

NBER WORKING PAPER SERIES

PRODUCTIVITY GROWTH AND THE PHILLIPS CURVE

Laurence Ball
Robert Moffitt

Working Paper 8421
<http://www.nber.org/papers/w8421>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2001

This paper is part of the Sustainable Employment Initiative of The Century and Russell Sage Foundations. We are grateful for research assistance from Cristian de Ritis, Daniel Leigh, Kevin Moore, Yelena Takhtamanova, Robert Tchaidze, Gergana Trainor, and Huiyan Zhang. We received helpful suggestions from Martin Baily, Susanto Basu, Hasan Bakhshi, Francisco Gallego, N. Gregory Mankiw, Stephen Nickell, Robert Solow, David Weiman, and many seminar participants. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2001 by Laurence Ball and Robert Moffitt. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Productivity Growth and the Phillips Curve
Laurence Ball and Robert Moffitt
NBER Working Paper No. 8421
August 2001
JEL No. E31, E32, J60

ABSTRACT

We present a model in which workers' aspirations for wage increases adjust slowly to shifts in productivity growth. The model yields a Phillips curve with a new variable: the gap between productivity growth and an average of past wage growth. Empirically, this variable shows up strongly in the U.S. Phillips curve. Including it explains the otherwise puzzling shift in the unemployment-inflation tradeoff since 1995.

Laurence Ball
Department of Economics
Johns Hopkins University
Baltimore, MD 21218
and NBER
lball@jhu.edu

Robert Moffitt
Department of Economics
Johns Hopkins University
Baltimore, MD 21218
and NBER
moffitt@jhu.edu

I. INTRODUCTION

The "New Economy" in the U.S. since the mid-1990s has featured surprisingly benign behavior of inflation and unemployment. Before this experience, most estimates of the NAIRU -- the non-accelerating inflation rate of unemployment -- were in the neighborhood of six percent. Yet unemployment has fallen far below this level, reaching 4.2% in 2000, and inflation has not risen substantially. This paper presents an explanation for the apparent improvement in the unemployment-inflation tradeoff. We argue that it is caused by another feature of the new economy: the rise in the growth rate of labor productivity.

Our argument builds on an old idea: workers' wage aspirations adjust slowly to shifts in productivity growth. As a result, such shifts produce periods when aspirations and productivity are out of line, causing the Phillips curve to shift. Authors such as Grubb et al. (1982) use this idea to argue that the productivity slowdown of the 1970s caused an unfavorable Phillips-curve shift. Authors such as Blinder (2000) and Council of Economic Advisors (2000) suggest that this process worked in reverse in the late 1990s, with a productivity speedup causing a favorable Phillips-curve shift. This paper presents new evidence that changes in productivity growth do indeed affect the Phillips curve. In addition to documenting this idea in general, we show that it explains most of the Phillips curve puzzle since 1995.

Our argument proceeds in several steps. In Section II, we discuss the ideas about wage determination that underlie our story. We draw on previous research suggesting that concepts of fairness

affect wage setting, and that perceptions of fair wage increases are tied to past wage increases.

Section III embeds these ideas in an otherwise standard model of the Phillips curve. In the model, an increase in productivity growth feeds one-for-one into lower price inflation for given wage inflation. It has less effect on wage inflation, which is determined largely by past wage increases. Wage inflation also depends negatively on unemployment. Combining these assumptions yields a Phillips curve in which the change in inflation depends on unemployment and the difference between current productivity growth and past real-wage growth. Shifts in productivity growth cause shifts in the unemployment-inflation relation for a period while wage aspirations are adjusting.

Section IV discusses the measurement of key variables in our model, and Section V presents our central empirical results. We estimate alternative Phillips curves with annual U.S. data from 1962-1995, and then use these equations to forecast inflation over 1996-2000. We first confirm previous findings that a conventional Phillips curve overpredicts inflation after 1995. We then estimate the Phillips curve from our model and find that the new variable -- the gap between productivity growth and past real wage growth -- has the effect predicted by our theory. When this variable is included, the overprediction of inflation since 1995 disappears. Section VI discusses extensions of the analysis, such as the addition of traditional "supply shock" variables to the Phillips curve.

Sections VII and VIII leave aggregate U.S. data to look for

other evidence for our theory. Section VII is a case study of Chile in the 1990s. This episode is another one in which a productivity acceleration appears to have caused a favorable Phillips-curve shift. Section VIII examines micro data from the U.S. Current Population Survey. Here, we show that our model helps explain differences in wage growth across workers as well as movements in aggregate variables.

Section IX concludes the paper.

II. WAGE ASPIRATIONS

It is clear that real wages are tied closely to labor productivity in the long run. Consequently, our model will have the feature that productivity, real wages, and real-wage aspirations all grow at the same rate in a steady state. We consider the possibility, however, that a shift in productivity growth is not matched immediately by a shift in wage aspirations, because these are tied partly to past wage increases. Many authors have suggested ideas along these lines; recent examples include Blanchard and Katz (1997), Stiglitz (1997), Blinder (2000), and DeLong (2000). However, these authors seldom justify their ideas about wage aspirations in much detail. We will not attempt a full theory of aspirations, but we will briefly review some relevant literature.

By "wage aspirations" we mean the real wages that workers consider fair. Our model rests on two assumptions about aspirations: that they affect the actual wages that workers receive, and that they are tied to past wage increases. We

discuss these points in turn.

The assumption that wages depend on what workers consider fair is a departure from neoclassical microeconomics, but one with strong empirical support. Akerlof and Yellen (1990) discuss a likely channel: workers reduce their effort if they perceive wages as unfair, making it in firms' interests to pay fair wages. An experimental literature in psychology (surveyed by Akerlof and Yellen) shows that workers' performance deteriorates when they believe wages are unfair. Management textbooks such as Milkovich and Newman (1996) stress the importance of paying fair wages to elicit effort. Bewley's (2000) field research suggests the similar idea that firms pay fair wages to maintain worker morale.

What wages do workers consider "fair?" The psychology literature suggests that workers judge the fairness of their wages by comparing them to "reference transactions" -- certain wages they have observed in the past (see Kahneman et al., 1986, Oswald, 1986, and Elliot, 1991). Researchers disagree about which wage payments are the reference transactions for a given worker. One possibility is wages paid to the same worker in the past, and another is wages paid to other workers of the same type. When we examine microeconomic data in Section VIII, we will ask whether a worker's wage is more closely tied to his own past wage or to others' wages. However, this distinction is not crucial at an aggregate level. If wage setters base their actions on past wages, aggregating across the economy yields a relationship between current and past wage increases, regardless of whose past

increases are relevant to individuals.¹

III. THE PHILLIPS CURVE AND THE NAIRU

This section embeds our ideas about wage aspirations in a canonical model of wage- and price-setting and derives a Phillips curve. Specifically, the model follows Blanchard and Katz (1997) and Katz and Krueger (1999) except for our treatment of productivity and aspirations.

A. Deriving the Phillips Curve

We denote inflation by π and wage inflation by ω , so real wage growth is $\omega - \pi$. We assume that wage setters have a target for real-wage growth given by

$$(1) \quad (\omega - \pi)^* = \alpha - \gamma U + \delta \theta + (1 - \delta)A + \eta, \quad \alpha, \gamma > 0, \quad 0 \leq \delta \leq 1,$$

where U is unemployment, θ is labor-productivity growth, A is an aspiration wage increase, and η is an error term. This equation makes the conventional assumption that higher unemployment reduces target real-wage growth. The target also depends on an average of productivity growth and the aspiration wage increase, which is given by

$$(2) \quad A = \frac{1 - \beta}{\beta} \sum_{i=1}^{\infty} \beta^i (\omega - \pi)_{-i}$$

¹ Note we assume that ideas about fairness concern wage increases rather than wage levels. This seems natural because, with productivity increases and life-cycle wage growth, workers are accustomed to fairly steady increases rather than steady levels. We have, however, explored a version of our model in which workers care about levels as well as growth rates. In this case, the Phillips curve includes an "error-correction" term, the lagged difference between the levels of productivity and real wages. This variable is never significant in our regressions.

