1 Should Keynesian Economics Dispense with the Phillips Curve?

Lawrence H. Summers

HARVARD UNIVERSITY and NBER

1 INTRODUCTION

My title should surprise the reader. That is certainly its intent. The presence of some sort of Phillips curve describing the process of sluggish price adjustment is often regarded as a defining characteristic of Keynesian models. Leading Keynesian macroeconomics textbooks all assign a central role to wage and price rigidity and to 'natural rate' Phillips curves describing the adjustment of wages and prices. On both sides of the Atlantic, discussions of macroeconomic policy assign a prime role to the concept of the non-accelerating inflation rate of unemployment (NAIRU) and debate its level. Keynesian economists in the USA point to the stability of the Phillips curve in recent years as decisive evidence upholding their position and refuting the views of new-classical economists. In Britain, Keynesians dolefully track the NAIRU as it continues its upward march into double digits, while remaining resolute in their devotion to the Keynesian paradigm.

In these remarks I shall try to make the case that models containing standard Phillips curves depicting the sluggish adjustment of prices are fatally flawed as depictions of Keynes's vision of the economy or of reality. More fundamentally, I will argue that the premise common to both Keynesian and classical macroeconomic models, that a downwards sloping aggregate demand schedule and an upwards sloping aggregate supply schedule intersect to determine uniquely and sharply the level of output and prices is untenable. Instead, I believe that models allowing for hysteresis effects – models in which equilibria are fragile and history-dependent – offer the best prospect for redeeming the promise of Keynesian macroeconomics.

I believe that accounting for hysteresis effects will require

3

revolutionary and not merely evolutionary changes in the way Keynesian (and classical) macroeconomists view the world. If this judgement is correct, there is nothing to be gained from my pretending that only minor modifications in textbook treatments are necessary, or that the points made here are already widely appreciated. On the other hand, if my judgements are wrong, there is little cost to my stating them in as vivid and bold a way as I can. These remarks, therefore highlight the flaws in conventional models and the promise of new approaches, but do not provide balance by stressing the scientific successes that have led sticky price models to be enshrined in textbooks.

Before proceeding further I want to stress one set of considerations that greatly reinforces my convictions. Even its friends must acknowledge that the textbook Keynesian view of aggregate supply possesses many of the attributes that Thomas Kuhn has ascribed to dying scientific paradigms. Two aspects are most obvious – its proponents maintain a wholly defensive posture and it is subject to regular and substantial amendment. I comment on these two points in turn.

Empirical work within successful scientific paradigms is outward-looking. It articulates the paradigm by resolving anomalies, or by demonstrating the paradigm's application to new phenomena. Keynesian empirical work on issues relating to wages and prices is usually inward- if not backward-looking. Many studies have been directed at defending the fundamental premise that wages and prices are rigid, at demonstrating the continued validity of an equation estimated several years earlier, or more frequently at finding out why an equation estimated several years earlier went off track. While words like menu costs, and overlapping contracts are often heard, little if any empirical work has demonstrated any connection between the extent of these problems and the pattern of cyclical fluctuations. Nor have these concepts been successfully related to phenomena other than cyclical fluctuations. It is difficult to think of any anomalies that Keynesian research in the 'nominal rigidities' tradition has resolved, or of any new phenomena that it has rendered comprehensible.

More striking evidence of the barrenness of the textbook Keynesian paradigm comes from scientific statesmen's overviews of the state of the science. It is difficult to find one in the Keynesian tradition that is constructive in charting past triumphs and pointing towards future challenges. Rather, prominent Keynesians' evaluations of the state of the field are destructive – being primarily comprised of attacks on the doctrines of the New Classical or monet-

arist schools. I think of the AEA Presidential addresses of James Tobin (1982), and Franco Modigliani (1977) or of Alan Blinder's recent evaluation of Keynes's contributions (1988). Similar overviews by economists of the New Classical school, notably Robert Lucas and Thomas Sargent (1981) have a much more constructive and confident tone. This does not guarantee that New Classical economists are right, and indeed they are not. But Keynesian economics should aspire to more than Churchill's defence of democracy as the best of bad alternatives.

