The conquest of American inflation: A summary

By Thomas J. Sargent and Ulf Söderström
Sargent is Professor of Economics at Stanford University and Senior Fellow at the Hoover Institution in Stanford, California; Söderström is a visiting researcher at the Research Department of Sveriges Riksbank

This essay compares two different interpretations of postwar U.S. inflation. In both stories, the government learns a version of the natural unemployment rate hypothesis: in the first, the correct rational expectations version, in the second, an approximating adaptive expectations version. Although the first story is more popular among modern macroeconomists, it suffers from contradictions and loose ends. Therefore the second story is considered. This story, which captures important features of policy making at the Federal Reserve, is more successful than the first in explaining the rise and fall of American inflation.

The rise and fall of U.S. inflation

Figure 1 plots the annual rate of inflation in the U.S. since World War II. Inflation was low during the late 1950's and early 1960's, swept upward into the 1970's, and then fell abruptly with Chairman Paul Volcker's stabilization in the early 1980's. If we take for granted that inflation is under the control of the Federal Reserve, how can we explain these observations?

This essay evaluates two interpretations of the U.S. post-war inflation history based on policy makers' beliefs about the Phillips curve. In both stories, the Federal Reserve authorities learn the natural rate of unemployment theory from a combination of experience and a priori reasoning.[1] The stories differ in how the

This is a summary of Sargent (1999), which in turn is based on the Marshall lectures at the University of Cambridge in October 1996 and the Nemmers inaugural lecture at the Northwestern University in May 1997. See also Cho and Sargent (1999) for formal proofs of the main arguments. Anders Vredin and Marianne Nessén have provided helpful comments on early drafts of this summary.

[1] The natural-rate theory was developed and refined by Phelps (1967), Friedman (1968), and Lucas (1972). It implies that the average rate of unemployment is determined by real forces, and not altered by permanent changes in the inflation rate.

Figure 1. U.S. post-war inflation

Monthly CPI inflation, 13-month centered moving average, percent

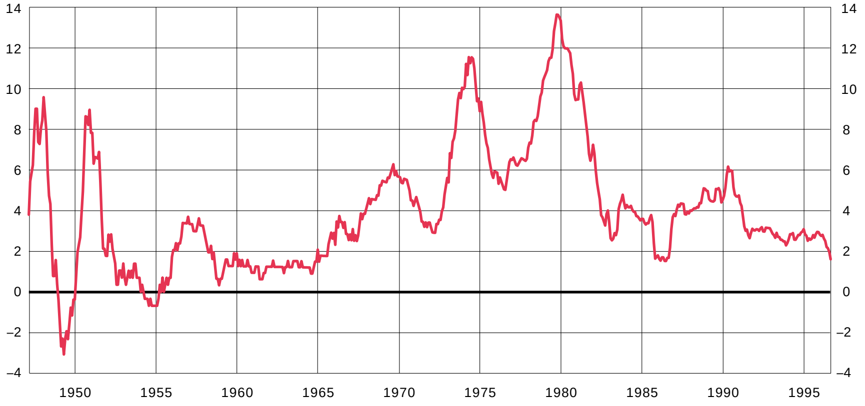

natural-rate theory is cast. In the first story, the government learns the correct rational expectations version of the theory, whereas in the second story, it learns an approximating adaptive version, adjusting its policy behavior as the economy develops over time. The first story will be called the *triumph of natural-*

This essay evaluates two interpretations of the U.S. post-war inflation history based on policy makers' beliefs about the Phillips curve. The first story will be called the *triumph of natural-rate theory* and the second one the *vindication of econometric policy evaluation.*

Figure 2a. Monthly unemployment rate and inflation

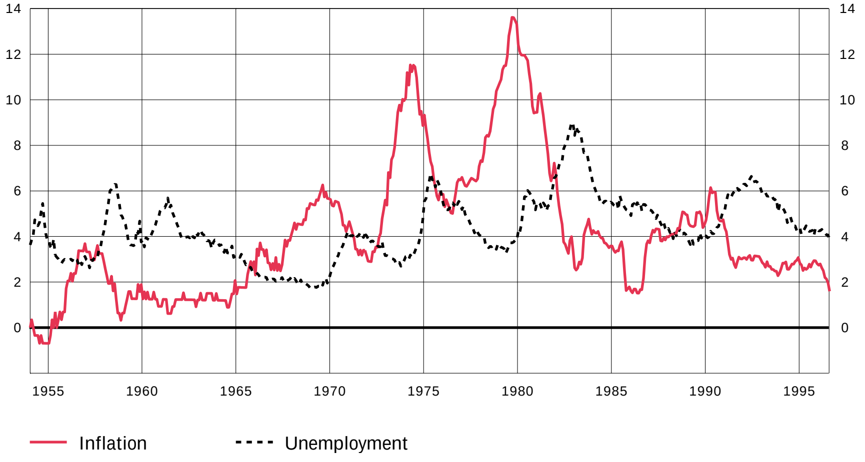

Figure 2b. Business cycle components of unemployment and inflation

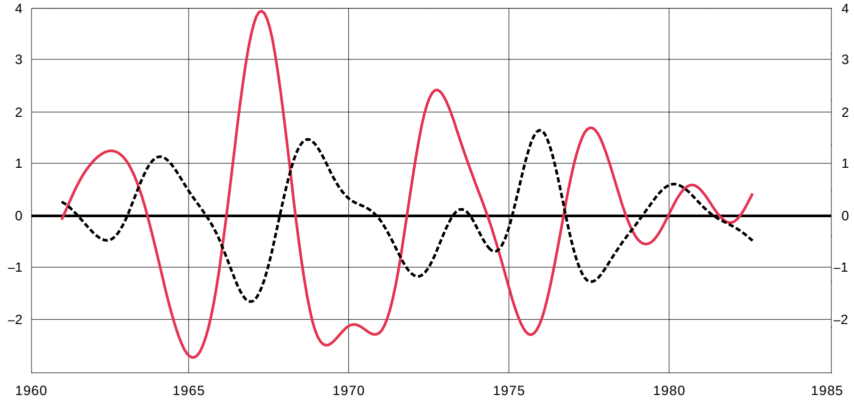

Figure 3a. Scatter plots of unemployment (U) and inflation (y) 1960–1982

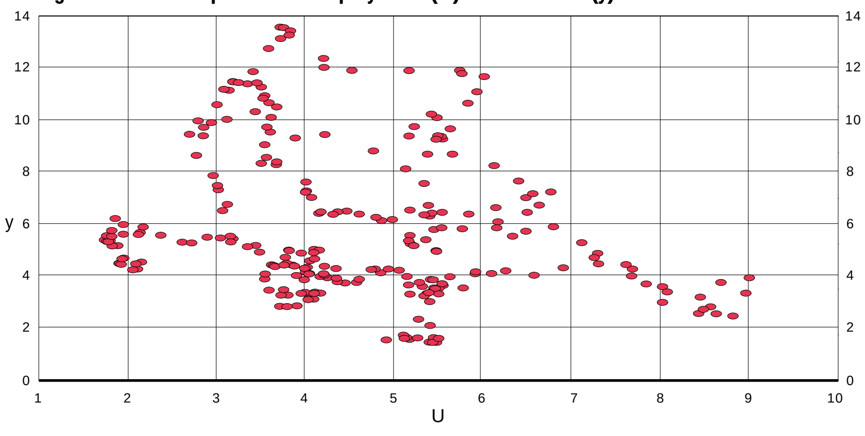

Despite its disrepute within important academic and policymaking circles, the Phillips curve persists in U.S. data.

rate theory and the second one the *vindication of econometric policy evaluation*.

Despite its disrepute within important academic and policymaking circles, the Phillips curve persists in U.S. data. Figure 2a plots the unemployment rate for white men over 20 years of age against the CPI inflation rate, and Figure 2b plots business cycle components of inflation and unemployment.[2] Likewise, Figure 3a plots monthly inflation against monthly unemployment for the 1960–82 subperiod, which interests us most, and Figure 3b plots the business cycle components of these two series. An examination of these figures

[2] These business cycle components have been calculated using methods described by Baxter and King (1999).

Figure 3b. Scatter plots of business cycle components of unemployment (U) and inflation (y) 1960–1982

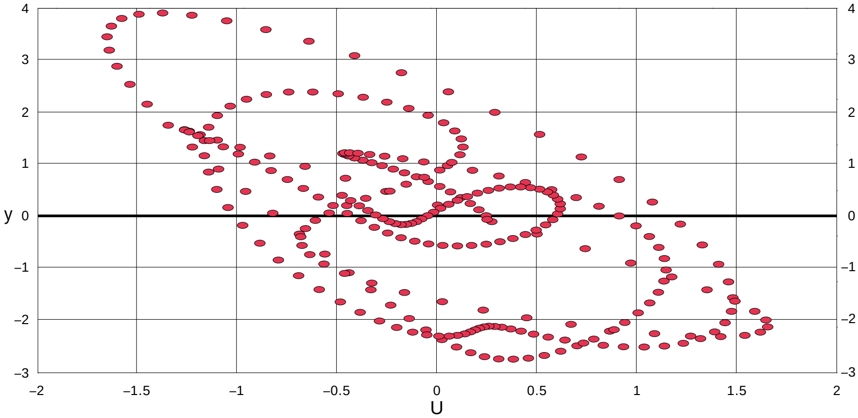

allows the eye to spot an inverse relationship between inflation and unemployment at business-cycle frequencies – a Phillips curve.

