Queen's Economics Department Working Paper No. 1155

US Inflation Dynamics 1981-2007: 13,193 Quarterly Observations

Gregor W. Smith
Queen's University

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

2-2008

US Inflation Dynamics 1981–2007: 13,193 Quarterly Observations

Gregor W. Smith†

January 2008

Abstract

The new Keynesian Phillips curve (NKPC) restricts multivariate forecasts. I estimate and test it entirely within a panel of professional forecasts, thus using the time-series, cross-forecaster, and cross-horizon dimensions of the panel. Estimation uses 13,193 observations on quarterly US inflation forecasts since 1981. The main finding is a significantly larger weight on expected future inflation than on past inflation, a finding which also is estimated with much more precision than in the standard approach. Inflation dynamics also are stable over time, with no decline in inflation inertia from the 1980s to the 2000s. But, as in historical data, identifying the output gap is di cult.

JEL classification: E31, E37, C23

Keywords: forecast survey, new Keynesian Phillips curve

†Department of Economics, Queen's University; smithgw@econ.queensu.ca. I thank the Social Sciences Research Council of Canada and the Bank of Canada research fellowship programme for support of this research. The opinions are the author's alone and are not those of the Bank of Canada. I thank James Nason and Henry Siu for directing me to sources and the Department of Economics at UBC for hospitality while this research was undertaken.

1. Introduction

Recent years have seen a boom in statistically estimating inflation dynamics, often represented by the new Keynesian Phillips curve (NKPC). The findings from these exercises then play an important role in larger models (for example including dynamic *IS* curves and policy rules) which then can be used to assess macroeconomic history or design good policy. Yet at the same time the actual persistence and variation in inflation seems to have declined in a number of countries, in part because of monetary policies that target inflation, whether explicitly or implicitly. Identifying and estimating inflation dynamics thus is challenging, as recent research on inflation-forecasting and on weak instruments has shown.

I provide a new may to measure inflation dynamics, by estimating the US NKPC entirely in panels of forecasts. The underlying idea is simple. If there is a stable pattern of persistence to inflation and a stable relationship between inflation and some output-gap indicator, then those links should show up in professional forecasts. If there are forecasters who are unaware of these links, one might even imagine that their resulting loss of accuracy would lead them to revise their forecasting methods or to exit.

The forecasts come from the *Survey of Professional Forecasters* (SPF). The variable being explained is CPI inflation. There are 104 quarterly observations since 1981:3. Combining forecasts from 254 forecasters and 4 combinations of forecast horizons (albeit with many missing observations) yields 13193 observations. Outsourcing the forecasting in this way avoids the search for instrumental variables and greatly increases the precision of the economic findings. Those are: (a) there is a much larger weight on expected future inflation than on past inflation in explaining current inflation (a contrast with some previous research); (b) it is di cult to find a stable e ect of the unemployment rate on aggregate inflation; and (c) parameters are stable over time. In particular, there is no evidence that inflation inertia is lower in the 2000s than it was in the 1980s.

Section 2 outlines the NKPC and provides references to derivations and inter-

pretations. Section 3 introduces the statistical method and compares it to standard approaches that use instrumental variables estimation or involve the median forecast from a survey. Section 4 describes the SPF data. Section 5 gives the empirical findings. They are presented for various time periods and horizons and also disaggregated by forecaster. Section 6 compares the findings to those from traditional approaches with historical data. Section 7 interprets the lack of significance of the unemployment rate, and assesses how it may a ect confidence in the other NKPC coe cients. Section 8 summarizes the findings.

2. Economic Context

A range of pricing environments with frictions give rise to a hybrid NKPC that describes inflation, π_t, like this:

$$\pi_t = \lambda_b \pi_{t-1} + \lambda_f \mathrm{E}_t \pi_{t+1} + \gamma x_t, \qquad (1)$$

where x_t denotes real aggregate demand (either real marginal cost or an output gap or minus the unemployment rate). The studies by Roberts (1997), Fuhrer and Moore (1995), and Galí and Gertler (1999) contain examples of these environments. The dynamics can reflect smooth adjustment with quadratic costs (as introduced by Rotemberg, 1982), a variation of Calvo's pricing model (with or without firm-specific capital) in which some price-setters are backward-looking, or a model in which inflation (rather than the price level) is sticky. This form also is consistent with the dynamic indexing model studied by Christiano, Eichenbaum, and Evans (2005) or sticky-information models like those of Devereux and Yetman (2003) and Mankiw and Reis (2002) or with environments that include real rigidities like that of Blanchard and Galí (2005). Dennis (2007) and Woodford (2007) provide up-to-date reviews and assessments. Because the reduced form (1) may be consistent with various pricing or information schemes I focus on its parameters and specifically on single-equation estimation and testing.

Here are four statements about empirical evidence on the NKPC. These four observations concern the economic context, while the next section of the paper looks at

the econometric methods. In each case I give only a few citations, though many more are possible.

2.1 There is ongoing debate about the relative sizes of the parameters on lagged inflation and on expected future inflation (λ_b and λ_f respectively). For example Fuhrer (1997) and Rudd and Whelan (2005) find that lagged inflation dominates, while Galí and Gertler (1999), Galí, Gertler, and Lopez-Salido (2005), and Sbordone (2005) find that expected future inflation dominates. Moreover, a large value for λ_b (also known as inflation inertia) seems inconsistent with a number of studies of price-setting frequency using data from individual firms.

2.2 The mixture of backward-looking and forward-looking dynamics matters to the design of optimal policy, whether under discretion or under commitment. The new Keynesian Phillips curve generally is included in a new Keynesian model in order to study optimal policy. In that context, Clarida, Galí, and Gertler (1999) note that:

> With inertia present, adjustments in current monetary policy affect the future time path of inflation. As consequence, policy now responds not only to current inflation but also to forecasts of inflation into the indefinite future. [p 1692]

They also show that inertia may lead optimal policy to involve a more aggressive response to shocks that affect inflation, for if it is not stopped, its persistence leads to ongoing price-adjustment costs. Woodford (2003a, chapter 8.3.2; 2003b) shows that, if the inertia in inflation stems from the underlying pricing model, then the degree of inertia can affect all of (a) the optimal target, (b) the optimal response to shocks, and (c) how to implement commitment. Walsh (2003, section 11.3.7) gives a numerical example of how the optimal policy changes with λ_b.

Even if there is no inflation inertia, so $\lambda_b = 0$, the value of λ_f generally affects optimal policy. For example, in the Calvo pricing model this parameter depends on the underlying fraction of firms that can adjust their prices each period. Schmitt-Grohé and Uribe (2007) construct a quantitative general equilibrium model with such pricing and find that the Ramsey policy is very sensitive to the value of this fraction.

2.3 Inflation inertia may not be structural. Under some interpretations of the NKPC

the coe cient on lagged inflation may vary with the policy rule or the inflation environment. Thus, the implications of the inertia for welfare and optimal policy may depend on its source, for example whether is stems from indexation or from information. Woodford (2007) gives several examples in which the inertia is not structural and should not be taken into account in designing policy. But since the inflation environment changed from the 1980s to the 2000s (as described by Nason, 2006, for example) it is thus interesting to check for stability of the NKPC parameters across these time periods.

2.4 There is ongoing debate about the way to measure real aggregate demand x_t and about its significance. Candidates for x_t include real marginal cost (as represented by the labor share of income), an output gap, or the unemployment rate. In new Keynesian models, the output gap is defined as the di erence between output and its value without nominal rigidities. It remains an open question how closely traditional measures such as detrended GDP coincide with this concept, though existing research does not suggest a very high correlation between the two. The measurement of x_t of course may a ect conclusions about λ_b and λ_f. But there is no clear pattern across studies on this connection. For example, Galí and Gertler (1999) find a large role for expected future inflation, while Nason and Smith (2006) — who also use marginal cost — find little evidence of forward-looking dynamics.

