**cef.up working paper
2020-01**

**FORECASTING INFLATION WITH THE NEW KEYNESIAN
PHILLIPS CURVE: FREQUENCY MATTERS**

**Manuel M. F. Martins
Fabio Verona**

Forecasting inflation with the New Keynesian Phillips Curve: frequency matters*

Manuel M. F. Martins[†] Fabio Verona[‡]

18 April 2020

Abstract

We show that the New Keynesian Phillips Curve (NKPC) outperforms standard benchmarks in forecasting U.S. inflation once frequency-domain information is taken into account. We do so by decomposing the time series (of inflation and its predictors) into several frequency bands and forecasting separately each frequency component of inflation. The largest statistically significant forecasting gains are achieved with a model that forecasts the lowest frequency component of inflation (corresponding to cycles longer than 16 years) flexibly using information from all frequency components of the NKPC inflation predictors. Its performance is particularly good in the returning to recovery from the Great Recession.

Keywords: inflation forecasting, new Keynesian Phillips curve, frequency domain, wavelets

JEL codes: C53, E31, E37

* Manuel M. F. Martins gratefully acknowledges the hospitality of the Research Unit of the Bank of Finland at the outset of this research project. We thank Fabio Canova for useful comments. The views expressed are those of the authors and do not necessarily reflect those of the Bank of Finland. Any remaining errors are the sole responsibility of the authors. Research at *CEF.UP* has been financed by Portuguese public funds through FCT -- Fundação para a Ciência e a Tecnologia, I.P., in the framework of the project with reference UIDB/04105/2020.

[†] University of Porto, Faculty of Economics and *CEF.UP* (mmfmartins@fep.up.pt).

[‡] Bank of Finland – Monetary Policy and Research Department, and University of Porto – *CEF.UP* (fabio.verona@bof.fi)

1 Introduction

The expectations-augmented Phillips curve has long been the key model for explaining inflation in macroeconomic analyses and models. However, irrespective of its many specifications over time, it has overall performed poorly in forecasting inflation out-of-sample, typically failing to beat simple time series models (see e.g. Canova, 2007, Stock and Watson, 2008, Faust and Wright, 2013, and Berge, 2018). Currently, the Phillips curve is not considered a useful model for inflation forecasting. In this paper we establish the ability of an empirical state-of-the-art New Keynesian Phillips Curve (NKPC) to forecast inflation, once frequency-domain information is taken into account.

While there is a large literature on the variation of Phillips curves coefficients over time (see e.g. the surveys by Faust and Wright, 2013 and Mavroeidis, Plagborg-Moller, and Stock, 2014), changes of Phillips curves coefficients across frequencies is much less explored. Yet, a few studies have detected time and frequency dependent coefficients in Phillips curves (e.g. Gallegati, Gallegati, Ramsey, and Semmler, 2011 and Aguiar-Conraria, Martins, and Soares, 2019). Moreover, frequency dependence allows for modelling Phillips curve nonlinearities (e.g. Ashley and Verbrugge, 2009), which have theoretical motivations and supportive literature (see e.g. the survey in Aguiar-Conraria, Martins, and Soares, 2019). Finally, the data indicate frequency variation of NKPC coefficients (see section 2.1). This paper is the first to forecast inflation with a Phillips curve that features frequency-dependent coefficients.

Our approach builds on a literature that improves forecasts of a variable of interest by forecasting its different frequency components rather than the aggregate (see e.g. Faria and Verona, 2018). In the case of inflation, there are at least two reasons why such approach is useful. First, it seems crucial to forecast well the low-frequency components of inflation. On the one hand, as it is well established that good inflation forecasts must account for a slowly varying local mean for inflation (Faust and Wright, 2013). On the other hand, as about two thirds of the variance of inflation is due to its low-frequency fluctuations (see Table 2). Second, it seems important to take into

account information from other frequencies of the predictors while forecasting the low-frequency component of inflation. This is because some predictors of inflation – proxies of slack and of cost-push shocks – have large variance at medium and high frequencies that surely have important explanatory power for inflation.

Our method consists of decomposing the time series of the NKPC into several frequency bands with the maximal overlap discrete wavelet transform, forecasting separately each frequency component of inflation, and then summing up those forecasts to obtain the forecast for aggregate inflation. In our approach, each frequency component of inflation may depend on other frequency components of the NKPC inflation predictors. The flexibility of allowing for – but not imposing – this frequency interaction is a key feature of our approach, given our interest in producing forecasts for the low-frequency component of inflation but not losing useful information from expectations, slack, or supply shocks at high or medium frequencies.

Our central result is that the larger statistically significant forecasting gains (for both 4-quarter and 8-quarter ahead horizons) are achieved with a model that forecasts the lowest frequency component of inflation (corresponding to cycles above 16 years) flexibly using information from all frequency components of the NKPC inflation predictors. Its performance is particularly good in the turning from the Great Recession to its recovery. Our results are consistent with those of an empirical literature focusing on the relevance of a slowly time-varying trend inflation (Chan, Clark, and Koop, 2018 and Hasenzagl, Pellegrino, Reichlin, and Ricco, 2018). They also support recent findings from theoretical Dynamic Stochastic General Equilibrium (DSGE) literature about the importance of modelling well the low-frequency movements of inflation (Del Negro, Giannoni, and Schorfheide, 2015) as well as of capturing the interactions between macroeconomic variables at different frequencies (Comin and Gertler, 2006 and Beaudry, Galizia, and Portier, 2020).

The remaining of this paper is organized as follows. In section 2 we ground our paper in the two literatures to which it is related – the specification of the Phillips curve and the frequency-domain approach to forecasting. In section 3 we present the data. In section 4 we describe the forecasting

methods. In section 5 we present the results. Section 6 concludes.

2 Motivation

This paper is related with two literatures – one on the specification of the New Keynesian Phillips curve, and the other on the frequency-domain approach to forecasting. In this section we briefly locate our paper within these two literatures.

2.1 The New Keynesian Phillips curve

Empirically, the history of the Phillips curve is one of “seemingly stable relationships falling apart upon publication” (Stock and Watson, 2010), and one of strong specification and sampling uncertainty (Mavroeidis, Plagborg-Moller, and Stock, 2014).

