Fifty Years of the Phillips Curve: A Dialog on What We Have Learned

Robert M. Solow and John B. Taylor
N. Gregory Mankiw, Moderator

Greg Mankiw: It's a great honor to be here to moderate this dialog between two of the most important macroeconomists of the past half century. It is very intimidating for me. But I've come to realize that all the hours I've spent watching TV talk shows has now finally produced some useful human capital for me. I've been trying to figure out who to emulate. I don't think I'm urbane enough to be Dick Cavett. I don't think I'm funny enough to be Johnny Carson. I'm going to aim for Jerry Springer. So, Bob and John, if you want to throw a chair at each other, just let me know and I'll duck.

The topic for today is the Phillips curve. I remember thinking a lot about the Phillips curve as a student. I remember thinking at the time that this was an incredibly important macroeconomic relationship. It almost defined macroeconomics and explained why classical economic principles didn't exactly apply in the short run.

The Phillips curve also made no sense to me as a student. I remember being very puzzled by it. I thought this was a great topic to do research on. I thought that if I start working on it, maybe, I'll figure out what's really behind this concept. I think I've done that for 20 years now. I still find it very frustrating and puzzling. I'm glad that we are all here today to finally settle the issue.

Let me start with a little history. Bob Solow was one of the people to bring the Phillips curve to the United States. With Paul Samuelson he wrote one of the great macroeconomic papers, a classic article titled "Analytical Aspects of Anti-Inflation Policy" and published in the *American Economic Review* in 1960, that applied the Phillips curve idea. So, I'm curious to first hear from Bob what he thought about the Phillips curve when he first saw it and what it was like then. Did you realize that we'd still be talking about it 50 years later?

Robert Solow: Before I get to answer that question directly, I would like to take a minute or two to say a word about Bill Phillips. This is a Phillips curve retrospective, after all. Am I the only person in the room who knew Phillips reasonably well, personally? One? Good; I'm not alone then.

Well, let me just remind everybody that Phillips was born and bred in New Zealand. He loved the place. When he spent the 1962–1963 academic year at the Massachusetts Institute of Technology, he was a neighbor of ours, as he lived close by in Concord, Massachusetts. Naturally I spent some time with him, including several evenings just poring over picture books and photographs of the New Zealand landscape. Phillips simply loved it; he was truly nostalgic about New Zealand. I can only confirm that it was very beautiful.

He was trained as an electrical engineer, and eventually wound up in the British armed forces as a technical officer with a unit of the Royal Air Force. He was taken prisoner by the Japanese in Indonesia, having escaped to that place during the Japanese conquest of Singapore. He then spent several years in a Japanese prison camp. He was not abused, as far as I know, but it cannot have been nice. According to all reports, Bill Phillips behaved extraordinarily well, was an important person to his co-prisoners. At great risk he jigged up radios out of odd pieces of wire and stolen parts so that they could know what was going on in the world. Then, after the war, he found his way to the London School of Economics (LSE), thought about studying sociology, thought better of it, and turned to economics. The rest, I hope, you know.

I learned about the Phillips curve the old-fashioned way. In those days, in 1958, you went to the library and looked at journals when they came out. There were only six or seven journals that an English speaker would be looking at; *Economica*, published at the LSE, was one of them. When I picked up the November 1958 issue to look at the table of contents, presumably in early 1959, here was this article on the relation between unemployment and the rate of change of money wages in the United Kingdom. I read it out of plain curiosity and I thought it was remarkable, because of this amazingly stable empirical relationship that he found. I took it out of the library and showed it to Paul Samuelson.

Now, it happens that we had already agreed to give a joint talk at the American Economic Association (AEA) meetings in December 1959—indeed, a talk about the analytical foundations of anti-inflation policy. You have to think about the setting. In 1958 and 1959 everyone was thinking about what was then called "creeping inflation." During the recessions of 1954 and 1958 the price level in the United States had continued to rise slowly at a rate that would be lost in the noise today, a mere 1 or 2 percent a year. The big debate was whether this inflation could be explained by "demand-pull" or "cost-push." Paul and I began to think about that argument and the various empirical tests that had been proposed as a way of distinguishing between them. We concluded that once you had a glimmer of general equilibrium thinking in your head, all those tests were wrong, and could not make the distinction that was wanted. We finally decided that a much more sensible idea to bring to bear was the Phillips curve. We thought that a more useful distinction was between movements along the Phillips curve and shifts of the Phillips curve. So that is what we did.

Greg Mankiw: Did you coin the term, "the Phillips curve," in that paper?

Robert Solow: I believe that Paul and I coined the phrase, "the Phillips curve." Obviously, this phrase doesn't appear in Bill Phillips's article—I think we invented the name. He invented the curve.

Greg Mankiw: There is one footnote on the New Zealand fact. In the first edition of my macroeconomics textbook, I referred to Phillips as a British economist because he was from the LSE, and I got quite a few annoyed letters from New Zealand. So that errata got quickly corrected in the second edition.

Now, John, you started grad school in the late 1960s. Can you give us some sense of what the academic thinking on the Phillips curve looked like at the time? That was also the time when the Phillips curve was still looking okay, but some people were starting to worry about it. Can you give us a sense of the intellectual climate back then?

John Taylor: I'd also like to preface my remarks, like Bob did, with some broader statements about Phillips. While I never met him personally, I felt like I knew him personally because I benefited from so much of his work outside of the Phillips curve early in my life. I graduated from college in 1968, the academic year that Milton Friedman gave his AEA presidential address in December 1967. So, in some sense I'm post-Friedman here, but my senior thesis in college used a macroeconomic model based on Phillips's work. The model had a Phillips curve in it, of course. But it also had an investment equation and a money demand equation. Why was I doing that? I was simulating different types of policy rules to see how they would work in this model. Actually, in this respect I was following very much the line of research that Phillips established long before based on his training in engineering. There were these various methods to control dynamic systems that came out of the engineering literature—thermostatic policies, actually. You had proportional control, derivative control, and integral control. Phillips had written some papers on these control policies, and what I wanted to do was look at those policy questions in a model that involved both cycles and economic growth, which was in a paper Phillips published in *Economica* in 1961.