"The triumph of natural-rate theory"

The first story that explains the inflation history of the U.S. builds on work by Kydland and Prescott (1977) and Barro and Gordon (1983). Adherence to the gold standard and then to the rules of Bretton Woods gave the U.S. low inflation and low expectations of inflation. In 1960, Paul Samuelson and Robert Solow found a Phillips curve in the U.S. time series for inflation and unemployment (Samuelson and Solow, 1960). They argued that the negative relationship between inflation and unemployment was exploitable and suggested raising inflation to reduce unemployment. Soon, their recommendation was endorsed by many macroeconomists and implemented by policy makers. To everyone's dismay, over time the Phillips curve shifted adversely: inflation rose, but unemployment on average did not fall.

In the meantime, the concept of the natural rate of unemployment was created and refined. This theory, which assigns a central role to people's expectations about inflation in locating the Phillips curve, allowed only a temporary trade-off between inflation and unemployment and could explain the observed adverse shifts in the Phillips curve. As inflation rose, unemployment would be

The concept of the natural rate of unemployment allowed only a temporary trade-off between inflation and unemployment and could explain the observed adverse shifts in the Phillips curve.

temporarily reduced. Eventually, however, as the public adjusted its expectations to the new level of inflation, unemployment would move back towards the natural rate. The rational expectations version of the theory meant that policy makers should ignore any temporary Phillips curve trade-off and strive only for low inflation. These ideas spread among academics, then influenced policy makers, and ultimately promoted the lower inflation rates of the 1980's and 1990's. Thus, events were shaped by policy makers' beliefs – some false, others true – and the actions those beliefs inspired.

"The vindication of econometric policy evaluation"

The alternative interpretation ascribes Volcker's conquest of inflation in the 1980's partly to the success of the econometric and policymaking procedures that Robert Lucas challenged in his famous Critique.[3] This story also assumes that the data conformed to the natural-rate hypothesis, whether or not the policy makers realized it. Policy makers accepted the Phillips curve as an exploitable trade-off; they also adopted their methods for learning from data and for deducing policy recommendations (in the way criticized by Lucas). Recurrently, they re-estimated the Phillips curve and used it to reset the inflation and unemployment targets, ignoring the effects of inflation expectations on the Phillips curve. That method revealed a shifting Phillips curve, which, when interpreted mechanically, led policy makers to pursue lower inflation.

The vindication story describes the post-war history of U.S. inflation in terms of an adaptive theory of policy.

To complete the vindication story, this essay describes the post-war history of U.S. inflation in terms of an adaptive theory of policy. The theory originates with a minimal departure from rational expectations and accounts for features of the data that rational expectations theory misses. Although the rational expectations revolution of the 1970's left adaptive expectations outmoded, the story recalls adaptive expectations in a modern form, incorporating forecasting functions like those in rational expectations models but with coefficients that adapt to fit recent data. Thus, adaptive expectations play an essential role in generating the inflation observations and in improving theoretical outcomes, making them fit more closely to the observed history.

[3] Lucas (1976) used Samuelson and Solow's method for deducing policy recommendation from a statistical Phillips curve as an example of erroneous methodology. He concluded that since agents' behavior varies with changes in the government's policy rules, "...comparisons of alternative policy rules using current macroeconometric models are invalid regardless of the performance of these models over the sample period or in ex ante short-term forecasting" (p. 41).

This interpretation accounts for the post-war inflation in terms of an adaptive government adjusting its naive view of a Phillips curve in the light of recent evidence, a procedure that makes the government's inflation policy change over time. The account of the inflation process denies that inflation policy is conducted in a vacuum or occurs as a natural experiment, as in the Lucas (1972, 1976) model where policy is exogenous. Instead, it asserts that inflation policy emerges gradually from an adaptive process. This story is more attractive than the triumph of the natural-rate story, since that story has contradictions, loose ends, and elements of adaptation. At the same time, although the vindication story backs away slightly from rational expectations, it imposes more restrictions on government policy than does the triumph of the natural-rate story, since policy, instead of being set exogenously, depends on the behavior of the economy.

Although the vindication story backs away slightly from rational expectations, it imposes more restrictions on government policy than does the triumph of the natural-rate story.

Reader's guide

The outline of this essay is as follows. First, the Lucas Critique is reviewed and modified, to set the stage for the adaptive models in the rest of the essay. The following section describes rational expectations models and the literature on the credibility problem of monetary policy. Using the basic model of Kydland and Prescott (1977), it is shown how attempts to exploit the temporary Phillips curve trade-off can lead to higher inflation, with no gains in terms of unemployment.

The remainder of the essay explores alternative modifications of the basic model that might produce the observed history. We start our departure away from rational expectations by returning to the origin of the natural-rate hypothesis in the 1950's and 1960's. That model is then modified to include misspecification, showing that if agents have only a slightly mistaken view of the Phillips curve, this can substantially influence the outcomes of inflation and unemployment. Then the concept of self-confirming equilibria is introduced. In such an equilibrium, the government's beliefs about the Phillips curve affect its policy choices, which in turn makes agents act in such a way that the government's beliefs are confirmed. Although these equilibria are not sufficient to generate lower inflation than in the previous rational expectations model, the idea that the government and private

The idea that the government and private agents learn about the Phillips curve by observing past data provides the basis for our vindication of econometric policy evaluation.

agents learn about the Phillips curve by observing past data provides the basis for our vindication of econometric policy evaluation.

Ignoring the Lucas Critique

This essay resurrects econometric and policy evaluation procedures that were strongly criticized by Lucas (1976). It emphasizes a neglected aspect of Lucas's Critique: drifting coefficients. In the adaptive models presented here, the government ignores the Critique and its implications for the government's econometric and policy procedures. Instead, these procedures make coefficients in the Phillips curve change over time, which in turn affects outcomes.

The procedures of policy evaluation using econometric models in the tradition of Tinbergen (1952) and Theil (1961) assume that private agents' behavior rules are fixed while the government considers variations of its policy rule. Lucas noted that if private agents solve intertemporal optimization problems, then their actions depend on the government's policy rule: if the policy rule is changed, agents will adjust their behavior. Because the Tinbergen-Theil formulation misses this dependence, it will not give reliable policy advice.

Lucas wrote the Critique when Keynesian macroeconometric models were highly regarded as tools for quantitative policy evaluation. He stressed that the methods used in econometric forecasting contradicted the assumption that agents' behavior rules were fixed. Instead, the coefficients in important forecasting equations were frequently adjusted. His interpretation of those adjustments was that the models were typically only approximations of the economy, with the relationships changing over time, that is, with drifting coefficients.

However, Lucas left the drift in coefficients unexplained. Neither the macroeconomic theory nor the rational expectations econometrics constructed after Lucas's Critique explains such drift. Each of these traditions focused on environments with fixed behavior. Yet coefficients continue to drift for macroeconometric models. Although the econometric forecasting literature has taken coefficient drift increasingly seriously, it typically offers no economic explanation of parameter drift.

Our interpretation of the U.S. inflation history starts with parameter drift, treating it as a smoking gun.

Our interpretation of the U.S. inflation history starts with parameter drift, treating it as a smoking gun. It is the key piece of evidence that the government's beliefs about the economy and therefore its policy toward inflation have evolved over time. In the model, government decisions are made in the Tinbergen-Theil tradition. However, the government's econometric procedures include drifting coefficients: as the economy develops, the government re-estimates the Phillips curve, leading to new coefficients, which are used as inputs in its policy decision. The exercise is constrained by assuming that a rational expectations version of the natural-rate hypothesis is true (although policy makers are not aware of it).

Focusing on the government's learning behavior raises an issue discussed in the wake of Lucas's Critique. If the fundamental factors of the environment are stable over time, the government's econometric estimates and its decision rule will eventually converge. In the limit, the economy can reach a self-confirming equilibrium, where the government's estimates of its econometric model are reinforced jointly by its own behavior and the private sector's reaction to it.

In a self-confirming equilibrium the government's estimates of its econometric model are reinforced jointly by its own behavior and the private sector's reaction to it.

Within such an equilibrium, the relevance of some aspects of the Lucas Critique disappear. First, although the assumption that private agents' decision rules are independent of the government's behavior is incorrect, the government will not be disappointed in the outcomes (as suggested by Lucas), since its beliefs are confirmed.