In this paper we use a state-of-the-art empirical NKPC, in the spirit of Coibion and Gorodnichenko (2015), Fuhrer (2017), and Coibion, Gorodnichenko, and Kamdar (2018), that explains inflation (π_t) with households survey expectations of inflation (π_{t+1}^e), the unemployment gap ($ugap_t$), and energy inflation (en_t) as control for supply shocks:

$$\pi_t = c + \alpha_1 \pi_{t+1}^e + \alpha_2 ugap_t + \alpha_3 en_t + \varepsilon_t . \quad (1)$$

Households survey inflation expectations, given by the Michigan survey of consumers, became the state-of-the-art in recent empirical NKPC for several reasons, both theoretical and empirical. Most notably, it has been shown that such expectations are the closest the possible to firms’ expectations – which, according to theory, are the relevant inflation expectations in the NKPC (see Coibion and Gorodnichenko, 2015, Coibion, Gorodnichenko, and Kumar, 2018, and Pfajfar and Roberts, 2018). Another key property of the households’ Michigan survey of inflation expectations is their

intrinsic inertia, due to the micro-founded inefficiency with which agents revise expectations, thus avoiding the need of any ad-hoc inertial mechanisms (Fuhrer, 2017 and Coibion, Gorodnichenko, and Kamdar, 2018).¹

In this paper we are interested in forecasting inflation with the NKPC. The literature of inflation forecasting based on Phillips curves has mostly used some version of its accelerationist specification (see e.g. Stock and Watson, 1999, Canova, 2007 and Dotsey, Fujita, and Stark, 2018). The notable exception is Berge (2018), who forecasts U.S. inflation using several alternative Phillips curves, including some versions of the NKPC akin to ours in that expectations of inflation are taken from the Michigan survey of consumers. The key difference between our paper and Berge (2018) is that we explore frequency-domain information, which is beyond his purely time series approach and, indeed, has never been explored before in the literature of inflation forecasting.²

While (1) stands as a constant coefficients relation, there is a large literature on the time variation of Phillips curve coefficients – especially its slope (α_2 in (1)) – either due to abrupt or gradual structural breaks, or to nonlinearities (see e.g. Stock and Watson, 2007, 2008 and Dotsey, Fujita, and Stark, 2018). In the context of forecasting, time variation of the Phillips curve coefficients arises naturally, as out-of-sample forecasts are produced recursively using either rolling or expanding window methods. Frequency dependence, in contrast, does not arise naturally.

The literature is relatively scarce as regards changes of Phillips curves coefficients across frequencies. Yet, recent research has emphasized that Phillips curves coefficients may change both over time and across frequencies (e.g. Gallegati, Gallegati, Ramsey, and Semmler, 2011 and Aguiar-Conraria, Martins, and Soares, 2019). In particular, Aguiar-Conraria, Martins, and Soares (2019) show that, for U.S. data, there is considerable variation of coefficients of the NKPC (1) along frequencies, in addition to over time.

¹ See Aguiar-Conraria, Martins, and Soares (2019) for more details on the foundations and empirical advantages of using the Michigan households survey of expected inflation and, overall, a more complete discussion of the empirical specification of a state-of-the-art empirical NKPC.

² Another (less relevant) difference is that we also include a proxy for supply shocks.

Furthermore, the data clearly indicate frequency variation of the NKPC coefficients, as can be seen in the in-sample estimates of (1) reported in Table 1 and Figure 1. The first row of Table 1 shows standard time series estimates of (1) that are consistent with the literature. The remaining rows present estimates of (1) for 3 different frequency bands – high frequency (HF, cycles of period between 2 and 8 quarters), business cycles (BCF, cycles of period between 2 and 8 years), and low frequency (LF, cycles longer than 8 years). Overall, the estimates of most coefficients are substantially different across cyclical frequencies.³

Figure 1 focuses on the slope of the NKPC, depicting estimates from expanding windows that start with the sample 1978Q1-1999Q4 and recursively add one quarter through 2019Q4. The upper left graph shows some variation over time of the NKPC slope, especially in the latter stages of the Great Recession and its recovery. The remaining graphs clearly show that the variation of the NKPC slope over time differs substantially across frequencies, and that statistical significance (dotted points) is much more pervasive once the NKPC is estimated separately for each frequency. The striking evidence of in-sample frequency dependence suggests that it may be relevant to consider frequency variation of the NKPC coefficients also in out-of-sample forecasting. Frequency dependence has never been explored so far for the purpose of forecasting inflation with the NKPC. Hence the major contribution of this paper.

2.2 Frequency-domain forecasts

Our approach builds on the literature that uses discrete wavelet methods to forecast out-of-sample economic and financial time series. Examples include Rua (2011, 2017), who forecast GDP growth and inflation using a factor-augmented wavelets approach; Zhang, Gençay, and Yazgan (2017) and Faria and Verona (2017, 2018, 2020b), who focus on forecasting stock market returns; Caraiani

³ Details on our data and on our econometric approach are presented in section 3. Specifically, the HF component is the sum of frequencies D_1 and D_2 , the BCF component is the sum of frequencies D_3 and D_4 , and the LF component is the sum of frequencies D_5 and D_6 . Gallegati, Gallegati, Ramsey, and Semmler (2011) found a similar result for the wage Phillips curve.

(2017), who forecasts exchange rates; and Faria and Verona (2020a), who forecast the bond risk premium and the equity risk premium.

The wavelet method used in this paper is known as wavelet multiresolution analysis. It allows for decomposing any time series into a trend (or permanent) component and cyclical (or transitory) movements in a way that is similar to the traditional time series trend-cycle decomposition approach (e.g. Beveridge and Nelson, 1981), or other filtering methods like the Baxter and King (1999) bandpass filter or the Hodrick and Prescott (1997) filter. In particular, the wavelet multiresolution decomposition is additive and therefore allows for extracting, and subsequently forecasting separately, each frequency component of the time series of interest.⁴

Our paper improves upon Faria and Verona (2018). In that paper, each frequency component of the variable to be forecasted depends only on the corresponding frequency component of the predictor variables. In our paper we also allow for each frequency component of inflation to depend on other frequency components of the predictors. In particular, we allow for the business-cycle frequencies or medium-term frequencies of, say, the unemployment gap, to affect low-frequency fluctuations of inflation. Such generalization of the wavelet-based forecast approach is particularly important in the case of inflation, as argued above and confirmed below by our results.

3 Data

Our data are U.S. quarterly time series for 1978Q1-2019Q4 of inflation, the unemployment gap, expectations of inflation, and energy inflation. Inflation is the annualized quarterly rate of growth of the consumer price index (CPI) provided by the U.S. Bureau of Labor Statistics, and energy inflation is the annualized quarterly rate of growth of the respective component of the CPI. Inflation expectations are the median expected changes in prices on average during the next 12 months reported by households in the Michigan survey of consumers (MSC). The unemployment gap is

⁴ See Verona (2020) for a description of the advantages of wavelet filters over other band-pass filtering techniques.

the difference between the (quarterly average of the) civilian unemployment rate provided by the U.S. Bureau of Labor Statistics and a linear trend.

Let P_t be the price index in quarter t . The quarterly rate of inflation at an annual rate is computed as $\pi_t = 400 \ln(P_t/P_{t-1})$, and the h -period inflation as $\pi_t^h = \frac{1}{h} \sum_{i=0}^{h-1} \pi_{t-i}$. In this paper we focus on two forecasting horizons – 4-quarters ($h=4$) and 8-quarters ahead ($h=8$) – that are arguably the most relevant for policymakers.

