Article:

Cornea-Madeira, Adriana orcid.org/0000-0002-0889-7145 and Rodrigues Madeira, Joao Antonio orcid.org/0000-0002-7380-9009 (2022) Econometric analysis of switching expectations in UK inflation. Oxford Bulletin of Economics and Statistics. pp. 651-673. ISSN 0305-9049

https://doi.org/10.1111/obes.12479

Econometric analysis of switching expectations in UK inflation

Adriana Cornea-Madeira and João Madeira*

November 2021

Abstract

We estimate with UK data a Phillips curve model with backward-looking and forward-looking methods of determining inflation expectations and with agents switching between these based on their recent performance. We find that, while on average backward-looking and forward-looking methods have about equal weight, there are considerable movements in the weight given to each method. We show this model has better in-sample fit than other Phillips curve models and this is robust to the methodology chosen. The model out-of-sample forecasts on certain dates do better than other Phillips curve models and the Atkeson and Ohanian model.

Total number of words: 7634

JEL Classification: C22, E31, E37.

Keywords: Phillips curve, heterogeneous expectations, inflation forecasting.

*Adriana Cornea-Madeira: University of York; Email: adriana.cornea-madeira@york.ac.uk. João Madeira: University of York; Email: joao.madeira@york.ac.uk. We would like to thank the editor Francesco Zanetti and three anonymous reviewers for useful comments and suggestions. Any errors are ours alone.

1 Introduction

In this paper, we examine how inflation expectations have evolved over time in the UK by switching between backward (the forecast corresponds to lagged inflation) and forward-looking or fundamentalist (the forecast is based on a discounted sequence of expected future output gaps, as implied by a purely forward-looking Phillips curve) rules.[1] We do this by estimating with UK data on inflation and output gap a model of inflation dynamics, developed by Cornea-Madeira et al. (2019), that allows for endogenous switching between different expectation rules. The model allows agents to switch between backward-looking and forward-looking rules based on their recent relative forecast performance (as in Brock and Hommes, 1997).

Our motivation is to show not only that endogenous switching between different expectation rules is a better description of human behavior (as documented for example in lab experiments, see Hommes, 2011) than the extreme cases of rational expectations (which assumes agents have the best possible understanding of how the economy works) or adaptative/naive expectations (which assumes that agents have no knowledge of how the economy works) but also that endogenous switching will improve the Phillips curve fit to inflation data (in-sample and out-of-sample).

We find that in the UK on average backward-looking and forward-looking behavior have about equal weight.[2] There are however substantial fluctuations in how expectations are formed over time.[3] In some periods forward-looking behavior is dominant (with the weight on the fundamental forecasting rule on occasion exceeding 0.8) while in others backward-looking

[1]There have been previous studies which have studied UK inflation using forward-looking expectations (see for example: Zanetti, 2011, and Faccini et al., 2013) or both backward and forward-looking expectations (see for example: Batini et al., 2005) but to the best of our knowledge no previous work has considered endogenous switching between forecast rules.

[2]For the US (see Cornea-Madeira et al., 2019) the backward-looking rule has on average about twice the weight of the forward-looking rule (one reason for this is that inflation is much less persistent in the UK, with first order autocorrelation only about 0.7, whereas in the US the value is about 0.9 for the same period).

[3]This is consistent with experimental evidence at the micro level that found individuals adjust expectations in response to past forecast errors (see Bloomfield and Hales, 2002, and Hommes, 2011) and with the finding of multiple structural breaks in the coefficient estimates from US macro data of the backward-looking and forward-looking terms of Phillips curve models (see Zhang et al., 2008, and Hall et al., 2012).

behavior is dominant (with the weight on the fundamental forecasting rule on occasion below 0.1). For example, we found that after the 1973 and 1979 oil price crises that there was an increase in backward-looking behavior. We also found a large and persistent increase in the backward-looking behavior in the early years of the XXI century, which seems to be due to a prolonged period in which UK inflation was below fundamentals (as determined by the present-value of expected future output gaps). Finally, we show that the cyclical behavior in adoption of forecast rules changed over time in the UK. Prior to the inflation targeting regime (which started at the end of 1992) the estimated series for the weight on the fundamentalist rule was much more volatile. This indicates that inflation expectations became more anchored under inflation targeting.

The estimation of the Phillips curve model with endogenous switching is shown to be quite robust to inclusion of different variables in estimation of the fundamentalist rule forecast, use of different number of lags in estimation of the fundamentalist rule forecast and changes to the driving variable in the vector autoregressive (VAR) model used in estimation of the fundamentalist rule forecast.

The Bayesian information criterion and Akaike information criterion both show that the Phillips curve model with endogenous switching in forecast rules accounts better for in-sample inflation dynamics than other Phillips curve models. We also show that the Phillips curve model with endogenous switching in forecast rules forecasts better out-of-sample in some periods than any of the competing Phillips curve models. Also, the Phillips curve model with endogenous switching in forecast rules does better in out-of-sample forecasting in some periods than the random walk model of Atkeson and Ohanian (2001), which has been shown to be one of the best econometric models in forecasting US inflation (Faust and Wright, 2013). Moreover, the Phillips curve model with endogenous switching in forecast rules has the advantage over econometric models that it is not as vulnerable to the Lucas (1976) critique (which argues that it can be misleading to try to predict the effects of a change in economic policy entirely on the basis of relationships observed in historical data).

Aside from Cornea-Madeira et al. (2019) our work is related to other empirical applications of the Brock and Hommes (1997) mechanism for endogenous switching in expectation rules. Other applications include the aggregate stock prices (Boswijk et al., 2007), commodity markets (Westerhoff and Reitz, 2005), oil prices (Ellen and Zwinkels, 2010) and exchange rates (Jongen et al., 2012).

The remainder of the paper is structured as follows. Section 2 consists of a review of related Phillips curve literature. Section 3 presents the Phillips curve models which we afterwards fit to UK data. The estimation methodology and results are in Section 4. Concluding remarks are in Section 5.

2 Literature review

The analysis by Phillips (1958) revealed a correlation between inflation and the level of economic activity in the United Kingdom (UK). This relation became known as the Phillips curve and has been confirmed to be valid in UK data in more recent work (see for example: Batini et al., 2005, Zanetti, 2011, and Faccini et al., 2013). However, most research on the Phillips curve has been done with United States (US) data (for a survey see Fuhrer et al., 2009) and to a lesser extent with Euro area data (see for example: Galí et al., 2001).

A key question in the literature and of significant relevance to monetary policy makers is how agents form expectations of future inflation. The purely forward-looking Phillips curve has emerged as the dominant framework and the assumption of rational expectations has become ubiquitous in macroeconomics (Coibion et al., 2018). Nonetheless, studies with US macro data are divided between those who conclude expectations are mostly forward-looking (see for example Galí and Gertler, 1999) and those who conclude expectations are mostly backward-looking (see for example Rudd and Whelan, 2006). Micro-level survey-based data on subjective expectations of individualshas revealed that expectations deviate from rational expectations in systematic and quantitatively important ways including forecast-error predictability and bias (see Mankiw et al., 2003). Inflation expectation surveys also support the view of heterogeneity in inflation expectations (see Branch, 2004, and Madeira and Zafar,

2015) and that agents inflation expectations change in response to economic conditions (see Carroll, 2003).

Coibion et al. (2018) consider that the expectation survey findings “highlight the need for using alternative frameworks (e.g., sticky information, noisy information, bounded rationality, models of learning) in describing how expectations are formed”. Coibion et al. (2018) argue that such alternative frameworks may help resolve some of the identified shortcomings of the Phillips curve and “call for careful consideration of expectation formation processes”. The analysis in this paper aims to do precisely that, by comparing several different ways of modelling UK inflation expectations (including applying for the first time to UK inflation the boundedly rational approach of Brock and Hommes, 1997) and their ability to account for the data. This adds to a growing literature which considers imperfect information (Mackowiak and Wiederholt, 2009, and Okuda et al., 2019), measuring expectations using survey data (Coibion and Gorodnichenko, 2015) and other methods to alleviate many of the previously identified puzzles in the Phillips curve literature.

3 The models

We start with the traditional Phillips curve (TPC) which relates price inflation ($\pi_t = p_t - p_{t-1}$) to a measure of economic activity, which we will refer as the output gap (y_t), and lagged inflation:

$$\pi_t = \beta\pi_{t-1} + \gamma y_t, \tag{1}$$

where β is a parameter close or equal to 1. The TPC implies inflation is determined by backward-looking behavior. For a recent example using the above TPC specification see Roeger and Herz (2012). The TPC has been able to characterize both US (see Rudebusch and Svensson, 1999) and Euro (Galí et al., 2001) inflation reasonably well and has proven relevant for forecasting (see Stock and Watson, 1999).

