Alternative approaches to macroeconomic theory: a partial view

ROBERT M. SOLOW / Massachusetts Institute of Technology

Abstract. The paper begins by discussing and comparing informally the pictures of the aggregate economy that underlie three current approaches to theoretical macroeconomics: equilibrium theory, the 'post-Keynesian approach,' and fixed (or sticky) price models. Regarding the last as much the most promising, the paper then describes a handful of hypotheses about the labour market, more complementary than rival, that might explain the (downward) stickiness of wages. A recurrent theme is that the background assumption of optimizing agents is capable of yielding quite unclassical results if the utility function and the perceived constraints are slightly unconventional.

Différentes approches à la théorie macroéconomique: une vue partielle/partiale. L'auteur présente sans cérémonie une discussion et une comparaison des images de l'économie globale qui sous-tendent les approches en vogue en macroéconomie théorique: la théorie de l'équilibre, l'approche post-keynesienne, et les modèles à prix plus ou moins rigides. Comme l'auteur considère ce dernier groupe comme le plus prometteur, il s'attache à examiner quelques hypothèses bien davantage complémentaires que rivales pour expliquer la rigidité à la baisse des salaires dans le marché du travail. L'un des leitmotivs de cet article est que le postulat conventionnel d'agents économiques optimisateurs peut fort bien engendrer des résultats tout à fait non-classiques si la fonction d'utilité ou les contraintes perçues par les agents sont elles mêmes inhabituelles.

I wonder if anyone has noticed that the title of this lecture contains a pun. There are indeed several alternative approaches to macro-theory now current, and I want to say something about them. Perhaps I seem to have promised a partial view in the sense that to do the opposite would be to offer a 'complete' view. Well, I do mean that; but even more than that I mean a partial view as opposed to an 'impartial' one.

I want to begin by characterizing the main schools of thought in contem-

porary macroeconomics, mainly because I think it is very important to keep straight the preconceptions that underlie each one. Otherwise it is too easy to accept quite powerful conclusions without a clear grasp of the assumptions to which they are logically bound. The mass media and the business press do that all the time when they popularize economic ideas. Then, I want to concentrate on the particular approach that seems to me the most sensible and trustworthy and say a little about the state of theoretical research within that approach and about some analytical trails that seem to need investigation.

Macroeconomics is now, as it has always been, the subject of intense controversy. There must be several different reasons for this: for instance the policy stakes are large; the ultimate test of any theory, conformity with the facts, often turns out to be inconclusive and perhaps indecipherable; and it is even possible that the institutional environment changes, so that a theory adequate for one decade may fail in the next. You may think of still other reasons. In any case one can distinguish three or four general approaches or stances or theoretical positions that compete for attention and allegiance today.

I

According to one of these positions the economy is characteristically in or near equilibrium. The concept of macroeconomic equilibrium is itself problematical, but the school of thought I am now describing seems to mean something like full general equilibrium, in the more or less Walrasian sense that economic agents are carrying out the actions that they regard as optimal under the circumstances that they perceive to be prevailing.

Now if I stopped there, as I think many of the protagonists of this approach might do,[1] I would leave you with the impression that I had made a strong and specific statement. But that would be misleading, and in fact what I have specified so far is quite innocuous and almost without content. There is an old joke about a man who fell from the top of a fifty-storey building. As he passed the twentieth floor, someone shouted out a window: 'How are things going?' The reply was 'OK so far.' You could describe that poor fellow as being in full equilibrium, indeed as carrying out those actions that he regards as optimal under the circumstances that he perceives to be prevailing. The point is not that those circumstances are bound to change drastically; I could have made the same statement about someone who was merely chained permanently to a tree.

The real point is different. What makes the concept of general equilibrium (and its Pareto-optimality) so powerful is that the circumstances (constraints) subject to which economic agents optimize – apart from prices and budgets – have to do with technology, tastes, and basic social and legal institutions, all

---

1 Lucas (1977) is probably the best single source for anyone who wants to see the picture of the economy that underlies the technicalities with which it is usually surrounded.

of which are conceived to be fundamentally non-economic. The school of thought that wishes to regard macroeconomic events, including the business cycle, as equilibrium phenomena must have something very like that in mind, or else the statement is pretty empty. So it really is assumed, in this style of macroeconomics, that what you see when you look out the window is an economy in ordinary general equilibrium. (There is one important exception to this statement, to be noted in a minute.)