Our method to forecast inflation relies on wavelet filtering methods to filter the data. Using the wavelet multiresolution analysis (MRA) it is possible to decompose a time series into its constituent multiresolution (frequency) components. Given a time series y_t , its wavelet multiresolution representation is:

$$y_t = \sum_{j=1}^J y_t^{D_j} + y_t^{S_J} , \quad (2)$$

where $y_t^{D_j}$, $j = 1, 2, \dots, J$, are the J wavelet detail components and $y_t^{S_J}$ is the wavelet smooth component. Expression (2) shows that the original time series y_t can be decomposed in different components, each of them representing the fluctuations of the original variable within a specific frequency band. In particular, for small j , the j wavelet detail components represent the higher frequency fluctuations of the time series (*i.e.* its short-term dynamics). As j increases, the j wavelet detail components represent lower frequencies movements of the series. Finally, the wavelet smooth component captures the lowest frequency fluctuations (*i.e.* its trend).⁵

In this paper, we perform a wavelet decomposition analysis by applying the maximal overlap discrete wavelet transform (MODWT) multiresolution analysis. We apply a $J=5$ level MODWT MRA to our time series using the Haar wavelet filter.⁶ As we use quarterly data, the first component

⁵ A more detailed analysis of wavelets methods can be found in Percival and Walden (2000).

⁶ Examples of papers using the MODWT MRA decomposition include Gallegati, Gallegati, Ramsey, and Semmler (2011), Crowley and Hughes Hallett (2015), Berger (2016), and Gallegati, Giri, and Palestini (2019), among others. The Haar filter is simple, widely used, and makes a neat connection to temporal aggregation as the wavelet

(D₁) captures fluctuations with a period between 2 and 4 quarters, while the components D₂, D₃, D₄ and D₅ capture fluctuations with a period of 1-2, 2-4, 4-8 and 8-16 years, respectively. Finally, the smooth component S₅, which in what follows we re-denote D₆, captures fluctuations with a period longer than 16 years.⁷

Figures 2-5 plot the time series of the variables (top pictures) and of their frequency components (middle and bottom pictures). These figures illustrate that the original time series are the result of the aggregation of several underlying frequency components that exhibit very different dynamics. In particular, and as expected, the lower the frequency, the smoother the resulting filtered time series.

Table 2 reports the variance decomposition by frequency of all variables of interest – 4-quarter inflation, 8-quarter inflation, and the predictors. Almost two thirds of the variance of both inflation time series occurs at the D₆ frequency band, and almost all their variance occurs at the three lower frequency bands (D₄ to D₆). The MSC inflation expectations feature a variance frequency decomposition quite similar to that of the 4-quarter inflation, as expected. The unemployment gap exhibits a more even distribution of its variance across the three lower frequency bands. Energy inflation, in contrast, has almost all its variance in the three higher frequency bands.

4 Forecasting methods

Our out of sample (OOS) forecasts are direct forecasts produced with a sequence of expanding windows. We use an initial sample (1978Q1-1999Q4) to make the first OOS forecast. The sample is then increased by one observation and a new OOS forecast is produced. This procedure is repeated until the end of the sample. Hence, the full OOS period runs from 2000Q1 to 2019Q4.

coefficients are differences of moving averages (see Bandi, Perron, Tamoni, and Tebaldi, 2019 and Lubik, Matthes, and Verona, 2019).

⁷ In the MODWT, each wavelet component at frequency j approximates an ideal high-pass filter with passband $f \in [1/2^{j+1}, 1/2^j]$, hence they are associated to fluctuations with periodicity $[2^j, 2^{j+1}]$ (quarters, in our case).

We thus have 77 observations of 4-quarter-ahead forecasts and 73 of 8-quarter-ahead forecasts.

4.1 Some traditional forecasting models

Following Atkeson and Ohanian (2001), the AO random walk model has typically been considered the benchmark model for inflation forecasting. With this model, the forecasts of h-quarter-ahead inflation are:

$$\hat{\pi}_{t+h}^h = \pi_{t-1}^h ,$$

where π_{t-1}^h is the previous period's h-quarter average inflation rate.

In addition to the AO model, the literature of inflation forecasting typically considers two simple univariate time series models – the AR(p) model and the ARMA(1,1) model.

The AR(p) model consists of:

$$\pi_t^h = \alpha + \varphi_1 \pi_{t-1}^h + \varphi_2 \pi_{t-2}^h + \dots + \varphi_p \pi_{t-p}^h + \varepsilon_t ,$$

where p denotes the lag-length for the autoregressive process.⁸ The forecasts of h-quarter-ahead inflation produced by the AR(p) model are given by:

$$\hat{\pi}_{t+h}^h = \hat{\alpha} + \hat{\varphi}_1 \pi_{t-1}^h + \hat{\varphi}_2 \pi_{t-2}^h + \dots + \hat{\varphi}_p \pi_{t-p}^h .$$

The ARMA(1,1) model consists of:

$$\pi_t = \alpha + \varphi \pi_{t-1} + \phi \varepsilon_{t-1} + \varepsilon_t ,$$

⁸ Typically, an information criterion such as the AIC is used to select the lag length; we follow that criterion in this paper, allowing for a maximum of six possible lags.

and its forecasts of h-quarter-ahead inflation are given by:

$$\hat{\pi}_{t+h}^h = \hat{\alpha} \left(1 - \hat{\phi}\right) + \left(\hat{\varphi} + \hat{\phi}\right) \pi_t^h - \hat{\phi}\hat{\varphi}\pi_{t-1}^h .$$

Besides simple univariate time series models, it has been increasingly standard in the forecasting literature to consider the expectations of inflation stated in surveys. In the case of our paper, this is particularly relevant, as our NKPC includes inflation expectations given by a survey of households, namely the MSC. The survey-based forecasts are given by $\hat{\pi}_{t+h}^h = \pi_{t+1}^e$, where π_{t+1}^e is the MSC median expected change in prices during the next 12 months. Given that the MSC does not elicit the 8-quarter-ahead forecasts, we use their 4-quarter-ahead inflation expectations as forecasts of both 4-quarter and 8-quarters ahead inflation. This model is denoted Survey.

Finally, our NKPC (1) combines the unemployment gap ($ugap_t$) with MSC expected inflation (π_{t+1}^e) and energy inflation as a proxy for supply shocks (en_t). At each step of the OOS period, for each h we first estimate a regression like:

$$\pi_t^h = c^h + \alpha_1^h \pi_{t+1}^e + \alpha_2^h ugap_t + \alpha_3^h en_t + \varepsilon_{t+h}^h , \quad (3)$$

and then compute the forecasts as:

$$\hat{\pi}_{t+h}^h = \hat{c}^h + \hat{\alpha}_1^h \pi_{t+1}^e + \hat{\alpha}_2^h ugap_t + \hat{\alpha}_3^h en_t . \quad (4)$$

It is important to note that, at each step t of the expanding window forecast, we compute the unemployment gap (needed to to estimate 3 and, then, to compute forecasts with 4) by fitting a linear trend to the unemployment rate data from the beginning of the sample through quarter t. Such procedure assures that there is no “look-ahead” bias in the predictive regression forecast based on the NKPC. This model is denoted NKPC_TS.