Second, in addition to imposing rational expectations for private agents, a self-confirming equilibrium restricts the government's econometric model and its behavior. It is an equilibrium where agents have rational expectations, but one with fewer parameters than those in the models of Lucas (1972, 1976), which had parameters describing government policy.

Third, since a self-confirming equilibrium does not admit regime changes and drifting coefficients, but we observe these in practice, convergence to such an equilibrium must be resisted through some mechanisms. Assuming that the government suspects that the environment is unstable (although in fact it is not), and therefore uses a learning rule that discounts past observations, weakens the tendency of the economy to converge to a self-confirming equilibrium. This makes regime shifts possible.

Ironically, the procedures that violate the Lucas Critique yield better outcomes than ones that respect it. If the government discounts past observations when estimating the Phillips curve, temporary disturbances may change the government's estimates, leading it to lower the rate of inflation. Thus the government's adaptive model can lead to better outcomes than under rational expectations.

The procedures that violate the Lucas Critique yield better outcomes than ones that respect it.

To expand on this point, the next section reviews the kinds of models and procedures that respect the Critique. To summarize important developments in macroeconomic thinking from the 1980's and early 1990's, the concept of Nash equilibrium is applied to a model with a natural unemployment rate. These rational expectations models respect the Lucas Critique. Models that challenge and extend the Critique appear later.

The credibility problem

This section first describes the basic one-period expectational Phillips curve model, modifications of which underlie both story lines. Then a version of this model if presented where the government also cares about the future, so its reputation in sustaining low inflation becomes important. The model formalizes the temptation to inflate unleashed by the discovery of the Phillips curve and the value of a commitment technology for resisting that temptation. It also shows that reputational mechanisms are not very successful as substitutes for the possibility to make a commitment.

A simple model

We first describe a version of the one-period model of Kydland and Prescott (1977).[4] A government decides the inflation rate, and private agents set their inflation expectations. The unemployment rate is assumed to deviate from the "natural" unemployment rate only if the actual inflation rate deviates from the expected rate; if there is surprise inflation, the unemployment rate is lower than the natural rate. The government's preferences are specified such that it would prefer to set both inflation and unemployment at zero, although the natural rate of unemployment is positive. Finally, we assume that private agents have rational expectations, so they understand the government's motivations, and on average set their inflation expectations equal to the actual inflation rate.

The inflation outcome depends crucially on whether the government takes into account the fact that its rule for the inflation rate will affect private agents' inflation expectations or not.

Note first that rational expectations imply that the unemployment rate cannot deviate from the natural rate in the long run, since private agents understand what inflation rate the government will choose, and adjust their expectations accordingly. The inflation outcome of this model then depends crucially

[4] See the Appendix for a simple formal presentation.

on whether the government takes into account the fact that its rule for the inflation rate will affect private agents' inflation expectations or not.

The *Nash outcome* is characterized by a positive inflation rate, but unemployment equal to the natural rate.

First assume that the central bank does not take the effects of its policy rule on inflation expectations into account, but instead takes the expected inflation rate as given. Then the government wants to lower the unemployment rate below the natural rate by setting actual inflation above the expected inflation rate, creating surprise inflation. However, since agents have rational expectations, they will see through the government's incentives, and set the expected rate of inflation equal to the actual rate. The equilibrium is reached at the level of inflation where the government has no incentives to change the inflation rate and private agents' expectations are fulfilled. This *Nash outcome* is characterized by a positive inflation rate, but unemployment equal to the natural rate.[5]

If the government realizes that the unemployment rate cannot deviate from the natural rate in the long run, the optimal rate of inflation is zero, leading to the *Ramsey outcome*.

If instead the government takes the effects on inflation expectations into account when setting the inflation rate, it realizes that the expected rate of inflation is equal to the actual rate in equilibrium, and that the unemployment rate cannot deviate from the natural rate in the long run. Then the optimal rate of inflation is zero, leading to the *Ramsey outcome* of zero inflation and unemployment equal to the natural rate.[6]

Figure 4 shows a graphical representation of the Kydland-Prescott model, where inflation (y) is measured along the vertical axis and unemployment (U) along the horizontal axis. The straight solid lines are a family of short-run Phillips curves, each associated with a different level of inflation expectations (higher expected inflation leads to an upwards shift of the short-run Phillips curve). In the long run, when inflation expectations are correct, the Phillips curve is vertical (the dashed line), positioned where unemployment is equal to the natural rate. The curved dashed lines are indifference curves for the government, that is, combinations of inflation and unemployment that yield the same level of utility. The fur-

[5] In game theoretic terms, a pair of strategies is a *Nash equilibrium* if A's choice is optimal given B's choice and vice versa. Then neither player has any incentives to change his strategy. See, for example, Varian (1987, ch. 27). The equilibrium outcome in the Kydland-Prescott model is referred to as "Nash" since the public's expectations are fulfilled and the government has no reason to change its policy choice.

[6] This outcome is termed "Ramsey" since it is the outcome that would be chosen by a central planner maximizing social welfare. This problem, although in the context of consumption and investment decisions, was first analyzed by Ramsey (1928).

Figure 4. Nash and Ramsey outcomes in the Kydland-Prescott model

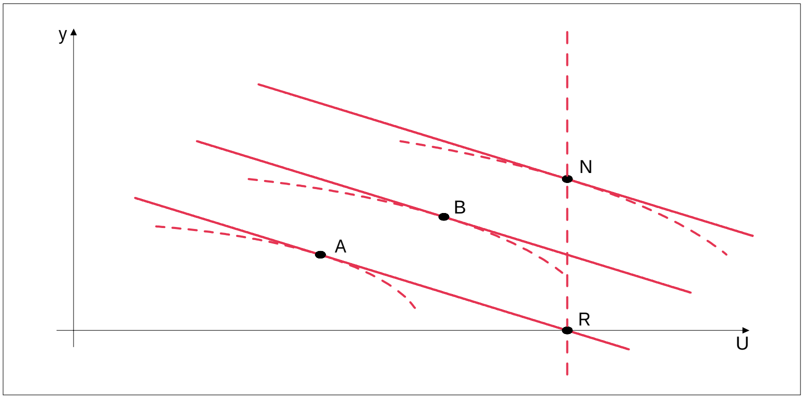

ther out from the origin (the government's preferred outcome) is the indifference curve, the lower is the government's utility.

Suppose first that inflation expectations are zero, so the short-run Phillips curve is given by the lowest solid line. Given this Phillips curve, the government chooses the combination of inflation and unemployment that maximizes its utility, leading to point *A*, where the Phillips curve is a tangent to one of the government's indifference curves.

At *A*, however, the public's expectations are not fulfilled (it expected zero inflation, but the government set a positive rate of inflation). Therefore *A* cannot be a rational expectations equilibrium. If the public had expected higher inflation, the short-run Phillips curve would have shifted upwards, leading the government to choose a higher rate of inflation (for example at *B*). The only Nash equilibrium in this model is given by the point *N*, where the public's expectations are fulfilled, and the government maximizes its utility, given the short-run Phillips curve. At this equilibrium, since inflation is equal to expected inflation, unemployment is equal to the natural rate. Thus, the Nash equilibrium is situated along the vertical long-run Phillips curve.

Given this long-run Phillips relationship, the government would have preferred to set inflation to zero, reaching the Ramsey outcome *R*. It is the inability to commit to this solution, and abstain from the short-run gains of lowering unemployment, that leads to the Nash outcome.[7] This outcome is clearly worse for the government than the Ramsey outcome, since it has higher inflation but

[7] Note that the assumption that the government aims for unemployment lower than the natural rate is crucial for these results. If the government's target for the unemployment rate coincides with the natural rate, it chooses the Ramsey outcome directly, and has no reason to exploit the Phillips curve trade-off.

the same rate of unemployment. The addition of a constraint to the government's problem in the Ramsey plan makes the government achieve better outcomes by taking into account how its actions affect the public's expectations. This is how Kydland and Prescott (1977) reached the pessimistic conclusion that a benevolent and knowledgeable government would set inflation too high because it makes decisions sequentially, not once-and-for-all, and cannot commit to keeping inflation low.[8]

A benevolent and knowledgeable government would set inflation too high because it makes decisions sequentially, not once-and-for-all, and cannot commit to keeping inflation.

The Nash outcome does not depend crucially on the assumption that the public's expectations are fully rational. If instead it forms its expectations adaptively, depending on the past inflation rate, one can imagine a dynamic version of the model. The government sets the inflation optimally given the public's expectations. In the next period, the public observes the inflation rate and updates its expectations, and the government resets inflation at a new level, after which the public's expectations are changed, etc. This leads to a process that eventually converges to the Nash outcome, where the private sector's expectations are realized, and the government sets the inflation rate optimally, given the expected rate of inflation. The speed of convergence is determined by the public's rate of learning, or updating of expectations. The more important is the most recent observation of inflation (the faster agents learn), the faster is convergence to the Nash outcome.

