

Which Nonlinearity in the Phillips Curve? *

The Absence of Accelerating Deflation in Japan

Emmanuel De Veirman

Reserve Bank of New Zealand

January 14, 2007

Abstract

It is standard to model the output-inflation trade-off as a linear relationship with a time-invariant slope. We assess empirical evidence for three types of nonlinearity in the short-run Phillips curve. At an empirical level, we aim to discover why large negative output gaps in Japan during the period 1998-2002 did not lead to accelerating deflation, but instead coincided with stable, be it moderately negative inflation. We document that this episode is most convincingly interpreted as reflecting a gradual flattening of the Phillips curve. The broader relevance of our analysis lies in its attempt to shed light on the determinants of such time-variation in the Phillips curve slope. Our results suggest that, in any economy where trend inflation is substantially lower (or substantially higher) today than in past decades, time-variation in the slope of the short-run Phillips curve has become too important to ignore.

JEL codes: C22, C32, E31, E32

Keywords: nonlinear Phillips curve, time-varying parameter models.

*This paper is based on the first chapter of my dissertation at the Johns Hopkins University. I am grateful to my advisor, Laurence Ball, and to Alan Ahearne, Carl Christ, Robert Davies, Hali Edison, Jon Faust, Yasuo Hirose, Michael Kiley, Takeshi Kudo, Kenneth Kuttner, Douglas Laxton, Andrew Levin, Louis Maccini, Athanasios Orphanides, Adrian Pagan, Erwan Quintin, John Roberts, Jirka Slacalek, Tsutomu Watanabe, Isamu Yamamoto, Naoyuki Yoshino, and seminar participants at the Federal Reserve Bank of Dallas, the International Monetary Fund, the Canadian Economics Association, the Japan Economic Seminar, and Johns Hopkins University for valuable comments and suggestions. Any errors are mine. I am currently an economist in the Research Department of the Reserve Bank of New Zealand. Contact: Reserve Bank of New Zealand, PO Box 2498, Wellington 6140, New Zealand. E-mail: deveirman@jhu.edu

1 Introduction

The original Phillips curve was nonlinear: Alban W. Phillips (1958) estimated a nonlinear relationship between nominal wage inflation and the unemployment rate in the United Kingdom. Since that time, it has become standard to model the short-run Phillips curve as a linear relationship with a time-invariant slope. The present paper argues that this simplifying assumption is not as innocent as it seems.

Our paper assesses the empirical performance of three classes of models in which the slope of the Phillips curve varies endogenously over time. The model classes differ according to the set of variables determining the slope of the output-inflation trade-off.

In papers such as Laxton, Meredith, Rose (1995), the size of the output gap determines the slope of the Phillips curve. In particular, the output-inflation trade-off becomes steeper as the output gap approaches the capacity constraint, which is the maximum possible level of output that firms can supply in the short run. As such, the short-run Phillips curve is convex, with a vertical asymptote at the capacity constraint.

In Ball, Mankiw, Romer (1988) and Dotsey, King, Wolman (1999), trend inflation is among the determinants of the Phillips curve slope. In these models of costly price adjustment, the frequency of price adjustment depends on firms' optimizing decisions. A decrease in trend inflation, for one, causes firms to adjust prices less frequently, which in turn implies a flatter Phillips curve.

In Lucas (1973), the slope of the Phillips curve depends on the volatility of aggregate demand and supply shocks. For instance, if aggregate volatility decreases, a larger fraction of any change in the overall price level is misperceived by firms as being a change in their relative price. In that scenario, any change in aggregate demand has a larger impact on firms' production, and a smaller effect on inflation. That is to say, the Phillips curve flattens.

Throughout this paper, we refer to the three classes of models as implying different types of nonlinearity in the Phillips curve. Strictly speaking however, only the first of the above model classes implies that the short-run Phillips curve is nonlinear at a given point of time. In the other cases, the Phillips curve is linear at any point of time, but its slope changes over time as a consequence of changes in trend inflation or aggregate volatility.

To test these theories of nonlinearity, we gather evidence from Japan. The period 1991-2002 in Japan can be characterized as a succession of recessions, interrupted only by brief or limited recoveries. Standard estimates suggest that the output gap was negative for most of that period. Initially, inflation declined, with core CPI inflation reaching the zero-level in the mid-nineties, and turning negative in the second half of the nineties. After 1998 however, annual core CPI inflation remained fairly stable at moderately negative levels, reaching its trough at -0.79% in 2002.

As we document in our paper, the fact that deflation remained surprisingly mild notwithstanding a relatively long period of negative output gaps presents a puzzle to anyone who takes a standard linear Phillips curve literally. This makes Japan a particularly interesting test case for assessing the nature of the output-inflation trade-off.

An advantage of using recent data for Japan is that, unlike the samples typically used in earlier empirical tests of the three types of nonlinearity, our sample includes a fairly large number of observations from the region of the Phillips curve at which inflation is near-zero or negative. The inclusion of low- and negative- inflation observations increases our chances of obtaining precise results as to the existence and type of nonlinearity. Our results are instructive for other economies in which a comparably long deflationary period has not occurred in the post-1945 period, yet which bear sufficient resemblance to Japan in that inflation has tended to decline and/or become less volatile over the course of recent decades.

Our main results, and their relationship to previous papers' findings, can be summarized as follows.

We find evidence for a statistically significant, gradual flattening in a linear Phillips curve which has been occurring since before the nineties. This finding is related to existing research on the flattening of the Japanese Phillips curve,¹ which however typically focuses on documenting a one-time structural 'break' in the Phillips curve under the assumption of a known break date.

Given that we observed a gradual flattening in the Japanese Phillips curve, we assess whether such time-variation is consistent with any of the theories of nonlinearity. This part of our paper is related to earlier empirical evidence on nonlinearity in the Phillips curve.² We find that each of the three types of nonlinearity is consistent with the data.

Our paper moves beyond documenting mere consistency of the nonlinear theories with the data. We find that each of the nonlinear models performs significantly better, in an econometric sense, than an atheoretical benchmark model in which the Phillips curve is linear, but its slope varies over time as a random walk. This implies that the time-variation in the Phillips curve slope has been largely systematic.

Moreover, we perform a series of non-nested model hypothesis tests to assess the relative performance of the three types of nonlinearity. Our results favor the hypothesis that declining trend

¹Nishizaki and Watanabe (2000) and Mourougane and Ibaragi (2004).

²A vast body of papers including Lucas (1973), Alberro (1981) and Kormendi and Meguire (1984) investigates whether the Phillips curve tends to be steeper in economies with high aggregate volatility, without however examining the determinants of changes in the Phillips curve slope over time. Froyen and Waud (1980) and Ilmakunnas and Tsurumi (1985) are more closely related to our paper in that they provide some intertemporal evidence.

DeFina (1991), Hess and Shin (1999), and Kiley (2000) further test the Ball-Mankiw-Romer theory. The first two papers as well as BMR provide some intertemporal evidence on the relation between changes in the Phillips curve slope and changes in trend inflation. See Bonomo and Carvalho (2004) for a recent theoretical contribution.

In addition to Laxton-Meredith-Rose, papers which support the possible existence of asymmetries in the Phillips curve include Turner (1995), Debelle and Laxton (1996), Laxton, Rose, Tambakis (1999), and Dolado, Maria-Dolores, Naveira (2005). See Schaling (2004) for a related theoretical contribution.

Dotsey-King-Wolman and subsequent papers provide model simulations, but no empirical evidence on the relation between trend inflation and the Phillips curve slope. Recent theoretical contributions include Golosov and Lucas (2003), Burstein (2006), Gertler and Leahy (2006), and Dotsey, King, Wolman (2007).

inflation caused firms to set their output prices less frequently, which would explain the observed gradual flattening in the Phillips curve. All but one of our tests lend equally strong support to the hypothesis that a decline in aggregate inflation volatility exacerbated firms' misperceptions about relative prices, implying a flatter Phillips curve. While stories in which capacity constraints engender a convex short-run Phillips curve are consistent with the data, they perform poorly in comparison with the two other models.

A few of the papers providing empirical evidence on the relation between the Phillips curve slope and average inflation / aggregate volatility do so in the time dimension (see footnote 2). Typically, these papers examine whether changes in the Phillips curve slope across subsamples are positively related to cross-subsample changes in average inflation and/or aggregate volatility. Such procedures assume that, if the Phillips curve slope changes over time, it does so in the form of a sudden jump at the sample split point. Our paper is written in the belief that the actual Phillips curve slope is more likely to vary gradually over time. For one, our finding that the Phillips curve gradually flattened is based on a time-varying coefficients model which yields an estimate for the Phillips curve slope at every point of time rather than only an estimate per subsample. This is logically consistent with the fact that our regressions testing the nonlinear theories similarly do not restrict time-variation in the Phillips curve slope to occur as a one-time jump.

Our paper is structured as follows. Section 2 documents that, assuming a standard linear relationship between the output gap and inflation, the size of the negative output gaps in Japan would have warranted accelerating deflation in the period 1998-2002. Section 3 gives arguments against popular explanations for the absence of accelerating deflation in Japan which do not imply time-variation in the slope of the short-run Phillips curve.

The remainder of the paper focuses on time-variation in the slope of the short-run Phillips

curve. In section 4, we detect a gradual, significant decline in the slope of the Phillips curve which has been occurring since before the nineties. Sections 5 through 7 investigate the determinants of the flattening of the Phillips curve. In section 5, we find that all three above-mentioned types of nonlinearity are consistent with the data. In section 6, we find that each of the theories of nonlinearity outperforms a model in which the Phillips curve is linear, yet its slope evolves over time as a random walk. Section 7 evaluates the relative performance of the three types of nonlinearity. Section 8 concludes and presents policy implications.

2 Background

This section deals with three issues. First, we discuss the output gap series used in this paper and its relation to other output gap estimates for Japan. Second, we document that in Japan, the output gap and inflation tended to comove positively through 1997, to such an extent that their relationship could be reasonably well approximated by a standard linear Phillips curve. Third, we show that, in the period 1998-2002, the output gap was sufficiently negative for a linear Phillips curve to predict accelerating deflation, a prediction which is at odds with the data.

Figure 1 documents the evolution of Japan's real Gross Domestic Product, along with potential real output as estimated for Japan by the US Federal Reserve. It is evident that average economic growth since the stock market crash of December 1989 has been lower than it was in any of the previous two decades.

The potential output series in Figure 1 corresponds to the Fed's estimates through 1998. Because the Fed's recent estimates of potential are confidential, we extrapolate potential output for 1999Q1-2004Q4. In doing so, we use the IMF staff estimates/forecasts of potential growth, as quoted in Bayoumi (2000), as a guideline.

According to our measure of potential, annual potential output growth has tended to slow down gradually from 3.88% in 1990 to 1.20% in 1998. Subsequently, potential output growth continued to decline, but at a slow pace, until it reached 0.83% in 2004.

The top panel of Figure 2 graphs the output gap series implied by the actual and potential output data from the previous figure. Potential output growth turns out to have been sufficiently high for a relatively large negative output gap to exist over most of the period 1993-2003. However, since estimates of potential output are typically associated with a high degree of uncertainty, we compare our output gap series with other existing output gap measures.

