



The U.S. Phillips curve: The case for asymmetry

Douglas Laxton, David Rose, Demosthenes Tambakis*

Research Department, International Monetary Fund, Washington DC 20431, USA

Accepted 20 November 1998

Abstract

Recent statistical rejections of convexity in the Phillips curve have been uninformative because researchers have employed measures of business cycle gaps that are inconsistent with the implications of convexity. The paper shows that identifying convexity in the Phillips curve will become even more difficult if policymakers are successful in avoiding large boom and bust cycles. To the extent that convexity in the Phillips curve is used as a rationale for stabilization policy, our findings present an interesting conundrum because successful policymakers will further weaken the empirical evidence on which such policies are based. © 1999 Elsevier Science B.V. All rights reserved.

JEL classification: C51; E31; E52

Keywords: Phillips curve; Unemployment; Monetary policy

1. Introduction

The recent literature on the U.S. Phillips curve has featured a debate over the shape of the trade-off between inflation and the degree of excess demand, and the implications for monetary policy. A number of papers have suggested that the U.S. Phillips curve has convex shape in terms of unemployment, meaning that as unemployment falls below its sustainable level, the upward pressure on inflation rises increasingly, on the margin (Turner, 1995; Clark et al., 1996; Debelle and Laxton, 1997).¹

* Corresponding author. E-mail: dlaxton@imf.org

¹ In some cases, these authors studied the question using output rather than unemployment gaps; the issue is the same. Convexity was an important feature of the original Phillips (1958) curve and Lipsey's (1960) exploration of its foundations. Macklem (1996) and Clark and Laxton (1997) review the evolution of views on this issue. They argue that curvature was overshadowed by analytical tractability and other issues such as introducing expectations.

In contrast, Stiglitz (1997), citing work at the Council of Economic Advisers, and Eisner (1997) have maintained that the U.S. Phillips curve may not be convex, but may even be *concave*. Finally, Gordon (1997) has recently repeated his argument that the U.S. Phillips curve is linear and, indeed, ‘resolutely linear’ (p. 26). See also Fuhrer (1996).

With various economists arguing for convexity, concavity and linearity, what is the policymaker to think? This question is clearly important, since Stiglitz, then Chairman of the Council of Economic Advisers, also urged the Fed to be willing to experiment to test the limits of capacity on the grounds that, owing to the *absence* of the ‘traditional’ convexity in the U.S. case, the costs of error would be small.

In this paper, we explain how there can be such difference of opinion on the empirical evidence and we make the case that the Fed should presume the traditional convexity in the face of the empirical uncertainty, since the costs will be relatively high if there is convexity and policy decisions do not take this into account.

We suggest a methodology for estimating the Phillips curve that imposes modest convexity and also provides estimates of the nonaccelerating-inflation rate of unemployment (NAIRU) that are consistent with this hypothesis. We argue that estimating the NAIRU and the Phillips curve simultaneously is important for the identification of the latter. We contrast the historical performance of this model with that of a linear alternative, with and without the use of a survey measure of inflation expectations. We show that the convex form fits the U.S. data a bit better than a linear alternative, but not sufficiently better to be conclusive from the perspective of standard classical statistical tests.

We then focus on the power of econometric techniques to identify the correct structure. We specify a small macro model that includes our estimated convex Phillips curve. Repeated stochastic simulations of the model provide hypothetical data, which we then use to assess what would be concluded by a researcher estimating a Phillips curve with various methods for measuring the degree of excess demand and inflation expectations. The results suggest that econometric methods, especially when applied in the traditional way, i.e., with lags to capture inflation expectations and with predetermined measures of excess demand, will have very low power to identify the convexity in typical samples.

We report a number of types of sensitivity analysis. We find that the low power of traditional tests for convexity persists even for extreme changes to the assumptions.

2. Implications of alternative Phillips curves for stabilization policy

Fig. 1 shows the essential reason why stabilization policy matters when the Phillips curve is convex. The vertical axis shows inflation net of expected




Fig. 1. Implications of convexity in the Phillips curve.

inflation.² The horizontal axis shows the rate of unemployment, u . Convexity means that the cyclical trade-off between inflation and unemployment worsens on the margin as the latter is pushed below the point u^* .

We call the variable u^* the *deterministic* NAIRU or DNAIRU. It is the level of u at which there is no systematic pressure for inflation to rise or fall, relative to expectations, etc., *in the absence of shocks* (hence deterministic). An important point is that the DNAIRU is not a feasible stable-inflation equilibrium in a *stochastic* economy with convexity. The average level of u consistent with an expectations equilibrium must lie above this level. We illustrate this in Fig. 1, assuming that the Fed constrains net inflation to lie between plus

²This is a simplification. In general, the vertical axis also excludes any other terms in lagged inflation in the Phillips curve.

and minus one percentage point of expectations, which results in equilibrium (average) unemployment \bar{u} ; $E[u] = \frac{1}{2}(u_1 + u_2)$.³

One could call \bar{u} the ‘natural’ rate, on the grounds that, for given institutions and stabilization policy rules, this is the rate that is ‘ground out,’ in Friedman’s (1968) terms, on average, by markets. One could also use the term NAIRU for this value, because this is the level of u where there will be no acceleration (or deceleration) of inflation in a *stochastic* setting. We adopt this latter convention and call the stochastic equilibrium rate the NAIRU.

It is important to note that it is the equilibrium level of u that will be reflected in the data. Filtered measures of u from the data will not yield u^* , but rather something higher that is associated with the equilibrium rate. This has important implications for the methodology of estimating the Phillips curve and testing for convexity, which we consider later.

The shape of the Phillips curve has some striking consequences for stabilization policy. In Fig. 1, we show the difference between the NAIRU and the DNAIRU as α . The size of α will depend on the degree of convexity and the range of observation, as measured by, say, the dispersion of u . Thus, in the convex case, stabilization policy has first-order welfare effects because success in reducing the variability of u will also lower its mean value.⁴ In the linear world, by contrast, certainty equivalence results hold; the NAIRU equals the DNAIRU and the variability of u will not affect its average value, all else equal. In the linear case, moreover, there is no first-order cost of *delaying* response to boom conditions, or more generally to allowing a boom and bust cycle to develop, since the costs of eliminating inflation are about the same regardless of when action is taken, whereas in the convex case delay will raise equilibrium unemployment (Clark et al., 1996).

There may also be implications for *disinflation* strategies. Clark and Laxton (1997) argue that, unlike the linear case, the convex model suggests gradualism; Isard and Laxton (1996) argue that convexity provides a rationale for *opportunistic* disinflation strategies.⁵

3. Phillips curves and Phillips lines: The burden of proof

As noted in Section 2, there are two broad determinants of how important convexity will be from a policy perspective: the degree of convexity itself, and the

³To simplify the presentation, Fig. 1 assumes a simple distribution, where inflation net of expectations is either 1 percentage point or -1 percentage point. The value for \bar{u} will not generally be the point where the line LL cuts the horizontal axis.

⁴Faruqee et al. (1997) apply this idea to explain part of the rise in trend unemployment in Europe.

⁵See Orphanides and Wilcox (1996) and Orphanides et al. (1996) for analysis of opportunistic disinflation strategies, where the monetary authority does not actively seek disinflation, but acts to lock in lower inflation when negative shocks occur.

range of variation in the observed outcomes. The two have analogues in the econometric identification problem. The greater the degree of convexity in the Phillips curve, the easier it will be to identify that convexity with standard techniques. Also, for a given hypothetical convex Phillips curve, the greater the dispersion in the joint distribution of inflation and unemployment, the greater the chance of identifying the convexity.⁶

The behavior of policymakers plays a key role in determining the extent of the identification problem. It is well-established in the extensive literature on policy rules (e.g., Bryant et al., 1993) that the choice of a rule can have an important effect on the dynamic properties of an economy. Thus, the nature of the joint distribution of inflation and unemployment is, to an important degree, a matter of policy choice. Ironically, it is policy errors, especially big ones that create boom and bust cycles, that help the econometrician.