In the last two decades the new Keynesian Phillips curve (NKPC) model became a key reference to research on inflation dynamics (see Galí and Gertler, 1999, Rudd and Whelan, 2006, among others):

$$\pi_t = \beta E_t\{\pi_{t+1}\} + \gamma y_t. \tag{2}$$

Iterating (2) forward yields what researchers (see for example Galí and Gertler, 1999, and Cornea-Madeira et al., 2019) refer to as "fundamental inflation":

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t y_{t+k}. \tag{3}$$

The NKPC model shown above is purely forward-looking. In particular the model described in (2) and (3) cannot account for the apparent inertia in inflation present in the data (see Rudd and Whelan, 2006). To better capture inflation inertia researchers developed a "hybrid" new Keynesian Phillips curve (HNKPC) which includes both backward-looking and forward-looking behavior:

$$\pi_t = \beta[(1-\theta)\pi_{t-1} + \theta E_t\{\pi_{t+1}\}] + \gamma y_t. \tag{4}$$

Recently Cornea-Madeira et al. (2019) have developed a "behavioral" new Keynesian Phillips curve (BNKPC) which generalizes the hybrid model by allowing the weights on backward-looking and forward-looking behavior to vary endogenously over time:

$$\pi_t = \beta[(1-\theta_t)E_t^n\{\pi_{t+1}\} + \theta_t E_t^f\{\pi_{t+1}\}] + \gamma y_t + \xi_t, \tag{5}$$

where E^n denotes the forecast of agents who adopt a backward-looking rule and E^f is the forecast of agents who adopt a forward-looking rule. The term ξ_t represents the difference between the agents' forecast of individual prices and the average forecast of the average price. Since we have no data on such deviations, we treat ξ_t as part of the error term and use in our regressions standard errors that are robust to heteroskedasticity and autocorrelation (HAC) constructed using the Bartlett kernel (also known in econometrics as the Newey-West kernel, for details see Newey and West, 1987). Following Kurz et al. (2013), it is possible to aggregate the individual pricing rules:

$$\bar{E}_t\pi_{t+1} = [(1-\theta_t)E_t^n\{\pi_{t+1}\} + \theta_t E_t^f\{\pi_{t+1}\}], \tag{6}$$

where $\bar{E}_t$ denotes the average expectation of individuals.

The backward-looking rule or naive expectations consists of

$$E_t^n\{\pi_{t+1}\} = \pi_{t-1}. \tag{7}$$

Despite its simplicity, the naive rule is near to optimal in a forecasting sense if inflation is close to a unit root process (previous research with US data suggests this is the case, see Stock and Watson, 2007). The rule is also close to optimal in a price setting sense because it implies no persistent deviations between the rule and optimal behaviour, that is in a steady state equilibrium the rule is consistent with optimal behavior (see Gali and Gertler, 1999).

The forecast of agents who adopt a forward-looking or fundamentalist rule consists of

$$E_t^f\{\pi_{t+1}\} = \gamma \sum_{k=1}^{\infty} \beta^{k-1} E_t^f y_{t+k}. \tag{8}$$

Equation (8) is derived by leading (3) one period ahead and applying the fundamentalist expectation operator E^f to both sides.[4]

It is important to note that agents adopting the fundamental rule in (8) differ from agents with rational expectations. In particular, agents who choose the fundamental rule do not take into account the existence of agents who adopt the naive rule. That is, fundamental expectations (unlike rational expectations) are not model consistent. In order to have model consistent expectations agents would need to obtain an incredible amount of information, including details about the beliefs of other agents (and their number) in the economy. More realistically, agents in our framework only have knowledge of their objectives and of the constraints that they face, and therefore they do not have a complete model of determination

[4]In obtaining (8) we made use of the law of iterated expectations at the individual level as is standard in the learning literature (see for example Evans and Honkapohja, 2001, and Branch and McGough, 2009).

of aggregate variables.[5]

Agents in the BNKPC model are allowed to switch between forecasting rules based on their recent empirical performance. We define the absolute forecast error (F) in the previous K periods as:

$$F_t^i = \sum_{k=1}^{K} \left| E_{t-k}^i\{\pi_{t-k+1}\} - \pi_{t-k+1} \right|, \tag{9}$$

with $i = f, n$. We select a number of $K = 4$ lags for measuring past performance (because we are considering quarterly frequency, this means agents evaluate rules by looking at the forecast errors over the past year).[6]

The fraction of agents choosing the fundamentalist rule (θ_t) is updated every period according to the following evolutionary fitness measure:

$$V_t = -\frac{F_t^f}{F_t^f + F_t^n}. \tag{10}$$

Agents compare, at the beginning of every period t, the realized relative performances of the fundamentalist and naive strategies. As in Brock and Hommes (1997) the endogenous selection of expectation rules is made according to the following discrete choice model with multinomial logit probabilities (see Manski and McFadden, 1981):

$$\theta_t = \frac{\exp(\phi V_{t-1})}{\exp(\phi V_{t-1}) + \exp(-\phi \frac{F_{t-1}^n}{F_{t-1}^f + F_{t-1}^n})} = \frac{1}{1 + \exp(\phi(\frac{F_{t-1}^f - F_{t-1}^n}{F_{t-1}^f + F_{t-1}^n}))}, \tag{11}$$

where $\phi \geq 0$ determines the intensity of choice and reflects the sensitivity of agents to selecting the optimal prediction strategy. Equation (11) implies that rules with higher recent fitness attract more individuals. For $\phi = 0$ differences in fitness cannot be observed, therefore no switching occurs, implying that θ_t is constant and equal to 1/2. If $\phi = \infty$ the fitness can

[5]The method in Carriero (2008) clearly rejects that agents form expectations in a model consistent manner as assumed by rational expectations. However, the author also warns that his results do not exclude other forms of forward-looking behavior and advocates researchers to explore such possibilities (something which we do here).

[6]In the online appendix we also consider $K = 6$, 8, 10 and 12.

be observed perfectly and everyone always chooses the best predictor.

The TPC lacked theoretical support since it was not explicitly derived from optimizing behavior (making it subject to the Lucas, 1976, critique). The use of microfoundations led to the development of the NKPC. The HNKPC augments the NKPC with a backward-looking rule which while ad hoc has, as argued by Galí and Gertler (1999), several appealing features: i) if inflation is stationary then deviations to optimal behavior are not persistent; and ii) the rule implicitly incorporates information about the future since past inflation is also determined by forward-looking agents. Similarly, the BNKPC also adds to the microfounded environment of the NKPC some ad hoc elements: a) the choice of forecast rules (we adopted the same rules as in Cornea-Madeira et al., 2019, which implies the model becomes the same as the TPC if $\theta_t = 0$ and the same as the NKPC if $\theta_t = 1$); b) the number of forecast rules (as in prior literature we consider only a backward and forward-looking rule but the model could allow for more); and c) the number of lags for measuring past performance (as in Cornea-Madeira et al., 2019, we considered the forecast errors over the past year because it is behaviorally plausible given many prices are reviewed on an annual basis, but we analyse the robustness of our findings to other choices). However, the BNKPC also adds a new microfounded element to the NK framework since the endogenous selection of expectation rules in (11) can be obtained from optimization using the rational inattention approach, as shown by Matêjka and McKay (2015), in which case ϕ is inversely related to the shadow cost of allocating attention to decision making.

4 Estimation

4.1 Data

We estimate the Phillips curves models shown in the previous section with UK data from 1966:1 to 2016:4. Inflation is obtained as four times the log difference of the gross domestic product (GDP) deflator (this is given by the ratio of nominal to real GDP) while the output gap (y_t) is measured as the one-sided Hodrick-Prescott (HP) filtered log of real GDP constructed using the Kalman filter (see Stock and Watson, 1999).[7] In constructing measures of fundamental inflation we will also use the log of the labor share of income (s_t) constructed as total compensation divided by nominal GDP. Because the Phillips curve models in Section 3 are written in log deviations from the steady state, we demean all variables prior to estimation. Nominal GDP, real GDP and total compensation were obtained from the Office for National Statistics (ONS) with CDIDs (ONS's four-letter variable identifiers) given respectively by ABMI, YBHA and DTWM.

4.2 Methodology

Our regression method is nonlinear least squares (NLS) using the Levenberg-Marquardt algorithm (Seber and Wild, 2003). The TPC shown in equation (1) can be easily estimated from the lagged value of inflation. However, we need a series for expectations of one period ahead inflation of forward-looking agents to estimate the NKPC, the HNKPC and the BNKPC shown in equations (2), (4) and (5) respectively. We do this by using a multivariate vector autoregressive (VAR) model:

$$\mathbf{Z}_t = \mathbf{A}_t \mathbf{Z}_{t-1} + \boldsymbol{\epsilon}_t, \tag{12}$$

where $\mathbf{Z}_t$ is a vector of variables used to predict future values of the output gap (which constitutes the first variable in the VAR). The matrix $\mathbf{A}_t$ is the matrix of coefficients and is estimated at each point in time t by Ordinary Least Squares (OLS) using data from 1 to t, $t = 1, \cdots, T$, where T is the total sample size. Fundamentalists do not have perfect foresight of the future. Therefore, they are only able to estimate the VAR model using the information set available at date t and $\mathbf{A}_t$ is then updated every period as new information becomes available. The VAR estimation is initialized using available presample data starting in 1958:2.

[7] For a detailed exposition of the HP filter see Cornea-Madeira (2017).

We can then, in the spirit of Campbell and Shiller (1987), rewrite (8) as:

$$E_t^f\{\pi_{t+1}\} = \gamma \mathbf{e}_1'(\mathbf{I} - \beta\mathbf{A}_t)^{-1}\mathbf{A}_t\mathbf{Z}_t, \tag{13}$$

where $\mathbf{e}_1$ is a suitable defined unit vector and the discount factor β is fixed at 0.99 as standard in the literature (see for example Rudd and Whelan, 2007, and Cornea-Madeira et al., 2019).