4.2 NKPC-wavelet forecasting method

Our NKPC-wavelet forecast method builds on the NKPC (1) and on the filtered data obtained with the MODWT multiresolution decomposition. Importantly, as the MODWT MRA is a two-sided filter, we compute it recursively at each iteration of the OOS forecasting process by using data from the start of the sample through the quarter of forecast formation. This ensures that our method does not have a “look-ahead” bias as the forecasts are made with current and past information only. We deal with boundary effect using a reflection rule, i.e. extending the time series symmetrically at the right boundary before computing the MODWT MRA.

For each frequency component D_j , $j=1,\dots,6$, we have a Phillips curve such as:

$$\pi_{t+h}^{h,D_j} = c^{h,j} + \alpha_1^{h,j} \pi_{t+1}^{e,D_j} + \alpha_2^{h,j} ugap_t^{D_j} + \alpha_3^{h,j} en_t^{D_j} + \varepsilon_{t+h}^{h,j} . \quad (5)$$

This specification implies that each frequency of inflation π_{t+h}^{h,D_j} , $j=1,\dots,6$, depends only on the frequencies of the predictors at the same frequency j . However, it is of interest to analyze the case when some of the other frequencies of the predictors, D_k with $k \neq j$, are useful to forecast π_{t+h}^{h,D_j} . This would be the case if there are important interactions across frequencies of inflation and frequencies of predictors.

More formally and generally, at each step of the OOS period, we first estimate the following system of equations:

$$\begin{bmatrix} \pi_t^{h,D_1} \\ \pi_t^{h,D_2} \\ \pi_t^{h,D_3} \\ \pi_t^{h,D_4} \\ \pi_t^{h,D_5} \\ \pi_t^{h,D_6} \end{bmatrix} = \begin{bmatrix} c^{h,D_1} \\ c^{h,D_2} \\ c^{h,D_3} \\ c^{h,D_4} \\ c^{h,D_5} \\ c^{h,D_6} \end{bmatrix} + \alpha_1^h \begin{bmatrix} \pi_{t+1}^{e,D_1} \\ \pi_{t+1}^{e,D_2} \\ \pi_{t+1}^{e,D_3} \\ \pi_{t+1}^{e,D_4} \\ \pi_{t+1}^{e,D_5} \\ \pi_{t+1}^{e,D_6} \end{bmatrix} + \alpha_2^h \begin{bmatrix} ugap_t^{D_1} \\ ugap_t^{D_2} \\ ugap_t^{D_3} \\ ugap_t^{D_4} \\ ugap_t^{D_5} \\ ugap_t^{D_6} \end{bmatrix} + \alpha_3^h \begin{bmatrix} en_t^{D_1} \\ en_t^{D_2} \\ en_t^{D_3} \\ en_t^{D_4} \\ en_t^{D_5} \\ en_t^{D_6} \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,t}^{h,D_1} \\ \varepsilon_{2,t}^{h,D_2} \\ \varepsilon_{3,t}^{h,D_3} \\ \varepsilon_{3,t}^{h,D_4} \\ \varepsilon_{3,t}^{h,D_5} \\ \varepsilon_{3,t}^{h,D_6} \end{bmatrix} , \quad (6)$$

where α_1^h , α_2^h and α_3^h are 6x6 matrices.⁹ We consider two cases of this wavelet-based model.

In a first case, matrices α_m^h , with $m = 1, 2, 3$, are diagonal, that is, we assume that only the frequency components of the predictors at level D_j are used to forecast the frequency component of inflation at the same level. We denote this model as NKPC_WAV_diag.

In a second case, we allow for interactions between predictors and inflation across frequencies, and therefore all the coefficients in matrices α_m^h , $m = 1, 2, 3$, are allowed to be different from 0. We denote this model as NKPC_WAV_all.

The choice between NKPC_WAV_diag and NKPC_WAV_all is ultimately empirical. NKPC_WAV_all is more general, as it includes 18 predictors, and should lead to better in-sample fit. However, it remains to be assessed whether relations between predictors and inflation across different frequencies are empirically that relevant that the improved in-sample fit does not harm the OOS performance. NKPC_WAV_diag, in turn, is more parsimonious and may achieve better OOS performance. Finding which specification predicts inflation more effectively and robustly, for our U.S. data, is a key contribution of our approach.

The forecasts of each frequency component of inflation are then computed as:

$$\begin{bmatrix} \hat{\pi}_{t+h}^{h,D_1} \\ \hat{\pi}_{t+h}^{h,D_2} \\ \hat{\pi}_{t+h}^{h,D_3} \\ \hat{\pi}_{t+h}^{h,D_4} \\ \hat{\pi}_{t+h}^{h,D_5} \\ \hat{\pi}_{t+h}^{h,D_6} \end{bmatrix} = \begin{bmatrix} \hat{c}^{h,D_1} \\ \hat{c}^{h,D_2} \\ \hat{c}^{h,D_3} \\ \hat{c}^{h,D_4} \\ \hat{c}^{h,D_5} \\ \hat{c}^{h,D_6} \end{bmatrix} + \hat{\alpha}_1^h \begin{bmatrix} \pi_{t+1}^{e,D_1} \\ \pi_{t+1}^{e,D_2} \\ \pi_{t+1}^{e,D_3} \\ \pi_{t+1}^{e,D_4} \\ \pi_{t+1}^{e,D_5} \\ \pi_{t+1}^{e,D_6} \end{bmatrix} + \hat{\alpha}_2^h \begin{bmatrix} ugap_t^{D_1} \\ ugap_t^{D_2} \\ ugap_t^{D_3} \\ ugap_t^{D_4} \\ ugap_t^{D_5} \\ ugap_t^{D_6} \end{bmatrix} + \hat{\alpha}_3^h \begin{bmatrix} en_t^{D_1} \\ en_t^{D_2} \\ en_t^{D_3} \\ en_t^{D_4} \\ en_t^{D_5} \\ en_t^{D_6} \end{bmatrix}. \quad (7)$$

Finally, given that $\pi_{t+h}^h = \sum_{j=1}^J \pi_{t+h}^{h,D_j}$, then the h-quarter-ahead inflation forecast is given by the

⁹ As for the forecast with the NKPC_TS model, the unemployment gap is recomputed at each step of the OOS period by fitting a linear trend to the data up to the quarter when the forecast is made.

sum of the h-quarter ahead forecasts given by each frequency component of the NKPC:

$$\hat{\pi}_{t+h}^h = \sum_{j=1}^J \hat{\pi}_{t+h}^{h,D_j} . \quad (8)$$

In the context of forecasting stock market returns, Faria and Verona (2018) showed that forecasts of the variable of interest may be improved if the forecasts corresponding to some specific frequencies are ignored. In particular, they found that disregarding the forecasts of high-to-medium frequency components, and only retaining the forecasts for lower frequencies, usually results in an improved forecasting performance. In our context, this means that the h-quarter-ahead inflation forecast given by

$$\hat{\pi}_{t+h}^h = \sum_{j=1}^J \kappa_j \hat{\pi}_{t+h}^{h,D_j} , \quad (9)$$

where κ_j could be 0 for some values of j , may be better than the case when we include all the frequencies of inflation (equation 8).

