

Kiefer, David

Working Paper

Alternative Phillips curves models with endogenous real-time expectations

Working Paper, No. 2010-03

Provided in Cooperation with:

Department of Economics, The University of Utah, Salt Lake City

Suggested Citation: Kiefer, David (2010) : Alternative Phillips curves models with endogenous real-time expectations, Working Paper, No. 2010-03, The University of Utah, Department of Economics, Salt Lake City, UT

This Version is available at:

<https://hdl.handle.net/10419/64442>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DEPARTMENT OF ECONOMICS WORKING PAPER SERIES

**Alternative Phillips Curves Models with Endogenous Real-Time
Expectations**

David Kiefer

Working Paper No: 2010-03

January 2010

University of Utah
Department of Economics
260 S. Central Campus Dr., Rm. 343
Tel: (801) 581-7481
Fax: (801) 585-5649
<http://www.econ.utah.edu>

Alternative Phillips Curves Models with Endogenous Real-Time Expectations

David Kiefer
University of Utah
kiefer@economics.utah.edu

Abstract

Originally presented as an empirical regularity, a variety of microeconomic derivations of the Phillips tradeoff between inflation and real output have been developed. Since these new Phillips curve models are expressed in terms of unobserved variables and expectations, we develop estimates of these unobservables using a state space characterization of the short-run political-economic equilibrium. This method is appropriate because it yields recursive forecasts based on contemporaneous information, and because we apply it to a real-time data set in order to accurately measure available information. Although none of the new Phillips curve tested are completely adequate, we find that Calvo's sticky price formulation provides the best fit for US data. It is inadequate because the estimate coefficient for the driving variable (either the output gap or the marginal cost) is essentially zero.

Keywords: new Phillips curve, microfoundations, real-time data

JEL Classification: E3, E6

1. Introduction

The Phillips curve tradeoff has evolved considerably since it was proposed over 50 years ago. Initially presented as an empirical regularity between the wage inflation and unemployment rates, it has been generalized to account for the related tradeoff between price inflation and the GDP gap. Friedman and Phelps augmented the equation to account for shifts in expectations. Beginning with Fischer (1977) and Taylor (1980), the more recent literature focuses on microeconomic models that could explain this macroeconomic regularity. These models invoke a variety of mechanisms from overlapping nominal wage contracts, to stochastic price resetting, to costly price adjustment, to stochastic updating of information. A direct empirical comparison of these new Phillips models is challenging because they all invoke unobserved variables.

We attempt such a comparison in two steps. Our method coheres with the new Keynesian vision of stabilization policy in which the government leans against the macroeconomic wind. Our first step supposes a quadratic government objective function in inflation-GDP gap space and an old, backward-looking Phillips curve. We use the state space methodology to model the unobserved potential rate of growth as a random walk. Formalizing the relation between observables and unobservables enables recursive Kalman filter estimates of the potential growth rate and other unobservables, conditioned on the observations that were available at each point in time. In the second step we estimate four new Phillips curve models. We find that the sticky price and adjustment cost models give the best fit of the data, while sticky information and overlapping contract models do less well. However, consistent with the literature, none of the new specifications provide an adequate explanation for the inflation-output tradeoff.

2. Microfoundations of the Phillips curve

Phillips' original idea is an inverse relation between wage inflation and the unemployment rate. Textbooks often specify the real side of the equation in terms of the gap between actual and natural unemployment. Frequently the output gap (defined as $x_t \equiv \ln(Y_t) - \ln(Y_t^*)$) where Y_t is real aggregate

output and Y_t^* is potential output) is substituted for the unemployment gap as the measure of macroeconomic disequilibrium. Friedman and Phelps augment the equation by added an adjustment for inflation expectations, so that

$$\pi_t = e_t + \psi x_t + \varepsilon_t, \quad (1)$$

where π_t is the inflation rate, e_t denotes expected inflation and ε_t an exogenous shock. We term (1) the old Phillips curve.

Theorists have developed a variety of explanations for this relation. This section outlines four alternative microfoundation formulations of the new Phillips curve. Even though these derivations can be found in the literature,¹ we begin with this outline in order to modify it below to more accurately model real-time information.

First, Calvo's (1983) stochastic price adjustment model assumes that $(1 - \eta)$ is the probability that a firm can adjust its price in the current quarter; this is called the "sticky price" model. It is assumed that the optimal price for the typical firm varies linearly with the aggregate price and the marginal costs.² The profit-maximizing price for each good is a markup of marginal cost determined by demand elasticity. Furthermore, under certain conditions it can be argued that the deviation from steady-state marginal costs is proportional to the aggregate output gap; thus the optimal price depends on the output gap.³ The parameter $\theta > 0$ specifies the price-gap relation.

Since they may not be able to change their prices for some time, those currently resetting want to forecast future market conditions. Firms average observable conditions with their forecasts, weighted

¹ See for example Froyen and Guender (2007).

² Some authors (for example, Gali (2008)) develop further microfoundations at this point, assuming an economy of monopolistically competitive firms providing a continuum of differentiated consumer goods.

³ There is doubt in the empirical literature about whether the conditions necessary for the cost-gap link hold. Gali and Gertler (1999) report consistent results for a measure of marginal cost, but not for the output gap, while neither variable can explain observed inflation in the Rudd and Whalen (2006) study.

according to the probability that their price will remain fixed in each quarter; accordingly, the profit-maximizing price is specified as⁴

$$p_t^* = (1-\eta) \sum_{\tau=0}^{\infty} \eta^\tau E_t (p_{t+\tau} + \theta x_{t+\tau}) + \mu_t.$$

An exogenous price shock μ is added to account for all other factors affecting the pricing decision. The subscript on the expectations operator gives the date of the forecast. Letting current expectations of current prices be correct, $E_t p_t = p_t$ and $E_t x_t = x_t$. The aggregate price level combines the firms who reset their price in the current quarter with those who set prices in previous quarters according to a geometric distribution. It can be shown that that aggregate inflation is determined as

$$\pi_t = E_t \pi_{t+1} - (1-\eta) E_t \mu_{t+1} + \frac{(1-\eta)^2 \theta}{\eta} x_t + \frac{(1-\eta)}{\eta} \mu_t. \quad (2)$$

An important revision in (2) is that it is forward looking. It includes the forecast of inflation one quarter into the future based on current information, as contrasted to a backward-looking interpretation of the e_t term in (1) as a forecast based past information. Equation (2) also involves a forward-looking forecast of the exogenous shock.

Rotemberg's (1982) model of inflation can be derived by assuming that the typical firm adjusts its price every period to minimize a quadratic function of the costs of changing its price plus the costs of deviation from the optimal price.⁵ The typical firm's expected cost is given by

$$E_t \sum_{\tau=1}^{\infty} \left[(p_\tau - p_\tau^*)^2 + c(p_\tau - p_{\tau-1})^2 \right],$$

⁴ It is appropriate for firms to discount future profits. But since this considerably complicates the result, we follow the much literature by weighting all quarters equally, except for the probability of price resetting. Since we focus on short-run decisions here, a complete neglect of discounting is a reasonable simplification of this model and the other three below. Below we estimate that the average length of price fixity is about 2.5 quarters.

⁵ This can interpreted as equivalent to maximizing a Taylor approximation of a more general profit function.

where p_t is the price chosen and p_t^* is the optimal price in the absence adjustment costs. The firm's optimal price is defined as the observed current price plus a response to the output gap plus a shock, conveniently expressed as

$$p_t^* = p_t + c(\theta x_t + \mu_t)$$

Using the first-order condition that results from minimizing the quadratic cost function with respect to p_t , it can be shown that

$$\pi_t = E_t \pi_{t+1} + \theta x_t + \mu_t. \quad (3)$$

This result is very similar to the sticky price one (2), indistinguishable when future shocks are unpredictable.