This view has obvious (and intended) affinities to nineteenth-century economic thought, Say's Law, and all that. Like that tradition, the new equilibrium school faces a basic problem: how can it account for the 'obvious' large-scale divergences from equilibrium that we think we see, especially in prolonged depressions? I am no student of the history of economic thought, I fear, but in my potted picture of nineteenth-century economics, the main argument used then was to attribute the bad times to disturbances focused on major sectors of the economy, usually caused by good or bad harvests, wars, or major technological revolutions. These exogenous events could be responsible for prolonged episodes of disequilibrium, or at least partial disequilibrium, during a reasonable period of adjustment, sometimes worsened and prolonged by financial panic. That escape route seems no longer to be routinely open to equilibrium theorists. Harvests are not so important, except very occasionally; wars have become less frequent; and technological events on the scale of the classics like railroad-building, electrification, and the coming of the automobile happen only very rarely. When such events do occur, they can serve the same analytical purpose they used to serve, as an occasion for disequilibrium phenomena. But nowadays something more is needed. We see too much of economic fluctuations to be fully satisfied with that story any more.

Contemporary equilibrium theorists solve the problem by depending primarily on expectational errors as the prime source of divergences from full equilibrium. Economic agents optimize subject to what they perceive to be their circumstances. But the relevant circumstances certainly include prices that will rule in the future, as well as other facts about the future that cannot now be known. Besides, there are plenty of present circumstances that may not be clearly visible to a participant in a complicated and opaque economy. (Thus some search theorists attribute ignorance about the prevailing real wage to the individual worker.) Agents have to form expectations about these unknown or imperfectly known circumstances. One necessary part of the definition of equilibrium in this kind of world is that those expectations be confirmed, at least in some reasonable statistical sense. The way is now open to explain major departures from equilibrium as mainly the result of unusually large and/or unusually prolonged expectational errors.

Here we come to an important point. It is a by-product of this view that macroeconomic policies can move the economy away from its Walrasian equilibrium only by creating and preserving errors of expectation. Notice that

this conclusion is essentially independent of the way in which agents actually form their expectations. The further assumption that expectations are formed 'rationally' in the sense of Muth and Lucas is merely a refinement. The hypothesis of rational expectations makes it even harder to pursue effective macro-policy, because it becomes well-nigh impossible for expectational errors to persist, unless the public authorities adopt the distasteful strategy of deliberately misleading private decision-makers, and even then they may fail. There is much misunderstanding on this point. At least in the United States the business press, which is all smiles when it hears of any idea that tends to downgrade public policy, has propagated the impression that the hypothesis of rational expectations provides the power in this demonstration of the impossibility of macro-policy. Actually, as I have tried to indicate, nearly all the work is done by the much more sweeping (and even less credible, if I may interject a little partiality) assumption that the economy is almost always in equilibrium.

I gather that German literary scholars and art historians have started a field that they call *Rezeptionsgeschichte*; they study the way in which a particular work or school has been received by various parts of the public as a source of insights into the work's social and artistic significance. Something like that should perhaps be done here. I have already hinted that I do not find this way of thinking about macroeconomics overwhelmingly persuasive. It is especially puzzling that a theory depending so much on ubiquitous misinformation about the state of the economy should have so much success without, so far as I know, making any attempt to verify its central hypotheses by direct observation. Is it true, for example, that the 'unemployed' systematically believe the prevailing real wage to be lower than employed workers do? Or do they have systematically lower expectations about current and future real interest rates, as some members of the school have proposed? You probably gather from my tone of voice that I find these propositions very hard to believe, and I am not sure why anyone should believe them in the absence of any evidence.

II

The second general approach to macro-theory I want to mention is one that may have been adopted by J.M. Keynes himself. At least it is repeatedly imputed to him by the school of thought that self-consciously describes itself as post-Keynesian. Since the group includes some of the members of Keynes's own circle in the Good Old Days, I am not about to tell them What Keynes Really Thought. I am uncultured enough, however, to think that it is a matter of mostly antiquarian interest. Keynes claimed, in *The General Theory*, to have proved the possibility of macroeconomic equilibrium with involuntary unemployment. I don't know, and hardly care, whether this was a considered opinion or a remark designed primarily to make the orthodox sit up and take notice. Posterity seems to have come to regard Keynes as a theorist

of persistent disequilibrium. From the standpoint of policy prescription it hardly matters whether the economy is at rest in a 'bad equilibrium' or merely takes a desperately long time to get back to a 'good equilibrium.' In either case public policy, and probably the same public policy, seems like a good idea. To justify laissez faire you have to go along with the root-and-branch equilibrium theorists.