Frequent *ad hoc* adjustments to account for embarrassing realities were a hallmark of Ptolemaic astronomy. It is sad but true that the half-life of various Keynesian views about the aggregate supply curve has been little more than a decade. In *The General Theory* (1947) Keynes proposed that the aggregate supply curve drawn in unemployment-price space was L-shaped. This view was falsified by the coincidence of inflation and less than full employment in the late 1940s and 1950s. By the early 1960s, a derivative was slipped and Keynes's view had given way to the Phillips curve vision of a stable downward-sloping relationship between unemployment and the rate of inflation. This view remained popular for not much more than a decade. The stagflation of the 1970s led to the slipping of another derivative and the widespread acceptance of the view that there existed a natural rate of unemployment, which was the only rate at which inflation could remain stable. On this 'expectations augmented' Phillips curve view, there is a trade-off not between current inflation and current unemployment but between permanent inflation and current unemployment.

A decade has now passed since the natural-rate hypothesis came to be widely accepted. In what follows, I will argue that it is again time for a major change in Keynesian conceptions. Secion 2 of this paper lays out the arguments on normative, logical and empirical grounds against the sticky price Phillips curve approach to economic fluctuations. Section 3 briefly makes it clear why New Classical theories are hopeless as descriptions of real world economic fluctuations, especially in Europe. It then demonstrates how hysteresis models can resolve the problems with sticky price formulations and at the same time account for the empirical observations that motivate the Keynesian approach to macroeconomics. Section 4 concludes the paper by discussing some policy implications of hysteresis models.

2   THE KEYNESIAN ORTHODOXY

The orthodox Keynesian view of economic fluctuations goes something like this. Real factors uniquely determine an equilibrium level of output and employment in an economy. However wages and prices are sticky because of long-term contracts and sluggish expectations and so can temporarily diverge from their equilibrium values. As a consequence of price stickiness, changes in aggregate demand, typified by an increase in the money stock, affect the quantity of output and employment in the short run. In the long run, purely nominal changes do not have real effects.

The proposition that changes in nominal magnitudes do not have real effects in the long run implies that the long run Phillips curve trade-off is vertical. Increases in the permanent anticipated rate of inflation do not affect the level of unemployment or output. In the short run, however, because wages and prices are inertial, there is a trade-off between inflation and unemployment represented by the short run Phillips curve. This view implies that when disinflationary policies are pursued, as they were at the beginning of the 1980s in both the USA and the UK, output falls and the rate of inflation slowly declines. Conversely, expansionary policies can temporarily but not permanently increase output. The extent and magnitude of nominal effects on output and employment will depend on the importance of the factors leading wages and prices to be rigid.

These orthodox views are flawed in three important respects. First, they are dispiriting and discouraging. If they were valid, there would be very little scope for macroeconomic policy to increase (or reduce) economic welfare. Second, they are logically deficient in failing to consider seriously the implications of wage and price rigidities for choices about quantities. Third, they are empirically refuted by the great persistence of unemployment and output fluctuations, and by the very substantial variability of output even in settings where wages and prices are highly flexible. Let me develop each of these points.

2.1   The Natural Rate Hypothesis is Dispiriting

Contemporary Keynesian views about the inflation–output trade-off are well captured by the slightly stylised Phillips curve relation:

$$P_t = b*GAP + P_{t-1} \tag{1}$$

where *GAP* represents the difference between output or employment and some natural or equilibrium value consistent with steady inflation. This equation holds that the rate of change of inflation depends on the output gap. Similar expressions may be found in leading Keynesian textbooks like those of Dornbusch and Fischer (1984), Gordon (1987), and Hall and Taylor (1988). Its striking implication for the efficacy of stabilisation policies may be seen by summing it over time and rearranging:

$$GAP = (P_T - P_0)/bT \tag{2}$$

Equation (2) holds that over any interval, the average value of *GAP* is proportional to the change from beginning to end in the rate of inflation. Over any period when the rate of inflation does not change, the average value of the output gap must equal zero. Macroeconomic policies which do not raise the long-run inflation rate cannot affect the average level of output and employment in the economy. Put differently, stabilisation policies can only mitigate recessions to the same extent that they also limit expansions. Perhaps more strikingly, bad macroeconomic policies cannot raise the average level of unemployment in an economy over any interval as long as the rate of inflation at the end of the interval is no less than the rate of inflation at the beginning.

