

The Famous Phillips Curve Article

Richard G. Lipsey

in Robert Leeson, ed.

A. W. H. Phillips: Collected Works in Contemporary Perspective

Cambridge University Press, 2000

Chapter 25, Phillips' essay on wages and unemployment was, for better or for worse both in its direct contributions and in the reactions that it provoked, one of the seminal articles of the last half of the twentieth century. Its theoretical origins lay in Phillips' work on stabilisation policy while its empirical origins lay in a casual comment by one of his colleagues on the LSE staff.

Phillips himself was one of the most remarkable persons I have ever met. He saw the economy as a dynamic system whose behaviour could not be understood using neoclassical static analysis – which, as someone who had been strongly influenced by Schumpeter in my student days, was a view that drew me to him. Although he had very little time for the comparative statics which was the stock in trade of conventional economists at the LSE in the 1950s and 1960s, he was always polite to us and never abrasive in any way. I believe he was proud of his varied career, his wartime accomplishments, his survival of a Japanese prisoner-of-war camp, his knowledge of languages, and his broad experiences of the world. But never did he show a suggestion of snobbery or condescension to we lesser mortals. He spoke with great authority and profoundly influenced many of us who came into close contact with him. As far as I knew, he had no strong political views. Certainly he never expressed opinions on the political matters that were constantly discussed in the LSE common room. His passion was for understanding the economy wherever that might lead him, and in that, he was in sympathy with those of us who were members of the LSE Staff Seminar on Methodology, Measurement and Testing in Economics (the M²T seminar) which he did not regularly attend.

Theoretical roots

When Phillips began his work on stabilisation policy, the Keynesian models of the IS-LM variety that were taught in British universities were typically closed with a reverse L-shaped, upward-ratcheting, kinked, aggregate supply curve relating the price level to national income.¹ Below full-employment income (to use the terminology of the time), the price level was given at its historically determined value and all fluctuations in aggregate demand caused fluctuations in real national income. When aggregate demand exceeded full-employment income, the price level rose until the excess aggregate demand was removed. If aggregate demand then fell, real income would fall along the AS curve which, below full-employment income, was always horizontal at the current price level. In other words, the behaviour of the economy was dichotomised so that, below full-employment national income, fluctuations in aggregate demand caused real income to vary, while, at full-employment income, increases in aggregate demand caused the price level to rise, ratcheting up the horizontal portion of the kinked AS curve.

Behind this AS curve, lay a micro base of product and factor pricing. Product prices were determined by price setting, oligopolistic firms who met variations in aggregate demand with variations in quantity at all levels of output below capacity. Voluminous empirical evidence suggested that marginal cost curves were flat below capacity and that mark-ups were fairly rigid, or at least not variable enough to cause major changes in the price level as output fluctuated below capacity. In the labour market, wages were assumed to be inflexible downwards, rising in the face of excess demand and staying constant in the face of excess supply. There were no strong theoretical underpinnings for this assumption about the behaviour of labour markets which was based mainly on the correct empirical observation that wages fell much more slowly in the face of excess supply than they rose in the face of excess demand.

At the time, this aggregate supply closure was understood to be unsatisfactory for at least two reasons (quite apart from the reason now commonly advanced, but seldom heard among my contemporaries, that it was a travesty of Keynes' own thinking on the subject). First, it was generally appreciated that the economy did not go sharply at one fixed level of national income from a state characterised by a stable price level and a variable real national income to a state characterised by a fixed real income and a variable price level. Second, although there was a large theoretical and empirical literature to buttress the assumptions about product pricing,² the assumption of fixed money wages seemed rather arbitrary.

Phillips found the kinked AS curve too restrictive for his models of stabilisation policy so he used a relation in which the rate of change of the price level was a non-linear function of the GNP gap, indicated by $Y - Y_f$. When there was zero excess aggregate demand, national income would be at its full-employment level and the price level would remain constant. As excess aggregate demand increased without limit, real national income would asymptotically approach its theoretical upper limit of a few percentage points above full-employment income while the rate of inflation would increase without limit. When excess aggregate supply developed, the price level would fall but, as excess supply increased and national income fell up to twenty or thirty percentage points below full-employment income, the rate of decrease of the price level would asymptotically approach a floor of some small negative value, say, -1 or -2 per cent.