Mankiw and Reis (2002) describe their alternative as the “sticky information” model: all firms change their prices in every quarter, but only a fraction uses up-to-date information. Now the term $(1 - \eta)$ defines the probability that a firm sets its price using current information. For the rest, the fraction using information from each earlier date follows a geometric distribution, with a small fraction using a price based on very old information. The aggregation of all pricing decisions is given by

$$p_t = (1 - \eta) \sum_{\tau=0}^{\infty} \eta^\tau E_{t-\tau} p_t^*,$$

where p_t^* denotes the profit-maximizing price, assumed again to respond to the current output gap plus an exogenous price shock as

$$p_t^* = p_t + \theta x_t + \mu_t.$$

This new Phillips curve is forward looking, but backward dated, with a geometric distribution of past expectations of current inflation and output conditions. It can be shown that

$$\pi_t = (1 - \eta) \left\{ \sum_{\tau=0}^{\infty} \eta^\tau E_{t-1-\tau} (\pi_t + \theta(x_t - x_{t-1}) + \mu_t - \mu_{t-1}) + \frac{\theta}{\eta} x_t + \frac{\mu_t}{\eta} \right\}. \quad (4)$$

Finally, Taylor's (1980) overlapping contracts model focuses on wage-setting behavior, although it can also be developed in terms of overlapping price contracts. It is supposed that a fixed nominal labor contract covers several quarters, so that part of the workforce is locked into contracts signed in the past.

This formulation is based on the impact of the macroeconomic conditions on the wage-bargaining outcomes (instead of pricing-setting). When the fixed nominal wage contract w_t^c last two periods, it is supposed that the current contract is set as the average,

$$w_t^c = \frac{p_t + \theta x_t + E_t p_{t+1} + \theta E_t x_{t+1}}{2} + \mu_t.$$

The foundation of this model is that wage bargaining responds to price levels and output gaps (proportional to the parameter $\theta > 0$) and to their expectations, plus an exogenous shock. The two-quarter contract can be rewritten as

$$w_t^c = \frac{2p_t + E_t \pi_{t+1} + \theta x_t + \theta E_t x_{t+1}}{2} + \mu_t.$$

Assuming that the price level is a simple mark up of the wages level, we write the current inflation rate as

$$\pi_t = \frac{1}{2} (E_{t-1} \pi_t + E_t \pi_{t+1} + \theta(x_{t-1} + x_t + E_{t-1} x_t + E_t x_{t+1})) + \mu_{t-1} + \mu_t.$$

More generally, if contracts cover n quarters with $(\frac{1}{n})^{\text{th}}$ of the labor force negotiating a new contract currently, by the same logic the contract wage is

$$w_t^c = \frac{1}{n} \left[np_t + \theta x_t + \sum_{j=1}^{n-1} ((n-j) E_t \pi_{t+j} + \theta E_t x_{t+j}) \right] + \mu_t.$$

Averaging the n overlapping contracts, gives the average wage as

$$w_t = \frac{1}{n} \sum_{k=0}^{n-1} \left\{ \frac{1}{n} \left[np_{t-k} + \theta x_{t-k} + \sum_{j=1}^{n-1} ((n-j) E_{t-k} \pi_{t-k+j} + \theta E_{t-k} x_{t-k+j}) \right] + \mu_{t-k} \right\}.$$

We can write the current inflation rate as

$$\pi_t = \frac{1}{n-1} \left\{ - \sum_{k=1}^{n-2} (n-k-1) \pi_{t-k} + \sum_{k=0}^{n-1} \left[\sum_{j=1}^{n-1} \left(\frac{n-j}{n} E_{t-k} \pi_{t-k+j} + \frac{\theta}{n} E_{t-k} x_{t-k+j} \right) + \frac{\theta}{n} x_{t-k} + \mu_{t-k} \right] \right\} \quad (5)$$

This rather complicated result includes current and past observations along with forecasts. It has a MA(n) error structure.

3. The real-time information available to agents

The derivations above inaccurately model the timing of public information. Since we model the decisions made by workers and firms, it is appropriate to study the information that was available at the time these decisions were taken. The real-time data published by the Philadelphia Federal Reserve Bank accurately reports information availability; it is constructed as dated histories from contemporaneous public information.⁶ These data are quarterly cohorts, each of which is revised over time. The most recent entries in each cohort are the Bureau of Economic Analysis' (BEA) "advance estimates." Since the BEA publishes advance estimates for the $(t-1)^{st}$ quarter in the middle of the t^{th} quarter, public information about current conditions is lagged by one quarter.

Consistent with this sequence, it is appropriate to lag the date of the expectation operator in Calvo's current optimal price model by one quarter,

$$p_t^* = (1-\eta) \sum_{\tau=0}^{\infty} \eta^\tau E_{t-1}(p_{t+\tau} + \theta x_{t+\tau}) + \mu_t;$$

by the same derivation equation (2) is replaced by

$$\pi_t = \frac{(1-2\eta)}{(1-\eta)} E_{t-1}\pi_t + \frac{\eta}{(1-\eta)} E_{t-1}\pi_{t+1} - \eta E_{t-1}\mu_{t+1} + (1-\eta)\theta E_{t-1}x_t + \mu_t. \quad (6)$$

When $E_{t-1}\pi_t = \pi_t$, $E_{t-1}x_t = x_t$, $E_{t-1}\pi_{t+1} = E_t\pi_{t+1}$ and $E_{t-1}\mu_{t+1} = E_t\mu_{t+1}$, (6) reverts to (2). To the extent that Calvo model appeals to theorists because of its functional simplicity, (6) is disappointingly complicated.

When expectations date from the $(t-1)^{st}$ quarter, the Rotemberg's cost adjustment curve (3) becomes

$$E_{t-1}\pi_t = E_{t-1}\pi_{t+1} + \theta E_{t-1}x_t + E_{t-1}\mu_t. \quad (7)$$

On average the typical firm's expectations should be correct, $E_{t-1}\pi_t = \pi_t$, giving a Phillips curve equation.

Likewise, the sticky information curve (4) can also be re-derived to reflect this lag in available information,

⁶ See Croushore and Stark, (2001).

$$\pi_t = (1-\eta) \left\{ E_{t-1} \pi_t + \eta \sum_{\tau=0}^{\infty} \eta^\tau E_{t-2-\tau} (\pi_t + \theta(x_t - x_{t-1}) + \mu_t - \mu_{t-1}) + \theta E_{t-1} x_t + E_{t-1} \mu_t \right\}. \quad (8)$$

And finally, the real-time contract wage for the overlapping-contract is rewritten as

$$w_t^c = \frac{1}{n} \left[n E_{t-1} p_t + \sum_{j=1}^{n-1} (n-j) E_{t-1} \pi_{t+j} + \sum_{j=0}^{n-1} \theta E_{t-1} x_{t+j} \right] + \mu_t,$$

and after some tedious algebra, Taylor's curve (5) becomes

$$\pi_t = \frac{1}{n} \left\{ - \sum_{k=1}^{n-1} (n-k) \pi_{t-k} + \sum_{k=1}^n E_{t-k} \pi_{t-k+1} + \sum_{k=0}^{n-1} \left[\sum_{j=1}^{n-1} \frac{n-j}{n} E_{t-k-1} \pi_{t-k+j} + \sum_{j=0}^{n-1} \frac{\theta}{n} E_{t-k-1} x_{t-k+j} + \mu_{t-k} \right] \right\}. \quad (9)$$

If current conditions are public knowledge, then the new Phillips curves (2)-(5) result, but since real-time expectations date from the previous quarter, the (6)-(9) specifications are more appropriate. The theoretical literature unrealistically dates expectations to the current time period, ignoring the implications of this assumption for model derivation. The econometric literature customarily ignores the measurement error introduced by BEA revisions that may not be published for more than a year.⁷ Many papers have focused on equations with the form of (3), often adding lagged inflation variables to examine the issue of whether this *ad hoc* extension is needed to explain the inflation persistence. We do not estimate such *hybrid* models, concentrating instead on the alternative forms (6) through (9).