I suspect the truth of the matter is that there are several strands in *The General Theory*; and Keynes need not have been conscious that they are only partially consistent, or even not consistent at all. One strand leads to an emphasis on the stickiness of the money wage. The economy can be in prolonged disequilibrium – and even appear to be at rest – because wages, and perhaps prices, do not adjust quickly to eliminate excess supplies and demands, especially excess supplies. This line of thought found its appropriate expression in the Hicks-Hansen IS-LM apparatus. With some extension and elaboration it became the standard textbook doctrine of the past thirty years. It is in fact the third general approach on my list, and I will come back to it later.

Many of the most self-conscious Keynesians or post-Keynesians are very hostile to that strand in his thought. I think they may fear it allows orthodoxy to save face too easily: 'Oh well, we knew all along that price rigidity can interfere with the normal equilibrating process of the market. If unemployed workers will not cut wages aggressively to seek jobs, it's their own fault. What else is new?' Of course that smacks of theology: to ignore an essential aspect of economic life is just as destructive of good economic analysis as not knowing about it at all, maybe worse.

But of course there *was* another strand to Keynes's thought. There are, for instance, passages in which he remarks that wage-cutting would do no good anyway because the price of goods would simply follow the wage down, without the reduction in the real wage that he then thought was necessary for any increase in employment. A later generation christened as 'the Keynes effect' the indirect stimulation of aggregate demand that went from general deflation to an increased real money-supply, to lower interest rates, to increased investment.

Modern post-Keynesians go further. They seem to say that Keynes's basic contribution to macro-theory was the rejection of the equilibrium concept altogether as hopelessly ahistorical. The textbook model, in its relentless addiction to neat model-building, has 'confused historical time and logical time' and ignored the essential importance of finance. Now some of this is almost incomprehensible to me, and the part I do understand strikes me as way out of proportion. It is certainly the case that much equilibrium theory and even disequilibrium theory of the IS-LM variety is negligent of historical detail and of the irreversibilities that history builds into the economy. But I cannot imagine any IS-LM-connoisseur who, confronted with that argument, would not reply: yes, that's fair comment. The comparative advantage of the

model-building method is certainly not in its ability to reflect the texture of historical detail. Macro-theory is unlikely to find a place for the Penn Central. In using a model to interpret a particular historical incident, one should certainly leave plenty of room for those nuggets of context that may condition the whole episode and may even determine important aspects of the outcome. It is harder for me to see how this point can be erected into a principle of doctrine.

My earlier allusion to the consequences of wage-cutting indicates yet another line of thought to be found in *The General Theory*. Keynes thought that wage flexibility might lead to instability, as just described, rather than to a quick return to equilibrium. He suggested that nominal wage rigidity might be a good thing precisely because it gives the price level an anchor. Some of the post-Keynesians seem to regard this emphasis on dynamics, on the instability or at least non-stability of an unmanaged capitalist economy, as the essence of the doctrine. But I cannot see how to associate this view with the rather violent attacks on the North American style of model-building that usually accompany it. One would think that instability or non-stability would be a likely candidate for model-building. The proper way to do macroeconomics can hardly be *all* historical context and *no* analytical structure. Unfortunately the school has provided no systematic description or example of what it conceives to be the right way to do macroeconomic theory. Thus far so-called post-Keynesianism seems to be more a state of mind than a theory.[2] I am reminded of Roy Campbell's doggerel:

> You praise the firm restraint with which they write.
> I'm with you there, of course.
> They've got the snaffle and the curb, all right,
> But where's the bloody horse?

III

The third general approach to macro-theory is, as I have said, the one that probably underlay the textbook IS-LM model all along without being fully articulated. In this view the economy is characteristically away from macroeconomic equilibrium. Adjustment processes are at best slow; for better or worse, nominal wages and prices are not very flexible in response to excess supplies. They may even be imperfectly flexible in the face of excess demands, but casual observation suggests that wages and prices rise more freely

2 The difficulty I experienced in trying to find a representative work in this tradition made me realize how little unity and structure it has. Davidson (1978) will convey some notion of the ideas, but is far from representative of the writings of Eichner, Minsky, Kregel, not to mention the English branch of the school. The curious reader might look at the first few issues of the *Journal of Post-Keynesian Economics* or at the series of articles published in *Challenge* during 1977–8.

than they fall. This presumed asymmetry has a long history in the tradition of Keynesian economics.