Thus, both in its simplest form and in more general forms, the Kydland-Prescott model leads to a pessimistic conclusion: unless the government somehow can commit not to try to lower unemployment by creating surprise inflation, the average inflation rate will be higher, but with no gains in terms of unemployment.

This model can thus explain why inflation increased in the late 1960's and early 1970's after the Phillips curve trade-off was discovered. To explain the stabilization of inflation in the early 1980's, however, we must accept that the government eventually learned the natural rate hypothesis, and abstained from trying to lower unemployment. But the lesson from Kydland and Prescott (1977) is that the government must find a commitment mechanism to do so. Since such mechanisms are rarely seen in practice, this simple story of

This model can thus explain why inflation increased in the late 1960's and early 1970's after the Phillips curve trade-off was discovered.

[8] The Nash equilibrium can be seen as a setup where the government chooses after the private sector sets its expectations. In the Ramsey equilibrium, on the other hand, the government chooses first, knowing that its decisions affect the private sector's expectations.

the conquest of inflation is not all that attractive. We will therefore search for other types of explanations.

Reputational mechanisms

In the simple Kydland-Prescott model, the government acts as if there is only one period, or it forgets that the economy lasts for more than one period. However, better outcomes can occur if the government plans for the future. When the economy lasts for several periods, the reputation of the government is important: by setting inflation lower than is optimal in the one-period model, it affects inflation expectations in the future, which can lead to better outcomes.

If the game between the government and the private sector is repeated indefinitely, Barro and Gordon (1983) have shown that better outcomes can be sustained if the government cares sufficiently about the future. This is because the gains from repeating the Ramsey outcome indefinitely may exceed the gains from fooling the private sector once by setting the optimal inflation rate in one period and then always repeating the Nash outcome.

Even if reputational forces are important, they provide a weak foundation for anti-inflation policy.

It is a well-known result from game theory, however, that such a situation has many equilibria, some better and some worse than the Nash outcome. Even if reputational forces are important, they therefore provide a weak foundation for anti-inflation policy. This essay now explores a different route, beginning with the notion that the multiplicity of equilibria stems from the rationality assigned to all participants in the system. To eradicate multiplicity, we will retreat from perfection and move to models in which some people have a more limited understanding of the economy.

Adaptive expectations (1950's)

The government sets inflation optimally, given its information about the economy and the public's expectations.

In 1967, Edmund Phelps formulated a theory of the natural-rate model based on the premise that the government sets inflation optimally, given its information about the economy and the public's expectations. He dropped rationality for the public, but not for the government, and assigned the public a particular mechanical forecasting rule known to the government: the public is assumed to form their inflation expectations as a weighted average of past observed inflation rates. This mechanism implies that as inflation moves around, the public's inflation expectations are not fulfilled, but if the government sets the inflation rate at the same level in every period, the public's expectations will eventually come close to the actual inflation rate. In formal terms, the expectation mechanism is characterized by an "induction property".

The government uses an empirical Phillips curve to compute an optimal policy rule for inflation, and is assumed to maximize a similar payoff function as before, except that it now covers the entire future. If the government does not discount the future (that is, it cares as much about the future as about today), it will eventually set inflation to zero, whereby the economy converges to the Ramsey outcome. How fast the economy converges depends on the rate at which the public updates its expectations, and although the inflation expectations are incorrect along the transition path, they are eventually correct. This is because of the induction property, which plays an important role in this model. In the more sophisticated models to be examined below, the rate of updating – or learning – is determined within the model (along with other aspects of behavior), but the induction property will still have a beneficial effect on the outcome.

The induction property has also been important in the empirical testing of the natural-rate hypothesis. There are two different interpretations of the Phillips curve: according to a Keynesian view, inflation is determined by the rate of unemployment, whereas a classical interpretation sees the relation in the opposite direction, from inflation to unemployment. In a traditional Keynesian Phillips curve, therefore, the current inflation rate depends on the expected rate of inflation and the unemployment rate. If inflation expectations are determined as in the Phelps model, as an average of past inflation rates, the induction property implies that the coefficients on lagged inflation in an estimated Phillips curve sum to unity. In the classical Phillips curve, on the other hand, where unemployment depends on the expected and actual inflation rates, the coefficients on current and lagged inflation sum to zero.[9] The induction hypothesis thus restricts the weights on lagged inflation in both versions of the Phillips curve.

When the empirical Phillips curve fulfills the induction hypothesis, the Phelps problem recommends lower inflation rates than when it doesn't.

Solow (1968) and Tobin (1968) proposed an empirical test of the natural-rate hypothesis by estimating the Keynesian Phillips curve and testing whether the sum of the coefficients on lagged inflation were equal to unity or not. They interpreted a finding of a sum of coefficients below unity as indicating a long-run trade-off between inflation and unemployment, a trade-off that could be exploited by the government. If, on the other hand, the sum of coeffi-

[9] See the Appendix for more details.

cients were equal to unity, the natural-rate would be supported, and there would be no trade-off to exploit. Thus, when the empirical Phillips curve fulfills the induction hypothesis, the Phelps problem recommends lower inflation rates than when it doesn't, since there is no trade-off between inflation and unemployment. This econometric test will turn out to play an important role in the complete model below.

Figure 5. Optimal inflation with differently sloped Phillips curves

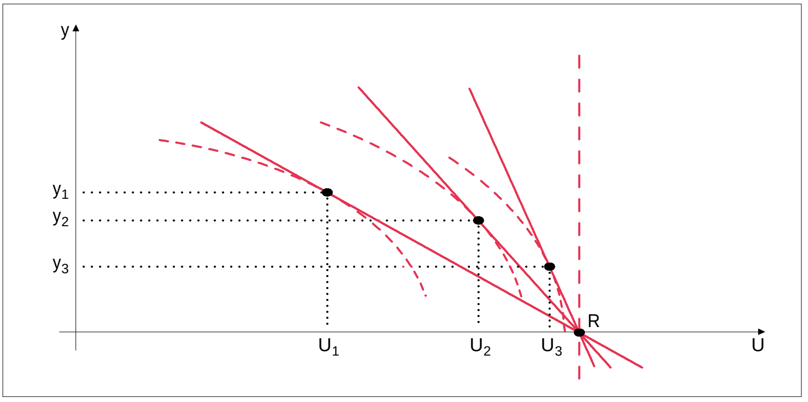

As the estimated Phillips curve moves around over time, so does the government's optimal rate of inflation.

In terms of the graphical representation of the Kydland-Prescott model, Figure 5 shows how the optimal inflation rate depends on the slope of the short-run Phillips curve. The steeper the Phillips curve, the lower is the inflation rate chosen by the government. This is because the trade-off between inflation and unemployment is less favorable, so to achieve a given decrease in unemployment, a larger increase in inflation is needed. In the extreme case when the short-run Phillips curve is vertical, the government does not identify any trade-off, so it sets inflation at the Ramsey outcome. Thus, as the estimated Phillips curve moves around over time, so does the government's optimal rate of inflation.

These results point us in one possible direction. What we are aiming at is a model that can explain why the government's estimates of the Phillips curve have shifted over time. If we can find such a model, we might be able to explain both the acceleration and the stabilization of inflation in the U.S.

Self-confirming equilibria

This section takes up the quest for models that depart minimally from the basic Kydland-Prescott model, but that also can replicate a 1960's acceleration of inflation followed by a Volcker stabilization in the early 1980's. Ideas from two literatures are combined to build a model with imperfect rational expectations that lead to equilibria where the government's view of the Phillips curve shifts over time. Similar models were initially constructed in the 1970's and 1980's in response to Lucas's Critique.

Ideas from two literatures are combined to build a model with imperfect rational expectations that lead to equilibria where the government's view of the Phillips curve shifts over time.

Assuming that expectations are adaptive, given by an average of past observations, if agents learn in a particular way (weighting all observations equally), we have seen how the Kydland-Prescott model predicts that the economy eventually converges to the same equilibrium as if expectations had been fully rational. However, if agents instead choose to discount past observations and weight new observations more heavily, this convergence can be arrested. Instead of converging to a constant rate of inflation, inflation expectations may converge to a process that moves around over time. In such a situation, if agents' beliefs are slightly wrong, although they learn in an optimal way, the outcome can be substantially different than in the rational expectations model. It will be argued that this is a plausible story about the U.S. inflation experience.

If agents' beliefs are slightly wrong, the outcome can be substantially different than in the rational expectations model.

Self-confirmation under different directions of fit

The story draws first on a literature that verifies the persistence of the Phillips curve in post-war U.S. data. King and Watson (1994) carefully document how an estimated Phillips curve is consistent with very different behavior of the economy over time, depending on the direction one estimates the Phillips curve: whether one regresses unemployment on inflation (as in the classical model) or inflation on unemployment (as in the Keynesian model). This implies that even if one can reach agreement about the existence and the slope of the Phillips curve, one may draw very different conclusions depending on one's interpretation of the curve.