4. Endogenous stabilization as a model of expectations

Many empirical studies adopt an instrumental variables approach to modeling expectations about inflation and the output gap. Invoking rationality, they assume that inflation forecasts are on average accurate. But since replacing $E_t \pi_{t+1}$ with π_{t+1} in (3) introduces an endogenous variable on the right hand side, they specify a list of lagged, but usually not real-time, instruments to mitigate endogeneity bias; for example Gali and Gertler (1999).

⁷ An exception is Brissimis and Magginas (2008) who use real-time forecasts of inflation from the Survey of Professional Forecasters and the Philadelphia Fed's real-time GDP. They also pay appropriate attention to publication lags in their choice of instruments.

Conventional measures of the expected output gap also suffer for methodological shortcomings. This variable is usually measured by detrending current, but not real-time, observations of real GDP. Popular detrending procedures use either the CBO's estimate of potential GDP or a Hodrick-Prescott-filtered time series to define the zero point; for example Roberts (2005). Both methods are omniscient in the sense that they are based on future as well as past observations. Rudd and Whalen (2006) pay attention to available information by specifying the detrended gap as one element of a VAR(1), and inferring forecasts of future gaps from the dynamic properties of the estimated VAR(1) system. This has the advantage of incorporating a model of the gap, albeit an agnostic one.⁸ New Phillips curve studies have yielded mixed results, ranging from support by Gali and Gertler to rejection by Rudd and Whalen.

We develop a procedure that endogenizes both inflation and gap expectations, beginning with the assumption that real potential growth, $g_t^* \equiv \ln(Y_t^*) - \ln(Y_{t-1}^*)$, is a random walk,

$$g_t^* = g_{t-1}^* + v_t, \quad (10)$$

This model recognizes that the underlying growth rate of the US economy has changed over time, but that the next turning point is unpredictable. We assume that the potential growth shocks follow the normal distribution $v_t \sim N(0, \sigma_v^2)$, that they are serially independent, $\text{cov}(v_t, v_{t-s}) = 0$, and independent of inflation shocks $\text{cov}(v_t, \varepsilon_t) = 0$. The level of potential GDP and the output gap are defined recursively,

$$\begin{aligned} \ln(Y_t^*) &= \ln(Y_{t-1}^*) + g_t^* \\ x_t &= \ln(Y_t) - \ln(Y_t^*) \end{aligned} \quad (11)$$

We combine this model of the unobservables with a model of the observables based on an assumption of activist stabilization policy. A simple possibility supposes that the government's goals are given by a quadratic function of the output gap and inflation,

$$U_t = -\left(x_t^2 + (\pi_t - \hat{\pi})^2\right), \quad (12)$$

⁸ Their implementation uses final rather than real-time data.

where $\hat{\pi}$ is the inflation target.⁹ Quadratic forms, such as (12), are tractable because they result in linear solutions.¹⁰ Within the quadratic family, a variety of alternatives are plausible. Ours has circular indifference curves, but these can be made elliptical by adding a parameter to reflect the relative weight of inflation versus output goals. Some studies consider parabolic indifference curves.¹¹ Differing targets for inflation could account for ideological differences. Often the output target exceeds zero.¹² Kiefer (2008) estimates several different quadratic forms. He confirms the conventional wisdom that it is not possible to statistically separate goal weights, the inflation target and the output target.¹³ Thus, our inflation target parameter can be interpreted as a composite measure of weights and targets.¹⁴

Although the government has limited options in this model, it exploits information and implementation advantages to lean against the macroeconomic wind. Nevertheless its goals ($x_t = 0$ and $\pi_t = \hat{\pi}$) may be unattainable. To derive the government's policy, we use the old Phillips curve (1) to substitute for x_t and in (12). This asserts that policymakers believe that a backward-looking Phillips curve prevails. Maximizing with respect to π_t , the government's preferred policy (also known as its reaction function) is

$$\pi_t = \frac{e_t + \psi^2 \hat{\pi}}{1 + \psi^2} + \frac{\varepsilon_t}{1 + \psi^2}. \quad (13)$$

Using (1), we find that the preferred output gap is

$$x_t = \frac{\psi(\hat{\pi} - e_t)}{1 + \psi^2} - \frac{\psi\varepsilon_t}{1 + \psi^2}.$$

⁹ The modeling of collective objectives is controversial. Textbooks often define social welfare as an aggregation of individual preferences. Woodford (2003) establishes microfoundations for several close relatives of this function form as an approximation to the utility of a representative consumer. (1999).

¹⁰ Ruge-Murcia (2003) questions the conventional linearity assumption. He develops an alternative where the government's inflation preferences are asymmetrical around its target.

¹¹ See, for example, Alesina et al. (1997).

¹² Barro and Gordon (1983) assume a zero inflation target and an unemployment target below the natural rate.

¹³ Also see Ireland (1999).

¹⁴ Objectives might also include the discounted value of expected future outcomes. The government might plan for its current term of office only, or it might plan to be in office for several terms, discounting the future according to the probability of holding office. Furthermore, it might weigh pre-election years more heavily. Here we assume that only current conditions matter.

Output gap and the growth rate are equivalent measures of stabilization policy because the growth rate can be defined in terms of the output gap as $g_t \equiv \ln(Y_t) - \ln(Y_{t-1}) \equiv x_t - x_{t-1} + g_t^*$. We rewrite the government's output policy in terms of the growth rate as

$$g_t = g_t^* - x_{t-1} + \frac{\psi(\hat{\pi} - e_t)}{1 + \psi^2} - \frac{\psi \varepsilon_t}{1 + \psi^2}. \quad (14)$$

This has the advantage of putting an observable variable on the left-hand-side. Equations (13) and (14) imply that observed inflation and growth depend on shocks, conditions inherited from the past, expectations and policy targets.¹⁵ We assume that the government can implement its policy through fiscal, monetary and other policy instruments, and that the various government agencies (central banks and treasuries) pursue this common policy.¹⁶

Altogether this model of policy is defined in state space as a function of unobserved state variables and lagged observed variables. Our state equations are (10) and (11) and our observation equations are (13) and (14).¹⁷

5. Step one of two

We develop a two-step comparison of the various new Phillips curve models: first we use the model of the previous section to estimate unobserved variables and expectations, and then we use these forecasts to evaluate various new Phillips models. This methodology differs markedly from the literature. Most empirical studies do not model expectations, and customarily the output gap is measured in a deterministic fashion, not as an element of a general macroeconomic equilibrium. Only rarely does the econometric literature study real-time information.

