

Which non-linearity in the Phillips curve? *

The absence of accelerating deflation in Japan

Emmanuel De Veirman

Johns Hopkins University

E-mail: deveirman@jhu.edu

November 10, 2005

Abstract

It is standard in the literature to model the output-inflation trade-off as a linear relationship with a time-invariant slope. Our paper argues that the assumption of linearity is not as innocent as it seems. We assess empirical evidence for three types of non-linearity in the short-run Phillips curve. We gather such evidence from Japan, an economy for which a linear Phillips curve appears to have broken down in the late nineties. 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 a flattening of the Phillips curve. The broader theoretical 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.

Keywords: Deflation, non-linear Phillips curve,
Kalman filter, time-varying coefficients.

JEL Classification: C22, C32, E31, E32.

*I am grateful to my advisor, Laurence Ball, and to Alan Ahearne, Christopher Carroll, Carl Christ, Hali Edison, Jon Faust, Yasuo Hirose, Michael Kiley, Charles Kramer, Takeshi Kudo, Douglas Laxton, Andrew Levin, Louis Maccini, Athanasios Orphanides, Adrian Pagan, Erwan Quintin, John Roberts, Jirka Slacalek, Tsutomu Watanabe, Mark Wynne, Isamu Yamamoto, Naoyuki Yoshino, and seminar participants at the Federal Reserve Bank of Dallas, the Bank of Japan, the International Monetary Fund, the Canadian Economics Association, Georgetown University, and Johns Hopkins University for valuable comments. Contact: Johns Hopkins University, Department of Economics, 3400 N. Charles Street, Baltimore, MD 21218.

1 Introduction

1.1 Theoretical framework

The original Phillips curve was non-linear: Alban W. Phillips (1958) estimated a non-linear 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. Our paper argues that this simplifying assumption is not as innocent as it appears to be at first sight.

In the present paper, we gather evidence from Japan, an economy in which the Phillips curve appears to have flattened. To detect the determinants of such flattening, we empirically evaluate three theories, each of which posits a different set of factors affecting the slope of the short-run Phillips curve.

First, in 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.

Second, in Ball, Mankiw, Romer (1988) as well as in 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.

Third, 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 demand has a larger impact on firms' production, and a smaller effect on inflation. That is to say, the Phillips curve flattens.

Strictly speaking, only the first of the above theories implies that the short-run Phillips curve is non-linear. In the second and third theories, the Phillips curve is actually linear: its slope depends on exogenous variables other than the output gap or inflation. Nevertheless, we will mostly refer to the three theories as implying different types of non-linearity in the Phillips curve.

1.2 Japan as a test case

Our interest in Japan follows from the apparent breakdown of a linear Phillips curve in that country since the mid-nineties. We document that a linear Phillips curve fails to capture the coincidence of protracted negative output gaps in Japan during the period 1998-2002 with stable, be it moderately negative inflation. Instead, actual output was sufficiently below potential output in that period for a linear model to predict accelerating deflation.

Finding out why deflation didn't accelerate in Japan is our empirical tool to shed light on the nature of the output-inflation trade-off. We show that the absence of accelerating deflation in Japan cannot be adequately explained by popular explanations which assume a linear short-run Phillips curve with a time-invariant slope. On the contrary, we find that the three above-mentioned theories of non-linearity are consistent with the data. Moreover, each of the non-linear models outperforms an atheoretical benchmark model in which the Phillips curve is linear, but its slope is allowed to vary over time as a random walk.

Among the theories of non-linearity, we find most evidence for the hypothesis that a decrease in trend inflation created an environment in which prices became gradually more sticky, which in turn caused the Japanese Phillips curve to flatten.

1.3 Relation to existing literature

By assessing evidence for the three theories of non-linearity as well as for the flattening of the Japanese Phillips curve, our study aims to add a new perspective to both areas of research.

First, empirical evidence for the three theories of non-linearity was thus far mostly based on samples which did not contain a large number of deflation observations.¹ Since our study considers data for Japan beyond the mid-nineties, our sample includes a fairly large number

¹The studies below find evidence for a particular type of non-linearity based on multi-country studies, with samples which do not extend beyond the mid-nineties.

The empirical evidence for convexity in the Phillips curve in Laxton, Meredith, Rose (1995) is based on pooled data for the G-7 economies.

Ball, Mankiw, Romer (1988) regress individual Phillips curves for a large number of industrialized economies. Similar samples are used in subsequent robustness tests: DeFina (1991), Kiley (1996), and Yates and Chapple (1996).

Lucas (1973) also runs individual regressions of the output-inflation trade-off for a fairly large number of individual countries. A follow-up investigation, Alberro (1981), adopts a similar approach.

Dotsey, King, Wolman (1999) contains model simulations, but no empirical evidence. Much the same holds true for more recent literature on this model, such as Bakhshi, Burriel-Llombart, Khan, Rudolf (2003), Bakhshi, Khan, Rudolf (2004), and Ho and Yetman (2005).

of observations from the relatively unexplored region of the Phillips curve at which inflation is near-zero or negative. In that way, our sample includes a broad spectrum of values for inflation and the output gap, which is desirable for any study aiming to discover the true nature of the output-inflation relationship.

Second, while papers detecting a flattening in the Japanese Phillips curve do mention some of the theories of non-linearity which we described, they never explicitly estimate any of those theories. In this paper, we explicitly estimate three theories of non-linearity, assess whether they add information beyond an atheoretical time-varying coefficients model, and perform a series of tests to evaluate the relative performance of the non-linear models.

In fact, we are aware of only two papers which focus on the flattening of the Japanese Phillips curve. Using a dummy variable approach, both papers detect a kink in an otherwise linear Phillips curve. Nishizaki and Watanabe (2000) find that the Japanese Phillips curve has been significantly flatter below an annualized inflation rate of 3% than it was above that threshold. Mourougane and Ibaragi (2004) discover a kink in Japan's Phillips curve at a threshold level of 2% inflation. Furthermore, other results from both studies suggest that the Japanese Phillips curve was flatter in the nineties than in earlier decades.

Consistent with these two studies, our paper detects a significant flattening in the Japanese Phillips curve over time. Yet, since we estimate the slope of the Phillips curve as a time-varying coefficient in a Kalman filter framework, we obtain a smoothly declining slope, rather than a slope which is subject to sudden jumps when inflation transgresses its threshold level.

1.4 Structure of the paper

Section 2 documents that the output-inflation comovement in Japan during the episode 1998-2002 presents a puzzle to anybody who takes a linear Phillips curve literally.

In section 3, we evaluate popular explanations for the absence of accelerating deflation in Japan, other than those focusing on time-variation in the slope of the short-run Phillips curve. Neither of those explanations appears to be a major factor behind the absence of massive deflation. Most importantly, we find evidence against downward nominal wage rigidity, output gap mismeasurement, and the possibility that inflation expectations failed to turn negative. We also explain that the linear model implies accelerating deflation even in the presence of massively expansionary monetary policy.

Given our results in section 3, the remainder of the paper focuses on time-variation in the slope of the short-run Phillips curve. In section 4, we estimate a linear Phillips curve, but allow its slope to vary over time as a state variable in the Kalman filter. In particular, we model each of the coefficients on the output gap as a random walk. We detect a gradual, significant decline in the slope of the Phillips curve which has been occurring since before the nineties. This atheoretical model of time-variation in the Phillips curve slope is meant to be a benchmark for the three theoretically founded models of non-linearity.

Sections 5 through 7 investigate the determinants of the flattening in the Phillips curve which we documented in section 4. In section 5, we find that all three above-mentioned theories of non-linearity are consistent with the data.

In section 6, we estimate models that nest the random walk model of section 4 and one of the non-linear models of section 5, and repeat this exercise for each of the theories of non-linearity. We test the appropriate restrictions to discover how the theoretical models of non-linearity perform relative to the atheoretical benchmark. We find that each of the theories of non-linearity outperforms the random walk model.

Given that each of the non-linear models is consistent with the data and adds information beyond that contained in the atheoretical benchmark, section 7 evaluates the relative performance of the three types of non-linearity. Our results favor the hypothesis that declining trend inflation caused firms to set their output prices less frequently, which would explain the observed flattening in the Phillips curve modeled as a linear, but time-varying relationship. 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.

Section 8 concludes and presents policy implications. For instance, our results suggest that, in any country in which trend inflation is substantially lower today than it was in previous decades, the Phillips curve is currently flatter than a linear model would indicate. Equivalently, the output costs of disinflation are larger than the standard model would suggest. Hence, the optimal inflation target exceeds the level of inflation which would appear optimal from a linear model.

2 Background

This section provides background on the output–inflation comovement in Japan. In the first subsection, we present a concise economic history of Japan since the early nineties. The main point of that description is that the nineties can be characterized as a succession of recessions. In the final subsection, we document that, within this period of negative output gaps, the episode 1998-2002 presents the most striking puzzle in a linear framework: it is the episode with the largest negative output gaps, yet it is among the episodes with the most stable inflation rates. Before laying out the puzzle since 1998, we document that on average over the period until right before that time, the comovement of the output gap and inflation was such that it could be reasonably well approximated by a linear function.

2.1 Concise recent history

Figure 2.1 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 31, 1989 has been lower than it was in any of the previous two decades. The period after the stock market crash reads as a succession of recessions, interrupted only by brief or limited recoveries. The Heisei recession, which took root around 1991, was followed by a modest recovery between the end of 1993 and the beginning of 1997. The Asian crisis, along with the April 1997 increase in consumption taxes, contributed to a second recession in 1997-98. A new recovery halted in 2001, plausibly at least in part due to the dampening effect on exports of a slowdown in the United States and East Asia, and to the Bank of Japan’s brief interruption of the zero interest rate policy between August 2000 and March 2001. Since 2003, there are signs of a new upturn.