The issues raised by King and Watson will be considered in the context of another literature that is concerned with different learning mechanisms when agents do not have the right model in mind, but use information in a rational

In a *self-confirming equilibrium* the government uses the wrong model of the economy, but fits its model to match the data as well as possible.

manner. Then a particular kind of imperfect rational expectations equilibrium can be defined – a *self-confirming equilibrium* – where the government uses the wrong model of the economy, but fits its model to match the data as well as possible. Two examples are described which differ in the direction that the government estimates the Phillips curve: the government either makes a Keynesian or a classical interpretation. Because the government's beliefs affect its behavior, the direction of estimation affects the outcomes of the model. All along, it is assumed that the government overlooks the econometric details, sees the Phillips curve as an exploitable relationship, and sets the inflation rate in the way assumed by Phelps. The true model is assumed to be given by the classical Phillips curve, but the government uses a misspecified econometric model, irrespective of whether it makes a Keynesian or a classical interpretation. In particular, when estimating the Phillips curve, it ignores the fact that changes in inflation expectations shift the curve up or down. These assumptions about how the government analyzes the economy do not seem entirely unrealistic.

Under the classical direction of fit, when the government regresses unemployment on inflation, the average outcome is equal to the Nash equilibrium of the simple Kydland-Prescott model. Under the Keynesian direction of fit, regressing inflation on unemployment, the estimated Phillips curve is flatter, so the government believes the trade-off is more favorable. As a consequence, it sets inflation higher on average (see Figure 5). This reflects the fact that the classical identification is closer to the model assumed to be true. The Keynesian identification scheme gives the government the mistaken interpretation of the relationship between inflation and unemployment, and that worsens the outcome.

Each of these self-confirming equilibria gives a mean outcome worse than the Ramsey outcome.

However, each of these self-confirming equilibria gives a mean outcome worse than the Ramsey outcome. What fails here is the induction property: since the sum of coefficients on lagged inflation are not restricted to unity, the government identifies a trade-off between inflation and unemployment that it wants to exploit (although the size of the trade-off varies between the different interpretations). As a consequence, although inflation differs depending on the government's interpretation of the Phillips curve, inflation is always higher than in the Ramsey outcome.

Equilibrium with misspecified beliefs

To explore how imputing a different wrong model might improve upon the Nash outcome, suppose instead that the public, not the government, makes a subtle specification error. The government solves the problem considered by Phelps, leading to a feedback rule for inflation as a function of the public's expectations. The public does not have rational expectations, but forms its expectations adaptively, so its inflation expectations are given by past inflation. However, it sets its rate of updating to fit the observed data.[10] Together, the government's optimal policy rule and the public's optimal rate of learning imply that both the true and the misspecified models can be reached as an equilibrium.

Together, the government's optimal policy rule and the public's optimal rate of learning imply that both the true and the misspecified models can be reached as an equilibrium.

Simulations of this model show that it would take many observations for the people living in the economy to detect that their model is wrong. Nevertheless, the implied average inflation rate is substantially lower than the Nash outcome, and the more the government cares about the future, the closer the average inflation rate gets to the Ramsey outcome of zero inflation. This improved outcome is delivered by the induction hypothesis incorporated in the adaptive expectations scheme – together with a government that cares about the future. Because the rate of learning is an outcome of the model, instead of being set exogenously as assumed in Phelps's model with adaptive expectations, the present model is a more realistic description of the workings of the induction hypothesis in this economy.

These results give some grounds for optimism: after disappointments from our self-confirming equilibria, which lead to the same inflation rate as in the Nash outcome, the equilibrium with forecast misspecification does support better outcomes. The equilibrium we reach embodies a type of self-confirmation, but with agents using the wrong model. This shows that there is a nearly self-confirming model with much better outcomes than the Nash outcome. If we can combine this adaptive learning model with the type of self-confirming model we used in the previous section, we might end up with a model that can explain the conquest of American inflation. This is what we shall try to do next.

There is a nearly self-confirming model with much better outcomes than the Nash outcome.

[10] In essence, the public chooses the rate of updating to minimize the error when forecasting inflation, given the government's behavior.

Adaptive expectations (1990's)

This section modifies our self-confirming models to attain an adaptive version and studies whether these models converge to self-confirming equilibria. Now knowledge of the exact model is withheld from all agents – the government as well as the public – and they are required to learn by updating their estimated regressions as time passes. The government sets the inflation rate at the recommendation that comes out of its optimization problem, given the current estimate of the Phillips curve.

If we make the agents discount past observations, new outcomes can emerge. Agents' discounting of past observations can sustain paths that look like Volcker terminating inflation in the early 1980's.

Starting from the self-confirming equilibrium models of the previous section, we in effect alter the rate at which past observations are discounted in agents' learning process. If the rate is set to implement "least squares learning", so all observations are weighted equally, we eventually get nothing new from these models, because they typically converge to self-confirming equilibria with the Nash outcome. However, if we make the agents discount past observations, so more recent observations are weighted more heavily than older ones, new outcomes can emerge.[11] Agents' discounting of past observations arrests convergence to a self-confirming equilibrium and can sustain paths that look like Volcker terminating inflation in the early 1980's. Empirically, discounting is not entirely unreasonable, rather, it is a good idea to weight recent observations more heavily than past ones when one believes that the Phillips curve wanders around over time.

The main purpose is to study outcomes that emerge when the government is using a learning algorithm that impedes convergence to a self-confirming equilibrium. We are as interested in movements away from a self-confirming equilibrium as in those toward one.

In the type of self-confirming equilibria considered earlier, the government solves its optimization problem only once, at the equilibrium values of its perceived Phillips curve. It then implements the recommendations of this unique policy rule in every period as time passes. In contrast, the government is now assumed to continuously adapt to new information about the Phillips curve, and solve a new problem in every period. As a consequence, its policy rule for inflation will vary over time, depending on its latest estimates of the Phillips curve.

The learning algorithm that discounts past observations arrests the force for

[11] See the Appendix for details on the learning mechanisms used.

convergence to a self-confirming equilibrium, both under the classical and under the Keynesian identification of the Phillips curve. This opens the possibility that the model will produce different outcomes than the Nash outcome. This issue is explored using computer simulations.

Figure 6. Simulation of classical adaptive model

Inflation over time

Sum of coefficients on inflation

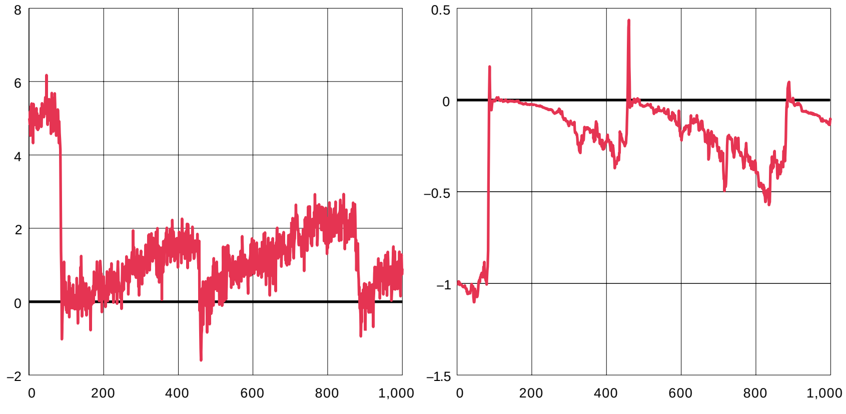

Figure 6 shows the results after simulating the classical model for 1,000 periods: the first panel shows the path of inflation over time, and the second panel shows the sum of coefficients on inflation in the government's estimated Phillips curve. In the first panel, inflation starts near the self-confirming Nash equilibrium value (which has inflation equal to 5 percent), but then drops almost to zero and stays there for a long time. Over time, inflation slowly heads back toward the self-confirming value, only to be propelled back toward zero again. Hence, the dynamics that pull the system toward the self-confirming equilibrium (the government's attempts to exploit the Phillips curve trade-off) are opposed by a recurrent force (learning about the natural rate) that sends the inflation rate close to the Ramsey outcome of zero inflation.

The dynamics that pull the system toward the self-confirming equilibrium are opposed by a recurrent force that sends the inflation rate close to the Ramsey outcome of zero inflation.

In the second panel of Figure 6, we see that during the first dramatic stabilization episode, the sum of coefficients on current and lagged inflation in the government's estimate of the classical Phillips curve jumps from its self-confirming value of – 1 to nearly zero. In the classical model, a value near zero of the sum of coefficients on inflation activates the induction hypothesis and makes the government reduce inflation, since it cannot find evidence for a Phillips curve trade-off. This makes the government more open to interpreting the data as consistent with the natural-rate hypothesis. After the stabilization begins, the government's behavior generates a series of observations that affirm the induction hypothesis. Evidently, the stabilization generates observations that temporarily add credibility to the induction hypothesis that prompted it. Our earlier analysis shows that this situation is not self-confirming in the technical sense. Nevertheless, it is reinforcing, due to the government's changing behavior in response to the new observations.