To interpret equations (1)-(2), consider first the special case of $\delta=1$. This is a neoclassical benchmark in which productivity growth feeds one-for-one into wages, and aspirations are irrelevant. At the other extreme of $\delta=0$, productivity is irrelevant and wage increases are based on aspirations. This period's aspiration for a real-wage increase is a weighted average of past increases, with exponentially declining weights. The aspiration real-wage increase can also be written recursively as $A = \beta A_{-1} + (1-\beta)(\omega-\pi)_{-1}$. As this shows, aspirations adjust over time in response to the most recent wage increase. The adjustment is fast if β is small and slow if β is close to one.²

Our model nests the two special cases, allowing both productivity growth and past real-wage growth to influence wage setting. Note we assume that these two variables have coefficients that sum to one. This implies that the target depends one-for-one on productivity growth in a steady state with real-wage growth equal to productivity growth.

Wage setters must choose nominal wages one period in advance. They choose a nominal increase ω equal to their target real wage increase, $(\omega-\pi)^*$, plus expected inflation. Expected inflation equals last period's inflation, π_{-1} . Combining these assumptions with equation (1) yields a "wage Phillips curve":

$$(3) \quad \omega = \alpha + \pi_{-1} - \gamma U + \delta \theta + (1-\delta)A + \eta .$$

² In our empirical work, we have experimented with alternatives to the exponentially-declining weights in equation (2). For example, we have defined A as a simple moving average of past real-wage changes. These

Wage inflation depends on past price inflation, unemployment, and an average of θ and A .

We complete the model with a standard equation for price inflation:

$$(4) \quad \pi = \omega - \theta + v ,$$

where v is another error. Price inflation depends one-for-one on the increase in unit labor costs, which is wage inflation minus productivity growth. Substituting the wage Phillips curve into (4) yields a "price Phillips curve":

$$(5) \quad \pi = \alpha + \pi_{-1} - \gamma U - (1-\delta)(\theta-A) + \varepsilon ,$$

where $\varepsilon = \eta + v$. This Phillips curve will be the centerpiece of our empirical analysis.

B. Discussion

To interpret our Phillips curve, we again start with the case of $\delta=1$: target real-wage increases depend on productivity growth but not on aspirations. In this case, the $\theta-A$ term drops out of (5), and the equation reduces to a conventional Phillips curve. For $\delta=1$, productivity growth has a negative effect on price inflation given wage inflation, but it has a fully-offsetting positive effect on wage inflation. Thus productivity growth has no role in the Phillips curve. Since $\delta=1$ is a natural neoclassical baseline, this result explains why research on the Phillips curve does not usually emphasize productivity growth.

Productivity growth does matter if wage growth is partly tied

variations have little effect on our results.

to past wage growth, i.e. $\delta < 1$. Productivity growth is still irrelevant in a steady state with $\theta = A$. But if productivity growth accelerates or decelerates, A does not adjust immediately, and $\theta - A$ moves in the direction of θ . A productivity acceleration causes a favorable shift in the unemployment-inflation relation and a slowdown causes an unfavorable shift. The shift can last a long time if wage aspirations adjust slowly -- if β is close to one.

While the aspiration variable A can differ from productivity growth, the actual growth of real wages cannot. Inverting the price equation (4) gives a formula for actual real-wage growth: it equals $\theta + v$. In equilibrium, this fact is reconciled with the behavior of wage setters by movements in unemployment or inflation. During a productivity slowdown, target wage growth rises relative to productivity growth for given unemployment, but higher unemployment offsets this effect or accelerating inflation reduces actual real-wage growth below the target. Thus the model captures the stylized fact that U.S. wages are closely tied to labor productivity, as shown by the near-constancy of labor's share of income.

We define the NAIRU in our model as the level of unemployment consistent with stable inflation and $\theta - A = 0$, which must hold in steady state. The NAIRU equals α/γ , the ratio of the constant in the Phillips curve to the unemployment coefficient. If a productivity acceleration raises $\theta - A$ above zero, we will say that

unemployment can fall below the NAIRU temporarily without accelerating inflation, not that the NAIRU itself has fallen. In other words, we treat movements in θ -A as "supply shocks" that shift the unemployment-inflation tradeoff for a given NAIRU.

IV. DATA AND MEASUREMENT

Our measurement of inflation and unemployment follows previous work, especially Blanchard-Katz (1997) and Katz-Krueger (1999). The data are annual. The inflation rate π is the change in the log of the consumer price index, and the wage-inflation rate ω is the change in the log of compensation per hour in the business sector. Unemployment is the unemployment rate for all civilian workers. All of these series are produced by the Bureau of Labor Statistics.

The rest of this section describes construction of the two key variables in our theory: the growth rate of labor productivity, and aspirations for real-wage growth.

A. Measuring Productivity Growth

Our starting point for measuring productivity growth is the change in the log of output per hour in the business sector, from the BLS. As shown below, this series captures both the productivity slowdown of the 1970s and the speedup since 1995. For our present purposes, the reasons for these productivity shifts are not important. For example, we need not take a stand on whether the recent acceleration in productivity reflects rapid TFP growth or capital deepening.

A practical issue in measuring productivity is cyclical adjustment. Output per hour is an imperfect measure of labor productivity because labor input varies through shifts in worker effort as well as measured hours. In particular, productivity growth is overstated in expansions because effort rises. In our underlying theory, price- and wage-setting depend on true rather than measured productivity, so we need to adjust our productivity variable to eliminate the effects of cyclical movements in effort.

Our approach to measuring true productivity follows Basu and Kimball (1996), who build on Bils and Cho (1994). Basu and Kimball assume that, over the business cycle, effort moves proportionately with average weekly hours of employed workers. This relationship follows from a model in which firms costlessly adjust both effort and weekly hours when they need more labor input (but adjusting employment may be costly). Empirically, a close link between effort and weekly hours is supported by time-motion studies that directly measure effort (Schor, 1987). Given this link, we can use variation in weekly hours as a proxy for variation in effort. To purge productivity fluctuations of the part caused by changes in effort, we regress measured productivity growth on the change in the log of weekly hours. We use the residuals from this regression to measure true productivity growth θ , adding a constant to make the mean of θ equal the mean of measured productivity growth.

For 1962-2000, regressing measured productivity growth on the change in log hours yields a coefficient of 0.66. The \bar{R}^2 is only

0.06, however, which means our cyclical adjustment removes a small fraction of productivity fluctuations. As a result, the adjusted and unadjusted series for θ , shown in Figure 1, are not very different. Our results confirm previous findings that measured labor productivity is only mildly cyclical (unlike total factor productivity).³

The series in Figure 1 capture the broad phenomena of the productivity slowdown and the recent acceleration. With cyclically-adjusted data, θ averages 3.3% over 1962-1973, 1.4% over 1974-1995, and 2.7% over 1996-2000. However, these broad trends do not fully explain the data. There is considerable year-to-year variation in productivity growth, even after our cyclical adjustment.

B. Wage Aspirations

The most novel variable in our analysis is A , which determines workers' aspirations for real-wage increases. In each period, A is an exponentially-weighted average of past real-wage increases (equation (2)). Two issues arise in constructing A : the choice of the weighting parameter β , and the need to approximate the infinite sum in the definition of A . We begin with the second issue.

In principle, A depends on real-wage increases back to the infinite past. In practice, our data on real-wage growth start in 1948. To address this problem, we make a reasonable guess of the

³ As this fact suggests, our results below do not change much if we use the unadjusted productivity-growth series. Similarly, changing the coefficient of 0.66 in our procedure does not make much difference.

value of A in 1948. Given this value, we can derive A for 1949, 1950, ... using the recursive definition, $A_t = \beta A_{t-1} + (1-\beta)(\omega - \pi)_{t-1}$. That is, we assume an A in 1948 and update A in each year based on the evolution of real wages.

Specifically, we set A for 1948 equal to trend real-wage growth in that year, as measured by the Hodrick-Prescott filter over 1948-2000 with smoothing parameter 1000. This yields $A=4.2\%$. The implicit assumption is that wage aspirations in 1948 were close to the actual trend in real wages: 1948 was not a time like the 1970s or late 90s when aspirations and actual wage-growth diverged. Fortunately, our results are not very sensitive to the choice of A for 1948, because our regressions use data starting in 1962. The 1948 value of A has a weight of only β^{14} in the A for 1962, and smaller weights in later A's.⁴

The exponential parameter β can in principle be estimated from the data. Our estimates are imprecise, however, and so we end up imposing values that are plausible a priori and not rejected by the data. Figure 2 shows the series for actual real-wage growth from 1948 through 2000 and for A with various values of β . Real-wage growth fluctuates around a trend that is stable until the late 1960s and then declines as a result of the productivity slowdown. For most values of β , A follows the downward trend in real-wage growth with a lag. Real-wage growth

⁴ We add a constant to the series on real-wage growth to make its mean equal the mean of productivity growth. That is, we impose the restriction that there is no trend in labor's share of income. The means of real-wage growth and productivity growth differ in the raw data, mainly because these variables are constructed from price indices with different trends.

rises sharply and aspirations modestly at the end of the sample.