That model is related to your question because, as I say, a Phillips curve is in the model. It is a macroeconomic model, a little dynamic model, with differential equations. But one of the equations was a Phillips curve in the sense that it related inflation to the output gap. Rather than the unemployment rate, he had the output gap, just like we frequently do in modern times. And he had potential output, natural output, moving around according to a simple production function. It was sophisticated, quite frankly. But the Phillips curve wasn't an *expectations-augmented* Phillips curve. It was the kind of Phillips curve that Bob just was referring to. What I remember so much is that in my work on this at the time, for some reason, I didn't think for a minute of trying to exploit the fact that there was a long-run trade-off. In other words, all the policy simulations were concerned with stabilizing GDP around the natural or potential rate. For some reason, I never exploited the long-run features, which I think to some extent reflected Phillips's own intuition here. He didn't think of the curve as something that should be exploited that way. Even though, if you took the algebra literally, without the expectations term, you'd argue it should be exploited by bringing GDP above potential and having a higher inflation rate. So, anyway, I didn't exploit it. I don't think he ever did either, but we could debate that.

With respect to the developments at the time, as I said, Milton's presidential address was in December 1967. I think that it changed thinking pretty rapidly. By the time I got to graduate school, that was the way we were thinking about things. It was an expectations-augmented curve. Of course you had Ned Phelps's important parallel work, and the Phelps volume of research following on very quickly. My recollection is that we all knew about the curve shifting over time and there was no long-run trade-off between inflation and unemployment. The change in academic thinking came very rapidly.

What was most striking to me in a policy context, I guess, is how this change didn't seem to influence policy very rapidly. Here is the reason I think there was a delay. People had been thinking in the 1960s about the long-run trade-off and that we could get more—let's say, less unemployment—with a higher inflation rate. Of course, by the late 1960s, early 1970s, mid-1970s, we were getting the high inflation rate but we weren't getting the low unemployment rate—so the Friedman-Phelps critique was validated. What happened with policy, however, is that actions were not taken instantly to bring the inflation rate down, even though the economy wasn't benefiting with lower unemployment. The new policy dilemma was that it was going to be costly to *reduce* the inflation rate. The expectations-augmented Phillips curve was put into the model, as Milton suggested, but with adaptive expectations—slowly adaptive expectations. Suddenly the policy trade-off about the level of inflation became a policy trade-off about the cost of disinflation. I think that "cost of disinflation" concept influenced economists all through the 1970s. It suddenly became fashionable to think that we can't get rid of inflation because it is going to be too costly. Those were the days that you had President Gerald Ford with the "Whip Inflation Now" buttons and his speech to a joint session of Congress. "How to Whip Inflation Now" was the theme at a White House conference, I think, in 1975. If you look through the transcript of that meeting, the only one at the time who was aggressive in terms of inflation reduction was Milton. So, it is a very interesting development, the interaction of the theory and the policy.

Greg Mankiw: John raised Milton Friedman's classic 1967 AEA presidential address. Bob, I was wondering what your reaction was to that paper? Did you think that address provided a fair assessment of what had come before? Or did you think it was a caricature of things that were said before? I went back some years ago and reread Samuelson and Solow and a lot of the caveats regarding the Phillips curve that we now talk about were in that 1960 article published in the *American Economic Review*. We just talked about expectations. In the paper you also talked about effects which today we call hysteresis.

See, a lot of the caveats were in that paper. I was wondering how you reacted to Milton Friedman's presidential address.

Robert Solow: It is true that our 1960 paper made all those allowances. We said explicitly that it is unlikely that one could successfully exploit the Phillips curve in the long run. We even mentioned the possibility that it was inflationary expectations that would shift the curve adversely if one tried. But I think we had something more general in mind: that the mere experience, however you process it, whether through expectations or the development of norms or behavior, would have that effect. So when I read Milton's address, that part didn't come as much of a surprise, though Milton dwelt on that point much more than we had thought to do so.

What did come as a surprise, and still comes as a surprise, though I didn't realize it right away, was that Milton had done something much more subtle and important, without explicitly saying so at all.

Let me go back to Phillips for a moment. Phillips's 1958 paper is purely empirical. All the theory in it is contained in the first two sentences of the paper, and what they say this: we are all used to the idea that excess demand in a market will cause the price to rise, and excess supply in a market will cause the price to fall. So why shouldn't the same idea apply to the labor market? That's all the theory there is. Evidently in Phillips's mind the arrow of causality clearly runs from the unemployment rate to the rate of inflation. He is thinking of the unemployment rate as a measure of the supply-demand balance in the labor market, an indicator of disequilibrium, and it would push the rate of wage inflation in the same way that you would expect excess supply or demand for peanut butter to push the price of peanut butter. (I don't remember any explicit discussion about why the rate of inflation rather than the price level.)

We, as sophisticated as we are, realize that all of this is slightly dangerous talk because those are both endogenous variables in whatever model we have in the back of our mind. So saying that A causes B rather than B causing A is tricky. On the other hand, in our less-sophisticated moments we know damn well what Phillips meant. He meant that the causality ran primarily from disequilibrium in the labor market to wages or wage inflation. What Milton did without ringing any bells to warn you, was simply to take it that the causality ran the other way, that it's the deviation of the rate of inflation from the expected rate of inflation that pushes the unemployment rate away from the "natural" rate. Phillips is about disequilibrium in the labor market. There is no question about that, just from the first couple of sentences from the paper that I paraphrased from memory. After Milton's address, everybody treated this as an equilibrium matter, looking in the reverse direction.