Neiss and Nelson (2005) conclude that there is more evidence of a role for marginal cost than for output gaps as traditionally measured. They replace detrended output with a theoretically-determined model of the output gap and find more evidence for its role in influencing inflation. Blanchard and Galí (2005) also show that the conventionally measured output gap may not be consistent with the theoretically correct one. In both papers the inappropriateness of a traditional output gap measure arises because of real rigidities in the form of sticky wages. (Blanchard and Galí also show that real rigidities increase inflation persistence, as measured by λ_b.)

From the perspective of this paper, a key finding of Blanchard and Galí is that — with real rigidities (sticky real wages) and staggered price-setting — there are two

relevant x_t variables: the unemployment rate and the change in the price of non-produced inputs (like oil prices for example). As they note, this leads to a Phillips curve with traditional exogenous variables. They estimate this version of the Phillips curve in US annual data and find a negative, though imprecisely estimated, $\hat{\gamma}$.

Orphanides and Williams (2002, 2005) estimate a new Keynesian Phillips curve that identifies x_t as di erence between the unemployment rate and various measures of the natural rate. Their specification thus resembles a traditional Phillips curve too, except for the inclusion of $E_t \pi_{t+1}$. And they too find a significant, negative coe cient in US quarterly data.

I use only SPF data in this paper. That data source does not include forecasts for the output gap or for a wage series needed to construct the labor share of income (a standard measure of marginal cost). But it does include forecasts for the unemployment rate, so I estimate NKPCs with that as the measure of x_t. I also investigate including both the current and lagged unemployment rates, as implied by the price-setting model of Fuhrer and Moore (1995) that also is studied by Roberts (1997). Though measures of supply shocks (as emphasized by Blanchard and Galí) and a time-varying natural rate of unemployment (as emphasized by Orphanides and Williams) are not available in a forecast panel, at a minimum this exercise demonstrates the method and compares it to standard estimation with the same variables. Conditional on this traditional measure of real aggregate demand, I measure the mix of backward and forward-looking inflation dynamics and also see whether they are estimated with greater precision than in standard methods.

3. Statistical Methods

This study focuses on single-equation or limited-information estimation of the NKPC. The advantage of this approach of course is that its findings apply regardless of the characteristics or parameter values in the rest of the economic model. And its disadvantage is that identification may be more di cult and statistical e ciency less than in a systems approach.

The relationship studied includes a constant term, to give:

$$\pi_t = \lambda_0 + \lambda_b \pi_{t-1} + \lambda_f \mathrm{E}_t \pi_{t+1} + \gamma u_t. \qquad (2)$$

Because of the endogeneity of $\mathrm{E}_t \pi_{t+1}$ the estimating equations are the sample versions of:

$$\mathrm{E}\left[(\pi_t - \lambda_0 - \lambda_b \pi_{t-1} - \lambda_f \pi_{t+1} - \gamma u_t) \cdot z_t\right] = 0. \qquad (3)$$

in which z_t is a list of instruments that lie within the information set used by market participants.

Instruments must be (a) uncorrelated with the residual and (b) relevant in predicting π_{t+1}. To satisfy criterion (a), researchers often have used only lagged instruments, z_{t-1}. The idea is that this step may give consistent estimates even if there is a correlation between a shock or residual and the unemployment rate, u_t, in other words if the unemployment rate also is endogenous. Lagging instruments also may provide a consistent estimator if there is some measurement error or if price-setters are missing some current information. Galí and Gertler (1999), Galí, Gertler, and López-Salido (2005), Neiss and Nelson (2005), and Jondeau and Le Bihan (2005) all use only lagged instruments.

Finding instruments is difficult. To be relevant, an instrument dated $t - 1$ must help predict π_{t+1} in part independently of all of π_t, π_{t-1}, and u_t. Stock and Watson (1999) reported that few variables have power to forecast postwar U.S. inflation once lagged inflation and the unemployment rate are accounted for, while Orphanides and van Norden (2005) and Stock and Watson (2007) illustrate and explain the ongoing challenges of inflation forecasting. This difficulty in predicting π_{t+1} suggests that estimating and testing with the NKPC may be subject to the effects of weak identification, a syndrome under which the instrumental variables estimator is biased towards OLS, its distribution is non-normal, and standard confidence intervals can be misleading. Ma (2002), Mavroeidis (2005), Dufour, Khalaf, and Kichian (2006), and Nason and Smith (2007) all reach this conclusion about the NKPC. In fact, Nason and Smith show analytically that there will be no valid instruments available in data generated from

the three-equation, new Keynesian model (with persistent shocks). The logic is that as one lags instruments enough for them to be uncorrelated with the residual they also become irrelevant to predicting future inflation. Dufour, Khalaf, and Kichian (2006) and Nason and Smith (2007) provide tests of the NKPC that are robust to weak identification, but valid confidence intervals constructed using these robust test statistics can still be wide.

One way to avoid the problem of weak instruments is to represent $E_t \pi_{t+1}$ with the median forecast from a survey. Ang, Bekaert, and Wei (2007) concluded that this series is the best predictor of annual US inflation. They ran a tournament among forecast models that included survey measures, time-series models, models with real-side variables, and arbitrage-free models of the term structure. Their main conclusion is that the median professional forecast (from the Livingston survey or SPF) is the best predictor of annual inflation. They also allowed for forecast combination or pooling, using least-squares and other methods. They found that little weight was attached to any other candidate besides the median professional forecast in their pooling exercises.

In the NKPC, Roberts (1995) pioneered the use of the survey median in estimation. Orphanides and Williams (2002, 2005), Adam and Padula (2003), and Zhang, Osborn, and Kim (2006) use measures such as the median SPF forecast or the Greenbook forecasts of the Federal Reserve Board. For example, Adam and Padula (2003) used the mean SPF forecast for the US. They found that either unit labour costs or detrended output is significant in the NKPC when this survey measure of expected inflation is adopted.

Smith (2007) describes how to combine the median or mean forecast with other sources of information such as the actual, realized series or forecasts from individual forecasters. But obviously these cannot all be included in the single-equation estimation without exhausting degrees of freedom. Thus an objection to the use of the median forecast is that it does not use all information in the cross-section of forecasters. A second objection is that the series of median forecasts does not represent the

expectations of any specific forecaster, and so may not have some of the properties one would expect of an individual's forecasts.

This study uses the complete cross-section of forecasts in the panel. This approach is feasible because the NKPC is linear in observed forecasts. Once one follows the existing literature and uses z_{t-1}, it is clear that the variables on *both* sides of the NKPC are being forecasted. I simply replace these with professional forecasts, with one equation for each forecaster. They all predict the same group of variables (and one can restrict the estimation to use the same NKPC parameters) but they have different information sets, so their forecasts need not be equal. They do the work of finding instruments.

This method was introduced by Smith and Yetman (2007) who applied it to the CCAPM. The idea is that if a given structure holds in the economy (and the law of iterated expectations applies) then it should be reflected in professional forecasts. To my knowledge there are no previous uses of panels of multivariate professional forecasts alone to estimate parameters of economic models, though Rudin (1992) measured the implicit views of forecasters on the univariate properties of output.