This was a major improvement over the kinked AS curve. It allowed a continuous variation in the division of the effects of a change in aggregate demand between changing real national income and changing the price level. The higher the existing level of national income, the more the effects were on the price level and the less on real income.

This was all Phillips needed for his early models since his interest was in stabilisation policy. In his first essay on this subject (chapter 16), he was fairly optimistic about having an effective policy. In his second essay (chapter 17), however, he allowed for more complex and seemingly more realistic lags in response functions and became much more pessimistic about designing stabilisation policies that actually would stabilise. There is some debate as to how Phillips interpreted his models, but, as one of his junior colleagues, I took away from the second article a message of profound caution about the alleged beneficial effects of fine tuning. Long before I heard Milton Friedman on the subject, I was alerted by Phillips that some seemingly innocuous lags could render many 'stabilisation' policies counterproductive.

Empirical roots

I heard Phillips tell the following story of his development of the Phillips Curve many times. His colleague at the LSE, Professor Henry Phelps Brown, a great empirical economist in his time, suggested to Bill in the common room one Friday that he could illustrate his price-income curve by using historical data that had been gathered by Beveridge for unemployment and Phelps Brown and Hopkins for wage rates. It was accepted that the behaviour of the price level was largely driven by the behaviour of wages, both by cost-push advocates, of which Phelps Brown was one,

and by demand–pull advocates. The latter used the Keynesian model in which markups were constant and excess aggregate demand first raised wages and the resulting cost increases were passed on into prices by oligopolistic price-setting firms. Phillips got the data and went home over the weekend to plot wage changes against unemployment. At first, the plots all looked a mess but then it occurred to him to do what others who had looked at the same data had not thought to do: he treated each cycle separately and joined up the points in chronological order. What then leapt to the eye were regular loops around what appeared to be a fairly stable average relation between money wage changes and unemployment. This subsequently led him to fit his famous curve and to interpret it as a stable relation around which the actual data fluctuated cyclically.

There has been some debate as to whether or not Henry Phelps Brown was the source of Phillips' empirical excursion. Yet Phillips told the above story often in my hearing. I find it more plausible that Phelps Brown forgot making a casual comment than that Phillips forgot the source of a comment that had such momentous consequences for him. In any case, the identity of the person who gave Phillips the lead is unimportant, except as an historical curio. Certainly, whomever it was played no role in Phillips' subsequent development of his curve, which was done as a one-man operation. It is highly likely, however, that it was one of his LSE colleagues because this was not the sort of data that Phillips' own education or previous research would have been likely to have drawn to his attention.

The relation between the Phillips and the Lipsey pieces

Phillips wrote his article quite independently of our group of young Turks who, at the time, were worrying about the methodology of economics. We had formed the M²T seminar and were trying to apply Popperian methodology to a subject that was then dominated, in the UK at least, by the Austrian–Robbinsian–Euclidean methodology in which theories were judged by the intuitive plausibility of their assumptions. According to this methodology, if the assumptions were plausible, the logical deductions based on them must be correct and, if the facts disagreed with the theory's predictions, the facts must be wrong. In this view, facts were used to illustrate, not to test theories. I have discussed these issues and our group's reaction to them in more detail elsewhere³ but it should be clear that, as recent converts to the Popperian methodology, Phillips' Curve seemed an excellent test case. Indeed, if one reads

the first few pages of Lipsey (1960) and Lipsey and Steuer (1961) one sees the strong influence of Popperian methodology.

We first became aware of Phillips' empirical work when we read it in manuscript, after Phillips had seen galley proofs from *Economica*. It was too late for any of our M²T group to contribute to the original article, but Phillips did encourage the group in general, and me in particular, to study his relationship in more depth. Once we had all read Bill's article, we put it on the agenda for study in the M²T group. For a year, I struggled to understand the Phillips Curve. Finally, I came up with the results, both empirical and theoretical, that are reported in my 1960 article.

The two essays by Phillips (chapter 25) and Lipsey (1960) really formed the unit that made the curve famous. Phillips discovered the relation and did a rudimentary empirical job on it, while I applied standard statistical procedures and tested a number of *ad hoc* hypotheses that Phillips had formulated. I also tried to give the curve a micro-theoretical explanation but was hampered by not having a good model of expectations. As a result, I related money wages to the price level by a catch-up rather than an expectations variable. Since the catch-up variable did not perform well, I dropped it and was left with a simple, stable Phillips Curve.