In this new story, the only way you can get the unemployment rate to depart from the "natural" rate is to create an inflation rate different from the expected rate. That kind of causality puzzled me, when I realized what had been done. Why do I say it "puzzled" me? What I mean is that I don't believe a word of it, and I find it strange that anyone does. When we come to talk about the so-called New Keynesian Phillips curve, we may find an intellectual advance, but that will arise in due course. Anyway, slowly, adaptively, I gradually realized that what Milton had done was to change the Phillips curve from what Phillips meant to this altogether contradictory kind of arrangement.

Greg Mankiw: When I teach the history of the Phillips curve, I start off with Phillips and then go on to the modern attempts to try to understand this trade-off between inflation and unemployment. One of the very important things that we go through in class is the Taylor model, which then evolves into the Calvo model and then it evolves into the New Keynesian Phillips curve. So, I was wondering if you would each say what you think about the current state of this work. I think of the New Keynesian Phillips curve as it comes down from the Taylor and Calvo versions as being the canonical model in the literature. This is where we are, perhaps with some bells and whistles attached, like adding lags of inflation, or indexation, or something. Do you think we are doing well, meaning heading in the right direction, or did we make a wrong turn somewhere along the way?

John Taylor: I think the research on price and wage adjustment has been very helpful, but actually to answer your question I would go back to the period well after Milton's critique, after the introduction of rational expectations, to talk about this phenomenon because Milton's model was all adaptive expectations, right? That adaptive expectations assumption is where the costly trade-off issue came from. After Milton's work you have the introduction of rational expectations by Bob Lucas, and of course Bob was using equilibrium models where the surprise change in inflation or in the price level generated deviations of unemployment from the natural rate or of GDP from trend.

Of course the implication of that rational expectation assumption, which Tom Sargent, Neil Wallace, and others emphasized, was that monetary policy was not going to be effective unless it tried to surprise. If it is at all predictable—which we thought it should be—then it wasn't going to have any impact. So, there was a real puzzle here, or a real vacuum or void, if you like, in the literature. It seemed to me that it had to be addressed and a lot of people started looking at it—Stanley Fischer, Phelps, Jo Anna Gray, myself. Originally this research focused on trying to build in a kind of price adjustment equation with rational expectations and sticky prices, and that is where I start to answer your question of whether we are on the right track. Those original "sticky price" models really lacked a lot of persistence. The models just basically jump back to equilibrium after the one period during which the prices were set or wages were set. The models didn't look that much different from what we call the new classical model. So, you had to do something else to get an empirically realistic model, and I'd like to explain the way I thought about it. This may be too long of a story, but I'll try to be quick.

We began to look at what was causing the persistence of inflation or price change. First we came up with the ideas that prices don't change every instant or every quarter or even every half-quarter, but they last for a while. They last for more than one period in our discrete time models, two periods rather than one, three periods rather than one.

But as soon as you do that, you begin to recognize that there is going to be some built-in inertia because not everybody is setting prices at the same time. You have this staggering nonsynchronization of price setting or wage setting, whatever it happens to be. So, I think that is the nonsynchronization phenomenon that people tried to think about: how do we model it? My own perspective is that it is a very hard problem to model. To me, it required a much different type of economic theory than we are used to teaching our students. So to get my hands around that, what I did was make a simple assumption that wages last for two periods, and half the workers change their wage every other period. Then there was the Calvo version of that, and now there are many others. But I think that approach generated—and this is why I think this literature is important—some important implications for policy.

Sometimes I think the value of a theory should be measured by what we get out of it. So, let me just mention what I think are five or six things you get out of that theory, which is continuing to evolve. First of all, you get some simple equations. When you are sitting in an office for a while, staring at the ceiling and trying to model some macroeconomic phenomenon, coming up with an equation is a big deal. You can estimate it and analyze it and compare different time periods, so the first thing is getting an equation. Second, you get some expectations of the *future* inflation mattering for the first time. In the Friedman-Phillips model, the augmented expectations were not expectations of the future. Similarly, in the "Lucas surprise" model, expectations are contemporaneous, not of the future. Now, for the first time, expectations of the future matter for prices today because of that staggering in the price setting. When they are setting their prices today, firms have to look to the future, and so for the first time you are getting expectations of the future mattering for inflation. I think that has a lot to do with the rationale for inflation targeting. Third, you get inertia. Inertia is built in because firms have to look back at past decisions, so you get a phenomenon that inertia lasts longer than the length of the longest contract. Fourth, you get a prediction that economic policy will affect the inertia. The more aggressive policy is, the less inertia or the lower inflation persistence will be. Fifth, you get a trade-off not between the level of inflation and the level of unemployment, but between the volatility of those two things, which has been useful. Besides the Calvo model, Jeff Fuhrer in his work with George Moore developed a model to explain persistence more fully. But the results and the things you take away from the Fuhrer-Moore model are very similar. More recently, people are looking at models where you have state-dependent pricing rather than just time-dependent pricing, and those models are more general and interesting, but what you are getting out of them is very similar. So, I think it is a very powerful way to think about the economy and, if you believe what I say, it has had a huge influence. So, I'll leave it at that and maybe Bob and I can talk a little more.

Greg Mankiw: You point out a couple of features of these classic models that some people view as bugs. For example, the fact that expectations of future inflation matter has been turned on its head. Larry Ball says well, gee, in these models, therefore a fully credible disinflation can cause an output boom rather than a recession. It is the same feature of the model that creates persistence problems. The model creates persistence in the price level, but you don't get persistence in the inflation rate, which is very different. I think when people look at the data, you have persistence in the inflation rate as well. That is why people go to alternative models like the Fuhrer-Moore model, and so on.

So, I think the literature is still a little inconclusive—I'm trying to figure out which of these features are features and which of them are bugs.