Let j index forecasters, numbered from 1 to J. Then let $\mathrm{E}_{jt-1}x_t$ denote a one-step-ahead prediction by forecaster j for variable x_t. The estimating equations now are:

$$\mathrm{E}_{jt-1}\pi_t = \lambda_0 + \lambda_b \mathrm{E}_{jt-1}\pi_{t-1} + \lambda_f \mathrm{E}_{jt-1}\pi_{t+1} + \gamma \mathrm{E}_{jt-1}u_t. \qquad (4)$$

Each of the J equations links forecasts made by the same forecaster at the same time. But forecasters in the SPF also make predictions for the same variables but at various horizons, here indexed by h. For any variable, say π, the forecast of the value at time t by forecaster j, h quarters in advance is denoted $\mathrm{E}_{jt-h}\pi_t$. Thus even more information can be brought to bear, using restrictions across horizons to give the full set of estimating equations:

$$\mathrm{E}_{jt-h}\pi_t = \lambda_0 + \lambda_b \mathrm{E}_{jt-h}\pi_{t-h} + \lambda_f \mathrm{E}_{jt-h}\pi_{t+1} + \gamma \mathrm{E}_{jt-h}u_t. \qquad (5)$$

The parameters $\{\lambda_0, \lambda_b, \lambda_f, \gamma\}$ can be estimated by ordinary least squares in a panel

of at most $J \times H \times T$ observations. In practice the panel is much smaller than this because of missing observations, but still involves many more than T observations.

This approach has two other advantages over the traditional methods. First, it uses real-time data. One need not worry that expectations are implicitly being modelled using data not available at the time actual forecasts were being made. Second, the availability of additional cross-forecaster and cross-horizon data means that estimation and testing can take place in shorter time series. That is useful in the case of inflation where policy rules may have changed over time and where one may want to test for stability of the NKPC parameters under di erent regimes.

It is an open question whether using these data provides greater precision and aids identification. But it is promising that the disagreement among inflation forecasters has been documented by Mankiw, Reis, and Wolfers (2004) and Capistran and Timmermann (2006). This estimation method takes advantage of that heterogeneity in forecasting methods or information (and hence in the survey data) to help identify the parameters and estimate them precisely.

4. SPF Forecast Data

The panel data come from the *Survey of Professional Forecasters* (SPF) conducted by the Federal Reserve Bank of Philadelphia (`www.philadelphiafed.org/econ/spf/`). I study forecasts for two variables (listed with their SPF codes in brackets): π, the CPI inflation rate, quarter-to-quarter, seasonally adjusted, at annual rates, in percentage points (`cpi`); and u, the quarterly average unemployment rate, seasonally adjusted, in percentage points (`unemp`). (Surveys of predicted inflation measured with the PCE deflator began only in 2007).

The data are quarterly, and run from 1981:3 to 2007:3. Although the SPF continued the ASA-NBER survey that began in 1968, CPI inflation was not included in the survey until the third quarter of 1981. Quarters, indexed by t, run from 1 to $T = 104$. Forecast horizons also are quarterly. Forecasts are reported for the previous quarter, the current quarter, and the following four quarters. Horizons are indexed by h,

which counts from 0 (applicable to the previous quarter) to $H = 5$.

The survey uses a cross-section of forecasters, indexed by j which runs from 1 to J. For these variables the SPF includes reports from $J = 254$ forecasters. The survey documentation notes that there may be cases in which a forecaster left the panel and that forecaster's identification number was later assigned to a di erent forecaster. And some judgement is used in deciding whether to associate identification numbers to firms or to individuals (who may switch firms). These factors do not a ect the pooled findings in this study though they may a ect some of the evidence for individual forecasters.

The number of observations is much less than $J \times H \times T$ because forecasters leave and join the panel and do not make forecasts for all time periods and horizons. In practice, the total number of observations, or jht combinations, is 13,193. This total is 118 times greater than the number of quarterly time series observations.

5. Empirical Evidence

Table 1 presents estimates of the parameters of the Phillips curve (5). The first row gives OLS estimates, followed by OLS standard errors and then by heteroskedasticity-consistent standard errors. The coe cient on lagged inflation is $\hat{\lambda}_b = 0.286$ while that on future inflation is $\hat{\lambda} = 0.672$. Both are estimated with great precision. However, $\hat{\gamma}$, the coe cient on the unemployment rate, is positive, although statistically insignificant at conventional levels of significance. The final column shows that $\overline{R}^2 = 0.82$. In addition, tests for residual autocorrelation find none of significance, so that the dynamics of the right-hand-side forecasts match those of the left-hand-side one. The overall conclusion, then, is that inflation forecasts are quite well 'explained' by forecasts for prior and subsequent periods.

The price-setting model of Fuhrer and Moore (1995) implies that both u_t and u_{t-1} should appear on the right-hand side of the estimating equations. To check on this specification in the forecast panel, I included $\mathrm{E}_{jt-h}u_{t-1}$ as an additional regressor and also estimated the equations with the average of the current and lagged unemploy-

ment rates. The results (not shown) were very similar to those in table 1. Neither unemployment rate nor their average was statistically significant at conventional levels.

The next set of estimates in table 1 is found by weighted least squares, with the weight on each forecaster in proportion to the square root of the number of reports to the survey questions made by that forecaster in the panel. And the final rows show what happens when forecaster-specific intercepts, λ_{0j}, are included in the panel estimation. Neither modification a ects the conclusions.

Table 2 shows results disaggregated by horizon. The first row repeats the result from table 1 as a benchmark. Again the main conclusions do not change, with one exception. When I estimate with only one-step ahead forecasts $(h = 1)$ and include forecaster-specific intercepts, I find a significant, negative $\hat{\gamma}$. In all cases $\hat{\lambda}_f > 0.5 > \hat{\lambda}_b$, though the values do vary by horizon.

Table 3 shows results for various time periods. The main motivation for this exercise was to see whether the estimates are stable over time. But in addition there seemed to be some outlier reports for inflation forecasts for the 1980s and also for 2006-2007, so those time periods are omitted from some of the samples. The first row again repeats table 1, for the entire 1980:1–2007:3 sample, as a benchmark. Later rows present results for the entire sample but omitting 2006 and 2007 and then results for the 1980s, 1990s, and 2000s. Again the main conclusions do not change.

Zhang, Osborn, and Kim (2006) estimate the output-gap version of the NKPC using both GMM estimation and OLS estimation with the median SPF forecast. They present evidence of instability in the parameters, with λ_f higher and λ_b lower after 1981. Roberts (2006) focuses on the conditional correlation between the unemployment rate and the inflation rate and finds that it falls over time. He shows that this change can arise with a stable, underlying NKPC and a change in the monetary policy rule. With panel data beginning in 1981, I cannot o er new evidence on stability with that as a break date. But table 3 shows that there is remarkable stability in λ_b and λ_f for this forecast-based estimation since 1981. An advantage of using the cross-forecaster and

cross-horizon dimensions of the panel is that one can test for stability with a relatively short time period. With this method there are 3810 observations with which to study inflation dynamics only in the 2000s.

It is notable that the estimated inflation inertia has not fallen, despite the change in the inflation environment from the 1980s to the 2000s. As Woodford (2007) explains, whether this inertia changed or not should shed light on various explanations for its presence. But its stability here suggests that perhaps it should be accounted for in designing policy.

Finally, I also disaggregate by forecaster, to find whether individual forecasters have implicit Phillips curves of this form. Only forecasters who made predictions for at least 20 quarters are included. For each such forecaster j I estimate:

$$\mathrm{E}_{jt-h}\pi_t = \lambda_{0j} + \lambda_{bj}\mathrm{E}_{jt-h}\pi_{t-h} + \lambda_{fj}\mathrm{E}_{jt-h}\pi_{t+1} + \gamma_j\mathrm{E}_{jt-h}u_t. \qquad (6)$$

Now all the parameters have a j subscript. Figure 1 contains the results, in the form of box plots for each of $\hat{\lambda}_{bj}$, $\hat{\lambda}_{fj}$, $\hat{\gamma}_j$, and $\overline{R}^2_j$.