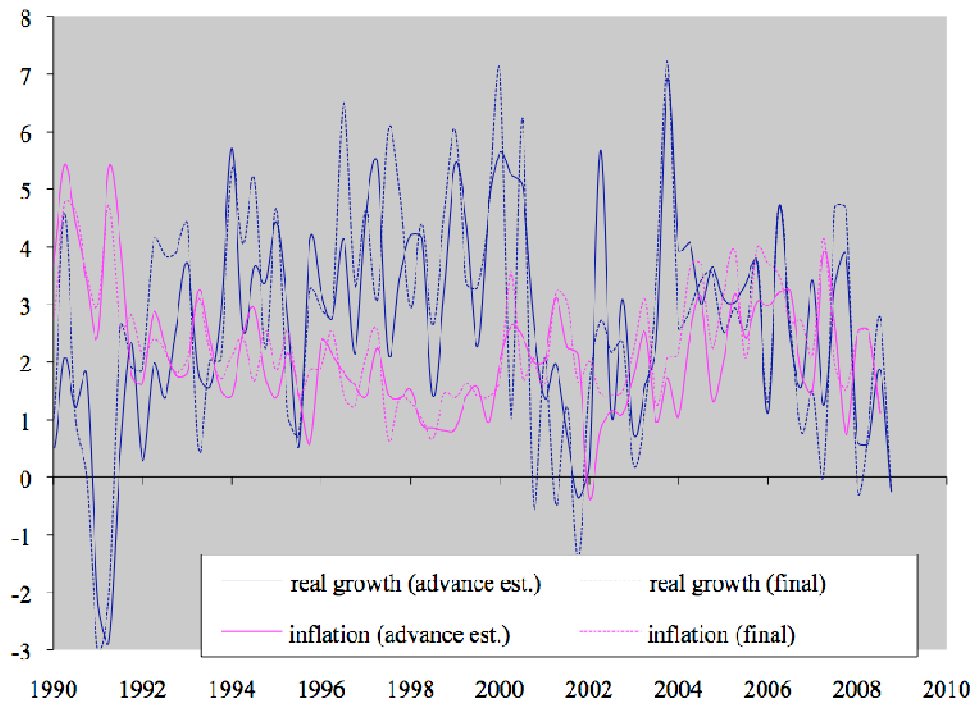
¹⁵ Rational agents come to understand that a policy of $\hat{\pi} > 0$ implies inflation. In the absence of shocks or uncertainty, the time-consistent equilibrium inflation rate should occur where inflation is just high enough so that the government is not tempted to spring a policy surprise. This equilibrium is the potential output, potential growth and a politically determined rate of inflation, $x = 0$, $g = g^*$, $\pi = \hat{\pi}$.

¹⁶ Because our goal is to model inflation and gap expectations, it is not necessary to treat the instrument question. As such our model can be seen to be the first two equations of Carlin and Soskice's (2005) three-equation model, dropping the *IS* equation.

¹⁷ See Hamilton (1994) for a textbook presentation of the Kalman filter methodology.

We take the dependent variable in (13) to be the real-time advance estimate as calculated from the GDP deflator levels for the t^{th} and $(t-1)^{\text{st}}$ quarters, as reported in the middle of the $(t+1)^{\text{st}}$ quarter, labeled π_t . Likewise, we take the dependent variable in (14) to be the same advance estimate of the real growth rate. Figure 1 shows that the BEA's measurement errors can be considerable for both inflation and growth rates, although larger revisions are apparent for growth.¹⁸

Figure 1. Comparing advance estimates with final values over the past two decades



The observation equations (13) and (14) are reduced forms determined by e_t , x_{t-1} , g_t^* and ε_t . These are linear in the variables, but nonlinear in coefficients. For the expectation term e_t , we use a conventional backward-looking assumption,

$$e_t = \frac{1}{T} \sum_{\tau=1}^T \pi_{t-\tau}. \quad (15)$$

¹⁸ We take the final values from the 2008Q3 data cohort.

We refer to this moving average of lagged observations covering the T previous quarters as the $MA(T)$ specification. We select the length of this average empirically. Although many economists view such backward-looking models with suspicion because they lack microfoundations and because their forecasts can be irrational, they are well known to provide a good empirical fit.

We assume that the government has pre-publication knowledge of outcomes. In fact, we assume that the government actively determines inflation and growth (random errors still occur) three months before the official estimates are made public. This timing is consistent the assumption that governments have an information advantage over agents. We could push this advantage assumption further: we could suppose that governments actively sets the “final” values of inflation and growth, statistics that will only be revealed in future BEA revisions. We justify as more realistic our assumption that the government and BEA both make the same errors.

By contrast firms and workers are only aware of current public information, that of the $(t-1)^{st}$ quarter. The backward-looking sum (15) reflects what agents currently know about inflation; it begins with π_{t-1} , the advance estimate released in the middle of the t^{th} quarter. Likewise, we define π_{t-2} as the advance estimate published in the $(t-1)^{st}$ quarter, π_{t-3} from the $(t-2)^{nd}$ quarter and so on. This definition of the e_t term is thus the average of current and out-of-date advance estimates of inflation, which we label *out-of-date* expectations.

An alternative definition of e_t consists of these same inflation rates all taken from the t^{th} quarter publication, which we label *up-to-date* expectations. With this definition only the most recent is an advance estimate, while $\pi_{t-2}, \pi_{t-3}, \dots, \pi_{t-T}$ have undergone updating by the BEA. This specification would be appropriate if all agents reset prices every quarter, while paying attention to the forecasting value of previous inflation rates. The out-of-date specification would be appropriate if prices are sticky in the Calvo sense so that some firms their set prices in the past.

Unfortunately, the estimation of our policy model does not converge. We can obtain convergence by adding an additional error term ξ_t to the growth equation (14), a plausible generalization of our *ad hoc*

approach. Perhaps the addition of a growth shock accounts for policy-implementation errors. Some studies add an inflation policy error to their specifications,¹⁹ although this differs from the growth error that we add here. We take our two shock terms, the original ε_t and the newly added ξ_t , to be exogenous, but not white noise. The inflation and growth shocks could plausibly be correlated, either contemporaneously, serially or both. Serial correlation may reflect the lags inherent in policy implementation emphasized by Carlin and Soskice (2005).

The likelihood statistics in Table 1 compare the goodness of fit of the out-of-date expectations assumption (the left panel) with the up-to-date assumption (the right panel). As a benchmark, we achieve a log likelihood of -695 from a standard VAR(1) model of π_t and g_t using the same sample. All of our *ad hoc* models are more likely to have generated these data. All our models assume VAR(1) errors,

$$\begin{bmatrix} \varepsilon_t \\ \xi_t \end{bmatrix} = \begin{bmatrix} \rho_{\pi\pi} & \rho_{\pi g} \\ \rho_{g\pi} & \rho_{gg} \end{bmatrix} \begin{bmatrix} \varepsilon_{t-1} \\ \xi_{t-1} \end{bmatrix} + \begin{bmatrix} \vartheta_t \\ \zeta_t \end{bmatrix},$$

with $\text{cov}(\vartheta_t, \zeta_t) \neq 0$.²⁰

Table 1. Comparing log likelihood statistics:
ad hoc new Keynesian models, 165 observations, 1967Q3-2008Q3

	out-of-date expectations			up-to-date expectations		
	$\sigma_v^2 = 0.01$	$\sigma_v^2 = 0.04$	$\sigma_v^2 = 0.25$	$\sigma_v^2 = 0.01$	$\sigma_v^2 = 0.04$	$\sigma_v^2 = 0.25$
MA(1)	-674	-675	-679	-674	-675	-679
MA(2)	-676	-677	-681	-676	-677	-681
MA(3)	-667	-669	-673	-666	-668	-672
MA(4)	-663	-665	-669	-661	-663	-667
MA(5)	-666	-668	-672	-664	-665	-670
MA(6)	-667	-668	-673	-666	-667	-672
MA(7)	-671	-673	-677	-670	-672	-676
MA(8)	-670	-672	-676	-668	-670	-674

¹⁹ See for example Ireland (1999).

²⁰ We find that generalizing the specification of ε_t and ξ_t to VAR(2) does not improve the fit of the *ad hoc* model. Its log likelihood statistic is -663, and none of the second-order parameters are statistically significant.