Figure 7. Simulation of Keynesian adaptive model

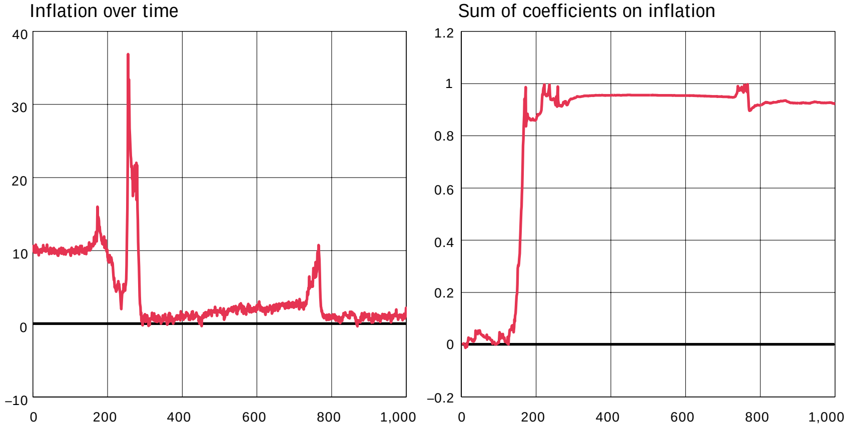

Under the Keynesian identification in Figure 7, the story is much like that for the classical identification scheme. The early part of the sample has inflation near the self-confirming value of 10 percent.[12] But the data from this period foster growing doubt about the location of the Phillips curve and put higher weight in the direction of the induction hypothesis, which manifests itself when the sum of the weights on lagged inflation equals one (see the second panel). Eventually, by chance (through shocks to the economy) some observations arrive that push the government's estimated Phillips curve toward the induction hypothesis. As the government solves the Phelps problem, it induces a stabilization, and the inflation rate settles below the Nash outcome of 10 percent.

Apparently, there is a mechanism that drives the economy away from the Nash outcome. The economy starts out at the Nash outcome (such as point *N* in

[12] As before, when the government uses the Keynesian interpretation of the Phillips curve, the Nash inflation rate is higher, since it interprets the relationship between inflation and unemployment in the wrong direction. (Recall that the classical interpretation is assumed to be true.)

Figure 4). At the same time, the economy is constantly hit by shocks to the unemployment rate. If, by chance, some of these shocks make the government's estimate of the slope of the Phillips curve steeper (remember that these new observations are weighted more heavily than past observations supporting the flatter Phillips curve), it will set the inflation rate lower than before (see Figure 5). But as the government sets inflation lower, the new observations created by the altered policy rule makes the next estimate of the Phillips curve even steeper, so inflation is set even lower. Eventually, we might reach the point where the government believes the Phillips curve is vertical, so there is no trade-off between inflation and unemployment, and it sets inflation equal to the Ramsey outcome (at *R* in Figures 4 and 5). The government's continuous updating of its estimated Phillips curve in combination with some initial favorable shocks have thus made it possible to escape from the Nash outcome and reach the optimal Ramsey outcome.

Apparently, there is a mechanism that drives the economy away from the Nash outcome.

However, even though inflation in both models settles close to the Ramsey outcome, this is not a stable equilibrium. Since the true model has a short-run trade-off in the Phillips curve, the government soon identifies it, after a sufficient number of new observations have arrived. It then wants to exploit the trade-off and set inflation higher to lower unemployment. As a consequence, the economy moves toward the Nash outcome again. If it eventually reaches the Nash outcome, the government gets caught in an "experimentation trap": if inflation expectations would vary sufficiently, the government would understand the natural-rate hypothesis and choose a lower inflation rate. But at the Nash outcome, inflation expectations are constant, since the government sets a constant inflation rate. Thus, for the government to start learning its version of the natural-rate hypothesis, the economy must be hit by shocks, which initiates the reinforcing dynamics towards the Ramsey outcome.

Consequently, the simulations have shown that for long periods, adaptive governments can learn to generate better than Nash or self-confirming outcomes. These results come from the recurrent dynamics induced by adaptation. The dynamics that drive the system toward a self-confirming equilibrium continue to operate under adaptation, but shocks let the adaptive system recurrently escape from a self-confirming equilibrium. Starting from a self-confirming equilibrium, an adaptive learning algorithm gradually makes the govern-

For long periods, adaptive governments can learn to generate better than Nash or self-confirming outcomes.

ment put enough weight on the induction hypothesis that eventually promotes better than Nash outcomes.

The learning mechanism used by agents is also a crucial ingredient in the model. If agents learn using the least squares learning algorithm that weights all observations equally, the model always converges to the inferior self-confirming Nash equilibrium. But if agents suspect the environment is changing, and therefore discount past observations when updating their expectations, better outcomes are possible. If, starting out from the Nash outcome, shocks hitting the model make the government decrease inflation (after solving the Phelps problem), the new observations will lend more support to the steeper Phillips curve, and the Phelps model will recommend even lower inflation. Eventually, this process can make the government learn the natural-rate hypothesis, and choose the Ramsey outcome. However, since this outcome is not an equilibrium (because the government cannot credibly promise to maintain zero inflation), eventually the government will increase inflation again, and the model moves towards the Nash outcome.

Our adaptive models thus contain basic support for our story explaining the conquest of American inflation during the 1980's. The econometric policy evaluation procedures criticized by Lucas are vindicated.

Our adaptive models thus contain basic support for our story explaining the conquest of American inflation during the 1980's as the success of the policy procedures of the 1960's and 1970's. Thus the econometric policy evaluation procedures criticized by Lucas are vindicated. It is time to leave the laboratory and turn to history. In the next section the historical data are taken as inputs and used to generate parameter estimates and residuals. How the model matches the data, and how it misses, will vindicate or indict econometric policy evaluation.

Econometric policy evaluation

This section reports estimates of adaptive versions of our model under both classical and Keynesian identification schemes. The primary purpose is to use our econometric results to assess whether and how they might vindicate our model government's econometric policy evaluation process. The results will show why the Volcker stabilization occurred and why it was postponed until the early 1980's.

For the classical and Keynesian models, respectively, Figures 8 and 9 display the one-step ahead prediction for inflation, the actual inflation rate and the estimated sum of coefficients in the empirical Phillips curve. The top panels reveal

Figure 8. Actual inflation and recommendation from classical model

Predicted and actual inflation

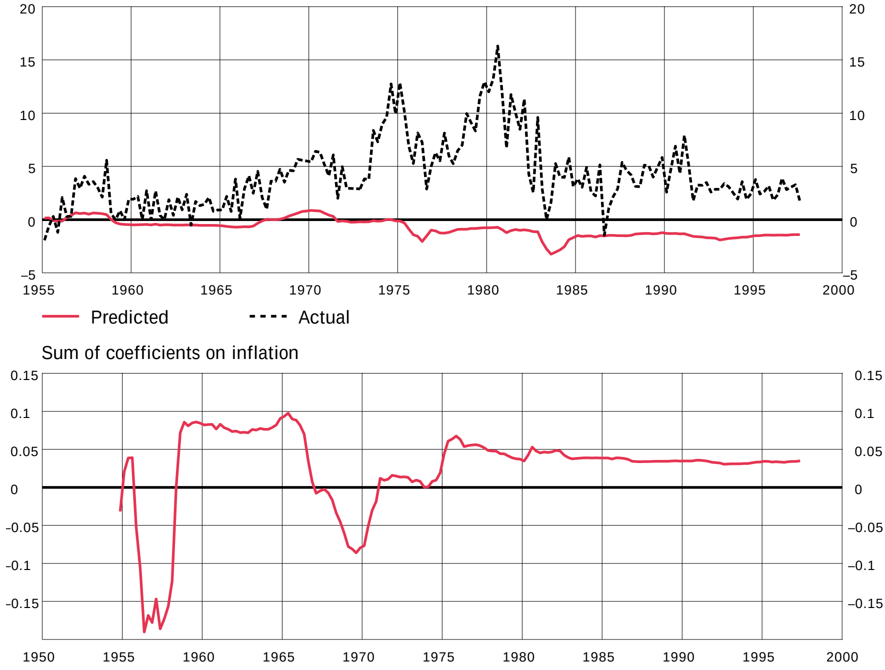

that both versions of the model fit the inflation process badly. The fit is appreciably worse for the model under the classical identification scheme. The Keynesian identification scheme leads to a more promising reflection of the inflation pattern, although the gap between predicted and actual inflation is large from 1973 until 1990. The Keynesian model to some extent matches the acceleration of inflation in the 1970's but underestimates inflation for the next 15 years. The classical model fails even to match the acceleration in inflation leading up to 1970. (The Keynesian model's better match to data caused its adoption in the U.S. Phillips curve literature.)