Much of our analysis will focus on the case of $\beta=0.95$. A fairly high β captures Stiglitz's (1997) suggestion that the adjustment of aspirations to the 1970's productivity slowdown continued into the 1990s. Moreover, values of β that are much smaller than 0.95 or very close to one are unappealing. As illustrated in Figure 2, values of 0.8 or below imply that aspirations fluctuate substantially in response to year-to-year movements in real-wage growth. It seems unlikely that concepts of fair wages fluctuate so much. At the other extreme, a β of one implies that workers still want the wage increases they received in the 1950s. In this case, the real-wage growth of the last five years falls short of aspirations, even though it is high compared to the previous 25 years.

For $\beta=0.95$, Figure 3 shows the difference $\theta-A$, the new term that appears in our Phillips curve, for 1962-2000. To isolate long-run trends, the Figure also presents a smoothed version of the series based on the Hodrick-Prescott filter with a parameter of 1000. The recent "New Economy" can be seen in the high values of $\theta-A$ for 1996-2000: the average value for this period is the highest for any five-year period since 1948. $\theta-A$ was high after 1995 because θ rose sharply and A reached low levels after finally adjusting to the productivity slowdown. θ was higher in the 1950s and 60s, but then it was balanced by high wage aspirations.

V. ESTIMATES OF THE PHILLIPS CURVE

This section estimates the Phillips curve from our model, equation (5), with annual U.S. data. We examine the general performance of the equation by estimating it with data from 1962 through 1995. We then perform out-of-sample forecasts to see whether the equation explains inflation in the post-1995 New Economy.

A. A Conventional Phillips Curve

As a benchmark, we first examine a Phillips curve that lacks our new variable θ -A. This is a simple textbook equation: the change in inflation depends on a constant and unemployment. As discussed above, this equation follows from our model if wage growth depends one-for-one on productivity growth and aspirations have no effect.

For 1962-1995 -- the Old-Economy period -- ordinary-least-squares estimation of the Phillips curve yields

$$\Delta\pi = 4.41 - 0.710U, \quad \bar{R}^2=0.34,$$

(1.14) (0.161)

where standard errors are in parentheses. These results look reasonable. One point-year of unemployment reduces inflation by seven tenths of a percent. The NAIRU -- the ratio of the constant to minus the unemployment coefficient -- is 6.2%.

Using these estimates, we next compute forecasts of inflation over 1996-2000, given the actual evolution of unemployment. Figure 4 plots the forecasts along with two-standard-error bands, and compares them to actual inflation. This Figure shows why many

authors have suggested that a New Economy has arrived. Since unemployment falls far below the NAIRU estimate of 6.2%, predicted inflation rises rapidly and reaches 8.3% in 2000. In contrast, actual inflation fluctuates mildly and ends at 3.3%. The overprediction of inflation with a 6.2% NAIRU suggests that the NAIRU has fallen for some reason.

B. The Phillips Curve with θ -A

We now estimate the Phillips curve from our model, equation (5). This is the conventional Phillips curve estimated above with the addition of the term θ -A.

Our modification of the Phillips introduces the parameter β , the weighting factor in the formula for A. Table 1 presents Phillips-curve estimates for 1962-1995 with different values of β imposed. The Table also reports joint estimates of β and the Phillips-curve coefficients obtained by non-linear least squares. The NLLS estimate of β is imprecise: a two-standard-error confidence interval runs from 0.01 to 1.03. This reflects the fact that a wide range of β 's fit the data equally well: the \bar{R}^2 's are close when different values of β are imposed. The point estimate of β is 0.52, which is far from the value of 0.95 that we suggested a priori. However, there is little evidence against $\beta=0.95$: an F-test of this hypothesis yields $F=2.24$ ($p=0.15$).⁵

⁵ This F-test compares the sum of squared residuals with and without the restriction that $\beta=0.95$. Following Staiger et al. (1998), we use this test because it appears more accurate than a test based on the asymptotic standard error.

Fortunately, we can draw conclusions from the data without knowing the value of β . As illustrated in Table 1, the coefficient on θ -A is significantly negative for all β 's from zero to one. Thus, as implied by our model, a rise in productivity growth relative to wage aspirations has a negative effect on inflation. The coefficient on θ -A is usually near -0.6. In terms of underlying parameters, this means that the aspiration term A has a weight of 0.6 in the formula for target wage-growth (equation (1)) and productivity growth has a weight of 0.4. The \bar{R}^2 's for the various β 's lie between 0.5 and 0.6, compared to 0.34 for the equation without θ -A. Thus our new variable explains a significant part of inflation variation over 1962-1995.

Figure 5 shows forecasts of inflation over 1996-2000 for various values of β . In most cases, adding θ -A to the Phillips curve greatly improves the accuracy of forecasts. For β 's ranging from 0.5 to 0.95, predicted inflation stays close to actual inflation throughout the period, and ends up lower by statistically insignificant amounts. For $\beta=0.95$, predicted inflation in 2000 is 2.1%. Thus our model eliminates the overprediction of inflation that arises with the usual Phillips curve. Our equation predicts that inflation stays low despite low unemployment because the productivity acceleration produces high values of θ -A.

The only qualification is that our equation overpredicts inflation if β is very close to one. As discussed above, $\beta=1$ means

that wage aspirations over 1996-2000 are still tied to the rapid wage growth of the 1950s. In this case, $\theta-A$ is negative for most of 1996-2000, so adding it to the model does not reduce inflation forecasts. Our story about the New Economy depends on the assumption that $\beta < 1$: there must be some adjustment of aspirations over time.

C. Short-Run and Long-Run Variation in $\theta-A$

Our results partly reflect broad trends in the data. In the early 1970s, the productivity slowdown reduced $\theta-A$, and the unemployment-inflation tradeoff worsened; these facts help produce the negative coefficient on $\theta-A$ in the pre-1996 Phillips curve. Similarly, the success of our model over 1996-2000 reflects the fact that $\theta-A$ rose while the output-inflation tradeoff improved. However, these broad trends are not the only reason for our model's success. As shown in Figure 3, there is considerable year-to-year variation in $\theta-A$ because of fluctuations in θ . These movements also help explain shifts in the U/π relation.

To make this point, we decompose the variable $\theta-A$ (for $\beta=0.95$) into two components: a trend, given by the HP-filter in Figure 3, and deviations from the trend. For 1962-1995, entering these components separately in the Phillips curve yields the regression

$$\Delta\pi = 3.19 - 0.719U - 1.080(\theta-A)^T - 0.568(\theta-A)^D ,$$

(1.15) (0.210) (0.412) (0.174)

where $(\theta-A)^T$ is the trend component of $\theta-A$ and $(\theta-A)^D$ is the deviation from trend. Both components have statistically

significant effects. The point estimate is higher for the trend component, but one cannot reject the hypothesis that the two coefficients are equal ($p=0.17$). Thus both long-term and short-term movements in θ -A have the effects predicted by our theory.

Researchers often give different interpretations of long-term and year-to-year shifts in the U/π relation. The former are interpreted as shifts in the NAIRU, and the latter as "supply" or "inflation" shocks. This is the case, for example, in the Kalman-filter approach to estimating time-varying NAIRUs (e.g. Gordon, 1998). In contrast, our results suggest that parts of the short-term and long-term shifts in the U/π relation have a common explanation.

D. Is Low Unemployment Sustainable?

This paper is written for a conference on the "sustainability" of today's low unemployment. At first glance, our analysis appears to have pessimistic implications about sustainability. The Phillips curve has shifted favorably because a productivity acceleration has produced positive values of θ -A. But when productivity growth stabilizes, aspirations for real-wage growth will eventually adjust to the new trend. In the long run we must see values of θ -A that average to zero, implying a worse U/π tradeoff than in the recent period of positive θ -A's.