2.2 Earlier on, output and inflation appeared to comove linearly

Figure 2.2 enables us to gain our first insights about the comovement of the output gap and inflation in Japan. The top panel graphs the output gap series implied by the actual and potential output data from the previous figure. The lower panel graphs annualized quarterly inflation in the Consumer Price Index excluding fresh foods, which the Bank of Japan adjusted for consumption

tax reforms.² 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 in 1974Q1 and 1980Q2. For now, a casual inspection of figure 2.2 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 inflation tended to increase in booms. As we will discuss below, the relationship became gradually less clear in the nineties, in particular during the period 1998-2002.

To characterize the output-inflation comovement through 1997 somewhat more formally, we regress the following linear Phillips curve using quarterly time series data for 1971Q2-1997Q4:

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

Observations for 1970Q2-1970Q1 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 import prices of petroleum, coal, and natural gas.³ Since oil import prices are on Yen basis, this variable effectively accounts for exchange rate movements, to the extent that there is no pass-through from exchange rate movements into domestic oil prices. The error term e_t is assumed to be i.i.d..

In equation (1), inflation expectations are proxied by lags of inflation. In section 3.1, we document that the data for Japan indeed suggest that inflation expectations tracked lagged inflation closely. In that section, we also briefly mention the results which we obtained when estimating forward-looking and hybrid Phillips curves, and Phillips curves augmented with lagged (rational) expectations of current inflation.

The lag structure in equation (1) removes all serial correlation from the error term, but is sufficiently parsimonious for our estimations involving time-varying output gap coefficients and/or non-linearities 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.⁴

²A consumption tax of 3% was introduced in April 1989. That sales tax was increased to 5% in April 1997.

³In 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.

⁴Augmented 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

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, which reflects the total effect of any output gap on inflation after two quarters, 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 tended to coincide with disinflations.

2.3 Seemingly puzzling output-inflation comovement since the nineties

Figure 2.2 suggests that the relationship between the output gap and inflation became gradually less pronounced. In fact, this phenomenon even seems to have been present during the bubble period in the late eighties: some observers note that inflation remained fairly moderate during that period, notwithstanding large, positive output gaps.⁵ For most of the nineties, negative output gaps coincided with moderate declines in inflation. 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 more formally, figure 2.3 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 2000Q4.⁶ This suggests that deflation over that period 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.⁷

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.

⁵See, for instance, Kumar e.a. (2003), p.22.

⁶We equally obtain a massive deflation forecast from a similar equation with the GDP deflator. Since the substitution bias is present in the CPI at all times, while it is not in the GDP deflator, using the latter merely shifts the puzzle down. With the GDP deflator, both actual and predicted inflation are more negative than in the CPI case discussed in the body of the text.

⁷Unsurprisingly, the forecast errors of the dynamic out-of-sample forecast widen over the forecast window, at least as far as the forecast errors account for uncertainty about future residuals. When we instead use a dummy

Our forecast of accelerating deflation is related to the predictions by many economists, including at the Bank of Japan, that Japan was about to enter a deflationary spiral in 1998. We find some clues to the Bank of Japan's thinking in 1998 in speeches by BOJ officials, including then-Governor Hayami, which allude to the possibility of a negative interaction between financial sector stress and slow economic growth.⁸ Only two years later, official Bank of Japan publications explicitly state that 'Japan's economy was on the brink of a deflationary spiral' in 1998, in connection with the possibility of negative interaction between financial pressures and slow output growth.⁹

However, our paper does not attempt to provide a complete answer to the question why a deflationary spiral did not occur. Out of all possible interlinkages between weak output growth, deflation, and financial sector vulnerabilities which plausibly motivated the Bank of Japan's deflationary spiral forecast,¹⁰ we focus on one step which did not occur: sizeable and protracted negative output gaps did not lead to massive deflation.

3 Capturing the output-inflation comovement in a linear setting?

What appears to be break in a linear Phillips curve in section 2, is in fact consistent with three possible situations regarding the nature of the output-inflation trade-off. First, Japan's short-run Phillips curve may have been linear both before and after 1998. This could be true if the actual comovement of output and inflation after 1998 does turn out to reflect a linear relationship once, say, any output gap mismeasurement is eliminated, or the possible existence of downward nominal wage rigidity is accounted for in the Phillips curve. Second, the relationship between the output gap and inflation may have changed drastically after 1998, perhaps in a fashion which could

variable approach, similar to that in Nishizaki and Watanabe (2000) and Mourougane and Ibaragi (2004), we find that the sum of the output gap coefficients is significantly flatter, at the 5% level, in the period 1998Q1-2004Q4 than it was in the period 1971Q2-1997Q4. All of our results are available upon demand.

⁸A speech by Governor Hayami on July 29, 1998 refers to the 'very close relation between economic conditions and the financial system'. A speech by Policy Board Member Taketomi on July 3, 1998 states that 'there is a chance that our economy can avoid a deflationary spiral.' Governor Hayami's speech is available on <http://www.boj.or.jp/en/press/98/ko9809a.htm>. Board Member Taketomi's speech is available on <http://www.boj.or.jp/en/press/98/ko9807b.htm>.

⁹Bank of Japan, Research and Statistics Department (2000). See also Hayakawa and Ugai (2000).

¹⁰A likely scenario, which materialized in part, was that the financial crises of 1997 and 1998 would bring about a credit crunch, depressing economic growth, which in turn would increase the share of non-performing loans even further and thus lead to a further deterioration of banks' balance sheets. Under that doom scenario, low economic growth would feed into massive deflation, which could in turn have detrimental effects on output growth and financial sector health.

not be explained by existing theories. Third, the apparent breakdown of a linear Phillips curve may suggest that the linear model was misspecified all along, while the misspecification became apparent only once the economy entered a region of the Phillips curve where the non-linearity is particularly pronounced. In other words, output and inflation may have moved together according to the same laws all along, but their relationship may have been of a different nature than the standard, linear model presupposes.

In its entirety, our paper finds evidence for the third hypothesis. In the present section, we disprove the first hypothesis. In particular, we document that the absence of accelerating deflation in Japan cannot be adequately addressed by popular explanations which assume a linear, time-invariant short-run relation between the output gap and inflation.¹¹ Based on our conclusions in the present section, we focus on the path and determinants of the Phillips curve slope in sections 4 through 7, which constitute the body of this paper.

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 a priori plausible that inflation expectations did in fact not turn negative in the period 1998-

¹¹In the present footnote, we assess evidence for two conjectures which are not commonly thought of, but which are a priori plausible explanations for the absence of accelerating deflation. First, did public price regulation in particular sectors dampen the fall in the aggregate price level? Second, is the virtual stability of service prices a reason behind the absence of accelerating deflation?

First, sectors in which prices were publicly regulated over the period 1998-2002, such as education, public transport, and health services, experienced positive, fairly high inflation over that same period. However, when we add relative inflation in the prices of regulated goods as a regressor in equation (1), this variable does not enter significantly. Furthermore, the out-of-sample forecast is the same as in the baseline scenario. This suggests that price regulation is not a major cause for the absence of accelerating deflation in the aggregate CPI.

[To carry out the above test, we constructed two indices capturing price patterns in regulated goods. The first included prices for public services (such as education and health services), and electricity, gas and water. The second included the same public service prices, in addition to the price of rice, a good for which there has traditionally been some price control in Japan. The difference between inflation in those indices and core CPI inflation yielded two measures for relative inflation in price-regulated goods, both of which we tested for in equation (1).]

Second, there has been virtually no decline in service prices (the trough was at -0.1% in 2001), while there has been substantial deflation in goods prices (which was lowest at -1.8% in 2002). This difference may be due the larger impact of any downward nominal wage rigidity on output prices in the service sector, where labor costs typically account for a larger share of total input costs. However, including relative service price inflation as a regressor does not change the Phillips curve regression, nor does it change the deflation forecast. This is consistent with our findings, in subsection 3.4, that downward nominal wage rigidity was not present during the period 1998-2002. In conclusion, relatively high service price inflation cannot be the explanation for the absence of accelerating deflation in Japan.

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.¹² Under that hypothesis, negative output gaps would imply negative, but stable inflation. This is consistent with the Japanese situation in 1998-2002.

However, neither data inspection nor basic econometric analysis provides much evidence for the claim that inflation expectations failed to turn negative in Japan. For instance, 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. This finding is confirmed by qualitative price expectations data in the Tankan business survey, and by inflation forecasts from the OECD, the US Federal Reserve, as well as by four-quarter ahead Consensus forecasts for inflation. All of these data sources suggest that inflation expectations tracked lagged inflation quite closely over the sample.

Furthermore, we do not find strong econometric evidence for a break in the process through which inflation expectations are formed. In particular, when we regress quarterly Consensus forecasts, or annual OECD forecasts, on a constant and lagged inflation, we do not find that the coefficient on lagged inflation is significantly larger after the hypothesized break date than it was before. This result holds for all potential break dates starting in 1995Q1.¹³

Since the data suggest that expectations tracked lagged inflation closely, it appears reasonable to proxy inflation expectations by lagged inflation in the Phillips curve. Nevertheless, we examine robustness across different ways to account for inflation expectations. We estimated New Keynesian Phillips Curves (NKPC) with current expectations of future inflation, as in Roberts (1995). We also estimated the Fuhrer and Moore (1995) hybrid Phillips curves with a forward- and backward-looking component. Finally, we regressed a more traditional expectations-augmented Phillips curve with lagged expectations of current inflation. Every time, we experimented with two types of expectations measures on the right-hand side: OECD forecasts, and a Two-Stage Least Squares approach, where expectations were proxied by the fitted value from a regression on a time trend, the right-hand side variables in the Phillips curve, and further lags of inflation and the output gap. In most cases, we encountered weak instruments issues, or estimated insignifi-

¹²To 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 $E_{t-1}(\pi_t) = 0$. From the above equation, it can be seen that in that case, negative output gaps tend to coincide with negative, but stable inflation.