As in Cornea-Madeira et al. (2019) we choose a four lag VAR. The variables in our baseline VAR specification will be the output gap and the labor share. The use of the labor share to forecast the output gap in the VAR is justified since under certain conditions the output gap is proportional to the labor share and simulations suggest that even if those conditions are not met "that the relation remains very close to proportionate" (Galí and Gertler, 1999:201). In the estimation of (1), (2), (4) and (5) the fitted value of the output gap from the VAR is used to avoid endogeneity concerns.

4.3 Main results

The estimation results for the models in Section 3 with the baseline VAR specification are shown in Table 1. The results are favorable to the idea that agents switch forecasting methods when making expectations of inflation, since: a) the parameter estimate in the BNKPC model for ϕ, which governs the intensity of choice, is statistically significant at the 1% level; and b) the BNKPC is the model favored by the Bayesian information criterion (BIC), Akaike information criterion (AIC) and Adj-R^2. The value obtained for γ of the BNKPC is 0.390. This positive value, which is statistically significant at the 1% level, confirms the presence of a short-run trade-off between inflation and the level of economic output.

Table 2 shows the conclusions above are robust to using: different number of lags in the VAR (in robustness check 1 the VAR has 8 lags), different driving variable in the VAR (in robustness check 2 the driving variable is the labor share instead of the output gap), both different number of lags in the VAR and different driving variable in the VAR (in robustness check 3 the VAR has 8 lags and the labor share as driving variable), and different sets of variables in the VAR (in robustness check 4 we add lagged inflation to the VAR).

In all robustness checks we obtain that in the BNKPC model the intensity of choice parameter ϕ is highly statistically significant.[8] In all robustness checks we obtain a positive and statistically significant value for γ (which measures the trade-off between inflation and output) for the BNKPC model (and also HNKPC model) but not for the TPC (not significant in any of the checks) and NKPC (not significant in robustness check 4) models. Table 2 also shows that the BNKPC model is favored by the BIC, AIC and Adj-R^2 in all robustness checks.

In summary, we show that a behavioral model of inflation accounts better for the Phillips curve relation.

4.4 Inflation expectation dynamics in the behavioral model

In Figure 1 we show the realized one quarter forward inflation rate (π_{t+1}) and the expected one quarter forward inflation rate of the BNKPC model ($\hat{E}_t^b\{\pi_{t+1}\}$) given by:

$$\hat{E}_t^b\{\pi_{t+1}\} = (1-\hat{\theta}_t)E_t^n\{\pi_{t+1}\} + \hat{\theta}_t\hat{E}_t^f\{\pi_{t+1}\}, \tag{14}$$

where E_t^n is as in equation (7), and $\hat{\theta}_t$ and $\hat{E}_t^f$ are obtained from equations (11) and (13) using the estimates in Table 1.

Figure 1 reveals that the expected inflation rate resulting from the BNKPC is fairly close to the realized future inflation rate. However, two significant distinctions between $\hat{E}_t^b\{\pi_{t+1}\}$ and π_{t+1} can be easily observed from Figure 1. The first is that realized inflation is considerably more volatile than the expected inflation series of the BNKPC (the standard deviation of the former is 0.058 whereas that of the latter is only 0.041). The second is that expected inflation from the BNKPC model only caught up with a delay the sharp and large

[8] It is relevant to note however that the absolute value of ϕ is difficult to interpret since it depends on the functional form of the evolutionary fitness measure V_t (Cornea-Madeira et al., 2019). Moreover, large changes in ϕ cause only small variation in the estimation of the series θ_t (see page 36 of Hommes and Wagener, 2008).

increases in the inflation rate which occurred in the mid and late 70s.

Figure 1 also suggests that UK inflation is substantially more volatile (the standard deviation is 0.058 for UK inflation between 1966:1 and 2016:4 which is more than twice that of US inflation in the same period) and less persistent than in the US (the first order autocorrelation is about 0.7 for UK inflation between 1966:1 and 2016:4 whereas for the US it is about 0.9 in the same period).

Figure 2 displays the expectation error of the fundamental rule ($\pi_{t+1}- \hat{E}_t^f\{\pi_{t+1}\}$) and the naive rule ($\pi_{t+1}- E_t^n\{\pi_{t+1}\}$). The fundamental rule shows large expectation errors in the 60s but afterwards the series is centered around zero. The summary statistics are shown in Table 3. We can see that the values for the median, maximum and correlation with the output gap (both have low correlation with y_t) are similar for the two rules. However, the naive rule is on average substantially more accurate (-0.0% mean) than the fundamental rule (-0.9% mean). We also see that the skewness value is -2.136 for the fundamental rule which indicates that the tail on the left side of the probability density function is longer or fatter than the right side (this is consistent with the fact that the minimum value is much higher in absolute terms than the maximum value). The fundamental rule therefore is more likely to predict a lower value than realized inflation. For the naive rule skewness is only 0.111 (consistent with the fact that the maximum and minimum values are similar in absolute terms). There are also differences with respect to the persistence of the expectation error of the two rules. The first order autocorrelation for the fundamental rule is 0.804 and for the naive rule is -0.014.

Figure 2 and Table 3 make it clear that the fundamental rule leads to systematic (average of -0.9%) and predictable (substantial degree of autocorrelation) errors. This can seem puzzling at first but recall that fundamentalists use a VAR, see (12), to forecast future values of the output gap (not future values of inflation), as can be seen in (7). Therefore, the econometric methodology does not ensure that the fundamentalist rule is an unbiased predictor of future inflation. The finding that the fundamental rule leads to systematic and predictable errors is consistent with the results from Table 1 and Table 2 which indicate that the pure forward-looking Phillips curve (the NKPC) is the model with worse fit to the data. While our results indicate that the present-value of future values of the output gap is by itself a poor empirical model of inflation (which is also the case in US data, for example in Rudd and Whelan, 2007, it is shown to have an $R^2 = 0.13$ whereas an autoregressive, AR, process for inflation with 4 lags would have an $R^2 = 0.75$), our analysis also indicates that it has some worth, since the models with best fit to the data (the HNKPC and BNKPC) make use of the fundamental rule.

We now study the evolution of adoption of forecast rules in the behavioral model. Figure 3 shows the estimated time series for the fraction of agents adopting the fundamental rule ($\hat{\theta}_t$). The series for $\hat{\theta}_t$ clearly has substantial fluctuations, with at times the economy being dominated by forward-looking behavior and at other times by backward-looking behavior. For instance, one can observe a large increase in the share of agents adopting the fundamental rule at the onset of the two oil crises in the 70s. The first oil crisis started in late September/October of 1973 and ended in March 1974 (although energy prices continue to increase after this and the same happened to inflation which was on average 8.7% in 1973, 18% in 1974 and 22% in 1975).[9] The fraction of fundamentalists at the start of 1973 was 0.55 but by the first quarter of 1974 it had reached 0.82 (and fell quickly after that, by the first quarter of 1975 it was only 0.27). This is consistent with the results of Cornea-Madeira et al. (2019) for US data who also document a large increase in the weight of the fundamental rule in the early 70s (peaking in 1974) which was followed by a rapid fall. The second oil crisis occurred in 1979 in the wake of the Iranian Revolution. At the start of the year the fraction of fundamentalists was below 0.48 but by the third quarter (when inflation reached its peak at 29%) it was 0.65 (and fell quickly after that). Oil prices also increased significantly in the

[9]In September 15 of 1973 the Organization of Petroleum Exporting Countries (OPEC) declared a negotiating front, consisting of the 6 Persian Gulf States, to pressure for price increases and an end to support of Israel, based on the 1971 Tehran agreement. In October of that year, the forces of Egypt and Syria attempted to overwhelm the state of Israel in an offensive later known as the Yom Kippur War. In October 1973, the Organization of Arab Petroleum Exporting Countries (OAPEC, consisting of the Arab majority of OPEC plus Egypt and Syria) declared significant production cuts and an oil embargo against the United States and other industrialized nations that supported Israel in the Yom Kippur War. The oil embargo ended in March 1974 following intense diplomatic activity.

first half of 2009 (at the end of 2008 the Brent crude oil price - Europe was 35 dollars per barrel and by early August it reached 75 dollars per barrel and then stabilized for the next year). In this case too an increase on the weight of the fundamental rule occurred (from 0.320 in the last quarter of 2008 to 0.629 in the third quarter of 2009).

Figure 3 also shows that in the UK there was a quite significant and persistent fall in the fraction of agents adopting the fundamental rule during the first decade of the XXI century. The reason for this seems to be that during that decade realized UK inflation was systematically below fundamental inflation (this can be observed from Figure 2, which shows that there were very few positive values for the expectation error of the fundamental rule in that decade). Figure 2 in Cornea-Madeira et al. (2019) shows that the association of a low fraction of fundamentalists with realized inflation systematically below fundamental inflation also occurred in the US during most of the 90s.