The large errors from our adaptive models are disappointing if we measure success by a good period-by-period fit. There is a long string of misses in the form of underpredictions of inflation during the 1970's, even for the Keynesian model. But such misses do not necessarily fail to vindicate econometric policy evaluation. To the

Unreconstructed Phillips curve fitters would have recommended lowering the inflation rate when inflation started to increase in the late 1970's.

Figure 9. Actual inflation and recommendation from Keynesian model

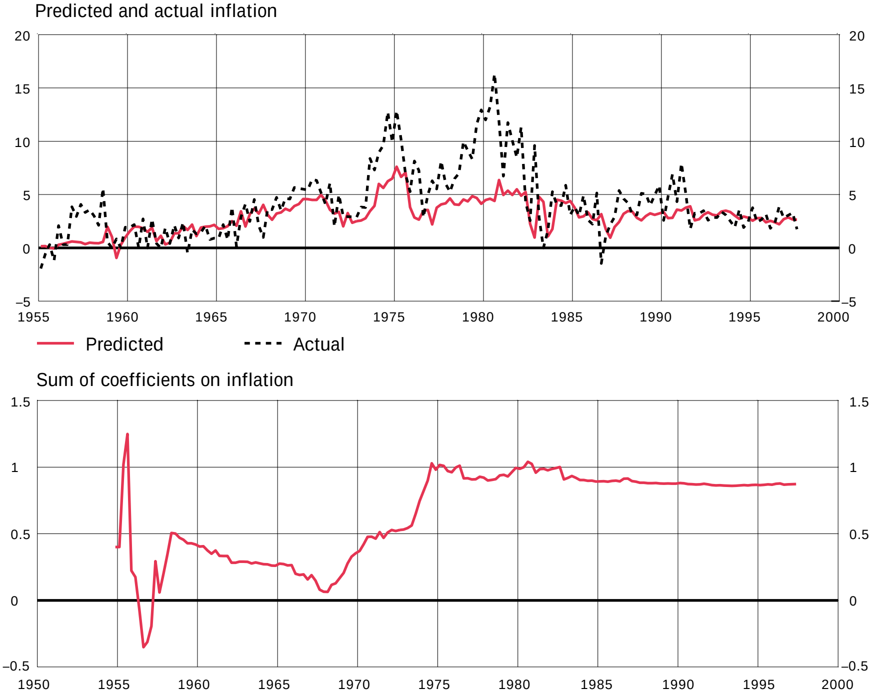

contrary, the pattern of misses from the estimated models favors vindication. Remember that the fitted value from the adaptive model is the government's recommendation of inflation as time passes. These recommendations under the Keynesian identification scheme actually confirm our story. The econometric estimates in Figure 9 tell us that unreconstructed Keynesian Phillips curve fitters would have detected the adverse shift in the empirical Phillips curve, and would have recommended lowering the inflation rate when inflation started to increase in the late 1970's. Those quantitative policy evaluators would not have concurred with the loosely argued recommendations current in the 1970's that long lags in expectations made it too costly to disinflate. Our results in Figure 8 say that recommendations under the classical identification would have been to lower inflation even earlier than under the Keynesian interpretation.

Triumph or vindication?

Expectations and the Lucas Critique

The contest between our two accounts of post-1960 U.S. inflation – the triumph of natural-rate theory and the vindication of econometric policy evaluation – raises various issues about rational expectations models of macroeconomic policies. We began from two benchmark models: (1) a natural unemployment rate model with adaptive expectations for the public, but an optimal policy for the government (Phelps's model); and (2) a natural-rate model with rational expectations for the public, but an exogenous and arbitrary government policy (the Kydland-Prescott model). Lucas recommended replacing the first benchmark model with the second. Coming to grips with our two stories about post-1960 inflation has caused us to propose other models that make various compromises between these two benchmarks.

After giving up a promise to offer recommendations, the theory of credible policy yields weak predictions about outcomes.

Models of credible government policies (such as those Kydland and Prescott, 1977, and Barro and Gordon, 1983) impose rationality on both the government and the private sector. In the end, it seems that after giving up a promise to offer recommendations, the theory of credible policy yields weak predictions about outcomes: even if concerns about the government's reputation might lead to better outcomes than in the one-period model, there are many possible outcomes, and it is not clear what will be the end result. This makes any declaration of the triumph of the natural-rate theory doubtful.

As an alternative, history was approached from the opposite pole, turning back from the Lucas Critique and beginning from Phelps's benchmark model. Starting with the problem he considered, we assumed that the government's model of the private sector's behavior is not arbitrary, based on some mechanical adaptive rule, but is chosen to fit historical data.

Via its connection to a self-confirming equilibrium, our 1990's adaptive model satisfies the desideratum that it should converge to rational expectations under tranquil conditions. But to match the data, our main interest has been in the recurrent dynamics contributed by adaptation. Suspecting that the Phillips curve is prone to wander, the government uses a learning algorithm that discounts past observations. Since the government's beliefs influence inflation via its optimization problem, its estimation choice makes a specification with time-varying coefficients worthwhile for both the government and the private sector.

Though both the public and the government behave in a stable way, our adaptive models generate simulations that exhibit abrupt stabilizations of inflation. It is the system's nonlinearities, rather than large shocks, that explain its behavior.

This brings us to regime shifts and nonlinearities. Though both the public and the government behave in a stable way, our adaptive models generate simulations that exhibit abrupt stabilizations of inflation. Regime shifts occur, not from a change in the government's econometric or policy-making procedures, but from disturbances and changes in beliefs created by the government's econometric procedures. But it is the system's nonlinearities, rather than large shocks, that explain its behavior.

This returns us to the origin of the induction hypothesis. This hypothesis was incorporated almost without comment by Friedman (1957) and Cagan (1956) in formulating the adaptive expectations hypothesis. It was also the basis of Solow's and Tobin's early tests of the natural-rate hypothesis. Cast as a villain in Lucas's Critique, the induction hypothesis emerges as a hero in delivering the superior long term outcomes in our simulations and the timely recommendations to stabilize inflation that emerge from our econometric estimates.

Reservations

The essay has compared two histories of postwar U.S. inflation: the *triumph of the natural-rate theory* and the *vindication of econometric policy evaluation*. Each history has the government learning and using a version of the natural unemployment rate hypothesis, either the correct rational expectations version in the triumph story or the approximating adaptive expectations version in the vindication story. The first history is more popular among modern macroeconomists than the second, which seems to defend discredited methods. The second story was considered partly because the first account has contradictions, loose ends, and elements of adaptation, and partly because the vindication story captures features of policy making at the Federal Reserve. The contest between the two histories is not rational expectations versus an alternative, because both selectively apply and withdraw from rational expectations.

The vindication of econometric policy evaluation is an exercise in positive economics, not normative economics. But because it produces near Ramsey outcomes for long periods, we might be tempted to transform it into a normative analysis recommending its econometric policy evaluation procedures. To dampen that temptation, we should recall the simulations presented above. The econometric policy evaluation methods would have yielded sound advice because the

U.S. data activated the induction hypothesis and recommended lowering inflation. However, the simulations contain episodes that resemble Arthur Burns as well as ones that look like Paul Volcker. In general, when estimates nearly affirm the induction hypothesis, the dynamics of the model point away from the induction hypothesis and toward regions where the government's estimates recommend resuscitating inflation to lower unemployment. It can take a long time to push the system back to the self-confirming equilibrium with high inflation, but it is bound to happen.

The vindication of econometric policy evaluation is an exercise in positive economics, not normative economics. The simulations contain episodes that resemble Arthur Burns as well as ones that look like Paul Volcker.

For this reason, the exercise in positive economics is not enough to commend its underlying policy making procedures. Theoretical work after Kydland and Prescott has insisted that anti-inflation policy is about designing and adhering to mechanisms that prevent the monetary authorities from choosing sequentially, and from even thinking about the possibility of lowering unemployment through inflation. That work seeks a secure foundation for assuring low inflation under fiat monetary systems. It rejects the idea suggested here that chance will lead policy makers armed with an approximate model eventually to learn to do approximately the right thing.

The exercise in positive economics is not enough to commend its underlying policy making procedures.

In the end, though our simulations and econometric evidence bolster the vindication of econometric policy evaluation story, we hope that it is the wrong story. We hope instead that policy makers somehow have learned a correct rational expectations version of the natural rate hypothesis and found devices to commit themselves to low inflation. Otherwise, the dynamics governing adaptation threaten eventually to rekindle inflation.

Though our simulations and econometric evidence bolster the vindication of econometric policy evaluation story, we hope that it is the wrong story.