¹³The results are available from the author.

cant or even negative Phillips curve slopes. However, the sum of the output gap coefficients was positive and significant at the 5% level or better for the purely forward-looking NKPC and for the expectations-augmented Phillips curve, as far as we used OECD forecast data to proxy for inflation expectations.

3.2 Was the output gap less negative than Fed and BOJ data suggest?

In section 2, we discussed that in Japan, most of the period since the 1989 stock market crash is characterized by a succession of recessions and low growth episodes. Figures 2.1 and 2.2 suggest that potential output was growing at a sufficiently fast pace for a large negative output gap to persist over the period 1998-2002. However, if true potential output growth were slower than suggested by these estimates, the output gap would not be as negative. Obviously, if the true output gap were close to zero, not even a linear model would have yielded a forecast of accelerating deflation.

There are several reasons for why potential output growth may have slowed down in Japan since the early nineties. For instance, potential growth may have slowed because much of the physical capital stock built up during the bubble period in the late eighties has become obsolete, and thus has gradually become less useful for production. Furthermore, some observers note that potential output growth plausibly slowed due to a decline in labor productivity in Japan's aging society.

Before we address this hypothesis, note that we use US Federal Reserve estimates of the Japanese output gap as a baseline. Since the Fed's recent output gap estimates are confidential, we compute an output gap series for 1999Q1-2004Q4 based on data for actual output, and our own extrapolation of potential output. As a guideline in the extrapolation, we use the figure which Bayoumi (2000) quotes as the IMF staff estimates/forecasts of potential growth.¹⁴

Of course, there is no perfect certainty about the true level of potential output, in particular at the end of the sample. However, both the growth rate of our measure of potential, and a comparison with other available output gap estimates, suggest that we do not overestimate

¹⁴In particular, we fit ARMA-processes on potential output growth through 1998Q4, and perform a dynamic out-of-sample forecast of potential output growth for 1999Q1-2004Q4. We use the IMF staff estimates/forecasts of potential output growth, as found in Bayoumi (2000), to choose among possible ARMA-processes. According to the IMF estimates, potential output grew on average 1.09% per year in the period 1999-2002. From an AR(3) process, we predict potential output growth to average 0.96% annually during that period. (Note that this is for a period ending in 2002, unlike the figure discussed in the body of the text.)

potential output growth, and hence do not overestimate the size of the output gap in absolute value.

First, notice, in figure 2.1, that estimated potential output growth has been gradually slowing down since the 1989 stock market crash. Our estimates of potential output grow on average 0.92% per year in the period 1999Q1-2004Q4, while average potential growth in 2004 is only 0.82%.

Second, our output gap series displays a very similar pattern as the Bank of Japan's production function based output gaps. Kamada (2005) describes six types of output gaps computed at the Bank of Japan. Three of these output gaps, including the BOJ's most widely used 'prototype' output gap, are based on a measure of potential which is estimated as the maximum level of production attainable given total factor productivity, the labor force, and the available capital stock. By construction, these output gaps are negative at all times, unlike the Fed output gap, which tends to be positive in booms and negative in recessions. Figure 3.1 documents that, when we rescale the BOJ's prototype output gap such that its average equals the average of the Fed output gap over the period of BOJ data availability, the BOJ estimates are slightly more negative than the Fed estimates at the end of the sample, relative to earlier years.¹⁵

Two other procedures to estimate potential output, which are not based on an aggregate production function, are not adequate for our purposes.

First, univariate smoothing methods such as the Hodrick-Prescott filter yield an output trend which is biased to the actual data at the starting point and the endpoint of the sample. Hence, such estimates of potential output are automatically close to actual output whenever the latter stagnates for a fairly long time. We reject unadjusted univariate smoothing methods as an unnecessarily ad hoc way to compute potential output. Figure 3.2 documents that the HP-filter, unsurprisingly, does not yield a large negative output gap at the end of the sample. Our graph is in line with HP-filter estimates of the output gap, starting in 1985, as the fourth output gap in Kamada (2005).

Second, it would not be accurate to estimate potential output purely as the level of output at which inflation is stable by means of an inverted-Phillips curve approach. For instance, Kuttner (1994) develops a methodology for estimating the non-accelerating inflation level of output

¹⁵The BOJ's two other production-function based output gaps display a very similar pattern. Each of the BOJ's three production-based output gap estimates, as reviewed in Kamada (2005), implies that the average output gap in the period 1998-2002 is substantially larger than the output gaps at any other time since 1985.

(NAILO) as a latent variable using the Kalman filter. As another example, Hirose and Kamada (2003) develop a methodology for simultaneous estimation of the NAILO and the Phillips curve coefficients. In their approach, the NAILO is the level of output that minimizes a value function consisting of the sum of squared residuals of the Phillips curve, and a smoothness term. Applying this methodology to our data, we find that the Hirose-Kamada output gap moves around zero at the end of the sample, thus diverging from the Fed output gap, as shown in figure 3.3. This result is in line with the time-varying NAIRU and NAILO output gap estimates since 1985 in Kamada (2005), which are the fifth and sixth output gaps in his paper. The outcome is not surprising. In fact, at times when inflation is fairly stable, output is by definition near its stable-inflation level. A linear Phillips curve with a NAILO-based output gap as a regressor will always fit inflation accurately, merely because such output gap is endogenous to the behavior of inflation.

3.3 Did expansionary monetary policy prevent massive deflation?

In a textbook world, fast money growth exerts upward pressure on inflation. In Japan, the monetary base and M1 have grown at massive rates, especially after March 2001, when the Bank of Japan began to target the reserves ('current account balances') of private banks at the Bank of Japan. Is this among the factors which prevented deflation from accelerating?

On the one hand, fast growth in narrow monetary aggregates has not translated into fast growth rates of broader aggregates such as M2, a fact which is plausibly related to a decline in bank lending since 1997. 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. Yet, any such effects would already be reflected in the output gap and inflation expectations data which we discussed in the previous two subsections. Inflation expectations did turn moderately negative, notwithstanding expansionary monetary policy. Similarly, the output gap did grow sufficiently negative to warrant accelerating deflation in a linear framework.

As such, the path of monetary policy does not change our conclusions from subsections 3.1 and 3.2.¹⁶

¹⁶To be sure, note that authors suggesting ways for Japan to escape the liquidity trap, such as Krugman (1998), Bernanke (1999), McCallum (2001), and Svensson (2001), do not allow for a direct effect of monetary aggregates on inflation. These authors invariably consider options for monetary policy to close the output gap and/or to raise inflation expectations.

3.4 Downward nominal wage rigidity?

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 particular, they detect skewness in the wage distributions of different groups of workers, classified by gender and part-time/full-time status. However, in a more recent study, 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, 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. At the same time, there is a large wage gap between regular and non-standard employees. In 2004, the hourly base wages of part-time workers were only 40.5% those of regular workers.¹⁷ 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 cannot be the explanation for the absence of accelerating deflation in Japan.

¹⁷The overall monthly cost (including bonuses, fringe benefits, social security contributions, and training expenses) of employing a part-time worker is 36.9% that of employing a full-time worker.

4 Evidence for a flattening Phillips curve

In section 3, we refuted hypotheses for the absence of accelerating deflation which do not necessarily imply time-variation in the slope of the short-run Phillips curve. From this point on, we focus on the path and determinants of the slope of the short-run output-inflation tradeoff. In the present section, we model each of the output gap coefficients as a random walk in a Kalman filter framework. Our results suggest that the Phillips curve has flattened significantly since the early seventies. The atheoretical framework of the present section will serve as a benchmark for three theoretically founded models of non-linearity in the short-run Phillips curve, which we will assess in sections 5 through 7. That model assessment will enable us to draw conclusions about the likely reasons for the flattening of the Phillips curve.

4.1 The model in state-space form

Similar to equation (1), we specify the Phillips curve as a linear relationship, at a quarterly frequency. However, this time we allow for time-variation in the output gap coefficients:

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

As motivated for equation (1), we impose the sum of the four inflation lag coefficients to be equal to one, and omit the constant. In addition, 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. Our new model nests equation (1), in which the output gap coefficients were time-invariant, and thus trivially related to each other by a constant factor of proportionality. After writing the model in state-space form, we will motivate the assumption of proportionality, and discuss econometric evidence that imposing it does not significantly worsen the fit.

Now, we write the model in state-space form.¹⁸ After imposing the assumption of proportionality, the measurement equation in effect depends on only one state-variable, $\gamma_{1,t}$:

¹⁸The format corresponds to that in Harvey (1994).

$$\pi_t = [ygap_{t-1} + p.ygap_{t-2}] \cdot \gamma_{1,t} + \begin{bmatrix} \beta_1 & \beta_2 & \beta_3 & (1 - \beta_1 - \beta_2 - \beta_3) & \delta \end{bmatrix} \cdot \begin{bmatrix} \pi_{t-1} \\ \pi_{t-2} \\ \pi_{t-3} \\ \pi_{t-4} \\ impoil_t \end{bmatrix} + e_t \quad (3)$$

Where the error term e_t is normally distributed with mean zero and variance σ_e^2 .