The summary statistics for the estimated time series weight for the fundamental rule ($\hat{\theta}_t$) are in Table 4. For the whole sample (1966:1 to 2016:4) both the mean and median for $\hat{\theta}_t$ are about 1/2. This is substantially higher than the 0.353 mean and 0.276 median that Cornea-Madeira et al. (2019) found for the US. The standard deviation of $\hat{\theta}_t$ from 1966:1 to 2016:4 was 0.228. In the US case, the value of standard deviation of the weight for the fundamental rule was a considerably higher 0.282 (see Cornea-Madeira et al., 2019). Other significant differences between the results we obtain with UK data and previous research with US data is with respect to skewness. We document a skewness value of -0.231 for the whole sample. This indicates that the tail on the left side of the probability density function of $\hat{\theta}_t$ is longer or fatter than the right side. Whereas Cornea-Madeira et al. (2019) found a distribution with positive skew (0.418) in the US case. This means that, while in the US the median of fundamentalists is substantially lower than in the UK, periods with a fraction of fundamentalists above the median are more frequent. One aspect in which both the US and UK cases seem quite similar is with respect to persistence. We obtain a first-order autocorrelation value of 0.836 for $\hat{\theta}_t$ in the whole sample. In the US case Cornea-Madeira et al. (2019) also found the weight for the fundamental rule to exhibit high persistence (0.887

first-order autocorrelation).

Table 4 also shows how the fraction of fundamentalists changed with the adoption of inflation targeting (which started in the UK in 1992:4). Table 4 shows that with the inflation targeting regime the weight on the forward-looking term became much less volatile (the standard deviation was 0.261 prior to inflation targeting and 0.185 afterwards). This suggests that inflation expectations became more anchored under the inflation targeting period and agents started switching less between different inflation expectation rules. The average weight on the forward-looking rule is slightly smaller in the inflation targeting period. This is in contradiction with the results of Benati (2008), who finds an association between higher weights on forward-looking behavior and inflation targeting. However, it is in agreement with the findings of Stock and Watson (2007) who show that the backward-looking rule has become very hard to outperform in recent decades.

In summary, we find substantial variation over time in the adoption of inflation forecast rules in the UK. We also find that UK inflation dynamics differ substantially from what is observed for the US (the former are more volatile and less persistent) and so do inflation expectations (in particular UK inflation expectations have a larger average weight on the fundamentalist rule).

4.5 Out-of-sample forecasting

In this section we compare the point forecast accuracy of our BNKPC model with the TPC, NKPC and the HNKPC models. In the following we construct iterated point forecasts from these models based on the recursive window scheme (the estimates are updated at each forecast origin using all available information). To obtain the forecasts we use real time data from the ONS available from January 1990 until September 2016 for real GDP, GDP price deflator, income based gross value added and total compensation of employees (we use these to obtain the output gap, inflation and labor share series in real time). We obtain quarterly data by choosing the data in March, June, September and December in each year. In forming the first forecast, we estimate the models using data from 1990:1 to 1996:1 and then form forecasts for horizons (quarters) $h = 1, \ldots, 16$. We then move forward one quarter and estimate the model from 1990:1 to 1996:2, and form forecasts for horizons 1–16. We continue in a similar fashion through the rest of the sample until 2016:2.

The iterative forecasts from the BNKPC baseline model are based on the one-step ahead forecasts from this model. Denote by $\hat{\beta}$, $\hat{\gamma}$, $\hat{\mathbf{A}}_t$ the parameter estimates based on the sample up to time t. Notice that the model delivers also an estimate of π_t, i.e.,

$$
\begin{aligned}
\hat{\pi}_t &= \beta(\hat{\theta}_t \hat{E}_t^f\{\pi_{t+1}\} + (1-\theta_t)E_t^n\{\pi_{t+1}\}) + \hat{\gamma}\hat{y}_t \\
&= \beta\left(\hat{\theta}_t \hat{\gamma}\mathbf{e}_1'(I - \beta\hat{\mathbf{A}}_t)^{-1}\hat{\mathbf{A}}_t\mathbf{Z}_t + (1-\hat{\theta}_t)\pi_{t-1}\right) + \hat{\gamma}\hat{y}_t, \qquad (15)
\end{aligned}
$$

where $\hat{\theta}_t = \left(1 + \exp\left(\hat{\phi}\left(\frac{F_{t-1}^f - F_{t-1}^n}{F_{t-1}^f + F_{t-1}^n}\right)\right)\right)^{-1}$ and $\hat{y}_t$ is the fitted output gap from the VAR.

The one period ahead ($h = 1$) forecast of inflation, $\hat{\pi}_{t+1}$, is obtained as:

$$
\begin{aligned}
\hat{\pi}_{t+1} &= \beta(\hat{\theta}_{t+1}\hat{E}_t^f\{\pi_{t+2}\} + (1-\hat{\theta}_{t+1})\hat{E}_t^n\{\pi_{t+2}\}) + \hat{\gamma}\hat{y}_{t+1} \\
&= \beta(\hat{\theta}_{t+1}\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_t)^{-1}\hat{\mathbf{A}}_t^2\mathbf{Z}_t + (1-\hat{\theta}_{t+1})\hat{\pi}_t) + \hat{\gamma}\hat{y}_{t+1}, \qquad (16)
\end{aligned}
$$

where $\hat{y}_{t+1}$ is the first variable from $E_t\mathbf{Z}_{t+1} = \hat{\mathbf{A}}_t\mathbf{Z}_t$ (the output gap) and $\hat{\theta}_{t+1} = \left(1 + \exp\left(\hat{\phi}\left(\frac{F_t^f - F_t^n}{F_t^f + F_t^n}\right)\right)\right)^{-1}$.

The forecast errors at time t are:

$$
\begin{aligned}
F_t^n &= |\pi_{t-5} - \pi_{t-3}| + |\pi_{t-4} - \pi_{t-2}| + |\pi_{t-3} - \pi_{t-1}| + |\pi_{t-2} - \hat{\pi}_t|, \qquad (17) \\
F_t^f &= |\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_{t-4})^{-1}\hat{\mathbf{A}}_{t-4}\mathbf{Z}_{t-4} - \pi_{t-3}| + |\hat{\gamma}\mathbf{e}_1'(I - \beta\hat{\mathbf{A}}_{t-3})^{-1}\hat{\mathbf{A}}_{t-3}\mathbf{Z}_{t-3} - \pi_{t-2}| \\
&+ |\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_{t-2})^{-1}\hat{\mathbf{A}}_{t-2}\mathbf{Z}_{t-2} - \pi_{t-1}| + |\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_{t-1})^{-1}\hat{\mathbf{A}}_{t-1}\mathbf{Z}_{t-1} - \hat{\pi}_t|. \qquad (18)
\end{aligned}
$$

The two-period ahead ($h = 2$) forecast of inflation, $\hat{\pi}_{t+2}$ is obtained as:

$$
\begin{aligned}
\hat{\pi}_{t+2} &= \beta(\hat{\theta}_{t+2}\hat{E}_t^f\{\pi_{t+3}\} + (1-\hat{\theta}_{t+2})\hat{E}_t^n\{\pi_{t+3}\}) + \hat{\gamma}\hat{y}_{t+2} \\
&= \beta(\hat{\theta}_{t+2}\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_t)^{-1}\hat{\mathbf{A}}_t^3\mathbf{Z}_t + (1-\hat{\theta}_{t+2})\hat{\pi}_{t+1}) + \hat{\gamma}\hat{y}_{t+2}, \qquad (19)
\end{aligned}
$$

where $\hat{y}_{t+2}$ is the first variable from $E_t\mathbf{Z}_{t+2} = \hat{\mathbf{A}}_t^2\mathbf{Z}_t$; $\hat{\theta}_{t+2} = \left(1 + \exp\left(\hat{\phi}\left(\frac{F^f_{t+1} - F^n_{t+1}}{F^f_{t+1} + F^n_{t+1}}\right)\right)\right)^{-1}$, and the forecast errors are:

$$F^n_{t+1} = |\pi_{t-4} - \pi_{t-2}| + |\pi_{t-3} - \pi_{t-1}| + |\pi_{t-2} - \hat{\pi}_t| + |\pi_{t-1} - \hat{\pi}_{t+1}|, \quad (20)$$

$$\begin{aligned} F^f_{t+1} &= |\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_{t-3})^{-1}\hat{\mathbf{A}}_{t-3}\mathbf{Z}_{t-3} - \pi_{t-2}| + |\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_{t-2})^{-1}\hat{\mathbf{A}}_{t-2}\mathbf{Z}_{t-2} - \pi_{t-1}| \\ &+ |\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_{t-1})^{-1}\hat{\mathbf{A}}_{t-1}\mathbf{Z}_{t-1} - \hat{\pi}_t| + |\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_t)^{-1}\hat{\mathbf{A}}_t\mathbf{Z}_t - \hat{\pi}_{t+1}|. \end{aligned} \quad (21)$$

