate more than a small part of the uninsurable risk borne at the individual level. I will return to the issue of social insurance later on, in section V, when we have invested in a framework more suitable for posing questions about individual earnings risk and ways of dealing with it.

An economic system is a collection of people and serious evaluation of economic policy involves tracing the consequences of policies back to the welfare of the individuals they affect. Without saying much more about the nature or workings of the economy than this, we can get a good if rough idea of the potential benefit of policies that alter individual consumption streams in various ways. I have run through two exercises to assess the potential welfare gains of policies that affect the growth of consumption and policies that affect the variability of consumption about its trend, not by describing policies that would have these effects, but simply by imagining that these effects somehow come about. It is worth re-emphasizing that these calculations rest on assumptions about preferences *only*, and not about any particular mechanism – equilibrium or disequilibrium – assumed to generate business cycles.

I find the exercise instructive, for it indicates that economic instability at the level we have experienced since the Second World War is a minor problem, even relative to historically experienced inflation[5] and certainly relative to the costs

5 In Robert E. Lucas, Jr, 'Discussion of: Stanley Fischer, "Towards an understanding of the costs of inflation: II"', *Carnegie–Rochester Conference Series on Public Policy* **15** (1981), pp. 43–50, I estimated the annual social cost of a sustained 10 per cent inflation in the USA to be 0.5 per cent of national income. The estimate follows the method used in Martin J. Bailey, 'The welfare cost of inflationary finance', *Journal of Political Economy* **64** (1956), pp. 93–110, and uses the interest elasticity of money demand estimates in





of modestly reduced rates of economic growth. This is not to say that economic fluctuations are a trivial problem, for fluctuations at the pre-Second World War level, especially combined as they were with an absence of adequate programs for social insurance, were associated with large costs in welfare. But it suggests that the main social gains from a deeper understanding of business cycles, whatever form this deeper understanding may take, will be in helping us to see how to avoid large mistakes with policies that have minimally inefficient side-effects, not in devising ever more subtle policies to remove the residual amount of business-cycle risk.

Phillip Cagan, 'The monetary dynamics of hyperinflation', in Milton Friedman (ed.) *Studies in the Quantity Theory of Money* (University of Chicago Press, Chicago, 1956).

# IV

It is possible to get a rough idea of the order of magnitude of the potential gains in welfare from stabilization policies by the method of the preceding discussion, and I thought this exercise would be a useful preliminary. But to go beyond the calculation of crude upper bounds it will be necessary to take a position on the way the economy actually works, or to construct a positive theory of business cycles. In this section, I will turn to the discussion of a simple prototype model that I think holds much promise for future developments.

There are many interesting prototype business cycles models in existence now, and I have no doubt that many of them will contribute, in different ways, to the development of improved future models. Of these, the most useful for the present discussion is one introduced recently by Kydland and Prescott.[1] This model focuses exclusively on

1 Kydland and Prescott, 'Time to build and aggregate fluctuations'. Technically, the immediate ancestor of Kydland and Prescott is William A. Brock and Leonard Mir-





real (as opposed to monetary) neoclassical considerations, which I think is a mistake, but it is the only model I know of that is theoretically coherent in the sense we discussed in section II, while yet having been developed to the point where its implications can be compared to observed time series in a quantitatively serious way.[2]

The Kydland and Prescott model is a highly simplified, competitive system, in which a single good is produced by labor and capital with a constant returns technology. All consumers are assumed to be infinitely-lived and identical. The *only* 'shocks' to the system are exogenous, stochastic shifts in the production technology. Kydland

man, 'Optimal economic growth and uncertainty: the discounted case', *Journal of Economic Theory* **4** (1972), pp. 479–513.

