DNB Working Paper
No. 735 / January 2022

How Does the Phillips Curve Slope Vary with Repricing Rates?

Emmanuel de Veirman

DeNederlandscheBank
EUROSYSTEEM

How Does the Phillips Curve Slope Vary with Repricing Rates?

Emmanuel de Veirman*

* Views expressed are those of the authors and do not necessarily reflect official positions of De Nederlandsche Bank.

Working Paper No. 735

De Nederlandsche Bank NV
P.O. Box 98
1000 AB AMSTERDAM
The Netherlands

January 2022

How Does the Phillips Curve Slope Vary with Repricing Rates?

Emmanuel De Veirman*

January 2022

In sticky price models, the slope of the Phillips curve depends positively on the probability of price adjustment. I use a series for the empirical frequency of price adjustment to test this implication. I find some evidence that the Phillips curve slope depends positively on the repricing rate. My results support the implication from New Keynesian theory with Calvo pricing that the Phillips curve slope is a convex function of the frequency of price adjustment. However, at all observed values of the frequency of price adjustment, the empirical Phillips curve relation is much flatter than the New Keynesian Phillips Curve at standard parameter values would imply.

Keywords: Inflation, Phillips curve, price setting

JEL codes: C22, E31

*Research Department, De Nederlandsche Bank. E-mail: `manu.veirman@dnb.nl`. I thank seminar participants at the European Central Bank, as well as Guido Ascari, Larry Ball, Paolo Bonomolo, Maurice Bun, Luca Dedola, Peter Karadi, Michele Lenza, Anton Nakov, Giovanni Olivei, Giorgio Primiceri, Jirka Slacalek, Jón Steinsson and Ludwig Straub for input. I carried out a substantial part of the work on this paper while I was seconded to the European Central Bank (DG-Research). The views expressed in this paper are those of the author and do not necessarily reflect those of De Nederlandsche Bank, the European Central Bank or the Eurosystem. Any errors are my own.

1 Introduction

A cornerstone implication of sticky price models is that when firms change their prices more frequently, aggregate demand fluctuations have a larger effect on inflation in the short run, i.e. the short-run Phillips curve is steeper. In particular, the New Keynesian Phillips Curve (NKPC), which is derived from a New Keynesian model with pricing as in Calvo (1983), implies that the slope is a positive and convex function of the probability of price adjustment.

In this paper, I examine empirically how the Phillips curve slope depends on the frequency of price adjustment, using the series for the adjustment frequency that Nakamura, Steinsson, Sun and Villar (2018) computed from price-level data.

To my knowledge, the present paper is the first to examine the relation between the Phillips curve slope and the frequency of price adjustment empirically. Earlier papers that test implications of sticky price models for the Phillips curve slope, such as Ball, Mankiw and Romer (1988), DeFina (1991), De Veirman (2009) and Ball and Mazumder (2011), examine the role of variables such as trend inflation or aggregate volatility, which in theory affect the Phillips curve slope through the frequency of price adjustment. Against the background that a long time series for the frequency of price adjustment was not yet available at that time, these papers did not explicitly examine the influence of the adjustment frequency on the Phillips curve slope.[1]

When I model the Phillips curve slope as a linear function of the adjustment frequency, I detect a statistically significant positive relation between the slope and the frequency when I use an inflation measure that arguably has the least amount of noise. My findings suggest

[1] Bils and Klenow (2004) and Nakamura and Steinsson (2008) were among the first to provide comprehensive evidence on the frequency of price adjustment in the United States.

that the Phillips curve slope was essentially zero at most times since the late 1990s due to relatively infrequent price adjustment at that time. For the period from the Great Recession onwards, this is consistent with the fact that inflation remained stable notwithstanding a long sequence of negative output gaps.

These results relate to an empirical literature that finds that the Phillips curve has flattened in the United States. See, for instance, Ball and Mazumder (2011), Blanchard (2016) and Del Negro, Lenza, Primiceri and Tambalotti (2020).[2] My results suggest that the Phillips curve may have flattened endogenously, as a result of declining repricing rates.

Next, I estimate Phillips curves with a slope that depends on the frequency of price adjustment according to the same non-linear functional form as that of the NKPC. In line with the New Keynesian model with Calvo pricing, I find that the Phillips curve slope is a convex function of the adjustment frequency.

I also characterize the level of the slope which the NKPC implies given the empirically observed frequencies of price adjustment, with other structural parameters set at standard values. Given the same empirical adjustment frequencies, the slope implied by the structural parameters is much steeper than the slope implied by my estimates at all times in the sample period. On average, the implied theoretical Phillips curve slope is about ten times steeper than the implied empirical slope specified as a non-linear function of the frequency. This finding implies that the Calvo model overstates the short-run response of inflation to aggregate demand fluctuations and understates the real effects.

The NKPC assumes that firms set prices as in Calvo (1983), which is a prominent type

[2]See Costain, Nakov and Petit (2021) for an explanation of this flattening based on a state-dependent pricing model. See De Veirman (2009), Blanchard, Cerutti and Summers (2015) and Okuda, Tsuruga and Zanetti (2021) for evidence of such flattening in other economies.

of time-dependent pricing. With Calvo pricing, firms have to wait until they are given the opportunity to adjust their price so as to react to a change in aggregate demand. In time-dependent pricing models such as those with Calvo price setting, firms that adjust their price cannot self-select into being those that desire the largest price changes. On the contrary, such a selection effect is present in state-dependent pricing models. For the same frequency of price adjustment, this implies that inflation responds more strongly to shocks in state-dependent pricing models than in time-dependent pricing models. Combined with my finding that the Calvo model overstates the speed of adjustment of prices relative to the empirical macro relations, this suggests that state-dependent pricing models would overstate the empirical response by even more.

The remainder of the paper is organized as follows. Section 2 shows that empirically, the Phillips curve slope depends positively on the frequency of price adjustment. Section 3 documents that the relation between the slope and the frequency is convex both in the NKPC and in the data, but that the theoretical slope implied by standard values for the structural parameters is much steeper than the slope of the empirical Phillips curve. Section 4 concludes.

2 Does the Phillips curve slope depend on the frequency of price adjustment?

In this section, I show that empirically, there is evidence for a positive relationship between the Phillips curve slope and the frequency of price adjustment, but this relationship is significant only with trimmed mean PCE inflation.