In principle, one could optimize the forecast, grid-searching the weights for each frequency that minimize the RMSFE for the entire OOS period. However, such procedure would not be implementable in real-time. Instead, a simple and robust rule of constant weighting of frequencies is feasible, and fair to use. As discussed above, it turns out that the very simple rule of considering only forecasts for frequency D_6 has strong motivations in the literature of wavelet-based forecasts and of inflation forecasts. Moreover, it has also motivation in the data: as shown in Table 2, almost two thirds of the variance of the two inflation time series that we want to forecast occurs at the D_6 frequency band.

Therefore, in addition to the more general models NKPC_WAV_diag and NKPC_WAV_all, we consider the case in which $\hat{\pi}_{t+h}^h = \hat{\pi}_{t+h}^{h,D_6}$, that is, $\kappa_j = 0 \forall j = 1, \dots, 5$ in (9). In short, we test whether the forecast of inflation may be improved by using only the forecast of its lowest frequency component. As in the case of the general model, we consider a model strictly based

on the forecasts from the D_6 frequency of predictors – which we denote NKPC_WAV_diag (only D_6) – but also consider a model in which the whole elements of the last row of matrices α_m^h , with $m = 1, 2, 3$, are allowed to be non null – which we denote NKPC_WAV_all (only D_6). This model is richer, as it allows for some influence from the higher-frequency fluctuations of the predictors into the low frequency of inflation – which Table 2 suggests that may be especially useful in the case of the unemployment gap and of energy inflation. Hopefully, it remains parsimonious enough to effectively improve the OOS forecasts of inflation.

5 Results

As common in the literature, we use the root mean squared forecast errors (RMSFE) as the indicator of forecast accuracy. We report our results in Table 3. Panel (a) reports the RMSFE of the AO random walk model, the usual benchmark for inflation forecasts. In panels (b) and (c) we report, for each model j , its RMSFE relative to that of the AO model, computed as $\text{RMSFE}_j / \text{RMSFE}_{\text{AO}}$, for the traditional and the wavelet-based models, respectively. A value less than one indicates that model j outperforms the benchmark. As we are also interested in assessing the performance of the NKPC-wavelet forecasting method relatively to the NKPC_TS forecasts, we report, in panel (d), the RMSFEs of the NKPC-wavelet models relative to those of the NKPC_TS model, computed as $\text{RMSFE}_j / \text{RMSFE}_{\text{NKPC-TS}}$, for model j . Asterisks indicate statistical significance of the Diebold and Mariano (1995) test of relative predictive accuracy at the 10% (*), 5% (**), and 1% (***) levels.

Panel (b) shows that the traditional models typically fail to outperform the AO benchmark model, sometimes by a substantial margin. In fact, none of the model produces a below one and statistically significant relative RMSFE at both forecasting horizons.

As regards the relative forecasting performance of our NKPC-wavelet forecasts, the two first lines of panels (c) and (d) indicate that at both the 4 and 8-quarter horizons, the forecasts of our NKPC-

wavelet models usually outperform those of the benchmark as well as those of the NKPC_TS model. In particular, the NKPC_WAV_diag model produces forecasts that are statistically better than those of the AO (NKPC_TS) model at the 4-quarter (8-quarter) horizon. Our results further indicate that there is a loss of forecasting accuracy when interactions between different frequencies of inflation and of its predictors are allowed for, as the forecasts obtained from the NKPC_WAV_all model have higher RMSFEs than those of the NKPC_WAV_diag model at both horizons.

The third and fourth lines of panels (c) and (d) report the key results in this paper. Consistently with our conjecture, focusing only on the low-frequency component of inflation (D_6), which corresponds to cycles longer than 16 years, produces forecasts for inflation that are markedly and statistically more accurate than those of both the AO and the NKPC_TS model.

The RMSFEs of the NKPC_WAV_diag (only D_6) model are 76% and 78% of those of the AO benchmark, at the 4-quarter and 8-quarter horizon, respectively, which are statistically significant at the 5% level. Moreover, they are 74% and 54% of those of the NKPC_TS at the 4-quarter and 8-quarter horizon, respectively, both statistically significant.

The NKPC_WAV_all (only D_6) model produces even better forecasts than the NKPC_WAV_diag (only D_6), meaning that NKPC-wavelet forecasts of the lower frequency of inflation are improved by allowing (but not imposing) for influences from other frequencies of the predictors (namely, medium and high frequencies). The RMSFEs from the NKPC_WAV_all (only D_6) model are 74% and 77% of those of the AO benchmark, at the 4-quarter and 8-quarter horizon, respectively (significant at the 5% level); and they are 72% and 53% of those of the time series NKPC at the 4-quarter and 8-quarter horizon, respectively.

The latter is a particularly large gain in accuracy, and highly relevant as it occurs at the horizon (two year) that is typically the relevant horizon for monetary policymakers. Furthermore, this result is quantitatively unmatched by any paper of the literature of Phillips curve-based forecasts

of inflation – either empirical NKPCs as in Berge (2018), accelerationist Phillips curves as in e.g. Dotsey, Fujita, and Stark (2018), or Phillips curves decomposed into components according to its cyclical sensitivity as in e.g. Tallman and Zaman (2017).

Having established that the NKPC_WAV_all (only D_6) model significantly outperforms both the benchmark and the time series NKPC, we now explore the timing of such outperformance. To do so, in Figure 6 we report the cumulative differences between the squared forecast errors (SFE) of the NKPC_WAV_all (only D_6) model and those of the AO and the NKPC_TS during the entire OOS period. The pictures on the left side report the cumulative differences of SFE relative to the AO model for 4-quarters-ahead (in the top) and 8-quarters-ahead (in the bottom) forecasts. The pictures on the right side report the cumulative differences of SFE relative to the NKPC_TS model, for 4-quarters-ahead (in the top) and 8-quarters-ahead (in the bottom) forecasts. The plots in Figure 6 should be interpreted as follows: when the line increases (decreases), the predictive regression of the NKPC_WAV_all (only D_6) model outperforms (underperforms) the alternative model (either AO or NKPC_TS).

The top left picture in Figure 6 shows that, at the 4-quarters ahead horizon, the largest forecast improvements of the NKPC_WAV_all (only D_6) model relative to the AO model occur essentially during the later stages of the Great Recession and in the following years until about 2012. It further shows that this model rarely underperforms the AO benchmark. The bottom left picture shows that, at the 8-quarters forecast horizon, the forecast improvements of the NKPC_WAV_all (only D_6) model relative to the AO are less continuous, with some episodes in which the benchmark outperforms our model. However, from the later stages of the Great Recession through about 2015, our model clearly improves on the forecasts of inflation two-year ahead produced by the benchmark model.