The Fed's estimates are directly comparable to those output gap estimates for Japan which are based on an estimate of potential output derived from a production function involving the capital stock, the labor stock, and their respective long-run factor utilization rates. In particular, the Bank of Japan (2006) has recently developed a production function based procedure, designed to minimize any upward bias in potential output growth which may have existed in its earlier production function based estimates, as reviewed in Kamada (2005).³

Like the Fed estimates which we use, the Bank of Japan (2006) estimates suggest that, even when accounting for a sizeable decline in potential output growth from the early nineties to the mid-nineties, Japan experienced relatively large negative output gaps for most of the nineties. The Bank of Japan output gap implies that, if anything, the nineties was even worse a decade for Japan than the Fed estimates suggest, relative to past output gaps.

Unlike the production function approach, two other standard procedures for estimating potential output do not yield large negative output gaps. However, we do not consider output gaps

³For instance, the new measure treats the following two developments as structural factors, and in so doing reduces the estimate of potential labor input which enters the production function: a decline in working hours, among others due to labor law changes at the end of the eighties, and a decline in the labor force participation rate since the mid-nineties, among others due to population aging.

based on these procedures to be valuable tools for assessing time-variation in the relation between the output gap and inflation in Japan.

First, univariate smoothing methods such as the Hodrick-Prescott filter yield an output trend which is automatically close to actual output whenever the latter stagnates for a fairly long time at the end of the sample. Unsurprisingly, we find (not shown here) that proxying potential output by a HP-filter trend does not yield large negative output gaps at the end of the sample, as confirmed by the HP-filter based output gap in Kamada (2005).

Second, we applied the methodology of Hirose and Kamada (2003) to estimate potential output as the level of output at which inflation is stable. We find (not shown here) that the Hirose-Kamada output gap moves around zero at the end of the sample. This outcome is not surprising: at times when inflation is fairly stable, output is by definition near its stable-inflation level. In general, there will be little to no time-variation in the slope of the relationship between inflation and an output gap measure which is precisely constructed to fit inflation accurately at all times.

We are now ready to gain our first insights about the comovement of the output gap and inflation in Japan. The lower panel of Figure 2 graphs annualized quarterly inflation in the Consumer Price Index excluding fresh foods, which the Bank of Japan adjusted for consumption tax reforms.⁴ Note that a simple comparison between the output gap and inflation is clouded by supply shocks, such as the oil price shocks which led core CPI inflation to spike up in 1974Q1 and 1980Q2. For now, a casual inspection of Figure 2 suggests that the relationship between the output gap and inflation was fairly well-behaved throughout the seventies and eighties, in the sense that inflation declined when the output gap was negative, and tended to increase in booms.

To characterize the output-inflation comovement through 1997 somewhat more formally, we

⁴A consumption tax of 3% was introduced in April 1989. That sales tax was increased to 5% in April 1997.

regress the following linear Phillips curve using data for 1971Q2-1997Q4:

$$\pi_t = \beta_1 \pi_{t-1} + \beta_2 \pi_{t-2} + \beta_3 \pi_{t-3} + \beta_4 \pi_{t-4} + \gamma_1 ygap_{t-1} + \gamma_2 ygap_{t-2} + \delta impoil_t + e_t \quad (1)$$

Observations for 1970Q2-1971Q1 are used to construct lags. Annualized CPI inflation excluding fresh foods is a function of four inflation lags and two output gap lags. To control for supply shocks in the seventies, we include relative inflation in the import prices of petroleum, coal, and natural gas.⁵

In equation (1), inflation expectations are proxied by lags of inflation. In section 3.1, we document that inflation expectations in Japan indeed tracked lagged inflation relatively closely.

The lag structure in equation (1) removes all serial correlation from the error term e_t , but is sufficiently parsimonious for our estimations involving time-varying output gap coefficients and/or nonlinearities in the Phillips curve in sections 4 through 7. We restrict the sum of the inflation lag coefficients to equal one, and set the constant to zero.⁶

A linear Phillips curve estimated through 1997Q4 fits the data well: the adjusted R-squared is 0.83. The sum of the output gap coefficients is positive (with a point estimate of 0.21) and significant at the 5% level. This confirms that, during a typical episode in the period 1971Q2-1997Q4, positive output gaps exerted upward pressure on inflation, while negative output gaps

⁵Oil import prices are on Yen basis. In equation (1) as in all Phillips curve specifications below, the results are comparable when we include relative inflation in *general* import prices instead. Both supply shock measures are obtained from the Bank of Japan.

⁶Augmented Dickey-Fuller tests reject a unit root in the output gap, and in relative oil import prices, at the 1% level. We cannot statistically reject a unit root in inflation, but the change in inflation is stationary. Moreover, in an unrestricted regression with a constant, the sum of the inflation lag coefficients is not significantly different from one. These considerations lead us to restrict the sum of the inflation lag coefficients to equal one, which is analogous to rewriting (1) as an equation for the change in inflation, with three lags of the change in inflation on the right-hand side. Since our Phillips curve is effectively written in terms of changes in inflation, excluding the constant is necessary to avoid the possibility of a long-run trend in inflation. If we do include a constant, it is virtually zero and insignificant.

tended to coincide with disinflations.

The relationship between the output gap and inflation became gradually less pronounced. In particular, we focus on the episode 1998-2002 because it constitutes the most striking puzzle. It is the episode with the largest negative output gaps, yet it is among the episodes with the most stable inflation rates. Over the period 1998-2002, actual output was on average 2.97% below potential. Meanwhile, annual core inflation did fall from 0.82% in 1997 to -0.35% in 1998, but from that point on declined only marginally until it reached its trough of -0.79% in 2002.

To illustrate this point, Figure 3 shows the result of a dynamic out-of-sample inflation forecast from equation (1) for the period 1998Q1-2004Q4, contrasted with actual inflation. Predicted deflation accelerates to -8.36% in 2002Q3, while actual annualized inflation fell below -1% in only two quarters, reaching -1.58% in 2000Q4.⁷ This suggests that deflation was milder than one would have expected conditional on the large, negative output gaps, and assuming a linear relationship between the output gap and inflation.

The out-of-sample forecast of accelerating deflation does not, by itself, constitute conclusive evidence for a statistically significant break in a standard linear Phillips curve. We do find evidence for statistically significant structural change in the Phillips curve slope in section 4. Before turning our attention to the slope of the Phillips curve however, we first evaluate candidate explanations for the absence of accelerating deflation which do not rely on time-variation in the Phillips curve slope.

⁷We equally obtain a massive deflation forecast from an analogous equation with the GDP deflator. With the GDP deflator, both actual and predicted inflation are more negative than in the CPI case.

3 Capturing output-inflation comovement without time-varying Phillips curve slope?

We give arguments against three possible explanations for the absence of accelerating deflation in Japan, none of which implies time-variation in the slope of the short-run Phillips curve.

3.1 Did inflation expectations fail to turn negative?

The forecast of accelerating deflation in section 2 originated from an accelerationist Phillips curve, in which inflation expectations were proxied by lagged inflation. Thus, equation (1) implicitly assumes that inflation expectations turned moderately negative, along with actual inflation. It is possible that inflation expectations did in fact not turn negative in the period 1998-2002, even at times of deflation in the actual core CPI. If Japanese inflation expectations hovered around zero, the Phillips curve would lose its accelerationist feature, as a passage in Blanchard (2000) explains. Under that hypothesis, negative output gaps would imply negative, but stable inflation.⁸ This would accurately capture the output-inflation comovement in Japan in the period 1998-2002.

However, every known measure of inflation expectations in Japan suggests that inflation expectations did turn negative. The one-shot 2002 METI survey finds that only 5.6% of firms, and only 3.0% of consumers, expected deflation to end within one year. The December 2001 Consensus forecasts predict headline CPI inflation of -0.9% for 2002. The finding that inflation expectations turned negative is confirmed by qualitative price expectations data in the Tankan business survey, and by inflation forecasts from the OECD and the US Federal Reserve.

⁸To see this, write the Phillips curve as $\pi_t = \beta.E_{t-1}(\pi_t) + \gamma_1.ygap_{t-1} + \gamma_2.ygap_{t-2} + \delta.impoilt + e$, where $E_{t-1}(\pi_t)$ stands for lagged expectations of current inflation. If inflation expectations remain at zero, this implies that $E_{t-1}(\pi_t) = 0$. From the above equation, it follows that in that case, negative output gaps tend to coincide with negative, but stable inflation.

Basic econometric analysis confirms that inflation expectations continued to track lagged inflation relatively closely even as lagged inflation turned negative. Our results suggest that, if there was any structural change at all in the process of expectations formation, inflation expectations turned even more negative in the period since the mid-nineties than would otherwise have been warranted by lagged inflation.⁹

3.2 Did expansionary monetary policy prevent massive deflation?

In a textbook world, fast money growth exerts upward pressure on inflation. Between March 2001 and March 2006, the Bank of Japan targeted the reserves ('current account balances') of commercial banks at the Bank of Japan, which at times resulted in massive growth in the monetary base and M1.¹⁰ Is this among the factors which prevented deflation from accelerating?

On the one hand, high growth in narrow monetary aggregates has not translated into high growth rates of broader aggregates such as M2, a fact which is plausibly related to a decline in bank lending which continued for several years after the banking crises of 1997 and 1998.¹¹ On the other hand, we cannot exclude the possibility that the Bank of Japan's policy of massive quantitative easing did prevent the output gap from becoming even more negative, and/or did keep agents from expecting more extreme deflation in the future. However, any such effects would already be reflected in the output gap and inflation expectations data which we discussed in the previous two subsections. As we documented, inflation expectations did turn moderately negative, notwithstanding expansionary monetary policy. Similarly, the output gap did grow sufficiently

⁹We regress Consensus forecasts or OECD forecasts on a constant and lagged inflation, and test for all potential break dates starting in 1995. The result is subject to data limitations: quarterly Consensus forecasts are only available from about 1990, and the OECD forecasts pertain to annual inflation.

¹⁰As it did before March 2001, the Bank of Japan now uses the uncollateralized overnight call rate as its main policy instrument.

¹¹Growth in lending by domestic commercial banks has been negative throughout the period 1998-2004, and only turned unambiguously positive in early 2006.

negative to warrant accelerating deflation in a linear framework.

3.3 Downward nominal wage rigidity?

The explanation in this subsection deals with the specification of the Phillips curve, but it alters the standard model in a different way than by allowing for time-variation in the slope of the short-run Phillips curve.

Akerlof, Dickens, and Perry (1996) develop a model in which downward nominal wage rigidity implies a convex *long-run* Phillips curve at inflation rates below 3%. The lower the inflation rate, the larger is the fraction of firms which can implement desired real wage cuts only through a reduction in the nominal wage. In the presence of downward nominal wage rigidity (DNWR), a lower inflation rate thus implies that a larger fraction of firms is forced to pay real wages exceeding the wage which they deem optimal. In the model of Akerlof, Dickens, Perry (1996), this increases the long-run sustainable level of unemployment, an effect which becomes more pronounced as inflation falls further below 3%.

For Japan, this story implies that, if DNWR exists, actual unemployment does not exceed its long-run rate by as much as unemployment gap estimates based on the assumption of a vertical and linear long-run Phillips curve would indicate.