Figure 1 plots the data. The top panel shows the median quarterly frequency of consumer price adjustment in the United States from Nakamura, Steinsson, Sun and Villar (2018), with the blue line representing an unsmoothed quarterly series and the black line representing a backward-looking four-quarter moving average. Aiming to reduce the impact of any measurement error, I perform analysis with the smoothed series throughout this paper. Note that Nakamura e.a. (2018) focus on an annual series, plausibly for the same reason.

While Nakamura e.a. (2018) report the frequency as the fraction of prices that change per month, I express it as the fraction of price changes per quarter so as to allow for a comparison with the quarterly probability of price adjustment from sticky price models in Section 3.[3]

The middle panel of Figure 1 plots the output gap, which I computed as the percentage difference between real Gross Domestic Product (GDP) and Congressional Budget Office estimates of real potential output.

The bottom panel plots inflation in the deflator for Personal Consumption Expenditures (PCE) excluding food and energy (in blue), as well as trimmed mean PCE inflation from the Federal Reserve Bank of Dallas (the thick black line). Trimmed mean PCE inflation is less volatile than PCE inflation ex food and energy. In particular, the former measure has smaller short-run fluctuations than the latter. Plausibly due to this feature, it turns out that my Phillips curve regressions for trimmed mean PCE inflation provide a better fit and tighter estimates than those with core PCE inflation. As I discuss later in this section, this is why

[3]Figure XIV in Nakamura, Steinsson, Sun and Villar (2018) plots an annual series for the frequency of price adjustment. I thank Jón Steinsson for sending me the underlying unsmoothed quarterly series $\widetilde{freq_t^m}$ expressed as the fraction of price changes per month. From this series, I first compute the frequency expressed as the fraction of price changes per quarter $\widetilde{freq_t^q}$ through the formula $(1 - \widetilde{freq_t^m})^3 = 1 - \widetilde{freq_t^q}$. I then compute the four-quarter moving average as $freq_t = (1/4)\sum_{i=0}^{3} \widetilde{freq_{t-i}^q}$.

I use trimmed mean PCE inflation as the baseline.

Beyond these two measures for core PCE inflation, I also report results using inflation in the constant methodology Bureau of Labor Statistics research series for the Consumer Price Index (CPI) excluding food and energy, and for inflation in the GDP deflator. Throughout, I use quarter-on-quarter annualized inflation.

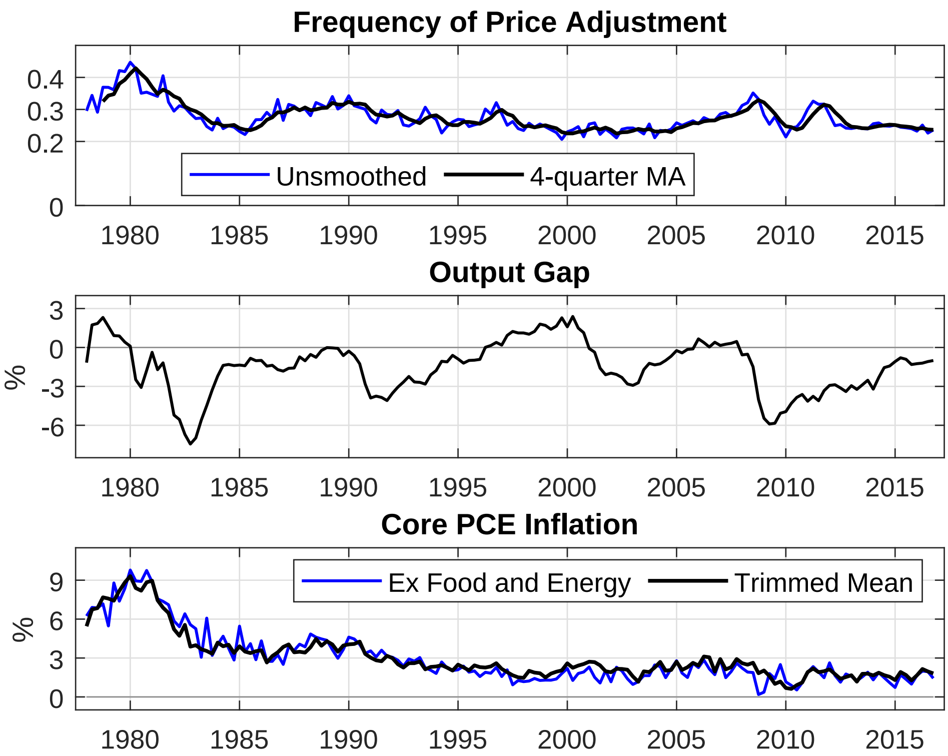

FIGURE 1. DATA

Note: This figure plots quarterly time series for the United States. The top panel plots the median frequency of consumer price adjustment excluding sales, expressed as the share of prices that change per quarter. These data underly Nakamura, Steinsson, Sun and Villar (2018). From the unsmoothed frequency (in blue), I computed the four-quarter backward-looking moving average (the thick black line). The middle panel plots the output gap, defined as the percent deviation of real Gross Domestic Product from the Congressional Budget Office estimate of real potential output. The bottom panel plots inflation in the deflator for Personal Consumption Expenditures ex food and energy (in blue) and trimmed mean PCE inflation from the Federal Reserve Bank of Dallas (the thick black line). Both cases pertain to annualized quarter-on-quarter inflation.

In this section, I specify the Phillips curve slope as a linear relationship of the frequency of price adjustment. With trimmed mean PCE inflation, I estimate the following regression:

$$\pi_t = \beta(L)\pi_t + (a + b\ freq_t)(ygap_t + \sum_{i=1}^{5} p_i\ ygap_{t-i}) + \varepsilon_t \tag{1}$$

where π_t is inflation, $freq_t$ is the frequency of price adjustment, $ygap_t$ is the output gap and ε_t is the residual. Equation (1) assumes adaptive inflation expectations. In particular, $\beta(L)\pi_t = \sum_{i=1}^{7} \beta_i \pi_{t-i}$. I impose $\beta_7 = (1 - \sum_{i=1}^{6} \beta_i)$, i.e., the inflation lag coefficients sum to one. This is equivalent to specifying the equation in terms of changes in inflation. I omit the intercept to avoid the possibility of a long-run trend in inflation. At the 5% level, neither of these two restrictions is rejected. In combination, these two assumptions yield a standard accelerationist Phillips curve.