The contemporary version breaks with the tradition in one important way, however. In the standard intermediate textbook model it is only the nominal wage that is sticky and only the labour market that fails to clear. The price level for goods is supposed usually to equal marginal cost at a level of output and employment corresponding to aggregate demand at that same price level. In other familiar words, the price level for goods is determined at the intersection of aggregate supply and aggregate demand curves, conditional on the money wage. In that sense the market for goods clears. In most of the newer models the goods market is treated like the labour market. The price level is also sticky; at least excess supply, and perhaps excess demand, is likely to persist in the goods market. In such periods of generalized excess supply, price exceeds marginal cost.[3]

In this approach disequilibrium can persist for a long time. A further implication of the same line of thought is that the economy will very likely adapt slowly to external shocks, whether they originate on the demand side or on the supply side. Macroeconomic equilibrium is thus vulnerable to all sorts of disturbances. Here I can offer an olive branch to the post-Keynesian school. One of the ways in which history and historical time matter for macroeconomics is that they can affect the adjustment process very intimately. Financial and other commitments from the past, and expectations and apprehensions about the future, can easily interfere with the economy's ability to adapt to changes in aggregate demand and supply.

Notice that in this view too economic agents are optimizing under the circumstances that they perceive to be prevailing. The difference is that among the constraints that they perceive are *market constraints*, an inability to sell (or to buy) what they would like to sell (or buy) at going prices. This point is important only because it has been misunderstood. The difference between the equilibrium view and the disequilibrium view is not that in one theory agents are assumed to optimize and in the other they are not. The difference is in the constraints they are assumed to take into account. In the equilibrium view a worker who is not employed must have chosen not to work. One explanation offered is that he or she believes the real wage for his or her kind of work to be lower than it actually is and has therefore decided to supply less labour than before. Another possibility is that he or she is 'investing in search' or in other self-improvement because he or she believes the real interest rate to be lower than it actually is, or than it was, and is responding in the natural way. In the disequilibrium view the unemployed worker may know perfectly well what the real wage is in the normal range of jobs but simply

---

3 I need hardly give a reference for all this. My own try at sticky-price macroeconomics (although we did not quite realize it at the time) is in Solow and Stiglitz (1968). A much clearer version of what is essentially the same model is Malinvaud (1977). See also Muellbauer and Portes (1978).

cannot find one at that wage, probably because there isn't one. From this point of view macroeconomic disequilibrium is the consequence of 'wrong' prices and interest rates; and the economy has an inadequate capacity to generate the 'right' prices and interest rates.

Now this line of thought has some serious intellectual problems of its own. The basic question it must answer is why wages and prices are sticky. The persistence of disequilibrium prices and interest rates means that there are Pareto-improvements available that are not being exhausted. Somewhere there are simple or complicated bargains that could be struck from which all participants would gain. Why do those transactions fail to occur? The equilibrium school finds that possibility so hard to accept that they insist the unexhausted Pareto-improvements are not there, or at least not there in the eyes of those who could collect the rents from eliminating them. But of course the 'paradox' of the business cycle has always been exactly that – one could see simple rearrangements that would unambiguously improve the situation but that the economic system seemed incapable of generating by itself. That incapacity, if it is real, needs to be explained, not explained away.

I V

My plan is to devote the rest of this talk to a partial – and now I do mean incomplete – sketch of some of the mechanisms that have been proposed to explain the stickiness, especially the downward stickiness, of wages and, to a lesser extent, of prices. I want to emphasize two general points. First, there may be many such mechanisms, and they need not be mutually exclusive or competitive with one another. Secondly, much of the theorizing that has been done in this direction underscores the emptiness of talk about 'optimizing behaviour' that does not specify the objectives and constraints. Unconventional objectives and constraints can lead to unconventional results. Perhaps there is a third point that needs – surprisingly – to be emphasized. If wages and prices *are* sticky, that fact and its consequences do not go away merely because we have not yet settled on a universally satisfactory theory about *why* they are sticky.

It is sometimes forgotten that Keynes – or at least the Keynes who believed that the nominal wage was sticky downwards in the face of involuntary unemployment – himself proposed an explanation of this characteristic of the labour market (see Trevithick, 1976). His argument started from the idea that workers have profound convictions about appropriate *relative* wages among occupations or industries. It does not matter much whether these convictions reflect feelings of equity, traditional survivals, or a successful con-game. Decentralized wage bargains, both informal decisions and formal collective bargains, can only determine the absolute wage in one occupation or industry at a time. Suppose, then, that there is a touch of involuntary unemployment. Each individual worker resists wage-cutting, not because he or she would be

unwilling to accept a universal reduction in the nominal wage but because he or she sees only his or her own corner of the labour market and observes that to accept a wage cut is to accept a reduced relative wage by the normal standards of comparison (what modern labour economists call a 'wage contour'). Since each corner of the labour market looks the same, the attempt by each to resist disturbance of traditional relativities results in universal stickiness of the nominal wage level. (It is obvious that reactions to tight and slack labour markets are unlikely to be symmetrical.)