Applied econometricians tend to view estimation and testing in classical terms. An effect must be statistically significantly different from zero, at a high level of confidence, for it to be retained.⁷ Students are instructed in the risk of statistical errors of type 2, but in practice much more emphasis is put on minimizing the chances that a false effect will be retained than on guarding against exclusion of a true effect. There are dangers in this approach when the answer matters for policy or welfare questions. Results that indicate that a parameter is not significantly different from zero generally also indicate that the parameter is not statistically distinguishable from other values that may have major implications for a policy issue.

If one thought that empirical research could settle the issue, that would be the end of it. We will argue, however, that the U.S. data are not capable of providing a clear answer. Our Monte Carlo evidence suggests that this may be a generic problem. Since the policy costs of presuming linearity when the truth is convexity may be high, while the costs of presuming convexity when the truth is linear are low, the burden of proof should fall on those who would use linear forms to show that their restriction is valid, with a high level of confidence.⁸

⁶ Of course, the covariance structure of innovations matters. If positive shocks to excess demand are associated with negative shocks to prices, then higher dispersion alone will not be as helpful. The obvious practical example is a 'supply' shock that pushes prices up and at the same time creates unemployment, such as happens in the U.S. case with an oil price shock. This confounds, rather than helps, the identification of a Phillips curve.

⁷ Some economists have gone further, discarding statistically significant evidence of convexity to keep policy analysis simple. See, for example, Chadha et al. (1992).

⁸ See Laxton et al. (1993), Laxton et al. (1995), Clark et al. (1996) and Clark and Laxton (1997) for further discussion of these issues.

4. A model of U.S. inflation with modest asymmetry

In this section, we present a simple model of the U.S. inflation process that captures some key features of the interactions linking excess demand, inflation and monetary policy. The model consists of three estimated behavioral equations: a Phillips curve with modest convexity, the dynamics of aggregate demand, and an equation describing expectations formation. The model is closed with a monetary policy reaction function.

This simple model characterizes the dynamics of inflation as primarily dependent on excess demand (unemployment gap) and on inflation expectations. The fundamental role of the Fed is to ensure that the economy has a nominal anchor, which we treat as a target rate of inflation, but it also acts with an eye on unemployment. The policy instrument is the short-term interest rate (the Federal Funds rate), which has an effect on inflation through aggregate demand and employment. The monetary control mechanism is not perfect, because the economy is subject to shocks that cannot be foreseen, and because the influence of monetary policy on aggregate demand and employment operates with a lag.

4.1. The Phillips curve

Our Phillips curve is a variant of the specification in Debelle and Laxton (1997):

$$\pi_t = \lambda \bar{\pi}_t^e + (1 - \lambda)\pi_{t-1} + \gamma(u_t^* - u_t)/(u_t - \phi_t) + \varepsilon_t^\pi, \quad (1)$$

where u_t^* is the DNAIRU (we will be more precise on this below), and where

$$\bar{\pi}_t^e = \left(\sum_{i=0}^{12} \pi_{t-i}^e \right) / 12, \quad (2)$$

where, for estimation, π_j^e is the one-year-ahead expectation of inflation, held at j , from the Michigan survey data.

The model of how expectations influence inflation dynamics reflects a bargaining framework, as in Fuhrer and Moore (1995).⁹ There is also an element of intrinsic dynamics, represented in the separate term in lagged inflation. One can think of this as reflecting costly quarterly price adjustment by firms, between contracts. We estimate the weights, λ and $(1 - \lambda)$ applied to these two terms. Imposing the constraint that these weights sum to unity ensures that no

⁹ We assume that each contract has a 3-year horizon, and that 1/12th of the contracts are renegotiated each quarter. Hence, in any quarter, expectations formed up to 12 quarters ago will have an effect on the dynamics.

long-run trade-off exists between the levels of inflation and unemployment.¹⁰ The effective lag length has been made reasonably long, reflecting the stylized fact that the U.S. inflation process has a lot of persistence.

Our Phillips curve has the convex form shown in Fig. 1. The parameter ϕ defines a lower bound on u , reflecting short-run constraints on how far rising aggregate demand can lower unemployment before capacity constraints become absolutely binding and inflationary pressure becomes unbounded. In our empirical work, we allow ϕ to be time-varying:

$$\phi_t = \text{MAX}(0, \bar{u}_t - 4), \tag{3}$$

where \bar{u}_t is a measure of trend unemployment, established by some filtering method. Thus, ϕ_t is constrained to be zero when the trend unemployment rate is at or below 4%. When the trend value rises above 4%, ϕ_t moves up with the trend value, keeping a fixed difference.¹¹ Finally, the parameter γ is the lower intercept, i.e., the maximum rate of deflation that would be generated as excess supply became unbounded.¹²

Recently, researchers have attempted joint estimation of a time-varying NAIRU and the parameters of the Phillips curve using a Kalman filter (DeBelle and Laxton, 1997; Gordon, 1997; Faruqee et al., 1997). This is what we do here, writing

$$u_t^* = u_{t-1}^* + \varepsilon_t^{u*}, \tag{4}$$

where ε^{u*} is a shock term – the NAIRU is modeled as a random walk.¹³

If we rewrite Eq. (1) for heuristic purposes as

$$D\pi_t = (\gamma u_t^*)(1/(u_t - \phi_t)) - \gamma(u_t/(u_t - \phi_t)) + \varepsilon_t^\pi \tag{5}$$

where $D\pi$ is the difference between π and the expectations and lag(s), and where $\delta_t = \gamma u_t^*$ is treated as a time-varying parameter, we have a (nonlinear) estimation problem that can be solved using the Kalman filter technique.¹⁴ The estimates of u_t^* are given by $\delta_t/\hat{\gamma}$.

¹⁰ If convexity in the short-run Phillips curve becomes greater at low inflation rates, then the long-run Phillips curve would not be vertical. See Akerlof et al. (1996).

¹¹ This rule was established by looking at a variety of possibilities for both the trigger point (4 in Eq. (3)) and the rate at which ϕ_t moves with \bar{u}_t , when the latter is above this value. The specification in Eq. (3) provided the highest likelihood value among the alternatives considered.

¹² See Clark and Laxton (1997) for a review of sources of convexity in the Phillips curve. See Faruqee et al. (1997) for a more formal derivation of this particular functional form, based on a model of labor markets taken from Layard et al. (1991).

¹³ Obviously, the NAIRU cannot literally follow a random walk, since it has to be bounded. This is not a practical problem in estimation, but can be in stochastic simulations. We have imposed boundary conditions on u^* , so Eq. (4) does not always hold strictly.

¹⁴ Kuttner (1992, 1994) has applied this idea to measuring potential output.

Table 1
Estimates of U.S. Phillips curves with model-consistent DNAIRUs

Convex model:	$\pi_t = \lambda \bar{\pi}_t^e + (1 - \lambda)\pi_{t-1} + \gamma(u_t^* - u_t)/(u_t - \phi_t) + \varepsilon_t^\pi$						
Linear model:	$\pi_t = \lambda \bar{\pi}_t^e + (1 - \lambda)\pi_{t-1} + \beta(u_t^* - u_t) + \varepsilon_t^\pi$						
Gordon's model:	$\pi_t = G(L)\pi + \beta(u_t^* - u_t) + \varepsilon_t^\pi$ $u_t^* = u_{t-1}^* + \varepsilon_t^{u^*}$						
π_t	percent change in the CPI, quarterly at annual rates						
$\bar{\pi}_t^e$	average expected inflation – see text						
$G(L)\pi_t$	24-quarter distributed lag on past inflation						
Φ_t	minimum unemployment rate						
u_t	unemployment rate						
u_t^*	estimated DNAIRU derived from Kalman filter						
γ	estimated coefficient on the unemployment gap in the convex model						
β	estimated coefficient on the unemployment gap in the linear model						
LLF	value of the log likelihood function						
σ^2	variance of residuals of the estimated Phillips curve						
$\sigma_{\Delta u^*}$	standard deviation of quarterly change in the DNAIRU (1972:Q1 to 1997:Q1)						
\bar{u}^*	average value of the DNAIRU (1972:Q1 to 1997:Q1)						
Estimation period: 1968: Q1 to 1997:Q1							
Model	λ	γ	β	LLF	σ^2	$\sigma_{\Delta u^*}$	\bar{u}^*
Nonlinear	0.65 (8.25)	4.71 (5.54)		- 232.47	2.23	0.11	6.4%
Linear	0.59 (7.33)		1.04 (5.33)	- 235.30	2.45	0.11	6.6%
Gordon			1.24 (5.62)	- 235.10	2.30	0.11	6.6%

Note: Figures in parentheses are *t*-statistics.