The sample is 1979Q1-2016Q4, with earlier quarters used for lags. Lag selection is based on the Akaike Information Criterion (AIC).

I measure the Phillips curve slope by the sum of the output gap coefficients. In equation (1), the sum of the output gap coefficients is $(a + b\ freq_t)(1 + \sum_{i=1}^{5} p_i)$. This specification allows the Phillips curve slope to vary over time due to changes in the frequency of price adjustment. Since the p's are time-invariant, it assumes that the coefficients on individual output gap terms remain proportional to one another.

First, I estimate a Phillips curve where I set $b = 0$, such that equation (1) reduces to a standard Phillips curve with a time-invariant slope. (I discuss the case with unrestricted b below.)

The adjusted R-squared of the regression with trimmed mean PCE inflation is 0.95. A Breusch-Godfrey LM test for serial correlation up to eight lags reveals no serial correlation in the residuals.

The leftmost numerical part of Table 1 reports results from Wald tests for the null hypothesis that $a(1+\sum_{i=1}^{5} p_i) = 0$, which is the sum of the output gap coefficients after setting $b = 0$.

Throughout this paper, I use Newey-West heteroskedasticity and autocorrelation robust standard errors.

As the first numerical row shows, the Phillips curve slope is 0.03 but statistically insignificant. This is also quite small in economic terms. In combination with the inflation lag coefficients, the Phillips curve slope implies that if the output gap is 1% for one year and 0 at all other times, annualized quarterly inflation increases by 0.11 percentage points in the long run.

Next, I allow the slope to depend on the frequency of price adjustment by estimating b along with the other coefficients. I perform Wald tests for the null hypothesis that $b(1 + \sum_{i=1}^{5} p_i) = 0$, i.e. that the frequency of price adjustment has no effect on the sum of the output gap coefficients.

As the first numerical row of the right part of Table 1 shows, I find that $b(1+\sum_{i=1}^{5} p_i) = 1.92$, which is statistically significant at the 1% level.

The remainder of the upper left numerical part of Table 1 documents that with other inflation measures, the slope of the accelerationist Phillips curve is also positive. It is statistically significant at the 5% level with inflation in the PCE deflator ex food and energy and in the constant methodology core CPI. It is insignificant with inflation in the GDP deflator.

Since the GDP deflator is the only one among the four inflation measures that captures headline inflation, it is also the only one for which I enter relative oil price inflation as a regressor. For every inflation measure, I use the lag specification that the AIC selects for that particular inflation measure.

The corresponding rows in the right part of the table show that with inflation measures

Table 1 Wald Tests: Phillips Curve Slope and Relation to Frequency of Price Adjustment

		Significant Phillips curve slope?			*Significant relation to frequency?*		
		H_0: $a(1+\sum_{i=1}^{5} p_i)=0$			H_0: $b(1+\sum_{i=1}^{5} p_i)=0$		
infexp	inflation	slope	stderr	pval	rel freq	stderr	pval
$\beta(L)\pi_t$	PCE_TM	0.03	0.02	0.11	1.92**	0.62	0.00
$\beta(L)\pi_t$	PCEX	0.06*	0.03	0.04	2.04	1.18	0.09
$\beta(L)\pi_t$	GDPDEF	0.04	0.03	0.24	-0.05	0.22	0.84
$\beta(L)\pi_t$	CPIX	0.07*	0.03	0.03	0.81	0.62	0.19
$\pi^e_{t\|t-1}$	GDPDEF	0.02	0.04	0.56	0.86	0.73	0.24
$\pi^e_{t+1\|t}$	GDPDEF	0.05	0.04	0.22	0.85	0.87	0.33
$\pi^e_{t+1\|t-1}$	GDPDEF	0.06	0.04	0.16	1.00	0.88	0.26

Notes: "Significant Phillips curve slope?" pertains to Wald tests for the null hypothesis that the sum of the output gap coefficients is zero after imposing $b = 0$. "Significant relation to frequency?" tests whether the sum of the output gap coefficients depends on a four-quarter moving average of the frequency of price adjustment. In each case, I report the value for the function of the coefficients that is indicated above the table, its standard error and the p-value of the F-statistic. $\beta(L)\pi_t$ stands for adaptive expectations as in equation (1). For any quarters τ_1 and τ_2, $\pi^e_{\tau_2|\tau_1}$ stands for Survey of Professional Forecasters forecasts formed in τ_1 for inflation in the GDP deflator at time τ_2. PCE_TM is trimmed mean PCE inflation from the Dallas Fed; PCEX is PCE inflation ex food and energy; GDPDEF is GDP deflator inflation; CPIX is inflation in the constant methodology research series for the CPI ex food and energy. Throughout, lag selection is based on the AIC. The sample is 1979Q1-2016Q4, with earlier quarters used for lags. I use Newey-West standard errors. * marks significance at the 5% level; ** at the 1% level.

other than trimmed mean PCE inflation, there is no statistically significant relation between the slope of the accelerationist Phillips curve and the frequency of price adjustment.

Taken together, the evidence in favor of the hypothesis that the slope depends on the price adjustment frequency is not overwhelming. One interpretation is that the specification with trimmed mean PCE inflation is better able to pinpoint the relation. Arguably, this inflation measure is better at cleaning out transitory relative price shocks than ex food and energy measures of core inflation are, such that variation in this inflation measure can be more easily explained by the regressors in the Phillips curve. In the accelerationist Phillips curve where the slope depends linearly on the frequency of price adjustment, the adjusted R-squared is 0.95 with trimmed mean inflation, as opposed to 0.88 for PCE inflation ex food and energy and somewhat lower still for the other inflation measures. The higher R-squared goes along with a smaller standard error of the regression, which tends to imply tighter confidence bands around the regression point estimates.[4]

Adopting the above interpretation, in much of the paper I focus on results with trimmed mean PCE inflation.

For the case with trimmed mean PCE inflation, the blue line in Section 3's Figure 3 plots the time variation in the Phillips curve slope implied by the coefficient estimates for equation (1) and the empirical frequencies of price adjustment. I find that the Phillips curve slope varied between -0.06 in 2002Q2 and 0.33 in 1980Q2. Relative to the constant-coefficients slope of 0.03, that variation is substantial. (The interpretation of the other lines in this figure will become clear in Section 3.)