2 Two other interesting prototypes that have been compared to actual time-series data are described in Thomas J. Sargent, 'A classical macroeconometric model for the United States', *Journal of Political Economy* **84** (1976), pp. 207–38; and John B. Taylor, 'Estimation and control of a macroeconomic model with rational expectations', *Econometrica* **47** (1979), pp. 1267–86. These models have the advantage of incorporating monetary as well as real disturbances, but since neither explicitly relates its structure to assumptions about technology and preferences, they are more difficult to relate to the theoretical structure I sketched in section II. On the other hand, illustrative models like that in Robert E. Lucas, Jr, 'Expectations and the neutrality of money', *Journal of Economic Theory* **4** (1972), pp. 103–24, are too abstract to be compared in any detail to observed aggregate time-series.

and Prescott ask the question: 'Can specific parametric descriptions of technology and preferences be found such that the movements induced in output, consumption, employment and other series in such a model by these exogenous shocks resemble the time series behavior of the observed counterparts to these series in the postwar, US economy?' This seems to me exactly the right question for macroeconomists to ask, and I want to discuss three aspects of Kydland and Prescott's answer to it: (1) the methods they used to carry out the simulation; (2) the way they make the critical term 'resemble' operational; and (3) the success the model can claim as an explanation for business cycles.

In describing the model, it will be easiest to begin with a system that is simpler than the one Kydland and Prescott used, and then to describe their model as a variation on this simpler one. Let the typical household be endowed with $\bar{n}$ units of time each period, and let its current period utility depend on goods $c_t$ consumed and 'leisure' $\bar{n}-n_t$, where $n_t$ is labor sold to firms. Preferences are assumed to be:

$$E\left\{\sum_{t=0}^{\infty} \beta^t U(c_t, \bar{n}-n_t)\right\}.$$

The technology is $F(k_t, n_t, x_t)$, where the value of $F$ is the units of output that can be produced with $k_t$ units of capital, and $n_t$ man-hours of labor





when the stochastic technology shock is $x_t$. These shocks $x_t$ follow a Markov process with transitions

$$G(x', x) = Pr\{x_{t+1} \leqslant x' \mid x_t = x\}.$$

$F$ is homogenous of degree one in $(k_t, n_t)$, so we can interpret all variables in per-household terms. Output is divided into consumption $c_t$ and gross investment $i_t$ and capital evolves according to

$$k_{t+1} = i_t + (1-\delta)k_t.$$

Households own all factors of production, renting them to profit-maximizing firms each period at wages and capital rentals $w_t$ and $u_t$ (with the price of current output normalized at unity). These markets are, as I have said, competitive. Households' expectations about future factor prices are rational. Once the functions $U$, $F$ and $G$ are specified (as Kydland and Prescott do) this is all that needs to be said: working out the predictions of the model is just a matter of technique.

But technique is interesting to technicians (which is what we are, if we are to be of any use to anyone) so let me go into a little more detail, by casting the decision problem faced by a typical household into the formalism I developed earlier. From the household's point of view, the state of this system is described at each date by three numbers: its own holdings of capital, $y$ (say), the capital stock in the economy as a whole, $k$, and the current technology shock, $x$. (In equilibrium, we know that $y$ and $k$ will have to be equal – otherwise the typical household would not be typical – but this equality cannot be *imposed* on the household: prices have to move to make equality *desirable* to it.) Then the state $s$ of the system will be this triple $(y, k, x)$. I will work toward a statement of the Bellman equation for the household's value function $v(s)$.

At each date, the action $a$ chosen by the household (since all households are alike, we don't need the subscript $i$) is the triple $(c, n, y')$ describing its consumption, labor supply and end-of-period capital holdings. The immediate return from any such action is just $R(a) = U(c, \bar{n} - n)$. The opportunity set $\Omega$ from which an action is selected is determined by the current factor prices $w(k, x)$, $u(k, x)$ (or $w(s)$, $u(s)$) which depend on the economy's state $(k, x)$, but *not* on the individual's holdings $y$: this is what competition means, in this context. Thus:

$$\Omega(y, k, x) = \{(c, n, y'): c + y' \leqslant w(k, x)n + u(k, x)y + (1-\delta)y,\ c \geqslant 0,\ 0 \leqslant n \leqslant \bar{n},\ y' \geqslant (1-\delta)y\}.$$