In a segmented labour market the only way a relative-wage-preserving reduction in real wages can be engineered is through a rise in the price level. That is why Keynes defined unemployment to be involuntary provided those who are experiencing it continue to supply labour at a slightly higher price level for goods – because then it can be said that the unemployed are not holding out for too high a real wage, though they may be unwilling to accept local nominal wage reductions for the reason just discussed.

That still seems to me a sensible argument; and it has the merit of fitting in with a respectable and long-standing body of ideas in labour economics.

This account has a general significance apart from its degree of particular truth or falsity. It is unconventional because relative wages figure as an object of preference; workers are supposed to want to protect traditional wage differentials. This is not any kind of money illusion; relative wages are real, not nominal, quantities. The unconventionality lies deeper. The implicit utility functions in this story are not exclusively individualistic, and this is not contradicted by the fact that most of us would be surprised if many people turned down a wage increase on the ground that it is not fair.

There is (at least) one major gap in this story. It explains why employed workers might resist wage reductions in recessions, and maybe – though not very convincingly – why the unemployed might be inhibited from undercutting the employed. It is not so good at explaining the behaviour of employers. It has always struck me as a more substantial mystery that employers do not more aggressively push for wage reductions in a buyers' market.

An extension of the argument can close the gap. Suppose employers feel, correctly, that anyone who actively tries to solicit or enforce wage-cutting acquires a reputation as a bad employer. Whether or not the attempt is successful, the first consequence might be deterioration of morale among the firm's workers, and this in turn could lead to reduced productivity, shoddy workmanship, or even mild sabotage. In the longer run, when the tone of the labour market improves, a bad employer might find the quality of his labour pool degenerating as the best workers in the relevant occupations and industry gravitate to good employers who do not take advantage of the competitive edge that unemployment presents to the aggressive employer.

It is obvious that this sort of thing would diminish employers' incentives to bring wage rates down in recessions. One can even formalize this insight by the unconventional device of including the wage as an argument in the firm's

production function to represent the morale, productivity, and quality effects in a summary way. The firm can then choose a package of wage-offer and volume of employment capable of producing the output that it can sell, or wants to sell, under current market conditions. The simplest question to ask of such a formal model is this: under what circumstances will a profit-maximizing (or merely cost-minimizing) employer choose a wage that is invariant to fluctuations in demand? In that case, it is in the firm's own interest to forgo wage-cutting when product demand falls and workers have to be laid off. The answer turns out to be this:[4] if, and only if, the wage enters the production function in an 'effort-augmenting' way, so that 'effective employment' is measured by the number of workers employed multiplied by an increasing function of the wage, then least cost is achieved always by offering the wage that minimizes the wage bill per unit of effort. In the effort-augmenting case the unit cost of effort depends only on the nature of the extended production function, and the optimal wage offer is thus independent of the state of demand.

We now have a consistent package. If workers resist wage-cutting in recessions for the reason that Keynes emphasized, their resistance could easily take the form that led me to this model of the firm's decision. Employers will not often try to do what workers would fight to keep them from doing.

I called attention earlier to a deeply unconventional aspect of the Keynesian relative-wage hypothesis: the non-individualistic character of the implicit utility function. Now I have done something equally unorthodox: introduced a price into the production function, as a carrier of 'psychological' or 'sociological' factors. This sort of unorthodoxy has a good side and a bad side. The bad side is that casual unorthodoxy has no particular merit. It cuts you off from the whole body of worked-out economic theory, and there is no assured reward, precisely to the extent that the unconventional assumptions are casual. The good side of course is that the unconventional assumptions may reflect something important about the real world; and to that extent they may allow us to model aspects of the real world that more conventional theory cannot account for. I have no settled opinion yet; but I have to admit that I find these bits of unorthodoxy incomparably more credible than the things that impeccably orthodox equilibrium theory asks me to believe about the world – to take the most egregious example, that what looks to the naked eye like involuntary unemployment is really a massive investment in leisure or search, misguided only because the searchers or resters believe the real rate of interest to be lower than it actually is. This sort of thing is an uncanny echo of a remark about the 1930s that I have heard quoted from Herbert Hoover's memoirs: 'Many persons left their jobs for the more profitable one of selling apples.'