John Taylor: I will try to include all of those in my comment. It seems to me the main message is what I just said, but yes, there are little differences one needs to worry about. First is a notion that somehow you could announce a perfect path for the money supply (usually it is the money supply, not the interest rate) that slows down gradually and maybe by two years later reaches a new level. And you assume everybody believes that. Then, yes, you can get a very costless disinflation. I think that is not how you use a model like that. That would be naïve. So, to me that is not a disadvantage of the model. Second, if you try to fit models without some sort of exogenous persistence built into them, I think they fit pretty well, quite frankly. In recent years, Luca Guerrieri at the Federal Reserve Board has shown that they fit quite well. So, I think it's not a slam dunk to show that these specific things are a problem with the models.

Greg Mankiw: Bob, do you have any comments?

Robert Solow: I don't have much to add to that. It still bothers me, I guess, that the standard equation for a New Keynesian Phillips curve has the property that if inflation is and is expected to be constant, then output settles at the natural rate. This happens immediately in the pure rational expectations case, and then only white noise deviations can occur, or it happens gradually if there is some inertia built in. That does not strike me as plausible, not remotely.

What I like about the New Keynesian Phillips curve is that it gets rid of this elementary causality business. By the way, a careful person, like Jordi Galí for instance, never says that this is a Phillips curve; he says that it looks like a Phillips curve. That is exactly the point: it looks like a Phillips curve, but it isn't. It is something very different. The output gap, as I understand it, enters only as a proxy for economy-wide marginal cost, so it isn't intended to do the disequilibrium work Phillips intended. The main point, however, is that a careful person like Jordi Galí (or John Taylor) always embeds this thing in a model. There is something like an IS curve. There is a Taylor rule equation instead of an LM curve. And then there is this New Keynesian Phillips curve, if we continue to call it that. Then at least you can talk rigorously about causality, which you couldn't otherwise.

Whether or not the current version of the New Keynesian Phillips curve can recreate the right amount of persistence appears to be doubtful. I don't keep up with the literature as closely as I would if I would 30 years younger, but the most recent item I happened to see (in the Federal Reserve Bank of Richmond's *Quarterly Review*) came to the conclusion that the standard version of the model, even with some exogenously imposed backward-looking effects, couldn't account for the degree of persistence that is found in the data. This may be inherent in stories that want to build so much around rational expectations. My old codger feeling is that the technology has made it so easy to do calculations and then write papers that we euphorically forget how hard it is to get time series properties right. It is too easy to go on to the next paper.

John Taylor: Could I just comment on that three-equation model? It usually has a very simple equation which looks like a Phillips curve. That equation has a coefficient (beta) times the expectation of the future inflation rate on the right-hand side. It seems to me that those models can't fit data. They shouldn't be meant to fit data, they are just little models we use to teach students and discuss theorems. But you have to have more than that for policy on empirical work. In some sense, I think I am sympathizing with Bob here about fitting these models to the data. I once built a complex multicountry model using these staggered wage-setting models. I had to have a different structure for Japan and a different structure for Europe. Sometimes, the contracts had to last two years. Sometimes they have different distributions because you need to do that to fit the data. How could you possibly think that one little price adjustment equation could fit all different sorts of data? It is useful for expository purposes and debating, but we shouldn't hold it to a tough empirical standard. Although, probably people do to some degree.

Greg Mankiw: Bob raised the question of the natural rate hypothesis, and he said that he was skeptical that the long-run unemployment rate was invariant to whatever inflation rate the Fed decides to peg. Do you have a view as to what the current state of play is on that question?

John Taylor: Yes, I'm very much of the view that the natural rate of employment, or whatever we want to call it, is invariant to monetary policy. Let's put it this way. The idea of the classical dichotomy is very useful. I like that principle. Obviously, it's an approximation, but it seems to me that it is an important idea for describing long-term economic growth. You have the Solow model. You have trend productivity growth that is separate from monetary policy. You are going to have deviations from that, and monetary policy is important for minimizing the deviations. Similarly with the real interest rate; similarly with the natural rate of unemployment. Those are things that are from the real economy, and the more we can think about it that way, and teach and convince people, the better, because otherwise it is so confusing. When you read about monetary policy in the press, it always gets it mixed up. The more we can emphasize the invariance of the natural rate to monetary policy, the better. I think it is a good approximation. I don't have any problems with it.

Robert Solow: Everything John says is sensible; this is an experience I have had before. Let me try to explain what nags at me in all this. I'll say it first generally, and then come to something particular. We are left here with a theory whose two central concepts, the natural rate of unemployment or output and the expected rate of inflation, have three suspicious characteristics in common. They are not directly observable. They are not very well defined. And so far as we can tell, they move around too much for comfort—they are not stable.

I suspect this is an intrinsic difficulty. I have no wish to minimize the importance of, say, inflationary expectations. But we are faced with a real problem: here is a concept that seems in our minds to play an important role in macro behavior, and yet it's very difficult to deal with because it escapes observation and it even escapes clear definition. On the natural rate of unemployment, I think the behavior of the profession exhibits problems. In order to make sensible use of this kind of theory, you want the natural rate of unemployment to be a fairly stable quantity. It won't do its job if it jumps around violently from one year to the next. But that's what seems to happen. We, the profession, are driven to explaining events by inventing movements of the natural rate, which we have not observed and have not very well defined. The issue came up first in the passage of the big European economies from 2 percent unemployment, on average, to 8 or 9 percent unemployment, on average, within a few years. The only way to explain that within the standard model is to say that the natural rate of unemployment must have increased from something like 2 percent to something like 8 or 9 percent. The actual facts that could account for any such dynamics never seemed to me or to any critical person to be capable of explaining so big a change. So we are left with inventing changes in the natural rate of unemployment to explain the facts, and it is all done in our heads, not in any tested model. I regret to say that you often find this kind of reasoning: the inflation rate is increasing because the unemployment rate is below the natural rate. How do you know that the unemployment rate is below the natural rate? Because the inflation rate is increasing. I think we are all good enough logicians to realize that this is exactly equivalent to saying that the rate of inflation is increasing, and nothing more.