In an equilibrium, the next period capital stock for the economy as a whole, $k'$, will be some function $h(k, x)$ of today's state. Rational expectations implies that this function is known by agents, along with the functions $G$, $w$ and $u$. Then $v(s)$ must satisfy:





$$v(s) = \max_{(c,n,y')\in\Omega} \{U(c, \bar{n} - n) + \beta \int v(y', h(k, x), x')dG(x', x)\}, \tag{9}$$

where $\Omega$ is the set defined above. Equation (9) thus describes the decision problem faced by a household deciding on its consumption, labor supply and savings, given current factor prices $w(s)$ and $u(s)$ and given expectations about the way these prices will behave in the future. These expectations, in turn, can be calculated from knowledge about the current state of the system, $(k, x)$, the way the distribution $G$ of future exogenous shocks depends on this state, and the way the capital stock of the economy evolves, $h$.

If households behave according to (9), their decisions $(c, n, y')$ will be fixed functions of the state $(y, k, x)$ that sets the terms of this maximum problem. In particular, the household's own capital holdings will evolve according to a difference equation:

$$y_{t+1} = y(y_t, k_t, x_t),$$

where the function $y(\cdot)$ can be obtained from the study of the maximum problem (9). The economy's capital stock, we have assumed, evolves according to:

$$k_{t+1} = h(k_t, x_t).$$

In equilibrium, it must be the case that $k_t = y_t$ for all $t$ (that is, that all capital is held) which thus translates into the condition that $y(k, k, x) = h(k, x)$, identically in $(k, x)$. That is to say, the system is in a rational expectations equilibrium when the savings behavior each household *believes* others will follow coincides with the savings behavior each household finds it optimal to follow, given its expectations about others.

In fact, Kydland and Prescott did not proceed directly to solve (9) and to construct the equilibrium function $h$ in the way I have just sketched. Instead, they calculated solutions to the planning problem:

$$f(k, x) = \max_{c, n, k'} \{U(c, \bar{n} - n) + \beta \int f(k', x') G(dx', x)\} \tag{10}$$

subject to $c \geqq 0, 0 \leqq n \leqq \bar{n}, k' \geqq (1 - \delta)k$, and

$$k' + c' \leqq F(k, n, x) + (1 - \delta)k.$$

Equation (10) describes the behavior of a planner deciding on households' consumption, labor supply and capital accumulation, given constraints imposed by the technology $F$ and the current shock $x$. It is a classical fact that the optimal capital accumulation behavior $k_{t+1} = h(k_t, x_t)$ for this hypothetical planning problem *coincides* with the competitive equilibrium accumulation behavior I have just described.[3] The study of (10) thus

3 Gerard Debreu, 'Valuation equilibrium and Pareto optimum', *Proceedings of the National Academy of Sciences* **40** (1954), pp. 588–92.





provides an indirect method for calculating the function $h$ describing the competitive equilibrium motion of the endogenous state variables in an inexpensive way, sidestepping the simultaneity I discussed above.

Having obtained numerical solutions to the model in this way, the equilibrium behavior of the capital stock can be simulated by drawing shocks $\{x_t\}$ from the assumed distribution $G(x', x)$ and running the difference equation $k_{t+1} = h(k_t, x_t)$. Since the values of consumption, employment and factor prices are all given by the theory as functions of $(k_t, x_t)$, the model generates time series for these variables as well. This is precisely the method Slutsky used in his 1927 paper,[4] in which he demonstrated for the first time that stochastic difference equation systems could generate behavior that closely resembled economic time series. The difference is that in this case the function $h(\cdot)$ has a clear economic interpretation in terms of preferences and technology.

The artificial time series so generated by the theoretical model 'look like' economic time series in the sense that the series Slutsky generated did: the variables show erratic, serially correlated fluctuations about their mean values. This much could be guessed from the economics of the model's

4 Eugenio Slutsky, 'The summation of random causes as the source of cyclic processes', *Econometrica* **5** (1937), pp. 105–46.

structure: a favorable technology shock shifts out current production possibilities; this induces high capital accumulation which spreads this benefit forward into future periods. But a more detailed comparison of the artificially generated series with their observed counterparts is not so encouraging. The employment movements predicted by the model have lower amplitude (relative to output movements) than do actual employment variations. Consumption is more volatile and investment much less so in the model as compared to actual data.