This result is very general. It should be obvious that adding lagged values of *GAP* in order to capture persistence or rate of change effects, or allowing for a more elaborate lag structure on inflation would affect it. Some economists prefer to replace lagged inflation in (1) by expected inflation, or by a lag distribution of expected inflation as of various points in time. In these cases, it is easy to demonstrate that policy cannot affect the average level of the output gap over any interval sufficiently long that surprises average zero. The conclusion that policy cannot affect average output is likely to be a feature of any model that postulates a unique equilibrium level of output around which output fluctuates.

If the natural rate Phillips curve (1) is accepted as a description of reality, it seems as if Keynesians are fighting for the low ground in their running battles with classical economists. If increasing output in one period requires acceptance of an equal output reduction in another, it is hard to see why it matters very much whether policy can increase output for one period, as in the classical model, or for

several periods as in the Keynesian model. Remember that no one really knows how much calendar time corresponds to one of a model's periods.

The natural rate Phillips curve hypothesis also implies that the social gains from macroeconomic stabilisation policies are not very large. Even assuming that the marginal utility of income diminishes very rapidly, Robert Lucas (1981) has shown that the social gains from stabilising consumption around a fixed mean are likely to be very small. Thinking about other aspects of fluctuations, it is far from obvious that having 8 million workers unemployed for 1 year is worse than having 4 million workers unemployed for 2 years. Certainly, the burden of unemployment is likely to be borne more broadly in the former case than in the latter. Stabilisation policies have costs. If they really could do nothing to increase the average level of output, it is doubtful that they could make much contribution to social welfare.

While many Keynesian economists accept equation (1) at least as a first approximation, they shrink from its normative implication that policy cannot affect the average level of output over long periods of time. Instead they regularly write and speak as if it were possible to fill in troughs without shaving peaks or accepting ever-accelerating inflation. Certainly this was how Keynes saw the proper objective of macroeconomic policy. Since Keynes wrote, criticism of avoidable recessions has been far more common than criticism of inappropriate expansions. Indeed, while American Keynesians condemn the 'three Eisenhower recessions', and the recessions of 1975 and 1982, as the result of excessively contractionary policies, there is no peacetime period when any consensus regarded policy as too expansionary.

Keynesian, and for that matter monetarist, tracts invariably leave the impression that the Depression was avoidable, and that avoiding it would not have saddled current generations with a higher permanent rate of inflation. As I discuss below, I think these convictions are correct. But they cannot be defended within the context of the current mainstream Keynesian model.[1] Justifying activist policy will necessitate looking elsewhere.

2.2   The Logic of Wage and Price Rigidity

Keynesian discussions of the Phillips curve assign a pivotal role to the sluggish adjustment of wages and prices. The idea is that because of wage contracts, menu costs or slowly adapting expectations, wages and prices remain stuck for a time at disequilibrium levels. As a

consequence, the argument goes, employment and output are determined not by the intersection of demand and supply curves but by demand alone along the predetermined level of wages and prices. This story makes sense as a depiction of contractions caused by unexpected decreases in demand. In the short run, output falls as suppliers of labour are constrained by the sticky downwards wage.

But the mainstream Keynesian model has no coherent explanation at all for booms. Suppose an economy with rigid wages but flexible prices is initially in equilibrium and then the money supply unexpectedly increases.[2] Then the notional demand for labour will exceed the notional supply. Economists concerned with rationing in markets that do not clear worked out the solution to this problem long ago. One would expect to observe the level of employment generated by the supply curve and realised real wage – a level that must be below the equilibrium level. This of course is not what we observe. We observe that monetary expansions raise employment, contradicting what the standard analysis of markets where prices are rigid would predict. The mainstream Keynesian model passes over this difficulty by simply assuming that output is always determined by demand.