Our model of potential growth is smoothed by imposing an arbitrary variance on its random step. Table 1 includes results for the rather volatile assumption that $\sigma_v^2 = 0.25$ (a standard deviation of 1/2% per quarter), and repeats the analysis with stronger smoothing $\sigma_v^2 = 0.04$ (a standard deviation of 1/5% per quarter) and also $\sigma_v^2 = 0.01$ (a standard deviation of 1/10% per quarter).²¹ Relaxing the fixed σ_v^2 assumption, the estimate of σ_v^2 converges to zero. Although the restriction that $\sigma_v^2 = 0.01$ gives the highest likelihood statistic in the table, it is suspect because the estimation does not succeed in actually imposing this restriction; in this case the estimated variance of σ_v^2 is about 0.04.²²

²¹ By comparison the CBO's estimate to the variance of quarter-to-quarter change in potential real GDP is only 0.008 over this sample period.

²² Perhaps this is feature of the eViews software used to compute these estimates. The table below compares specified variances with realizations measured from the estimates of potential growth. Obviously, the software fails to accurately enforce the specification with small variances.

specified variance	estimated variance
0.0025	0.0299
0.0100	0.0356
0.0400	0.0604
0.0900	0.1041
0.1600	0.1668
0.2500	0.2483

Table 2. Estimation details for selected backward looking, four-quarter moving average, *ad hoc* models, 1967Q3-2008Q3 (z statistics in parentheses)

	(a)	(b)	(c)	(d)	(e)	(f)
expectation assumption	up-to-date	up-to-date	up-to-date	up-to-date	out-of-date	out-of-date
potential growth variance σ_v^2	0.04	0.04	0.01	0.04	0.04	0.04
Phillips curve slope ψ	0.406 (6.120)	0.430 (3.942)	0.419 (3.894)	0.490 (2.774)	0.396 (3.678)	0.467 (2.438)
inflation target $\hat{\pi}$	3.324 (3.364)	3.434 (2.660)	3.420 (2.578)	ave: 3.946 na	3.884 (2.980)	ave: 4.296 na
price shock variance σ_ϑ^2	2.440	1.974	1.946	2.147	1.931	2.127
growth shock variance σ_ζ^2	11.931	6.784	6.802	6.742	6.711	6.667
shock covariance $\text{cov}(\vartheta, \zeta)$	1.868	0.452	0.428	0.522	0.319	0.408
autocorrelation parameter $\rho_{\pi\pi}$		0.202	0.201	0.199	0.228	0.225
autocorrelation parameter $\rho_{\pi g}$		0.144	0.143	0.143	0.126	0.126
autocorrelation parameter $\rho_{g\pi}$		-0.258	-0.256	-0.290	-0.304	-0.335
autocorrelation parameter ρ_{gg}		0.785	0.782	0.785	0.785	0.786
log likelihood	-729	-663	-661	-664	-665	-665

Table 2 reports detailed results for some of the more likely specifications. Model (a) restricts all the serial correlation parameters to be zero. By comparison to the VAR(1) of model (b), serial correlation among the error terms improves the goodness-of-fit substantially, although it does not have much effect of the other parameter estimates.²³ Our estimates of the variance ζ_t of the growth error are large compared to those for ϑ_t ; probably this result explains why we are unable to obtain convergence with inflation errors only. Although very similar, the results favor up-to-date expectations averaged over the past four quarters. Models (b) and (c) report different smoothing assumptions for the up-to-date specification of the expectation term, and model (e) mirrors model (b) for out-of-date expectations. Two additional models, (d) and (f), specify a random walk in the inflation target, and are discussed below. The results are consistent with theory. In all cases the estimated slopes of the Phillips curve are positive and statistically significant and the estimated target variable implies a time-consistent inflation rates of around 4%. Our results are robust to these modeling variations.

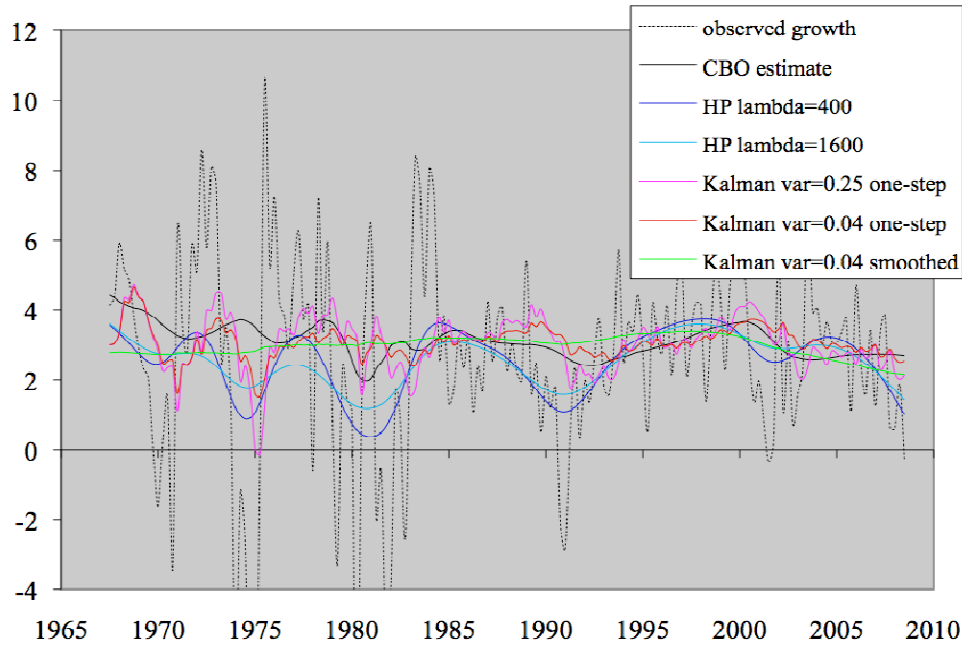
²³ By introducing serial correlation we may have created a problem of endogeneity, to the extent that e_t and x_{t-1} are jointly determined with ε_t . The insensitivity of models (a) and (b) to this specification change suggests that such endogeneity bias is unimportant.

The literature reports other methods of estimating potential growth. Figure 2 compares two of our Kalman filter estimates with two alternatives: the Hodrick-Prescott filter and the estimate published by the Congressional Budget Office (2001). The popularity of the HP filter may be due to its simple agnostic formula.²⁴ The CBO estimate is more complicated, using a growth accounting method inspired by the Solow growth model. This method combines estimates of the trends in the labor force, the capital stock and technological progress. Cyclical components of the labor supply and productivity are removed from observed statistics using the CBO's estimate of the non-accelerating inflation rate of unemployment, constraining potential labor and productivity growth rates to be constant over the business cycle. All methods illustrate the conclusion that the underlying growth rate of the US economy has changed over time. Recently, it shows a peak during the late 1990s followed by a decline over the past decade.

²⁴ It estimates of the potential rate series by minimizing the expression

$$\sum_{t=2}^{T-1} \left\{ (g_t - g_t^*)^2 + \lambda \left[(g_{t+1}^* - g_t^*) + (g_t^* - g_{t-1}^*) \right]^2 \right\},$$

where λ is an arbitrary smoothness parameter that penalizes sharp curves in the g_{it}^* series. It is conventional to set $\lambda=1600$ for quarterly data; as a comparison, Figure 2 also shows the estimate using $\lambda=400$.