To deal with these discrepancies, Kydland and Prescott's published model used a modified formulation of both preferences and technology. For preferences, the 'leisure' argument $\bar{n} - n_t$ in current utility was replaced with a distributed lag of current *and* past values of $\bar{n} - n_t$. This has the effect of increasing the degree of *intertemporal substitutability* of leisure without altering the assumed intertemporal substitutability of consumption.[5] On the technology side, the assumption

5 Gary D. Hansen, 'Indivisible labor and the business cycle', *Journal of Monetary Economics* **16** (1985), pp. 309–28; and William Rogerson, 'Indivisible labor, lotteries, and the business cycle', unpublished University of Rochester working paper (1985), have modified the preference structure described above in a different direction, by assuming that household must supply labor in an all-or-nothing fashion, so that a 10 per cent decrease in aggregate labor supply comes about by a 100 per cent reduction in the labor supplied by





that investment at date $t$ augments the stock of productive capital at date $t+1$ was replaced with a gestation lag scheme with the property that an investment project initiated at date $t$ comes to completion at date $t+\tau$, $\tau$ fixed, with investment expenditures being incurred at all dates in between. The nature of the technology shock was also altered somewhat, so that it is assumed to consist of 'permanent' (that is, highly persistent) and transient components, in a mix that cannot be observed by agents. The effect of these modifications is to increase the dimension of the state space in which $s_t$ lies, and markedly to improve the model's ability to fit data. The underlying logic of the model and its solution method remains, however, exactly as I have described.

It is instructive to run a simulated 'boom' through the Kydland and Prescott model. Suppose a high technology shock occurs, increasing the current productivity of both capital and labor. This makes the current period attractive to work and produce, relative to conditions that are expected to prevail in future periods, so both employment and output rise. It also *may* signal high productivity in

10 per cent of households, rather than a 10 per cent reduction by all households. It is remarkable (but true) that this variation, in some ways more realistic, can still be analyzed within an identical-household framework. Hansen's work shows that a Kydland–Prescott model, so modified, permits much higher employment variability than even the modified Kydland–Prescott model as reported in tables 3 and 4.

future periods, and the only way for firms to hedge against this (attractive) contingency is to initiate investment projects now. The projects so initiated will operate to increase output and employment until they are completed, spreading the effects of this shock – even if it should turn out after the fact to be transient – forward into future periods. They also carry within them the seeds of a future downturn, both because they increase the capital stock – possibly inappropriately – and because workers will be less willing to supply labor in future periods, having extended themselves in sustaining the boom.

Can such a scenario generate the investment-dominated fluctuations we seem to observe? It is hard to tell: the interactions are too complicated, even in so stripped-down a model, to work out in one's head or with pencil and paper. To pursue this question more systematically, Kydland and Prescott began by estimating as many parameters as possible from a wide variety of out-of-sample evidence. For example, the fact that people work about one-third of the time pinned down one preference parameter; the observation that investment projects take something like a year to complete was used to fix a technological parameter; and so forth. Having estimated as many parameters as they could in this way, without even looking at the time series they were attempting to fit, the number of free parameters – including the critical parameters characterizing the technology shocks





that drive the system – was reduced to about six. Kydland and Prescott then chose values for these remaining parameters so as to make certain low order moments (variances, covariances, autovariances) predicted by the model 'match' the corresponding moments from the collection of time series in the sample they used. The result of this last step completed the estimation, and the matches between the theoretical and actual moments they reported are the only reported 'test' of the model's ability to 'fit' these series.