The transition equation:

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

Where $v_{1,t}$ is normally distributed with mean zero and variance $\sigma_{v_1}^2 = \sigma_e^2 \cdot q_1$. The parameter q_1 is the signal-to-noise ratio for the coefficient on the output gap's first lag. Implicitly, we model the second output gap coefficient as $\gamma_{2,t} = \gamma_{2,t-1} + v_{2,t}$, where $v_{2,t}$ is normally distributed with mean zero and variance $\sigma_{v_2}^2 = \sigma_e^2 \cdot q_2$. Modelling the output gap coefficients as a random walk is a standard approach.¹⁹

4.2 Motivating the assumption of proportional output gap coefficients

Next, we motivate the assumption that the output gap coefficients are related by a time-invariant factor of proportionality. We are primarily interested in the sum of the output gap coefficients, which indicates the full effect of a given output gap observation on inflation after two quarters. The precise way in which this sum is allocated over the two individual output gap coefficients is less important for our purposes. The estimation results are similar whether we impose proportionality or not, with that difference 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, the assumption of proportional output gap coefficients implies two restrictions on

¹⁹See, for instance, Faust, Rogers, Wang, Wright (2003). To make entirely sure that writing the coefficients as non-stationary processes does not bias our results towards finding a decline in the sum of the output gap coefficients, we simulated our model with different types of artificially generated coefficients. When the 'true' coefficients were constant, or stationary around a mean, so were the Kalman smoothed estimates.

the hyperparameters, which can be tested by a likelihood ratio test. As for the variances of the shocks to the output gap coefficients, the assumption of proportionality implies that:²⁰

$$\sigma_{v_2}^2 = p^2 \cdot \sigma_{v_1}^2 \quad (5)$$

As for the covariance between the shocks to the two output gap coefficients, proportionality implies that:

$$\text{cov}(v_1, v_2) = p \cdot \sigma_{v_1}^2 \quad (6)$$

When estimating the model with two distinct time-varying output gap coefficients, i.e. without assuming proportionality, we estimated $\sigma_{v_1}^2$, $\sigma_{v_2}^2$, and $\text{cov}(v_1, v_2)$ without imposing these two restrictions. The likelihood ratio statistic, computed as twice the difference between the log likelihood functions in the unrestricted and restricted case, is distributed $\chi^2(2)$. We find that imposing restrictions (5) and (6) does not significantly worsen the fit.²¹

In conclusion, we impose proportionality in this model, as well as in any other models which we estimate in the remainder of this paper.

4.3 Estimation procedure and results

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 hyperparameters²² by means of the Matlab-function `fminunc`. 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 4.1 compares the results of the time-varying coefficients linear Phillips curve with those of a linear Phillips curve estimated by OLS. Both of these estimations, as well as all other estimations

²⁰The assumption of proportional output gap coefficients implies $v_{2,t} = p \cdot v_{1,t}$, which in turn implies the restrictions (5) and (6) on the hyperparameters of the model.

²¹Note that the assumption of proportionality implies restrictions (5) and (6), yet those two restrictions do not necessarily imply proportionality of the output gap coefficients. Restrictions (5) and (6) do, however, imply proportionality of the *change* in the output gap coefficients. To see this, note that the combination of (5) and (6) implies that $\rho_{v_1, v_2} = 1$, i.e. the shocks to the output gap coefficients are perfectly correlated, and thus proportional to one another. Given the random walk assumption, this is equivalent to saying that the changes in the output gap coefficients are perfectly correlated, and thus proportional. In practice, the estimation results are similar, whether we impose proportionality of the output gap coefficients, or only restrictions (5) and (6).

²²The only exception is the signal-to-noise ratio, q_1 , which we set to 1/1600 in the baseline. In any case, the parameter estimates reported in table 4.1 are robust for all values of q_1 up to 1/25.

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. Notice that the standard errors are smaller in the MLE case, which is indeed the most efficient estimator on theoretical grounds, since it takes the entire distribution of the data into account. 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, while the average standard error is again smaller in the MLE case. 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 corresponds to earlier estimates of the Japanese sacrifice ratio in Ball (1994) and Zhang (2001).²³

Figure 4.2 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 over the sample. The graph suggests that the flattening is significant at the 5% level. Notice that the flattening has occurred since before the nineties. Note that the absence of accelerating deflation in 1998-2002 is only one among the episodes consistent this result. For example, the finding that the Phillips curve has flattened since before the nineties is also in line with the absence of massive inflation during the bubble period in the second half of the eighties.

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

In section 4, we documented a gradual, significant flattening in a linear Phillips curve, where each of the output gap coefficients was modeled as a random walk. The flattening of a linear Phillips curve may suggest that the output-inflation trade-off should actually be modeled as a non-linear relationship. In the present section, we estimate Phillips curves implied by three different theories, each of which posits a different set of factors determining the slope of the Phillips curve. These theories can be thought of as implying different types of non-linearity in the Phillips curve. We find that each of the three models is consistent with the data. In section 6, we will show that each of the three non-linear models adds information beyond that contained in the random walk

²³Ball (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 (2001) computes a sacrifice ratio of 1.85% when accounting for long-lived effects. Both estimates are based on six disinflation episodes.

model. In section 7, we assess the relative performance of the three theories.

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

Laxton, Meredith, Rose (1995) and related papers²⁴ allow for convexity in the short-run Phillips curve. In Laxton, Meredith, Rose (LMR), capacity constraints constitute the economic rationale for non-linearity. 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 capacity constraint, where the slope becomes gradually steeper as the economy moves in the direction of 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 the convexity 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 removed 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 at that time. 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,²⁵ but such reasonings are not contained in LMR's original paper.

In the present section as well as in section 6, 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

²⁴Selected papers in this strand of the literature: Turner (1995), Debelle and Laxton (1996), Debelle and Vickery (1997), and Schaling (1999).

²⁵We propose two tentative reasonings. First, the fraction of firms which are capacity constrained may decrease as the overall economy moves further away from the capacity constraint, yet it is possible that even at negative output gaps, a small fraction of firms operates near full capacity. Second, persistent negative output gaps plausibly increase long-run unemployment, where the long-run unemployed do not exert as much downward pressure on wage and hence price inflation as would appear from a traditional model.

assessment in section 7.2.

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

$$\pi_t = \beta(L).\pi_t + \gamma_1 \cdot \left[\left(\frac{\phi.ygap_{t-1}}{\phi - ygap_{t-1}} \right) + p \cdot \left(\frac{\phi.ygap_{t-2}}{\phi - ygap_{t-2}} \right) \right] + \delta.impoil_t + e_t \quad (7)$$

Apart from the non-linear functional form as reflected in the output gap terms, we impose exactly the same restrictions as in equation (2). Yet, in this case, we assume that γ_1 is time-invariant.²⁶ That is, we do not allow for time-variation in the Phillips curve slope beyond that implied by the non-linearity of the functional form. (We will relax this assumption in section 6).

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 non-linearity 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. If, instead, ϕ is virtually infinite, the capacity constraint is that far away from the origin that the Phillips curve reduces to a linear Phillips curve. Hence, the linear Phillips curve in equation (1) is a special case of equation (7).

The rightmost column of table 5.1 presents results from estimating (7) over 1971Q2-2004Q4. As a basis for comparison, we include results from a purely linear Phillips curve estimated over the same sample period. The linear Phillips curve essentially imposes the restriction that $\phi = \infty$. In the potentially non-linear case, ϕ is precisely estimated, with a point estimate of exactly 10.00. This suggests that the economy would reach the capacity constraint if ever 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 (7), we graph the sum of the output gap terms in figure 5.2, as a function of the output gap. In particular, the bold curve in figure 5.2 graphs $\gamma_1 \cdot \left[\left(\frac{\phi.ygap_{t-1}}{\phi - ygap_{t-1}} \right) + p \left(\frac{\phi.ygap_{t-2}}{\phi - ygap_{t-2}} \right) \right]$ with respect to the output gap, where we impose that $ygap_{t-1} = ygap_{t-2}$. Thus, the function graphs the total effect of a given output gap level on inflation after two quarters, abstracting from past inflation and supply shocks. For comparison, the 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

²⁶In analogy with equation (2), we substituted $\gamma_2 = p.\gamma_1$. Yet, in the case of equation (7), this does not impose any actual restriction. Equation (7) is equivalent to an equation with γ_1 and γ_2 appearing independently.

degree of non-linearity in the Phillips curve.²⁷ 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 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.

In conclusion, we do find that the Laxton, Meredith, Rose hypothesis is consistent with the Japanese data.

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

In this subsection, we explicitly model the slope of the Phillips curve as a function of trend inflation. By doing so, we wish to assess the empirical validity of two theoretical models 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.²⁸

The two models measure the frequency of price adjustment in a slightly different way. On the one hand, Ball, Mankiw, Romer (BMR) focus on the length of time between price changes, as in Taylor (1980). On the other hand, Dotsey, King, Wolman (DKW) focus on the probabilities of price adjustment, as in Calvo (1983). BMR and DKW share the common feature that the frequency of price adjustment is endogenous, unlike pricing in the Taylor- and Calvo-models.

In 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. Under high trend inflation, any firm expects its relative price to change rapidly over time, which in turn leads the firm to expect

²⁷The 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 non-linearity parameter. Note that even in that case, the function still looks distinctly non-linear.