Acceptance

I suspect that my more orthodox statistical treatment of the curve did quite a bit to still some of the many early criticisms, helping the curve to gain acceptance within the profession.⁴ There were, however, two more basic reasons for the wide degree of acceptance that the curve achieved. First, it was a distinct improvement on the old, kinked AS curve with its dichotomy between periods of stable prices and variable unemployment and stable unemployment and variable prices. Second, it allowed the profession to go beyond the *ad hoc* assumption of wage rigidity. The new curve related wage changes to aggregate demand and provided an empirical basis for a reaction of wages to excess supply (as proxied by the rate of unemployment) that was slow relative to their reaction to excess demand. The elimination of the aggregate supply curve, however, had momentous consequences in the long term.

Early reactions

Many early reactions to the curve were hostile, as shown by the list of articles published in the first two years after the original. British economists with left wing leanings were particularly critical. Richard Kahn was

reputed to be giving a series of lectures at Cambridge attacking the curve on theoretical and empirical grounds.

I was surprised by his hostility – although Corry (1995) explains why I should not have been. But, at the time, I was interested in academic economics and did not take much interest in British economic performance or policy. (My interest in those subjects was first aroused by my experience with the reactions to the Phillips Curve and then developed through the 1960s.)

At the time, I identified two reasons for the hostility. The first was a feeling that such a relation made no theoretical sense and was merely a statistical artifact. Since the equivalent relation in price–income space had been around for a long time, playing, for example, an important part in Bent Hansen's *Theory of Inflation* (1951), I could not understand this line of objection. Leaving expectations aside by concentrating on the nineteenth century when the price level was relatively stable, there seemed to be very good precedent for assuming that the rate at which prices would change in the face of disequilibrium would be an increasing function of the magnitude of the disequilibrium.⁵

The second source of opposition seemed to be the old conflict between judging a theory on the standard grounds of its logical consistency and its empirical relevance and judging it for its policy implications. Many left-leaning economists were profoundly worried that the curve would provide strong support for those who would raise unemployment in order to suppress inflation. At the time, there was a great debate on the relative social values to be placed on unemployment and inflation, with the left tending to put a high value on avoiding unemployment and the right a high value on avoiding inflation. While the reverse L-shaped AS curve gave no trade-off, the Phillips Curve did. According to that curve, inflation could be reduced by raising unemployment, and the Phillips Curve suggested that rates of unemployment around 3 per cent would be sufficient to hold inflation at 1 per cent. This unemployment figure was high by post-war UK standards, but might be acceptable to right-leaning politicians. Hence there was great hostility to the suggestion of an effective trade-off involving levels of unemployment that would be acceptable to some policy makers.⁶

Early cost–push interpretations

From its inception, the Phillips Curve was often misinterpreted as a cost–push phenomenon. Of the many misinterpretations, my favourite is Meiselman's (1968, 745):

If it were not for the fact that many of these Phillips Curve studies are taken seriously, I would also be quite amused by their implied monetary theory. Because the price level measures the value of money, by tying the price level to the level of money wages, these models essentially present a labour theory of the value of money – one which contains neither a supply of nor a demand for money!

The most influential misinterpretation, however, must have been Samuelson's. In the 1961 edition of his world-famous text book, he introduced the curve and asserted that it depicted a cost-push phenomenon.

In fact, when embedded in an IS-LM structure, the curve had quite orthodox implications. It determined the rate at which wage costs, and hence the price level, changed in the face of macroeconomic disequilibrium. Given a fixed money supply, the rate of change of the price level then determined the rate at which the LM curve shifted – left in the face of excess aggregate demand and inflation and right in the face of excess supply and deflation. The curve thus accommodated a Keynesian demand-pull version of inflation in which full-cost prices marked up increases in wage costs which were, in turn, caused by excess demand in the labour market which, in its turn, was generated by firms attempting to raise output in the face of rising aggregate demand. The curve was also orthodox in denying money illusion in equilibrium. Given a constant money supply and no growth, the IS-LM model closed by the Phillips Curve provided unique equilibrium levels of real income and price level.⁷

Later interpretations

Over the subsequent years, many people have sought to infer what Phillips really had in mind in his curve. See for example several of the essays in Cross (1995), including Desai's contrast between what he believes Phillips actually meant and how he thinks Phillips was interpreted by myself, Samuelson and Solow (1960), and others.