The rationale for this assumption is rarely provided. Sometimes vague reference is made to contracts entitling employers to force their workers to work overtime. This cannot be important. Apart from the difficulty of enforcing contracts that call for people to work against their will, there is the prominent fact that most cyclical employment gains take the form of more people working, not people working more hours. Another suggestion is that employment gains come because workers are somehow fooled and do not realise that real wages are lower. This suggestion has more of a classical than a Keynesian flavour. More importantly, observation suggests that booms cause few regrets. Somehow there are few complaints after cyclical expansions by people who wish that they had not been tricked into working. Perhaps the Keynesian position can be defended by some sort of argument suggesting that demand expansions reduce frictional unemployment – unemployment that is notably absent from the aggregate supply–demand diagrams found in the textbooks.

I am not aware of any convincing answer that those who ascribe cycles to nominal rigidities can give to the problem of explaining booms. This difficulty is really symptomatic of a general problem plaguing all attempts to explain fluctuations in terms of wage and price rigidities. Any serious thought about the rigidities leads one to

despair of using standard supply and demand curves along with disequilibrium prices to determine the level of output. Suppose, for example, that firms and workers agree to long-term contracts fixing nominal wages and that there is no possibility of renegotiating them while they are in force. Is there any reason to expect firms to operate along their labour demand curves? If agreements about employment can be negotiated, then one certainly would not expect them simply to allow firms to move regularly along their demand curve. Even if they cannot be negotiated, firms are likely to set employment recognising that their choices will affect subsequent wage bargains in a variety of ways.

Take another example. It is often noted that firms raise prices infrequently for fear of alienating customers and that this makes the price level more sticky. Grant for a moment that this is an important aspect of firm pricing policies. Does it make sense to suppose that their customers always operate along a Walrasian demand curve that makes no allowance for the alienating effects of price changes? Probably not.

These examples could be multiplied. It think it will usually be found that whatever logic explains wage or price rigidity also undercuts the use of standard supply and demand curves to determine quantities. In every other part of economics price rigidities lead to too little being bought or sold. Only in Keynesian macroeconomics do wage and price rigidities lead half the time to quantities in excess of their equilibrium level.

2.3 Nominal Rigidities and the Real World

In addition to the logical arguments, there are important empirical problems with the nominal rigidities model as an explanation of economic fluctuations. First, the nominal rigidities explanation is less plausible in the current era of secular inflation than it might have been at an earlier time. The original Phillips curve could be thought of as capturing *tâtonnement* effects – more demand pressure meant more rapidly rising prices. The pattern of high unemployment and high inflation observed during the 1970s made it clear that prices could rise rapidly even in the absence of abnormally strong demand conditions. This renders the whole idea of sluggish price adjustment to demand conditions less plausible, and suggests instead that inflation is better thought of as generating a sequence of equilibrium price levels.

Second, there is even at this late date no concrete empirical evidence linking the extent of nominal rigidities and the extent of cyclical fluctuations. A number of less than satisfactorily controlled comparisons point in the opposite direction. Countries with high inflation where wages and prices change extremely frequently, have especially volatile economies. The decline in cyclical variability in the USA after the Second World War coincided with the introduction of three-year union contracts and other institutions often thought to generate nominal rigidities. Employment is most variable among secondary workers, whose wages are not set by contract and are subject to wide variations. Across a sample of OECD countries Sushil Wadwhani and I (1988) recently found a positive association between wage flexibility and output variability.

Third, an essential feature of the mainstream view is the idea that economic fluctuations represent transitory movements away from equilibrium. This idea receives little empirical support. While evidence of very substantial persistence in output can be explained by arguing that technical progress today suggests greater growth in the future, it is much more difficult for the mainstream view to account for great persistence in unemployment. Yet Olivier Blanchard and I (1986) find that in recent years unemployment in a number of European countries has followed a process very close to a random walk. Over the past century, the first autocorrelation of unemployment for both the USA and the UK is over 0.9. I am aware of no other evidence suggesting a tendency for output or employment in any country to demonstrate a strong equilibrium-reverting tendency.

These empirical considerations as well as the logical difficulties with nominal rigidity theories and their disquieting normative implications lead me to conclude that macroeconomists should look elsewhere in trying to account for economic fluctuations. *A fortiori*, classical approaches are not the way to go. This point as well as the profound problems with the mainstream Keynesian model is driven home by recent British experience.