[4]This interpretation is akin to that of Ball and Mazumder (2019a) on the measurement of the Phillips curve slope with median inflation.

The implied changes in the Phillips curve slope are at times swift, a type of instability in the output-inflation trade-off which can greatly alter the effects of aggregate demand fluctuations.

A first example of this is the Volcker disinflation. The Phillips curve in which the slope depends linearly on the frequency implies that the slope was 0.33 in 1980Q2. Together with the inflation lag coefficients, this implies that if repricing rates would have stayed as high as in 1980Q2 and in the scenario that the output gap was -1% for a year and zero at all other times, that would have implied a long-run reduction in inflation by 1.16 percentage points. However, the frequency of price adjustment quickly declined from its 1980Q2 peak, such that on average in 1982, the implied Phillips curve slope was 0.13. With this slope, an output gap of -1% for one year implies a long-run reduction in inflation by 0.44 percentage points. This suggests that the output gap, which reached its trough at an average of -6.23 percent in 1982, had an effect on inflation that was less than half of what it would have been if the slope had stayed at its 1980Q2 level.

The Great Recession provides a second example of swift implied changes in the Phillips curve slope. The implied Phillips curve slope was 0.14 in 2008Q4. However, the slope declined swiftly after that, to essentially zero on average in 2010, where it stayed for most of the remainder of the sample period. This suggests that the output gap, which was -4.18% on average in 2010, had virtually no downward effect on inflation, as opposed to what would have occurred if the slope had stayed at 0.14. In the latter case, a -1% output gap for one year would have implied a long-run decline in inflation by 0.50 percentage points. This suggests that the negative output gaps in 2010 would have tended to imply a long-run decline in inflation by $4.18 * 0.50 = 2.09$ percentage points. Against the background that trimmed mean PCE inflation was 1.68% on average in 2008Q1-2016Q4, the implied differences in the effects of the output gap on inflation are substantial.

To further check robustness, I estimate Phillips curves where I relax the assumption of adaptive expectations. In particular, I regress over 1978Q4-2016Q4:

$$\pi_t = \beta \pi^e_{t+1|t} + (a + b\ freq_t)\ ygap_t + \sum_{i=0}^{6} \gamma_i\ (\pi_{o,t-i} - \pi_{t-i}) + \varepsilon_t \tag{2}$$

where $\pi^e_{t+1|t}$ is the forecast from the Survey of Professional Forecasters (SPF) of GDP deflator inflation in quarter t+1, formed in quarter t. $\pi_{o,t} - \pi_t$ is relative inflation in the West Texas Intermediate spot crude oil price. As before, lag selection is based on the AIC.

An advantage of using survey expectations is that this flexibly deals with any structural break in the process by which inflation expectations are formed.[5]

Since the one-quarter ahead forecast is only available over the required time span for GDP deflator inflation, in this context I use inflation in the GDP deflator as the dependent variable. With $\pi^e_{t+1|t}$, the timing of inflation expectations is the same as that of the New Keynesian Phillips curve, which features in Section 3.

The timing of the deadlines by which SPF forecasters need to submit their forecasts and those of the Bureau of Economic Analysis publications of the GDP deflator are such that when SPF forecasters make their forecast in quarter t, they know the preliminary estimate of the GDP deflator in t-1, but do not know the GDP deflator in t. In this context, equation(2)

[5]The results from Ball and Mazumder (2019b) suggest that US inflation expectations became anchored from 1998 onwards.

treats $\pi^e_{t+1|t}$ as exogenous.

The row marked $\pi^e_{t+1|t}$ in Table 1 summarizes the results. When I set $b = 0$ in equation (2), the point estimate for the Phillips curve slope is 0.05, but insignificant.

When I estimate b along with the other coefficients in equation (2), I do not find a significant relation between the Phillips curve slope and the frequency of price adjustment.

The rows $\pi^e_{t|t-1}$ and $\pi^e_{t+1|t-1}$ show that the results are similar when I instead use, respectively, SPF forecasts for quarter t formed in quarter t-1 and SPF forecasts for quarter t+1 formed in t-1.[6]

3 How does the Phillips curve slope depend on the frequency of price adjustment? Theory vs. empirics

In this section, I document that both in the New Keynesian model and in the data, the Phillips curve slope is a convex, increasing function of the repricing rate. I also show that at all frequencies of adjustment that occurred in the United States, the slope of the empirical Phillips curve is much flatter than the model-based slope at standard values of the structural parameters.

I first turn to theory, so as to set up the stage for a comparison with the empirical estimates that I present later in this section. In sticky price models, the short-run Phillips curve is steeper when firms reprice more frequently. This intuition is formalized in the New Keynesian Phillips Curve (NKPC). In the canonical NKPC, inflation π_t depends on the

[6]For each of the three types of survey expectations, I choose the lag specification based on the AICs that apply for that particular type of expectations.

contemporaneous output gap $ygap_t$ and on current expectations of future inflation $E_t\pi_{t+1}$:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa \; ygap_t \tag{3}$$

where β is the household time discount factor, and where expectations are rational. As Galí (2008) shows, equation (3) follows from a New Keynesian model with Calvo pricing. The coefficient on the output gap is:

$$\kappa \equiv \left(\frac{(1-\theta)(1-\theta\beta)}{\theta}\right)\left(\frac{1-\alpha}{1-\alpha+\alpha\epsilon}\right)\left(\frac{\phi+\alpha}{1-\alpha}+\sigma\right) \tag{4}$$

where θ is the probability that a firm cannot reset its price, $1-\alpha$ governs the marginal product of labor, ϵ governs the price elasticity of demand, ϕ governs the elasticity of labor supply and σ is the inverse of the household's intertemporal elasticity of substitution.

To examine what this implies for the level of the Phillips curve slope and for its relation to the adjustment probability, I now calibrate the structural parameters in (4). I calibrate most of the structural parameters to the full information estimates of Smets and Wouters (2007). In particular, I set $\beta = 0.9984$, $\alpha = 0.19$, $\phi = 1.83$ and $\sigma = 1.38$. Furthermore, I set $\epsilon = 6$.

In Smets and Wouters (2007), all firms adjust prices every quarter, with a fraction of firms setting their prices optimally and a fraction indexing partially to inflation. Therefore, the Smets-Wouters model does not feature a parameter that could directly be compared to the empirical frequency of price adjustment.