However, wages have not been downwardly rigid in Japan during the period of our focus. At a micro level, Kuroda and Yamamoto (2003 a,b) find evidence for DNWR with data spanning 1992-1998. In a more recent study however, Kuroda and Yamamoto (2005) find no evidence for downward rigidity in the nominal wages of full-time workers during the period 1998-2001. Since full-time workers' nominal wages started being cut in 1998, downward nominal wage rigidity can hardly explain the absence of accelerating deflation, which became a puzzle at exactly that time.

At a macro level, Japan's wages are even less rigid. As described in Morgan (2005), the fraction of non-standard employees, such as part-time and temporary workers, has increased from 19.4% in 1995 to 29.0% in 2004. Furthermore, there is a large wage gap between regular and non-standard employees. In 2004, a part-time worker's hourly base wage was only 40.5% that of a typical regular worker.¹² Hence, even if no single group of workers had experienced nominal wage cuts, the shift from regular to non-standard workers had led to a decline in the aggregate wage.

Since both micro-economic and macro-economic data suggest that wages were not downwardly rigid during our period of interest, any story involving downward nominal wage rigidity is unlikely to explain the absence of accelerating deflation in Japan.

4 Evidence for a flattening Phillips curve

From this point on, our paper focuses on the path and determinants of the slope of the short-run output-inflation tradeoff.

The present section presents two findings. First, structural stability tests suggest that the slope of the Phillips curve has changed over the sample in a statistically significant fashion. Given that result, we estimate the Phillips curve slope as a time-varying parameter using the Kalman filter. Our results suggest that the Phillips curve has flattened over the sample, where much of the flattening occurred before the nineties.

4.1 Significant structural change in the Phillips curve slope

Redefining $\gamma_2 = p\gamma_1$, we rewrite the linear Phillips curve from equation (1) as:

¹²The overall monthly cost (including bonuses, fringe benefits, social security contributions, and training expenses) of employing a part-time worker was 36.9% that of employing a full-time worker. Data are from Morgan (2005).

$$\pi_t = \beta_1 \pi_{t-1} + \beta_2 \pi_{t-2} + \beta_3 \pi_{t-3} + \beta_4 \pi_{t-4} + \gamma_1 (ygap_{t-1} + p ygap_{t-2}) + \delta impoil_t + e_t \quad (2)$$

In order to assess the presence of structural change in the slope of the Phillips curve, we test for the stability of γ_1 , while assuming that p remains constant at its estimated value. This procedure is directly comparable to that of subsection 4.2., where we model γ_1 as a time-varying parameter, but continue to estimate p as being time-invariant. We motivate the assumption of time-invariance in p in subsection 4.2.

First, we use a standard dummy variable procedure to test for a structural break in the Phillips curve slope, at a hypothesized break date of 1990Q1.¹³ The test result suggests that, conditional on the absence of structural change in the other parameters, γ_1 was significantly smaller, at the 1% level, from 1990 onwards than before that time. Earlier research on the flattening of the Japanese Phillips curve, such as Nishizaki and Watanabe (2000) and Mourougane and Ibaragi (2004), similarly implemented dummy variable tests and found that the Japanese Phillips curve was significantly flatter in the nineties than it was in earlier decades.

Second, we implement a test for structural change which, unlike the above test, does not require us to assume any particular break date. In particular, we apply the Nyblom (1989) test, in the version developed by Hansen (1992). We reject the null hypothesis of time-invariance in γ_1 at the 1% level.¹⁴ It deserves emphasis that, unlike what is the case for other structural stability tests such as the above dummy variable test, rejection of the null hypothesis does not necessarily

¹³We regress $\pi_t = \beta(L) \pi_t + (\gamma_1 + \gamma'_1 breakdum) (ygap_{t-1} + p ygap_{t-2}) + \delta impoil_t + e_t$, where $breakdum = 1$ for all quarters starting in 1990Q1, and 0 for all earlier observations. γ'_1 is estimated to be negative and significant.

¹⁴The joint test statistic for all model parameters suggests significant structural change in the overall Phillips curve, at the 1% level. Furthermore, the relevant individual test statistic suggests significant change in the variance of the error term, at the 1% level. On this point, note that the Hansen (1992) test is asymptotically robust to heteroskedasticity. Throughout the paper, our OLS/NLS regressions use heteroskedasticity robust standard errors.

imply a one-time jump in the Phillips curve slope, yet could just as well reflect gradual structural change. In fact, the Nyblom test has optimal power against the hypothesis that a parameter follows a martingale. Given the test result, we model γ_1 as a random walk in the next subsection.

4.2 State-space form of model with time-varying Phillips curve slope

We write the model in state-space form. The measurement equation, in the form of Harvey (1994):

$$\pi_t = [ygap_{t-1} + pygap_{t-2}] \gamma_{1,t} + [\beta(L) \pi_t + \delta impoil_t] + e_t \quad (3)$$

Where $\gamma_{1,t}$ is the state variable. The error term e_t is normally distributed with mean zero and variance σ_e^2 .

In equation (3), we impose the restriction that the two output gap coefficients are proportional at any point of time, i.e. $\gamma_{2,t} = p \gamma_{1,t}$, where p is a time-invariant parameter to be estimated by Maximum Likelihood. The estimation results are similar whether we impose proportionality or not, except for the fact that the sum of the output gap coefficients is imprecisely estimated if the assumption of proportionality is not imposed. We simply do not have enough observations to obtain precise estimates for the sum of two distinct time-varying coefficients. In any case, we are primarily interested in the sum of the output gap coefficients, and less so in the precise way in which this sum is allocated over the two individual output gap coefficients.

While equation (3) represents the measurement equation, the transition equation is as follows:

$$\gamma_{1,t} = \gamma_{1,t-1} + v_{1,t} \quad (4)$$

Where $v_{1,t}$ is normally distributed with mean zero and variance $\sigma_{v_1}^2 = \sigma_e^2 q_1$. The parameter

q_1 is the signal-to-noise ratio for the coefficient on the output gap's first lag. The path of the coefficient on the second output gap lag follows from $\gamma_{2,t} = p \gamma_{1,t}$.

4.3 Estimation procedure and results

We apply Maximum Likelihood to estimate the model constituted by equations (3) and (4). As in Harvey (1994) and Kim and Nelson (1999), we compute the log likelihood function, in its prediction error decomposition form, from the Kalman filter prediction errors and their variances. We maximize the log likelihood function with respect to the hyperparameters.¹⁵ Finally, we use the Kalman filter run that maximized the likelihood in order to compute Kalman smoothed estimates of the time-varying output gap coefficients and their sum.

Table 1 compares the results of the time-varying coefficients linear Phillips curve with those of a linear Phillips curve estimated by OLS. Both estimations, as well as all other estimations in the remainder of this paper, are carried out over 1971Q2-2004Q4, where data for 1970Q2-1971Q1 are used to construct lags. The MLE estimates of the time-invariant parameters are comparable to their OLS counterparts. Similarly, the average of the sum of the output gap coefficients is virtually identical to the sum of the output gap coefficients implied by the OLS estimation. The sacrifice ratio's are plausible in both cases. In the MLE case, a disinflation of one percentage point requires output to be 1.39% below potential for four quarters. This is in line with earlier estimates of the Japanese sacrifice ratio in Ball (1994) and Zhang (2005).¹⁶

Figure 4 graphs Kalman smoothed estimates of the output gap coefficients and their sum, along with a 95% confidence interval. The sum of the output gap coefficients declines gradually

¹⁵We use the Matlab-function `fminunc` to optimize the log likelihood function. We set the signal-to-noise ratio, q_1 , to 1/1600 in the baseline. The parameter estimates reported in Table 1 are robust for all values of q_1 up to 1/25.

¹⁶Ball (1994) computes a sacrifice ratio for Japan of 0.93%. Our slightly larger estimate is in line with a continued flattening of the Phillips curve after 1994. Zhang (2005) computes a sacrifice ratio of 1.85% when accounting for long-lived effects.

over the sample. The results suggest that much of the flattening occurred before the nineties. The absence of accelerating deflation in 1998-2002 is only one among the episodes consistent with the time-path of the Phillips curve slope. For example, the finding that the Phillips curve was already relatively flat during the bubble period in the late eighties is in line with the fact that inflation remained surprisingly moderate at that time, notwithstanding large positive output gaps.

5 Why did the Phillips curve flatten? Candidate types of non-linearity

The flattening of a linear Phillips curve may suggest that the output-inflation trade-off should actually be modeled as a nonlinear relationship. In this section, we assess the empirical performance of three types of nonlinearity. The coefficient estimates from nonlinear Phillips curve regressions are in line with each of the nonlinear theories. Moreover, while section 4 detected statistically significant structural change in the output gap coefficients of a linear Phillips curve, there is no significant structural change in the coefficients for any of the three types of nonlinearity.

5.1 A convex short-run Phillips curve due to capacity constraints?

Laxton, Meredith, Rose (1995) and related papers¹⁷ allow for convexity in the short-run Phillips curve. In Laxton, Meredith, Rose (LMR), capacity constraints constitute the economic rationale for nonlinearity. Suppose that at the current level of output, firms are operating near the capacity constraint. In such a situation, any increase in aggregate demand can hardly be met by increased production. As such, the increase in demand translates almost uniquely into an increase in inflation, even in the short run. Hence, the Phillips curve is nearly vertical near the

¹⁷See footnote 2 for references to papers related to any of the theories of nonlinearity.

capacity constraint, where the slope becomes gradually steeper as the economy moves towards the capacity constraint. This story implies a vertical asymptote in the Phillips curve at the capacity constraint. The baseline functional form which LMR use implies that, if convexity is present, it exists along the entire Phillips curve.

Note that it is not obvious whether the presence of capacity constraints can be a rationale for convexity in the Phillips curve in regions which are far away from the capacity constraint. The answer to this question is particularly important for our purposes: the LMR model's predictions for Japan are that, since the economy was far from the capacity constraint in 1998-2002, Japan was on a flatter part of a convex Phillips curve during that period. That would explain the flattening which we observed in section 4. Yet, if convexity is not present at negative output gaps, the Japanese economy would have moved along a linear part of the Phillips curve for most of the nineties, such that the LMR model could not explain any time-variation in the Phillips curve slope during that period. There are surely ways to motivate convexity at negative output gaps,¹⁸ but such reasonings are not contained in LMR's original paper.

We follow LMR in using a functional form which implies that the Phillips curve is either convex in all regions, or linear everywhere. The absence of a clear theoretical motivation for convexity at negative output gaps will enter our overall model assessment in section 7.2.

We estimate a potentially nonlinear Phillips curve by Nonlinear Least Squares, where the functional form of the output gap terms is equivalent to that in LMR:

$$\pi_t = \beta(L) \pi_t + \gamma_1 \left[\left(\frac{\phi ygap_{t-1}}{\phi - ygap_{t-1}} \right) + p \left(\frac{\phi ygap_{t-2}}{\phi - ygap_{t-2}} \right) \right] + \delta impoil_t + e_t \quad (5)$$

¹⁸In the presence of sectoral heterogeneity, it is possible that even at negative output gaps, a small fraction of firms operates near full capacity. If so, it is plausible that the fraction of capacity-constrained firms increases as the output gap becomes less negative (or more positive).

