Frank Hahn and Robert Solow

A Critical Essay on Modern Macroeconomic Theory

A Critical Essay
on Modern
Macroeconomic
Theory

Frank Hahn and Robert Solow

Contents

1   Introduction

It may help the reader to understand the direction taken by the following chapters if we say at the outset that the originating impulse that led to this book was fundamentally *negative*. We found that we shared a profound disagreement with the main trend of macroeconomic theory in the early 1980s, and we wanted to create some sort of respectable theoretical resistance to it.

We had better describe the trend we were resisting. It is roughly the line of thought that runs from Lucas's famous price-misperception model of 1974 to the "real business cycle theory" of Prescott, Kydland, and others today. Its essential characteristic is *not* that it "pays attention to micro foundations." As many people have noticed, macroeconomics has always done that, at least in the sense that aggregative relationships have always been explicated and justified by reference to microeconomic behavior. (Just think of the transactions motive, the precautionary motive, and the speculative motive in the demand for money, or of the "psychological law" that is supposed to underlie the propensity to consume.) The distinctive flavor of post-Lucas macro theory comes from a more special and more powerful intellectual commitment.

There seem to be two parts to it. The first is the belief that a valid macro model should be the exact aggregation of a micro model, or very nearly that. It is hard to argue with that belief except by looking at its consequences. Any macroeconomist would like to be reassured about aggregation biases. So no issue of principle can be involved. The only question is whether incremental relaxation of the principle buys enough in the way of scope, realism, flexibility, and tractability to be worthwhile. The second part is the belief that the only appropriate micro model is Walrasian or intertemporal-Walrasian or, as the saying goes, based exclusively on intertemporal utility maximization subject only to budget and technological constraints. This commitment, never really argued, is the rub.

The conjunction of these two beliefs leads to the sort of model we were, and are, reacting against. It proposes that the actual economy can be read as if it is acting out or approximating the infinite-time discounted utility maximizing program of a single, immortal "representative agent."[1] The only admissible constraints come from initial resources, a supply of labor, and a well-behaved technology for turning produced means of production and labor into consumer goods and produced means of production. There is simply no possibility of coordination failure. That means the economy accurately carries out the wishes of the representative agent. Under favorable conditions, of course, production decisions can be decentralized. The optimizing program will be the competitive equilibrium for an economy equipped with a complete set of Arrow-Debreu markets, all open at time zero. Alternatively it must be equipped with perfect price foresight for each state of nature and a full set of insurance possibilities.

The temptation has proved irresistible. It has become good form to treat just such a model as a descriptive macro model that need only be estimated or calibrated and then directly applied to this economy or that. We have a macroeconomics squarely based on perfect foresight, infinite time optimization, and universal perfect competition. What Ramsey took to be a normative model, useful for working out what an idealized omniscient planner should do, has been transformed into a model for interpreting last year's and next year's national accounts.

Of course that is the economics of Dr. Pangloss, and it bears little relation to the world. In a decade that has seen vast progress in our study of asymmetric information, "missing markets," contracts, strategic interaction, and much else precisely because those aspects are regarded as real phenomena that require analysis, macroeconomics has ignored them all. The consequence is this: no account has been given of how and why a decentralized economy could behave as if guided by a Ramsey maximizer. It is true that an Arrow-Debreu equilibrium is an allocation that maximizes a special social welfare function, but that is not the case, for instance, when some insurance markets are absent, or indeed when any even mildly realistic phenomena are included.

The irony here is that macroeconomics began as the study of large-scale economic pathologies: prolonged depression, mass unemployment, persistent inflation, etc. This focus was not invented by Keynes (although the depression of the 1930s did not pass without notice). After all, most of Haberler's classic *Prosperity and Depression* is about ideas that were in circulation before *The General Theory*. Now, at last, macroeconomic theory has

as its central conception a model in which such pathologies are, strictly speaking, unmentionable. There is no legal way to talk about them.

We found that we could not swallow this way of doing macro-economics, precisely because those umnentionables need mentioning, even if they should turn out in the end to be illusions. In the modern spirit, however, resistance has to begin with alternative micro foundations. And that explains the plan of this book. We have not tried for completeness, but are content with a few critical and constructive essays. In some ways, of course, we aim to preserve or restore the respectability of "Keynesian" ways of thinking. But we want to emphasize that defending Keynesian doctrine per se is not part of our intention. In many respects we think Keynes was on the right track in *The General Theory*, although he lacked— as did everyone in 1936—the techniques and concepts that could have let him clinch the case for his insights (e.g., for the existence of unemployment equilibrium). If we can preserve those insights, so much the better. But ideological piety is not our aim or motive.