On the other hand, it will not be necessary for future unemployment to rise back to the level thought to be the NAIRU in the mid-1990s. The apparent NAIRU has fallen in 1996-2000 relative to 1962-1995 both because θ -A has been positive in the

recent period and because it was negative on average in the earlier period. The average θ -A before 1996 was negative because, as shown in Figures 1-3, A lagged behind the falling θ during the productivity slowdown. In steady state, the economy must give up the gains from today's positive θ -A's, but not the gains from eliminating negative θ -A's. In other words, the true NAIRU is higher than the apparent NAIRU of today, but lower than the apparent NAIRU before 1996, when unemployment was raised by slow adjustment of aspirations to the productivity slowdown.

Specifically, recall that a Phillips curve for 1962-1995 without the θ -A term yields a NAIRU estimate of 6.2%. In contrast, the equation with θ -A implies a NAIRU of 5.1% (for $\beta=0.95$). 5.1% is our estimate of the unemployment rate consistent with stable inflation when θ -A equals zero. If the true Phillips curve has not shifted since 1995, our equation implies that unemployment must eventually rise to 5.1% from its 2000 level of 4.2%. But it need not rise to the 6.2% level suggested by a conventional Phillips curve.⁶

VI. EXTENSIONS

This section considers various extensions of our time-series analysis.

A. The Wage Phillips Curve

⁶ Following Staiger et al., we can construct confidence intervals for the NAIRU by performing a series of F-tests for whether the NAIRU equals various values. A 95% confidence interval is (3.5, 5.9). This confidence interval becomes (3.8, 6.1) when lagged inflation changes are added to the Phillips curve to eliminate

So far we have focused on our model's implications for price inflation. To further test the model, we now turn to the wage Phillips curve, equation (3). Recall that wage inflation depends on lagged price inflation, unemployment, and a weighted average of θ and A . We also consider the neoclassical special case in which the weight on θ is fixed at one.

Table 2 presents estimates of wage Phillips curves for 1962-1995 ($\beta=0.95$). These estimates support the model. The estimated weights on θ and A are 0.16 and 0.84; the weight on θ is smaller than the weight implied by the price Phillips curve, but the difference is not statistically significant. The hypothesis of a unit weight on θ is strongly rejected. When we relax the restriction that the θ and A coefficients sum to one, it is not rejected ($p=0.76$).

Using the estimates for 1962-1995, Table 2 also reports forecast errors for $\omega-\pi_1$ after 1995. The results parallel those for price Phillips curves. The neoclassical equation overpredicts wage inflation relative to π_1 by a total of 6.4 percentage points. This equation assumes that wage growth rises one-for-one with the productivity acceleration, when in fact the effect was much smaller. Our wage Phillips curve is more accurate: it underpredicts wage growth by an insignificant amount.

B. Additional Phillips-Curve Variables

Most authors who estimate Phillips curves include additional

serial correlation in the errors (see Section VIB).

variables, in particular lags of unemployment and inflation changes and measures of supply shocks (e.g. Gordon, 1998; Staiger et al., 1997). Here we check the robustness of our conclusions to adding such variables. We experiment with two lags of the change in inflation; unemployment lags are never significant, so we omit results with these variables. We measure supply shocks with three standard variables: the change in the relative price of food and energy, the change in the trade-weighted real exchange rate, and Gordon's dummy for the Nixon price controls.⁷

Table 3 presents estimates of our generalized Phillips curves for 1962-1995. We estimate equations with and without the three supply shocks, with and without the two $\Delta\pi$ lags, and with and without $\theta-A$, in all possible combinations. In all cases, we set $\beta=0.95$ in calculating A . There are two robust conclusions. First, the three supply shocks are jointly significant and so are the two $\Delta\pi$ lags, regardless of whether $\theta-A$ is included. The various coefficients have reasonable signs and magnitudes. Including all the variables (column (8)) yields an \bar{R}^2 of 0.81.

Second, the term $\theta-A$ remains significant in all the specifications. However, the magnitude of the coefficient falls when additional variables are included. In the most general specification, the coefficient is -0.32 ($t=3.3$), compared to -0.62

⁷ The change in the relative price of food and energy is the log change in the food-energy component of the CPI minus the log change in the CPI. The exchange-rate variable is the change in the log of the trade-weighted real exchange rate from Data Resources, Inc. Following Gordon, we add constants to these variables to make their means equal zero. The Nixon dummy equals 0.5 in 1972 and 1973, -0.3 in 1974, and -0.7 in 1975.

when the supply shocks and $\Delta\pi$ lags are excluded.⁸

Figure 6 shows forecasts of inflation for 1996-2000 based on the 1962-1995 estimates. The four panels give results with and without the supply shocks and with and without inflation lags. In each case, we show actual inflation and forecasts that arise when θ -A is included and when it is excluded. The forecasts vary across specifications, but again broad conclusions emerge.

First, if one leaves θ -A out of the Phillips curve, accounting for supply shocks reduces the overprediction of inflation by only a moderate amount. When supply shocks are included, predicted inflation stays low through 1998, because the dollar appreciates and energy prices fall in 1998. But predicted inflation rises sharply in 1999-2000 as the appreciation slows and energy prices rise. In the most general specification without θ -A, predicted inflation reaches 6.4% in 2000, compared to 8.3% in the simplest Phillips curve.

Second, including θ -A always reduces predicted inflation by a large amount. In most cases in Figure 6, adding θ -A turns an overprediction of inflation into a fairly accurate prediction. In one case, it turns a moderate overprediction into a moderate underprediction.

Finally, our most general specification - the one including θ -A, supply shocks, and $\Delta\pi$ lags -- produces remarkably accurate

⁸ The proper interpretation of the lower coefficient is not clear. According to our model, it implies a lower coefficient on A in the target-wage equation and a higher coefficient on θ . However, estimating these coefficients from wage-

forecasts throughout the 1996-2000 period. In the first three years, the combination of the productivity acceleration and favorable supply shocks more than offsets the effect of falling unemployment, and inflation is predicted to fall modestly. In 1999 and 2000, when productivity growth stays high but the supply shocks reverse, inflation is predicted to rise modestly. Actual inflation follows a path very close to this predicted one.

C. Time-Varying NAIRUs

The recent behavior of unemployment and inflation has suggested to many observers that the NAIRU has fallen. This idea has increased interest in estimating Phillips curves with time-varying NAIRUs (e.g. Staiger et al. and Gordon). So far this paper has estimated constant-NAIRU models. However, our idea that such a model forecasts inflation better when $\theta-A$ is included can be turned around to say there is less time-variation in the NAIRU once $\theta-A$ is included. In particular, the NAIRU falls less since 1995 if we account for the anti-inflationary role of the productivity acceleration. Here we explore this version of our story.

We estimate time-varying NAIRUs in the following way. We start with the simple Phillips curves we have already estimated. Shocks to these equations cause fluctuations in the level of unemployment consistent with stable inflation. For example, in 1974 it would have taken very high unemployment to offset the OPEC shock and keep inflation stable. As discussed earlier, economists

rather than price-Phillips curves yields an A coefficient above 0.7 regardless of whether supply shocks and inflation lags are included.

generally do not interpret such shocks as year-to-year fluctuations in the NAIRU. Instead, they assume the NAIRU changes gradually, and interpret shifts in the U/π relation as NAIRU shifts only if they appear persistent. In this spirit, we define the time-varying NAIRU as the long-term component of movements in the U/π relation.

Specifically, consider two Phillips curves:

$$(6a) \quad \Delta\pi = -\gamma(U-U^N) ;$$

$$(6b) \quad \Delta\pi = -\gamma(U-U^N) + (1-\delta)(\theta-A) .$$

If U^N is a constant, these reduce to the Phillips curves with and without $\theta-A$ that we estimate above. We impose values of γ and $(1-\delta)$ obtained by estimating constant- U^N equations over 1962-2000: $\gamma=0.636$ in (6a) and $\gamma=0.668$, $(1-\delta)=0.550$ in (6b). Given these coefficients and the data on $\Delta\pi$, U , and $\theta-A$, each equation defines a series for U^N over 1962-2000. In (6a), U^N is the unemployment rate that would produce stable inflation; in (6b) it is the unemployment rate that would produce stable inflation if $\theta-A=0$. Finally, we extract a long-term trend from each U^N series using the Hodrick-Prescott filter with parameter 1000. These smoothed series are our measures of time-varying NAIRUs.