Figure 2. Alternatives estimates of the US potential growth rate g_t^* 

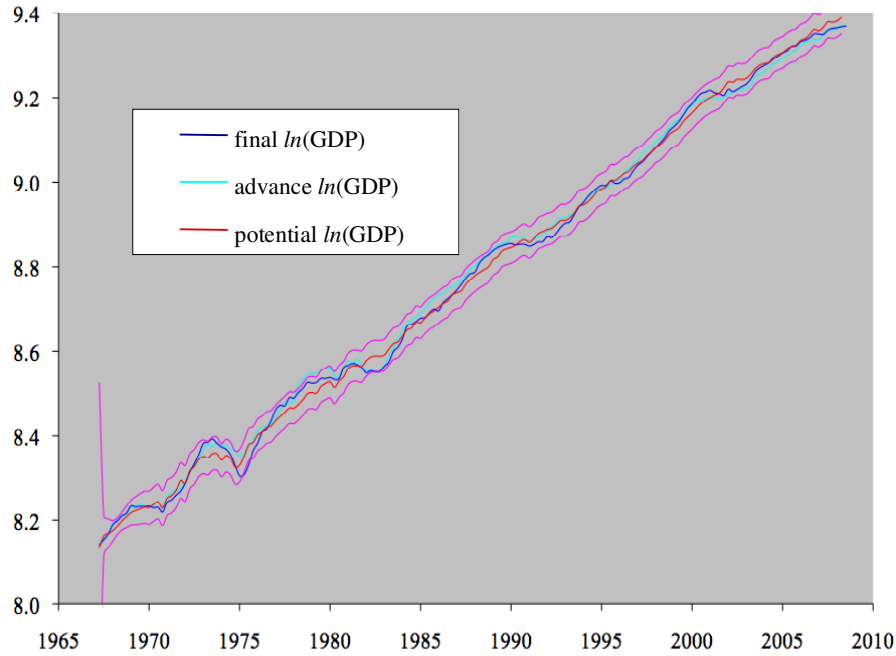
Clearly the HP and CBO estimates are smoother than our model (b) estimate. The difference reflects different assumptions about potential growth as well as different methods of estimation. Conventionally, the potential level changes over time as technology advances, as capital is accumulated and as the labor force grows. Assuming that these influences evolve slowly and independently of business cycles, both the HP filter and CBO estimate impose a gradually evolving process, without large shifts. On the other hand, our generating process is typified by small random shifts that can be occasionally large.

An appealing feature of this state space formulation is that potential growth is a one-step forecast based on only available information. This explains why it becomes smoother and converges with the others as more observations become available.²⁵ On the other hand the HP and CBO methods are omniscient in the sense that they include both past and future observations; they are more comparable to the “smoothed” Kalman prediction of the potential growth conditioned on the entire data set. Figure 2 shows the evident

²⁵ We specify $g_{110}^* = 3$ with a variance of 9 as a plausible prior for the potential growth, and set $\ln(Y_{110}^*)$ equal to the observed value of $\ln(Y)$ in 1967Q3 with a variance of 0.04.

differences between one-step Kalman $E_{t-1}g_t^*$ and smoothed $E_T g_t^*$ forecasts. Although the smoothed estimate is not always close to the HP and CBO estimates, it is obviously less volatile.²⁶ We prefer the one-step forecast as the more appropriate characterization of the information available to decision makers in real time.

Figure 3. Observed real GDP $\ln(Y_t)$ and Kalman predictions of its potential level $E_{t-1}\ln(Y_t^*)$, model (b)

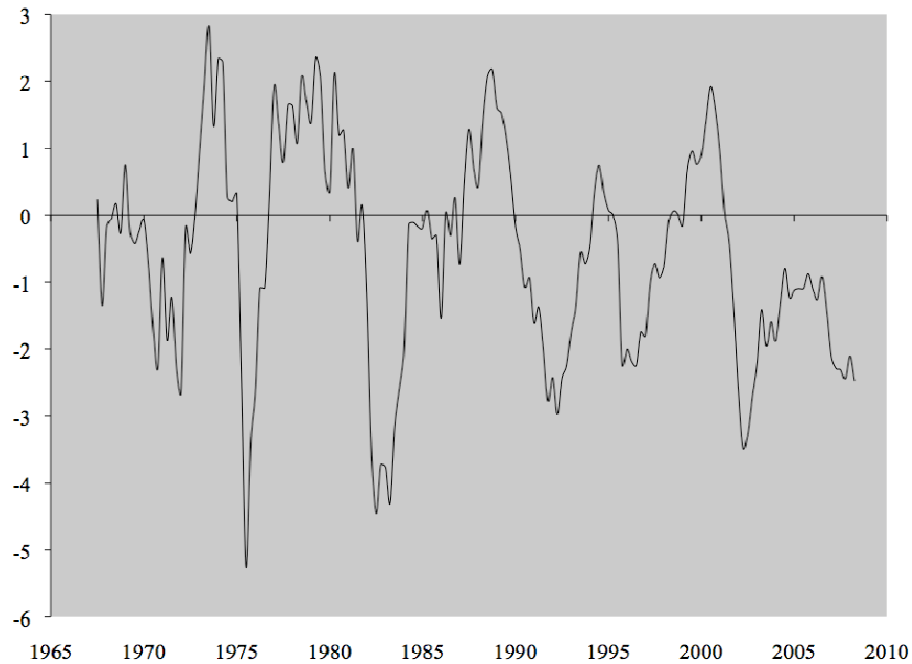


Conditional on current observations and our model specification, the Kalman filter defines recursive forecasts of the unobserved state variables. Figure 3 compares the model (b) estimate of the potential output and its 95% confidence interval with the BEA's advance and final estimates. We prefer (b) because of its goodness-of-fit and its relatively smooth potential growth series. Although Figure 1 shows substantial measurement error in the growth rate series, Figure 3 shows smaller errors in the output series. Usually they are also smaller than our confidence interval for potential output. This plot also indicates how quickly experience comes to dominate our prior mean and variance for potential output.

²⁶ Orphanides and van Norden (2002) study alternative methods of estimating the output gap using these real-time data.

The Kalman methodology also defines forward and backdated forecasts of inflation, the output gap (see Figure 4) and growth. Given the model specification these are Bayesian updates, weighted averages of a previous forecast and the most recent observation. Although we present no evidence that firms, workers or governments learn according to Bayes rule, we interpret these predictions as rational. They estimate what the agents and policymakers might have thought at the time that decisions were taken, conditional on available information.

Figure 4. Real-time GDP gap estimates $E_{t-1}x_t^*$, model (b)



Of course, these forecasts are also conditional by the restrictions, approximations and theory assumed by our model. For example forecasts are conditional on two parameters, the inflation target and the slope of the Phillips curve; we obtain maximum likelihood estimates of these parameters based on the entire sample. A shortcoming of this procedure is that it assumes that agents know ψ and $\hat{\pi}$ with certainty, and that this knowledge is based on the entire sample, not just contemporaneous real-time information. We investigate the seriousness of this limitation by redefining the target as a random coefficient, $\hat{\pi}_{t+1} = \hat{\pi}_t + \omega_t$, where $\omega_t \sim N(0, \sigma_\omega^2)$. This evolving target generalization can be specified by the addition of another state

variable, namely $\hat{\pi}_{t+1}$.²⁷ Our results are plausible: starting from a prior located at $\hat{\pi}_1 = 4$, the Kalman estimate of the target rises to about 8 by 1981, but then declines to around 2 at the end of the sample period. Our finding that model (b) is slightly more likely to have generated the data than the random coefficient model (d), supports the conclusion that the fixed-and-known target assumption is not an important limitation of our method.²⁸

6. An empirical comparison of new and old models

Since the new Phillips theories claim to be an explanation of actual inflation, we use the observed final figure as the dependent variable (taken to be the 2008Q4 data cohort).²⁹ In the case of (9), lagged inflation terms are also measured as final values. Expectation variables are measured as Kalman forecasts according to model (b). We use the up-to-date expectation specification because it fits the data slightly better than the out-of-date definition in Table 1.³⁰ In the case of (8), we approximate the infinite series with an 8-quarter power series, maintaining the requirement that the weights add to unity. Namely, the τ^h stickiness weight is approximated as

$$(1-\eta)\eta^\tau \approx \frac{\eta^\tau}{\sum_{s=0}^7 \eta^s}.$$

Table 3 reports the estimation results for (6)-(9). The estimation period is shortened to 1969Q3-2007Q2 (152 observations) to insure that all the required forecasts are defined and to enhance the “finality” of the 2008Q4 cohort. We take all expectations of shocks, found in (6), (7) and (8), to be zero, despite the fact that we could substitute the Kalman forecast of $E_{t-\tau}\varepsilon_t$ from model (b) for $E_{t-\tau}\mu_t$. The zero expectation

²⁷ We specify $\hat{\pi}_{10} = 4$ with a variance of 9 as a plausible prior. We arbitrarily restrict $\sigma_\omega^2 = 0.09$. Estimation of this random coefficient model cannot be further generalized to include an estimate of σ_ω^2 . When we attempt such a generalization, the model converges to the fixed target model (b).