Where γ_1 is time-invariant. For equation (5), the Nyblom test fails to reject the null hypothesis of structural stability in γ_1 , at the 10% level. That is, there appears to be little to no time-variation in the Phillips curve slope beyond that implied by the nonlinearity of the functional form.

The crucial parameter to be estimated is ϕ . This parameter indicates the level of the output gap at which the economy reaches the capacity constraint. By the same token, ϕ governs the degree of nonlinearity in the Phillips curve. The smaller the point estimate for ϕ is, the smaller the distance between the zero output gap and the capacity constraint will be. This in turn yields a higher degree of convexity.

The rightmost column of Table 2 presents estimation results for equation (5). We also include results from a purely linear Phillips curve, which essentially imposes the restriction that $\phi = \infty$. In the potentially nonlinear case, ϕ is precisely estimated, with a point estimate of exactly 10.00. This suggests that the economy would reach the capacity constraint if actual output were to exceed potential output by 10%.

To illustrate the degree of convexity implied by the estimates for γ_1 , p , and ϕ in equation (5), Figure 5 graphs the sum of the output gap terms as a function of the output gap. In particular, the bold curve in Figure 5 graphs $\gamma_1 [(\phi ygap_{t-1} / (\phi - ygap_{t-1})) + p (\phi ygap_{t-2} / (\phi - ygap_{t-2}))]$ with respect to the output gap, where we impose that $ygap_{t-1} = ygap_{t-2}$. For comparison, the thin solid line in the same figure graphs the same function, with exactly the same values for γ_1 and p , but imposing that $\phi = \infty$. Visually, we see a fairly strong degree of nonlinearity in the Phillips curve.¹⁹ In other words, booms in real activity increase inflation by more than recessions decrease it. The asymmetry in the effects of demand shifts becomes more pronounced as one moves further from the zero output gap in either direction. For instance, an output gap of -5% tends to lead to

¹⁹The dotted line also uses the same values for γ_1 and p , but assumes a value for ϕ which is the upper bound of the 95% confidence interval around the estimated nonlinearity parameter.

a disinflation of 0.53 percentage points after two quarters, while the total impact of a 5% output gap is to increase inflation by 1.60 percentage points.

5.2 A flatter Phillips curve due to a lower frequency of price adjustment?

In this subsection, we assess the empirical validity of two theories in which costs of price adjustment lead firms to adjust their output prices infrequently: Ball, Mankiw, Romer (1988), and Dotsey, King, Wolman (1999). In both models, lower trend inflation decreases the frequency of price adjustment. Less frequent price adjustment in turn reduces the effect of aggregate demand shifts on inflation. That is to say, the Phillips curve is flatter at lower rates of trend inflation.

In Ball, Mankiw, Romer (BMR), firms, when setting their price, also choose the length of time over which their price will be in effect. Firms minimize a loss function which depends on the average cost of price adjustment per period, and on deviations of their actual nominal price from the profit-maximizing nominal price over the course of the period that the price is in effect. When trend inflation is high, any firm expects its relative price to change rapidly over time, which in turn leads the firm to expect a rapid change in its profit-maximizing nominal price. Thus, the forward-looking firm will not fix its actual price for a long time. Instead, the firm opts for more frequent price adjustment, thus paying a higher per-period cost of price adjustment, in order to avoid large deviations of its future prices from their profit-maximizing levels.

In Dotsey, King, Wolman (DKW), higher steady-state inflation implies that any firm's relative price has been eroded to a larger extent since its last price adjustment. This implies that for a larger fraction of firms, the benefit of price updating will exceed the (labor) cost of price adjustment. In conclusion, higher steady-state inflation leads to higher steady-state probabilities of price adjustment.

These two models' prediction for Japan is that, as trend inflation gradually decreased over the sample, the frequency of price adjustment declined, which in turn led to a gradually flattening Phillips curve. At the time of writing, we are not aware of any publicly available time series data on the average frequency of price adjustment for Japan. We therefore cannot model the frequency of price adjustment explicitly, much like what was the case for the empirical analyses in the earlier studies referenced in the introduction. However, we can test whether the slope of the Phillips curve depends positively on trend inflation.

Inspired by DeFina (1991), we adopt a one-step approach.²⁰ That is to say, we estimate a Phillips curve in which the slope depends on the absolute value of trend inflation.²¹

$$\pi_t = \beta(L) \pi_t + [a + b |\bar{\pi}_t|] [ygap_{t-1} + p ygap_{t-2}] + \delta impoil_t + e_t \quad (6)$$

We generate trend inflation at time t as a geometric average of J quarters of past inflation:

$$\bar{\pi}_t = \frac{1 - \theta}{\theta - \theta^{J+1}} \sum_{j=1}^J \theta^j \pi_{t-j} \quad (7)$$

In the baseline, $\theta = 0.93$ and $J = 71$. Note that trend inflation does not depend on current inflation, so as to avoid endogeneity issues in equation (6). The factor in front of the summation sign ensures that the sum of the weights on the past inflation terms is equal to one.

In equation (6), the Nyblom test does not detect any structural change in a or b individually,

²⁰ Empirical findings on the relationship between the slope of the Phillips curve and trend inflation or aggregate volatility, as referenced in the introduction, have mostly been based on a two-step approach. In a time-series setting, it is undesirable to enter the Phillips curve slope as a left-hand side variable in a second-stage regression, among others because the time-varying Phillips curve slope (obtained, say, from rolling windows regressions) is likely to be nonstationary.

²¹ We take the *absolute value* of trend inflation based on the intuition that the effects of more pronounced deflation should affect firms' relative prices, and hence the frequency of price adjustment, in much the same way as an increase in inflation does. Neither DKW, BMR, nor DeFina (1991) take the absolute value, yet this can be attributed to their dealing with economies in which negative trend inflation could hardly be imagined at that time.

or a and b jointly, at the 10% level. This suggests that there is little to no time-variation in the output gap coefficients beyond that associated with changes in trend inflation.

Table 3 displays the estimation results for equation (6). For comparison, we include results from a Phillips curve in which the coefficient on trend inflation is set to zero. In equation (6), the coefficient on trend inflation, b , is positive and significant at the 1% level. This result is in line with the Ball-Mankiw-Romer and Dotsey-King-Wolman theories.

From the estimates for a , b , and p , and our series for trend inflation, we computed the implied output gap coefficients and their sum. Given that the Phillips curve slope is a linear transformation of trend inflation, it displays a similar pattern over time as trend inflation itself. In particular, the sum of the output gap coefficients (not graphed here) increases until 1976, and decreases quickly through the late eighties. From the early nineties on, the Phillips curve slope still tends to decrease, but at a slower pace. It falls below zero in 1994, yet from 1996 on remains fairly stable at moderately negative levels.

5.3 A flatter Phillips curve due to a decline in aggregate volatility?

In Lucas (1973), the slope of the Phillips curve depends on the volatility of aggregate demand and supply shocks. Firms set quantities produced based on their perceived relative price. As the variance of aggregate shocks decreases relative to the variance of firm-specific shocks, a larger fraction of any change in the overall price level is misperceived by firms as being a change in their relative price. In that way, lower aggregate volatility implies that any change in aggregate demand has a larger effect on a typical firm's production, and thus on aggregate output. Correspondingly, demand shifts have a smaller impact on inflation. In conclusion, low levels of aggregate volatility imply a flatter Phillips curve.

A testable implication of this model for Japan is that a decrease in the variance of aggregate demand and/or supply shocks would have been associated with the flattening of the Phillips curve which we documented in section 4.²²

We capture aggregate volatility by the variance of inflation.²³ We estimate a Phillips curve in which the slope is explicitly modeled as a function of the variance of inflation:

$$\pi_t = \beta(L) \pi_t + [c + d \text{var}_t(\pi)] [ygap_{t-1} + p ygap_{t-2}] + \delta \text{impoil}_t + e_t \quad (8)$$

We generate the variance of inflation at time t as a geometrically weighted average of past squared deviations of inflation from its trend:

$$\text{var}_t(\pi) = \frac{1 - \theta}{\theta - \theta^{J+1}} \sum_{j=1}^J \theta^j (\pi_{t-j} - \bar{\pi}_t)^2 \quad (9)$$

Where trend inflation $\bar{\pi}_t$ is computed as in equation (7). Again, the baseline values are $\theta = 0.93$ and $J = 71$.

In equation (8), the Nyblom test rejects the null of no structural change in c individually, and in c and d jointly, but only at the 10% level. It fails to reject the hypothesis of time-invariance in d . This suggests that changes in the variance of inflation explain most, but not all, of the time-variation in the sum of the output gap coefficients.

Table 4 contains the estimation results for equation (8). As predicted by the Lucas-theory, the coefficient d on inflation volatility is positive and significant at the 1% level.

²²Ball, Mankiw, Romer (1988) equally implies that a decrease in the variance of aggregate shocks leads to a flatter Phillips curve. Yet, in BMR, the mechanism works through the frequency of price adjustment: declining aggregate volatility, which reduces uncertainty about future optimal prices, enables firms to set their prices for a longer period of time. A lower frequency of price adjustment in turn implies a flatter Phillips curve.

²³The other standard candidate, the variance of nominal GDP, would not be as appropriate a measure to capture both supply and demand shocks. For instance, if aggregate demand is unit-elastic, aggregate supply shocks have no visible impact on nominal GDP, since their effect on prices is exactly offset by their effect on real activity.

The sum of the output gap coefficients implied by equation (8) increases steeply from 1973 to 1975, then decreases quickly through the late eighties. From the early nineties on, the implied Phillips curve slope decreases only slightly. It stays positive at all times.

6 Do the nonlinear models beat the random walk model?

In the present section, we estimate models in which the Phillips curve slope depends on a random walk term as well as on a function implied by a particular theory of nonlinearity.

In the cases where we test for it, we find that we can omit the random walk term from the encompassing model without engendering a significant decline in the value of the likelihood function. This is related to our finding, in section 5, of no structural change in the coefficients on the output gap terms for each of the nonlinear models. Beyond that, the present section's results imply that adding any of the three types of nonlinearity to a pure random walk model yields a significant improvement in the fit.

6.1 The trend inflation model beats the random walk model

We estimate a model which encompasses the random walk model and the Ball-Mankiw-Romer / Dotsey-King-Wolman (BMR / DKW) trend inflation model, and test for the statistical relevance of the random walk term on the one hand, and the trend inflation term on the other hand.