Tables 3 and 4 summarize some of the main features of the comparisons reported by Kydland and Prescott. Both tables compare sample moments to moments implied by the estimated model for deviations of the indicated series about a fitted trend. In the simulations underlying these tables, the variance of the technology shocks was chosen so as to make the standard deviations

TABLE 3 Auto-correlations of output

| *Time lag (quarters)* | *Model* | *US economy 1950–79* |
|---|---|---|
| 1 | 0.71 | 0.84 |
| 2 | 0.45 | 0.57 |
| 3 | 0.28 | 0.27 |
| 4 | 0.19 | −0.01 |
| 5 | 0.02 | −0.20 |
| 6 | −0.13 | −0.30 |

Table 4 Standard deviations and correlations with real output

| *Variable* | *Model* | | *US economy* | |
|---|---|---|---|---|
| | *Standard deviation* | *Correlation with output* | *Standard deviation* | *Correlation with output* |
| **Real GNP** | 1.8 | – | 1.8 | – |
| **Consumption (non-durable)** | 0.9 | 0.66 | 0.6 | 0.94 |
| **Investment** | 6.4 | 0.80 | 5.1 | 0.71 |
| **Man-hours** | 1.0 | 0.93 | 2.0 | 0.85 |



(about trend) for real GNP for the model equal to its value for the post-war US economy: this is not a test of the model. The difficult moments to fit (simultaneously) are the standard deviations of consumption, investment and man-hours, and the formulations of preferences and technology that I described a moment ago are motivated exactly to bring these three predictions into closer correspondence to what we observe.

Whether these results are viewed as 'good' or 'bad' is a difficult question, as is the related question of which comparisons of theoretical to sample moments are most interesting. One could obtain a formal sharpening of these questions by using the discipline of classical hypothesis testing (and a recent paper by Sumru Altug shows that this route is indeed instructive[6]) but the interesting question raised by the Kydland and Prescott model is surely not whether it can be accepted as 'true' when nested within some broader class of models. Of course the model is not 'true': this much is evident from the axioms on which it is constructed. We know from the outset in an enterprise like this (I would say, in *any* effort in positive economics) that what will emerge – *at best* – is a workable approximation that is useful in answering a limited set of questions.

6 Sumru Altug, 'Estimation and Tests of an Aggregate Equilibrium Model', unpublished Federal Reserve Bank of Minneapolis working paper, 1985.

Kydland and Prescott do not say much about which questions they hope their model could simulate accurately, or with what level of accuracy, but the model is set up to focus on the way firms and consumers react to changes in the intertemporal pattern of actual and expected prices: hence their focus on the opportunities agents have for substituting across time – for working now and postponing leisure, for initiating investment now as opposed to waiting for more information, and so on. They parameterize as many of these trade-offs as they can in such a way as to facilitate bringing as wide a variety of evidence as possible to bear on these questions. Thus the coefficient of risk-aversion in consumer preferences can be deduced from cross-section patterns in risk premia as well as from investments in aggregate consumption, the gestation lags for investment projects can be observed directly, at least for some kinds of investment, and so on. This is the point of 'microeconomic foundations' of macroeconomic models: to discover parameterizations that have interpretations in terms of specific aspects of preferences or of technology, so that the broadest range of evidence can be brought to bear on their magnitudes and their stability under various possible conditions.

Kydland and Prescott have taken macroeconomic modeling into new territory, with a formulation that combines intelligible general equilibrium theory with an operational, empirical seriousness





that rivals at least early versions of Keynesian macroeconometric models. Exactly *because* their model carries predictions for so wide a range of evidence, it has been subjected to an unusually wide range of empirically-based criticism: here is a macroeconomic model that actually makes contact with microeconomic studies in labor economics! The chances that the model will survive this criticism unscathed are negligible, but this seems to me exactly what explicit theory is for, that is, to lay bare the assumptions about behavior on which the model rests, to bring evidence to bear on these assumptions, to revise them when needed, and so on.

The Kydland and Prescott model is another in a long and honorable (though recently dormant) line of real business-cycle models. Substantively, the model reopens a debate that played an important role in pre-Keynesian theory – a debate that Haberler surveyed so masterfully in *Prosperity and Depression.*[7] But this time around, the terms of the discussion are explicit and quantitative, and the relationship between theory and evidence can be (and is being) argued at an entirely different level. I would like to call this progress.

7 Gottfried Haberler, *Prosperity and Depression* (League of Nations, Geneva, 1937).