The top right picture in Figure 6 shows that, at the 4-quarters-ahead horizon, the forecasts from our NKPC_WAV_all (only D_6) model drastically outperform those of the NKPC_TS at the very end of the Great Recession and the first quarters of the recovery, and then substantially outperform

the time series NKPC forecasts from 2012 to about 2016. It further shows that our model rarely – and never substantially – underperforms the NKPC_TS model.

The bottom right picture in Figure 6 shows that the ability of our NKPC_WAV_all (only D_6) model to forecast inflation 8-quarters-ahead is essentially similar to that of the NKPC_TS until about 2009. Then, from the earlier stages of 2010 our model drastically outperforms the NKPC_TS model. Moreover, our NKPC_WAV_all (only D_6) model steadily keeps on outperforming the NKPC_TS model until about the end of 2016. Overall, it is particularly noteworthy the ability of our wavelet-based NKPC model to outperform the NKPC_TS when the forecasts are made during the turbulent times of the later stages of the Great Recession.

Finally, we ran several robustness checks. First, we replaced the proxy for slack and consider either the unemployment rate (in level) or the output gap (computed as the difference between real GDP and its linear trend). Second, we computed forecasts of the personal consumption expenditure (PCE) inflation (using the corresponding component for energy price inflation). Third, we used different wavelet filters (such as Daubechies and Coiflets of different lengths). Fourth, we computed the forecasts using rolling window estimates (with window size of 88 quarters, the same as our initial in-sample period). Results of these robustness checks are reported in Table 4. Overall, we find that our conclusions are qualitatively, and often quantitatively, robust to all these changes.

6 Concluding remarks

The expectations-augmented Phillips curve has long been the key model for explaining inflation, but has consistently failed to be useful in forecasting out-of-sample inflation. In this paper we show that a state-of-the-art empirical New Keynesian Philips Curve (NKPC) is able to outperform standard benchmarks in forecasting U.S. inflation. To reinstate the forecast ability of the NKPC, we combine a discrete wavelet approach with standard linear regression and forecast methods. In short, we decompose the time series of inflation and its NKPC predictors into a set of frequency

bands, forecast inflation with NKPCs estimated for each frequency band, and then obtain the forecast of inflation by aggregation of the forecasts from each cyclical frequency. Moreover, as motivated by the literature and the data, we consider a model that forecasts only the lowest frequency of inflation, with and without allowing for influences from other frequencies of its predictors.

Our evidence relates to the period 1979Q1-2019Q4, and consists of OOS forecasts of inflation 4- and 8-quarters ahead starting from 2000Q1. Our preferred NKPC wavelet-based model takes in only the forecasts from the lowest-frequency component of inflation, corresponding to cycles above 16 years, but flexibly allows for some influence from the higher-frequency fluctuations of the predictors into the low frequency of inflation. Furthermore, our NKPC wavelet-based model performs particularly better than traditional time series models in the turning from the Great Recession to the ensuing recovery.

Our results are consistent with those of an empirical literature focusing on the relevance of a slowly time-varying trend to successfully model the dynamics of inflation (Chan, Clark, and Koop, 2018 and Hasenzagl, Pellegrino, Reichlin, and Ricco, 2018). They also support recent findings from theoretical DSGE literature about the importance of modelling the low-frequency fluctuations of inflation to better capture its overall dynamics (Del Negro, Giannoni, and Schorfheide, 2015), as well as about the relevance of capturing interactions between macroeconomic variables at different frequencies (Comin and Gertler, 2006 and Beaudry, Galizia, and Portier, 2020).

Our procedure is theoretically and empirically grounded and may be used in real time forecasting. We thus conclude that the new Keynesian Phillips curve can be successfully used to forecast inflation if information from the frequency domain is taken seriously.

References

- AGUIAR-CONRARIA, L., M. M. MARTINS, AND M. J. SOARES (2019): “The Phillips Curve at 60: time for time and frequency,” Research Discussion Papers 12/2019, Bank of Finland.
- ASHLEY, R., AND R. VERBRUGGE (2009): “Frequency Dependence in Regression Model Coefficients: An Alternative Approach for Modeling Nonlinear Dynamic Relationships in Time Series,” *Econometric Reviews*, 28(1-3), 4–20.
- ATKESON, A., AND L. E. OHANIAN (2001): “Are Phillips curves useful for forecasting inflation?,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 25(2), 2–11.
- BANDI, F., B. PERRON, A. TAMONI, AND C. TEBALDI (2019): “The Scale of Predictability,” *Journal of Econometrics*, 208(1), 120–140.
- BAXTER, M., AND R. KING (1999): “Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series,” *Review of Economics and Statistics*, 81(4), 575–593.
- BEAUDRY, P., D. GALIZIA, AND F. PORTIER (2020): “Putting the Cycle Back into Business Cycle Analysis,” *American Economic Review*, 110(1), 1–47.
- BERGE, T. J. (2018): “Understanding survey-based inflation expectations,” *International Journal of Forecasting*, 34(4), 788–801.
- BERGER, T. (2016): “Forecasting Based on Decomposed Financial Return Series: A Wavelet Analysis,” *Journal of Forecasting*, 35(5), 419–433.
- BEVERIDGE, S., AND C. R. NELSON (1981): “A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the "business cycle",” *Journal of Monetary Economics*, 7(2), 151–174.
- CANOVA, F. (2007): “G-7 Inflation Forecasts: Random Walk, Phillips Curve Or What Else?,” *Macroeconomic Dynamics*, 11(1), 1–30.

- CARAIANI, P. (2017): “Evaluating exchange rate forecasts along time and frequency,” *International Review of Economics & Finance*, 51(C), 60–81.
- CHAN, J. C., T. E. CLARK, AND G. KOOP (2018): “A new model of inflation, trend inflation, and long-run inflation expectations,” *Journal of Money, Credit and Banking*, 50(1), 5–53.
- COIBION, O., AND Y. GORODNICHENKO (2015): “Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation,” *American Economic Journal: Macroeconomics*, 7(1), 197–232.
- COIBION, O., Y. GORODNICHENKO, AND R. KAMDAR (2018): “The formation of expectations, inflation, and the phillips curve,” *Journal of Economic Literature*, 56(4), 1447–91.
- COIBION, O., Y. GORODNICHENKO, AND S. KUMAR (2018): “How Do Firms Form Their Expectations? New Survey Evidence,” *American Economic Review*, 108(9), 2671–2713.
- COMIN, D., AND M. GERTLER (2006): “Medium-Term Business Cycles,” *American Economic Review*, 96(3), 523–551.
- CROWLEY, P. M., AND A. HUGHES HALLETT (2015): “Great moderation or "Will o' the Wisp"? A time-frequency decomposition of GDP for the US and UK,” *Journal of Macroeconomics*, 44(C), 82–97.
- DEL NEGRO, M., M. P. GIANNONI, AND F. SCHORFHEIDE (2015): “Inflation in the Great Recession and New Keynesian Models,” *American Economic Journal: Macroeconomics*, 7(1), 168–196.
- DIEBOLD, F. X., AND R. S. MARIANO (1995): “Comparing Predictive Accuracy,” *Journal of Business & Economic Statistics*, 13(3), 253–263.
- DOTSEY, M., S. FUJITA, AND T. STARK (2018): “Do Phillips Curves Conditionally Help to Forecast Inflation?,” *International Journal of Central Banking*, 14(4), 43–92.