The Phillips curve, alias measurement equation of the state-space model, is exactly the same as equation (3):

$$\pi_t = [ygap_{t-1} + pygap_{t-2}] \gamma_{1,t} + [\beta(L) \pi_t + \delta impoil_t] + e_t \quad (10)$$

The novelty lies in the transition equation. In the encompassing model, the output gap coeffi-

cient $\gamma_{1,t}$ is allowed to depend both on its own lag and on trend inflation. In the pure random walk model, $\gamma_{1,t} = \gamma_{1,t-1} + v_{1,t}$. On the other hand, in the pure BMR / DKW model, $\gamma_{1,t} = a + b\bar{\pi}_t$. Nesting these two yields the first row of the state equation:

$$\gamma_{1,t} = \lambda(\gamma_{1,t-1} + v_{1,t}) + (1 - \lambda)(a + b\bar{\pi}_t) \quad (11)$$

Note that trend inflation appears as an exogenous variable in the first row of the transition equation. Textbook treatments of the Kalman filter such as Harvey (1994), Hamilton (1994), or Kim and Nelson (1999) do not discuss solutions for how to enter an exogenous variable in the state equation. If we wish to enter $\bar{\pi}_t$ in the transition equation, we need to specify a transition process for trend inflation, and enter this process in the second row of the state equation. We derive such process from the definition of trend inflation in equation (7). For θ sufficiently small and J converging to infinity, we find:

$$\bar{\pi}_{t+1} = (1 - \theta)\pi_t + \theta\bar{\pi}_t \quad (12)$$

The transition equation thus becomes:

$$\begin{bmatrix} \gamma_{1,t} \\ \bar{\pi}_{t+1} \end{bmatrix} = \begin{bmatrix} (1 - \lambda)a \\ (1 - \theta)\pi_t \end{bmatrix} + \begin{bmatrix} \lambda & (1 - \lambda)b \\ 0 & \theta \end{bmatrix} \cdot \begin{bmatrix} \gamma_{1,t-1} \\ \bar{\pi}_t \end{bmatrix} + \begin{bmatrix} \lambda v_{1,t} \\ 0 \end{bmatrix} \quad (13)$$

First, we estimate the encompassing model, consisting of equations (10) and (13). In the unrestricted model, λ is estimated to be -0.51. Essentially, the weight on the BMR / DKW model in the transition equation exceeds unity.

Next, we restrict $\lambda = 0$, in which case the model reduces to the BMR / DKW trend inflation

model. We test that restriction by means of a likelihood ratio test. Since we are testing one restriction, the likelihood ratio statistic has a $\chi^2(1)$ distribution. According to the test result, relaxing the restriction that $\lambda = 0$ does not significantly improve the fit, not even at the 10% level.

Finally, we restrict the model such that equation (11) reduces to a random walk. In this case, the likelihood ratio has a $\chi^2(3)$ distribution.²⁴ Abstracting from the BMR / DKW term significantly deteriorates the fit at the 1% level.

On the one hand, we found that the pure BMR / DKW model does not perform significantly worse than the encompassing model. On the other hand, the pure random walk model does perform significantly worse than the encompassing model. We conclude that the BMR / DKW model provides a more accurate description of the data than the random walk model does.

6.2 The misperceptions model beats the random walk model

This subsection implements a similar procedure for the Lucas misperceptions model as the previous subsection did for BMR / DKW. The measurement equation is exactly the same as equation (10).

The state equation is analogous to equation (13), but is somewhat complicated by the fact that the transition process for the variance of inflation is more involved than the process for trend inflation. The first row of the transition equation is analogous to equation (11):

$$\gamma_{1,t} = \lambda(\gamma_{1,t-1} + v_{1,t}) + (1 - \lambda)[c + d \text{var}_t(\pi)] \quad (14)$$

From the definition of the variance of inflation in equation (9), we derive its transition process,

²⁴At first sight, the distribution appears to be nonstandard, since a and b potentially act as nuisance parameters. Yet, rewrite equation (11) as $\gamma_{1,t} = \gamma_{1,t-1} + v_{1,t} + (\lambda - 1)(\gamma_{1,t-1} + v_{1,t}) + (1 - \lambda)a + (1 - \lambda)b\pi_t$. Redefining $(1 - \lambda)a$ and $(1 - \lambda)b$ such that they are parameters in their own right, this equation is in effect linear in the parameters. It reduces to the random walk model after imposing three restrictions: $(\lambda - 1) = 0$, $(1 - \lambda)a = 0$, and $(1 - \lambda)b = 0$.

to be included in the second row of the state equation. For θ sufficiently small and J converging to infinity, we find that:

$$var_{t+1}(\pi) = (1 - \theta) X_t + \theta var_t(\pi) \quad (15)$$

Where $X_t = (2 - \theta) (\pi_t - \bar{\pi}_t)^2 - 2 \frac{1-\theta}{\theta} (\pi_t - \bar{\pi}_t) \sum_{j=1}^J \theta^j (\pi_{t-j+1} - \bar{\pi}_t)$.

The transition equation thus becomes:

$$\begin{bmatrix} \gamma_{1,t} \\ var_{t+1}(\pi) \end{bmatrix} = \begin{bmatrix} (1 - \lambda) c \\ (1 - \theta) X_t \end{bmatrix} + \begin{bmatrix} \lambda & (1 - \lambda) d \\ 0 & \theta \end{bmatrix} \cdot \begin{bmatrix} \gamma_{1,t-1} \\ var_t(\pi) \end{bmatrix} + \begin{bmatrix} \lambda v_{1,t} \\ 0 \end{bmatrix} \quad (16)$$

Where X_t is as defined under equation (15).

In the encompassing model, which consists of equations (10) and (16), the nesting parameter λ is not significantly different from zero, with a point estimate of -0.13. This is evidence in favor of the Lucas model, relative to the random walk model. As in the previous subsection, we find that imposing the restriction that $\lambda = 0$ does not significantly worsen the fit, while imposing restrictions such that equation (14) reduces to a random walk leads to a significant decline in the log likelihood function value at the 1% level.

In conclusion, the random walk model provides a significantly less accurate fit than the encompassing model, while the fit of the Lucas model is statistically indistinguishable from that of the encompassing model. Hence, the misperceptions model beats the random walk model.

6.3 Convexity even with independent time-variation in the Phillips curve slope

In section 5, we detected a strong degree of nonlinearity in a Phillips curve modeled as in Laxton, Meredith, Rose (LMR). However, that section's procedure is not designed to determine whether the nonlinearity is statistically significant. In the present subsection, we do find that a model

which allows for LMR-style nonlinearity captures the evolution of the Phillips curve slope in a significantly better fashion than a pure random walk model does.

We specify a Phillips curve which nests the linear time-varying coefficient model of equations (3) and (4), and the nonlinear LMR model of equation (5):

$$\pi_t = \beta(L) \pi_t + \gamma_{1,t} \left[\left(\frac{\phi ygap_{t-1}}{\phi - ygap_{t-1}} \right) + p \left(\frac{\phi ygap_{t-2}}{\phi - ygap_{t-2}} \right) \right] + \delta impoil_t + e_t \quad (17)$$

Where $\gamma_{1,t}$ evolves as a random walk:

$$\gamma_{1,t} = \gamma_{1,t-1} + v_{1,t} \quad (18)$$

This model collapses to the linear time-varying coefficients model if $\phi = \infty$, and reduces to the nonlinear model with time-invariant coefficients if $v_{1,t} = 0$ for all t .

We estimate two models, one in which the nonlinearity parameter ϕ is restricted to be a very large number,²⁵ and one in which ϕ is freely estimated.

In the unrestricted model, the nonlinearity parameter is small and precisely estimated, be it somewhat larger than in section 5.

We apply a likelihood ratio test to examine whether the model in which ϕ is freely estimated performs significantly better than the model in which linearity is imposed. The likelihood ratio statistic is distributed $\chi^2(1)$. The test result suggests that relaxing the assumption of linearity significantly increases the value of the likelihood function at the 1% level.

In conclusion, the LMR model adds information beyond that contained in the linear time-varying coefficients model.

²⁵We impose $\phi = 1E20$.

7 Which type of nonlinearity in the Phillips curve?

So far, we have found that each of the three models of nonlinearity not only is consistent with the data, but also outperforms the benchmark random walk model. In the present section, we compare the three nonlinear theories' success in explaining the flattening of Japan's Phillips curve.

7.1 Non-nested model fit comparison

We perform three hypothesis testing procedures to compare the performance of the nonlinear models.

First, we regress Phillips curves which nest two nonlinear models. For example, the following equation nests the Laxton, Meredith, Rose (LMR) model from equation (5) and the Ball-Mankiw-Romer / Dotsey-King-Wolman (BMR / DKW) model of equation (6):

$$\pi_t = \beta(L) \pi_t + [a + b |\bar{\pi}_t|] \left[\left(\frac{\phi ygap_{t-1}}{\phi - ygap_{t-1}} \right) + p \left(\frac{\phi ygap_{t-2}}{\phi - ygap_{t-2}} \right) \right] + \delta impoil_t + e_t \quad (19)$$

As it turns out, the coefficient on trend inflation b is positive and significant at the 1% level, which is in line with the BMR / DKW model. The nonlinearity parameter ϕ is estimated to be 12.94, with a somewhat larger standard error than in sections 5 or 6, which all in all suggests that the LMR-convexity still plays a role.

The results with a Phillips curve nesting the LMR- and Lucas-models are similar: there is evidence for both models. On the other hand, regressing a Phillips curve in which the output gap coefficients depend on both trend inflation and the variance of inflation does not yield conclusive results. To see why, note that the correlation between trend inflation and the variance of inflation is 0.96. In the presence of multicollinearity, it is not surprising that both variables enter

insignificantly.

Second, we discuss results from pairwise non-nested tests as developed by Davidson and MacKinnon (1981). The LMR convexity turns out to perform poorly relative to the other two models. The coefficient on the fitted value from the BMR / DKW model, when added to a LMR regression, is significant at the 5% level. This suggests that the BMR / DKW model adds information beyond that contained in the LMR model. An analogous result holds when we add the fitted value from the Lucas model to the LMR model. On the other hand, Davidson-MacKinnon tests favor the BMR / DKW and Lucas models. The fitted value from the LMR model does not enter significantly in either model.

Third, to assess which among the models in our model space is most likely to correspond to the truth, we apply Bayesian model averaging methods as in Brock, Durlauf, and West (2004). In particular, we use the procedure in Kiley (2005) to compute pseudo-posterior model odds based on a comparison of the Bayesian Information Criteria from the three nonlinear models and the linear model. This procedure assumes a uniform prior distribution over the model space. As Table 5 documents, the results strongly favor the BMR / DKW endogenous pricing model. According to the pseudo-posterior distribution, the probability that the BMR / DKW model is the true model is 81.42%. The pseudo-posterior model odds for the Lucas model are 18.58%. The probability for either the LMR model or the linear model to be the true model is virtually zero.

7.2 Assessment

Two out of three procedures yielded conclusive results. Davidson-MacKinnon tests, as well as the computation of pseudo-posterior model odds, suggested that the Laxton, Meredith, Rose (LMR) model provides a less accurate description of the data than the two other models.

In this subsection, we take a more detailed look at the regression results from section 5, so as to examine why the LMR model performed poorly in the non-nested model hypothesis tests. Before doing so, remember from section 5.1 that it is doubtful whether capacity constraints can be a rationale for convexity at regions of the Phillips curve which are far from the capacity constraint.

It turns out that the accurate fit of the LMR model in a regression of equation (5) over 1971Q2-2004Q4 is mostly driven by its superior fit around the time of the first oil price shock. Much of the nonlinearity seems to spring from the 1974Q1 observation, when oil import prices surged, annualized core CPI inflation spiked to 32%, and the pre-1974 boom suddenly halted. As a robustness test, we perform regressions for the three nonlinear models as in section 5, but over a sample which excludes all pre-1975 observations. It turns out that there is no evidence for LMR-type convexity over the sample 1975Q1-2004Q4. More precisely, the standard error on the nonlinearity parameter ϕ is that large that no inference can be drawn as to whether the Phillips curve is convex or linear. In contrast, the results for the BMR / DKW and Lucas models are robust to the exclusion of observations from the oil shock episode. Trend inflation and the variance of inflation, respectively, enter significantly at the 1% level even when pre-1975 observations are excluded.