I would be the first to admit that to allow yourself too free a hand with the

---

4 The proof is in Solow (1979).

stipulation of social conventions is to permit cheap 'proof' of almost anything. But that does not mean that there are not social conventions. George Akerlof (1978) has gone as far as anyone in taking this point of view seriously. He has shown that you can produce quite sophisticated models which admit a stable equilibrium in which a social convention is widely observed even though it would be to the pecuniary advantage of some participants to violate it. The stability is conferred by the following sort of mechanism: people are induced to obey the convention because so many people believe in it and take a dim view of violators, and people are induced to believe in the convention because there are so few violators. Of course such a model is likely also to admit another stable equilibrium in which the convention has fallen by the wayside: there are many violations because there are so few believers, and there are few believers because there are so many violators.

If the convention in question has something to do with a 'fair wage' – which is exactly the example Akerlof works out – then it is easy to see that shifts in demand can generate unemployment when the fair wage is 'too high.' Natur-ally one can say that this unemployment is 'voluntary,' and I suppose there is a sense in which that is so. But the label settles nothing. In the first place, the sense in which observing social conventions is a voluntary act is not exactly the sense in which choosing chicken fricassee rather than pork chops is a voluntary act; and it is a bit of a swindle to pretend that it is. Secondly, the important conclusion is that the full-employment no-convention equilibrium may be Pareto-superior to the unemployment equilibrium in which the con-vention prevails. The problem is that the choice between those two situations is never presented for a vote in just that way.

I may lose my licence for saying so; but I think that right now it would be a good idea to divert some research resources in this unorthodox direction, but carefully.

V

The next model of wage-stickiness I want to mention has had quite a bit of discussion since it first surfaced some five years ago under the general heading of 'implicit contract' theory.[5] It is rather more conventional than the ideas I have just been talking about, but it also has some aspects in common with them. The basic notion is that workers are strongly risk-averse, while em-ployers are much less so, by virtue either of self-selection or of better access to the capital market. If that is so, then formal or informal wage bargains are likely to evolve in such a way that employers assume most or all of the risks associated with demand fluctuations. In practice this means that workers buy

5 The basic texts are Baily (1974), Azariadis (1975), and Gordon (1974). The version described
here is the simplest teaching model to be abstracted from those papers and later ones.

insurance in the form of real-wage and employment stability, presumably in exchange for a lower average wage.

Which will it be: stability of wages or stability of employment? A theory that is supposed to be compatible with a sticky-wage account of unemployment had better be a theory of wage stability, not employment stability. I want to explore that question, at least informally, not primarily for its own sake but because some general strategic points will emerge from the analysis.

Imagine that a firm has a pool of essentially identical workers attached to it. Any year may be a good or a bad year, with known probabilities. In a good year all the workers in the pool will be employed – that defines the size of the pool. In a bad year only a known fraction of the pool can be employed. The lucky ones will be chosen by lot, and the rest are unemployed for that period. The firm is risk-neutral; it cares only about its expected wage-bill, since the revenue side is assumed to be beyond its control. Each worker has the same strictly concave utility function defined over the wage; the worker cares only about expected utility. For the worker there are three contingencies: employed in a good year, employed in a bad year, unemployed in a bad year. The explicit or implicit bargain has to specify the wage in a good year and the wage in a bad year. That tells us what we need to know about the employed workers in the pool. For the moment let us assume that an unemployed worker in a bad year achieves a utility level of zero.

In this setting any bargain with a higher wage in good years than in bad is dominated by a bargain in which the wage is the same in good and bad years. Start from a trial situation with the good-year-wage higher than the bad-year-wage. The risk-neutral firm would be just as happy with a bargain in which the good-year wage were a penny lower and the bad-year wage a bit more than a penny higher, so long as the expected wage-bill were unchanged, with account taken of the fact that employment is higher in good years than in bad. But those employment weights are exactly proportional to the weights attached to the two relevant contingencies by workers in their expected-utility calculation. In effect, the wage-change just described is like *undoing* fair game played between workers and management. Management, risk-neutral, is indifferent about playing a fair game; workers, risk-averse, would prefer to avoid a fair game. So a move towards wage-equalization accompanied by a tiny reduction in the expected wage would be advantageous to both.

So far so good: we have a rationalization for a wage insensitive to business-cycle fluctuations. Unfortunately, however, exactly similar reasoning will provide an exactly similar rationalization for *employment* insensitive to business-cycle fluctuations. Any contract implies an expected wage-bill for the firm; the firm would indifferently incur the same expected wage bill in any pattern the workers prefer. But left to themselves to divide up a given expected wage-sum, identical risk-averse workers would prefer to eliminate all uncertainty. They can do this, in effect, by accepting a constant wage low enough that the unemployed in bad years can be fully compensated without

exceeding the total wage-sum available on average. Alternatively, at that wage the firm would be prepared to engage in complete labour-hoarding in bad years. This is not so good: we set out to explain sticky wages in the face of unemployment and have ended up explaining that there cannot be any unemployment.