It seems to me that we ought to be thinking much more about the determinants of whatever you choose to call it. I hate to use the phrase "natural rate" but of course I do. It was a masterpiece of persuasive definition by Milton. Who could ever want an unnatural rate of unemployment? That was beautiful. But that issue is where attention ought to be directed. Similarly on the inflation side: why, for instance, are macroeconomists not talking about the shift from goods to services, and whether this changes price behavior very much, in a way that should be built into calculations? How exactly does international competition have an effect on movements of the domestic price level?

Greg Mankiw: One of the papers that influenced my thinking on this topic is by Jim Stock and Mark Watson. They wrote a paper about ten years ago trying to estimate how big the confidence intervals were around the NAIRU, as I think they called it, or the natural rate, which I'll take as a synonym. The answer was really big. At any point in time, you really don't know what it is. When Ben Bernanke goes to his staff and says, is the current unemployment rate higher or lower than the natural rate, they really need to say, "Who the hell knows?" But that does raise the question: is this concept useful for policymakers if one of the key parameters is immeasurable; that is, when we try to measure it, it has a large standard error? Is it useful merely as a theoretical construct to finish the Keynesian model, or is it practically useful when the Fed sits down and sets monetary policy if we can't estimate a key parameter in the equation?

John Taylor: Well, there is uncertainty, but it is not like it is useless. Bob is making a good point about how do you explain unemployment going from 2 to 7 percent in France, but that said, you have got to get out and do the empirical work. You get the macro data and you get the micro data. I don't think there is any substitute for that. But I don't want to give up on the concept of the Phillips curve. It seems to me that the concept is useful, but it's hard to estimate. That is just life. It has always been hard to estimate. It was hard at the beginning. It's a difficult policy problem, a difficult statistical problem. But why would you want to give up on the concept?

Greg Mankiw: No, I'm not giving up the concept, but it must be the case that the usefulness of the concept for making policy diminishes as the uncertainty about a key parameter increases.

John Taylor: Well, I'll give an example, just to respond. Some people, including my colleague Bob Hall, say that we should ignore the unemployment rate in the Taylor rule because it is so uncertain. Just look at inflation and ignore unemployment. I don't agree with that. I think that's a real mistake. Whether you want to use the GDP gap or other measures to measure capacity, you need to have some measure of capacity, some measure of where the overall economy is to do policy. So, I wouldn't go to the extreme of getting rid of all these concepts that try to measure where we are relative to normal.

Greg Mankiw: One of the things that we've been observing lately is tremendous increases in commodity prices. Some people may have noticed that. There has been a big literature that has tried to build in things like oil prices into the Phillips curve. There are other economists, and I think that Milton Friedman is one of them, who argue that is nonsense. The increase in the price of oil is an increase in the relative price of oil, and inflation is a monetary phenomenon. According to Friedman, relative prices have nothing to do with the inflation rate. But I think most of the people here would probably disagree with Milton's assessment. I don't know if you will agree with me on that. Do we have a good way to think about how relative price changes like the relative price of oil should fit into the overall inflation rate? What do you think is a good framework for thinking about that?

Robert Solow: Well, I don't know that we have an elegant framework for thinking about that at all, but I believe the point that Greg just made does correspond to the common sense of us all. One possible story about the current time is that the world is trying to bring about an increase in the relative price of oil. Maybe also the relative price of food, for all we know. It is one thing to be high and mighty and say, inflation is always and everywhere a monetary phenomenon, so therefore there is no reason why this event should affect the general price level at all. Other prices will fall far enough to offset the effect on the general price level of an increase in the price of oil, so there we are. We all know, whether we care to admit it or not, that there is some asymmetry about rising and falling prices, at least for many prices. It is going to be very difficult to generate the market-required rise in the price of oil without having the general price level rise too. Just to be paradoxical and annoy people, we could insist that this is not inflation if the price level rises just once to accommodate the change in the relative price of oil. A one-time rise in the price level is not inflation in the relevant sense. There is some truth to that view. If we had a clearer picture of the asymmetry of price behavior, if there is such asymmetry, then I think you could do as John described, and sit and stare at the ceiling and find an equation.

Greg Mankiw: So you think oil price increases and decreases are fundamentally very different? Your asymmetry story suggests that.

Robert Solow: This implies nothing theoretically special about oil, except that it is important. The general point is that important changes in relative prices cannot be brought about without changes in the price level because prices are not infinitely or equally flexible up and down.

Greg Mankiw: John, do you want to say anything about oil?

John Taylor: I think that, in principle, it is very hard to distinguish the price level shifts and the inflation changes, as Bob is mentioning. In fact, I think it's hard because there are these tendencies for prices to pass through to other prices, and that is the inflation dynamic that we are trying to study. We do have equations. I looked at the period from the late 1960s and compared different countries, and it seems to me that you got much less of a pass-through of oil prices to the general inflation rate in countries that had monetary policies that were focused on keeping the inflation rate low, and oil price shocks were much less costly in terms of output, too.

So, if you just simply let the oil prices just go through to the inflation rate, it can be quite dangerous in terms of leading to higher inflation. I like the theorem that showed that a more inflation-focused monetary policy leads to a smaller pass-through in oil prices. Plus it seems to hold empirically. So, that seems to me a very important message, and we shouldn't forget it.

Robert Solow: I agree with about two-thirds of what John just said, especially about the danger of just letting the oil price go. I agree that there will be a smaller effect on the price level the more restrictive monetary policy is. But I don't see why it should be less costly. I thought that Larry Ball had found, some time ago, in a cross-country comparison, that disinflation through monetary restriction was in fact quite costly.

John Taylor: Well, I'd say think of the cost in terms of how much better the business cycle has been in terms of smaller volatility of both output inflation, since central banks were more aggressive in terms of target inflation.

Greg Mankiw: Okay, I have one more question and then I want to open it up to the floor for questions for these two gentlemen. I started this discussion off by looking back, and now I want to look forward. What are the big unanswered questions that the next generation of macroeconomists should be focused on in this line of work?