Table 1 reports three sets of estimates, all with model-consistent measures of a time-varying DNAIRU. Inflation is measured using the overall CPI. The first results are for the convex model, estimated using the Michigan survey expectations data. All the coefficients have the expected signs and are statistically significant. The estimated function implies a modest, but important degree of convexity. The average value of *u* for this sample is 6.4%, while the average of our estimated u_t^* series is 6.1%. Thus, on average, the gap between the DNAIRU and the NAIRU has been about 0.3 percentage points.

The second set of estimates is for the linear model, again with the Michigan survey expectations, while the third set are for Gordon's (1997) linear model, i.e., with a 24-quarter distributed lag of past inflation.¹⁵ For these estimations, we

¹⁵ We use Gordon's restriction that fixes the 4 coefficients for the quarters of each lag year at the same value, so there are just 6 free parameters, not 24, in the estimated lag distribution.

impose that the standard deviation of the change in u^* is the same as estimated for the convex model. Neither linear model fits as well as the convex model, in terms of the log-likelihood for the system, but one cannot make a strong statistical case for convexity based on these historical estimates.

4.2. The dynamics of unemployment

The next equation provides the link between the instrument controlled by the monetary authorities and aggregate demand. We estimate an equation for u as a function of its lagged values and the real Fed Funds rate. The constant in the equation is allowed to be time-varying, following a random walk, to capture any combined changes in the trend levels of u and the real interest rate. It is also estimated using a Kalman filter, as follows:

$$u_t = c_t + 1.033u_{t-1} - 0.267u_{t-2} + 0.011r_{t-1} + 0.034r_{t-2} + 0.022r_{t-3} + \varepsilon_t^u, \quad (6)$$

$$c_t = c_{t-1} + \varepsilon_t^c, \quad (7)$$

$$r_t = rs_t - \pi 4_{t+4}^e, \quad (8)$$

where the lag structure of Eq. (6) is written in its final form, following testing down from specifications with longer lags. In Eq. (8), rs is the Fed Funds rate (annual rates), and $\pi 4_{t+4}^e$ is the Michigan survey measure of expected inflation over the next year held at t .¹⁶

Eq. (6) is written with the point estimates of the coefficients.¹⁷ Our results reflect two stylized facts concerning the ability of the Fed to control the economy. First, there are important lags between changes in interest rates and their effects on aggregate demand. Second, there is persistence in movements in the unemployment rate, implying that shocks to aggregate demand propagate into future periods. The coefficients on the unemployment lags imply some augmenting propagation, but with relatively speedy reversion to the mean. The dominant root is about 0.63. In terms of the monetary transmission mechanism, we have a sum of coefficients of 0.067, which is consistent with other estimates.¹⁸

¹⁶Fuhrer and Moore (1995) argue that it is longer-term interest rates that affect aggregate demand. To keep our model simple for the stochastic simulations, we use the interest rate controlled by the Fed. The nominal rate is converted to a real rate by subtracting the one-year-ahead inflation expectations taken from the Michigan survey. Thus, we treat the dynamic effects of expectations as much less sluggish in financial markets than in labor markets.

¹⁷See Laxton et al. (1998) for details on this estimation.

¹⁸For a recent summary of evidence on this issue, see Clark et al. (1997).

4.3. The monetary policy reaction function

To complete the model we specify a monetary policy reaction function that is a close cousin to the Taylor rule (e.g., Taylor, 1993).

$$rs_t = \pi 4_t + 0.5(\pi 4_t - \pi^*) + (u_t^* - u_t). \quad (9)$$

According to this simple rule, the monetary authority raises rs if inflation is above the target level or if the unemployment rate is below the DNAIRU. Following Taylor (1993), the inflation measure, $\pi 4_t$, is a 4-quarter rate of change of the CPI. However, while the Taylor rule uses a weight of 0.5 on the output gap, this rule uses a weight of 1 on the unemployment gap.¹⁹ Taylor (1998) shows that such rules provide a reasonable characterization of how the Federal Funds rate evolved after the great disinflation episode in the early 1980s and argues that following such a rule that responds to both inflation and the output gap is responsible for better macroeconomic performance over the last decade relative to earlier periods where monetary policy ‘errors’ resulted in significant boom and bust cycles.²⁰ We consider an alternative assumption in Section 5 where monetary policy is responsible for creating boom and bust cycles. This alternative rule is calibrated to create ongoing boom and bust cycles that have roughly the same magnitude as observed in the 1970s and early 1980s.

4.4. Expectations: Modeling the Michigan survey data

For the simulations, we need a *model* of how expectations evolve. An important question is the extent to which expectations are forward-looking and exploit the predictive content of current data. We define an instrument variable, a predictor of the future inflation, to use in a model of how expectations, as captured by the Michigan survey, are determined. Our auxiliary equation attempts to predict inflation over the next year using 4 lagged values of: the quarterly rate of inflation, the unemployment rate, a long bond rate, and the Federal Funds rate. The fitted values from this equation then serve as a proxy for a forward-looking component of expectations in the main regression, which

¹⁹ Because the unemployment gap tends to vary less than the output gap over the business cycle, it is necessary to use a higher coefficient to give this rule the basic properties of the Taylor rule. A coefficient of 1 is consistent with an Okun’s law coefficient of 2. The paper’s conclusions are not sensitive to the precise parameters chosen, as long as the rule does a reasonable job of stabilizing the business cycle.

²⁰ While his original weights were not formally estimated, Taylor (1998) provides weights on inflation and output gaps from equations estimated over the period from 1987Q1 to 1997Q3 that are very similar. Clarida et al. (1997) have estimated more complicated forward-looking reaction functions and argued that monetary policy has been used to stabilize the business cycle over this period. In the model considered here, forward-looking rules do an even better job of stabilizing the business cycle than simple Taylor rules.

explains the Michigan survey data using an own lag and the proxy variable. The results are as follows:

$$\pi 4_t^e = 0.267 \hat{\pi}_{t+4} + 0.733 \pi 4_{t-1}^e. \quad (10)$$

This equation is written with the point estimates, in the form we use for the simulations.²¹ We tested whether the instruments used to generate the proxy were significant in this regression. They were not. We also tested whether lagged inflation measures, in particular, have any explanatory power. They did not. The estimator preferred the lag of the dependent variable, conditional on the forward-looking proxy variable being there. We take this as strongly suggestive that expectations are not inherently backward-looking. The lags of inflation are useful in explaining the expectations data only to the extent that they help predict the future. A second important implication of these results is that there is inertia in expectations. There is a relatively high coefficient on the own lag, consistent with the view that expectations have important cycle propagation effects.

This estimated model serves to provide inflation expectations in the Monte Carlo exercise, where the forward component is solved as a model-consistent forecast. We assume for this research that the simulated measures are the true expectations and are observed by the econometrician. We leave for future work the question of what happens when expectations are observed with error. We do examine, however, what happens if the researcher ignores the expectations ‘data’ and uses lags to capture inflation expectations.

5. The power of econometric tests for convexity

In this section, we present the results of Monte Carlo experiments, where we ask whether an econometrician would be successful in uncovering the truth – that the Phillips curve is modestly convex. The data for this analysis are generated through repeated stochastic simulations of the hypothetical economy described by our model.