The three-period ahead ($h = 3$) forecast of inflation is:

$$\begin{aligned} \hat{\pi}_{t+3} &= \beta(\hat{\theta}_{t+3}\hat{E}^f_t\{\pi_{t+4}\} + (1 - \hat{\theta}_{t+3})\hat{E}^n_t\{\pi_{t+4}\}) + \hat{\gamma}\hat{y}_{t+3} \\ &= \beta(\hat{\theta}_{t+3}\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_t)^{-1}\hat{\mathbf{A}}_t^4\mathbf{Z}_t + (1 - \hat{\theta}_{t+3})\hat{\pi}_{t+3}) + \hat{\gamma}\hat{y}_{t+3}, \end{aligned} \quad (22)$$

where $\hat{y}_{t+3}$ is the first variable from $E_t\mathbf{Z}_{t+3} = \hat{\mathbf{A}}_t^3\mathbf{Z}_t$; $\hat{\theta}_{t+3} = \left(1 + \exp\left(\hat{\beta}\left(\frac{F^f_{t+2} - F^n_{t+2}}{F^f_{t+2} + F^n_{t+2}}\right)\right)\right)^{-1}$, and the forecast errors are:

$$F^n_{t+2} = |\pi_{t-3} - \pi_{t-1}| + |\pi_{t-2} - \hat{\pi}_t| + |\pi_{t-1} - \hat{\pi}_{t+1}| + |\hat{\pi}_t - \hat{\pi}_{t+2}|, \quad (23)$$

$$\begin{aligned} F^f_{t+2} &= |\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_{t-2})^{-1}\hat{\mathbf{A}}_{t-2}\mathbf{Z}_{t-2} - \pi_{t-1}| + |\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_{t-1})^{-1}\hat{\mathbf{A}}_{t-1}\mathbf{Z}_{t-1} - \hat{\pi}_t| \\ &+ |\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_t)^{-1}\hat{\mathbf{A}}_t\mathbf{Z}_t - \hat{\pi}_{t+1}| + |\hat{\gamma}\mathbf{e}_1'(\mathbf{I} - \beta\hat{\mathbf{A}}_t)^{-1}\hat{\mathbf{A}}_t^2\mathbf{Z}_t - \hat{\pi}_{t+2}|. \end{aligned} \quad (24)$$

Subsequent forecasts are obtained in a similar fashion.

Figures 4 to 7 show the results of the fluctuation test by Giacomini and Rossi (2010) for forecast horizons of 1 to 16 quarters ahead. The Giacomini and Rossi (2010) approach tests for equal accuracy at each point in time of the competitor model relative to the benchmark BNKPC model.

If the curve is within the 95% band shown in figures 4 to 7, then the two models (the BNKPC model and the competitor model) have equal forecasting performance. If the curve is above the upper limit of the 95% confidence interval, then the BNKPC model has superior forecasting performance. If the curve is below the upper limit of the 95% confidence interval, then the competitor model has superior forecasting performance relative to the BNKPC model.

Figure 4 shows the BHNKPC forecasts better than the TPC model in several years after 2010 for all forecast horizons. For most of the remaining dates the BHNKPC and the TPC do equally well (only occasionally does the TPC outperform the BHNKPC).

Figure 5 shows the BNKPC outperforms in the years around 2010 the NKPC for most short term horizons (1, 2, 3, 6 and 7 quarters ahead) and some medium to long term horizons (10, 12, 13 and 15 quarters ahead) while doing equally well on most other dates. The NKPC does better only at some medium and long term horizons (5, 9 and 14 quarters ahead). This indicates that fundamentalists are not very good at forecasting inflation at short horizons but that over the medium to long term inflation tends to revert to fundamentals.

Figure 6 shows the BNKPC model outperforms in the years around 2010 the HNKPC at horizons 1, 2, 3, 7, 10, 12, 13 and 15. The HNKPC does better in the years around 2005 at horizons 5, 6 and 9 and does better in the years after 2010 at horizon 14. At other horizons it is not clear which model does better.

In Figure 7 we also compare the BNKPC to the Atkeson and Ohanian (2001), hence AO, random walk model. We include the AO since Faust and Wright (2013) showed this is one of the best performing econometric models in forecasting inflation. Figure 7 shows the BNKPC outperforms the AO model at most forecast horizons in the years around 2010 and towards the end of our sample. The fact that the BNKPC model is competitive in relation to one of the best econometric models is quite notable because econometric models are not constrained by economic theory. Moreover, because the BNKPC is built on microfoundations it is less vulnerable to the Lucas (1976) critique, having therefore the advantage that it can be used to make forecasts conditional on a policy change.

It is curious that the BNKPC model outperforms the competitor models in the years around 2010. From mid 2007 to early 2009 the UK economy suffered significantly from the global financial crisis. However, the predicted disinflation from conventional Phillips curve models due to reduced demand never materialized (International Monetary Fund, 2013). The reason for this is that household inflation expectations rose due to a dramatic increase in oil prices in the first 8 months of 2009 (Coibion and Gorodnichenko, 2015). As discussed previously in Section 4.4 oil price shocks are associated with an increase in the weight on the fundamental rule in the BNKPC model and this too happened in 2009. By placing low weight on the fundamental rule prior to 2009 (in those years realized inflation was systematically below fundamentals as shown in Figure 2) and a high weight when the oil price shock occurred the BNKPC model was able to forecast inflation better than its competitors.

In summary, the out-of-sample forecast performance of the BNKPC model improves over the other Phillips curve (TPC, NKPC and HNKPC) models considered in certain periods. Moreover, the BNKPC model also outperforms on some dates the out-of-sample forecast ability of the AO model which constitutes one of the most favored econometric models by researchers.

4.6 Additional empirical exercises

In an online appendix we include several other estimation results. The first exercise (shown in Table A1) shows the parameters of the BNKPC model for the slope of the Phillips curve and intensity of choice remain statistically significant with different values of K, the number of periods used by agents to measure the empirical performance of forecast rules. The second exercise estimates the BNKPC model with inflation detrended using a one-sided HP filter (see Table A2). This is shown to lead to more switching between expectations rules because inflation becomes less persistent after detrending. As a result, the intensity of choice parameter is estimated to be higher in value but with lower statistical significance (because the data is more volatile). The third exercise uses Monte Carlo simulations to demonstrate the reliability of NLS in estimating the BNKPC model (see Table A3). The last empirical exercise in the online appendix compares the BNKPC model to a HNKPC model with time-varying weights. This is done by estimating the HNKPC model in a rolling sample. This approach is similar to Canova (2006), but instead of using the Bayesian approach we use a frequentist approach (we estimate the model for each sub-sample by NLS to be able to compare the results with the BHNKPC model). The results are shown in Table A4 and reveal the BNKPC model predicts inflation better (on average, the time-varying HNKPC model underpredicts inflation by about 0.8% per quarter whereas the BNKPC model only does so by about 0.1%).

5 Conclusion

In this paper we estimate with UK data a Phillips curve model which allows for endogenous switching between forecast strategies. We find that on average backward-looking and forward-looking inflation expectations have about equal weight. However, there are large swings back and forth between backward-looking and forward-looking behavior over time. The estimation results also show that the model has a better in-sample fit (as indicated by the BIC and AIC) to inflation dynamics than alternative Phillips curve models (specifically, a purely backward-looking model, a purely forward-looking model and a model both backward-looking and forward-looking but with constant weights) and this is shown to be robust to several methodological choices.

Finally, we show in a out-of-sample forecast exercise that the Phillips curve model with endogenous switching expectations does better than any of the other Phillips curve models considered and than a leading econometric model by Atkeson and Ohanian (2001).

References

[1] Atkeson, A., and Ohanian, L., 2001. Are Phillips curves useful for forecasting inflation? *Federal Reserve Bank of Minneapolis Quarterly Review* 25, 2-11.

[2] Batini, N., Jackson, B., and Nickell, S., 2005. An open-economy new Keynesian Phillips curve for the U.K. *Journal of Monetary Economics* 52 (6), 1061-1071.

[3] Benati, L., 2008. Investigating Inflation Persistence Across Monetary Regimes. *The Quarterly Journal of Economics* 123 (3), 1005-1060.

[4] Bloomfield, R., and Hales, J., 2002. Predicting the next step of a random walk: Experimental evidence of regime-switching beliefs. *Journal of Financial Economics* 65 (3), 397–414.

[5] Boswijk, H., Hommes, C., and Manzan, S., 2007. Behavioral heterogeneity in stock prices. *Journal of Economic Dynamics and Control* 31 (6), 1938-1970.

[6] Branch, W., 2004. The Theory of Rationally Heterogeneous Expectations: Evidence From Survey Data on Inflation Expectations. *Economic Journal* 114 (497), 592–621.

[7] Brock, W., and Hommes, C., 1997. A rational route to randomness. *Econometrica* 65 (5), 1059–1095.

[8] Campbell, J., and Shiller, R., 1987. Cointegration and tests of present value models. *Journal of Political Economy* 95 (5), 1062–1088.

[9] Canova, F., 2006. Monetary policy and the evolution of the US economy. CEPR discussion paper no. 5467.

[10] Carriero, A., 2008. A simple test of the New Keynesian Phillips Curve. *Economics Letters* 100 (2), 241-244.

[11] Carroll, C., 2003. Macroeconomic Expectations of Households and Professional Forecasters. *Quarterly Journal of Economics* 118 (1), 269–298.

[12] Coibion, O., and Gorodnichenko, Y., 2015. Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation. *American Economic Journal: Macroeconomics* 7 (1), 197-232.