Figure 8 presents the U^N and smoothed- U^N series for each equation. Note first that the average U^N is 6.0% when $\theta-A$ is excluded from the Phillips curve and 5.2% when it is included. This result confirms our earlier finding that including $\theta-A$ reduces the NAIRU when it is assumed to be fixed. The new result is that

adding θ -A also reduces the time-variation in the NAIRU. When θ -A is excluded, the smoothed U^N rises by 1.7 percentage points from 1962 to 1979, then falls by 1.9 points from 1979 to 2000. This hump-shaped path is similar to the NAIRU behavior estimated by previous authors. When θ -A is included, by contrast, the NAIRU rises only 0.7 points from 1962 to 1980, and remains almost constant thereafter. The NAIRU fall from 1990 to 2000 -- a rough measure of the New-Economy effect -- is 1.2 points without θ -A but less than 0.1 points with θ -A. Once our new variable is included, there is no need to search for explanations for a falling NAIRU.

The choice of a smoothing parameter for the HP filter is arbitrary. Reducing the parameter increases the time-variation in both NAIRU series, but does not change the result that the NAIRU is more stable when θ -A is included.

VII. THE CHILEAN MIRACLE

So far we have focused on the United States. It is natural to ask whether our theory also explains apparent Phillips-curve shifts in other countries. The experience of the 1970s suggests that it does. Productivity growth slowed throughout the OECD during the 70s, and the NAIRU appeared to rise in most countries. Grubb et al. (1982) and many others discuss this experience.

Unfortunately, it is difficult to produce international evidence for our theory beyond a broad observation about the 1970s. One might hope to find a cross-country relation between the size of productivity slowdowns or speedups and the size of

NAIRU shifts. A look at OECD data suggests, however, that no clear relation exists. The problem is that the NAIRU has moved sharply in many countries for reasons unrelated to our model, involving labor-market institutions and long-run effects of monetary policy (see Blanchard and Wolfers, 1999, and Ball, 1999). These NAIRU movements usually swamp the effects of productivity shifts that we would like to detect.

The good news is that the cross-country data yield one useful case study: Chile in the 1990s. Chile experienced a major productivity acceleration during this period, one which is usually attributed to economic liberalization. Figure 8 plots the growth rate of labor productivity in Chile for 1976-1997 (measured as the change in log output per worker, from World Development Indicators). Average productivity growth was 0.85 percent over the ten years from 1977 to 1987 and 4.96 percent over 1987-1997. The increase of 4.11 percent is much larger than the recent productivity acceleration in the United States.

Indeed, the Chilean episode is an outlier in international data. There are 40 countries for which 20 or more years of data on productivity growth are available from either the World Development Indicators or the OECD. (The starting dates range from 1961 to 1977 and the ending dates from 1992 to 2000.) For each of these countries, we compute the largest productivity acceleration, defined as the largest difference between average productivity growth in a ten-year period and the previous ten years. For Chile, the largest acceleration is the 4.11 percent increase between 1977-1987 and 1987-1997. This is the largest

acceleration for any country in the sample. The country with the next largest acceleration is Jamaica, with 3.27%, but this reflects an increase from -4.39 percent to -1.12 percent. After that comes Thailand, with an acceleration of 2.96 percent from 1976-86 to 1986-96. Only three other countries have accelerations above 2% starting from positive initial growth rates. Thus Chile's productivity acceleration is more than twice the largest one experienced by most countries, and more than a full point above the second-best in the sample (ignoring Jamaica).

If productivity shifts affect the Phillips curve, there should have been a favorable Phillips-curve shift in Chile. And there was. The shift took a different form than the recent shift in the U.S.: it showed up mainly as falling inflation with stable unemployment rather than vice-versa. That is, Chile had the rare experience of a costless disinflation. Research has shown that a substantial reduction in inflation almost always reduces output and raises unemployment in the short run. For example, Ball (1994) examines 28 disinflations in OECD countries and finds output losses in 27 of them. Dornbusch and Fischer (1993) find that disinflations from moderate levels reduce output in middle-income countries as well.

Chile is a stark exception to this stylized fact. Figure 9 plots inflation, unemployment, and output growth from 1985 through 1997 (after which the miracle was interrupted by the world financial crisis). As shown in the Figure, inflation peaked at 26% in 1990 and then fell steadily, reaching 3% in 1997. But one can see no adverse effects on the real economy. Unemployment fell

from 9.6% in 1990 to 6.6% in 1997. Output growth was 3.7% in 1990 and exceeded 5% in every year from 1991 through 1997.⁹

Thus the Chilean episode combined an unusual productivity acceleration with an unusual shift in the Phillips curve. It stands out from the cross-country data on both counts. Of course the Phillips curve might have shifted for some other reason, but we doubt it. A leading view within Chile is that inflation expectations shifted because the central bank introduced a credible inflation target (e.g. Corbo, 1998). However, other countries have adopted inflation targets, and research has not detected a favorable effect on the Phillips curve. Disinflations usually cause recessions even under inflation targeting (Bernanke et al., 2001).

VIII. MICRO EVIDENCE

So far we have examined aggregate relations among productivity growth, unemployment, and wage and price inflation. We now turn to micro evidence on wage changes for individual workers to corroborate our aggregate findings and to explore the formation of wage aspirations in more detail. Our model assumes that workers use lagged wage increases to form their wage aspirations, but at an individual level we must be more specific about which lagged wages are relevant. In the language of Kahneman et al. (1986), we are interested in who forms the "reference group" that a worker compares himself to in judging the fairness of his wage.

⁹ The data on inflation and output are from the Bank of Chile. The data on

We consider two alternative assumptions about reference groups. The first is that workers form aspirations based on the lagged wages of workers who have the same level of skill and belong to the same birth cohort. This idea generalizes the concept that workers use their own, individual lagged wages to form aspirations. (The idea that a worker examines only his own past wage and nobody else's seems overly narrow, and also requires panel data that are not available.¹⁰) Our second hypothesis is that workers form aspirations based on lagged wages of other cohorts at the same age and skill level; that is, a worker of age a in year t bases his aspirations on the wages of workers of age a in years $t-1$, $t-2$, and so on. The difference in these two hypotheses relates to the familiar demographic distinction between "cohort" and "period" effects.

Following Katz and Krueger (1999), we use individual data from the Current Population Survey (CPS). The May CPS is available from 1973-1978 and data for the Outgoing Rotation Group are available from 1979-1999. We use both hourly wage and weekly wage measures, for the latter are measured more reliably. Like Katz and Krueger, we measure skill by education level and consider four education groups: less than high school, high school, some college, and college degree or more. We use data on workers aged 25-64 over the 1973-1999 period, and group workers by five-year birth cohorts ranging from 1916-20 to 1971-75. Our data cover a total of 888 year-education-cohort cells.

unemployment (in Santiago) are from the University of Chile.

¹⁰ One might use the Michigan Panel Study of Income Dynamics (PSID), but its sample is too small to allow disaggregation by skill group.

The equation we estimate is a micro version of the wage Phillips curve presented earlier:

$$(7) \quad \omega(e,c,t) - \pi(t-1) = a(e) + b(t) + \tau_1(\text{age}) + \tau_2(\text{age})^2 \\ - \gamma U(e,t) + (1-\delta)A(e,c,t) ,$$

where $\omega(e,c,t)$ is wage inflation for education group e and cohort c in year t ; $\pi(t-1)$ is price inflation at $t-1$; $U(e,t)$ are BLS-published unemployment rates by education group; and $A(e,c,t)$ is an average of past wage growth. A is constructed in the same way as before, using a β of 0.95 and an HP-filtered value for the start of the process in 1974 -- but using, in one case, a cohort's own lagged-wage-growth profiles and, in the other, the wage growth of workers of the cohort's current age in past years. We include a quadratic in age to capture life-cycle patterns in wage growth, and dummies for education groups and for years. Including these variables means that the coefficient on A is determined by the cross-sectional relation between year-to-year changes in A for different education and birth-year groups and changes in real-wage growth.

The major difference between equation (7) and the aggregate equations estimated earlier is the absence of a productivity variable. Unfortunately, productivity data are unavailable for education groups and other disaggregate portions of the labor force, and hence it must be omitted. The education and year dummies and the age variables capture productivity growth that is common to all groups in each year (i.e., aggregate) as well as productivity growth that is common to each education and age group

in all years. It omits the portion of productivity growth that is specific to different education groups in different years. However, productivity shocks of this kind should be orthogonal to lagged wages and hence to A, and thus should not bias the coefficients.