- FARIA, G., AND F. VERONA (2017): “Out-of-sample time-frequency predictability of the equity risk premium,” Discussion paper, mimeo, previously circulated as "Forecasting the Equity Risk Premium with Frequency-Decomposed Predictors", Bank of Finland Discussion paper 1/2017.
- (2018): “Forecasting stock market returns by summing the frequency-decomposed parts,” *Journal of Empirical Finance*, 45, 228 – 242.
- (2020a): “Frequency-domain information for active portfolio management,” Research Discussion Papers 2/2020, Bank of Finland.
- (2020b): “The yield curve and the stock market: mind the long run,” *Journal of Financial Markets*, yy, xxx–xxx.
- FAUST, J., AND J. H. WRIGHT (2013): “Forecasting Inflation,” in *Handbook of Economic Forecasting*, ed. by G. Elliott, C. Granger, and A. Timmermann, vol. 2 of *Handbook of Economic Forecasting*, chap. 1, pp. 2–56. Elsevier.
- FUHRER, J. (2017): “Expectations as a source of macroeconomic persistence: Evidence from survey expectations in a dynamic macro model,” *Journal of Monetary Economics*, 86, 22–35.
- GALLEGATI, M., M. GALLEGATI, J. B. RAMSEY, AND W. SEMMLER (2011): “The US Wage Phillips Curve across Frequencies and over Time,” *Oxford Bulletin of Economics and Statistics*, 73(4), 489–508.
- GALLEGATI, M., F. GIRI, AND A. PALESTRINI (2019): “DSGE model with financial frictions over subsets of business cycle frequencies,” *Journal of Economic Dynamics and Control*, 100(C), 152–163.
- HASENZAGL, T., F. PELLEGRINO, L. REICHLIN, AND G. RICCO (2018): “A Model of the Fed’s View on Inflation,” CEPR Discussion Papers 12564, C.E.P.R. Discussion Papers.

- HODRICK, R. J., AND E. C. PRESCOTT (1997): “Postwar U.S. Business Cycles: An Empirical Investigation,” *Journal of Money, Credit and Banking*, 29(1), 1–16.
- LUBIK, T. A., C. MATTHES, AND F. VERONA (2019): “Assessing U.S. aggregate fluctuations across time and frequencies,” Research Discussion Papers 5/2019, Bank of Finland.
- MAVROEIDIS, S., M. PLAGBORG-MOLLER, AND J. H. STOCK (2014): “Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve,” *Journal of Economic Literature*, 52(1), 124–88.
- PERCIVAL, D., AND A. WALDEN (2000): *Wavelet methods for time series analysis*. Cambridge University Press.
- PEAJFAR, D., AND J. M. ROBERTS (2018): “The Role of Expectations in Changed Inflation Dynamics,” Finance and Economics Discussion Series 2018-062, Board of Governors of the Federal Reserve System.
- RUA, A. (2011): “A wavelet approach for factor-augmented forecasting,” *Journal of Forecasting*, 30(7), 666–678.
- (2017): “A wavelet-based multivariate multiscale approach for forecasting,” *International Journal of Forecasting*, 33(3), 581–590.
- STOCK, J. H., AND M. W. WATSON (1999): “Forecasting inflation,” *Journal of Monetary Economics*, 44(2), 293–335.
- (2007): “Why Has U.S. Inflation Become Harder to Forecast?,” *Journal of Money, Credit and Banking*, 39(s1), 3–33.
- (2008): “Phillips Curve Inflation Forecasts,” NBER Working Papers 14322, National Bureau of Economic Research.

- (2010): “Modeling inflation after the crisis,” in *Proceedings-Economic Policy Symposium-Jackson Hole*, pp. 173–220. Federal Reserve Bank of Kansas City.
- TALLMAN, E. W., AND S. ZAMAN (2017): “Forecasting inflation: Phillips curve effects on services price measures,” *International Journal of Forecasting*, 33(2), 442–457.
- VERONA, F. (2020): “Investment, Tobin’s Q, and Cash Flow Across Time and Frequencies,” *Oxford Bulletin of Economics and Statistics*, 82(2), 331–346.
- ZHANG, K., R. GENÇAY, AND M. E. YAZGAN (2017): “Application of wavelet decomposition in time-series forecasting,” *Economics Letters*, 158, 41–46.

Figure 1: Slope of the New Keynesian Phillips curve over time and across frequencies

Recursive estimates of the slope of the Phillips curve ((1)) for the original time series data (upper left graph) and filtered data for different frequency bands (remaining graphs). HF: high frequency, cycles of period between 2 and 8 quarters; BCF: business cycles, cycles of period between 2 and 8 years; LF: low frequency, cycles longer than 8 years. Sample periods are expanding windows starting in 1978Q1-1999Q4, recursively including one additional quarter through 2019Q4. Grey bars denote NBER-dated recessions. Statistically significant coefficients (at 5%) are reported with a circled marker.

Figure 2: Inflation: time series and MODWT MRA decomposition

Top graph: time series of inflation, for $h=4$, U.S. 1980Q1-2019Q4. Remaining graphs: six frequency-specific time series (D_j , $j = 1, \dots, 6$) resulting from the MODWT MRA decomposition of the inflation time series. Grey bars denote NBER-dated recessions.

Figure 3: Inflation expectations: time series and MODWT MRA decomposition

Top graph: time series of inflation expectations, U.S. 1980Q1-2019Q4. Remaining graphs: six frequency-specific time series (D_j , $j = 1, \dots, 6$) resulting from the MODWT MRA decomposition of the inflation expectations time series. Grey bars denote NBER-dated recessions.

Figure 4: Unemployment gap: time series and MODWT MRA decomposition

Top graph: time series of the unemployment gap, U.S. 1980Q1-2019Q4. Remaining graphs: six frequency-specific time series (D_j , $j = 1, \dots, 6$) resulting from the MODWT MRA decomposition of the unemployment gap time series. Grey bars denote NBER-dated recessions.

Figure 5: Energy inflation: time series and MODWT MRA decomposition

Top graph: time series of energy inflation, U.S. 1980Q1-2019Q4. Remaining graphs: six frequency-specific time series (D_j , $j = 1, \dots, 6$) resulting from the MODWT MRA decomposition of the energy inflation time series. Grey bars denote NBER-dated recessions.