In Galí and Gertler (1999), a fraction of firms hold prices fixed every period, with some of the firms that do adjust doing so optimally and the remainder indexing fully to lagged inflation. I calibrate $\theta = 0.834$, which is Galí and Gertler's estimate for the probability of non-adjustment.

I now plug the above values for the structural parameters into equation (4) to obtain the implied Phillips curve slope. Given that Smets and Wouters (2007) estimated their model on non-annualized quarterly data, I multiply the implied slope by four so as to allow for a comparison with the slopes from the empirical Phillips curves that I estimated with annualized quarterly inflation. I obtain $4\kappa = 0.21$. This is substantially steeper than any of the estimated Phillips curve slopes from Table 1 in Section 2.

Moreover, with Galí and Gertler's (1999) estimated $\theta = 0.834$, price adjustment is less frequent than in the data from Nakamura, Steinsson, Sun and Villar (2018). On average in my sample, the empirical frequency of non-adjustment, expressed as a rate per quarter, is 0.7239. This illustrates the fact that typically, model-based estimates of the probability of price adjustment differ from empirical evidence on the adjustment frequency. When I instead set $\theta = 0.7239$, I obtain $4\kappa = 0.68$, which is not in the ballpark of my empirical estimates of the Phillips curve slope.

Therefore, with an empirically realistic value for the average frequency of non-adjustment and other structural parameters at standard values, the slope of the New Keynesian Phillips curve is much steeper than what I find in the data.

Note that even in the case with forward-looking survey expectations $\pi^e_{t+1|t}$, the specification required by the data is not entirely the same as that of the New Keynesian Phillips curve. This is because empirical survey expectations typically do not conform to the assumption of rational expectations and because equation (2) controls for relative oil price inflation.

We are about to see that both the theoretical and empirical Phillips curve slopes are convex functions of the adjustment frequency. With convexity, Jensen's inequality implies in both the theoretical and the empirical case that the time average of the slopes implied by the empirical series of adjustment frequencies exceeds the slope evaluated at the average adjustment frequency. It is therefore important to evaluate the Phillips curve slope at a broader range of values for the frequency.

I do so now. The black curve in Figure 2 tracks the slope of the New Keynesian Phillips Curve as a function of the probability of price adjustment $1-\theta$, keeping other structural parameters at the above-mentioned values. The function is upward sloping, such that an increase in the adjustment probability implies a steeper Phillips curve. The function is convex, such that the slope varies more strongly with changes in repricing rates at times when many firms change their prices.

These are general features of the NKPC. By inspecting the first derivative of the slope of the NKPC from equation (4) with respect to the non-adjustment probability, we can see that the slope of the NKPC is decreasing in the probability of non-adjustment:

$$\frac{\delta\kappa}{\delta\theta} = \left(\frac{\beta\theta^2-1}{\theta^2}\right)\left(\frac{1-\alpha}{1-\alpha+\alpha\epsilon}\right)\left(\frac{\phi+\alpha}{1-\alpha}+\sigma\right) \tag{5}$$

If $0<\beta<1$ and $0<\theta<1$, then $\beta\theta^2-1<0$ and $\theta^2>0$. If $0\leq\alpha<1$ and $\epsilon\geq 0$, then $1-\alpha>0$ and $1-\alpha+\alpha\epsilon>0$. If, in addition, $\phi\geq 0$ and $\sigma\geq 0$, then $[(\phi+\alpha)/(1-\alpha)]+\sigma\geq 0$. As a result, $\delta\kappa/\delta\theta\leq 0$.

Since the slope of the NKPC is decreasing in the probability of non-adjustment θ, it is increasing in the probability of adjustment $1-\theta$.

Second, from the second derivative we can see that the slope of the NKPC is a convex function of the probability of non-adjustment:

$$\frac{\delta^2\kappa}{\delta\theta^2} = \left(\frac{2}{\theta^3}\right)\left(\frac{1-\alpha}{1-\alpha+\alpha\epsilon}\right)\left(\frac{\phi+\alpha}{1-\alpha}+\sigma\right) \tag{6}$$

If $\theta > 0$, then $\theta^3 > 0$ and $\delta^2\kappa/\delta\theta^2 \geq 0$.

The blue line in Figure 2 charts the relationship between the empirical Phillips curve slope from Section 2 and the frequency of price adjustment. As we found in Section 2, this relation is upward sloping. Recall that in that section, I imposed a linear relationship between the slope and the frequency.

In the United States in 1978Q1-2016Q4, the frequency of price adjustment ranged from 0.22 to 0.43 per quarter. Over this range, the empirical Phillips curve slope, specified as a linear function of the adjustment frequency, is always well below the slope implied by the New Keynesian model with Calvo pricing. In addition, at most observed values for the frequency, the slope of the New Keynesian Phillips curve varies much more strongly with changes in repricing rates than that of the empirical Phillips curve does.

To see whether the assumption of a linear relation between the slope and the frequency in Section 2 accounts for these differences, I now estimate empirical Phillips curves in which the slope depends on the frequency of price adjustment in a non-linear fashion akin to that of the NKPC. As before, I use trimmed mean PCE inflation.

Defining the probability of price adjustment $\zeta \equiv 1-\theta$, one can rewrite equation (4) as follows:

$$\kappa \equiv \left[(1-\beta)\left(\frac{\zeta}{1-\zeta}\right)+\beta\left(\frac{\zeta^2}{1-\zeta}\right)\right]\left(\frac{1-\alpha}{1-\alpha+\alpha\epsilon}\right)\left(\frac{\phi+\alpha}{1-\alpha}+\sigma\right) \tag{7}$$

such that the Phillips curve depends on two non-linear terms in the probability of price adjustment: $[\zeta/(1-\zeta)]$ and $[\zeta^2/(1-\zeta)]$.

In a first specification, I write the Phillips curve slope as an unrestricted function of these two non-linear terms, after replacing the probability of price adjustment ζ by the empirical frequency of price adjustment $freq_t$:

$$\pi_t = \beta(L)\pi_t + \left[c + d\ \left(\frac{freq_t}{1-freq_t}\right) + e\ \left(\frac{freq_t^2}{1-freq_t}\right)\right](ygap_t + \sum_{i=1}^{5} p_i\ ygap_{t-i}) + \varepsilon_t \tag{8}$$

where $\beta(L)$ is as defined under equation (1). I call this the unrestricted non-linear specification.