Robert Solow: I've already shot my wad on that. I think that there are changes in the structure of the economy, like the shift from goods to services, like the growing importance of import competition, and, presumably, some decrease in the domestic degree of monopoly, and I can't believe that changes of that kind and magnitude do not affect price behavior. That would include the part of price behavior that we try to capture in one version or another of the Phillips curve. Those things, and they are only the obvious ones, ought to be studied.

John Taylor: I would like to see much more work along the lines of testing the theories on micro data, getting into the details like Pete Klenow and Mark Bils are beginning to do it. I think it is very productive. You are able to discriminate against the different models. In my view, we need more work in the wage area, where we have very little micro-testing of the wage equations. The more that can be done along those lines, the better, and I think that will help lead to better theories. My second idea is to think more about the macro equations that we are trying to explain by the micro theories. Right now, we've got lots of micro theories out there for just about every macro equation that you can think of, whether money demand or price adjustment. None of them work perfectly. Maybe we are searching for something that we are never going to find, like *the* microeconomic theory of price adjustment. You know, there are lots of different models out there that explain different aspects of price setting, and I don't know if that is necessarily a bad thing. I tried to think of an analogy from physics or physical science, and here is my analogy. It is probably a terrible one, but one of the most famous macro equations in physics, if you like, is Newton's Law of Gravity. The force of gravity depends on the product of the masses of the two bodies divided by the distance between them squared. It's an amazing equation and it fits very well.

Yet, there is no microeconomic foundational equivalent for that equation. There is no overarching theory in the sense that we call microeconomics to explain that theory. So, in this sense, maybe what we should be looking for is better macro equations. By the way, physicists have tried to find micro theories for that macro equation, done lots of work, and they just haven't found it.

The bottom line: maybe we should be looking for robust macro price equations that incorporate many different types of price adjustment at the micro level.

Greg Mankiw: Okay, thank you very much. Let me start off with taking some questions from the floor now on anything we've said or anything we didn't say.

Jeff Fuhrer: Jeff Fuhrer with the Federal Reserve Bank of Boston. So, we talked a little bit about large relative price changes, such as the one we are witnessing in energy. We also talked about expectations and the importance of expectations. We also talked about the idea that we have a partial understanding of how expectations are formed, perhaps. Could you comment on, for current circumstances, how you would think about the ways in which large relative price changes like energy might or might not feed into expectations, notwithstanding the strong presumption we have that they shouldn't feed into expectations. Then, what would a monetary policymaker in real time do about that?

Robert Solow: First of all, I have some confidence that a large rise in the price of energy generates expectations of a large rise in the price of energy. This question comes back to a point that John was making. The question is: how does that micro expectation feed into an expectation about general inflation, or, more precisely, the general price level? John is surely right that an atmosphere of fairly restrictive monetary policy will limit the extent to which the initial price increase would be generalized into further inflation, into a further general rise in the price level. But how exactly it works and how either the actions or the statements of the central bank can influence how the public translates relative price changes into expectations about the consumer price index or the PCE (personal consumption expenditures) deflator, that's not obvious. The question alerts us to further conceptual problems about expected inflation. There are various interest groups in the economy: bankers, investors, savers, lenders, borrowers, buyers and sellers and what not. There is no reason for them to react in the same way. How does one aggregate expectations?

John Taylor: I agree with that. Just another thought is that while we talk a lot about globalization, I think in the current environment you really have to think of monetary policy in the world and not just one central bank. After all, oil is a globally priced commodity, and the pass-through and the price-setting expectations take place not just in the United States, but all over the world.

So, I've been thinking more and more about international coordination. We used to think that you didn't have to worry too much about coordination across countries, but now I think that this is an example where the pass-through not only depends on what the Fed does but on what the European Central Bank does and what other countries do. So, I'd like to think of it more on those terms.

Robert Leeson: Phillips was a very sophisticated theoretician, and in 1954 on a theoretical level developed the Phillips curve relationship between prices and output. In that model, published in the *Economic Journal*, there is a very precise role for inflationary expectations, such that when inflationary expectations become embedded in the system, the system becomes unstable. Now in Milton Friedman's model, inflationary expectations are equilibrating, and in Phillips's model these elements are destabilizing. Now, Phillips did try to find some empirical data looking at prices and output and couldn't find it. But Henry Phelps Brown gave him some other data and that is the origin of the empirical Phillips curve.

Now, my question is that Phillips never really properly related his empirical work to his theoretical work, but then nobody else did either. So, how come the sophisticated theoretical work on which John was working in the 1960s got disconnected from the empirical work into the seminal role that Phillips played in developing the theory of inflation expectations. In fact, he gave Milton the formula, the adaptive inflation expectations formula. How come that got washed out of the system as well, and then it appeared that Milton was kind of developing this new concept that in fact he derived from Bill in the first place?

Robert Solow: I can't answer the second half of that, but on the first half, I'm sure John will correctly deny that in his mind the theory ever got separated from the facts, and I will confirm in advance the truth of what he is about to say.

Greg Mankiw: There is one fact that connects with that, which is that high inflation countries tend to have a lot of inflation instability, and the Friedman story doesn't naturally lead you to that fact. If you already expect 25 percent annual inflation, there is no particular reason that it couldn't be stable at that rate, but we don't seem to observe that in practice. I don't think that we've completely got a handle on that phenomenon.

Robert Solow: I think that's a good point.

Greg Mankiw: Thank you.

Peter Hooper: Peter Hooper from Deutsche Bank. I have an empirical question for the panelists. Given the essential importance, I think, of inflation expectations and applying the augmented model, how would you judge the best way to try to measure this? Would you ask a small group of professional forecasters? Would you ask a somewhat larger group of households in a survey or would you depend on what financial markets are seeing? To add to that, how much weight might you put on lagged inflation in your measure?