For $\hat{\lambda}_{bj}$, the coe cient on the previous period's inflation, the median is 0.226. The second quartile (the box) begins at 0.148 and the third quartile ends at 0.297. The whiskers (the box plus and minus 1.5 times the interquartile range, *i.e.* the height of the box) span from 0.057 to 0.346.

For $\hat{\lambda}_{fj}$, the coe cient on future inflation, the median is 0.657, the box runs from 0.487 to 0.769 and the whiskers span 0.240 to 0.931. This second box lies entirely above the box and whiskers for $\hat{\lambda}_{bj}$. A typical value of $\hat{\lambda}_{fj}$ is roughly twice as large as a typical value of $\hat{\lambda}_{bj}$.

For $\hat{\gamma}_j$, the coe cient on the unemployment-rate forecast, the third box plot has a median of -0.014, consistent with a downward-sloping Phillips curve. But both the box, $\{-0.112, 0.028\}$, and the whiskers, $\{-0.359, 0.134\}$, include 0.

The fourth box plot, for $\overline{R}^2_j$, has a median of 0.65. Again this sends the same message as in the pooled estimation. Even though there is no strong correlation between forecasts for the unemployment rate and those for the CPI inflation rate, the

changes in earlier and later inflation alone capture a significant part of the variation in inflation forecasts.

One might wonder whether forecasters who place a large weight on future inflation also place a small weight on past inflation *i.e.* whether the result in the coe cient medians or box plots that $\hat{\lambda}_f > \hat{\lambda}_b$ holds forecaster-by-forecaster. In addition, there might be a relationship across forecasters between the weight on past inflation and the response to the unemployment rate. Figure 2 contains the scatter plot of $\hat{\lambda}_{fj}$ against $\hat{\lambda}_{bj}$ (the dark circles) as well as that of $\hat{\gamma}_j$ against $\hat{\lambda}_{bj}$ (the light circles). No correlations are apparent. But notice that the great majority of the $\hat{\lambda}_{fj}$ lie above the 45 degree line. Thus the result that the weight on future inflation exceeds the weight on past inflation does apply to almost all individual forecasters.

6. Historical Data

The forecast-based estimates can be put in context by comparing them with traditional estimates based on historical data. I estimated this equation:

$$\pi_t = \lambda_0 + \lambda_b \pi_{t-1} + \lambda_f \mathrm{E}_t \pi_{t+1} + \gamma u_t, \qquad (7)$$

for the same time period, 1981:3–2007:3. I used currently available data, though an alternative might be to use the real-time data of Croushore and Stark (2001). The inflation rate is the annualized, quarter-to-quarter growth rate in the CPI (all items) seasonally adjusted, averaged from monthly data: `cpiaucsl` from FRED. The unemployment rate is the civilian rate from the BLS, series `unrate` from FRED. These are the outcomes corresponding to the definitions forecasted in the SPF.

I replace the unobservable $\mathrm{E}_t \pi_{t+1}$ with π_{t+1}, then estimate by instrumental variables. Table 4 shows the sets of instruments considered, and presents the results. The first two rows use variables z_t known at time t. The third row uses only lagged instruments, z_{t-1}, which section 2 noted is a common method. In the fourth row the instrument set includes π_{t+1} so that instrumental-variables regression becomes ordinary least squares. It is well-known that, with weak instruments, estimates con-

verge to the inconsistent OLS values as instruments are added; this line is included to capture that syndrome.

A difference between the estimates from the forecast panel (in table 1) and those from the historical data (in table 4) is that the value of $\hat{\lambda}_b$, the estimated weight on lagged inflation, or degree of inflation inertia, is much larger in the traditional approach. As section 2 noted, conclusions about this parameter affect optimal monetary policy. The values of $\hat{\lambda}_b$ and $\hat{\lambda}_f$ are roughly similar in the historical data, whatever the instrument set. That is a contrast with the forecast-based estimates.

But this contrast vanishes when one looks at the standard errors from traditional, instrumental-variables estimation in table 4. They are large enough that one could not reject the hypothesis that the true parameters are those measured with the forecast data in table 1. So the real difference between the two approaches lies in the precision. Overall, the standard errors in the traditional approach are roughly ten times larger than those in the forecast approach.

The coefficient on the unemployment rate, $\hat{\gamma}$, is imprecisely estimated in each case. That is similar to the outcome in the forecast data. These two similarities — (a) an insignificant effect of the unemployment rate and (b) lead and lag coefficients that are not significantly different from those in table 1 — support the idea that the forecasts-only approach can add precision to estimates we usually find using realized data alone. In keeping with the argument of this paper, and following from the law of iterated expectations, estimation with forecast data may measure the same things as in realized data (but with greater precision) rather than simply describing the partial correlation between the forecasts and nothing more.

The last row of table 4 contains results from estimation by OLS but replacing $E_t\pi_{t+1}$ by the median, one-quarter-ahead inflation forecast from the SPF (series `CPI2`) from `MedianLevel.xls`). Estimation again uses data for 1981:3 to 2007:3, just as in the forecast-only exercise in the previous section. Here the results are very different from both the forecast-based estimates and from the IVE estimates. The coefficient $\hat{\lambda}_b$ is negative, the coefficient $\hat{\lambda}_f$ is greater than 1, and the coefficient $\hat{\gamma}$ is negative. The

parameters are estimated with greater precision and the goodness of fit is greater than with the historical data alone. But the goodness of fit is still less than with the forecast data, even when — as in the first row of table 1 — no cross-sectional heterogeneity is allowed for in the panel estimation.

Next, compare the point estimates from estimation with the median inflation forecast and historical data (the last row of table 4) by visualizing them in the distribution of forecaster-specific coe cients (in figure 1). This comparison shows that all three coe cients from estimation with median forecasts are very di erent from the corresponding medians of the forecaster-specific estimates. Estimation with the median is not like the median of estimations.

7. Interpretation

The idea underlying the NKPC is that inflation tracks the x_t variable. Given this idea, how can we interpret the fact that λ_b and λ_f (describing the inflation dynamics) are well-identified yet γ (describing the impact of the unemployment rate) is not? From the benchmark estimates in table 1 (with heteroskedasticity-consistent standard errors) the t-statistics for the three variables are 26, 52, and 1.2 respectively. Figure 3 shows the box plots of forecaster-specific t-statistics for $\hat{\lambda}_{bj}$, $\hat{\lambda}_{fj}$, and $\hat{\gamma}_j$. These plots show that this discrepancy in precision arises in part from the time-series dimension of the panel. The median values of the forecaster-specific t-statistics are 2.29, 4.83, and -0.24. Forecaster-by-forecaster there is more evidence of roles for lagged and future inflation than there is for the unemployment rate. By definition the remainder of the di erence in precision across variables in the pooled estimation of table 1 comes from the cross-sectional dimension of the data. Forecasters who predict high inflation for periods $t-1$ and $t+1$ tend to be forecasters who predict high inflation for period t. But for a typical quarter and horizon there is not much correlation between a forecasting entity's inflation forecast and its unemployment-rate forecast.

I also examined three other candidates from the SPF for the role of x_t, the variable driving inflation. These were (a) the forecasted quarter-to-quarter growth rate of real

GDP, constructed from the levels forecasts RGDP; (b) the probability of a recession in the forecasted quarter, RECESS; and (c) an approximation to the forecasted labour share of output (in turn approximating marginal cost), that uses forecasts of nominal GDP and corporate profits as follows: log(NGDP- CPROF) - log NGDP. I included each of these in the forecast Phillips curve in turn. The coe cient on real GDP growth was negative and the coe cient on the recession probability was positive; neither outcome is what one would expect of a Phillips curve. The coe cient on approximate marginal cost was positive, as one would expect. But none of these alternate measures was statistically significant. In addition, whatever the proxy variable for x_t, I always found a much larger coe cient on future inflation than on past inflation: $\hat{\lambda}_f > 0.5 > \hat{\lambda}_b$.