The top panel of Table 4 shows our initial estimation results. We denote the aspirations variable by AC when it is constructed from a cohort's own lagged wages, and by AA when it is based on wages of workers of the same age. The unemployment coefficients in the regressions are significantly negative, as expected, albeit smaller than in the aggregate results. Most important, the aspirations variables are all positive and significant. Thus the micro data corroborate our aggregate finding that wage growth is tied to lagged wages. The effect is significant when aspirations are measured by either AC or AA.

While significant, the coefficients on aspirations are smaller in our initial micro regressions than in our aggregate regressions. However, a common problem with micro data is that regressors based on lagged dependent variables are noisy and contain large random fluctuations. It is unlikely that individuals change their aspirations in response to these fluctuations and, indeed, a certain fraction represents sheer measurement error. The consequence of this problem, which is formally equivalent to an errors-in-variables problem, is downward bias in the coefficients. To remedy the problem, we replace the raw aspirations variables with variables that are smoothed over year, education, and age. The results are shown in

the lower panel of Table 4. The coefficients on A rise substantially, and reach magnitudes close to those obtained with aggregate data. This result strongly suggests the presence of errors-in-variables bias in the raw data.¹¹

Table 5 shows the results of including both aspirations variables, AC and AA, in the model at the same time. The two coefficients are both smaller than in Table 4, but they are both significant and they are close to each other in size. This result suggests that, in forming ideas about fair wage increases, workers put roughly equal weight on their own past experience and on the wage growth of similar workers in the past.

Tests reveal no significant differences across the four education groups in the coefficients on unemployment and aspirations. That is, while unemployment and lagged wages move in different ways for different groups, the effects of given movements on wage growth are the same.

Further inspection of the data reveals that A has been drifting upward for the more educated groups and downward for the less educated groups, thus producing very different patterns of wage growth. Note that A represents real-wage growth in the past, not the level of the wage, so this is not necessarily to be expected from the well-known increasing dispersion of wages by education level. Instead, it implies that the spreading out accelerated over most of the period we examine. Because A has declined so severely for the less educated group, the average A

¹¹ The smoothed series for A are fitted values from regressions of A on education-year polynomial interactions, education-age polynomial interactions, a quadratic in age, and year and education dummies.

has also fallen, consistent with the aggregate data. However, the less-educated groups experienced above-average real-wage growth in the second half of the 1990s, which slowed the decline in A for those groups.

IX. CONCLUSION

This paper proposes a new variable for the Phillips curve: the difference between productivity growth and an average of past real-wage growth. Theoretically, this variable appears if workers' aspirations for real-wage increases adjust slowly to shifts in productivity growth. Empirically, our new variable shows up strongly in the U.S. Phillips curve. Including it explains the otherwise-puzzling shift in the unemployment-inflation relation since 1995.

Our theory contributes to a parsimonious interpretation of macroeconomic history. It yields a unified explanation of why unemployment rose during the productivity slowdown of the 1970s and why it fell after 1995. The theory also explains part of the year-to-year fluctuations in the unemployment-inflation tradeoff as arising from fluctuations in productivity growth. Finally, our story links two features of the post-1995 New Economy. The Phillips curve shift was caused by the productivity acceleration rather than happening to occur at the same time for some other reason.

In the mid-1990s, the consensus estimate of the NAIRU was 6%. Since then, unemployment has fallen near 4%, and inflation has not risen substantially. Our results suggest that the non-

inflationary fall in unemployment is partly but not entirely sustainable. The economy has moved from a regime in which wage aspirations exceed productivity growth, raising unemployment, to one in which aspirations are below productivity growth. Eventually the economy must move toward a steady state in between. We estimate the NAIRU in this steady state to be around 5.1%.

REFERENCES

- Akerlof, George A., and Janet Yellen, "The Fair Wage-Effort Hypothesis and Unemployment," Quarterly Journal of Economics 105 (May, 1990), 255-283.
- Ball, Laurence, "Aggregate Demand and Long-Run Unemployment," Brookings Papers on Economic Activity 1999:2, 189-251.
- Ball, Laurence, "What Determines the Sacrifice Ratio?" in N. Gregory Mankiw (ed.), Monetary Policy, University of Chicago Press, 1994.
- Basu, Susanto and Miles S. Kimball, "Cyclical Productivity with Unobserved Input Variation," NBER Working Paper 5915 (February, 1997).
- Bernanke, Ben S., Thomas Laubach, Frederic S. Mishkin, and Adam S. Posen, Inflation Targeting: Lessons from the International Experience, Princeton University Press, January 2001.
- Bewley, Truman, Why Wages Don't Fall During a Recession, Harvard University Press, 2000.
- Bils, Mark and Jang-Ok Cho, "Cyclical Factor Utilization," Journal of Monetary Economics 33 (April, 1994), 319-54.
- Blanchard, Olivier and Justin Wolfers, "The Role of Shocks and Institutions in the Rise of European Unemployment: The Aggregate Evidence," NBER Working Paper 7282 (August, 1999).
- Blanchard, Olivier and Lawrence F. Katz, "What We Know and Do Not Know About the Natural Rate of Unemployment," Journal of Economic Perspectives 11 (Winter, 1997), 51-72.
- Blinder, Alan, "The Internet and the New Economy," Brookings Institution Policy Brief 60, June 2000.
- Corbo, Vitorio, "Reaching One-Digit Inflation: The Chilean Experience," Journal of Applied Economics 1 (November, 1998), 123-63.
- Council of Economic Advisors, Economic Report of the President 2000.
- DeLong, J. Bradford, "What Went Right in the 1990s? Sources of American and Prospects for World Economic Growth," in David Gruen and Sona Shrestha (eds.), The Australian Economy in the 1990s, Reserve Bank of Australia, 2000.
- Dornbusch, Rudiger and Stanley Fischer, "Moderate Inflation," World Bank Economic Review 7 (March, 1993), 1-44.
- Elliot, Robert F., Labor Economics: A Comparative Text, McGraw-

Hill, 1991.

Gordon, Robert J, "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU," Brookings Papers on Economic Activity 1998:2, 297-333.

Grubb, Dennis, Richard Jackman and Richard Layard, "Causes of the Current Stagflation," Review of Economic Studies 49 (Special Issue, 1982), 707-30.

Kahneman, Daniel, Jack L. Knetsch, and Richard Thaler, "Fairness as a Constraint on Profit-Seeking," American Economic Review 76 (1986), 728-41.

Katz, Lawrence F. and Alan B. Krueger, "The High-Pressure US Labor Market of the 1990s," Brookings Papers on Economic Activity 1999:1, 1-65.

Milkovich, George T., and Jerry M. Newman, Compensation (5th ed.), Irwin, 1996.

Oswald, Andrew, "Is Wage Rigidity Caused by 'Lay-offs by Seniority'?", in Wilfred Beckerman (ed.), Wage Rigidity and Unemployment, Gerald Duckworth and Co., 1986.

Schor, Juliette B., "Does Work Intensity Respond to Macroeconomic Variables? Evidence from British Manufacturing, 1970-86," Harvard University, 1997.

Staiger, Douglas, James H. Stock and Mark W. Watson, "How Precise Are Estimates of the Natural Rate of Unemployment?" in Christina D. Romer and David H. Romer (eds.), Reducing Inflation: Motivation and Strategy, University of Chicago Press, 1997.

Stiglitz, Joseph, "Reflections on the Natural Rate Hypothesis", Journal of Economic Perspectives 11 (Winter, 1997), 3-10.

Table 1. Phillips Curve Estimates, 1962-1995.

(Dependent Variable: $\Delta\pi$)

	β estimated	β imposed				
Value of β	0.515	0.500	0.750	0.900	0.950	1.000
	(0.255)					
Constant	0.0244	0.0245	0.0258	0.0303	0.0328	0.0377
	(0.0107)	(0.0099)	(0.0108)	(0.0115)	(0.0111)	(0.0096)
Coef. on U	-0.409	-0.410	-0.449	-0.565	-0.649	-0.777
	(0.166)	(0.157)	(0.171)	(0.180)	(0.175)	(0.154)
Coef. on $\theta-A$	-0.613	-0.609	-0.640	-0.623	-0.619	-0.498
	(0.205)	(0.160)	(0.192)	(0.200)	(0.187)	(0.152)
\bar{R}^2	0.552	0.566	0.547	0.534	0.539	0.508

Table 2. Wage Phillips Curves, 1962-1995.