There are 4 shocks, disturbances to: the DNAIRU, unemployment (a demand shock), inflation (a ‘supply’ shock that shifts the Phillips curve), and to inflation expectations. We assume that these shocks are independent, and take their base-case standard deviations from the estimation results, with some adjustments. For the u^* and u shocks we take the estimation results directly, using standard deviations of 0.11 and 0.17 percentage points. For the inflation shock, we reduce the standard deviation slightly from the 1.47 percentage points (quarterly, at annual rates) from the estimation, on the grounds that part of our

²¹ See Laxton et al. (1998) for details on this estimation.

estimate reflects omitted influences and measurement errors. We assume a value of 1.29 percentage points.²² For innovations to expectations we use 0.5 percentage points, about half the figure from the estimated model. We vary these assumptions as part of the sensitivity analysis.

We simulate quarterly data for 100 replications of samples of 100 years.²³ Each replication provides hypothetical data for an econometrician, who uses various methods to estimate the Phillips curve and test for convexity. To show the sensitivity of the results to sample size, we report statistics for samples of 100, 200, 300 and 400 quarters. The smallest of these samples, 25 years, is typical of many real-world empirical exercises; it is rare for empirical work on the Phillips curve to have the luxury of 50 years of quarterly data, our second smallest sample size. Many U.S. exercises have samples that fall between these sizes.

We consider three possible data-generating processes (DGPs). Our base case uses a policy rule that is relatively successful in stabilizing the business cycle.²⁴ We then make the economy much more convenient for the econometrician by raising significantly the signal-to-noise ratio in the Phillips curve by reducing the standard deviation of the inflation shock and raising the standard deviation of the demand (unemployment) shock. Finally, we consider the case where the Fed follows a rule that results in ongoing boom and bust cycles that are of the same order of magnitude as in the 1970s and early 1980s.

5.1. Base-case DGP: Policymakers care about stabilization objectives

In the base case, the Fed places twice the weight on unemployment gaps as on inflation gaps. The first and second moments that emerge with this reaction function are reported in the upper panel of Table 2.²⁵ As the estimated Phillips curve has modest convexity, this calibration of the Taylor rule is fairly effective

²² This is roughly the residual standard error reported by Gordon (1997) from his model with numerous dummy and other variables to capture supply shocks and other special effects.

²³ These simulations were done using the stacked-time algorithm in portable TROLL. See Armstrong et al. (1998) and Juillard et al. (1998) for details on TROLL's algorithms for solving forward-looking models. We actually compute more than 100 years of data to provide for a start-up period and to eliminate any problems with terminal conditions.

²⁴ We set the target rate of inflation at 5% in these simulations. This removes a potential difficulty in dealing with the Summers effect – the zero floor on nominal interest rates – which is not germane here. The average inflation rate ends up being somewhat higher than the target. Some experimentation with alternative reaction functions that eliminated this slight 'inflation bias' did not reveal any discernible effects on the results.

²⁵ The table reports the average values, across the 100 replications, of the statistics computed from the full 100 yr of data in each replication. Note that a 'standard deviation' is the estimated second moment for the variable, not the standard deviation for the estimated mean.

Table 2
Alternative data generating processes

Case 1: Policy rule that stabilizes the business cycle

$$rs_t - \pi 4_t = rreq + 0.5[\pi 4_t - \pi_t^*] + 1.0 ugap_t$$

True DGP has convex Phillips curve with $\phi_t = \text{Max}(0, u_t^* - 4)$

	u	u^*	π	π^c	rs	ε^u	ε^π	ε^{π^c}
Mean	5.06	5.02	5.34	5.32	7.96	0.00	0.00	0.00
Standard deviation	0.82	0.66	2.08	1.63	2.61	0.17	1.29	0.50

Case 2: Lower noise ratio

	u	u^*	π	π^c	rs	ε^u	ε^π	ε^{π^c}
Mean	5.08	5.02	5.32	5.23	7.92	0.00	0.00	0.00
Standard deviation	0.94	0.66	1.83	1.59	2.66	0.34	0.65	0.50

Case 3: Regime switching rule that generates boom-and-bust cycles

Inflation fighting regime: $rs_t - \pi 4_t = rreq_t + 0.5[\pi 4_t - \pi_t^*] + 1.0 ugap + 4$

High inflation regime: $rs - \pi 4 = rreq + ugap - 4$

	u	u^*	π	π^c	rs	ε^u	ε^π	ε^{π^c}
Mean	5.36	5.02	11.82	11.81	15.59	0.00	0.00	0.00
Standard deviation	1.56	0.66	3.80	3.30	7.27	0.17	1.29	0.50

at stabilizing the cycle in the simulations, and there is only a small deviation between the average unemployment rate and the average DNAIRU.

5.2. Alternative DGP: A substantially lower noise ratio

For the first alternative DGP, we reduce the standard deviation of the inflation shock by half and double the standard deviation of the unemployment rate shock. The results are reported in middle panel of Table 2. We refer to this case as ‘lower noise ratio’ because a much larger proportion of the variation in inflation is being driven by variation in the rate of unemployment as opposed to inflation shocks. Note, however, that while the overall variability in unemployment is higher, the increase is proportionally much smaller than the change to the standard deviation of the shock itself. This case highlights how severe the problems are with traditional methods for detecting evidence of convexity; these assumptions are unrealistically distorted in favor of increasing the power of the traditional tests.

5.3. Alternative DGP: A policy rule that generates boom-and-bust cycles

For the third case, reported in the third panel of Table 2, we use a different policy reaction function – designed to generate boom-and-bust cycles of approximately the same order of magnitude as was observed during the 1970s and early 1980s. The basic motivation is based on recent interpretation of historical monetary policy errors in the United States. The argument is that policymakers took risks with allowing the inflation rate to rise because they believed that there were small consequences associated with overheating. Indeed, there is some evidence that policy during the excessive monetary expansion and accommodation in the 1970s may have been based on a view that there was a long-run trade-off between inflation and unemployment – see De Long (1997), Romer and Romer (1997) and Taylor (1998). Looking back with the benefit of hindsight, Taylor (1998) and others have described the 1970s as a period where real interest rates were adjusted slowly in response to the buildup of inflationary pressures. Taylor (1998) also argues that the great disinflation in the early 1980s was a period where monetary policy was systematically too tight and that this was the main factor that led to the largest recession in post-war history, where unemployment increased a full 3 percentage points from 7.4% in 1980q4 to 10.4% in 1982q4.

In order to replicate boom and bust cycles of this order of magnitude, we use a simple regime-switching rule, where monetary policy is systematically too loose for 12 quarters and then becomes systematically too tight for 12 quarters. The first regime we describe as a *high inflation regime* because policymakers attempt to reduce unemployment by keeping the real interest rate systematically below the equilibrium real interest rate, which results in higher inflation. This is implemented by setting the parameter on the excess inflation term in the Taylor rule, to zero and subtracting a constant, k , to induce systematic negative bias on interest rate settings.

$$\text{High Inflation Regime: } rs_t - \pi 4_t = rreq_t + ugap - k.$$

The second regime is an *inflation fighting regime* implemented by adding a constant k to the stabilization rule to induce systematic positive bias to real interest rates so that monetary policy has a tendency to be too tight.

$$\begin{aligned} \text{Inflation Fighting Regime: } & rs_t - \pi 4_t \\ & = rreq_t + 0.5[\pi 4_t - \pi_t^*] + 1.0ugap + k. \end{aligned} \quad (12)$$

The value of 4 was chosen for k because this generates an average boom-and-bust cycle with a swing in the unemployment rate of about 3.4 percentage points from trough to peak – just a bit larger than the rise in unemployment during the 1980–1982 episode. It is in this sense that our monetary policy rule generates

boom-and-bust cycles that are representative of the large historical episodes that have been attributed to monetary policy errors.