How can the statistical method precisely identify the pattern of tracking without precisely identifying the object being tracked? The simplest explanation for the small coe cients on unemployment-rate forecasts is a bias towards zero in $\hat{\gamma}$ caused by measurement error: u_t does not accurately measure the output gap or marginal cost that drives inflation. This section investigates what this interpretation would then mean for (a) the statistical significance of the measured e ect $\hat{\gamma}$, and (b) the estimated inflation dynamics.

Standard analysis of the errors-in-variables problem shows that the least-squares estimator of γ is inconsistent and that $\hat{\gamma}$ will be biased towards zero. This is the usual attenuation e ect of measurement error. A less well-known consequence of measurement error is that it also biases down the associated t-statistic. Meijer and Wansbeek (2000) prove analytically that there is attenuation in the t-statistic too, when the associated regressor includes measurement error in the linear regression model.

But there are two problems with applying the standard, measurement-error story here. First, the coe cient $\hat{\gamma}$ is not just small; it also has the theoretically wrong sign in the pooled estimation and for a number of horizons and forecasters. Second, it takes some work to think of a stochastic structure of measurement error that would be present both in the actual data (used in table 4) and in the forecasts (used in tables 1-3). In particular, there must be an error at each horizon h.

What does this explanation for the insignificance of the unemployment rate forecast, $E_{jt-h}u_t$, mean for the interpretation of $\{\hat{\lambda}_b, \hat{\lambda}_f\}$? Here the ordinary-least-squares estimators will be unbiased and consistent if the associated variables, $E_{jt-h}\pi_{t-h}$ and $E_{jt-h}\pi_{t+1}$, are uncorrelated with the mis-measured variable $E_{jt-h}x_t$. But they will be biased and inconsistent if there is a correlation. This time the bias is upwards, if the correlation is positive. The appendix provides a proof.

An extreme case provides intuition into this conclusion. Suppose that the variance of the measurement error is so large that the correlation between $-E_{jt-h}u_t$ and $E_{jt-h}x_t$ falls to zero. This interpretation is probably the simplest one to apply to the results in tables 1-4. At that point the output-gap measure has been omitted from the OLS regression. Thus $\{\hat{\lambda}_b, \hat{\lambda}_f\}$ reflect omitted-variables bias. If the forecasts of lagged and future inflation are positively correlated with the omitted variable then their coe cients are biased upwards, as they partly capture the e ect of the missing $E_{jt-h}x_t$ on $E_{jt-h}\pi_t$. In this case the coe cients still measure the relationship between the inflation forecasts and the unemployment-rate forecasts. It is interesting to note from tables 1-3 that there is little connection between the two. But the coe cients then cannot be thought of as consistent estimates of the NKPC parameters. At the other extreme, though, if the forecasts of future and lagged inflation, on the one hand, and the output gap, on the other, are uncorrelated, then mis-measuring the output gap will not a ect the estimator of the inflation dynamics.

Not having data on $E_{jt-h}x_t$, I cannot tell what bias may be present in $\{\hat{\lambda}_b, \hat{\lambda}_f\}$. One way to shed light on this issue would be to simulate a general equilibrium model that includes all three variables, and then fit the estimating equations (3) to the simulated data. Another, traditional response to measurement error is to estimate by instrumental variables. In fact, the use of lagged instruments z_{t-1} in GMM estimation of the NKPC sometimes is explained on these grounds. But table 4 shows that $\hat{\gamma}$ remains insignificant (and positive) when estimated by IVE in the historical data. And it seems unlikely in practice that the measurement error is serially uncorrelated.

8. Summary

It is appealing to try to estimate the parameters of the new Keynesian Phillips curve without necessarily nesting it in a complete economic model. The new method in this paper estimates them using only a panel of professional forecasts. This approach follows the logic of GMM estimation of the NKPC, which typically uses lagged instruments. In that case the econometrician implicitly forecasts all of the variables while estimating the parameters. Another standard approach uses the median forecast of future inflation from a survey panel. By forecasting all variables — or rather drawing on the work of professional forecasters who have done so — I can use the cross-forecaster dimension of the panel and so take advantage of many more observations than are available in either of these traditional approaches.

The main findings are (a) a relatively large role for expected future inflation (summarized by $\hat{\lambda}_f$); (b) a relatively small role for lagged inflation ($\hat{\lambda}_b$) also known as inflation inertia; (c) much greater precision in these estimates than from standard methods; (d) an insignificant role for the unemployment rate, standing in for the measure of real aggregate demand that drives inflation; and (e) evidence of inflation inertia that, while smaller than in some previous studies, is stable over time and for the majority of individual forecasters. The estimated role of lagged inflation does not decline with the fall in the average US inflation rate from 1981 to 2007.

I compare this evidence with findings using GMM and historical data. There too there is little evidence of a role for the unemployment rate, and the estimated mix of backward and forward dynamics in inflation is insignificantly di erent from the one found with forecast data. But this approach gives much less precision (larger standard errors).

The simplest explanation for the insignificance of the unemployment-rate forecasts is that they are unrelated to the output gap or marginal cost variable that price-setters track (a limiting case of measurement error). But that explanation means that there may also be bias in the estimates $\{\hat{\lambda}_b, \hat{\lambda}_f\}$. If the associated inflation forecasts are positively correlated with the output gap, then these coe cients will be biased up.

Drawing on a general-equilibrium model in which inflation, the output gap, and the unemployment rate all are endogenous would be a useful next step. In such an environment, one could ensure that the NKPC holds and use the Monte Carlo method to see what an investigator would find when including the unemployment rate in estimating the dynamic Phillips curve.

Appendix: Output Gap Measurement Error and Estimated Inflation Dynamics

Forecasts follow a regression:

$$\mathrm{E}_{jt-h}\pi_t = \lambda_b \mathrm{E}_{jt-h}\pi_{t-h} + \lambda_f \mathrm{E}_{jt-h}\pi_{t+1} + \gamma \mathrm{E}_{jt-h}x_t, \qquad (A1)$$

with $\lambda_b, \lambda_f, \gamma > 0$. I ignore the intercept λ_0 with no effect on the conclusion. To keep notation simple, write $\pi \equiv \mathrm{E}_{jt-h}\pi_t$ and $x \equiv \mathrm{E}_{jt-h}x_t$. Collect the forecasts of lagged and future inflation in a vector $v \equiv (\mathrm{E}_{jt-h}\pi_{t-h}\ \mathrm{E}_{jt-h}\pi_{t+1})$ and the corresponding parameters in a vector $\lambda \equiv (\lambda_b\ \lambda_f)$. Then rewrite (A1) as:

$$\pi = \lambda v + \gamma x. \qquad (A2)$$

Again solely to keep notation as simple as possible, I examine the case in which v is a scalar. Let the data vector $(\pi\ v\ x)$ have population variance-covariance matrix:

$$\begin{bmatrix} \sigma_\pi^2 & \sigma_{\pi v} & \sigma_{\pi x} \\ \sigma_{\pi v} & \sigma_v^2 & \sigma_{vx} \\ \sigma_{\pi x} & \sigma_{vx} & \sigma_x^2 \end{bmatrix}. \qquad (A3)$$

For future reference, note that the population value of the first regression coefficient in (A2) is:

$$\gamma = \frac{\sigma_{\pi x}\sigma_v^2 - \sigma_{\pi v}\sigma_{vx}}{\sigma_v^2\sigma_x^2 - (\sigma_{vx})^2}. \qquad (A4)$$