8 Conclusion

At a direct empirical level, our paper investigates why deflation did not accelerate in Japan notwithstanding large negative output gaps during the period 1998-2002. We find that the absence of accelerating deflation cannot be adequately addressed by popular explanations which assume a linear short-run relation between the output gap and inflation with a time-invariant slope. Given that finding, the body of our paper focuses on the path and determinants of the Phillips curve

slope. We document a gradual, significant flattening of the Japanese Phillips curve which predates the nineties.

As for the determinants of such time-variation in the Phillips curve slope, our results favor the Ball-Mankiw-Romer / Dotsey-King-Wolman hypothesis that declining trend inflation created an environment in which prices became stickier, which in turn caused the Phillips curve to flatten. All but one of our tests lend equally strong support to the Lucas hypothesis that a decline in aggregate inflation volatility exacerbated firms' misperceptions about relative prices, implying a flatter Phillips curve. While stories such as Laxton-Meredith-Rose in which capacity constraints engender a convex short-run Phillips curve are consistent with the data, they perform poorly in comparison with the two other models.

On a broader level, our results are indicative for the appropriate theoretical framework to model the output-inflation trade-off. If it is indeed a general rule that the Phillips curve flattens as trend inflation declines, we see two implications for monetary policy makers in economies where trend inflation is low today relative to past experience.

First, the Ball-Mankiw-Romer / Dotsey-King-Wolman model implies that the Phillips curve in those countries is currently flatter than the standard linear model would suggest. All other things being equal, this implies a higher sacrifice ratio: if the central bank brings about a disinflation, the associated reduction in output will be larger than it would appear from the linear model.

Second, although the endogenous pricing models imply that the Phillips curve turns flatter as trend inflation declines to zero, these models do not predict that the risk of a deflationary spiral is negligible. On the contrary, both the Ball-Mankiw-Romer and Dotsey-King-Wolman models are symmetric around zero. Once trend inflation turns negative, these models imply that any further decrease in trend inflation is associated with an increase in the frequency of price adjustment.

This in turn means that any negative output gap has stronger deflationary effects, thus increasing the risk of a more pronounced negative interaction between deflation, real activity, and financial sector vulnerabilities.

Our analysis also lends some empirical support to the Lucas model. This model implies that, in economies where inflation is currently more stable than in earlier decades, the Phillips curve is flatter than the standard linear model would suggest. Hence, this model implies that in such economies, the short-run output costs of disinflation are higher than they would appear from a linear model.

References

- [1] Akerlof, George A., William T. Dickens, and George L. Perry (1996). “The Macroeconomics of Low Inflation.” *Brookings Papers on Economic Activity* 1, 1-76.
- [2] Alberro, Jose (1981). “The Lucas Hypothesis on the Phillips Curve: Further International Evidence.” *Journal of Monetary Economics* 7, 239-250.
- [3] Ball, Laurence (1994). “What Determines the Sacrifice Ratio?” In *Monetary Policy*, edited by N. Gregory Mankiw, pp. 155-193. Chicago, IL: University of Chicago Press.
- [4] Ball, Laurence, N. Gregory Mankiw, and David Romer (1988). “The New Keynesian Economics and the Output-Inflation Trade-Off.” *Brookings Papers on Economic Activity* 1, 1-82.
- [5] Bank of Japan (2006). “The New Estimates of Output Gap and Potential Growth Rate”, Bank of Japan Review Series No. 3.
- [6] Bayoumi, Tamim (2000). “Where Are We Going? The Output Gap and Potential Growth.” In *Post-Bubble Blues; How Japan Responded to Asset Price Collapse*, edited by Tamim Bayoumi and Charles Collins, pp. 89-106. Washington, DC: International Monetary Fund.
- [7] Blanchard, Olivier (2000). “Bubbles, Liquidity Traps, and Monetary Policy.” In *Japan’s Financial Crisis and its Parallels to U.S. Experience*, edited by Mikitani and Posen, pp. 185-193. Washington, DC: Institute for International Economics.
- [8] Bonomo, Marco, and Carlos Carvalho (2004). “Endogenous Time-Dependent Rules and Inflation Inertia.” *Journal of Money, Credit, and Banking* 36, 1015-1041.
- [9] Brock, William A., Steven N. Durlauf, and Kenneth D. West (2004). “Model Uncertainty and Policy Evaluation: some Theory and Empirics.” NBER Working Paper No. 10916.

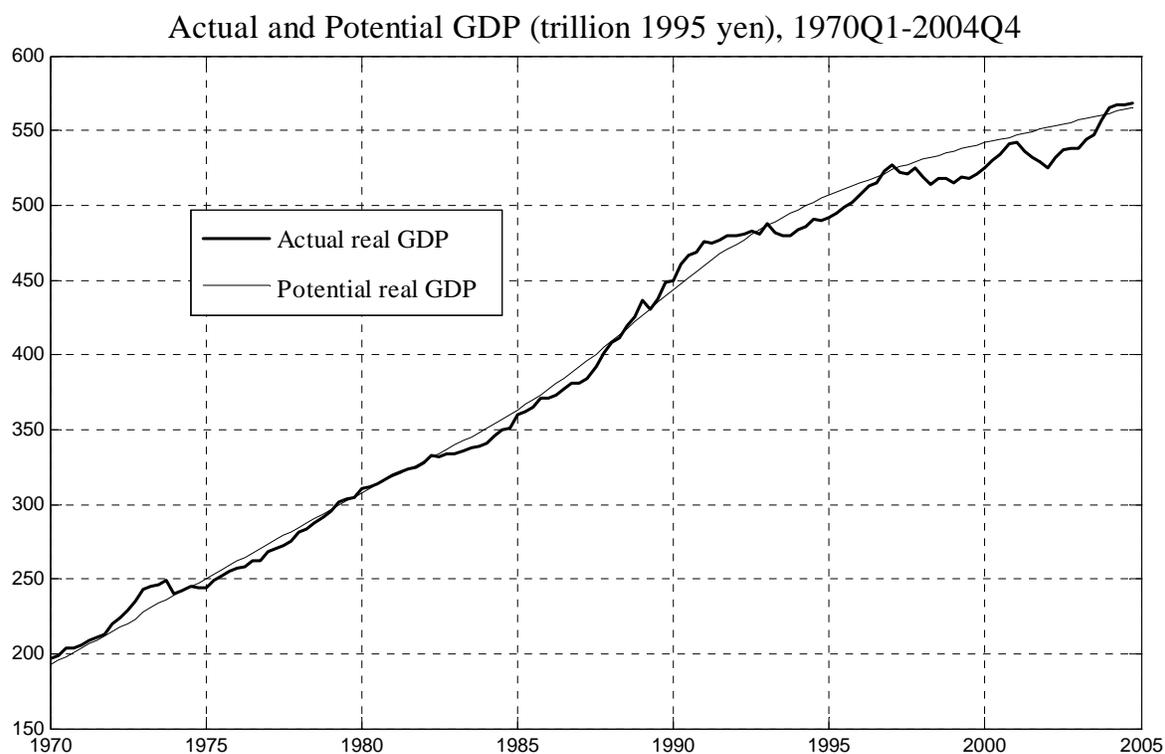
- [10] Burstein, Ariel T. (2006). "Inflation and Output Dynamics with State-Dependent Pricing Decisions." *Journal of Monetary Economics* 53, 1235-1257.
- [11] Davidson, Russell, and James G. MacKinnon (1981). "Several Tests for Model Specification in the Presence of Alternative Hypotheses." *Econometrica* 49, 781-793.
- [12] Debelle, Guy and Douglas Laxton (1996). "Is the Phillips Curve Really a Curve? Some Evidence for Canada, the United Kingdom, and the United States." IMF Staff Papers 44, 249-282.
- [13] DeFina, Robert H. (1991). "International Evidence on a New Keynesian Theory of the Output-Inflation Trade-Off." *Journal of Money, Credit, and Banking* 23, 410-422.
- [14] Dolado, Juan J., Ramon Maria-Dolores, and Manuel Naveira (2005). "Are Monetary Policy Reaction Functions Asymmetric? The Role of Nonlinearity in the Phillips Curve." *European Economic Review* 49, 485-503.
- [15] Dotsey, Michael, Robert G. King, and Alexander L. Wolman (1999). "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output." *Quarterly Journal of Economics* 114, 655-690.
- [16] Dotsey, Michael, Robert G. King, and Alexander L. Wolman (2007). "Aggregate Dynamics of Inflation and Real Activity with Firm-Level Productivity Shocks." Mimeo.
- [17] Froyen, Richard T., and Roger N. Waud (1980). "Further International Evidence on Output-Inflation Tradeoffs." *American Economic Review* 70, 409-421.
- [18] Gertler, Mark, and John Leahy (2006). "A Phillips Curve with an Ss Foundation." NBER Working Paper No. 11971.

- [19] Golosov, Mikhail, and Robert E. Lucas (2003). "Menu Costs and Phillips Curves." NBER Working Paper No. 10187.
- [20] Hamilton, James (1994). *Time Series Analysis*. Princeton, NJ: Princeton University Press.
- [21] Hansen, Bruce E. (1992). "Testing for Parameter Instability in Linear Models." *Journal of Policy Modeling* 14, 517-533.
- [22] Harvey, Andrew (1994). *Time Series Models*. Cambridge, MA: MIT Press.
- [23] Hess, Gregory D., and Kwanho Shin (1999). "Some Intranational Evidence on Output-Inflation Trade-offs." *Macroeconomic Dynamics* 3, 187-203.
- [24] Hirose, Yasuo, and Koichiro Kamada (2003). "A New Technique for Simultaneous Estimation of the Output Gap and Phillips Curve." *Bank of Japan Monetary and Economic Studies* 21, 93-112.
- [25] Ilmakunnas, Pekka, and Hiroki Tsurumi (1985). "Testing the Lucas Hypothesis on Output-Inflation Trade-offs." *Journal of Business and Economic Statistics* 3, 43-53.
- [26] Kamada, Koichiro (2005). "Real-Time Estimation of the Output Gap in Japan and its Usefulness for Inflation Forecasting and Policymaking." *North American Journal of Economics and Finance* 16, 309-332.
- [27] Kiley, Michael T. (2000). "Endogenous Price Stickiness and Business Cycle Persistence." *Journal of Money, Credit, and Banking* 32, 28-53.
- [28] Kiley, Michael T. (2005). "A Quantitative Comparison of Sticky-Price and Sticky-Information Models of Price Setting." *Computing in Economics and Finance* 2005 No. 183.