[13] Coibion, O., Gorodnichenko, Y., and Kamdar, R., 2018. The Formation of Expectations, Inflation, and the Phillips Curve. *Journal of Economic Literature* 56 (4), 1447-1491.

[14] Cornea-Madeira, A., 2017. The Explicit Formula for the Hodrick-Prescott Filter in a Finite Sample. *Review of Economics and Statistics* 99 (2), 314-318.

[15] Cornea-Madeira, A., Hommes, C., and Massaro, D., 2019. Behavioral Heterogeneity in U.S. Inflation Dynamics. *Journal of Business & Economic Statistics* 37 (2), 288-300.

[16] Ellen, S., and Zwinkels, R., 2010. Oil price dynamics: A behavioral finance approach with heterogeneous agents. *Energy Economics* 32 (6), 1427-1434.

[17] Faccini, R., Millard, S., and Zanetti, F., 2013. Wage Rigidities in an Estimated Dynamic, Stochastic, General Equilibrium Model of the UK Labour Market. *Manchester School* 81, 66-99.

[18] Faust, J., and Wright, J., 2013. Forecasting Inflation. In: *Handbook of Economic Forecasting*, edited by G. Elliott, C. Granger and A. Timmermann. Elsevier.

[19] Fuhrer, J., Kodrzycki, Y., Little, J., and Olivei, G. (ed.), 2009. *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective*. MIT Press Books.

[20] Galí, J., and Gertler, M., 1999. Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics* 44 (2), 195-222.

[21] Galí, J., Gertler, M., and Lopez-Salido, D., 2001. European inflation dynamics. *European Economic Review* 45 (7), 1237-1270.

[22] Giacomini, R., and Rossi, B., 2010. Forecast Comparisons in Unstable Environments. *Journal of Applied Econometrics* 25 (4), 595-620.

[23] Hall, A., Han, S., and Boldea, O., 2012. Inference Regarding Multiple Structural Changes in Linear Models With Endogenous Regressors. *Journal of Econometrics* 170 (2), 281–302.

[24] Hommes, C., and Wagener, F., 2008. Complex evolutionary systems in behavioral finance. CeNDEF Working Papers 08-05.

[25] Hommes, C., 2011. The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics and Control* 35 (1), 1-24.

[26] International Monetary Fund, 2013. The Dog That Didn't Bark: Has Inflation Been Muzzled or Was It Just Sleeping. Chapter 3 in *World Economic Outlook*, April 2013.

[27] Jongen, R., Verschoor, W., Wolff, C., and Zwinkels, R., 2012. Explaining dispersion in foreign exchange expectations: A heterogeneous agent approach. *Journal of Economic Dynamics and Control* 36 (5), 719-735.

[28] Kurz, M., Piccillo, G., and Wu, H., 2013. Modeling diverse expectations in an aggregated New Keynesian Model. *Journal of Economic Dynamics and Control* 37 (8), 1403-1433.

[29] Lucas, R., 1976. Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy* 1 (1), 19-46.

[30] Madeira, C., and Zafar, B., 2015. Heterogeneous Inflation Expectations and Learning. *Journal of Money, Credit and Banking* 47 (5), pages 867-896.

[31] Mankiw, N., Reis, R., and Wolfers, J., 2003. Disagreement about Inflation Expectations. In: Gertler, M., and Rogoff, K. (Eds.), NBER Macroeconomics Annual 2003, Volume 18, 209-270.

[32] Manski, C., and McFadden, D., 1981. Alternative estimators and sample designs for discrete choice analysis. In: *Structural Analysis of Discrete Data*, edited by C. Manski and D. McFadden. MIT Press.

[33] Matêjka, F., and McKay, A., 2015. Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. *American Economic Review* 105 (1), 272-298. January.

[34] Newey, W., and West, K., 1987. A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55 (3), 703-08.

[35] Okuda, T., Tsuruga, T., and Zanetti, F., 2019. Imperfect Information, Shock Heterogeneity, and Inflation Dynamics. Discussion Papers 1918, Centre for Macroeconomics (CFM).

[36] Phillips, A., 1958. The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957. *Economica* 25, 283–299.

[37] Roeger, W., and Herz, B., 2012. Traditional versus New Keynesian Phillips Curves: Evidence from Output Effects. *International Journal of Central Banking* 8 (1), 87-109.

[38] Rudd, J., and Whelan, K., 2006. Can Rational Expectations Sticky-Price Models Explain Inflation Dynamics? *American Economic Review* 96 (1), 303-320.

[39] Rudd, J., and Whelan, K., 2007. Modeling Inflation Dynamics: A Critical Review of Recent Research. *Journal of Money, Credit and Banking* 39 (s1), 155-170.

[40] Rudebusch, G., and Svensson, L., 1999. Policy Rules for Inflation Targeting. NBER Chapters in *Monetary Policy Rules*, edited by J. Taylor. University of Chicago Press.

[41] Seber, G., and Wild, C., 2003. *Nonlinear Regression*. Hoboken, NJ: Wiley-Interscience.

[42] Stock, J., and Watson, M., 1999. Forecasting inflation. *Journal of Monetary Economics* 44 (2), 293-335.

[43] Stock, J., and Watson, M., 2007. Why Has U.S. Inflation Become Harder to Forecast? *Journal of Money, Credit and Banking* 39 (s1), 3-33.

[44] Westerhoff, F., and Reitz, S., 2005. Commodity price dynamics and the nonlinear market impact of technical traders: empirical evidence for the US corn market. *Physica A: Statistical Mechanics and its Applications* 349 (3), 641-648.

[45] Zanetti, F., 2011. Labor market institutions and aggregate fluctuations in a search and matching model. *European Economic Review* 55 (5), 644-658.

[46] Zhang, C., Osborn, D., and Kim, D., 2008. The New Keynesian Phillips Curve: From Sticky Inflation to Sticky Prices. *Journal of Money, Credit and Banking* 40 (4), 667–699.

6 Figures

Figure 1: Expected inflation from the BNKPC model ($\hat{E}_t^b\{\pi_{t+1}\}$) and realized forward inflation (π_{t+1})