In a second specification, I estimate:

$$\pi_t = \beta(L)\pi_t + f\ \left[(1-g)\ \left(\frac{freq_t}{1-freq_t}\right) + g\ \left(\frac{freq_t^2}{1-freq_t}\right)\right](ygap_t + \sum_{i=1}^{5} p_i\ ygap_{t-i}) + \varepsilon_t \tag{9}$$

I call this the restricted non-linear specification due to the fact that I restrict the coefficients on the two non-linear terms to sum to one like in equation (7). The parameter f stands for $[(1-\alpha)/(1-\alpha+\alpha\epsilon)][(\phi+\alpha)/(1-\alpha)+\sigma]$ from that equation.

In Figure 2, the red and green lines chart how the Phillips curve slope depends on the adjustment frequency in, respectively, the unrestricted and restricted non-linear case. In both cases, the function is upward-sloping and convex. Recall that the empirical quarterly adjustment frequency ranges from 0.22 to 0.43 in my sample. For the higher adjustment frequencies within that range, the slopes from the non-linear specifications are closer to the

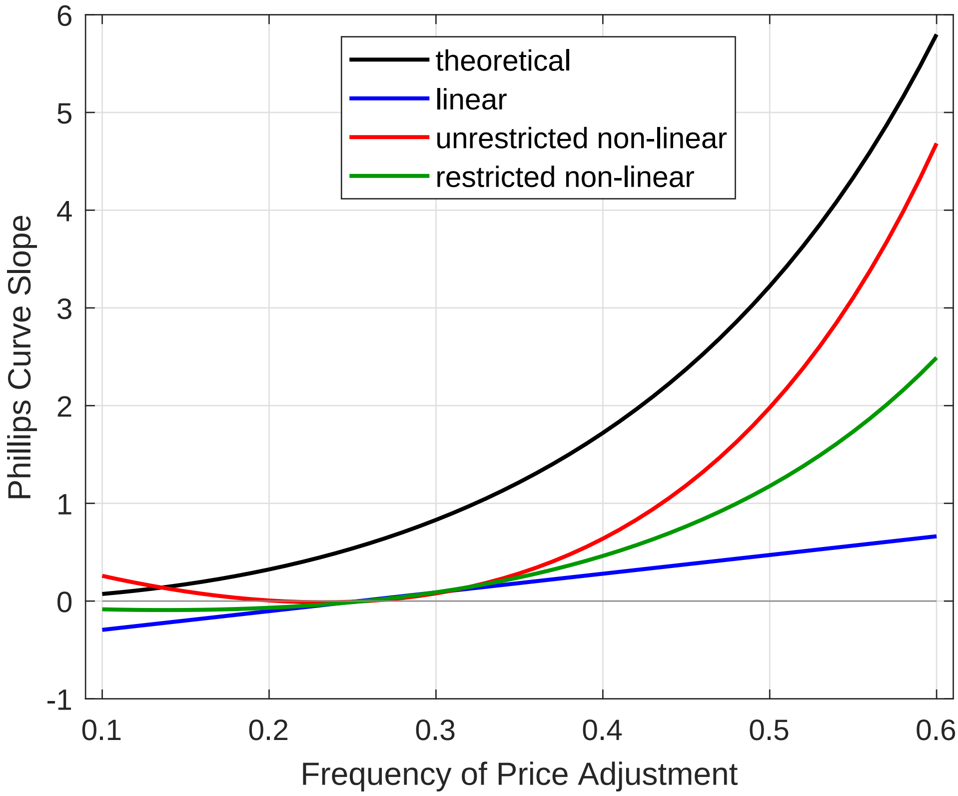

Figure 2. Phillips Curve Slope as a Function of the Frequency of Price Adjustment

Note: The black line charts the slope of the New Keynesian Phillips Curve as a function of the quarterly probability of price adjustment as in equation (4), with other structural parameters at standard values. The other lines indicate the relation between the empirical Phillips curve slope and the frequency of price adjustment in three cases. The blue, red and green lines pertain to the cases where the slope is, respectively, a linear, unrestricted non-linear, and restricted non-linear function of the frequency. The relevant equations are, respectively, (1), (8) and (9). In all cases, I use annualized quarter-on-quarter trimmed mean PCE inflation. In the US in 1978Q4-2016Q4, the adjustment frequency ranged from 0.22 to 0.43. In that range, all empirical slopes are well below the model-implied Phillips curve slope.

theoretical slope than the slope that I specified as a linear function of the frequency. Among the three empirical specifications, the unrestricted non-linear specification is closest to theory.

Still, the empirical Phillips curve slope remains clearly below that of the slope of the New Keynesian Phillips curve at all observed values for the frequency of price adjustment. To see this from another angle, Figure 3 shows the time path of the Phillips curve slope implied by the empirical values for the frequency of price adjustment. The solid black line shows the slope of the NKPC when setting the other structural parameters to their values estimated by Smets and Wouters (2007) and setting $\epsilon = 6$. Because price adjustment was very frequent around 1980, the NKPC implies that the Phillips curve was steep at that time, rising to 2.07 in 1980Q2. However, the convexity of the relationship between the slope and the frequency implies that the slope was particularly unstable at that time. By 1985Q3, it reached 0.47.

For the unrestricted and restricted non-linear case, respectively, the red and green lines in Figure 3 show that relaxing the assumption of a linear relation bridges a non-trivial part of the gap between theory and empirics around 1980, especially so for the unrestricted non-linear case. At all other times, the differences between the linear and non-linear specifications are minor. In all cases in Figure 3, the empirical Phillips curve slope is always well below the theoretical slope.

On average over the sample, the model-based slope plotted in Figure 3 is 0.71. Among the empirical specifications, the time average of the slope is steepest in the unrestricted non-linear case, at 0.07. In this sense, the theoretical slope is typically about ten times steeper than the unrestricted non-linear slope.

Looking at particular episodes, the empirical Phillips curve slopes imply that reducing inflation during the Volcker disinflation was much costlier than the model-based Phillips curve implies.