- [29] Kim, Chang-Jin and Charles Nelson (1999). *State-Space Models with Regime Switching*. Cambridge, MA: MIT Press.
- [30] Kormendi, Roger C., and Philip G. Meguire (1984). “Cross-Regime Evidence of Macroeconomic Rationality.” *Journal of Political Economy* 92, 875-908.
- [31] Kuroda, Sachiko, and Isamu Yamamoto (2003a). “Are Japanese Nominal Wages Downwardly Rigid?” *Bank of Japan Monetary and Economic Studies* 21, 1-30.
- [32] Kuroda, Sachiko, and Isamu Yamamoto (2003b). “Are Japanese Nominal Wages Downwardly Rigid?” *Bank of Japan Monetary and Economic Studies* 21, 31-68.
- [33] Kuroda, Sachiko, and Isamu Yamamoto (2005). “Wage Fluctuations in Japan after the Bursting of the Bubble Economy: Downward Nominal Wage Rigidity, Payroll, and the Unemployment Rate.” *Bank of Japan Monetary and Economic Studies* 23, 1-29.
- [34] Laxton, Douglas, Guy Meredith, and David Rose (1995). “Asymmetric Effects of Economic Activity on Inflation.” *IMF Staff Papers* 42, 344-374.
- [35] Laxton, Douglas, David Rose, and Demosthenes Tambakis (1999). “The U.S. Phillips Curve: The Case for Asymmetry.” *Journal of Economic Dynamics and Control* 23, 1459-1485.
- [36] Lucas, Robert E. (1973). “Some International Evidence on Output-Inflation Trade-Offs.” *American Economic Review* 63, 326-334.
- [37] Mourougane, Annabelle, and Hideyuki Ibaragi (2004). “Is There a Change in the Trade-Off between Output and Inflation at Low or Stable Inflation Rates? Some Evidence in the Case of Japan.” OECD Economics Department Working Paper No. 379.

- [38] Morgan, Peter (2005). "Employment - the Shift to Flexibility." HSBC special report, February 2005.
- [39] Nishizaki, Kenji, and Tsutomu Watanabe (2000). "Output-Inflation Trade-Off at Near-Zero Inflation Rates." *Journal of the Japanese and International Economies* 14, 304-326.
- [40] Nyblom, Jukka (1989). "Testing for the Constancy of Parameters Over Time." *Journal of the American Statistical Association* 84, 223-230.
- [41] Phillips, Alban W. (1958). "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957." *Economica* 25, 283-299.
- [42] Schaling, Eric (2004). "The nonlinear Phillips Curve and Inflation Forecast Targeting: Symmetric Versus Asymmetric Monetary Policy Rules", *Journal of Money, Credit, and Banking* 36, 361-386.
- [43] Turner, Dave (1995). "Speed Limit and Asymmetric Inflation Effects from the Output Gap in the Major Seven Economies." *OECD Economic Studies* 24, 57-87.
- [44] Zhang, Lawrence Huiyan (2005). "Sacrifice Ratios with Long-Lived Effects." *International Finance* 8, 231-262.

Figures and tables



Source: US Federal Reserve, Board of Governors

Figure 1

Note: since the Fed's recent potential output estimates are confidential, we extrapolate potential output for 1999Q1-2004Q4 using IMF staff estimates/forecasts for potential output growth, as quoted in Bayoumi (2000), as a guideline.

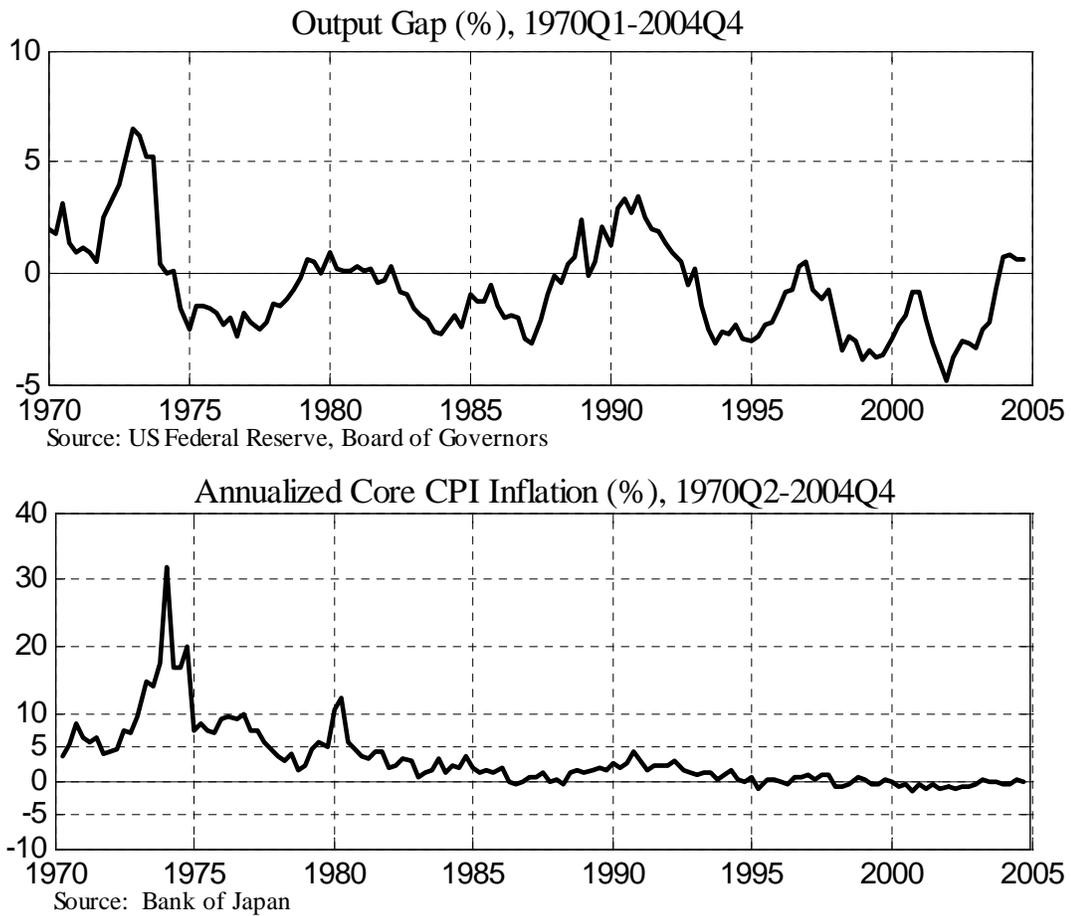


Figure 2

Note: 'Annualized Core CPI inflation' stands for annualized quarterly inflation in the Consumer Price Index excluding fresh foods, which the Bank of Japan adjusted for consumption tax reforms in April 1989 and April 1997.

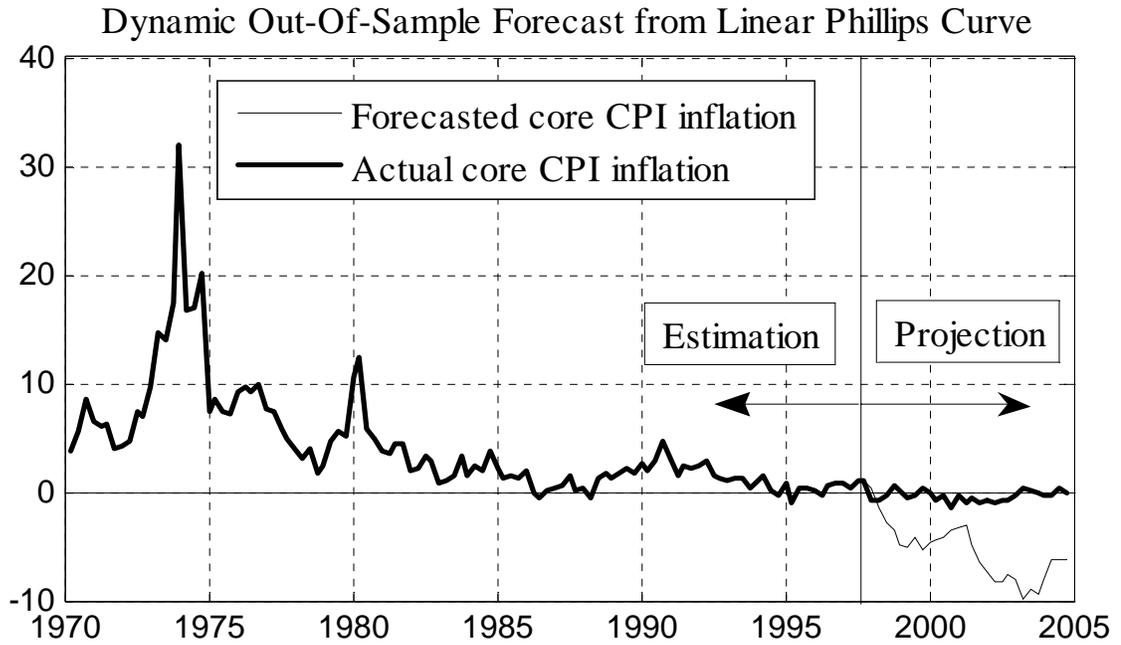


Figure 3

Note: The linear Phillips curve is estimated over 1971Q2-1997Q4; the forecast window is 1998Q1-2004Q4. The result suggests that, assuming a standard linear relationship between the output gap and inflation, the size of the negative output gaps in Japan would have warranted accelerating deflation in the period 1998-2002.

Linear Model: Time-invariant vs. Time-varying Output Gap Coefficients

$$\pi_t = \beta(L) \pi_t + \gamma_{1,t} (ygap_{t-1} + p ygap_{t-2}) + \delta impoil_t + e_t$$

Sample: 1971Q2-2004Q4	OLS (linear)	MLE (linear TV)
β_1	0.67*** (0.17)	0.67*** (0.08)
β_2	0.17 (0.16)	0.17** (0.08)
β_3	0.39*** (0.15)	0.37*** (0.08)
β_4	-0.23 (0.16)	-0.21*** (0.07)
$\gamma_{1,t}$	0.87** (0.36)	0.75*** (0.16) avg
$\gamma_{2,t} = p \cdot \gamma_{1,t}$	-0.70** (0.35)	-0.58*** (0.12) avg
δ	0.018* (0.009)	0.018*** (0.004)
σ_e	-	1.33*** (0.04)
p	-0.81*** (0.10)	-0.77*** (0.10)
Sum output gap coefficients	0.17** (0.07)	0.18*** (0.04) avg
Sacrifice ratio	1.47	1.39
Fit	$R^2 = 0.85$	LLF=-255.94
	$\overline{R}^2 = 0.84$	

Standard errors are in parentheses.

*** indicates significance at the 1% level; ** at 5% level; * at 10% level.

Table 1

Note: The rightmost column shows the results from estimating the state-space model consisting of equations (3) and (4) by means of the Kalman filter and Maximum Likelihood. The Phillips curve is linear, yet the output gap coefficients are allowed to vary over time as a random walk. 'avg' indicates the average of a time-varying coefficient and its standard error over the sample. For comparison, the middle column shows the results from estimating a standard linear Phillips curve with time-invariant output gap coefficients.