Appendix

The credibility problem

In a simple version of the Kydland and Prescott (1977) model, the government sets the inflation rate y, and private agents choose their inflation expectations y^e. The unemployment rate U is determined by the relationship

$$(1) \qquad U = \bar{U} - \theta(y - y^e),$$

where $\bar{U}$ is the natural rate of unemployment and $\theta > 0$. The government's payoff is assumed to be

$$(2) \qquad -\tfrac{1}{2}\left[U^2 + y^2\right],$$

and private agents have rational expectations, so on average they set their inflation expectations equal to the actual inflation rate; $y^e = y$.[13]

If the government takes the expected inflation rate as given, substituting the unemployment equation (1) into the payoff equation (2), its payoff is

$$(3) \qquad -\tfrac{1}{2}\left[(\bar{U} - \theta(y - y^e))^2 + y^2\right],$$

so the optimal inflation rate is, minimizing equation (3) with respect to y,

$$(4) \qquad y = \frac{\theta}{1+\theta^2}\bar{U} + \frac{\theta^2}{1+\theta^2}y^e .$$

Assuming rational expectations, then, setting $y^e = y$ in equation (4), the *Nash outcome* is

$$(5) \qquad y^N = \theta\bar{U},$$

$$(6) \qquad U^N = \bar{U}.$$

If instead the government realizes that $y^e = y$, and thus that $U = \bar{U}$, its payoff function is simply

[13] Since there is no uncertainty in this simple model, rational expectations coincide with perfect foresight.

(7) $\quad -\frac{1}{2}\left[\bar{U}^2 + y^2\right],$

and it chooses the *Ramsey outcome*

(8) $\quad y^R = 0,$

(9) $\quad U^R = \bar{U}.$

Consequently, $y^N > y^R$, but $U^N = U^R = \bar{U}$, so the Nash outcome is worse than the Ramsey outcome.

Note also that both the optimally chosen inflation rate and the Nash inflation outcome depend on the government's estimate of the slope of the Phillips curve (θ). As the Phillips curve becomes steeper in inflation – unemployment space (as in Figure 5), θ falls. For a given level of inflation expectations, the optimal inflation rate in (4) then decreases, as does the Nash outcome (5). The Ramsey outcome is of course not affected, since it is always zero.

The classical and Keynesian Phillips curves

A simple version of the classical Phillips curve sees unemployment in period t as determined by the natural unemployment rate and the difference between actual and expected inflation (as in the Kydland-Prescott model) plus a disturbance:

(10) $\quad U_t = \bar{U} - \theta(y_t - y_t^e) + \varepsilon_t^U.$

The Keynesian version of the Phillips curve instead sees the inflation rate as determined by the expected rate of inflation and the difference between the actual and the natural rate of unemployment:

(11) $\quad y_t = y_t^e - \gamma(U_t - \bar{U}) + \varepsilon_t^y.$

Suppose the expected rate of inflation is determined by past inflation rates, for example,

(12) $\quad y_t^e = \alpha_1 y_{t-1} + \alpha_2 y_{t-2}.$

The induction hypothesis implies that if the inflation rate is set to a constant in every period, inflation expectations eventually converge to that constant. In our example, this means that if $y_{t-1} = y_{t-2} = y$, then eventually $y_t^e = y$ as well, so the sum of coefficients in the expectations mechanism (12) must equal unity: $\alpha_1 + \alpha_2 = 1$.

Using this expectational mechanism in the two versions of the Phillips curve, we get

$$(13) \qquad U_t = \bar{U} - \theta\,(y_t - \alpha_1 y_{t-1} - \alpha_2 y_{t-2}) + \varepsilon_t^v$$

in the classical model, and

$$(14) \qquad y_t = \alpha_1 y_{t-1} + \alpha_2 y_{t-2} - \gamma(U_t - \bar{U}) + \varepsilon_t^U$$

in the Keynesian model. Therefore, if the induction hypothesis holds (so $\alpha_1 + \alpha_2 = 1$), the estimated coefficients in front of current and lagged inflation in the classical Phillips curve sum to $\theta(1 - \alpha_1 - \alpha_2) = 0$, whereas the coefficients in front of lagged inflation in the Keynesian Phillips curve sum to $\alpha_1 + \alpha_2 = 1$. This in turn makes sure that there is no long-run trade-off between unemployment and inflation, but that the unemployment rate is determined by real factors only. Thus, the natural-rate hypothesis holds: setting inflation to a constant, both models imply that unemployment on average is equal to the natural rate.

Learning mechanisms

In the simple adaptive expectations model, the expected rate of inflation in period t depends on past observed inflation according to

$$(15) \qquad \begin{aligned} y_t^e &= y_{t-1}^e + \lambda_t(y_{t-1} - y_{t-1}^e) \\ &= \lambda_t y_{t-1} + (1-\lambda_t) y_{t-1}^e, \end{aligned}$$

where $0 \le \lambda_t \le 1$ is a parameter, possibly time-varying, that determines the rate of updating of expectations (or learning). The larger is λ_t, the more weight do agents put on the most recently observed inflation rate, so the faster they update their expectations. In the extreme case of $\lambda_t = 1$, agents simply set their inflation expectations equal to the observed inflation rate, $y_t^e = y_{t-1}$.

If λ_t is set to a constant ($\lambda_t = \lambda$ for all t), we can repeatedly substitute for the past inflation expectations, and get

$$
\begin{aligned}
y_t^e &= \lambda y_{t-1} + (1-\lambda) y_{t-1}^e \\
&= \lambda y_{t-1} + (1-\lambda)[\lambda y_{t-2} + (1-\lambda)(\lambda y_{t-3} + \ldots)] \qquad (16) \\
&= \lambda \sum_{\tau=1}^{t} (1-\lambda)^{\tau-1} y_{t-\tau}.
\end{aligned}
$$

In this case, with a constant rate of learning (or "constant-gain learning"), past observations are discounted by $(1 - \lambda)$ for each period, so more recent observations are weighted more heavily than older observations.

An alternative learning scheme ("least squares learning") weights all observations equally by setting $\lambda_t = 1/t$, so past observations are not discounted. Then inflation expectations are set as a simple average of all past observations:

$$
y_t^e = \tfrac{1}{t} \sum_{\tau=1}^{t} y_{t-\tau}. \qquad (17)
$$

References

Barro, Robert J. and David B. Gordon, “Rules, discretion, and reputation in a model of monetary policy,” *Journal of Monetary Economics* 12 (1), July 1983, 101–121.

Baxter, Marianne and Robert G. King, “Measuring business cycles: Approximate band-pass filters for economic time series,” *Review of Economics and Statistics* 81 (4), November 1999, 573–593.

Cagan, Philip, “The monetary dynamics of hyperinflation,” in Friedman, Milton (ed.), *Studies in the quantity theory of money*, Chicago University Press, 1956.

Cho, In-Koo and Thomas J. Sargent, “Escaping Nash inflation,” manuscript, Stanford University, August 1999.

Friedman, Milton, *A theory of the consumption function*, Princeton University Press, 1957.

Friedman, Milton, “The role of monetary policy,” *American Economic Review* 58 (1), March 1968, 1–17.

King, Robert G. and Mark W. Watson, “The post-war U.S. Phillips curve: A revisionist econometric history,” *Carnegie-Rochester Conference Series on Public Policy* 41, December 1994, 157–219.

Kydland, Finn and Edward C. Prescott, “Rules rather than discretion: The inconsistency of optimal plans,” *Journal of Political Economy* 85 (3), June 1977, 473–491.

Lucas, Robert E., “Econometric testing of the natural rate hypothesis,” in Eckstein, Otto (ed.), *The econometrics of price determination*, Board of Governors of the Federal Reserve System, 1972.

Lucas, Robert E., “Econometric policy evaluation: A Critique,” *Carnegie-Rochester Conference Series on Public Policy* 1, 1976, 19–46.

Phelps, Edmund S., “Phillips curves, expectations of inflation and optimal unemployment over time,” *Economica* 2 (3), August 1967, 22–44.

Ramsey, Frank P., “A mathematical theory of saving,” *Economic Journal* 38, 1928, 543–559.

Samuelson, Paul A. and Robert M. Solow, “Analytical aspects of anti-inflation policy,” *American Economic Review* 50 (2), May 1960, 177–194.

Sargent, Thomas J., *The conquest of American inflation*, Princeton University Press, 1999.

Solow, Robert M., “Recent controversy on the theory of inflation: An eclectic view,” in Rousseaus, Stephen (ed.), *Proceedings of a symposium on inflation: Its causes, consequences, and control*, New York University, 1968.

Theil, Henri, *Economic forecasts and policy*, 2nd edition, North-Holland, Amsterdam, 1961.

Tinbergen, Jan, *On the theory of economic policy*, 2nd edition, North-Holland, Amsterdam, 1952.

Tobin, James, "Discussion," in Rousseaus, Stephen (ed.), *Proceedings of a symposium on inflation: Its causes, consequences, and control*, New York University, 1968.

Varian, Hal R., *Intermediate microeconomics*, Norton, New York, 1987.