At the end of the sample, the empirical slopes imply that to the extent that expansionary monetary policy is able to shift the aggregate demand curve to the right, this would stimulate

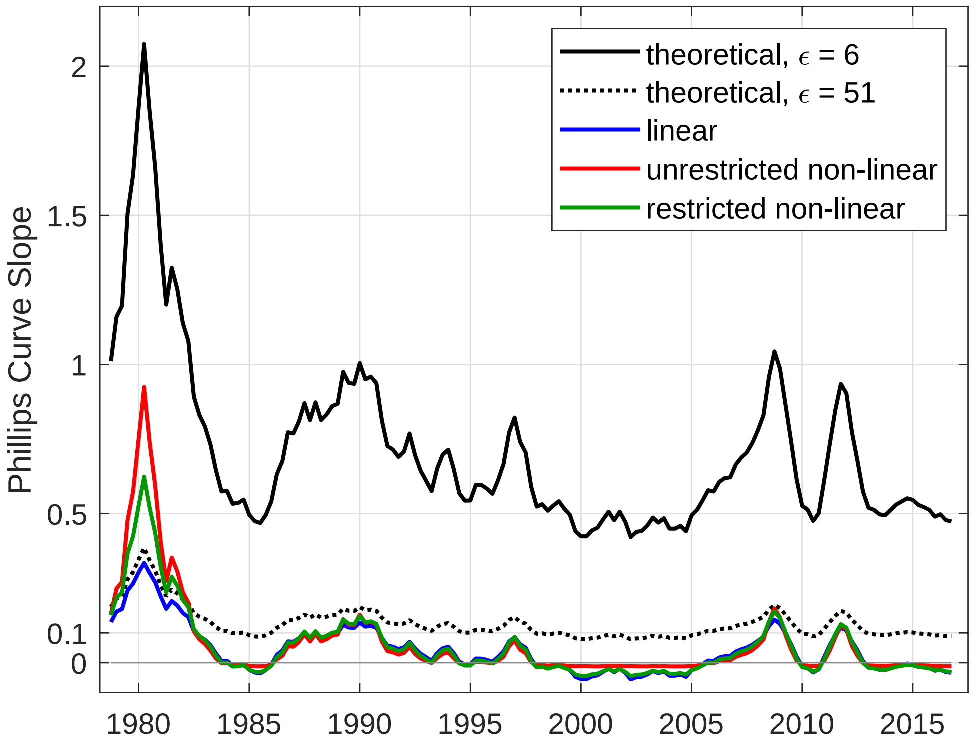

Figure 3. Phillips Curve Slopes Varying with the Empirical Frequency of Price Adjustment

Note: The solid black line charts the slope of the New Keynesian Phillips Curve over time given the changes in the frequency of price adjustment that occurred in the United States, with other structural parameters at standard values including $\epsilon = 6$. The dotted black line charts the same slope with $\epsilon = 51$. The other lines pertain to empirical Phillips curve slopes. The blue, red and green lines represent the case where the slope is, respectively, a linear, unrestricted non-linear, and restricted non-linear function of the frequency. The relevant equations are, respectively, (1), (8), and (9). I use annualized quarter-on-quarter trimmed mean PCE inflation. At all times, all three empirical Phillips curve slopes are well below the model-implied slope with $\epsilon = 6$. Assuming a very high degree of competition as in $\epsilon = 51$ makes the model-based slope much more similar to the empirical slopes.

real output to a substantial extent while it would have small short-run effects on prices. To be sure, the Phillips curve is only one part of the picture. The effect of monetary policy stimulus on aggregate demand depends on other factors, such as whether monetary policy is at the effective lower bound on nominal interest rates and on the extent to which monetary policy transmission through the banking system is operative.

In another effort to explain the discrepancy between theory and empirics, I alter the price elasticity of demand. In the New Keynesian model, a re-optimizing firm raises its price when it observes an increase in the demand for its product. At the same time, the firm takes into account that this price rise tends to dampen the increase in demand, which in turn dampens the optimal price increase. With a high price elasticity of demand ϵ, this dampening effect is large, such that the overall effect of the initial increase in demand on the optimal price is small. This implies that inflation does not increase as much in response to an increase in aggregate demand. In sum, an increase in the elasticity of demand tends to imply a flatter Phillips curve.

The black dotted line in Figure 3 represents the NKPC slope when I set $\epsilon = 51$, implying a very high degree of competition and therefore very elastic demand, while keeping the other structural parameters unchanged.

The increase in the elasticity of demand does imply a substantially flatter Phillips curve. The implied Phillips curve slope is much more similar to the empirical ones. In the first three years of the sample, the empirical slopes specified as a non-linear function of the frequency actually exceed the slope of the model-based Phillips curve with $\epsilon = 51$. At all later times, the theoretical slope still exceeds the empirical ones.

On average over the sample, the model-based slope with $\epsilon = 51$ is 0.13.

Therefore, by imposing a value for the elasticity of substitution that far exceeds the range of typical values, one could explain most of the discrepancy.

The finding that the empirical relationship is convex is robust to using inflation in the constant methodology core CPI and PCE inflation ex food and energy. With GDP deflator inflation, the relationship is estimated to be concave with the unrestricted non-linear specification and close to linear with the restricted non-linear specification, with any of the specifications for inflation expectations in table 1.

With all four inflation measures, all three specifications for survey expectations and at all times, the empirical Phillips curve slopes remain substantially below the theoretical slope at baseline parameters.

The evidence in this section suggests that when one matches the empirical frequency of price adjustment in a model where nominal rigidities follow from Calvo price setting, the model implies a Phillips curve slope that is much steeper than what one actually observes in the data. However, with Calvo pricing, the probability of non-adjustment is a structural parameter. In such a context, agents expect that the probability of price adjustment remains constant. This does not match the observation from Figure 1 that in the data, the frequency of price adjustment varies substantially over time.

It is therefore instructive to gauge what state dependent pricing models would mean for the slope of the Phillips curve. State-dependent pricing models stand a much better chance at explaining the observed short-run fluctuations in the adjustment frequency in that they imply that the frequency of price adjustment varies endogenously. Auclert, Rigato, Rognlie and Straub (2021) show that the implications of state-dependent pricing models such as those by Golosov and Lucas (2007) and Nakamura and Steinsson (2010) are observationally equivalent to those of canonical New Keynesian Phillips Curves with a higher probability of price adjustment. Therefore, to generate the implications from a state-dependent pricing model calibrated at the empirical frequency of price adjustment, one can use a New Keynesian Phillips Curve with higher-than-empirical adjustment frequencies. As I documented under equation (5), the slope of the New Keynesian Phillips Curve is in general increasing in the frequency of price adjustment. Therefore, one can infer that in state-dependent pricing models, the empirical series for the frequency of price adjustment would imply an even steeper output-inflation trade-off than the one implied by Calvo price setting.