Random Walk Model: Time-Varying Output Gap Coefficients

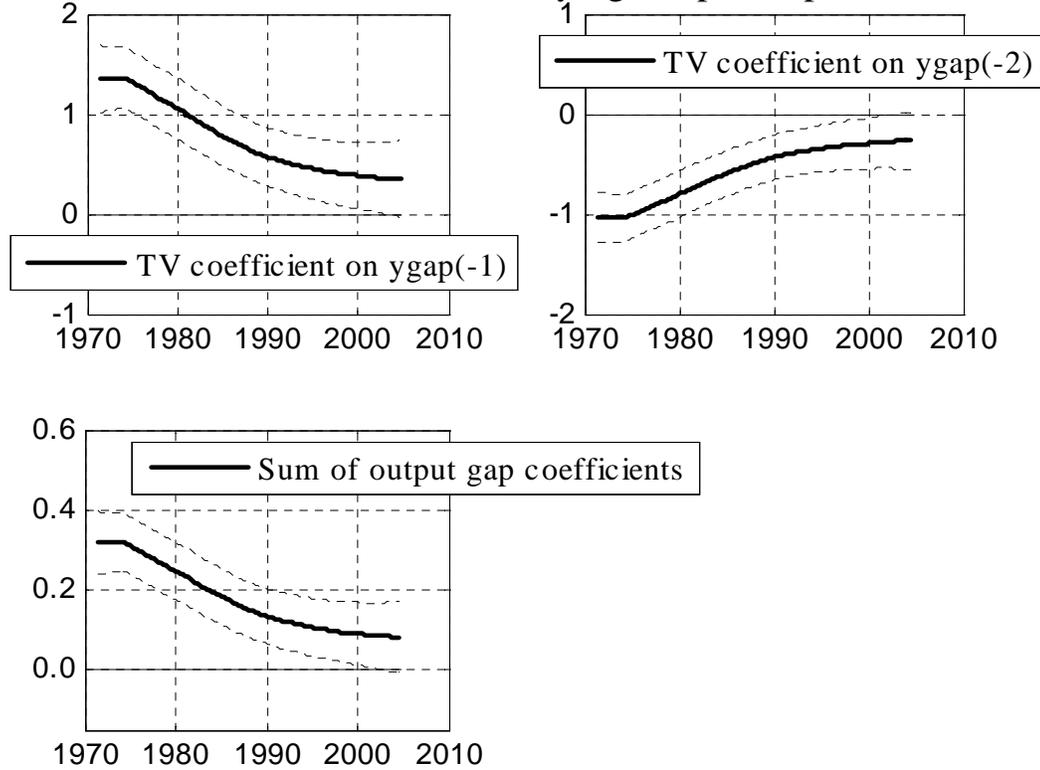


Figure 4

Note: This figure graphs the Kalman-smoothed time-varying output gap coefficients corresponding to the estimation results in Table 1, along with their 95% confidence interval. Note that the sum of the output gap coefficients is graphed on a different scale than the individual output gap coefficients in the top row.

Linear Phillips Curve vs. Laxton-Meredith-Rose Phillips Curve

$$\pi_t = \beta(L) \pi_t + \gamma_1 \left[\left(\frac{\phi ygap_{t-1}}{\phi - ygap_{t-1}} \right) + p \left(\frac{\phi ygap_{t-2}}{\phi - ygap_{t-2}} \right) \right] + \delta impoil_t + e_t$$

Sample: 1971Q2-2004Q4	OLS (linear)	NLS (LMR)
β_1	0.67*** (0.17)	0.70*** (0.14)
β_2	0.17 (0.16)	0.15 (0.14)
β_3	0.39*** (0.15)	0.37** (0.15)
β_4	-0.23 (0.16)	-0.22 (0.17)
γ_1	0.87** (0.36)	0.69** (0.30)
$\gamma_2 = p \gamma_1$	-0.70** (0.35)	-0.54* (0.29)
δ	0.018* (0.009)	0.019** (0.008)
ϕ	∞	10.00*** (1.83)
p	-0.81*** (0.10)	-0.77*** (0.12)
Fit	$R^2 = 0.85 / \bar{R}^2 = 0.84$	$R^2 = 0.87 / \bar{R}^2 = 0.86$
Q-stat [with p-value]: 4th lag	5.61 [0.23]	5.26 [0.26]
Q-stat [with p-value]: 12th lag	6.74 [0.16]	9.27 [0.68]

Huber-White standard errors are in parentheses.

*** indicates significance at the 1% level; ** at 5% level; * at 10% level.

Table 2

Note: The rightmost column shows the results from estimating a Laxton-Meredith-Rose Phillips curve. The nonlinearity parameter ϕ is precisely estimated. As Figure 5 demonstrates, its point estimate implies a fairly strong degree of convexity in the Phillips curve, with a vertical asymptote at an output gap of 10%. For comparison, the middle column in the table above graphs the results from estimating a linear Phillips curve.

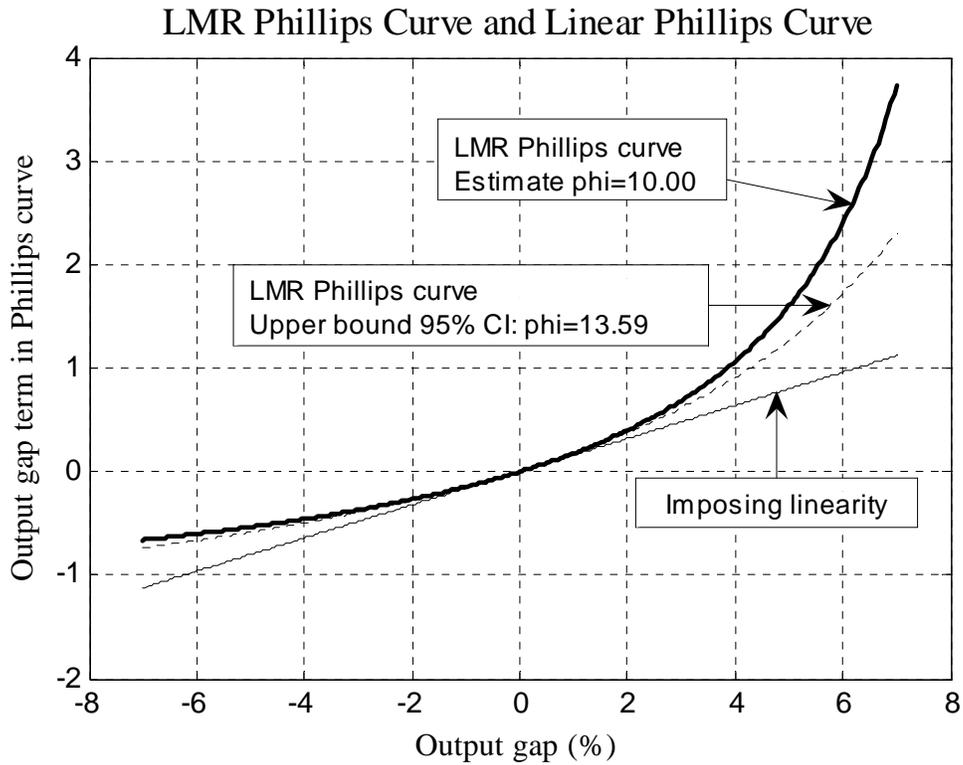


Figure 5

Note: This figure graphs $\gamma_1 [(\phi ygap_{t-1} / (\phi - ygap_{t-1})) + p (\phi ygap_{t-2} / (\phi - ygap_{t-2}))]$, as estimated in Table 2, with respect to the output gap. The dotted line graphs the same function imposing a value for ϕ which equals the upper bound of the 95% confidence interval around the point estimate for ϕ .

Standard Phillips Curve vs. Ball-Mankiw-Romer / Dotsey-King-Wolman Phillips Curve

$$\pi_t = \beta(L) \pi_t + [a + b |\bar{\pi}_t|] [ygap_{t-1} + p ygap_{t-2}] + \delta impoil_t + e_t$$

Sample: 1971Q2-2004Q4	OLS (linear)	OLS (BMR/DKW)
β_1	0.67*** (0.17)	0.77*** (0.10)
β_2	0.17 (0.16)	0.18 (0.12)
β_3	0.39*** (0.15)	0.22** (0.09)
β_4	-0.23 (0.16)	-0.17* (0.10)
γ_1	0.87** (0.36)	0.76 avg
$\gamma_2 = p \gamma_1$	-0.70** (0.35)	-0.63 avg
δ	0.018* (0.009)	0.019** (0.008)
a	γ_1	-0.53*** (0.19)
b	0.00	0.34*** (0.05)
p	-0.81*** (0.10)	-0.82*** (0.07)
Fit	$R^2 = 0.85 / \bar{R}^2 = 0.84$	$R^2 = 0.90 / \bar{R}^2 = 0.89$
Q-stat [with p-value]: 4th lag	5.61 [0.23]	0.61 [0.96]
Q-stat [with p-value]: 12th lag	6.74 [0.16]	4.29 [0.98]

Huber-White standard errors are in parentheses.

*** indicates significance at the 1% level; ** at 5% level; * at 10% level

Table 3

Note: The rightmost column contains the results from estimating a Phillips curve in which the slope depends on the absolute value of trend inflation. In line with Ball-Mankiw-Romer and Dotsey-King-Wolman, the coefficient b on trend inflation is positive and significant at the 1% level. For comparison, the middle column provides the results from a standard Phillips curve in which the coefficient on trend inflation is set to zero.

Standard Linear Phillips Curve vs. Lucas Phillips Curve

$$\pi_t = \beta(L) \pi_t + [c + d \text{var}_t(\pi)] [ygap_{t-1} + p ygap_{t-2}] + \delta \text{impoil}_t + e_t$$

Sample: 1971Q2-2004Q4	OLS (linear)	OLS (Lucas)
β_1	0.67*** (0.17)	0.83*** (0.12)
β_2	0.17 (0.16)	0.14 (0.11)
β_3	0.39*** (0.15)	0.22** (0.09)
β_4	-0.23 (0.16)	-0.19* (0.10)
γ_1	0.87** (0.36)	0.66 avg
$\gamma_2 = p \gamma_1$	-0.70** (0.35)	-0.52 avg
δ	0.018* (0.009)	0.020** (0.008)
c	γ_1	-0.05 (0.20)
d	0.00	0.04*** (0.01)
p	-0.81*** (0.10)	-0.79*** (0.10)
Fit	$R^2 = 0.85 / \bar{R}^2 = 0.84$	$R^2 = 0.90 / \bar{R}^2 = 0.89$
Q-stat [with p-value]: 4th lag	5.61 [0.23]	0.27 [0.99]
Q-stat [with p-value]: 12th lag	6.74 [0.16]	2.79 [1.00]

Huber-White standard errors are in parentheses.

*** indicates significance at the 1% level; ** at 5% level; * at 10% level.

Table 4

Note: The rightmost column contains results from a Phillips curve in which the slope depends on the variance of inflation. In line with the Lucas misperceptions theory, the coefficient d on the variance of inflation is positive and significant at the 1% level.

Bayesian Model Averaging: Pseudo-Posterior Model Odds

Ball-Mankiw-Romer / Dotsey-King-Wolman	81.42%
Lucas	18.58%
Laxton-Meredith-Rose	1.95E-06%
Linear	1.84E-09 %

Table 5

Note: This table displays the pseudo-posterior odds for each of the four listed models to be the true model, according to a Bayesian Model Averaging procedure as in Brock, Durlauf, West (2004). This procedure places equal prior probability on each of the models.