Figure 6: Cumulative differences in squared forecast errors

Left side graphs: cumulative difference between the squared forecast errors of the NKPC_WAV_all (only D_6) model and those of the AO model, for $h=4$ (top) and $h=8$ (bottom). Right side graphs: cumulative difference between the squared forecast errors of the NKPC_WAV_all (only D_6) model and those of the NKPC_TS model, for $h=4$ (top) and $h=8$ (bottom). Grey bars denote NBER-dated recessions.

	c	α_1	α_2	α_3	$\bar{R}^2$
NKPC	-1.4 (0.19)	1.27 (0.05)	-0.15 (0.05)	0.08 (0.005)	0.89
NKPC_HF	0 (0.05)	0.08 (0.22)	-1.08 (0.36)	0.09 (0.005)	0.79
NKPC_BCF	0 (0.02)	1.01 (0.05)	-0.19 (0.04)	0.08 (0.003)	0.94
NKPC_LF	-2.25 (0.12)	1.60 (0.04)	-0.30 (0.03)	-0.01 (0.01)	0.96

Table 1: Estimates of the New Keynesian Phillips Curve

Estimates of equation (1), U.S. data, sample period 1978Q1-2019Q4. First row: estimates obtained with the original time series data. Subsequent rows: estimates obtained from filtered data for different frequency bands. HF: high frequency, cycles of period between 2 and 8 quarters; BCF: business cycles, cycles of period between 2 and 8 years; LF: low frequency, cycles longer than 8 years. Standard errors in parenthesis.

	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆
Inflation (h=4)	2	4	7	10	16	61
Inflation (h=8)	1	2	4	10	19	64
Inflation expectations	2	3	5	10	21	59
Unemployment gap	1	3	9	20	32	35
Energy inflation	36	28	15	8	6	7

Table 2: Variance decomposition by frequency

Each row presents the percentage of the variance of the corresponding time series explained by each specific frequency band, for the U.S. 1980Q1-2019Q4. D₁: cycles of period between 2 and 4 quarters. D₂: cycles of period between 4 and 8 quarters. D₃: cycles of period between 8 and 16 quarters. D₄: cycles of period between 16 and 32 quarters. D₅: cycles of period between 32 and 64 quarters. D₆: cycles of period larger than 64 quarters.

Model	Forecasting horizon	
	h=4	h=8
a) Benchmark		
AO random walk	1.66	1.18
b) Traditional models (vs AO)		
AR(p)	0.95	0.96
ARMA(1,1)	1.18***	1.05**
Survey	1.01	1.22
NKPC_TS	1.02	1.44
c) Wavelet-based models (vs AO)		
NKPC_WAV_diag	0.92*	0.96
NKPC_WAV_all	0.99	1.17**
NKPC_WAV_diag (only D ₆)	0.76**	0.78**
NKPC_WAV_all (only D ₆)	0.74**	0.77**
d) Wavelet-based models (vs NKPC_TS)		
NKPC_WAV_diag	0.90	0.66**
NKPC_WAV_all	0.97	0.81
NKPC_WAV_diag (only D ₆)	0.74*	0.54**
NKPC_WAV_all (only D ₆)	0.72*	0.53**

Table 3: Relative out-of-sample root mean squared forecast errors

Panel a): Root mean squared forecast errors (RMSFEs) at different forecasting horizons (h=4 and h=8) for the AO model. Panels b) and c): RMSFEs relative to those of the AO model ($\text{RMSFE}_j / \text{RMSFE}_{\text{AO}}$ for model j). Panel d): RMSFEs relative to those of the NKPC_TS model ($\text{RMSFE}_j / \text{RMSFE}_{\text{NKPC-TS}}$ for model j). Asterisks indicate statistical significance of the Diebold and Mariano (1995) test of comparative predictive accuracy at the 10% (*), 5% (**), and 1% (***) levels, relative to the AO model (panels b and c) or the NKPC_TS model (panel d). The out-of-sample period is 2000Q1-2019Q4.

Model	Measure of slack				Measure of inflation		Wavelet filter		Forecasting method	
	Unemployment rate		Output gap		PCE inflation		Daubechies length 4		Rolling window	
	Forecasting horizon		Forecasting horizon		Forecasting horizon		Forecasting horizon		Forecasting horizon	
	h=4	h=8	h=4	h=8	h=4	h=8	h=4	h=8	h=4	h=8
a) Benchmark										
AO random walk	1.66	1.18	1.66	1.18	1.25	0.92	1.66	1.18	1.66	1.18
b) Traditional models (vs AO)										
AR(p)	0.94	0.96	0.94	0.96	1.04	0.97	0.95	0.96	0.93	0.91
ARMA(1,1)	1.18***	1.05**	1.18***	1.05**	1.11**	1.03	1.18***	1.05**	1.01	1.01
Survey	1.01	1.22	1.01	1.22	1.33**	1.68**	1.01	1.22	1.01	1.22
NKPC_TS	1.02	1.35	0.99	1.47	1.15	1.62**	1.02	1.44	0.99	1.05
c) Wavelet-based models (vs AO)										
NKPC_WAV_diag	0.92*	0.93	0.95	1.05	0.97	0.99	0.97	1.01	0.98	0.98
NKPC_WAV_all	1	1.16**	0.98	1.12	1.03	1.23**	1.01	1.17**	0.99	1.07
NKPC_WAV_diag (only D ₆)	0.77**	0.80*	0.76**	0.79**	0.83	0.91	0.77**	0.81*	0.74**	0.77**
NKPC_WAV_all (only D ₆)	0.75**	0.77**	0.74**	0.77**	0.79*	0.85	0.76**	0.81*	0.74**	0.77**
d) Wavelet-based models (vs NKPC_TS)										
NKPC_WAV_diag	0.90*	0.69***	0.95	0.71**	0.85**	0.61**	0.95	0.70*	0.99	0.94
NKPC_WAV_all	0.98	0.86	0.99	0.76	0.90	0.76	0.98	0.82	1	1.02
NKPC_WAV_diag (only D ₆)	0.76*	0.59**	0.76	0.54**	0.73*	0.56**	0.75*	0.56**	0.75	0.74
NKPC_WAV_all (only D ₆)	0.73*	0.57**	0.75*	0.53**	0.69*	0.52**	0.75*	0.56**	0.75	0.74

Table 4: Relative out-of-sample root mean squared forecast errors - robustness checks

Panel a): Root mean squared forecast errors (RMSFEs) at different forecasting horizons (h=4 and h=8) for the AO model. Panels b) and c): RMSFEs relative to those of the AO model ($\text{RMSFE}_j / \text{RMSFE}_{\text{AO}}$ for model j). Panel d): RMSFEs relative to those of the NKPC_TS model ($\text{RMSFE}_j / \text{RMSFE}_{\text{NKPC-TS}}$ for model j). Asterisks indicate statistical significance of the Diebold and Mariano (1995) test of comparative predictive accuracy at the 10% (*), 5% (**), and 1% (***) levels, relative to the AO model (panels b and c) or the NKPC_TS model (panel d). The out-of-sample period is 2000Q1-2019Q4.