My belief that I was reflecting Phillips' own interpretation of his curve is based on the following considerations. First, I was in close contact with Phillips during the year that I was working on my article. If he had thought my interpretation was at variance with his, I would have known it. Indeed, when I tried to work with a market-clearing interpretation in which each point on the curve was generated by the intersection of relevant demand and supply curves, Phillips told me forcibly that he thought I was on the wrong track because his curve was a disequilibrium phenomenon. Second, he constantly referred to Bent Hansen's *Theory of Inflation* as an antecedent of his work, and, in particular, to Hansen's reaction function which related excess demand to changes in the price

level. Third, his abiding interest was in short-run stabilisation policy, as shown by his Ph.D. thesis (1953) and his first two major published articles (chapters 16 and 17). He saw his wage curve as an empirical underpinning of the curve he used to close his stabilisation models. The latter clearly was a disequilibrium reaction function relating excess aggregate demand or supply to the price level. Some authors, such as Desai (1995, 347), have accused others of judging Phillips' 1958 essay (chapter 25) without reading it. Be that as it may, those who interpret Phillips Curve on the basis of this article alone often fail to read the earlier two pieces on stabilisation policy, although all three articles need to be seen as a unit.

Some have interpreted his use of averages rather than the raw data as a desire to establish equilibrium, or at least long-term, relations rather than a dynamic reaction function. I asked Phillips many times why he had not used more conventional statistical methods for his original article. He had two answers. Early on, he said that, since the curve had a logarithmic form and since there are no logs of negative numbers, he was forced to use unconventional methods of first averaging the data into a few points and then fitting a curve to those points by eye. When he first said this to me, I accepted the challenge of finding a functional form that would take on the same shape as Phillips Curve but that could be treated by conventional statistical methods. I looked up my old statistics text, *Methods of Correlation Analysis* by Mordecai Ezekiel, and found several likely forms of which $Y = \alpha X^{-1} + bX^{-2}$ seemed the most promising. I tried fitting the curve

$$\dot{W} = a + bU^{-1} + cU^{-2}$$

to Phillips' data and found that it could track Phillips' own curve very well. I knew that the two unemployment variables would not seem significant separately because they must be closely related to each other, but my concerns were with finding a curve that could take on the same shape as Phillips' own curve, and with the resulting curve's *overall* fit.

Phillips' second response, which he used more often after my work was completed, was that he saw no half-way house between really crude eyeballing of data and what he regarded as a fully satisfactory econometric treatment, which would take him well beyond the conventional statistical methods which I had been taught.

Rejection of the naive version

Problems arose when the Phillips Curve was taken as showing a stable, long-term trade-off between unemployment and inflation. Ned Phelps led the theoretical criticism of the naive relation between unemployment and

changes in money wages, which, when used at face value, suggested acute money illusion in the long run.

Since Phillips was mainly concerned with short-term stabilisation policy, it is not clear if he condoned the use of his curve as a long-term trade-off. But many others certainly did.⁸ In any case, the theoretical objections to the long-term trade-off were strongly supported when stable empirical Phillips Curves came unstuck empirically in the 1970s in many countries.

Replacement by the expectations augmented curve

In a classic example of Popper's methodology, the failure of the naive curve resulted in the minimum amendment of the Keynesian macro model that was needed to accommodate the new observations. The IS-LM part of the model was compressed into an aggregate demand (AD) curve and two aggregate supply curves were reintroduced to complete the model. The short run aggregate supply (SRAS) curve was based on fixed input prices and was positively sloped in price level/real income space. The long-run aggregate supply (LRAS) curve was assumed (in simple cases at least) to be vertical at the unique point of macroeconomic equilibrium. To the best of my knowledge, Gordon (1978) was the first macro text to be organised around the natural rate hypothesis, while Lipsey and Steiner (1980) and Baumol and Blinder (1979) were the first two elementary text books to use aggregate demand–aggregate supply treatments.⁹ In Lipsey and Steiner's version, the expectations-augmented Phillips Curve

$$\Delta W = F(U) + \Delta P^e$$

provided the explanation of how fast the SRAS curve shifted in the presence of macroeconomic disequilibrium. This model quickly found its way into many other text books.

The Keynesian AD–AS model closed by an expectations-augmented Phillips Curve is capable of explaining most of what we see better than the competing new classical model. The reintroduced SRAS curve allowed for cost–push through such forces as OPEC's massive increases in oil prices, and thus explained stagflation. The expectations-augmented Phillips Curve explained the speed of wage and price adjustment to macro disequilibrium. It also predicted accelerating inflation if the monetary authorities tried to hold unemployment below its natural rate.