The most resolute and right-wing government in Britain since the Second World War has been in power for a decade, during which time it has launched major attacks on trade unions and on market imperfections generally. Its commitment to disinflation is not in doubt. Yet more man-years of unemployment have been suffered during the Thatcher years than were experienced over the whole of the period after the Second World War and prior to 1979. In Britain today, about 60 per cent of all unemployment is composed of persons

in the midst of spells lasting two years or more. Can anyone seriously maintain that this outcome is the result of intertemporal substitution, misperceptions, or efficient search? For that matter, what microeconomic factors could possibly have doubled the NAIRU since 1979? Certainly continuing wage inflation belies the idea that unemployment is currently held far above its equilibrium level by rigid nominal wages.

These phenomena suggest the need to look beyond the mainstream Keynesian and classical models. I take up this challenge in the next section.

3   FRAGILE EQUILIBRIA AND HYSTERESIS

One of Keynes's distinctive contributions to the study of economic fluctuations was his stress on the possibility that they were caused by 'animal spirits', fluctuations in businessmen's and financiers' expectations about future prosperity unrelated to real events. His suggestion was that exogenous conditions determined the level of output with a degree of arbitrariness. As a result, there was scope for expectations about future output to be self-fulfilling. More recently the question of whether anticipated purely nominal changes have real effects has been treated as a litmus test for determining whether a model is 'Keynesian'. And Keynesian economists have produced evidence demonstrating that anticipated nominal changes matter empirically. In a sense these two ideas are closely related. The relevance of animal spirits and of purely nominal changes for the determination of output is difficult to understand in an environment where exogenous conditions sharply determine equilibrium. Were multiple equilibria possible, and so the level of output is in some sense arbitrary, it would be much easier to understand how extrinsic variables like the money stock or animal spirits could affect the level of output.

3.1   The Daylight Savings Time Example

My point is well made by example. Consider the problem of the way in which we measure time. Any competent economics graduate student would have no difficulty proving the following proposition: *The numbers attached by convention to moments when the sun is at different levels in the sky have no effect on an economy's real allo-*

cations. More informally, the time standard is a purely nominal variable that should have no effect on real outcomes. The proof mimics the standard demonstration that doubling the stock of nominal money should have no effect on the level of real output. A unique equilibrium may not exist. But to each equilibrium under time standard A, there will correspond an equilibrium under time standard B, in which all real variables (like the time at which things open relative to the time when the sun is highest in the sky) take on the same values. In theorists' language, altering the units in which we measure time does not change the set of equilibrium allocations that the price system can support.

There is one important thing to understand about this proposition. It is false as an empirical statement. Every spring we see daylight savings time imposed, and then observe people getting an extra hour of sunlight after they get home from work. Is there anyone who believes that if daylight savings time did not exist, somehow all opening and closing times in the economy would simultaneously be altered? Over the years, thousands of pages of Congressional testimony have been taken arguing the merits of daylight savings time. Considerations relating to energy conservation, school buses, and the safety and convenience of farm workers have all played a prominent role in these debates. I doubt very much that it has ever been argued that the choice of a time standard is of no consequence, because it is a purely nominal variable.

The reasons why the choice of a time standard has real effects are instructive. Imagine a store in a shopping centre. Its owner may care about how much time s/he gets in the sun after work. But s/he cares far more about opening and closing the shop at the same time as the other stores. As a consequence, many real equilibria in which all firms open and close at the same time are possible. While all store-owners might prefer to open and close an hour earlier in the summer time, it would be very difficult to coordinate this outcome in the absence of a change in the time standard. Any one store that changed its hours of operation would regret it, even though all would be better off if all changed. In this setting, daylight savings time is a constructive innovation that yields benefits by shifting the economy from less to more favourable equilibria.