(Dependent Variable: $\omega - \pi_t$)

		Our Model	Coef. on θ set at 1.0
- Constant		0.0388 (0.0068)	0.0541 (0.0103)
Coef. on U		-0.789 (0.121)	-0.871 (0.145)
Coef. on θ		0.163 (0.170)	1.000
Coef. on A		0.837	
Forecast Errors (Percentage Points)	1996	0.50	2.20
	1997	0.84	1.64
	1998	-1.43	0.22
	1999	-1.19	0.14
	2000	-0.61	2.25
	Sum for 1996-2000	-1.89	6.44
	(std. error for sum)	(1.53)	(2.00)

Table 3. Generalized Phillips Curves, 1962-1995.

(Dependent Variable: $\Delta\pi$)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	0.0441 (0.0114)	0.0328 (0.0111)	0.0367 (0.0079)	0.0299 (0.0077)	0.0274 (0.0055)	0.0215 (0.0038)	0.0319 (0.0083)	0.0271 (0.0066)
U	-0.710 (0.161)	-0.649 (0.175)	-0.585 (0.121)	-0.575 (0.122)	-0.437 (0.096)	-0.422 (0.064)	-0.505 (0.136)	-0.494 (0.096)
θ -A		-0.619 (0.187)		-0.501 (0.155)		-0.379 (0.107)		-0.316 (0.096)
$\Delta\pi_1$			0.281 (0.150)	0.184 (0.121)			-0.061 (0.117)	-0.083 (0.099)
$\Delta\pi_2$			-0.448 (0.115)	-0.362 (0.111)			-0.295 (0.069)	-0.253 (0.100)
Δ food/energy price					1.234 (0.191)	0.976 (0.139)	1.192 (0.186)	1.000 (0.150)
Δ exchange rate					-0.065 (0.035)	-0.080 (0.033)	-0.045 (0.033)	-0.061 (0.028)
Nixon					-0.010 (0.011)	-0.006 (0.007)	-0.020 (0.010)	-0.018 (0.008)
\bar{R}^2	0.343	0.539	0.513	0.636	0.684	0.781	0.770	0.812

Table 4. Wage Phillips Curves Estimated on Micro Data

Results with Raw Aspirations Variables

Dependent Vbl.	Hourly Wage Growth		Weekly Wage Growth	
Coef. on $U(e,t)$	-0.16	-0.17	-0.25	-0.27
	(0.07)	(0.07)	(0.08)	(0.08)
Coef. on $AC(e,c,t)$	0.31		0.25	
	(0.14)		(0.14)	
Coef. on $AA(e,c,t)$		0.38		0.38
		(0.07)		(0.07)

Results with Smoothed Aspirations Variables

Dependent Vbl.	Hourly Wage Growth		Weekly Wage Growth	
Coef. on $U(e,t)$	-0.34	-0.31	-0.40	-0.37
	(0.08)	(0.08)	(0.09)	(0.09)
Coef. on $AC(e,c,t)$	0.87		0.72	
	(0.20)		(0.19)	
Coef. on $AA(e,c,t)$		0.53		0.49
		(0.08)		(0.08)

Table 5. Wage Phillips Curves with Both Aspirations Variables (Smoothed).

Dependent Vbl.	Hourly Wage Growth	Weekly Wage Growth
Coef. on $U(e,t)$	-0.30 (0.08)	-0.35 (0.09)
Coef. on $AC(e,c,t)$	0.46 (0.21)	0.40 (0.20)
Coef. on $AA(e,c,t)$	0.46 (0.09)	0.44 (0.08)

Figure 1. Productivity Growth (θ).

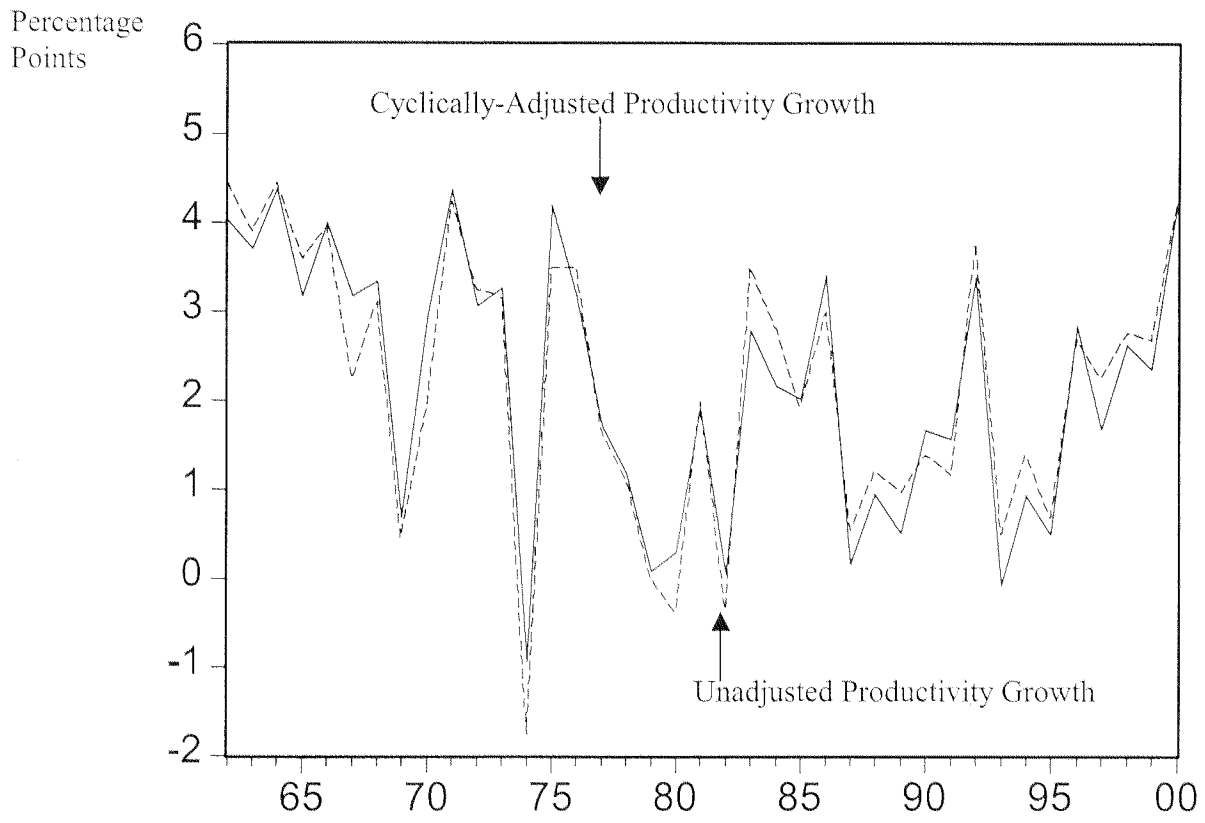


Figure 2. Real-Wage Growth and Aspirations.




Figure 3. The Gap Between Productivity Growth and Aspirations ($\beta=0.95$)




Figure 4. Dynamic Inflation Forecasts: Conventional Phillips Curve.