The average level of inflation rises from 5.34% in the baseline to 11.81% under this policy rule and the switching back and forth between the high and low inflation regimes results in an increase in the standard deviation on inflation from 2.08 to 3.80. This policy rule increases the standard deviation of u to 1.56 from 0.82 in the base case, and this has the effect of raising the average u to 5.36% from 5.06 in the base case. Thus, even modest convexity in the Phillips curve can imply significant first-order welfare effects as the average unemployment rate rises almost 0.3 percentage points relative to the world where policymakers attempt to stabilize the business cycle.²⁶

5.4. Traditional tests for asymmetry

Some recent tests for asymmetry in the Phillips curve (Gordon, 1997; Eisner, 1997) are a special case of the simple estimation and testing procedure described in Clark et al. (1996), henceforth CLR. The strategy suggested by CLR, is to estimate a piecewise linear approximation of a general convex function, which involves adding a separate term for the measure of excess demand when the latter is positive:

$$\pi_t = \lambda \pi_t^c + (1 - \lambda) \pi_{t-1} + \beta ugap^* + \gamma posugap^* + \varepsilon_t^\pi, \quad (13)$$

where $ugap^* = u^* - u$, $u^* = \bar{u} - \alpha$, where \bar{u} is the average value of u , and where $posugap^*$ contains the positive values of $ugap^*$.

In the CLR methodology, α is treated as a parameter to be estimated. We require $\alpha > 0$ for consistency with a convex form. An advantage of this approach is that it is easy to test the restriction to linearity, which requires that α and γ both be zero. It is also an advantage that this approach enables us to look directly at the results of the Gordon–Eisner methodology for testing for non-linearity, which is the same as CLR with α constrained to zero.

Consider, first, the results in Table 3, where we report results with Gordon's (1997) methodology, whereby the NAIRU is estimated simultaneously with the Phillips curve, under the presumption of linearity.²⁷ Then, a test is conducted, under the assumption that α is zero, as to whether the positive gaps enter the Phillips curve significantly, holding the gaps at the values inferred from the estimation of the linear model. The first row shows the percentage of replications in which the coefficient on the positive gaps, γ , is significant, based on

²⁶ It is interesting to note that our historical estimates imply an α shift of the same 0.3 percentage points.

²⁷ We are using Gordon's methodology only with respect to measuring the unemployment gap, not his 24 quarters of lags. Here, we assume that the econometrician knows the true expectations from the DGP. We consider the expectations issue later in the paper.

Table 3
Properties of traditional tests for asymmetry

Gordon's time-varying NAIRU method creates estimates for a linear Phillips curve				
Estimated equation				
$\pi_t = \lambda \bar{\pi}_t^c + (1 - \lambda) \pi_{t-1} + \beta \text{ugap}^* + \gamma \text{posugap}^* + \varepsilon_t^a,$ $\text{ugap}^* = \text{NAIRU} - \alpha - u$				
Percentage of draws where restriction is rejected				
	25 yr	50 yr	75 yr	100 yr
<i>Case 1: Policy rule that stabilizes the business cycle</i>				
<i>t</i> -Test $\gamma = 0 \mid \alpha = 0$	14	16	16	24
<i>t</i> -Test $\gamma = 0 \mid \alpha = \hat{\alpha}$	9	13	19	28
<i>F</i> -test $\gamma = \alpha = 0$	18	14	24	26
<i>Case 2: lower noise ratio</i>				
<i>t</i> -Test $\gamma = 0 \mid \alpha = 0$	43	61	74	84
<i>t</i> -Test $\gamma = 0 \mid \alpha = \hat{\alpha}$	30	64	78	89
<i>F</i> -test $\gamma = \alpha = 0$	44	69	87	92
<i>Case 3: Policy rule that generates boom-and-bust cycles</i>				
<i>t</i> -Test $\gamma = 0 \mid \alpha = 0$	13	31	35	49
<i>t</i> -Test $\gamma = 0 \mid \alpha = \hat{\alpha}$	40	71	90	97
<i>F</i> -test $\gamma = \alpha = 0$	36	70	86	92

a simple *t*-test, at the 5% significance level. The second row has the same statistics for the case where the CLR method is used and α is estimated along with the other parameters. The third row shows the percentage of replications where the restriction that both α and γ are zero, which is the linear model, is rejected at the 5% significance level using an *F*-test.

In our base case, when linearity is presumed in estimating the NAIRU, subsequent tests for convexity have very low power, even in samples as large as 100 yr. The results are much worse than using a coin toss to decide, even in very large samples. The CLR α correction does not help much, in this case, even in large samples.

As shown in the middle panel, when we lower the noise ratio, the presumption of linearity still resolutely confounds inference in typical samples. With 25-year samples, the test is still less powerful than a coin toss. However, the econometric test does beat the coin toss, on average, when 50 yr of data have accumulated, and the power of the test rises significantly with larger samples, especially if the α correction is made.

The third panel shows that Gordon's test still does worse than a coin toss in samples as large as 100 yr, even if the Fed allows boom-and-bust cycles to develop. Here, however, the CLR α correction does appear to systematically improve the power of the tests (although not enough to beat a coin toss in small samples).

Table 3 features time-varying estimates of the NAIRU, determined simultaneously with the Phillips curve parameters. This is a recent innovation. The traditional approach is to prefilter unemployment to obtain estimates of the NAIRU and the unemployment gaps. In Table 4, we report the results obtained when the Hodrick–Prescott filter is used in this way.²⁸ Again, traditional tests have low power to uncover the convexity in typical samples. Using CLR's α correction does not improve the power of the tests much when policy is successful in stabilizing the cycle, but it does when there are boom-and-bust cycles in the data.

5.5. Better tests for asymmetry

While the results in Table 3 have model-consistent estimates of the NAIRU, they are estimates based on the presumption of linearity, and the results are then used in subsequent tests for omitted asymmetry effects. This is problematic, because it ignores an important implication of the convex model in a test against the linear alternative (that the NAIRU must lie above the DNAIRU). We now apply the methodology we used for the historical estimation reported in Section 4, which allows for the possibility of convexity and derives time-varying estimates of the DNAIRU consistent with this possibility. For these experiments, we give the econometrician the correct functional form for the Phillips curve, i.e., the one actually used in the DGP, but with varying degrees of information on ϕ (the minimum unemployment rate).²⁹

The first entry in Table 5 provides a benchmark. Here, the econometrician is assumed to know the DNAIRU, the true value of ϕ , and expectations. Even in this unrealistic case, it is not certain that the econometrician will identify the convex truth with 25 yr of data. The 'true' model fits better in 91% of the replications, but the restrictions to linearity (likelihood ratio test, 5% significance level) are rejected in just 83% of the trials. However, when the sample size is increased to 50 yr, the truth is consistently recovered.

²⁸ For these results we use the standard smoothing parameter of 1600. The more complete discussion in Laxton et al. (1998) also reports results with a higher smoothing value, as well as results for the case where NAIRU is determined by a quadratic trend. The results for these alternatives are essentially the same as those reported here.

²⁹ We continue to assume that expectations are known. This is an extreme assumption, but it permits us to focus on the importance of uncertainty about the measure of excess demand and the importance of the assumptions maintained in providing estimates of its magnitude.