Note several features of this argument. First, the efficacy of daylight savings time depends on whether people care more about their relative time of opening than their absolute time of opening. In a community of hermits, there would be a unique equilibrium in which
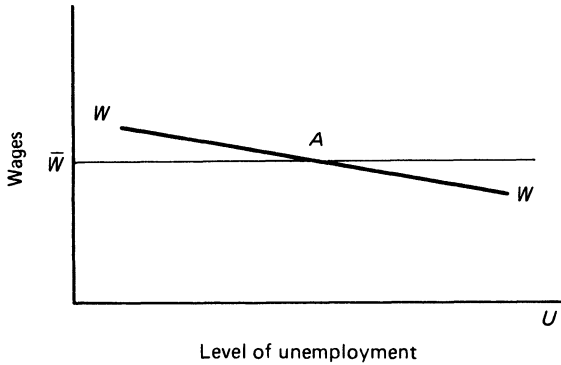
Figure 1.1 The determination of wages paid by a firm when other firms' wages are the dominant influence

everyone woke up at the time they most preferred, and daylight savings time would have no real effects. The impact of a purely nominal value – the time standard – is dependent on the fact that multiple real equilibria are possible. Second, sufficiently large changes in the time standard would cause firms to alter their stated opening and closing times. If the USA were put on Greenwich time, people would not find themselves leaving work before the sun had reached its highest point in the sky. Opening and closing times would adjust, and it is impossible to predict just what real equilibrium would be selected. Third, the efficacy of daylight savings time is related to the coordination problem arising because people care about relative rather than absolute time, not any nominal rigidity in opening and closing times. Stores open and close at different times on different days of the week and in different parts of the year. The costs of posting a sign with opening and closing times have nothing to do with the efficacy of daylight savings time.

3.2   Fragile Macroeconomic Equilibria

What does all this have to do with economic fluctuations generally or hysteresis specifically? I believe there is a very close analogy. In developing it, I will focus on the labour market and on the determination of wages and unemployment. Parallel arguments stressing product market considerations could probably be developed. Figures 1.1, 1.2 and 1.3 provide a plausible description of the way a given firm sets wages. Concerned about turnover, the problem of motivat-
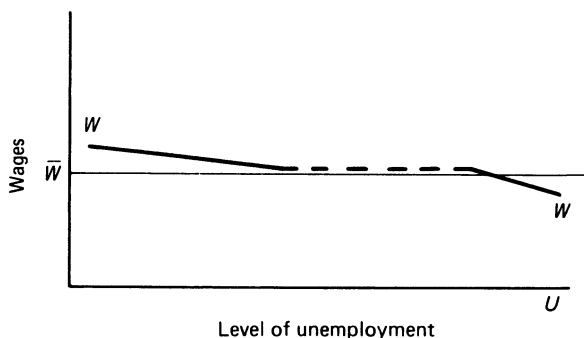
Figure 1.2 The determination of wages paid by a firm when all firms wish to pay similar wages
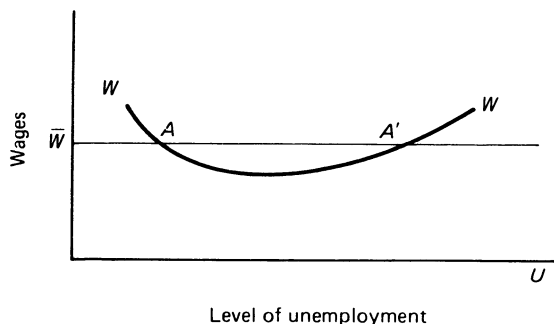

Figure 1.3 The determination of wages paid by a firm when extrinsic factors affect the relationship between wages and unemployment

---

ing its workers, the need to fill vacancies and so forth, it uses other firms' wages as a benchmark ($\overline{W}$) in setting its own wage. If unemployment is low, the firm is likely to decide to pay higher than prevailing wages. If unemployment is high it will decide to pay lower than prevailing wages. If all firms are symmetric, equilibrium is determined at point A, where each firm is happy to pay the average wage. As long as the typical firm would prefer to pay supra-normal wages in a hypothetical situation where unemployment is zero, unemployment will be observed in equilibrium.

The slope of the WW schedule in Figure 1.1 will depend on the relative importance firms attach to other firms' wages and unemployment in setting their own wages. If, as is plausible, other firms' wages

have the dominant influence, the *WW* schedule will be very flat, as in
the figure. While in this case there is a unique equilibrium unemploy-
ment rate, it will be extremely sensitive to anything that moves the
WW schedule. Moreover, it should be clear there will be a large
number of 'near equilibrium' unemployment rates where the wages
which firms actually pay differ only trivially from those they would
optimally pay. In such a setting, very small expectational errors are
likely to have large effects. In the limit, where firms always want to
pay the same wage as other firms over some range of unemployment
rates, a multiplicity of equilibria are possible, as in Figure 1.2, and so
unemployment is arbitrary. Just as in the case of the time standard,
when firms care a great deal about conformity, it will be possible for
extrinsic variables to have important real effects.