²⁸ This inference is also justified by a comparison of models (e) and (f) using out-of-date expectations.

²⁹ In unreported regressions we substitute the BEA’s advance estimate of inflation (used to forecast the unobservables in the first step) as the dependent variable in Table 3 to gauge the effect of measurement error in this study. No qualitative changes result.

³⁰ In unreported regressions we substitute the out-of-date procedure; the results differ only slightly from those in Table 3.

is consistent with our independent normal assumption for μ .³¹ Since this assumption removes all shocks from (7) and (8), we add an error term to these models.

Table 3. Regression results for alternative Phillips curve models, 1969Q3 through 2007Q2, the dependent variable is final inflation (z statistics in parentheses)

	(g)	(h)	(i)	(j)	(k)
	sticky prices	price adjustment costs	sticky information	overlapping contracts, 4 quarters	old Phillips curve
Phillips curve slope	-0.051 (-0.445)	-0.005 (-0.099)	0.056 (0.702)	-0.105 (-1.207)	0.225 (4.156)
stickiness parameter	0.604 (19.79)		0.237 (1.319)		
log likelihood	-209	-212	-231	-255	-240

Except for the overlapping contracts model, these specifications all surpass the likelihood of a benchmark forecast. Our benchmark model is $\pi_t = E_{t-1}\pi_t + \mu_t$, where $E_{t-1}\pi_t$ is the one-step Kalman forecast computed from model (b); its log likelihood is -232. The best performing microfoundations model is the sticky prices model (g), with Rotemberg's price adjustment cost formulation (h) a close second. It is a little disappointing that the data do not more clearly distinguish between these two approaches since they imagine such different behaviors: the Calvo model with prices fixed for long periods, and Rotemberg model with prices flexible in every period.

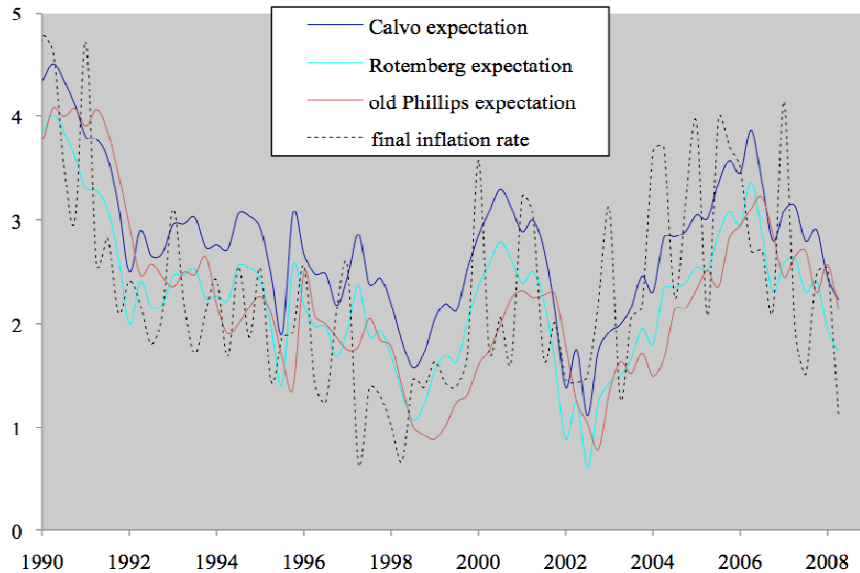
Perhaps our benchmark comparison is inappropriate since model (b) is computed from a two-equation model, while all regressions in Table 3 are single equations. The benchmark may improve the accuracy of its forecasts by applying a theory of government policy and information about the growth rate. For a more comparable benchmark we rewrite (1) to reflect our best-fitting backward-looking expectation and the up-to-date Kalman prediction of the GDP gap (the latter also appears in all four new models),

$$\pi_t = \frac{1}{4}(\pi_{t-1} + \pi_{t-2} + \pi_{t-3} + \pi_{t-4}) + \theta E_{t-1}x_t + \mu_t, \quad (16)$$

³¹ Unreported regressions using the assumption that $E_{t-\tau}\mu_t = E_{t-\tau}\varepsilon_t$ are qualitatively unchanged.

where again the dependent variable is final inflation from the 2008Q4 data cohort, but the four lagged inflations on the right-hand-side are real-time measures. As reported in the last column of Table 3, three of the four new models fit the data better than this old Phillips curve. Figure 5 plots the expectation terms (everything on the right-hand-side except the GDP gap term) from three models. Although not always closer to observed final inflation, sometimes the Calvo term appears to lead the Rotemberg term, while the old Phillips curve term lags both. We interpret this result as evidence that agents' expectations are forward looking.

Figure 5. Comparing expectation term estimates with inflation observations over the past two decades



Comparing the two sticky models, the sticky prices formulation appears much more likely than the sticky information version. In the former the estimated stickiness is plausible; the distributed lag of quarterly weights implied by model (g) are 0.40, 0.24, 0.15, 0.09, 0.05,, for an average length of fixity of about 2.5 quarters.³² The poorer-fitting sticky information model implies less stickiness. Of these new Phillips curve formulations, the overlapping contract model is the weakest fitting specification.³³

³² We note above that the out-of-date expectation assumption does not affect Table 3, and earlier that the out-of-date procedure is more consistent with the sticky price story. It thus noteworthy that we can slightly

7. Substituting marginal costs for the output gap

Most of our new curve estimates of the Phillips curve slope are negative, although statistically insignificant. Negative slopes are problematic because they imply that the short-run equilibrium is unstable; zero slopes are also an issue because they negate the ability of governments to lean against price shocks. Only the old model validates the theoretical expectation.

Other authors also document this puzzle. Gali and Gertler (1999) suggest that wrong sign for the Phillips curves slope is due to labor market imperfections that break the theoretical link between the marginal costs of the typical firm and the aggregate output gap. They report theoretically consistent new Phillips curve estimates after substituting a measure of marginal cost. Their measure is the log deviation of the labor share of GDP from its sample mean; this can be theoretically justified when the typical production function is Cobb-Douglas and when the sample mean approximates “steady-state” marginal costs.

The literature usually measures marginal costs by final statistics, and assumes that its omniscient sample mean measures the steady state; this is a questionable method. Following our model of natural growth, we conceptualize the evolution of the unobserved steady-state marginal costs as a random walk,

$$\ln(C_t^*) = \ln(C_{t-1}^*) + \delta_t. \quad (17)$$

We further assume that the potential cost shocks follow the normal distribution $\delta_t \sim N(0, \sigma_\delta^2)$, that they are serially independent, $\text{cov}(\delta_t, \delta_{t-s}) = 0$, and independent of other shocks: $\text{cov}(\delta_t, v_t) = 0$, $\text{cov}(\delta_t, \xi_t) = 0$ and $\text{cov}(\delta_t, \varepsilon_t) = 0$.³⁴ The real-time observations are modeled as an AR(1) process,

$$\ln(C_t) = \ln(C_{t-1}^*) + \beta c_{t-1} + \varepsilon_t \quad (18)$$

improve the fit of model (g) by assuming out-of-date expectations in step one. The resulting log likelihood increases to -208 with little change in the parameter estimates.