²⁸DeFina (1991), Kiley (1996), and Yates and Chapple (1996) confirm the robustness of the Ball, Mankiw, Romer result. Recent literature on the Dotsey, King, Wolman model includes Bakhshi, Burriel-Llombart, Khan, Rudolf (2003), Bakhshi, Khan, Rudolf (2004), and Ho and Yetman (2005).

a rapid change in its profit-maximizing nominal price. Thus, the forward-looking firm will not fix its actual price for a long time, else it incurs large deviations from the profit-maximizing price. 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.²⁹

In DKW, the firms' steady-state probability of price adjustment depends positively on steady-state inflation. Higher steady-state inflation implies that firms which set their prices several periods ago face more rapid erosion of their relative prices. Thus, higher steady-state inflation increases the benefit to firms of updating their prices. This implies that for a larger fraction of firms, the benefit of adjustment will exceed the (labor) cost of price adjustment. In short, higher steady-state inflation leads to higher steady-state probabilities of price adjustment.³⁰

These two models' prediction for Japan is that, as trend inflation 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 available data set for Japan which documents the evolution of the frequency of price adjustment over time. However, we can test the hypothesis whether the slope of the Phillips curves depends on trend inflation.

To this end, we estimate a Phillips curve in which the slope is explicitly modeled as a function of the absolute value of trend inflation:

$$\pi_t = \beta(L).\pi_t + [a + b. |\bar{\pi}_t|] . [ygap_{t-1} + p.ygap_{t-2}] + \delta.impoil_t + e_t \quad (8)$$

Sample period and coefficient restrictions are the same as in equation (2). While each of the output gap coefficients in equation (2) is modeled as a random walk, equation (8) makes the time-path of the output gap coefficients endogenous to the absolute value³¹ of trend inflation, $|\bar{\pi}_t|$. Notice that equation (8) is strictly speaking linear, although its slope depends on trend inflation. We generate trend inflation at time t as a geometric average of J quarters of past inflation:

²⁹In BMR, the frequency of price adjustment is also endogenous to aggregate demand shocks and firm-specific demand and cost shocks. An increase in the variance of aggregate or firm-specific shocks increases the firms' uncertainty about future profit-maximizing prices. In response, the firm set its price for shorter periods of time.

³⁰A firm's steady-state adjustment probability is increasing in the time that has elapsed since the firm's last price adjustment.

³¹Intuitively, 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 BMR nor DKW 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.

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

In the baseline, $\theta = 0.93$ and $J = 71$. Note that trend inflation does not depend on current inflation, so as to avoid endogeneity problems in equation (8). 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 (9), observe that inflation lags which are further in the past receive relatively smaller weights. This can be motivated from either of the two models. In the BMR model, firms choose the length of time over which a new price will stay in effect, based on their expectations of aggregate inflation. Assuming that firms form inflation expectations by looking at past inflation, and in doing so assign relatively more weight to recent inflation observations, it follows that recent inflation lags have a relatively larger impact on the frequency of price adjustment. In a DKW-setting, inflation in π_{t-j} will only affect the pricing behavior of firms which last updated their price at time $t - j$ or earlier. Hence, if $t - j$ lies in the distant past, only a small fraction of firms take π_{t-j} into account. This effect is compounded by the fact that, for any $t - j$, the fraction of firms which still charges a price which was optimal at $t - j$ is smaller the further $t - j$ lies in the past. In sum, distant inflation observations will only slightly affect the average frequency of price adjustment, and hence only have a minor impact on the slope of the Phillips curve.

Table 5.3 displays the estimation results for equation (8). For comparison, we include results for a Phillips curve in which the coefficient on trend inflation is set to zero. In equation (8), the coefficient on trend inflation, b , is positive and significant at the 1% level. This result is in line with the theories of Ball, Mankiw, Romer and Dotsey, King, Wolman.

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 in the appendix) increases until 1976, and gradually decreases from that point on.

In conclusion, our results are consistent with the hypothesis that declining trend inflation in Japan decreased the frequency of price adjustment, which in turn caused the Phillips curve to flatten.

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, 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.³²

A testable implication of this model for Japan is that a decrease in the variance of aggregate demand and/or supply shocks would have caused the flattening of the Phillips curve which we documented in section 4.³³

We focus on the variance of inflation to capture aggregate volatility. Such focus on inflation enables us to measure aggregate volatility in an analogous fashion as we measured trend inflation in equation (9). In any case, the variance of nominal GDP would not be as appropriate a measure to capture both aggregate demand and supply shocks. For instance, if aggregate demand is unit-elastic, the effects of an aggregate supply shock on prices will exactly offset its effect on real activity, thus leaving no visible impact on nominal GDP.

Does the slope of the Phillips curve indeed depend on the volatility of past inflation? To investigate this, we regress:

$$\pi_t = \beta(L).\pi_t + [c + d.var_t(\pi)] . [ygap_{t-1} + p.ygap_{t-2}] + \delta.impoil_t + e_t \quad (10)$$

Again, the sample is 1971Q2-2004Q4, and the coefficient restrictions are the same as in equation (2). This time, we write the Phillips curve slope explicitly as a function of the variance of inflation $var_t(\pi)$, where the variance at time t is computed as:

$$var_t(\pi) = \frac{1 - \theta}{\theta - \theta^{J+1}} \sum_{j=1}^J \theta^j . (\pi_{t-j} - \bar{\pi}_t)^2 \quad (11)$$

³² Alberro (1981) also presents empirical evidence for this hypothesis, using a similar approach as Lucas himself.

³³ Ball, 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: decreased aggregate volatility, and thus reduced uncertainty about future optimal prices, enables firms to set their prices for a longer period of time. (See footnote 29.) A lower frequency of price adjustment in turn implies a flatter Phillips curve.

That is, the variance of inflation at time t is a geometrically weighted average of past squared deviations of inflation from its trend, where trend inflation is computed as in equation (9). Again, the baseline values are $\theta = 0.93$ and $J = 71$. In equation (11), note that both $var_t(\pi)$ and $\bar{\pi}_t$ are weighted averages of data from quarter $t - J$ through $t - 1$. That is to say, the variance of inflation is computed based on inflation volatility over the same J quarters as the quarters which are used to compute trend (or average) inflation. Equation (11) has that feature in common with the standard formula for the sample variance. What differentiates the two is the geometric weighting in the former, while the latter is an unweighted average of squared deviations of inflation from its unweighted mean.

Table 5.4 contains the results for equation (10). As predicted by the Lucas-theory, the coefficient d on inflation volatility is positive and significant at the 1% level.

In analogy with the previous subsection, we computed the implied sum of the output gap coefficients. The sum displays the same pattern as inflation volatility: it increases steeply from 1973 to 1975, then decreases gradually over the sample.

The results presented in this subsection are consistent with the hypothesis that a decline in aggregate volatility exacerbated the extent to which firms wrongly interpret aggregate shocks as being firm-specific, which in turn caused the Phillips curve to flatten.

6 Do the non-linear models beat the random walk model?

In section 4, we showed evidence for a significant flattening in a linear Phillips curve. In section 5, we estimated three theoretical models of non-linearity, each of which was found to be consistent with the data. We thus gathered evidence in line with four models of time-variation in the Phillips curve slope, three of which are theoretically founded. In the present section, we investigate whether the three non-linear models add information beyond the atheoretical random walk model. To do so, we estimate models that nest the random walk model and one of the non-linear models, and repeat this exercise for each of the theories of non-linearity. As it turns out, the data favor the non-linear models over the pure random walk model. Given our results in the present section, we will assess the relative performance of the three non-linear models in section 7.

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

In subsection 5.1, we found evidence for convexity in the Phillips curve, of the form predicted by Laxton, Meredith, Rose (LMR). At that point, we did not allow for time-variation in the output-inflation trade-off beyond that implied by the non-linearity of the functional form. In the present subsection, we show that the LMR convexity plays a role even when we allow for independent time-variation in the coefficients on the output gap terms.

We specify a Phillips curve which nests the linear time-varying coefficient model of equations (3) and (4), and the non-linear LMR model of equation (7):

$$\pi_t = \beta(L).\pi_t + \gamma_{1,t} \cdot \left[\left(\frac{\phi.ygap_{t-1}}{\phi - ygap_{t-1}} \right) + p \cdot \left(\frac{\phi.ygap_{t-2}}{\phi - ygap_{t-2}} \right) \right] + \delta.impoil_t + e_t \quad (12)$$

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

$$\gamma_{1,t} = \gamma_{1,t-1} + v_{1,t} \quad (13)$$

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

We estimate this model by Maximum Likelihood Estimation in a Kalman filter framework, with the same distributional assumptions on the error terms as for equations (3) and (4). We estimate two models, one in which the non-linearity parameter ϕ is restricted to be a very large number,³⁴ and one in which ϕ is freely estimated. In the former case, the Phillips curve is essentially linear.

MLE estimates for the unrestricted model are shown in the rightmost column of table 6.1. The non-linearity parameter is small and precisely estimated, be it somewhat larger than in section 5. That is, some of the variation in the Phillips curve slope is now reflected in the path of $\gamma_{1,t}$, while the remaining portion still translates into a convex functional form.

Table 6.1 also compares the value of the log likelihood function for the model in which linearity is imposed with that of the model in which ϕ is freely estimated. Since the two models are nested, we can use a likelihood ratio test to examine whether relaxing the assumption of linearity significantly improves the fit. As we are testing one restriction, the likelihood ratio statistic is approximately distributed $\chi^2(1)$. The likelihood ratio statistic turns out to be 9.12, which exceeds the 1% critical value. Hence, relaxing the assumption of linearity significantly increases the value

³⁴We impose $\phi = 1E20$.

of the likelihood function at the 1% level, even when allowing for independent time-variation in the output gap coefficients.