It is probably a mistake to distinguish too sharply the cases de-
picted in Figures 1.1 and 1.2. Situations with multiple equilibria, and
with very weakly determined but unique equilibria are not likely to
be observationally different. Think about hemlines. Their determi-
nation could be analysed using Figures 1.1–1.3 by replacing unem-
ployment with the average hemline and representing the preferred
relative hemline on the vertical axis. When hem lengths are long, the
average woman prefers to wear shorter than average skirts and when
they are short the average woman prefers a longer than average
hemline. Equilibrium is attained at that hemline where the average
woman desires to have the average hemline. Whether there are
literally multiple equilibria or not is unclear. It is obvious that
extrinsic variables matter and that instability is likely.

There are other possibilities as well, besides those depicted in
Figures 1.1 and 1.2. Suppose that increases in unemployment are
associated with more generous unemployment benefits as society
recognises that long-term unemployment may not be the fault of the
unemployed, or that they reduce the stigma associated with losing a
job, or that, as Olivier Blanchard and I argue (1986) increases in
unemployment lead unions to set a lower employment target as they
bargain over wages. In any of these cases, the *WW* schedule might
look like that depicted in Figure 1.3. There will be two equilibrium
unemployment rates. Extrinsic variables can move the economy
between them.

So far, my argument has been that if firms care a great deal about
their relative wage, an economy's equilibrium unemployment rate
may be very sensitive to small changes in conditions, or that an
economy may actually possess many equilibrium unemployment

rates. In such settings, it is natural to expect that extrinsic variables could have real effects in moving the economy between equilibria. It is also plausible that unemployment would be unstable because of the possibility of movements between multiple equilibria as well as movements in equilibrium values. Finally, small changes in its determinants may lead to very large movements in equilibrium values. Think of a change in union power that lifted slightly the *WW* schedule in Figures 1.1–1.3. Fragile equilibria of this type seem to me to correspond much more closely to what is observed in the world than the unique well-defined equilibrium displayed in textbook supply and demand or *IS–LM* diagrams.

While I have been brief, I am confident that it is possible to provide detailed microeconomic justifications for a variety of shapes of the *WW* locus in Figures 1.1–1.3. The real challenge at this point is to provide some indication of how an equilibrium is selected, and how changes in extrinsic variables alter real equilibria. This is where the idea of hysteresis comes in. Surely history is what determines the equilibrium an economy selects. It is history that conditions the wages that firms expect other firms to pay, and expect workers to expect other firms to pay. Once an economy is in an equilibrium state, it may be reasonable for agents to form expectations that it will remain there. These expectations may prove to be self-fulfilling.

Unfortunately, there is only one conclusion of which I am confident at this stage. There is no reason to expect market forces automatically to select the best of many equilibria. Saying more will require us to develop a deeper understanding of hysteresis effects than we have at present.

3.3   Resolving the Problems with Keynesian Models

Let me conclude this section by making clear how models of multiple equilibria and hysteresis can resolve the problems with the mainstream Keynesian model that I stressed in the first section of this paper. From the perspective of these models, the economy does not fluctuate around a unique equilibrium. Instead, it is capable of settling at many different equilibria, one of which is best. Policies that improve today's outcome need not compel the selection of an inferior equilibrium tomorrow. We do not talk about cycles around a fixed mean in people's health. Instead, we talk about them being healthy or getting sick. The multiple equilibrium approach to fluctuations takes a similar view of periods of high and low employment.

It should be obvious from the perspective advocated here that booms are no mystery. They simply represent the attainment of desirable equilibria. Since there is chronic involuntary unemployment in the model sketched here, no one regrets the extra work done in a boom. Since the perspective taken here is that the economy is always in equilibrium, no problems of describing quantity choices in the presence of disequilibrium prices need to be faced.