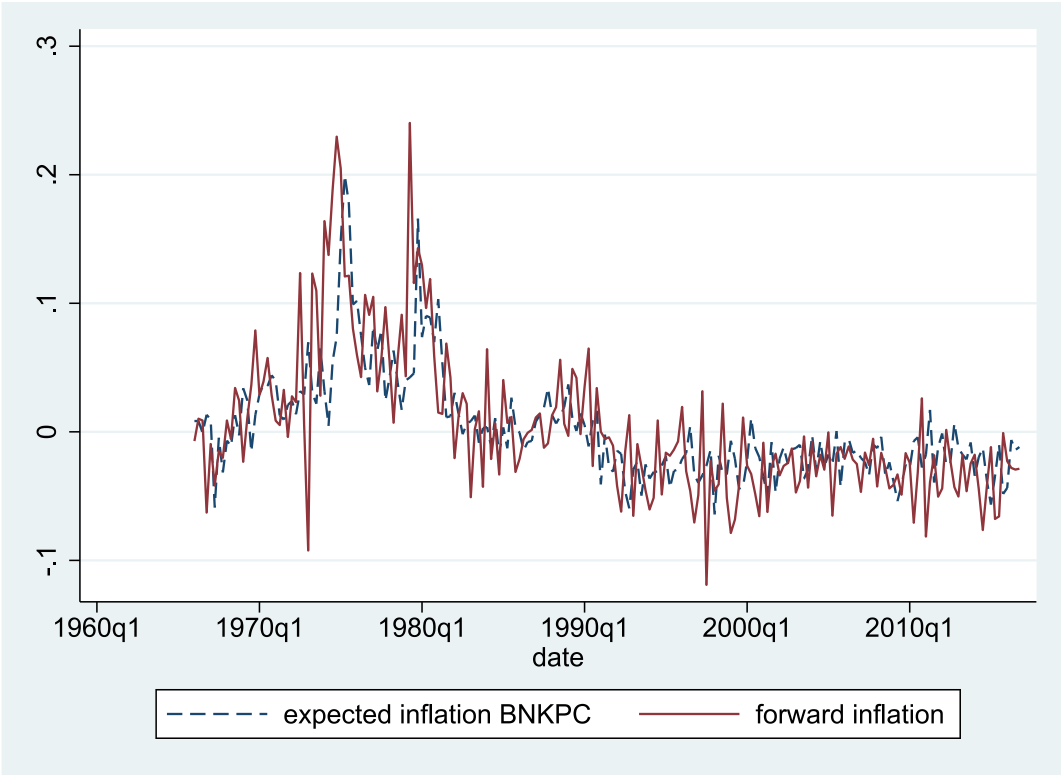

Figure 2: Difference between realized forward inflation and expected inflation

Figure 3: Fraction of individuals adopting the fundamental rule

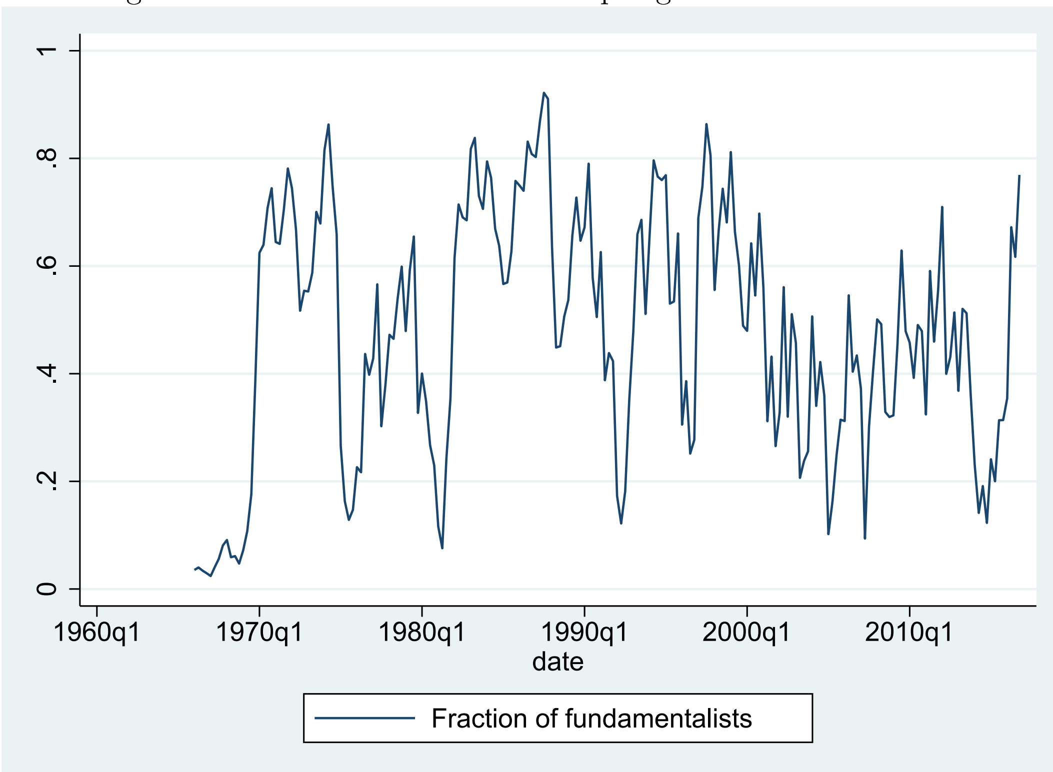

Figure 4: Giacomini and Rossi fluctuation test for horizon 1 (first plot in the first line), horizon 2 (second plot in the first line), until horizon 16 (last plot in the last line). TPC versus BNKPC

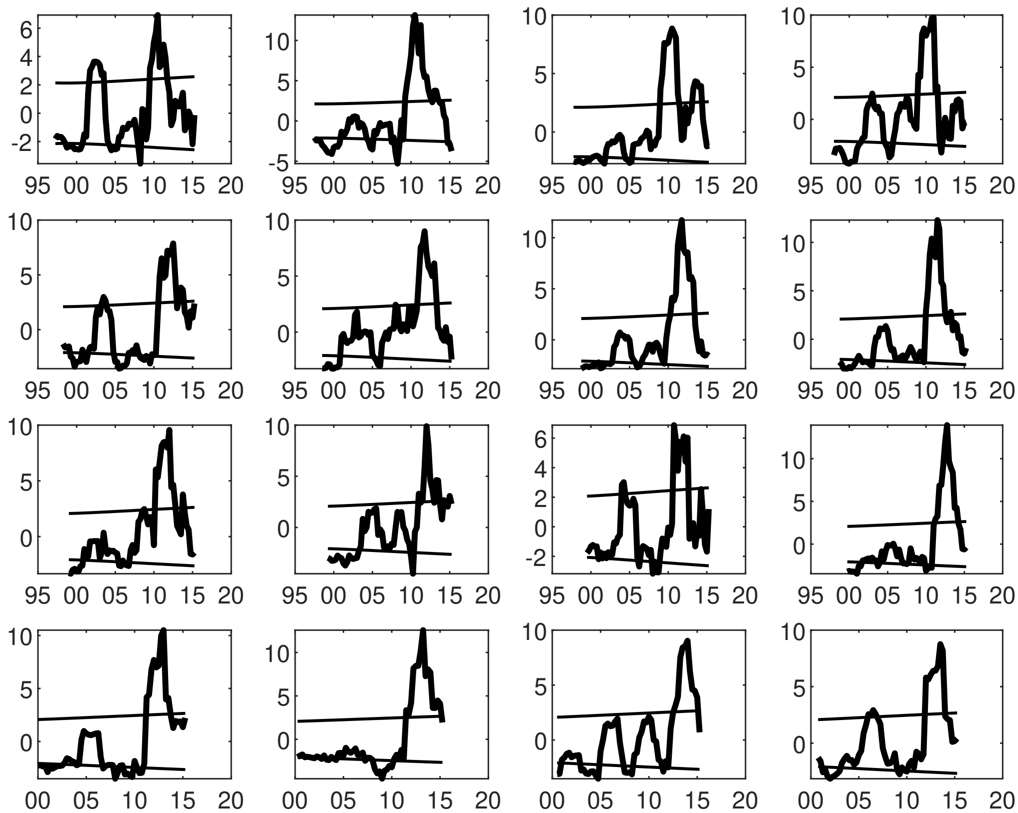

Figure 5: Giacomini and Rossi fluctuation test for horizon 1 (first plot in the first line), horizon 2 (second plot in the first line), until horizon 16 (last plot in the last line). NKPC versus BNKPC

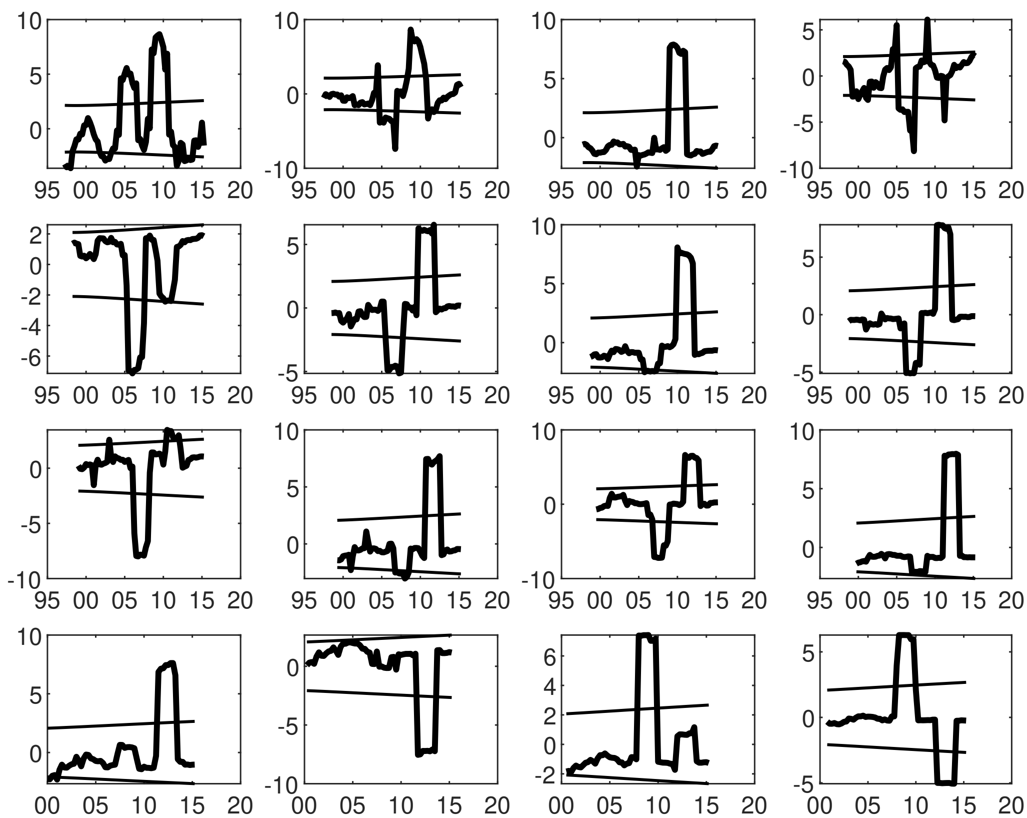

Figure 6: Giacomini and Rossi fluctuation test for horizon 1 (first plot in the first line), horizon 2 (second plot in the first line), until horizon 16 (last plot in the last line). HNKPC versus BNKPC