The investigator observes unemployment-rate forecasts, denoted u, that coincide with the output-gap forecasts, x, with classical measurement error,

$$u = x + \eta, \qquad (A5)$$

where the measurement error η has mean zero, variance σ_η^2, and covariance 0 with x and v. Thus

$$\sigma_u^2 = \sigma_x^2 + \sigma_\eta^2. \qquad (A6)$$

The investigator runs this linear projection:

$$\pi = b_\lambda v + b_\gamma u. \qquad (A7)$$

The population value of the second least-squares coefficient is:

$$\begin{aligned} \text{plim}\ \hat{b}_\gamma &= \frac{\sigma_{\pi u}\sigma_v^2 - \sigma_{\pi v}\sigma_{vu}}{\sigma_v^2\sigma_u^2 - (\sigma_{vu})^2} \\ &= \frac{\sigma_{\pi x}\sigma_v^2 - \sigma_{\pi v}\sigma_{vx}}{\sigma_v^2(\sigma_x^2 + \sigma_\eta^2) - (\sigma_{vx})^2}. \end{aligned} \qquad (A8)$$

The measurement-error a ects only the denominator, which yields the attenuation bias. The population value of the coe cient on the variable v that is measured without error is:

$$\begin{aligned} \text{plim}\ \hat{b}_\lambda &= \frac{\sigma_{\pi v}\sigma_u^2 - \sigma_{\pi u}\sigma_{vu}}{\sigma_v^2\sigma_u^2 - (\sigma_{vu})^2} \\ &= \frac{\sigma_{\pi v}(\sigma_x^2 + \sigma_\eta^2) - \sigma_{\pi x}\sigma_{vx}}{\sigma_v^2(\sigma_x^2 + \sigma_\eta^2) - (\sigma_{vx})^2}. \end{aligned} \tag{A9}$$

In this case the variance of the measurement error a ects both the numerator and the denominator. Straightforward di erentiation gives:

$$\frac{\partial \text{plim}\ \hat{b}_\lambda}{\partial \sigma_\eta^2} = \frac{\sigma_{vx}(\sigma_{\pi_x}\sigma_v^2 - \sigma_{\pi v}\sigma_{vx})}{[\sigma_v^2(\sigma_x^2 + \sigma_\eta^2) - (\sigma_{vx})^2]^2}. \tag{A10}$$

The denominator is positive. The term in brackets in the numerator is also the numerator of the population expression for γ (A4), which is positive because $\gamma > 0$, so

$$sgn\ \frac{\partial \text{plim}\ \hat{b}_\lambda}{\partial \sigma_\eta^2} = sgn\ \sigma_{vx}. \tag{A11}$$

Thus plim $\hat{b}_\lambda > \lambda$ (this coe cient is biased up) if $\sigma_{vx} > 0$. Measurement error in the output gap causes upward bias in the other coe cients when the other variables are positively correlated with the output gap. Notice also from (A9) that:

$$\lim_{\sigma_\eta^2 \to \infty} \text{plim}\ \hat{b}_\lambda = \frac{\sigma_{\pi v}}{\sigma_v^2}. \tag{A12}$$

As discussed in the text, if the variance of η becomes large enough then x is an omitted variable, and π is projected on v.

References

Adam, Klaus, and Mario Padula (2003) Inflation dynamics and subjective expectations in the United States. European Central Bank, Working Paper No. 222.

Ang, Andrew, Geert Bekaert and Min Wei (2007) Do macro variables, asset markets or surveys forecast inflation better? *Journal of Monetary Economics* 54, 1163-1212.

Blanchard, Olivier and Jordi Galí (2005) Real wage rigidities and the new Keynesian model. *Journal of Money, Credit and Banking* 39 (Supplement), 35-65.

Capistran, Carlos and Allan Timmermann (2006) Disagreement and biases in inflation expectations. mimeo, Department of Economics, University of California San Diego.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005) Nominal rigidities and the dynamic e ects of a shock to monetary policy. *Journal of Political Economy* 113, 1-45.

Clarida, Richard, Jordi Galí, and Mark Gertler (1999) The science of monetary policy: A new Keynesian perspective. *Journal of Economic Literature* 37, 1661–1707.

Croushore, Dean, and Tom Stark (2001) A real-time data set for macroeconomists. *Journal of Econometrics* 105, 111–130.

Dennis, Richard (2007) Fixing the new Keynesian Phillips curve. *Federal Reserve Bank of San Francisco Economic Letter* Number 2007-35, November 30.

Devereux, Michael B. and James Yetman (2003) Predetermined prices and the persistent e ects of money on output, *Journal of Money, Credit and Banking* 35, 729-41.

Dufour, Jean-Marie, Lynda Khalaf, and Maral Kichian (2006) Inflation dynamics and the New Keynesian Phillips Curve: An identification-robust econometric analysis, *Journal of Economic Dynamics and Control* 30, 1707-1727.

Fuhrer, Je rey C. (1997) The (un)importance of forward-looking behavior in price specifications. *Journal of Money, Credit and Banking* 29, 338-50.

Fuhrer, Je rey C. and George R. Moore (1995) Inflation persistence. *Quarterly Journal of Economics* 110, 127-160.

Galí, Jordi and Gertler, Mark (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics* 44, 195-222.

Galí, Jordi, Gertler, Mark, and López-Salido, J. David (2005). Robustness of the estimates of the hybrid New Keynesian Phillips curve. *Journal of Monetary Economics* 52, 1107-1118.

Jondeau, Eric and Hervé Le Bihan (2005) Testing for the new Keynesian Phillips curve. Additional international evidence. *Economic Modelling* 22, 521– 550

Ma, Adrian (2002). GMM estimation of the new Phillips curve. *Economics Letters* 76, 411-417.

Mankiw, N. Gregory 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.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers (2004) Disagreement about inflation expectations. *NBER Macroeconomics Annual 2003* 209-248.

Mavroeidis, Sophocles (2005). Identification issues in forward-looking models estimated by GMM, with an application to the Phillips curve. *Journal of Money, Credit, and Banking* 37, 421-448.

Meijer, Erik and Tom Wansbeek (2000) Measurement error in a single regressor. *Economics Letters* 69, 277–284

Nason, James M. (2006) Instability in U.S. inflation: 1967-2005. Federal Reserve Bank of Atlanta *Economic Review* 91(Q2), 39-59.

Nason, James M. and Gregor W. Smith (2006) Identifying the new Keynesian Phillips curve. forthcoming in the *Journal of Applied Econometrics*

Neiss, Katharine and Nelson, Edward (2005) Inflation dynamics, marginal cost, and the output gap: Evidence from three countries. *Journal of Money, Credit, and Banking* 37, 1019-45.

Orphanides, Athanasios and Simon van Norden (2005) The reliability of inflation forecasts based on output gap estimates in real time. *Journal of Money, Credit, and Banking* 37, 583-601.

Orphanides, Athanasios and John C. Williams (2002) Robust monetary policy rules with unknown natural rates. *Brookings Papers on Economic Activity* 2:2002, 63-118.

Orphanides, Athanasios and John C. Williams (2005) The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations. *Journal of Economic Dynamics and Control* 29, 1927-1950.

Roberts, John M. (1995). New Keynesian economics and the Phillips curve. *Journal of Money, Credit, and Banking* 27, 975-984.

Roberts, John M. (1997). Is inflation sticky? *Journal of Monetary Economics* 39, 173-196.

Roberts, John M. (2006) Monetary policy and inflation dynamics. *International Journal of Central Banking*, 2, 193-230.

Rotemberg, Julio J. (1982) Monopolistic price adjustment and aggregate output. *Review of Economic Studies* 49, 517-31.

Rudd, Jeremy, and Karl Whelan (2005) New tests of the new Keynesian Phillips curve. *Journal of Monetary Economics* 52, 1167-1181.