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

6.2 The trend inflation model beats the random walk model

In subsection 5.2, we discovered a positive, significant relationship between the slope of the Phillips curve and trend inflation, as predicted by Ball, Mankiw, Romer (BMR) and Dotsey, King, Wolman (DKW). In the present subsection, we compare the fit of the BMR / DKW model with that of the linear Phillips curve in which each of the output gap coefficients is modeled as a random walk. To do so, we estimate a model which nests both the trend inflation and random walk models, and test restrictions in both directions.

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

$$\pi_t = [ygap_{t-1} + p.ygap_{t-2}] \cdot \gamma_{1,t} + \begin{bmatrix} \beta_1 & \beta_2 & \beta_3 & (1 - \beta_1 - \beta_2 - \beta_3) & \delta \end{bmatrix} \cdot \begin{bmatrix} \pi_{t-1} \\ \pi_{t-2} \\ \pi_{t-3} \\ \pi_{t-4} \\ impoil_t \end{bmatrix} + e_t \quad (14)$$

The novelty lies in the transition equation. In the encompassing model, the output gap coefficient $\gamma_{1,t}$ is allowed to depend both on its own lag and on trend inflation. First, let us focus on the trend inflation part. Assume that the slope of the Phillips curve, i.e. the *sum* of the output gap coefficients, is a linear function of trend inflation. Assuming that the output gap coefficients are proportional, their sum equals $(1 + p) \cdot \gamma_{1,t}$. Thus, saying that the sum of the output gap coefficients is a linear function of trend inflation amounts to saying that $(1 + p) \cdot \gamma_{1,t} = k_1 + s_1 \cdot \bar{\pi}_t$. Thus, in the pure BMR / DKW model, $\gamma_{1,t} = \frac{k_1}{1+p} + \frac{s_1}{1+p} \cdot \bar{\pi}_t$. On the other hand, in the pure random walk model, $\gamma_{1,t} = \gamma_{1,t-1} + v_{1,t}$. Nesting these two yields the first row of the state equation:

$$\gamma_{1,t} = \lambda \cdot (\gamma_{1,t-1} + v_{1,t}) + (1 - \lambda) \cdot \left(\frac{k_1}{1+p} + \frac{s_1}{1+p} \cdot \bar{\pi}_t \right) \quad (15)$$

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 (9). For θ small and J converging to infinity, we find:

$$\bar{\pi}_{t+1} = (1 - \theta).\pi_t + \theta.\bar{\pi}_t \quad (16)$$

The transition equation thus becomes:

$$\begin{bmatrix} \gamma_{1,t} \\ \bar{\pi}_{t+1} \end{bmatrix} = \begin{bmatrix} (1 - \lambda).\frac{k_1}{1+p} \\ (1 - \theta).\pi_t \end{bmatrix} + \begin{bmatrix} \lambda & (1 - \lambda).\frac{s_1}{1+p} \\ 0 & \theta \end{bmatrix} \cdot \begin{bmatrix} \gamma_{1,t-1} \\ \bar{\pi}_t \end{bmatrix} + \begin{bmatrix} \lambda.v_{1,t} \\ 0 \end{bmatrix} \quad (17)$$

Table 6.2 first shows results for the estimation of the encompassing model, consisting of equations (14) and (17). In the unrestricted model, λ is estimated to be -0.52. Essentially, the weight on the BMR / DKW model in the transition equation exceeds unity. The following two columns indicate estimation results for the two restricted models. First we restrict $\lambda = 0$, in which case the model reduces to BMR / DKW. The likelihood ratio statistic, which again has a $\chi^2(1)$ distribution, reveals that relaxing the restriction that $\lambda = 0$ does not significantly improve the fit, not even at the 10% level. Next we restrict $\lambda = 1$, which amounts to estimating the random walk model. In this case, the likelihood ratio statistic suggests that relaxing the restriction that $\lambda = 1$ significantly improves 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.

6.3 The misperceptions model beats the random walk model

Our results in section 5.3 suggested that the slope of the Phillips curve is a positive, significant function of the variance of inflation, which is in line with the Lucas misperceptions model. We are about to test whether the misperceptions model provides a significantly tighter fit than the random walk model. The procedure is analogous to that in the previous subsection.

The measurement equation is exactly the same as equation (14).

The state equation is analogous to equation (17), 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 (15):

$$\gamma_{1,t} = \lambda.(\gamma_{1,t-1} + v_{1,t}) + (1 - \lambda).\left[\frac{k_2}{1+p} + \frac{s_2}{1+p}.var_t(\pi)\right] \quad (18)$$

From the definition of the variance of inflation in equation (11), we derive its transition process, to be included in the second row of the state equation. For θ small and J converging to infinity, we find that:

$$var_{t+1}(\pi) = (1 - \theta).X_t + \theta.var_t(\pi) \quad (19)$$

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:

$$\begin{bmatrix} \gamma_{1,t} \\ var_{t+1}(\pi) \end{bmatrix} = \begin{bmatrix} (1 - \lambda).\frac{k_2}{1+p} \\ (1 - \theta).X_t \end{bmatrix} + \begin{bmatrix} \lambda & (1 - \lambda).\frac{s_2}{1+p} \\ 0 & \theta \end{bmatrix} \cdot \begin{bmatrix} \gamma_{1,t-1} \\ var_t(\pi) \end{bmatrix} + \begin{bmatrix} \lambda.v_{1,t} \\ 0 \end{bmatrix} \quad (20)$$

Where X_t is as defined after equation (19).

Table 6.3 presents the estimation results for the encompassing model, which consists of equations (14) and (20). The nesting parameter λ is estimated to be 0.20 and not significantly different from zero. 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 $\lambda = 1$ 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.

7 Which type of non-linearity in the Phillips curve?

In section 5, we introduced three different theories, each of which posits a different set of determinants for the slope of the output-inflation trade-off. We can think of these models as implying different types of non-linearity in the Phillips curve. We found that each of the three models is consistent with the data. In section 6, we discovered that any of the non-linear models adds

information beyond a purely atheoretical random walk model. In the present section, we compare the three theories' success in explaining the flattening of Japan's Phillips curve. In that way, we obtain an insight as to which models are the most likely explanation for the absence of accelerating deflation in that country during the period 1998-2002.

7.1 Non-nested model fit comparison

We perform three procedures for non-nested hypothesis testing, all of which are based on OLS / NLS results. First, we regress models which nest two of the non-linear models, and repeat this exercise for each set of two models. Second, we perform Davidson-MacKinnon tests, which entail adding the fitted value from one model as a regressor to another model, and testing whether that fitted value is significant. Third, we apply Bayesian model averaging methods as in Brock, Durlauf, West (2004) to obtain the pseudo-posterior model odds of the three non-linear models (and the linear model).

First, we regress Phillips curves which nest two non-linear models. For example, the following equation nests the Laxton, Meredith, Rose (LMR) model from equation (7) and the Ball, Mankiw, Romer / Dotsey, King, Wolman (BMR / DKW) model of equation (8):

$$\pi_t = \beta(L).\pi_t + [a + b. |\bar{\pi}_t|] \cdot \left[\left(\frac{\phi.ygap_{t-1}}{\phi - ygap_{t-1}} \right) + p. \left(\frac{\phi.ygap_{t-2}}{\phi - ygap_{t-2}} \right) \right] + \delta.impoil_t + e_t \quad (21)$$

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 non-linearity parameter ϕ is estimated to be 12.94, with a somewhat larger standard error than in equations (7) or (12), 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 to find that both variables enter insignificantly. For now, note that the Bayesian model averaging method below does suggest a substantial difference between the performance of the BMR / DKW model and that of the Lucas model.

Second, we discuss results from pairwise non-nested tests as developed by Davidson and MacK-

innon (1981). To examine whether model A adds information beyond that contained in model B, we add the fitted value from model A as a regressor to model B, and test whether the fitted value enters significantly. We do so for each set of two models. The LMR convexity turns out to perform poorly relative to the other two models. That is, the coefficient on the fitted value from the BMR / DKW model, when added to a LMR regression, is significant at the 5% level. 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. Furthermore, we find that the fitted values from BMR / DKW do not enter significantly in the Lucas model, and vice versa.

Third, to obtain a measure for the probability that either of the three models is the true model, 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 non-linear models and the linear model. This procedure assumes a uniform prior distribution over the model space. As Table 7.1 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 that either the LMR or the linear model is 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. As for the relative performance of those two other models, the pseudo-posterior distribution appears to favor the Ball, Mankiw, Romer / Dotsey, King, Wolman (BMR / DKW) model over the Lucas model.

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

negative output gaps as an additional reason, beyond the test results in subsection 7.1, why the LMR story is not the most likely explanation for the flattening of Japan's Phillips curve through the nineties.

Upon closer consideration of the regression results from section 5, it appears that the accurate fit of the LMR model in a regression over 1971Q2-2004Q4 is mostly driven by its superior fit around the time of the first oil price shock. In fact, much of the non-linearity 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 non-linear 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 non-linearity 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 after pre-1975 observations are excluded.

8 Conclusions

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, time-invariant short-run relation between the output gap and inflation. 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 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 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.

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 three implications for monetary policy makers in economies where trend inflation is low today compared to past inflation experience.