Table 4
Properties of traditional tests for asymmetry

HP filter ($\lambda = 1600$) method for creating NAIRU				
Estimated equation				
$\pi_t = \lambda \bar{\pi}_t^c + (1 - \lambda)\pi_{t-1} + \beta \text{ugap}^* + \gamma \text{posugap}^* + \varepsilon_t^n$ $\text{ugap}^* = \text{NAIRU} - \alpha - u$				
Percentage of draws where restriction is rejected				
	25 yr	50 yr	75 yr	100 yr
<i>Case 1: Policy rule that stabilizes the business cycle</i>				
<i>t</i> -Test $\gamma = 0 \mid \alpha = 0$	12	15	17	23
<i>t</i> -Test $\gamma = 0 \mid \alpha = \hat{\alpha}$	11	23	31	37
<i>F</i> -test $\gamma = \alpha = 0$	10	20	25	32
<i>Case 2: lower noise ratio</i>				
<i>t</i> -Test $\gamma = 0 \mid \alpha = 0$	39	63	77	80
<i>t</i> -Test $\gamma = 0 \mid \alpha = \hat{\alpha}$	40	68	76	91
<i>F</i> -test $\gamma = \alpha = 0$	48	72	81	93
<i>Case 3: Policy rule that generates boom-and-bust cycles</i>				
<i>t</i> -Test $\gamma = 0 \mid \alpha = 0$	33	64	78	85
<i>t</i> -Test $\gamma = 0 \mid \alpha = \hat{\alpha}$	45	81	95	98
<i>F</i> -test $\gamma = \alpha = 0$	46	82	93	97

We now consider what happens when we remove the knowledge of the true value of ϕ . In the next block, the econometrician simply assumes $\phi = 0$. While the true convex function still generally fits the data better, there is a dramatic decline in the success of a classical test to discover the truth from 25 yr of data. In only 43% of the replications does the test reject the false linear alternative at the 5% significance level. While using model-consistent estimates of the DNAIRU raises the power of the test in small samples, relative to the results in Tables 3 and 4, that power remains low. Even with 50 yr of data, the test fails to find the truth in 8% of the replications. In the final block, the econometrician uses the same ϕ rule that we use in the DGP. The results are essentially the same.

In Table 6, we study what happens when the econometrician estimates a version of Gordon's (1997) linear model, where lags of inflation are used to capture expectations. Here, we meld the expectations and intrinsic dynamics of Eq. (1), as is necessary with this approach; the separate effects of expectations and intrinsic dynamics cannot be identified. As in the historical estimation, we use lags of up to 6 yr (24 quarters), imposing the restriction that the sum of the

Table 5
Inference with model-consistent NAIRU Estimates

Base case: The econometrician knows and uses the true inflation expectations
 Methodology: The econometrician generates model-consistent estimates of the DNairu
 The econometrician is assumed to know the second moment of the DNairu distribution (i.e., the variance of the u^* shocks)

Estimated equations

Convex (C) model: $\pi_t = \lambda \bar{\pi}_t^c + (1 - \lambda)\pi_{t-1} + \gamma(u_t^* - u_t)/(u_t - \phi_t) + \varepsilon_t^\pi$
 Linear (L) model: $\pi_t = A(L)\pi_{t-1} + \beta(u_t^* - u_t) + \varepsilon_t^\pi$
 DNairu process: $u_t^* = u_{t-1}^* + \varepsilon_t^{u^*}$

	25 yr	50 yr	75 yr	100 yr
<i>True DNairu and PHI</i>				
% of draws C model fits better	91	100	100	100
% of draws where L model is rejected	83	100	100	100
<i>Model consistent ($\phi = 0$)</i>				
% of draws C model fits better	89	100	100	100
% of draws where L model is rejected	43	92	98	100
<i>Model consistent (ϕ rule)</i>				
% of draws C model fits better	89	100	100	100
% of draws where L model is rejected	40	91	98	100

coefficients on these lags is 1.³⁰ Aside from the modeling of expectations, the experiments in Table 6 mirror those in Table 5.

If the econometrician knows the DNairu and ϕ , the test is relatively powerful, despite the imprecision on expectations, even with the smallest sample size considered. Literal *knowledge* of the DNairu is critical in this case, however. In the realistic case, where the econometrician must estimate the DNairu, the restriction to linearity is rejected in just 7% of the replications with a 25-year sample, regardless of which approach is taken to determining ϕ , and the rejection rate rises only modestly as the sample size is increased to 100 yr. Even with 100 yr of data, the test result is notably worse than deciding by a coin toss. Note, also, that the fit of the true model is not systematically better than that of the false linear model, and not just in small samples. Indeed, the false model fits better than the true model in close to half of replications with 100-year samples, when the econometrician simply assumes that $\phi = 0$. The

³⁰ We continue to impose Gordon’s identifying restriction to limit the number of free coefficients on the lags. There are 6 free parameters, one for each year in the lag distribution.

Table 6
Inference with model-consistent NAIRUS

Base case: The econometrician estimates Gordon’s 24-quarter distributed lag model
 Methodology: The econometrician generates model-consistent estimates of the DNAIURUS
 The econometrician is assumed to know the second moment of the DNAIURU distribution (i.e., the variance of the u^* shocks)

Estimated equation

Convex (C) model: $\pi_t = \lambda\pi_t^c + (1 - \lambda)\pi_{t-1} + \gamma(u_t^* - u_t)/(u_t - \phi_t) + \varepsilon_t^\pi$
 Linear (L) model: $\pi_t = A(L)\pi_{t-1} + \beta(u_t^* - u_t) + \varepsilon_t^\pi$
 DNAIURU process: $u_t^* = u_{t-1}^* + \varepsilon_t^{u^*}$

	25 yr	50 yr	75 yr	100 yr
<i>True DNAIURU and PHI</i>				
% of draws C model fits better	98	100	100	100
% of draws where L model is rejected	98	100	100	100
<i>Model consistent ($\phi = 0$)</i>				
% of draws C model fits better	61	59	57	53
% of draws where L model is rejected	7	22	28	33
<i>Model consistent (ϕ rule)</i>				
% of draws C model fits better	59	62	64	62
% of draws where L model is rejected	7	22	25	37

extra flexibility of the estimated lag parameters to adapt to the idiosyncracies of particular samples masks the truth to a substantial degree.

The results in Tables 5 and 6 are generated using the base-case calibration of the Taylor rule. If the data contained boom-and-bust cycles of the sort produced by our regime-switching formulation, one would expect the tests to be more powerful. This is indeed the case. For example, with the ϕ rule, the linear model is rejected in 66% of the replications with the short sample, compared with the 40% shown in Table 5. With the boom-and-bust formulation, the lags approach of Table 6 also does better, but remains systematically less powerful. It is striking that even with these extreme policy assumptions and the consequently much more volatile data, it is not at all certain that the convex truth will be recovered from typical samples, regardless of what methodology is used.

6. The implications of aggressive ‘probing’

The NAIRU cannot be determined precisely. It was not long ago that 6% was seen as a reasonable assumption for the United States, but estimates have been falling as inflation has remained basically stable while the actual rate of

unemployment has declined. Uncertainty about the NAIRU has important implications for the design of monetary policy; policymakers who care about unemployment may want to probe the economy's short-term capacity limits to minimize the deadweight losses incurred by operating with excess supply. However, the extent of such experimentation that is prudent will depend on the functional form of the Phillips curve as well as the preferences of policymakers.

Experimentation has been studied in the context of the linear model by Wieland (1998), who argues that uncertainty about the NAIRU creates a logic for modest experimentation.³¹ Wieland assumes quadratic preferences, where policymakers dislike reductions in the unemployment below the NAIRU as much as they dislike increases in unemployment above the NAIRU. If the political process imposed high discounting of the future or policymakers had asymmetric loss functions, the extent and type of 'optimal' experimentation could change dramatically. For example, if policymakers attached *benefits* to unemployment below the NAIRU, instead of Wieland's costs, there may be an incentive for policymakers to attempt to push unemployment down aggressively, even if they were convinced that it was already below the NAIRU. The resulting variation in the cycle would help improve econometric estimates, which would provide longer-term gains at little cost, since, at least in the popular 'integral gap' (IGAP) linear model (Summers, 1988), the integral of the excess demand gap is always the same.