Chapter 2 is a good example. We alter the standard assumptions in two normally acceptable ways. First, we think it requires no special defense to justify abandoning the immortal, all-seeing representative agent. That assumption is entitled to no priority, either in fact or in theory. We replace it by another conventional assumption, that of overlapping generations, with two-period lives and no bequests. We would not defend this formalization either, but our purpose is negative and we are entitled to see where it leads. Second, we insist on an essentially monetary economy, so that savers have to allocate their savings between holdings of outside money and bonds issued by firms to finance capital investment. At this stage we stick to the assumptions of universal perfect competition and perfect (but only short-term) foresight. The markets for goods and labor clear at every instant, through price and wage movements.

From this rather conventional model we derive a number of anti-Panglossian conclusions. To begin with, the adjustment dynamics are very badly behaved. When the model economy is disturbed from a steady-state equilibrium, it does not respond well. It may return to steady state—but equally well it may not. It may oscillate; it may go off on an unstable path. Interestingly, the mechanism at work here is one of Keynes's flashes of insight. A contractionary shock induces a fall in the nominal wage, and prices will normally follow. If the resulting deflation is sharp enough, the real interest rate must rise because the nominal interest rate cannot fall below zero. Investment is then depressed and the economy suffers an

unnecessary fall in output. Finally, we show that there is an appropriate fiscal and monetary policy that could, if operated in time, move the economy from its old steady state to a new one with minimal disruption. It may be very difficult to calculate the required policy and put it into effect quickly. But the model makes its point. There is a clear role for stabilization policy, even with perfect flexibility of wages and prices—in a sense *because of* wage and price flexibility.

In chapter 3 we produce another unconventional result from this fairly conventional model. We introduce some wage stickiness. For technical reasons, this is done through a real-wage Phillips curve. Of course unemployment becomes a possibility; the labor market does not clear instantaneously. On the other hand, we show that some stickiness of this kind will often be stabilizing. The model will be less likely to be unstable after a disturbance. One can imagine that the residents might actually prefer a little more stickiness to a little less, within limits. By the way, it can be proved that the sort of stabilizing policy shown to be possible in the perfect-flexibility economy remains possible in the imperfect-flexibility case, although there it is necessarily more complicated because it has to take account of the Phillips curve.

Chapter 4 then makes a more drastic break with Panglossian assumptions: we allow goods-producing firms to have increasing returns to scale and correspondingly some monopoly power. The market form is then large-group monopolistic competition. Otherwise the situation is much as in chapter 2. The overlapping-generations structure is maintained. Two main conclusions are obtained. First, it emerges quite naturally that certain income aggregates appear as shift factors in the firms' demand curves. As long as there is perfect foresight, this does not matter very much. But later, in chapter 6, in a closely related model we will allow imperfectly competitive firms to make mistakes in predicting the location of their demand curves. In any case, they will make their own plans contingent on what amount to aggregative forecasts. (This sounds comfortably like the sort of thing one observes.) Second, we show that the presence of increasing returns to scale and imperfect competition opens the way to multiple equilibria.

Before pursuing this line of thought further, we turn in chapter 5 to the labor market. All the earlier action is centered in the market for goods. Except for chapter 3 the labor market clears instantaneously, but even in chapter 3 it will clear eventually. Chapter 5 is motivated by the conviction that involuntary unemployment is more than a transitory phenomenon in modern industrial economies. In that case, a good macroeconomics

would model the labor market in a way that permits persistent unemployment. There are in the literature well-worked-out stories that do just that: efficiency-wage and insider-outsider models are the two standard examples, and we endorse both of them. We add a couple of ideas of our own. Perhaps the more interesting is a formalization of the notion that excess supply of labor can persist because workers (and employers) regard wage undercutting as a violation of a social norm—as "unfair." Just to say that is to say very little. We go further by exhibiting such behavior as an equilibrium strategy for a repeated game. We think of this game as metaphorical rather than actual. It is a way of giving an account of the possibility that a customary real wage level will not be eroded by the mere presence of some involuntary unemployment, so long as there is not too much of it. Employment can then vary in an (endogenously determined) interval, while real wages are sticky. A second idea is a modification of standard labor-market search theory. It is conventionally assumed that workers not currently employed can find offers of jobs more or less ad lib. The stock of unemployed workers is maintained because some of those offers are quite rationally rejected in the (rational) hope of getting a better one. We take more seriously the notion that some number of non-employed workers are without any job offer and would take any one that came along. In this setting, long-lasting jobs are desirable, but their prevalence may cause the stock of unemployed workers to be larger.