John Taylor: Well, I think you can't look at just one group. Actually, Bob made this point already. Different consumers are going to have different ways to process the expectation of inflation, so I think you have to look at all of them. If you are talking about looking at it for policy reasons, it seems to me that we are probably better off not trying to become too dependent on these expectations. Look for policies that are robust. Milton Friedman had this fixed money growth rule. Where did expectations appear in that?

So, it's not necessarily the case that you need to take these expectations into account that way when formulating policy. And we are going to hear more tomorrow about how miserable we are in forecasting inflation. How can you expect all of those other people out there that you are surveying to be that much better? So it seems to me that you'd have to have some humility in using both expectations of and forecasts of inflation to formulate policy and try, the best you can, to find ways to do it without using those measures.

Greg Mankiw: The empirical literature on the Phillips curve tends to find a pretty big role for lagged inflation. It's a bit of embarrassment from a theoretical standpoint since lagged inflation shouldn't necessarily be there once you have better proxies for expected inflation, but it is. So, people start putting ad hoc things into the theory. For people who are working in this literature, I think, it is no question an embarrassment.

Robert Solow: Once you start down that line, I think you come quickly to what I find embarrassing and difficult. Even if you were to divide the population into 30 different groups, each of whom hold some expectations about inflation, how do you weight them together? What weights do you apply, since you were not going to have a model with 30 or even 10 different expected rates of inflation? I think the better part of wisdom here may be not to pretend that that there is a precise concept, not to pretend that there is a numerical variable that you can fit into your model and call "the expected rate of inflation." Then you go on from there. If you do go on from there, there might easily be a role for lagged inflation as an indicator, as a proxy, as a last resort.

Michael Kiley: I'd like to follow up on Greg's comment that the usefulness of a theory must in some way be inversely related to our degree of uncertainty about the concepts that enter it. That idea must be at least somewhat important when we think about the natural rate of unemployment and how it links to inflation. When I think about writing down a model to help inform monetary policy, not just for forecasting purposes, I ask myself, well then what would it take for it not to be useful? To move to another model, you have to have the other model. So, if I weren't using that framework, what could I go to? I think we are at a point where this is the core theory that we have used to think about the links between unemployment and inflation and, like John said, we have uncertainty and we have to take that into account—but if there is no other option then we use what is there.

Greg Mankiw: Is that question for me?

Michael Kiley: Yes.

Greg Mankiw: I'm somewhat sympathetic to the argument that in the Taylor rule, for example, you give some weight to the unemployment gap, but the worse you are at measuring it, the lower the weight should be. What I learned from Mark Watson and Jim Stock is that we are very bad at measuring it. That is just a conjecture—I haven't written down the model to establish that notion. From a theoretical perspective, I have no problems with the idea of a Phillips curve. The question is, if the key parameters are very, very badly measured, then one has to question the utility for formulating policy. I think the Phillips curve becomes less useful. You end up setting policy by the seat of the pants.

V.V. Chari: V.V. Chari, University of Minnesota. Professor Solow suggested that there was a lot of asymmetry, and price increases in some sense being much more likely than price decreases. I think that sounds plausible and sounds right, but one of the valuable things I learned from Bils and Klenow's work is that just isn't so in the data. So, Mark and Pete looked at data at the level of individual commodities and at level of individual stores for a period in the United States when inflation was very moderate. What they found was that whenever prices changed, about half of the prices went up on average by about 9 percent, and half of the prices fell on average by about 8 percent. That came as a surprise to me. But I think it is worthwhile remembering that data.

My second point is that much has been made of the difficulties of getting inflation persistence. Recent work by Argia Sbordone, Tim Cogley, and Peter Ireland has argued that some of the findings of inadequate inflation persistence may come from misspecification of monetary policy. Conventional Taylor rule-type policies seem to not be using enough inflation, but formulations in which the monetary authority's inflation target is a random walk seem to do much better. That also seems to be much more consistent with data from the pricing of long- and short-term bonds. Those are two cautionary things worth keeping in mind.

Joseph Carson: Joe Carson from AllianceBernstein. I have a question about your relative price discussion. Now you are talking about relative prices moving because of oil, but before we used to talk about housing and stocks or whatever. Do you think the problem of relative prices versus absolute prices is that we do not have a broad price index to cover all of these price movements? Do we focus on too narrow a price measure?

Robert Solow: I don't know.

John Taylor: No.

Greg Mankiw: Well, this does raise a question of which price index we should be focusing on. For example, there is wages versus prices. Samuelson and Solow, if I remember, was prices; Phillips was wages. *The Economist* magazine suggests that we stick equity prices in the price index. I don't think a lot of people here propose that, but there are a lot of prices in the economy, so which prices should we be focusing on when we think about the Phillips curve?

Robert Solow: I don't think that that's an abstract question nor a theoretical question. I think that you are looking for robust relationships, and I would settle for any price index that did a good job.

Greg Mankiw: So, we should use a price index that maximizes the fit of the Phillips curve?

Robert Solow: No –

Greg Mankiw: That's an interesting project. I'm not even sure what that would look like.

Robert Solow: In all such problems, there is an interplay between common sense and goodness of fit. One wants both of those characteristics. If the price of thumb tacks happens to work best, surely you don't think I would go for it. (But I would sure ask myself why it works best.)

Greg Mankiw: The thumb tacks standard.

Zvi Bodie: Zvi Bodie, Boston University. I am curious as to why, particularly Bob, why you wouldn't consider the spread between the conventional Treasury rate and TIPS (Treasury Inflation-Protected Securities) as a reasonable proxy for—

Robert Solow: I thought of that, and as I turned things over in my mind, that might have been the best thing I could come up with. But it's still not clear that the ordinary TIPS spread would tell you anything about what is in the mind of those who are engaged in wage determination or in the setting of some other prices. The reason why I did not fix on that is because it is a fairly narrowly restricted market. We know from surveys that there is a wide spread of expectations if you ask individuals how much higher they expect the price level to be a year from now. Somehow that has to be weighted and averaged. Inside that spread, there are groups that may behave differently, and it's not clear to me that the people whose purchases and sales, or intentions to buy and sell, actually determine the price level also determine the TIPS spread, even indirectly.