In this section, we 'probe' the policy implications of the alternative views about the Phillips curve based on the estimated equations presented in Section 4. Fig. 2 provides some illustrative simulations of the IGAP model and the model based on the convex Phillips curve. The experiment in both cases consists of an aggressive probe where the Federal Funds rate is cut by 300 basis points for 8 quarters, after which the base-case policy rule operates. This is a deterministic analysis, so the NAIRU is equal to the DNAIUR, which is set at 5%. For the IGAP model, the unemployment rate declines by more than one percentage point at the peak of the boom and 4-quarter inflation peaks at about 2.2 percentage points above control. However, the secondary contraction that is necessary to eventually bring inflation back to the target level is quite modest and it would be difficult to argue that there are serious real costs to aggressive probing in this case, especially if policymakers prefer unemployment to be below the NAIRU and place a small weight on inflation rising persistently above the target. Indeed, while the IGAP model imposes a restriction that the cumulative effect on unemployment must be zero in the long run – as long as monetary

³¹ The rationale is to improve information and the speed of learning. If good outcomes from subsequently better models outweigh the transitory costs, it may be optimal for a monetary authority to destabilize the economy, for a time, to speed learning.

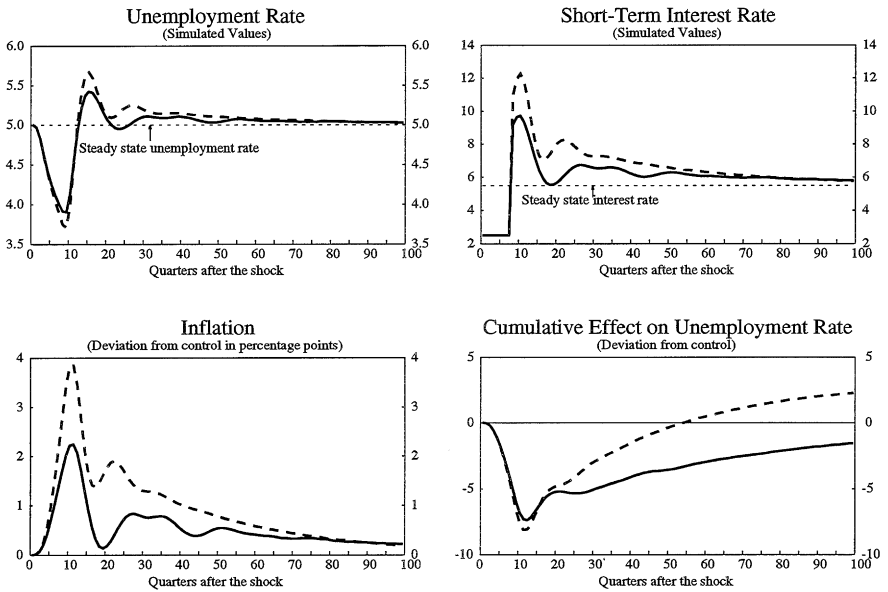


Fig. 2. Model responses to aggressive probing (300 basis points, 8 quarters): (—) integral gap model; (---) convex Phillips curve model.

policy ensures inflation returns to its target level – Fig. 2 shows that it remains negative for over 100 periods.³²

The same cannot be said for the simulations where we use our estimated convex form. Inflation peaks at almost 4 percentage points above control, and the subsequent monetary action necessary to bring it back to the target rate is severe; the Federal Funds rate surpasses 12% at its peak, over 650 basis points above the control level. Moreover, unemployment rises sharply in the secondary cycle, and there is a long period of excess supply – the cumulative effect on unemployment approaches +2 percentage points. In other words, the short-term gains from lower unemployment are substantially more than offset over the longer term.³³

Because the estimated nonlinear model is approximately linear in the region of equilibrium, the predictions of the two models are very similar for experiments that consider modest probing.³⁴ The lesson from these experiments is

³² The cumulative effect is a simple sum of the shock-control values of the unemployment gap. If we extend the horizon, the cumulative effect approaches zero.

³³ Recall that this is a deterministic analysis, where the NAIRU is unaffected by the experiment. In a stochastic setting, the extra volatility would also act to raise the NAIRU.

³⁴ The results of simulations with modest probing are not reported here but are available in the working paper from which this article is drawn (Laxton et al., 1998), which also discusses the implications of the ‘concave’ Phillips curve model with respect to probing.

clear. Modest convexity means little if policy acts to keep the economy in the region of equilibrium, and this conclusion extends to experiments with modest probing to test the short-term capacity limits of the economy.³⁵ However, any attempt to go beyond gentle probing to encompass long and protracted attempts to hold unemployment well below the NAIRU are another matter. Modest convexity bounds the limits of probing below the NAIRU, before the costs of error become large.

7. Conclusions

In this paper, we ask what form the Fed should presume for the U.S. Phillips curve, given conflicting empirical claims that it is ‘resolutely linear’ (Gordon, 1997), perhaps concave (Eisner, 1997; Stiglitz, 1997), and convex (Turner, 1995; Clark et al., 1996; Akerlof et al., 1996). Our answer is that the Fed should assume the traditional convex form. Our case has several parts.

The form of the Phillips curve has important implications for the conduct of monetary policy. If the Phillips curve is linear, positive and negative shocks to demand will have equal effects on inflation and the overall effect will average to zero, regardless of the response of monetary policy. Moreover, the timing of any monetary response will be of little consequence to the final outcome. Thus, in the linear world there is little incentive to move early to combat inflationary pressures, and every incentive to temporize with inflation by aggressively probing to ensure that the NAIRU has not been overestimated.

By contrast, in the case of an asymmetric (convex) Phillips curve, positive shocks to demand raise inflation to a greater extent than negative shocks of the same magnitude lower it. This property implies that early action to counteract emerging inflation pressures can reduce the need to take stronger disinflationary action later. Moreover, to the extent that a prompt monetary policy response can succeed in stabilizing employment, it will also *lower the average rate of unemployment*. Thus, in the convex world, success at stabilizing the cycle will generate first-order welfare gains – the NAIRU will be lower. Thus, the goal of a prudent monetary authority in a convex world will be to

³⁵ It is important to note that this analysis embodies very strong assumptions about monetary policy credibility and this tends to bias the results in favor of experimentation. Using an open-economy model with a convex Phillips curve and endogenous policy credibility, Isard and Laxton (1998) show that there can be significantly larger inflation costs from aggressive experimentation that attempts to minimize the average level of unemployment. For a discussion of alternative linear and nonlinear models of the inflation process with NAIRU uncertainty and endogenous policy credibility, see Isard et al. (1998).

avoid large cyclical excesses. Modest probing for the limits of productive capacity will have minor costs, but errors that lead to serious overheating will have major costs.

The concave world is quite different. The incentives become slanted strongly towards strong action – rapid disinflation, aggressive probing for the limits of capacity, and so on, because the consequences of overshooting decline on the margin. Moreover, in this case, *higher* cyclical volatility has the effect of lowering the NAIRU.

These diametrically different visions of the role of the prudent central banker make the tasks of assessing the empirical evidence and of establishing an appropriate research methodology especially important. Unfortunately, the task is not straightforward. The data cannot discriminate clearly among the competing forms. We show that a Phillips curve with traditional modest convexity fits the U.S. data as well or better than the linear alternative, but that no conclusive statistical case can be made that convexity is necessary to explain the data. However, the results are entirely consistent with the presence of modest convexity. The test results depend entirely on which hypothesis is treated as the null.

We also assess the power of econometric tests that seek to identify the presence of asymmetry in the Phillips curve. The identification problem is particularly severe for work on the Phillips curve because two key variables, inflation expectations and the degree of excess demand, are not directly observable. We show through Monte Carlo experiments that the data have little power to discriminate between alternate forms in typical samples, the more so if policy is successful in stabilizing the economy and avoiding observations of extreme outcomes. The root cause of this low power is that measures of the degree of excess demand are quite imprecise. Even when we assume precise knowledge of inflation expectations, the uncertainty surrounding the degree of excess demand confounds the econometrician and limits the power of tests for convexity in the Phillips curve.

Our conclusion is that standard empirical techniques are not likely to be capable of providing a reliable answer on functional form. Since the issue is of great importance for monetary policy, research leading to policy advice and policy action must be based on broader criteria. The final piece in our case for presuming convexity is that the costs of errors of incorrect presumption are far from symmetric. If there is convexity in the Phillips curve, but policy is based on the presumption of linearity, or worse, concavity, the consequences can be severe. Policy errors that lead to relatively severe overheating will be costly to correct, and the data will be characterized by boom-and-bust cycles with deep and protracted recessions. If the other error is made, that is, policy is based on the presumption of convexity when the truth is less challenging, there will be deadweight losses, but these will be relatively minor.