Models of multiple equilibria do not suffer the empirical defects of models emphasising nominal rigidities. They do not predict that systematic inflation should have systematic real effects, or that there should be a relationship between the costs of changing wages and prices and the extent of output variability. Most importantly, they do not carry any implication that economies should exhibit equilibrium reverting behaviour. As in any other social situation where individuals value conformity highly, the model sketched here suggests that outcomes should be both volatile and persistent.

Continuing inflation, rapid GNP growth, the fact that redundancy and short-time rates are low by historical standards, and the fact that overtime work is abnormally common by historical standards all suggest that it is not fruitful to think of the contemporary British economy as being far out of equilibrium. It is neither plausible nor bearable to think of this equilibrium as being unique. This too compels consideration of models with many equilibria.

4   CONCLUSIONS

Robert Lucas (1981) in his celebrated critique of econometric policy evaluation charged that 'it is only recently that the proposition that inflation and output are positively correlated and therefore more inflation is good has been elevated from obvious fallacy to the cornerstone of economic policy'. He was right. It is, however, equally fallacious to suppose, as do Lucas and mainstream Keynesians, that because steady inflation does not affect the average level of unemployment, nothing else can do so either. Purging this fallacious view will require us to eliminate the natural-rate Phillips curve from our models.

My daylight savings time example supports these conclusions. It is easy to imagine time policies that would be undone by the private sector and so have no real effects. Think of a policy of changing the clocks by three hours every week. Presumably under such a regime,

people would find ways of setting opening and closing times that do not lead to stores being open only before dawn. This proves only that there are limits to the equilibria that policy can impose on the economy, not that all policy is ineffective. And it shows that no simple equation can explain the process by which opening and closing hours are determined. But it hardly demonstrates the impossibility of government policies that help to solve the coordination problem involved in time setting.

The difficulty, of course, is in designing policies that can work. This will require improved hysteresis theories describing how history determines an economy's equilibrium. But a little bit can be said at a high level of generality. From the perspective of the view of cyclical fluctuations considered here, the problem of economic policy is very much like that of winning at poker. To succeed at poker one has to guess the endowments, intentions, and the guesses of others. A poker player can insure an average outcome by following a simple policy rule – not betting. This is what he will do if he assumes the universal rationality of expectations and decision-making that modern macroeconomists so blithely postulate. Many poker players do much better than breaking even. Likewise shrewd policy makers who, like successful poker players, make case-by-case judgements and do not shrink from bold actions will do much better than those who passively follow fixed rules. Improved theories of economic hysteresis will help them out.

Notes

1. One possible defence would stress non-linearities in the relation between GAP and changes in inflation. I am not aware of strong evidence demonstrating the existence of such non-linearities. While an asymmetry between upward and downward adjustment of prices is plausible, the idea of an asymmetry between upward and downward adjustment of the rate of inflation seems less compelling.
2. Similar arguments can be carried out assuming that prices are less than fully flexible and allowing for some flexibility in wages.

References

Blanchard, O. J. and Summers, L. H. (1986) 'Hysteresis and the European Unemployment Problem', *NBER Macroeconomics Annual*, vol. 1, pp. 15–78.

Blinder, A. S. (1988) 'The Challenge of High Unemployment', *American Economic Review*, vol. 78, pp. 1–15.

Dornbusch, R. and Fischer, S. (1984) *Macroeconomics*, 3rd edn (New York: McGraw-Hill).

Gordon, R. J. (1987) *Macroeconomics*, 4th edn (Boston: Little, Brown).

Hall, R. E. and Taylor, J. B. (1988) *Macroeconomics: Theory, Performance and Policy*, 2nd edn (New York: Norton).

Keynes, J. M. (1947) *The General Theory of Employment, Interest and Money* (London: Macmillan).

Lucas, R. E. and Sargent, T. J. (1981) (eds.) *Rational Expectations and Econometric Practice* (London: George Allen & Unwin).

Modigliani, F. (1977) 'The Monetarist Controversy or, Should We Forsake Stabilization Policies?', *American Economic Review*, vol. 67, pp. 1–19.

Summers, L. H. and Wadhwani, S. B. (1988) 'Some International Evidence on Labour Cost Flexibility and Output Variability', unpublished mimeo, Harvard University, Cambridge, Mass.

Tobin, J. (1982) 'Inflation and Unemployment', *American Economic Review*, vol. 62, pp. 1–18.