Greg Mankiw: There is one other problem, I think, which is there may be differences in risk and liquidity premiums associated with these different instruments. I believe that the Cleveland Fed has some corrected series where they can say that the spread that you are looking at has been corrected for fluctuations over time in some liquidity premium as they estimated from some model. Still, if you are thinking you have to do that, your answer is only as good as the model of liquidity premium. So, if you think that is an important correction that needs to be done, that raises the possibility of measurement error there.

Barry Bluestone: Barry Bluestone, Northeastern University. I know tomorrow morning we are going to have a session on this, but I'm wondering to what extent, particularly when we look at the shifts in the Phillips curve, it is possible to model some of the institutional factors that, Bob, you had started to talk about. I am thinking of changes in the strength of the trade union movement, union density, trade policy, environmental policy, industry deregulation, and so on. I wonder if you could just comment on how we might even bring that into a model and what the impact might be of this inclusion?

Robert Solow: I agree with the general point and in fact, 48 years ago, Samuelson and I mentioned those things as important things to think about. It's damned hard to do that sensibly. When you do try, it devolves too soon into storytelling, and that sort of turns you off. It may be, although I'm not optimistic, that international differences could play a role there. The small year-to-year changes in union density, for instance, are never going to show up as a variable with a significant coefficient in regressions of that kind. It is conceivable that you might be able to interpret country differences, where institutional differences are much larger as a clue. But then there are so many other international differences that could mess up the relationship, if any.

Greg Mankiw: Allan Meltzer.

Allan Meltzer: I think this has been a very useful discussion because it's brought up a lot of the major problems. If one is teaching macroeconomics, especially to undergraduates, then I would think you would surely want to use the Phillips curve to illustrate what was going on and why, and how they might think about macroeconomic issues. When it comes to the policy issues, I think it's important—it's not just my view, but the view of the two most successful chairmen of the Federal Reserve, Paul Volcker and Alan Greenspan—that the Phillips curve was not very useful. In fact, it was not useful at all. Both of them didn't use it. Their quotations on this score will be in my forthcoming history of the Federal Reserve.

Why is that? I think that among the things that are part of the problem is the one that Bob Solow talked about. There isn't enough made of the distinction between large relative price changes and changes in the rate of price change that will persist through time given the policies that we have. Another is that we don't pay much attention, to the extent to which people inform you of their expectations, trying to decide whether what they are observing is a temporary change or a persistent change—that is, whether it's one that is going to go on for a while. Surely in the history, one finds that 1973 and 1979 oil price changes largely ended by coming down; that is, the price of oil eventually came down. Much of the cost of those changes was born by wages; real wages fell or rose very slowly during that period.

The present change in the price of energy appears to be much more of a persistent change. Wouldn't you want to expect that this change is going to have a different effect than what the history would tell us about the 1970s? I think it should. I don't see where the models that we use do a very good job of trying to make this distinction between changes in relative prices and changes in the absolute price level and between persistent and relative changes. I think that affects not only what happens to prices, but it also affects what happens to output. For example, in the change that we observed in output in the 1990s, how long and how persistent would the change in productivity be that came at that point? Most of the models of the Phillips curve that we have do not have very explicit attention to details of that kind that would be important. Finally, I would say that much of the problem in translating the theory to the policy is it's not clear to me, it's never been clear to me, that the theory applies quarterly, but much of the policy is aimed at doing things quarterly, and that is a source, I believe, of errors.

I think economists probably have better ideas about what is going to happen over time than they do of what is going to happen from quarter to quarter. But that's not a view that my other colleagues necessarily share.

Greg Mankiw: Can I ask you a question about the two Fed chairmen that you mentioned? Did they think the Phillips curve wasn't very useful or that it was just a dumb idea? Did they think about it as theoretically useful, but not useful for me as a policymaker, or that it's kind of a goofy idea for you academics to think about?

Allan Meltzer: It's hard for me to pin down exactly what Volcker thought because he seemed to change his mind a lot. But with Greenspan, he was explicit. He said that this is a fine theoretical construct. It doesn't have much to do with what I have to do.

Robert Solow: On that issue I want to point out two things. First of all, Allan, it may be true that two very successful Fed chairman had no use for the Phillips curve. The same is true of the unsuccessful chairmen. So there's not a lot of power in that test.

The second point is that I'm not sure that it is the right question to ask. Playing with the Phillips curve algebraically or graphically is something that some people have a temperament to do and others not. But I wonder if you had asked either Paul Volcker or Alan Greenspan this question—"if you strongly suspected that excess demand was emerging in the labor market, would you think wages were likely to rise?"—how do you think they would have replied?

Allan Meltzer: That's not a theoretical question, that's an empirical question. It's a question about whether that theory applies to the labor market, and the answer they would give is yes. The question then becomes for them, does that explain to me what is going to happen in the near term to the rate of price change and the rate of wage change and the rate of output change and the unemployment rate, and the answer they gave to that question was no.

Robert Solow: That I don't understand.

John Taylor: Maybe I could comment here. I think people have different ways of talking about the Phillips curve. Sometimes people say that the Phillips curve means there is a permanent trade-off between inflation and employment, others that there is only a short-run such Phillips curve. Some people refer to one simple equation. Others contend that you have to bring in wages too, at least. Still others say that you have to bring in an interaction between wages and prices, and I tend to agree with that viewpoint. One equation just isn't going to make it. So, when you read in the press that somebody believes in the Phillips curve, you think he must be out of his mind. But that article is probably referring to one particular caricature of the Phillips curve or one particular aspect of it. Wage and price dynamics have been part of economics since the days of Hume, and it seems to me that's what people want to understand more about, whether you say they use the Phillips curve or not.

Greg Mankiw: With that, let me bring the session to a close. We thank both of you and the audience for the questions. Thank you very much.