The point of this chapter is twofold. Any approach to macroeconomic theory has to take a stand on the behavior of the labor market. Unemployment is too central a phenomenon to be ignored. An essential part of any non-Walrasian approach must be an account of why the labor market fails to clear. We gave the opposite assumption—market-clearing—a run in the earlier chapters. It is not enough to rescue a reasonable macro model from other pathologies. Our preference is obviously on the non-Walrasian side, either our own contribution or one of the others.

Any of those non-Walrasian models—and this is the second point of chapter 5—will yield a locus of real wage rates and employment levels that leave the labor market in equilibrium from the supply side. Almost always this equilibrium locus has positive slope: higher employment goes with a higher real wage. The reason differs from model to model. (For instance, in efficiency-wage theory, when there is less unemployment jobs are easier to find, and it takes a higher real wage to deter turnover or shirking.) Our stories also yield such a locus. One interesting feature of the fairness model is that the equilibrium locus will have a flat interval (horizontal if employment is measured horizontally) at the going wage.

Employment can vary, within limits, without disturbing the conventional wage. Since the conventional wage might be high or low or in-between, there is a sense in which the equilibrium locus is "thick," a two-dimensional rather than a one-dimensional curve. Interesting micro foundations make for interesting macroeconomics.

We use this construction, along with the results of chapter 4, in chapter 6, which is in a way our destination. There we try to distill a macro model that corresponds in a reasonable but not finicky way to the micro-economics of the earlier chapters. It has large-group monopolistic com-petition (with increasing returns to scale) as its market form. Imperfectly competitive firms have to predict the position of their demand curves, essentially by forecasting aggregate demand. (We assume that they know the constant elasticity of demand, but that is just a convenient simplifica-tion.) They must also form an expectation about the general price level for competing substitute goods. On the basis of these expectations, they choose a price and production plan and finance irreversible investment decisions by selling securities to savers. In the event, expectations may be disappointed or exceeded; realized demand, realized prices, and realized profits will differ from anticipated values. Firms must revise their expecta-tions—a process on which we have nothing new to say—and proceed. One sort of medium-run equilibrium is reached when expectations are confirmed by events.

It is a very simple macro model, and not one we would defend unto death. When combined with an appropriate model of the labor market, however, it is capable of doing some of the things we wanted it to do. It can generate fluctuating output and unemployment. It can exhibit a posi-tive correlation between the real wage and employment. It offers scope for corrective monetary policy. (It would surely do the same for fiscal policy, but as it stands the model has no government to tax, borrow, and spend.) Above all, it allows for a variety of short-run paths, depending on entre-preneurial beliefs, and also a multiplicity of medium-run equilibria, some of which are pretty clearly more desirable than others. The point to keep in mind is that these various outcomes are a consequence of the model's micro foundations, not a violation of them. It is all a matter of choosing interesting and plausible micro foundations. There is plenty of room for maneuver.

The appendix to chapter 6 continues by reporting some computer simu-lations of the out-of-equilibrium behavior of this model. The purpose is to get a first look at the way in which certain key parameters (degree of monopoly, degree of increasing returns to scale, characteristics of the de-

mand for equities) affect the short-run dynamics of the model. To carry out this exercise, we have to make some assumptions about the updating of point expectations for the strength of demand and for the general price level. For this purpose we use a simple error-correction model, just to get on with the job in a simple mechanical way. Our own conviction is that expectations are much more complex and may often depend on the theory of the economy that market participants have learned to accept—usually not from textbooks. This factor may be especially important when it comes to "the market's" response to policy actions. All that is a little too deep for us at this stage of the game.

Finally, chapter 7 steps back from the particular details and offers our reflections on what we think we have learned from this exercise about the right way to do macroeconomics.