But there is a way out. I casually suggested above that we normalize at zero the worker's utility when unemployed, that is, when he or she receives no wage income. That simplification turns out to be far from innocent. Suppose that an unemployed worker in a bad year receives an unemployment compensation benefit from some source other than the firm. This makes a big technical difference. Such a worker will not choose to work at a wage less than the unemployment compensation benefit. If the worker's utility is plotted as a function of the wage, it appears as a horizontal line (at the level of utility guaranteed by the unemployment compensation benefit) and rejoins the utility function proper only at wages higher than that. The result is that the utility function is no longer concave throughout, and we have to reconsider the arguments made earlier.

The argument for wage-constancy survives. Obviously the firm cannot offer a bad-year-wage smaller than the unemployment compensation benefit and hope to recruit any workers. Once that condition is met, the whole manœuvre takes place in the concave part of the utility function. It involves a fair game between the firm and its employed workers, and it goes through as before. So we still have a theory of wage-insensitivity in the face of fluctuations in employment.[6]

The argument for employment-constancy is on a different footing. It involves spreading the expected wage bill among all the workers. In view of wage-constancy, we can think of the expected wage bill as a constant wage times the expected number of *employed* workers. If that sum is spread over the whole labour pool, it implies a new wage equal to the old one multiplied by the ratio of average employment to the size of the labour pool. If the bad years are very bad, or if they happen a large fraction of the time, this new wage could actually be lower than the unemployment compensation benefit. In that case the argument for employment constancy fails. It would be better for workers and employer together to tap the unemployment compensation scheme even though all workers must then face ex ante uncertainty about income.

We can define the critical replacement ratio as the ratio of unemployment compensation benefit to going wage rate that just causes the stable-employment solution to break down in favour of one with lower employment in bad years. It depends on the shape of the utility function, of course, and on the probability that a member of the labour pool will be employed. (The latter

6 The difficulty, and the way out, are implicit in Baily's article. Many teachers, like me, must have come upon these observations themselves; I have seen them attributed at an early stage to T. Sargent's lecture-notes at Minnesota. An excellent and clear discussion is to be found in Akerlof and Miyazaki (1978)

is simply calculated in terms of the proportion of bad years and the proportion of the labour pool employed in bad years.) It is a straightforward calculation, but I will not bore you with the numbers. Let me just say that, with a logarithmic utility function and an employment-probability of 0.9, the critical replacement ratio is about one-third. With higher employment rates the critical replacement ratio rises. Since actual replacement ratios in the United States may be higher than one-half, they are clearly in the range where the model just outlined may be practically relevant.

Since I intend to use this politically loaded idea in yet another way, I had better comment on it now. Even if this theory were true – remember I have only shown it to be logically consistent – it does not necessarily follow that unemployment compensation benefits are 'too high.' If the underlying demand fluctuations were true 'states of nature' – traceable to variations in the weather, to natural catastrophes, or to technological uncertainties for example – it would be natural to ask whether social welfare would be increased by eliminating this source of variations in employment. But to the extent that the underlying demand fluctuations are just fluctuations in demand, the sort of thing that could be partially offset by fiscal and monetary policy, it seems much more natural to try to stabilize employment by stabilizing aggregate demand. That route has the virtue of leaving the unemployment compensation system to perform its intended insurance function. There is an unfortunate tendency for economists who study uncertainty to label all contingencies as 'states of nature' and to be misled by the label into treating them all as irreducible facts of life.

VI

I have one last analytical idea to suggest.[7] It too builds on the two institutional facts that underlie the implicit-contract model: segmentation of the labour market into pools tied to single firms or industries, at least in the short run, and the existence of an unemployment compensation system. This time imagine that the labour pool is organized and has the power to set the wage. In response, the employer can choose the level of employment. The organized labour pool, perhaps a trade union, is thus a monopoly seller. But it has many members and has to take their wishes into account.

If everyone in the labour pool is alike, each has the same probability of being unemployed if the wage is set so high as to induce the employer to reduce employment below its short-run maximum. On these simple assumptions it is in the unanimous interest of the workers to choose the wage to maximize their aggregate net gain from employment, i.e. the product of the number of workers hired and the amount by which the utility of the going wage exceeds the utility achievable from the unemployment compensation benefit.

7 This one was stimulated by my reading of a commentary by Ian McDonald of Melbourne on a paper by J.H. Moore (1975).

The 'union' – though I emphasize it need not be a formal trade union – would unanimously rather set that wage than any higher or lower one.