³³ The 4-quarter overlapping contract model is favored by the data, even though its fit is inferior to alternative formulations. In unreported regressions we estimate two-quarter, three-quarter and five-quarter overlapping contract models; all are markedly inferior to the four-quarter model reported in Table 3.

³⁴ Gali and Gertler (1999) assume instead that the equilibrium is the observed sample mean. They do not model the dynamics of marginal costs, assuming that firms are aware of the final version of observed labor share even though this is not yet public knowledge.

where $c_t = \ln(C_t) - \ln(C_t^*)$ is the deviation of costs from trend.

Estimation of (17) and (18) on our sample suggest that the generating process for real-time marginal costs is also well approximated as a random walk,³⁵

$$\ln(C_t) = \ln(C_t^*) + 0.97c_{t-1} \quad (42.3)$$

In an unreported estimate we also find that this labor cost model is not improved by adding the output gap, actually $E_{t-1}x_t$, to the right hand side of (18); thus we do not find any statistically significant relation between the observed labor share and predicted output gap. Perhaps the theoretically predicted link is absent because our measurement of marginal cost is erroneous. Gali and Gertler (1999) measure marginal costs as the final nonfarm labor share. Unfortunately, the nonfarm share is unavailable in the Philly real-time data. But, our use of the available aggregate labor share probably does not introduce much measurement error because the farm component is relatively small.³⁶ Roberts (2005) notes stock option accounting as another source of measurement error in his criticism of this variable. Although there are probably other sources of measurement error, the rejection of the theoretical link is surprising.

Our specification, (17) and (18), has the advantage that it is consistent with our approach because it permits Kalman predictions of expected costs $E_{t-\tau}c_t$, which we substitute for expected output gap terms in model (6) through (9). Table 4 repeats Table 3 with the substitution of expected marginal cost for the expectations of the output gap. Although it does not change the ordering of the five models in terms of likelihood, signs of Phillips curve slope estimates do not change. This finding casts doubt on the argument that the new Phillips curve is empirically supported when it returns to its marginal cost roots.

³⁵ In order to impose stability on the steady-state marginal cost, we restrict $\sigma_\delta^2 = 0.000001$. In an unreported estimate we find that this model differ little from one with its equilibrium fixed at the sample mean. Thus, our method is quite comparable to the literature.

³⁶ Gali, Gertler and Lopez-Salido (2001) also use the aggregate share to measure marginal cost in their extension to European data.

Table 4. Substituting labor cost for the output gap in alternative Phillips curve models
(*z* statistics in parentheses)

	sticky prices	price adjustment costs	sticky information	overlapping contracts, 4 quarters	old Phillips curve
Phillips curve slope	-0.039 (-0.443)	-0.014 (-0.392)	-0.012 (-0.258)	-0.139 (-2.688)	0.024 (0.552)
stickiness parameter	0.602 (19.44)		0.179 (1.008)		
log likelihood	-209	-212	-231	-254	-247

8. Conclusion

This paper develops an econometric comparison of four popular microfoundations models of the Phillips curve. We develop a two-step procedure to estimate unobserved expectations starting with a new Keynesian characterization of the political-economic equilibrium, an admittedly *ad hoc* model involving a old backward-looking Phillips curve, and a simplified characterization of stabilization policy. This first step is estimated on a real-time data set as a measure of the information that workers and firms might use when forming forecasts of inflation and the real economy. The second step uses estimates of the unobservable expectations from the first step to compare the goodness-of-fit of the new Phillips curve equations, adjusted to reflect available public information. The best-fitting new model is Calvo's sticky prices model, closely followed by Rotemberg's cost adjustment model, and the worst is the Taylor model of overlapping contracts. Consistent with the literature, our results are mixed. Generally the new models are more likely to have generated these data than the best-fitting old Phillips curve regression. However, the unexpected slope estimates support Mankiw's (2001) conclusion that the "so-called 'new Keynesian Phillips curve' ... is ultimately a failure."

References

- Alesina, Alberto and Nouriel Roubini with Gerald Cohen, (1997) *Political Cycles and the Macroeconomy*, Cambridge, MA, MIT Press.
- Barro, Robert J. and David B. Gordon, (1983) "A Positive Theory of Monetary Policy in a Natural-Rate Model," *Journal of Political Economy* 91: 598-610.
- Brissimis, Sophocles and Nicholas S. Magginas, (2008) "Inflation Forecasts and the New Keynesian Phillips Curve," *International Journal of Central Banking* 4: 1-22.
- Calvo, Guillermo A., (1983) "Staggered Prices in a Utility Maximizing Framework," *Journal of Monetary Economics* 12: 383-398.
- Carlin, Wendy and David Soskice, (2005) "The 3-Equation New Keynesian Model--a Graphical Exposition," *Contributions to Macroeconomics* 5: Issue 1, Article 13.
- Congressional Budget Office (2001) "CBO's Method For Estimating Potential Output," Washington D.C., Congress of the United States, www.cbo.gov
- Croushore, Dean and Tom Stark, (2001) "A Real-Time Data Set for Macroeconomists," *Journal of Econometrics* 105: 111-30, www.philadelphiafed.org/research-and-data/real-time-center
- Fischer, Stanley (1977) "Long-Term Contracts, Rational Expectations, and the Optimal Money Supply," *Journal of Political Economy* 85, 191-205.
- Froyen, Richard T. and Alfred V. Guender (2007) *Optimal Monetary Policy Under Uncertainty*, Northampton MA, Edward Elgar.
- Gali, Jordi (2008) *Monetary Policy, Inflation and the Business Cycle*, Princeton NJ, Princeton University Press.
- _____ and Mark Gertler (1999) "Inflation dynamics: A structural econometric analysis," *Journal of Monetary Economics* 44: 195-222
- _____, Mark Gertler and J. David Lopez-Salido (2001) "European inflation dynamics," *European Economic Review* 45: 1237-1270
- Hamilton, James D. (1994) *Time Series Analysis*, Princeton NJ, Princeton University Press.

- Kiefer, David (2008) "Revealed Preferences for Macroeconomic Stabilization," *Journal of Applied Economics* 11: 119-143.
- Mankiw, N. Gregory (2001) "The Inexorable and Mysterious Tradeoff between Inflation and Unemployment," *The Economic Journal* 111: C45-C61
- _____ and Ricardo Reis (2002) "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics* 117: 1295-1328
- Orphanides, Athanasios and Simon van Norden (2002) "The Unreliability of Output-Gap Estimates in Real Time," *Review of Economics and Statistics* 84: 569-583
- Roberts, John M. (2005) "How Well Does the New Keynesian Sticky-Price Model Fit the Data?" *Contributions to Macroeconomics* 5: Issue 1, Article 10.
- Rotemberg, Julio J., (1982) "Sticky Prices in the United States," *The Journal of Political Economy* 90: 1187-1211.
- Rudd Jeremy and Karl Whelan (2006) "Can Rational Expectations Sticky-Price Models Explain Inflation Dynamics?" *American Economic Review* 96: 303-320
- Ruge-Murcia, Francisco (2003) "Inflation Targeting under Asymmetric Preferences," *Journal of Money, Credit and Banking*, 35: 763-785
- Taylor, John B., (1980) "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy* 88: 1-22.