Rudin, Jeremy R. (1992) What do private agents believe about the time series properties of GNP? *Canadian Journal of Economics* 25, 369-391.

Sbordone Argia M. (2005) Do expected marginal costs drive inflation dynamics? *Journal of Monetary Economics* 52, 1183-1197.

Schmitt-Grohé, Stephanie and Martin Uribe (2007) Optimal inflation stabilization in a medium-scale macroeconomic model, in *Monetary Policy Under Inflation Targeting*, edited by Klaus Schmidt-Hebbel and Frederick Mishkin, Central Bank of Chile, Santiago, Chile, 125-186.

Smith, Gregor W. (2007) Pooling forecasts in linear rational expectations models. Queen's Economics Department working paper 1129.

Smith, Gregor W. and James Yetman (2007) The curse of Irving Fisher (professional forecasters' version). Queen's Economics Department working paper 1144.

Stock, James and Watson, Mark (1999). Forecasting inflation. *Journal of Monetary Economics* 44, 293-335.

Stock, James, and Mark Watson (2007) Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking* 39, 3-34.

Walsh, Carl E. (2003) *Monetary Theory and Policy*. MIT Press.

Woodford, Michael (2003a) *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

Woodford, Michael (2003b) Discussion 2, pp 393-402 in *Price Adjustment and Monetary Policy*, Bank of Canada.

Woodford, Michael (2007) Interpreting inflation persistence: Comments on the conference on "Quantitative evidence on price determination." *Journal of Money, Credit and Banking* 39 (Supplement), 203-210.

Zhang, Chengsi, Denise Osborn, and Dong Keon Kim (2006) The new Keynesian Phillips curve: From sticky inflation to sticky prices. forthcoming in the *Journal of Money, Credit and Banking*.

Table 1: Forecast Phillips Curve Estimation

$$E_{jt-h}\pi_t = \lambda_0 + \lambda_b E_{jt-h}\pi_{t-1} + \lambda_f E_{jt-h}\pi_{t+1} + \gamma E_{jt-h} u_t$$

Estimator	$\hat{\lambda}_b$ *(se)*	$\hat{\lambda}_f$ *(se)*	$\hat{\gamma}$ *(se)*	Obs.	$\overline{R}^2$
OLS	0.286	0.672	0.0089	13193	0.82
	(0.005)	(0.006)	(0.0047)		
HC	(0.011)	(0.013)	(0.0074)		
WLS	0.279	0.677	0.007	13193	0.88
	(0.005)	(0.006)	(0.0045)		
λ_{0j}	0.279	0.661	-0.0048	13193	0.82
	(0.011)	(0.014)	(0.007)		

Notes: $\{h,j,t\}$ index horizon, forecaster, and time period. Obs. is the number of observations. HC denotes heteroskedasticity-consistent standard errors. WLS denotes weighted least squares based on observations per forecaster. λ_{0j} refers to estimation with forecaster-specific intercepts. Estimation uses SPF data from 1981:3 to 2007:3.

Table 2: Results by Horizon

$$E_{jt-h}\pi_t = \lambda_0 + \lambda_b E_{jt-h}\pi_{t-1} + \lambda_f E_{jt-h}\pi_{t+1} + \gamma E_{jt-h} u_t$$

Horizons	$\hat{\lambda}_b$ *(se)*	$\hat{\lambda}_f$ *(se)*	$\hat{\gamma}$ *(se)*	Obs.	$\overline{R}^2$
$h = \{1, 2, 3, 4\}$	0.286 (0.011)	0.672 (0.013)	0.0089 (0.0074)	13193	0.82
$h = 1$	0.283 (0.17)	0.679 (0.26)	-0.028 (0.19)	3301	0.64
$h = 2$	0.213 (0.017)	0.730 (0.022)	0.023 (0.013)	3329	0.84
$h = 3$	0.344 (0.025)	0.625 (0.028)	0.021 (0.010)	3318	0.92
$h = 4$	0.463 (0.034)	0.512 (0.034)	0.010 (0.0089)	3245	0.93
$h = \{2, 3, 4\}$	0.281 (0.014)	0.673 (0.015)	0.023 (0.007)	9892	0.89
λ_{0j}; $h = 1$	0.266 (0.016)	0.698 (0.028)	-0.043 (0.019)	3301	0.64
λ_{0j}; $h = 2$	0.214 (0.016)	0.702 (0.022)	0.00026 (0.010)	3329	0.85
λ_{0j}; $h = 3$	0.346 (0.022)	0.603 (0.026)	0.014 (0.008)	3318	0.92
λ_{0j}; $h = 4$	0.449 (0.034)	0.520 (0.033)	0.010 (0.009)	3245	0.94

Notes: $\{h,j,t\}$ index horizon, forecaster, and time period. Obs. is the number of observations. Standard errors are heteroskedasticity-consistent. λ_{0j} refers to estimation with forecaster-specific intercepts. Estimation uses SPF data from 1981:3 to 2007:3.

Table 3: Results by Time Period

$$E_{jt-h}\pi_t = \lambda_0 + \lambda_b E_{jt-h}\pi_{t-1} + \lambda_f E_{jt-h}\pi_{t+1} + \gamma E_{jt-h} u_t$$

Time Period	$\hat{\lambda}_b$ *(se)*	$\hat{\lambda}_f$ *(se)*	$\hat{\gamma}$ *(se)*	Obs.	$\overline{R}^2$
1980:1–2007:3	0.286 (0.011)	0.672 (0.013)	0.0089 (0.0074)	13193	0.82
1980:1–2005:4	0.309 (0.011)	0.656 (0.012)	0.003 (0.0078)	12424	0.82
1980:1–1989:4	0.292 (0.017)	0.674 (0.017)	-0.090 (0.013)	3974	0.75
1990:1–1999:4	0.257 (0.015)	0.645 (0.020)	0.009 (0.009)	5409	0.66
2000:1–2007:3	0.280 (0.018)	0.667 (0.023)	0.022 (0.013)	3810	0.82

Notes: *{h,j,t}* index horizon, forecaster, and time period. Obs. is the number of observations. Standard errors are heteroskedasticity-consistent. Estimation uses SPF data.