First, the Ball, Mankiw, Romer and Dotsey, King, Wolman models imply higher costs of disinflation, at comparatively low initial rates of inflation, than those implied by the standard linear model. Meanwhile, these endogenous pricing models do not predict that the benefits from disinflating are any different from those implied by the linear model. For any given economy in which trend inflation is currently low compared to past experience, the endogenous pricing models thus imply higher optimal inflation targets. This is a generalization from Nishizaki and Watanabe (2000), who argue that central banks should target positive inflation rather than aiming to achieve price stability.

Second, a flatter Phillips curve implies that, absent any policy intervention, inflation becomes less volatile relative to output volatility. In that event, inflation becomes less informative as an indicator for the monetary policy maker. If the central bank's policy rule assigns too much weight to inflation stabilization relative to the weight on output, this may imply that the central bank is not actually stabilizing the real economy. That is, the central bank appears to stabilize inflation, but in truth merely observes low, stable inflation resulting from the flattening of the Phillips curve. A similar point is made in Nishizaki and Watanabe (2000).

Third, 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. If trend inflation falls below zero and grows more negative, these models predict that the frequency of price adjustment starts to increase again. 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.

It remains to describe future work. In sticky price models with endogenous pricing, such as Ball, Mankiw, Romer and Dotsey, King, Wolman, changes in the frequency of price adjustment constitute the mechanism through which changes in trend inflation affect the slope of the Phillips curve. In the present paper, we assess whether the slope of the Phillips curve depends positively on trend inflation, without however explicitly modeling the frequency of price adjustment. In future work, we intend to perform a more direct analysis of the relationship between trend inflation and

the frequency of price adjustment. Such study would focus on groups of countries for which data on the frequency of price adjustment are currently available, such as the euro area. In particular, we aim to investigate whether the response of the frequency of price adjustment to trend inflation is stronger, or weaker, at low inflation rates than at higher rates of inflation.

References

- [1] AKERLOF, GEORGE, WILLIAM DICKENS, AND GEORGE PERRY (1996): "The Macroeconomics of Low Inflation", *Brookings Papers on Economic Activity*, 1996 (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] BAKHSHI, HASAN, PABLO BURRIEL-LLOMBART, HASHMAT KHAN, BARBARA RUDOLF (2003): "Endogenous Price Stickiness, Trend Inflation, and the New Keynesian Phillips Curve", Bank of England Working Paper No. 191.
- [4] BAKHSHI, HASAN, HASHMAT KHAN, BARBARA RUDOLF (2004): "The Phillips Curve under State-Dependent Pricing", mimeo, Bank of England and Swiss National Bank.
- [5] BALL, LAURENCE (1994): "What determines the sacrifice ratio?", in N. Gregory Mankiw (ed.), *Monetary Policy*, University of Chicago Press.
- [6] BALL, LAURENCE, GREGORY MANKIW, AND DAVID ROMER (1988): "The New Keynesian Economics and the Output-Inflation Trade-Off", *Brookings Papers on Economic Activity*, 1988 (1), 1-82.
- [7] BANK OF JAPAN, RESEARCH AND STATISTICS DEPARTMENT (2000): "Price Developments in Japan - a Review Focusing on the 1990s", *Bank of Japan Monthly Bulletin*, October 2000.
- [8] BAYOUMI, TAMIM (2000): "Where Are We Going? The Output Gap and Potential Growth", in Bayoumi and Collins (eds.), *Post-Bubble Blues; How Japan Responded to Asset Price Collapse*, International Monetary Fund.
- [9] BERNANKE, BEN (1999): "Japanese Monetary Policy: A Case of Self-Induced Paralysis?", in Mikitani and Posen (eds.), *Japan's Financial Crisis and its Parallels to U.S. Experience*, Institute for International Economics.
- [10] BLANCHARD, OLIVIER (2000): "Bubbles, Liquidity Traps, and Monetary Policy", in Mikitani and Posen (eds.), *Japan's Financial Crisis and its Parallels to U.S. Experience*, Institute for International Economics.

- [11] BROCK, WILLIAM, STEVEN DURLAUF, AND KENNETH WEST (2004): "Model Uncertainty and Policy Evaluation: some Theory and Empirics", *NBER Working Paper* No. 10916.
- [12] CALVO, GUILLERMO (1983): "Staggered Prices in a Utility-Maximizing Framework", *Journal of Monetary Economics*, 12, 383-398.
- [13] DAVIDSON, R., AND J. G. MACKINNON (1981): "Several Tests of Model Specification in the Presence of Alternative Hypotheses", *Econometrica*, 49, 781-793.
- [14] 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 (2), 249-282.
- [15] DEBELLE, GUY AND JAMES VICKERY (1997): "Is the Phillips Curve a Curve? Some Evidence and Implications for Australia", Reserve Bank of Australia Research Discussion Paper No. 9706.
- [16] DEFINA, ROBERT H. (1991): "International Evidence on a New Keynesian Theory of the Output-Inflation Trade-Off", *Journal of Money, Credit and Banking*, 23 (3), 410-422.
- [17] DOTSEY, MICHAEL, ROBERT KING, AND ALEXANDER WOLMAN (1999): "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output", *Quarterly Journal of Economics*, May 1999, 655-690.
- [18] FAUST, JON, JOHN ROGERS, SHING-YI WANG, JONATHAN WRIGHT (2003): "The High Frequency Response of Exchange Rates and Interest Rates to Macroeconomic Announcements", International Finance Discussion Paper 784, Board of Governors of the Federal Reserve System (US).
- [19] FUHRER, JEFFREY AND GEORGE MOORE (1995): "Inflation Persistence", *Quarterly Journal of Economics*, 110, 127-159.
- [20] HAMILTON, JAMES (1994): *"Time Series Analysis"*, Princeton University Press, Princeton, New Jersey.
- [21] HARVEY, ANDREW (1994): *"Time Series Models"*, The MIT Press, Cambridge, Massachusetts.


- [22] HAYAKAWA, HIDEO, AND HIROSHI UGAI (2000): "Why Did Prices in Japan Hardly Decline During the 1997-98 Recession?", BIS Papers No. 3, 139-173.
- [23] HIROSE, YASUO, AND KOICHIRO KAMADA (2003): "A New Technique for Simultaneous Estimation of the Output Gap and Phillips Curve", *Monetary and Economic Studies*, 21 (2), 93-112.
- [24] HO, WAI-YIP ALEX AND JAMES YETMAN (2005): "The Long-Run Output-Inflation Trade-Off in the Presence of Menu Costs", mimeo, The University of Hong Kong.
- [25] 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*, forthcoming.
- [26] KILEY, MICHAEL (1996): "Endogenous Price Stickiness and Business Cycle Persistence", Finance and Economics Discussion Series 96-33, Board of Governors of the Federal Reserve System (US).
- [27] _____ (2005): "A Quantitative Comparison of Sticky-Price and Sticky-Information Models of Price Setting", mimeo, Board of Governors of the Federal Reserve System (US).
- [28] KIM, CHANG-JIN AND CHARLES NELSON (1999): *"State-Space Models with Regime Switching"*, The MIT Press, Cambridge, Massachusetts.
- [29] KRUGMAN, PAUL (1998): "It's Baaack: Japan's Slump and the Return of the Liquidity Trap", *Brookings Papers on Economic Activity*, 1998 (2), 137-205.
- [30] KUMAR, MANMOHAN, TAIMUR BAIG, JOERG DECRESSIN, CHRIS FAULKNER-MACDONAGH, TARHAN FEYZIOGLU (2003): "Deflation: Determinants, Risks, and Policy Options", *IMF Occasional Paper* No. 221.
- [31] KURODA, SACHIKO, AND ISAMU YAMAMOTO (2003a): "Are Japanese Nominal Wages Downwardly Rigid?", *Monetary and Economic Studies*, Bank of Japan, 21 (2), 1-30.
- [32] _____ (2003b): "Are Japanese Nominal Wages Downwardly Rigid?", *Monetary and Economic Studies*, Bank of Japan, 21 (2), 31-68.

- [33] _____ (2005): "Wage Fluctuations in Japan after the Bursting of the Bubble Economy: Downward Nominal Wage Rigidity, Payroll, and the Unemployment Rate", IMES Discussion Paper No. 2005-E-2, Bank of Japan.
- [34] KUTTNER, KENNETH (1994): "Estimating Potential Output as a Latent Variable", *Journal of Business and Economic Statistics*, 12 (3), 361-368.
- [35] LAXTON, DOUGLAS, GUY MEREDITH, AND DAVID ROSE (1995): "Asymmetric Effects of Economic Activity on Inflation", *IMF Staff Papers*, 42 (2), 344-374.
- [36] LUCAS, ROBERT E. (1973): "Some International Evidence on Output-Inflation Trade-Offs", *American Economic Review*, 63 (3), 326-334.
- [37] MCCALLUM, BENNETT T. (2001): "Inflation Targeting and the Liquidity Trap", *NBER Working Paper 8225*.
- [38] 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.
- [39] MORGAN, PETER (2005): "Employment - the Shift to Flexibility", HSBC special report, February 2005.
- [40] 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.
- [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 (100), 283-299.
- [42] ROBERTS, JOHN (1995): "New Keynesian Economics and the Phillips Curve", *Journal of Money, Credit, and Banking*, 27, 975-984.
- [43] SCHALING, ERIC (1999): "The Non-Linear Phillips Curve and Inflation Forecast Targeting", Bank of England.
- [44] SVENSSON, LARS (2001): "The Zero Bound in an Open Economy: A Foolproof Way of Escaping from a Liquidity Trap", *Monetary and Economic Studies*, Bank of Japan, special edition February 2001.