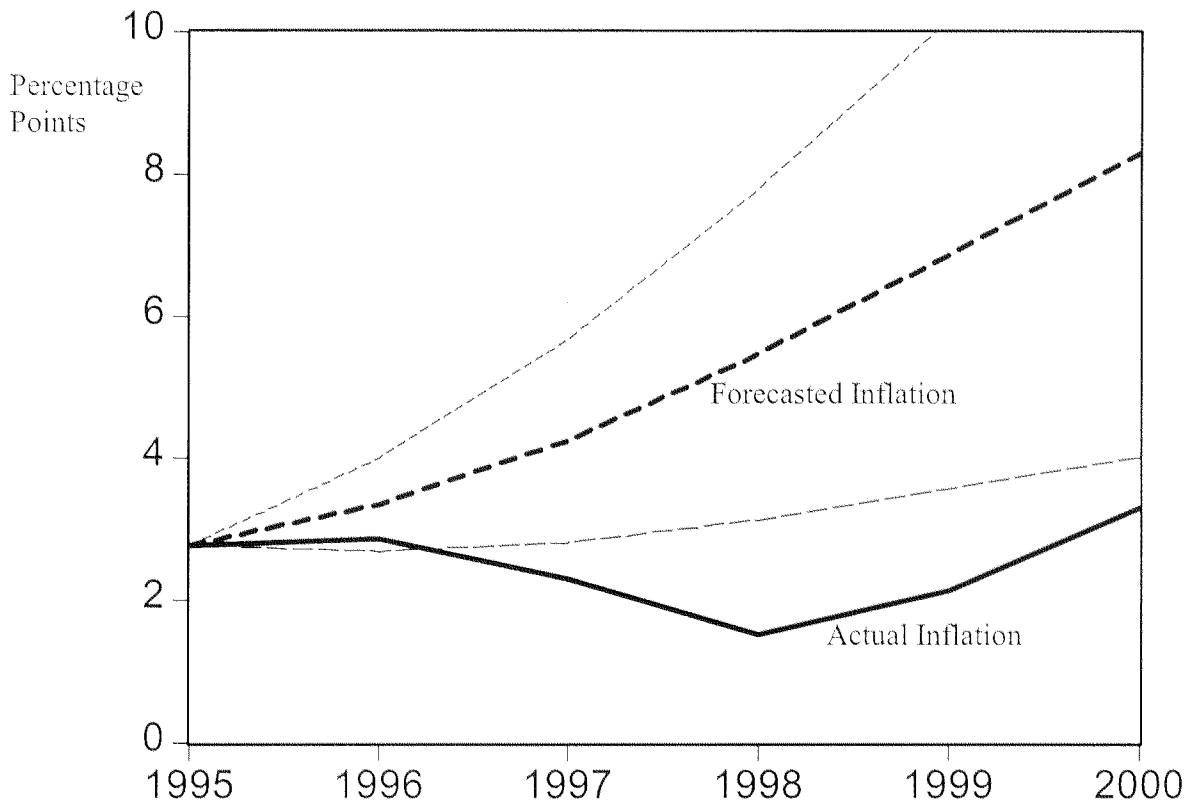


Figure 5. Dynamic Inflation Forecasts: Phillips Curve with θ -A







Figure 6. Dynamic Inflation Forecasts: Generalized Phillips Curves.
Simple Phillips Curves



Phillips Curves with Supply Shocks



Phillips Curves with Supply Shocks and $\Delta\pi$ Lags




Figure 7. A Time-Varying NAIRU?

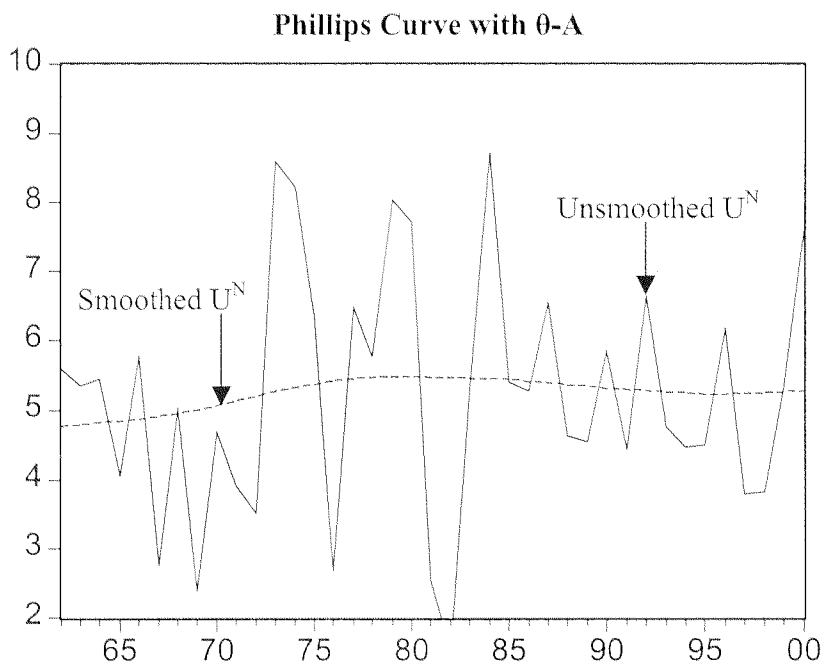
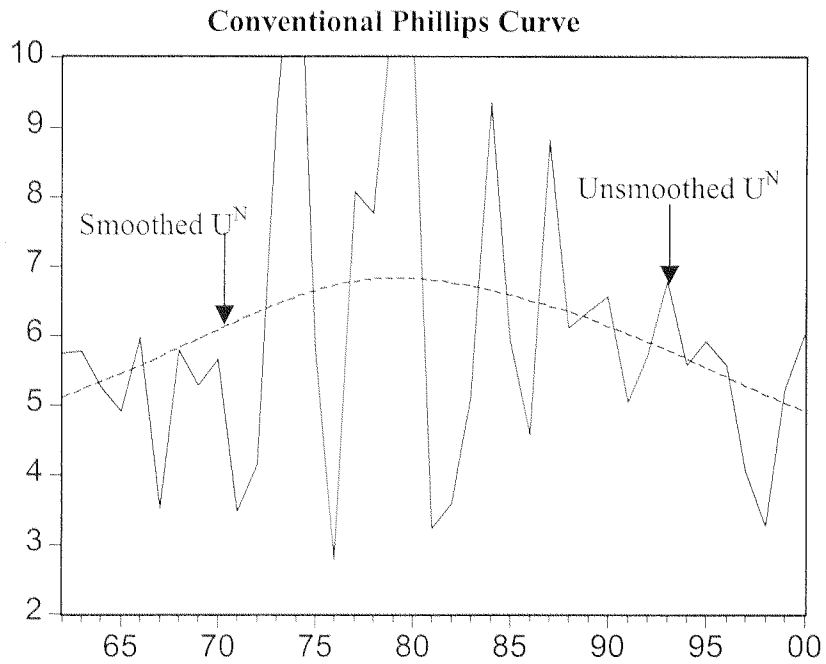


Figure 8. Productivity Growth in Chile.

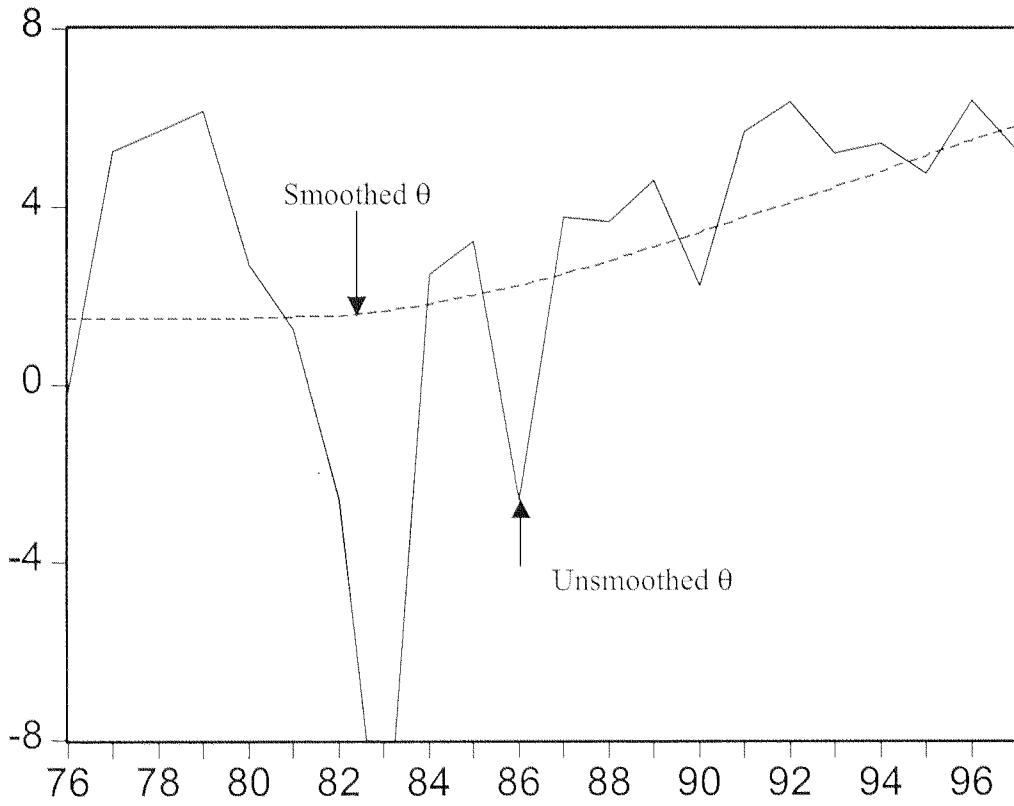


Figure 9. Chile's Phillips Curve Shift.