Figure 1: Forecaster-Specific Coefficients

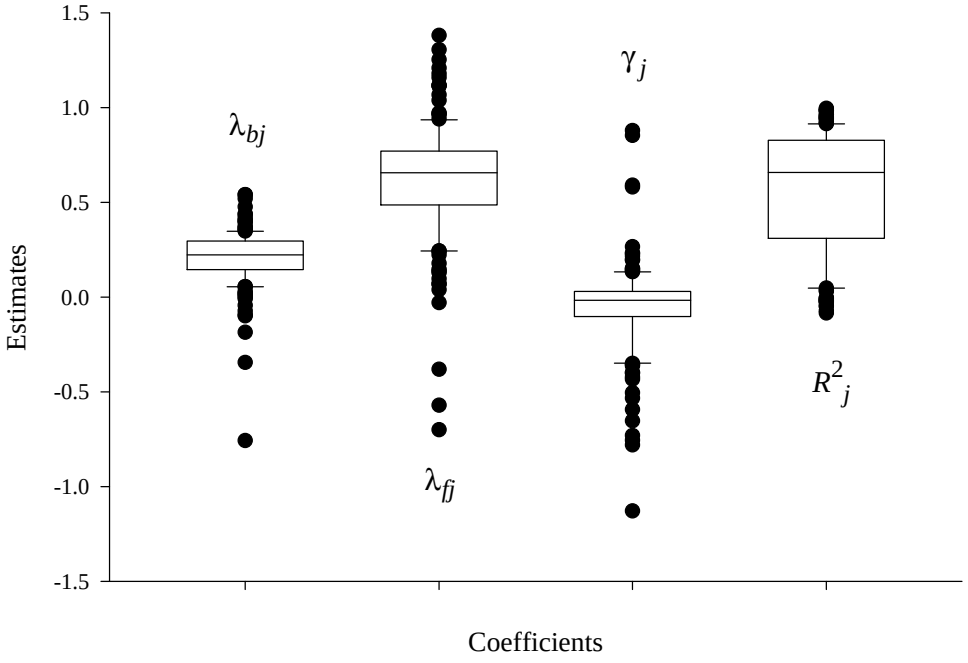

Figure 2: Coefficient Correlations

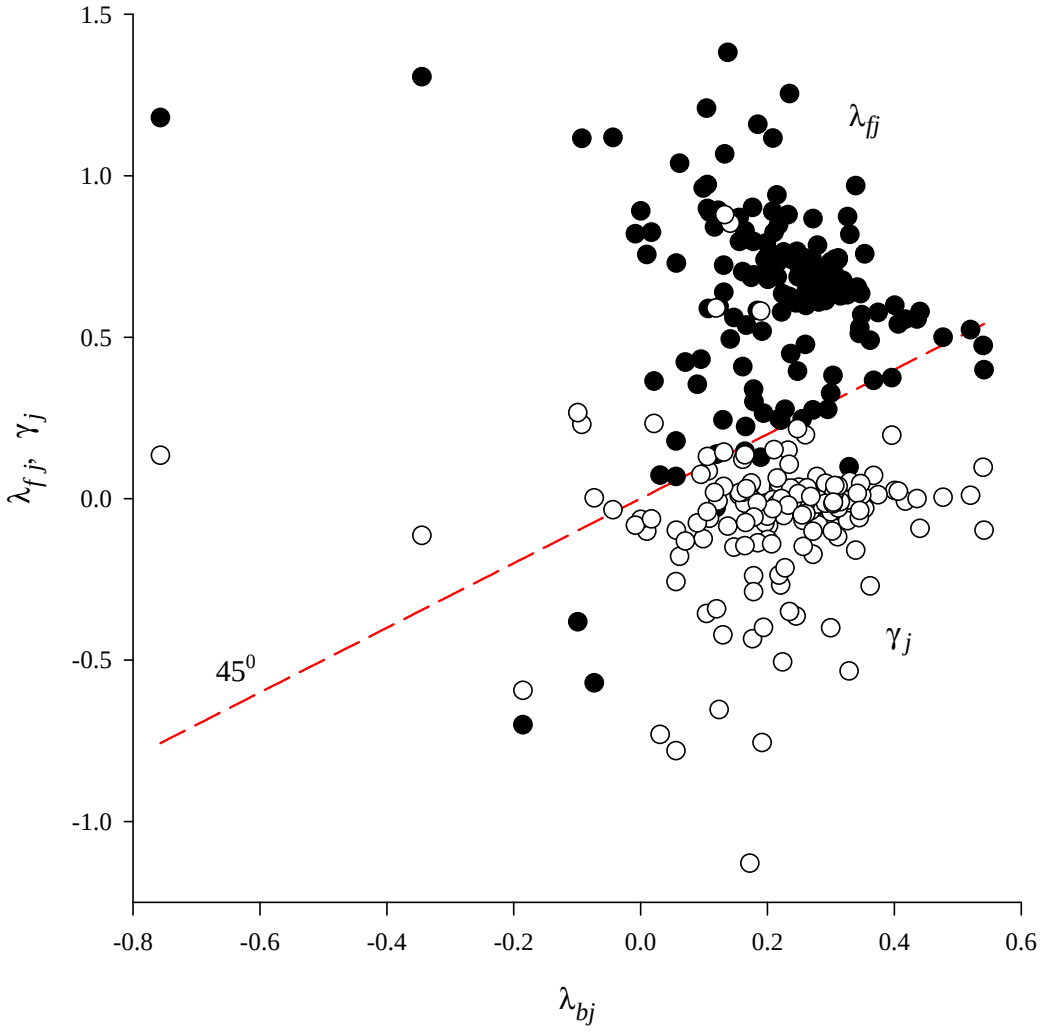

Table 4: Estimation with Historical Data

$$\pi_t = \lambda_0 + \lambda_b \pi_{t-1} + \lambda_f \mathrm{E}_t \pi_{t+1} + \gamma u_t$$

instruments	$\hat{\lambda}_b$ *(se)*	$\hat{\lambda}_f$ *(se)*	$\hat{\gamma}$ *(se)*	Obs.	$\overline{R}^2$
$\{\iota, \pi_{t-1}, u_t, \pi_{t-2}, u_{t-1}\}$	0.447 (0.104)	0.374 (0.213)	0.027 (0.138)	104	0.46
$\{\iota, \pi_{t-1}, u_t, \pi_{t-2}, \pi_{t-3}, u_{t-1}, u_{t-2}\}$	0.399 (0.100)	0.453 (0.233)	0.011 (0.138)	104	0.40
$\{\iota, \pi_{t-1}, \pi_{t-2}, \pi_{t-3}, u_{t-1}, u_{t-2}\}$	0.372 (0.104)	0.561 (0.254)	-0.019 (0.142)	104	0.38
$\{\iota, \pi_{t-1}, \pi_{t+1}, u_t\}$ OLS	0.382 (0.089)	0.380 (0.096)	0.032 (0.136)	104	0.46
Median SPF	-0.297 (0.067)	1.37 (0.101)	-0.234 (0.089)	104	0.73

Notes: ι is a vector of ones. Obs is the number of observations. Standard errors are heteroskedasticity- consistent. Estimation is for 1981:3–2007:3.
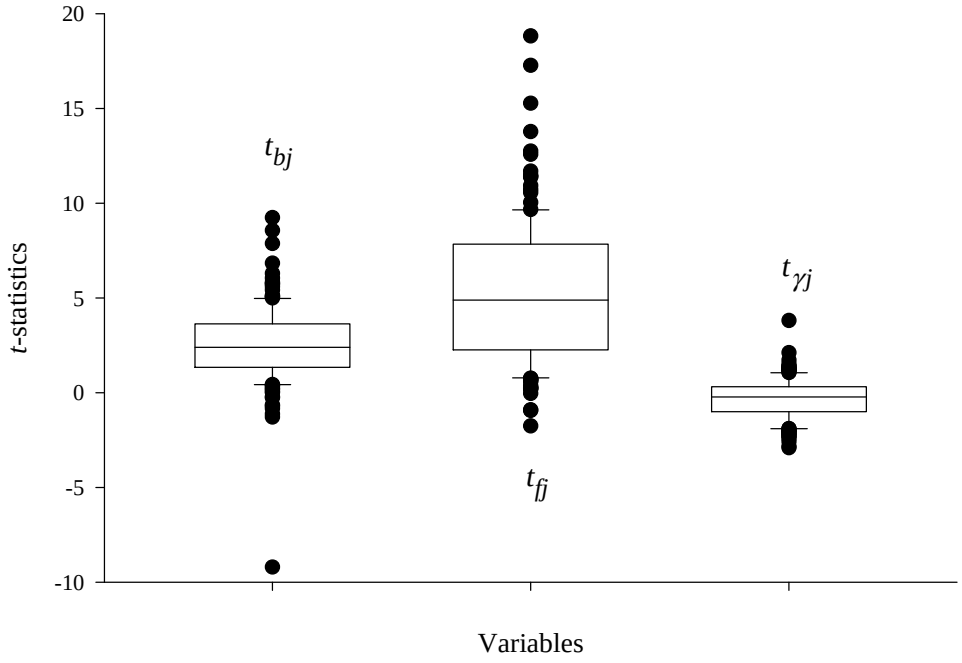
Figure 3: Forecaster-Specific *t*-statistics