2   Perfectly Flexible Wages

In the light of these considerations I am now of the opinion that the maintenance of a stable general level of money wages is, on balance of considerations, the most advisable policy for a closed system

—Keynes, *The General Theory*, p. 270

2.1   Where This Chapter Is Going

One classical technique of subversion is called "boring from within," and we try our hand at that first. Much contemporary macroeconomic theory leaves the impression that unemployment and recession are primarily the result of excessive rigidity of wages and prices, and perhaps the related immobility of labor. If only the artificial barriers to wage and price flexibility were removed—by the weakening of trade unions and the deregulation of industry and trade—the market mechanism would see to it that the labor market cleared. True unemployment would disappear and business cycle fluctuations would be minimal.

This sort of theory has practical consequences. Central bank governors and ministers of finance are given to saying in public, even while unemployment rates hover around 10 percent of the labor force, that they can do nothing about it and should do nothing about it. It is not their problem; the only proper policy is to chip away at obstacles to wage cuts and labor mobility. Wage flexibility will eventually do the rest. Presumably they have something more than competitive real depreciation in mind.

Another branch of macroeconomic theory holds that wages and prices are already adequately flexible, and that observed fluctuations in output and employment are not pathological at all. They are the economy's optimal response to unavoidable erratic shifts in tastes for goods and leisure and in the technology of production. The implication is that even if public

policy could do something to increase production and reduce unemployment, the temptation should be resisted. So far as we know, this view has not yet converted any central banks or ministries, but many up-to-date macroeconomists adhere to it.

We have no sympathy with either view. In this chapter we try to stay as close as we can to the assumptions and methods that characterize the schools of thought that we wish to subvert. We cannot possibly go all the way, however. For instance, we cannot adopt the "representative agent" approach that simply assumes the model economy to solve and carry out the infinite-time optimization problem of a single, immortal, foresighted worker-owner-consumer. That approach cannot seriously be said to *conclude* that economic fluctuations are nonpathological, because it has already assumed just that. Because we want to preserve at least the option of concluding that the economy may behave in a deplorable way, even if wages and prices are flexible, we have to choose some other line of argument.

We adopt instead the overlapping-generations formalism that is sometimes favored in current macro theory. This approach has its problems too. The main one is that, in actual fact, fluctuations in prices and output are of much shorter duration than a lifetime; but a model that could handle that fact would be hopelessly unwieldy. Nevertheless, we use the overlapping-generations model in the spirit of boring from within. In the same spirit we assume perfect foresight on the part of economic agents, and we allow that wages and prices are *perfectly flexible*, in the sense that they are at every instant at the values that equate supply and demand for everything in sight, including labor. Our immediate goal in this chapter is to show that even such an economy can easily follow unmistakably pathological paths. And not only that: these paths can be improved by the timely use of policies that are recognizably macroeconomic in character.

There is little or nothing that is specifically Keynesian in the story that we tell. Even so, we were led to it by one of those flashes of insight that litter *The General Theory*. Arguing in favor of some stability of nominal wages, Keynes observed that, because the nominal interest rate cannot be negative, severe deflation in a monetary economy must be accompanied by a high real interest rate that will necessarily discourage investment. Our unflinching devotion to wage and price flexibility guarantees that reduced investment will not lead to recession and unemployment; but it will certainly lead to lower productive capacity in the immediate future. We shall follow through this chain of events in grisly detail.

The idea of perfect foresight never occurred to Keynes, or if it did, must have been dismissed at once. That means, for instance, that he could point to the disastrous consequences of falling money wages to those with debts denominated in money. Under perfect foresight such difficulties can be ignored (although not forgotten). On the other hand, the assumption of perfect foresight allows us to study the relation between fully anticipated price changes and the asset choice between money and productive capital. If prices are expected to fall there is a prima facie case that there will be a shift to money from productive capital and so a declining marginal product of labor, with further consequences (under labor-market clearing) for the behavior of prices and the real economy.

Our procedure will be to examine what happens when an initial steady state is unexpectedly disturbed by an increase in the labor supply. The new steady state will have the same capital/labor ratio as the old (before labor was more plentiful). However, the economy will most naturally set out on a path in which, for a time at least, prices are falling and so less is invested. It will then begin by moving away from the steady state. But there are, as we shall show, many possibilities and some of them are not pleasant; others, though more desirable, require perfect foresight over the infinite future.

It is not easy or even in our power to give a complete global analysis of these paths. Because we are studying dynamics, it is not possible to confine ourselves to an intuitive, or indeed to a purely verbal, exposition. Some of the least attractive algebra has been confined to an appendix, but inevitably some algebra remains in the text. The reader who has understood the basic model may wish at first reading to skip the sections on equilibrium dynamics and take their conclusions on trust.