- [45] TAYLOR, JOHN (1980): "Aggregate Dynamics and Staggered Contracts", *Journal of Political Economy*, 88 (1), 1-23.
- [46] TURNER, DAVE (1995): "Speed Limit and Asymmetric Inflation Effects from the Output Gap in the Major Seven Economies", *OECD Economic Studies* 24, 57-87.
- [47] YATES, ANTHONY, AND BRYAN CHAPPLE (1996): "What Determines the Short-Run Output-Inflation Trade-Off?", Bank of England Working Paper No. 53.
- [48] ZHANG, LAWRENCE HUIYAN (2001): "Sacrifice Ratios with Long-Lived Effects", mimeo, Johns Hopkins University.

Figures and tables


The figures and tables are numbered according to the section of the paper which they pertain to. For instance, figure 4.2 indicates the second object in section 4.




Source: US Federal Reserve, Board of Governors

Figure 2.1

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.



Source: US Federal Reserve, Board of Governors



Source: Bank of Japan

Figure 2.2

'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 2.3

The linear Phillips curve is estimated over 1971Q2-1997Q4; the forecast window is 1998Q1-2004Q4. The result is a simple formalization of the perception (based on a linear, time-invariant Phillips curve) that deflation did not accelerate as much in Japan as warranted by the generally large, negative output gaps in the late nineties.




Figure 3.1

This figure compares the Fed output gap with the Bank of Japan's most widely used output gap, which is based on the production function approach. Since the BOJ output gap underlying the present graph is always negative by construction, we rescaled it such that its average equals the average of the Fed output gap over the period of BOJ data availability.




Figure 3.2

This figure contrasts the Fed output gap with a Hodrick Prescott-filter based output gap, where the latter is computed using a smoothing parameter of 1600. With a protracted low-growth period at the end of the sample, the HP-filter trend is automatically close to actual output at that time.




Figure 3.3

The dotted line graphs output gap estimates from applying the original Hirose and Kamada (2003) procedure. The solid, thin line graphs the results from the HK procedure after including relative oil import price inflation in the Phillips curve. Since inflation is fairly stable at the end of the sample, actual output is trivially close to its stable-inflation level, which is reflected in the fact that the HK-output gaps hover around zero at the end of the sample.

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

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)
$\gamma_{1,t}$		0.87** (0.36)	0.75*** (0.16) avg
$\gamma_{2,t} = p.\gamma_{1,t}$		-0.70** (0.35)	-0.58*** (0.12) avg
δ		0.017* (0.009)	0.018*** (0.004)
σ_e		-	1.33*** (0.04)
p		-0.81*** (0.10)	-0.77*** (0.10)
Implied	σ_e^2	-	1.77
variances	$\sigma_{v_1}^2$	-	0.0011
and	$\sigma_{v_2}^2$	-	0.0007
covariances	$cov(v_1, v_2)$	-	-0.0009
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 4.1

This table 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. In this table as in any future table, 'avg' indicates the average of a time-varying coefficient and its standard error over the sample.

Random Walk Model: Time-Varying Output Gap Coefficients




Figure 4.2

This figure graphs the Kalman-smoothed time-varying output gap coefficients corresponding to the estimation results in table 4.1, 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.

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


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)
γ_1	0.87** (0.36)	0.69** (0.30)
$\gamma_2 = p.\gamma_1$	-0.70** (0.35)	-0.54* (0.29)
δ	0.017* (0.009)	0.019** (0.008)
φ	∞	10.00*** (1.83)
p	-0.81*** (0.10)	-0.77*** (0.12)
Fit	$R^2 = 0.85 / \overline{R}^2 = 0.84$	$R^2 = 0.87 / \overline{R}^2 = 0.86$
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 5.1

The rightmost column shows the results from estimating a Laxton, Meredith, Rose-type Phillips curve. The non-linearity parameter φ is precisely estimated. As figure 5.2 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 graphs the results from estimating a linear Phillips curve.



This figure graphs $\gamma_1 \cdot \left[\left(\frac{\phi \cdot ygap_{t-1}}{\phi - ygap_{t-1}} \right) + p \cdot \left(\frac{\phi \cdot ygap_{t-2}}{\phi - ygap_{t-2}} \right) \right]$, as estimated in table 5.1, with respect to the output gap. The curve indicates the total effect of a given output gap level on inflation after two quarters, abstracting from past inflation and supply shocks. 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 .

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

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)
γ_1	0.87** (0.36)	0.76 avg
$\gamma_2 = p.\gamma_1$	-0.70** (0.35)	-0.63 avg
δ	0.017* (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 5.3

The rightmost column contains the results from estimating a Phillips curve in which the slope depends on 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.

$$\pi_t = \beta(L).\pi_t + [c + d.var_t(\pi)] . [ygap_{t-1} + p.ygap_{t-2}] + \delta.impoil_t + e_t$$

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)
γ_1	0.87** (0.36)	0.66 avg
$\gamma_2 = p.\gamma_1$	-0.70** (0.35)	-0.52 avg
δ	0.017* (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 5.4

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.

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

Sample: 1971Q2-2004Q4	MLE (linear TV)	MLE (LMR TV)
β_1	0.67*** (0.08)	0.76*** (0.08)
β_2	0.17** (0.08)	0.01 (0.08)
β_3	0.37*** (0.08)	0.48*** (0.08)
$\gamma_{1,t}$	0.75*** (0.16) avg	0.58*** (0.15) avg
$\gamma_{2,t} = p.\gamma_{1,t}$	-0.58*** (0.12) avg	-0.45*** (0.11) avg
δ	0.018*** (0.004)	0.018*** (0.004)
φ	1E20	12.32*** (2.55)
σ_e	1.33*** (0.04)	1.31*** (0.04)
p	-0.77*** (0.10)	-0.78*** (0.08)
Fit	LLF=-255.94	LLF=-251.38

Standard errors are in parentheses.

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

Table 6.1

The rightmost column shows the results from a model which nests the Laxton, Meredith, Rose model from table 5.1 with the random walk model from table 4.1. The middle column shows results from the model in which the Phillips curve is linear, but its slope varies over time as a random walk. Relaxing the assumption of linearity significantly increases the value of the likelihood function at the 1% level. The results in this table, as in the following two tables, are computed by means of the Kalman filter and Maximum Likelihood.

Sample:	MLE (encompassing)	MLE (BMR / DKW)	MLE (random walk)
1971Q2-2004Q4			
β_1	0.79*** (0.07)	0.77*** (0.07)	0.68*** (0.08)
β_2	0.16** (0.07)	0.17** (0.07)	0.17** (0.08)
β_3	0.24*** (0.07)	0.24*** (0.07)	0.36*** (0.08)
$\gamma_{1,t}$	0.78*** (0.08) avg	0.82*** (0.03) avg	0.73*** (0.18) avg
$\gamma_{2,t} = p \cdot \gamma_{1,t}$	-0.64*** (0.06) avg	-0.67*** (0.03) avg	-0.57*** (0.14) avg
δ	0.019*** (0.003)	0.019*** (0.003)	0.019*** (0.004)
k_1	-0.08** (0.04)	-0.08** (0.04)	-
s_1	0.06*** (0.02)	0.06*** (0.02)	-
λ	-0.52** (0.23)	0.00	1.00
σ_e	1.24*** (0.04)	1.24*** (0.04)	1.32*** (0.04)
p	-0.82*** (0.06)	-0.82*** (0.06)	-0.78*** (0.09)
Fit	LLF=-234.21	LLF=-235.45	LLF=-253.18

Standard errors are in parentheses.

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

Table 6.2

The column labeled 'MLE (encompassing)' contains results for a model which nests the Ball, Mankiw, Romer / Dotsey, King, Wolman model of table 5.3 with the random walk model of table 4.1. The next column restricts λ to be zero, such that the model reduces to BMR / DKW. The rightmost column sets λ to one, such that the model collapses to the random walk model. The former restriction does not lead to a statistically significant deterioration in the value of the log likelihood function, while imposing the latter restriction implies a significant deterioration in fit at the 1% level.

Sample: 1971Q2-2004Q4	MLE (encompassing)	MLE (Lucas)	MLE (random walk)
β_1	0.78*** (0.09)	0.67*** (0.09)	0.68*** (0.08)
β_2	0.25*** (0.07)	0.30*** (0.07)	0.17** (0.08)
β_3	0.11 (0.08)	0.18** (0.07)	0.36*** (0.08)
$\gamma_{1,t}$	0.63*** (0.06) avg	0.69*** (0.03) avg	0.73*** (0.18) avg
$\gamma_{2,t} = p \cdot \gamma_{1,t}$	-0.48*** (0.04) avg	-0.48*** (0.02) avg	-0.57*** (0.14) avg
δ	0.020*** (0.003)	0.020*** (0.003)	0.019*** (0.004)
k_2	-0.050 (0.050)	-0.008 (0.056)	-
s_2	0.012** (0.006)	0.013*** (0.005)	-
λ	0.20 (0.20)	0.00	1.00
σ_e	1.26*** (0.04)	1.26*** (0.04)	1.32*** (0.04)
p	-0.77*** (0.12)	-0.70*** (0.10)	-0.78*** (0.09)
Fit	LLF=-238.80	LLF=-239.72	LLF=-253.18

Standard errors are in parentheses.

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

Table 6.3

This table shows results analogous to the previous table, but for the Lucas model. Again, imposing $\lambda = 0$ does not significantly worsen the fit, while restricting $\lambda = 1$ leads to a significant decline in the log likelihood function value 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 7.1

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.