The best wage in that particular sense can be described in a simple way. It is the wage that makes the employer's elasticity of demand for labour equal to the elasticity of the worker's 'excess utility' with respect to the wage. Again, it is easy to tabulate the monopoly wage as a multiple of the unemployment compensation benefit for any given utility function and elasticity of demand for labour. At this stage that would be an exercise in misplaced concreteness. Let me just say that if I take that venerable trial horse, the logarithmic utility function, and if I suppose that the perceived elasticity of demand for labour is near one, then the best wage is around three times the unemployment compensation benefit. It is an interesting coincidence that this figure is very close to the critical replacement ratio calculated earlier for the implicit-contract model. I do not think that coincidence is very significant. I chose the elasticity of demand for labour to be near one because I imagine that small changes in the wage in real life have very small effects on the aggregate wage bill. If I had to make a second approximation, I would suppose the perceived demand elasticity to be smaller than one, in which case the monopoly wage would be rather more than three times the unemployment compensation benefit.

This model has an interesting implication for wage-stickiness, which is why I have taken the trouble to sketch it. Imagine, as before, that history consists of an irregular alternation of good years and bad years. Suppose that a labour pool fitting this model has achieved its preferred wage during a sequence of good years, and now comes a string of bad years. Will the preferred wage fall, or will it rise?

The effect of the recession appears in the model as a shift in the firm's demand curve for labour. Now the only aspect of the demand curve that appears explicitly is its wage-elasticity. If the recession demand is perceptibly more elastic than the prosperity demand, the preferred wage will fall. If on the other hand the perceived elasticity of demand is not much different, i.e. if the demand curve has merely shifted down more or less isoelastically, the 'union' will have no wish to set a different wage.

Needless to say, I would not ask you to believe this model of the labour market. You should not be happy with the assumption that all workers are alike, since much of what you see actually going on in labour markets is incompatible with that assumption. Nor should you be happy with the notion that the 'union' sets the wage and the firm chooses the level of employment, since wage bargaining is evidently a more complicated game than that. You should not be happy with the partial-equilibrium character of all this, for obvious reasons. But my ambitions here and throughout have been quite modest.

If I had to summarize in a few sentences the case I have been trying to make, it would go like this:
– If we need good micro-foundations for macroeconomics, we are equally in

need of good macro-foundations for microeconomics. Cant phrases about optimizing behaviour lead nowhere without a reasonable specification of what is being maximized and what constraints are perceived.

– If what we see in the world suggests unorthodox specifications, they are worth trying. The methods of textbook economics are easily applicable to a wide variety of contexts.

– If wages and prices are much less than infinitely flexible, it is important to explore with care the macroeconomics of sticky prices and the microeconomic circumstances that might possibly account for them.

– It is much too early to tear up the IS-LM chapters in the textbooks of your possibly misspent youth.

REFERENCES

Akerlof, G.A. (1978) 'A theory of social custom of which unemployment may be one consequence.' Special Studies Paper No. 118, Board of Governors of the Federal Reserve System

Akerlof, G.A. and H. Miyazaki (1978) 'The implicit contract theory of unemployment meets the wage bill argument.' Stanford University Center for Research on Economic Growth, No. 217

Azariadis, C. (1975) 'Implicit contracts and underemployment equilibria.' *Journal of Political Economy* 83, 1183–202

Baily, M.N. (1974) 'Wages and employment under uncertain demand.' *Review of Economic Studies* 41, 37–50

Davidson, P. (1978) 'Why money matters: lessons from a half-century of monetary theory.' *Journal of Post-Keynesian Economics* 1, 46–70.

Gordon, D.F. (1974) 'A neo-classical theory of Keynesian unemployment.' *Economic Inquiry*, 12, 431–59

Lucas, R.E. (1977) 'Understanding business cycles.' In K. Brunner and A. Meltzer, eds, *Stabilization of the Domestic and International Economy*, Supplement to the *Journal of Monetary Economics*

Malinvaud, E. (1977) *The Theory of Unemployment Reconsidered* (Oxford: Blackwell)

Moore, J. (1975) 'Uncertainty and sticky downward but upward mobile wages.' *Economic Inquiry* 13, 559–64

Muellbauer, J. and R. Portes (1978) 'Macroeconomic models with quantity rationing.' *Economic Journal* 88, 788–821

Solow, R. (1979) 'Another possible source of wage stickiness.' *Journal of Macroeconomics* 1, 79–82

Solow, R. and J. Stiglitz (1968) 'Output, employment and wages in the short run.' *Quarterly Journal of Economics* 82, 537–60

Trevithick, J. (1976) 'Money wage inflexibility and the Keynesian labour supply function.' *Economic Journal* 86, 327–32