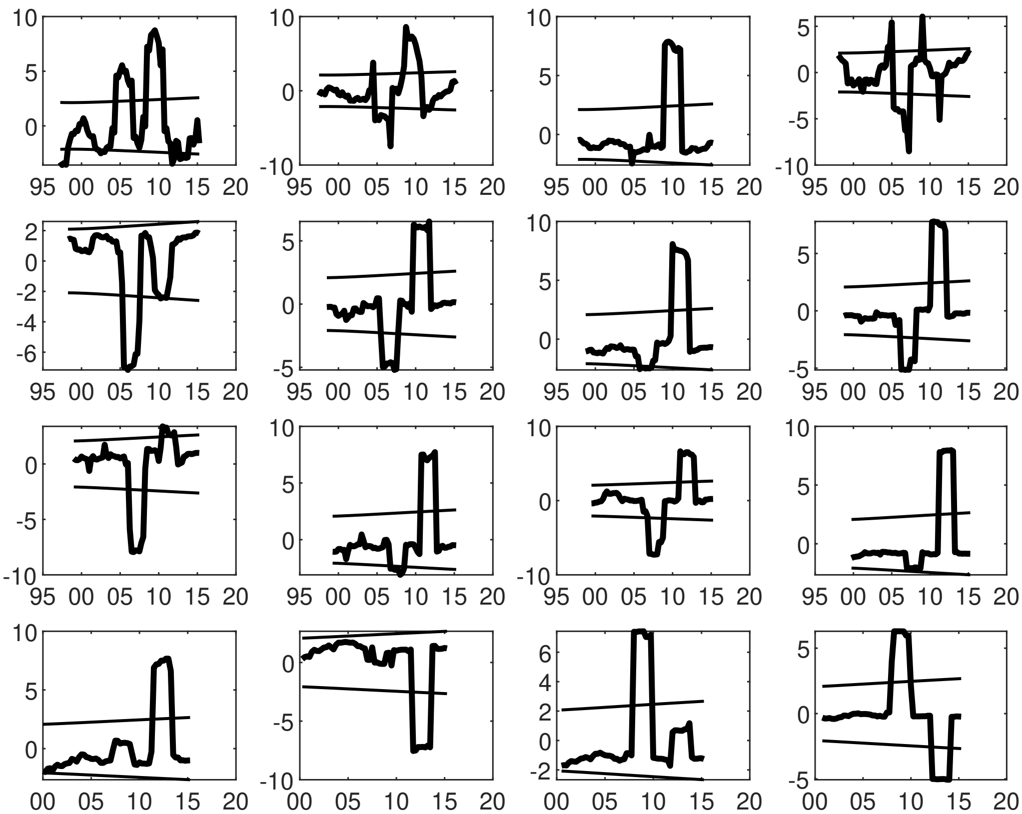

Figure 7: Giacomini and Rossi fluctuation test for horizon 1 (first plot in the first line), horizon 2 (second plot in the first line), until horizon 16 (last plot in the last line). AO versus BHNKPC

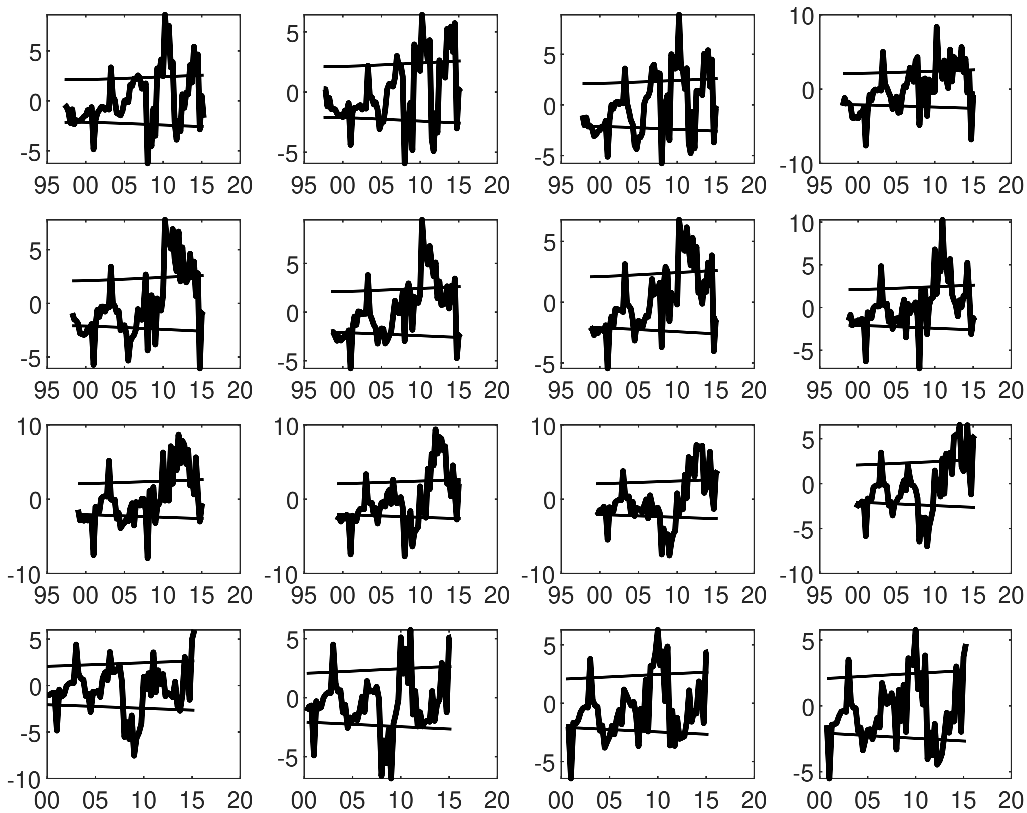

7 Tables

Table 1: NLS estimates of Phillips curve models for 1966:1–2016:4 data (baseline case)

	γ	θ	ϕ	BIC	AIC	Adj-R^2
TPC	0.148			-664	-668	0.345
	(0.161)					
NKPC	0.132*			-589	-592	0.052
	(0.068)					
HNKPC	0.158 **	0.344***		-698	-705	0.455
	(0.073)	(0.088)				
BNKPC	0.390***		4.142***	-712	-718	0.488
	(0.104)		(0.647)			

HAC standard errors are shown in parenthesis. *, **, *** denote significance at the 10%, 5%, and 1% level. In performing the estimation we set $\beta = 0.99$ and used a VAR to obtain a fundamental forecast of inflation (applying the Campbell and Shiller, 1987, method). The Phillips curve measure of inflation is obtained from the GDP deflator. The output gap measure is obtained from log of real GDP by applying the one-sided HP filter using the Kalman filter (Stock and Watson, 1999). In the regressions above the fitted value of the output gap from the VAR is used. The baseline VAR has 4 lags of the output gap and the labor share. The VAR is estimated from 1958:2. The absolute forecast error from (8) is based on $K = 4$.

Table 2: NLS Phillips curve models estimates for 1966:1–2016:4 data (robustness checks)

	Robustness check 1						Robustness check 2					
	γ	θ	ϕ	BIC	AIC	Adj-R^2	γ	θ	ϕ	BIC	AIC	Adj-R^2
TPC	0.032			-664	-668	0.344	-0.005			-664	-667	0.344
	(0.034)						(0.023)					
NKPC	0.024**			-595	- 598	0.079	0.004***			-605	-608	0.123
	(0.007)						(0.001)					
HNKPC	0.024***	0.354***		-699	-705	0.455	0.005***	0.375***		-703	-710	0.467
	(0.006)	(0.090)					(0.002)	(0.085)				
BNKPC	0.027***		3.371***	-708	-714	0.478	0.005***		3.714***	-715	-721	0.496
	(0.005)		(0.628)				(0.001)		(0.798)			

	Robustness check 3						Robustness check 4					
	γ	θ	ϕ	BIC	AIC	Adj-R^2	γ	θ	ϕ	BIC	AIC	Adj-R^2
TPC	0.000			-664	-667	0.344	0.037			-664	-668	0.344
	(0.024)						(0.037)					
NKPC	0.004**			-599	-603	0.099	0.006			-578	-581	-0.002
	(0.002)						(0.013)					
HNKPC	0.005***	0.365***		-702	-708	0.463	0.027**	0.323***		-697	-704	0.451
	(0.002)	(0.082)					(0.013)	(0.084)				
BNKPC	0.006***		3.732***	-713	-720	0.492	0.044***		4.143***	-710	-716	0.484
	(0.001)		(0.779)				(0.010)		(0.584)			

HAC standard errors are shown in parenthesis. *, **, *** denote significance at the 10%, 5%, and 1% level. In performing the estimation we set $\beta = 0.99$ and used a VAR to obtain a fundamental forecast of inflation (applying the Campbell and Shiller, 1987, method). Robustness check 1 has a VAR with 8 lags of the output gap and the labor share. Robustness check 2 has a VAR with 4 lags of the labor share (the driving variable in the VAR) and output gap. Robustness check 3 has a VAR with 8 lags of the labor share (the driving variable in the VAR) and output gap. Robustness check 4 has a VAR with 4 lags of the output gap, labor share and lagged inflation.

Table 3: Descriptive statistics of expectation errors for the fundamental and naive rules for 1966:1-2016:4

	fundamental rule	naive rule
Mean	-0.009	-0.000
Median	-0.005	-0.005
Maximum	0.217	0.202
Minimum	-0.580	-0.216
Standard Deviation	0.103	0.050
Skewness	-2.136	0.111
Kurtosis	12.326	5.372
First-order autocorrelation	0.804	-0.014
Correlation with output gap	-0.230	0.136

Table 4: Descriptive statistics of the weight on the fundamental rule

	1966:1-2016:4	1966:1-1992:3	1992:4-2016:4
Mean	0.476	0.487	0.465
Median	0.491	0.554	0.460
Maximum	0.922	0.922	0.864
Minimum	0.024	0.024	0.094
Standard Deviation	0.228	0.261	0.185
Skewness	-0.231	-0.393	0.125
Kurtosis	2.106	1.891	2.232
First-order autocorrelation	0.836	0.909	0.674
Correlation with output gap	0.292	0.362	0.160