2.2   The Model

The model we use tells, for the most part, a conventional two-period overlapping-generations story. Households are either young or old. When they are young they supply one unit of labor inelastically. By virtue of wage flexibility they are always fully employed, earning a competitively determined wage income. Part of this income is spent on current consumption of the single good; the rest is saved. The amount saved, plus any earnings in the form of interest or profit, is the household's only source of purchasing power when it is old. We exclude bequests, so an old household spends all of its available purchasing power on consumption.

Our main departure from the general run of overlapping-generations models is that young households can allocate their saving to either or both of two assets. One of these assets is a claim to a pro rata share of the profits of firms during the household's old period. The other asset is money, which might as well be currency. Cash balances earn no nominal return; the real return, positive or negative, depends on the change in the nominal price of the good between the two periods of a household's life. *Since we assume perfect foresight, there is no difference in riskiness between the two assets.*

A problem thus arises about the demand for cash balances: it could easily vanish. We dodge this problem by imposing a partial cash-in-advance constraint on the old. A young household planning to spend a certain (nominal) amount on consumption in its old period knows that it will have to make a cash deposit at the very beginning of that period, equal to a fixed fraction of its expenditure. The rest, of course, will be just settled by the old household's share of the profits earned in that period. So young households hold some of their saving in the form of cash balances.

That will certainly guarantee a demand for money. Such an economy will be in one of two possible phases or, improbably, on the borderline between them. In one phase—which we describe as *liquidity constrained*—the nominal net return on investment in firms is positive. Since there is no uncertainty, households would prefer to invest more and hold smaller cash balances. But then they would not be able to provide the required deposit on intended consumption. They hold as much money as they have to. In the other phase—*portfolio indifference*—the real return on the two assets is the same and so households are indifferent between them. They are perfectly willing to hold the given stock of money, and they do so.

This is, no doubt, a fairly artificial construction. No doubt also the right way would involve an explicit transaction technology, but monetary theory is not our goal. Barring that, the partial cash-in-advance constraint is not much more artificial than the other standard devices that involve entering holdings of money in the utility function of households or the production function for firms. We adopt the simplest device, even though it imposes a cost: we have to keep track of the two phases or regimes of the model.

The firms in this model are wage takers and price takers. (We give up perfect competition in chapter 4.) A typical firm sells shares to young *households and uses the proceeds to buy goods that it will use as capital in* the next period. In the later period it employs (young) workers, produces

and sells output, pays out wages, and distributes its gross profits to its shareholders (who are by then old). The only noteworthy thing is that we assume away durable capital; investment goods are used up in the period in which they are used. Thus shares are liquidated between one period and the next. This artificiality is made necessary by the simple overlapping-generations structure. The basic point we want to make would, we hope, survive into a model with a more realistic time structure.

In the rest of this section, we lay out the basic elements of the model in detail. Then we say what it would mean for this economy to be in equilibrium. After that, we consider its steady-state equilibria. This means that we put all the exogenous variables—like the size of the population and the stock of money—equal to constants and we look for configurations of real and nominal variables (consumption, saving, investment, wage rate, price level, interest rate) that are capable of sustaining themselves at endogenously determined constant levels.

A household born at t will be said to be of generation t, or simply G^t. It lives two periods, and we shall write c_{tj}, $j = t, t + 1$ to represent its consumption of the single good in period t and period $t + 1$. All generations have the same utility function $u(c_{tt}, c_{tt+1})$, which is monotone and concave. Also each young household is endowed with one unit of labor, which we assume to be supplied inelastically. For the moment the number of households is inessential, so we assume that there is one household in each generation.

A young household is paid a real wage w_t at t. With only one good, the real wage is well defined. Part of the real wage will be consumed (c_{tt}) and part will be saved (s_t, in real terms). The budget constraint at t is

$$c_{tt} + s_t \leq w_t. \tag{2.2.1}$$

We shall suppose that saving can take two forms: (a) the household can lend to the productive sector at a gross real rate of return R_t, and (b) the household can hold real money balances (m_t). Let

$$x_t = \frac{p_{t+1}}{p_t},$$

where p_t is the money price of the good at t and p_{t+1} is already known at t. Then $1/x_t$ is the gross (real) rate of return on money balances. We can now write the household's budget constraint for $t + 1$ as

$$c_{tt+1} \leq R_t s_t - m_t \left[R_t - \frac{1}{x_t} \right]. \tag{2.2.2}$$