When grafted on to Keynesian expenditure flow models, the expectations-augmented Phillips Curve performed satisfactorily by providing for a short-term trade-off and for the neutrality of money not only in the levels, as did the original Phillips Curve, but also in the rates of change. A

stable level of unemployment below its natural rate could only be sustained by an ever-accelerating inflation. So the Keynesian model proved quite resilient. It was able to incorporate an expectations-augmented curve and, through it, to explain stagflation, which seemed utterly paradoxical when it broke out seriously in the mid 1970s, as well as the existence of an inflation–unemployment trade-off in the short term and its absence in the long term.¹⁰

A paradigm shift

Then, in a classic demonstration that, when broad paradigms are in conflict, economics is to a significant extent a political exercise rather than a science, the failure of the Phillips Curve became a reason for rejecting the whole Keynesian model and replacing it with a rational-expectations, new-classical, GE model which was unable to explain the stagflation that had undone the naive Phillips Curve.¹¹

The victory of the new classical model had little to do with any empirical refutation of the Keynesian model (as amended to accommodate the new facts) in preference to a better, new classical explanation. The new-classical critics argued that the Keynesian model had no microeconomic underpinnings while their model did. This interesting claim turned out to mean that Keynesian macro relations could not be formally aggregated from the empirically based Keynesian underpinnings of oligopolistic, price-setting firms. Instead, the empirically unreal micro underpinnings of perfectly competitive firms could be used to make the desired aggregation formally.

I have pointed out elsewhere (Lipsey 1981), that there was a large body of evidence for the micro underpinnings of the Keynesian model based on factual observations of pricing behaviour of oligopolistic firms. The theoretical problem was that no one knew, and does not know today, how to aggregate this type of behaviour, which is observed in the real world, into consistent macro relations.¹² So, presented with empirical evidence that was theoretically messy to handle, many in the profession chose the logically consistent, new classical model whose relations could be aggregated from the empirically refuted micro model of pricing and employment under perfect competition.¹³

Conclusion

The expectations-augmented closure of the Keynesian model used the natural rate hypothesis with the location of the vertical long-run aggregate supply curve indicating the level of real income that corresponded to

the natural rate. It is this natural rate hypothesis, embedded in both the new-classical models and those Keynesian models that are closed by the expectations-augmented Phillips Curve, that is the Achilles heel of both. As the essays in Cross (1995) amply demonstrate, and as many earlier critics such as Tobin (1980) and Lipsey (1978) predicted, the natural rate has defied attempts both to establish it in theoretical models that capture well-established features of the labour market and to find solid empirical support for its existence and stability.

Notes

- 1 As Bernard Corry (1995) has emphasised, this formulation bore little relation to Keynes' much more complex views on the behaviour of the general levels of prices and wages. It was, however, the formulation that had entered the formal models that built on Hicks's IS-LM interpretation of Keynes. Even in my time as a Ph.D. student at the LSE (1953–5), most of us got our Keynesian economics from text books and lectures and few of us opened the covers of the *General Theory*.
- 2 I have discussed this literature in several places; Lipsey (1981) provides the fullest discussion.
- 3 See the introductory essay for Lipsey (1997).
- 4 I reference and assess these early criticisms in a series of footnotes in Lipsey (1960).
- 5 In Lipsey (1974), I elaborated on the kind of micro process that I had in mind for driving the relation.
- 6 For myself, by the time I wrote the Phillips piece I strongly advocated putting up with significant amounts of inflation as the price of keeping unemployment low – thus illustrating that not everyone's economic analysis is determined by their value judgements.
- 7 I have developed these points in more detail in Lipsey (1978).
- 8 In Lipsey (1964) I drew a stable Phillips Curve and used policymakers' indifference curves to establish the optimal combination of unemployment and inflation.
- 9 Given publication lags, it is clear that these two treatments were written independently of each other.
- 10 For an empirically oriented treatment in this vein see Eckstein (1981); and for a more general discussion see Lipsey (1981).
- 11 In chapter 32 in this volume, Geoffrey Harcourt makes the same point, as does Corry (1995).
- 12 See Hahn (1995).
- 13 I have discussed the clash between the New Classical and Keynesian world in much more detail in Lipsey (2000).