We conclude, therefore, that there is a clear and compelling case that the Fed should operate on the presumption that the Phillips curve has modest convexity.

8. For further reading

The following references are also of interest to the reader: Laxton et al. (1994) and Wieland (1996).

Acknowledgements

The views expressed in this paper are those of the authors and do not necessarily reflect those of the International Monetary Fund. We would like to thank Victoria Ejikeme, Dirk Muir, Susanna Mursula, and Ken Ling Hwang for their assistance. We would also like to thank, without implicating, Peter Clark, Hamid Faruquee, Josh Felman, Peter Isard and Volker Wieland for their comments, as well as two anonymous referees.

References

- Akerlof, G., Dickens, W., Perry, G., 1996. The macroeconomics of low inflation. *Brookings Papers on Economic Activity* 27, 1–26.
- Armstrong, J., Black, D., Laxton, D., Rose, D., 1998. A robust method for simulating forward-looking models. *Journal of Economic Dynamics and Control* 22, 489–501.
- Bryant, R.C., Hooper, P., Mann, C.L., 1993. *Evaluating Policy Regimes: New Research in Empirical Macroeconomics*. The Brookings Institution, Washington.
- Chadha, B., Masson, P., Meredith, G., 1992. Models of inflation and the costs of disinflation. *IMF Staff Papers* 39, 395–431.
- Clarida, R., Gali, J., Gertler, M., 1997. Monetary policy rules in practice: some international evidence, New York University, C.V. Starr Center for Applied Economics, Economic Research Report Number 97–32.
- Clark, P.B., Laxton, D., 1997. Phillips curves, phillips lines and the unemployment costs of overheating. *IMF Working Paper 97/17*, International Monetary Fund, Washington February.
- Clark, P.B., Laxton, D., Rose, D., 1997. Capacity constraints, inflation, and the transmission mechanism: Forward-looking versus myopic policy rules. *IMF Working Paper 95/75*, revised 1997, International Monetary Fund, Washington.
- Clark, P.B., Laxton, D., Rose, D., 1996. Asymmetry in the U.S. output-inflation nexus: Issues and evidence, *IMF Staff Paper 43*, International Monetary Fund, Washington.
- Debelle, G., Laxton, D., 1997. Is the phillips curve really a curve? Some evidence for Canada. The United Kingdom, and the United States, *IMF Staff Papers 44*, International Monetary Fund, Washington.
- De Long, J.B., 1997. *America's Peacetime Inflation: The 1970s, Reducing Inflation: Motivation and Strategy*. The University of Chicago Press, Chicago.

- Eisner, R., 1997. A new view of the NAIRU. In: Davidson, P., Kregel, J., (Eds.), *Improving the Global Economy: Keynesianism and the Growth in Output and Employment*. Edward Elgar, Cheltenham U.K. and Brookfield (original manuscript, Northwestern University, 1995) forthcoming.
- Faruqee, H., Laxton, D., Rose, D., 1997. Inflation and Unemployment in Europe and North America: Asymmetry Versus Hysteresis. International Monetary Fund, Washington, presented at the CEPR Conference on Unemployment Persistence and the Long Run: Re-evaluating the Natural Rate, Vigo, Spain.
- Friedman, M., 1968. The role of monetary policy. *American Economic Review* 58, 1–17.
- Fuhrer, J., 1996. The Phillips curve is alive and well. *New England Economic Review*, 41–56.
- Fuhrer, J.C., Moore, G.R., 1995. Monetary policy trade-offs and the correlation between nominal interest rates and real output. *American Economic Review* 85, 219–239.
- Gordon, R.J., 1997. The time-varying NAIRU and its implications for economic policy. *The Journal of Economic Perspectives* 11, 11–32.
- Isard, P., Laxton, D., 1996. Strategic choice in Phillips curve specification: what if Bob Gordon is wrong? presented at a Conference on European Unemployment: Macroeconomic Aspects, Florence, Italy.
- Isard, P., Laxton, D., 1998. Monetary policy and uncertainty in the unemployment-inflation process: how aggressively should policymakers experiment to find the NAIRU? IMF Working Paper, forthcoming.
- Isard P., Laxton, D., Ann-Charlotte Eliasson, 1998. Inflation targeting with NAIRU uncertainty and endogenous policy credibility. presented at the 4th Conference on Computational Economics, Cambridge, United Kingdom.
- Juillard, M., Laxton, D., McAdam, P., Pioro, H., 1998. An algorithm competition: first-order iterations versus Newton-based techniques. *The Journal of Economic Dynamics and Control*, forthcoming.
- Kuttner, K., 1992. Monetary policy with uncertain estimates of potential output, Reserve Bank of Chicago, *Economic Perspectives* 2–15.
- Kuttner, K., 1994. Estimating potential output as a latent variable. *Journal of Business and Statistics* 12, 361–368.
- Laxton, D., Meredith, G., Rose, D., 1995. Asymmetric effects of economic activity on inflation: evidence and policy implications. IMF Staff Papers 42, pp. 344–374.
- Laxton, D., Ricketts, N., Rose, D., 1994. Uncertainty, learning and policy credibility. *Economic Behavior and Policy Choice Under Price Stability*. Bank of Canada, Ottawa.
- Laxton, D., Rose, D., Tambakis, D., 1998. The U.S. Phillips curve: the case for asymmetry. IMF Working Paper, International Monetary Fund, Washington, forthcoming.
- Laxton, D., Rose, D., Tetlow, R., 1993. Monetary policy, uncertainty and the presumption of linearity. Technical Report No. 63, Bank of Canada, Ottawa.
- Layard, R., Nickell, S., Jackman, R., 1991. *Unemployment: Macro-Economic Performance and the Labour Market*. Oxford University Press, Oxford.
- Lipsy, R.G., 1960. The relation between unemployment and the rate of change of money wage rates in the United Kingdom 1862–1957: a further analysis. *Economica* 27, 1–31.
- Macklem, T., 1996. Asymmetry in the monetary transmission mechanism: what can we learn from VARs? Bank of Canada, unpublished.
- Orphanides, A., Wilcox, D., 1996. The opportunistic approach to disinflation. Discussion Paper 96-24, Board of Governors of the Federal Reserve System, Washington.
- Orphanides, A., Small, D., Wieland, V., Wilcox, D., 1996. A quantitative exploration of the opportunistic approach to disinflation. Unpublished paper, Board of Governors of the Federal Reserve System, Washington.
- Phillips, A.W., 1958. The relationship between unemployment and the rate of change of money wage rates in the United Kingdom, 1861–1957. *Economica* 25, 283–300.

- Romer, C.D., D.H. Romer, 1997, Institutions for Monetary Stability, Reducing Inflation: Motivation and Strategy. University of Chicago Press, Chicago.
- Summers, L. 1988. Should Keynesian economics dispense with the Phillips curve?. In: Rod Cross (Ed.), Unemployment, Hysteresis and the Natural rate Hypothesis. Basil Blackwell, Oxford, pp. 11–25.
- Stiglitz, J., 1997. Reflections on the natural rate hypothesis. *Journal of Economic Perspectives* 11, 3–10.
- Taylor, J., 1993. Discretion versus policy rules in practice. *Carnegie–Rochester Conference Series on Public Policy* 39, 195–214.
- Taylor, J., 1998. An historical analysis of monetary policy rules. Paper presented at an NBER Conference at the Cheeca Lodge, Islamorada, Florida, 17 January, 1988.
- Turner, D., 1995. Speed limit and asymmetric effects from the output gap in the seven major countries. *OECD Economic Studies* 24, 57–88.
- Wieland, V., 1996. Monetary policy and uncertainty about the natural unemployment rate. *Finance and Economics Discussion Series*, 98–104.