Therefore, I infer that for a large class of models with nominal rigidities, the slope of the output-inflation trade-off is steeper than with Calvo pricing, and therefore much steeper than that for the empirical Phillips curve.

4 Conclusion

I find some evidence that the slope of the Phillips curve depends positively on the frequency of price adjustment, although the evidence requires focus on a particular inflation measure for which the relation can arguably be better identified as it yields a better-fitting regression. Consistent with the New Keynesian model, the slope is a convex function of the adjustment frequency. My results suggest that at all times since the late 1970s, the Phillips curve was much flatter than the model-based Phillips curve calibrated at standard values for the structural parameters would imply.

I can only partially explain the discrepancy by specifying the Phillips curve slope as a non-linear function of the frequency of price adjustment akin to that of the New Keynesian Phillips Curve. When I assume an implausibly high degree of competition in the calibration of the slope of the New Keynesian Phillips Curve, this explains most of the difference.

Given that the New Keynesian Phillips curve is derived from Calvo pricing, and state-dependent pricing models imply even steeper Phillips curve slopes, I infer that, when calibrated to micro evidence on the frequency of price adjustment, a broad class of models with nominal rigidities imply steeper output-inflation trade-offs than what is the case in the data. Models with real rigidities[7] constitute a promising line of research aiming to reduce this gap between theory and empirics.

[7] For instance, see Höynck (2020) and Rubbo (2021).

References

[1] Auclert, Adrien, Rodolfo Rigato, Matthew Rognlie, and Ludwig Straub (2021). 'New Pricing Models, Same Old Phillips Curve?', mimeo, Harvard University.

[2] Ball, Laurence, N. Gregory Mankiw, and David Romer. (1988). 'The New Keynesian Economics and the Output-Inflation Trade-Off', *Brookings Papers on Economic Activity*, Spring 1988, pp. 1-82.

[3] Ball, Laurence and Sandeep Mazumder. (2011). 'Inflation Dynamics and the Great Recession', *Brookings Papers on Economic Activity*, Spring 2011, pp. 337-381.

[4] Ball, Laurence and Sandeep Mazumder. (2019a). 'The Nonpuzzling Behavior of Median Inflation', NBER Working Paper No. 25512.

[5] Ball, Laurence and Sandeep Mazumder. (2019b). 'A Phillips Curve with Anchored Expectations and Short-Term Unemployment', *Journal of Money, Credit and Banking*, vol. 51, pp. 111-137.

[6] Bils, Mark, and Peter J. Klenow. (2004). 'Some Evidence on the Importance of Sticky Prices', *Journal of Political Economy*, vol. 112, pp. 947-985.

[7] Blanchard, Olivier. (2016). 'The Phillips Curve: Back to the '60s?', *American Economic Review*, vol. 106, pp. 31-34.

[8] Blanchard, Olivier, Eugenio Cerutti, and Lawrence Summers. (2015). 'Inflation and Activity - Two Explorations and Their Monetary Policy Implications', NBER Working Paper No. 21726.

[9] Calvo, Guillermo A. (1983). 'Staggered Contracts in a Utility-Maximizing Framework', *Journal of Monetary Economics*, vol. 12, pp. 383-398.

[10] Costain, James, Anton Nakov, and Borja Petit. (2021). 'Flattening of the Phillips Curve with State-Dependent Prices and Wages', *Economic Journal*. pp. 1-36.

[11] De Veirman, Emmanuel. (2009). 'What Makes the Output-Inflation Trade-Off Change? The Absence of Accelerating Deflation in Japan', *Journal of Money, Credit and Banking*, vol.41, pp. 1117-1140.

[12] DeFina, Robert H. (1991). 'International Evidence on a New Keynesian Theory of the Output-Inflation Trade-off', *Journal of Money, Credit and Banking*, vol. 23, pp. 410-422.

[13] Del Negro, Marco, Michele Lenza, Giorgio E. Primiceri, and Andrea Tambalotti. (2020). 'What's Up with the Phillips Curve?', *Brookings Papers on Economic Activity*, Spring 2020, pp. 301-373.

[14] Galí, Jordi. (2008). *Monetary Policy, Inflation, and the Business Cycle. An Introduction to the New Keynesian Framework*, Princeton University Press: Princeton and Oxford.

[15] Galí, Jordi and Mark Gertler. (1999). 'Inflation Dynamics: a Structural Econometric Analysis'. *Journal of Monetary Economics*, vol. 44, pp. 195-222.

[16] Golosov, Mikhail, and Robert E. Lucas. (2007). 'Menu Costs and Phillips Curves', *Journal of Political Economy*, vol. 115, pp.171-199.

[17] Höynck, Christian. (2020). 'Production Networks and the Flattening of the Phillips Curve', mimeo, Universitat Pompeu Fabra.

[18] Nakamura, Emi, and Jón Steinsson. (2008). 'Five Facts about Prices: A Reevaluation of Menu Cost Models', *Quarterly Journal of Economics*, vol. 123, pp. 1415-1464.

[19] Nakamura, Emi, and Jón Steinsson. (2010). 'Monetary Non-Neutrality in a Multisector Menu Cost Model', *Quarterly Journal of Economics*, vol. 125, pp. 961-1013.

[20] Nakamura, Emi, Jón Steinsson, Patrick Sun, and Daniel Villar. (2018). 'The Elusive Costs of Inflation: Price Dispersion During the U.S. Great Inflation', *Quarterly Journal of Economics*, vol. 133, pp. 1933-1980.

[21] Okuda, Tatsushi, Tomohiro Tsuruga, and Francesco Zanetti. (2021). 'Imperfect Information, Heterogeneous Demand Shocks, and Inflation Dynamics', mimeo, University of Oxford.

[22] Rubbo, Elisa. (2021). 'Networks, Phillips Curves, and Monetary Policy', mimeo, University of Chicago.

[23] Smets, Frank, and Rafael Wouters. (2007). 'Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach', *American Economic Review*, vol. 97, pp. 586-606.

DeNederlandscheBank
EUROSYSTEEM

De Nederlandsche Bank N.V.
Postbus 98, 1000 AB Amsterdam
020 524 91 11
dnb.nl