Centro de Economia e Finanças da UPorto

cef.up working paper
2021-01

Inflation Dynamics and Forecast: Frequency Matters

Manuel M. F. Martins
Fabio Verona

Centro de Economia e Finanças da UP
www.fep.up.pt

Inflation dynamics and forecast: frequency matters*

Manuel M. F. Martins† Fabio Verona‡

June 8, 2021

Abstract

Policymakers and researchers see inflation characterized by cyclical fluctuations driven by changes in resource utilization and temporary shocks, around a trend influenced by inflation expectations. We study the in-sample inflation dynamics and forecast inflation out-of-sample by analyzing a New Keynesian Phillips Curve (NKPC) in the frequency domain. In-sample, while inflation expectations dominate medium-to-long-run cycles, energy prices dominate short cycles and business-to-medium cycles once expectations became anchored. While statistically significant, unemployment is not economically relevant for any cycle. Out-of-sample, forecasts from a low-frequency NKPC significantly outperform several benchmark models. The long-run component of unemployment is key for such remarkable forecasting performance.

* This is an extended and revised version of a paper that circulated previously under the title "Forecasting inflation with the New Keynesian Phillips Curve: frequency matters". 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 and seminar participants at the Bank of Finland 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)

Keywords: inflation dynamics, inflation forecast, New Keynesian Phillips Curve, frequency domain, wavelets

JEL codes: C53, E31, E37

1 Introduction

Policymakers typically see inflation characterized by i) a trend strongly influenced by inflation expectations that, in turn, are shaped by the conduct of monetary policy, and ii) deviations from that trend caused by persistently high or low resource utilization, as well as temporary movements in energy prices and other shocks (see e.g. Yellen, 2016).

In the economics literature, this general description of the inflation dynamic is referred to as the expectations-augmented Phillips curve. The state-of-the-art empirical (micro-founded) New Keynesian Phillips Curve (NKPC) relates inflation (π_t) to inflation expectations (π^e_{t+1}), a measure of economic activity (such as the unemployment gap, $ugap_t$), and measures of supply shocks (such as energy inflation, ss_t) (see e.g. Coibion and Gorodnichenko, 2015):

$$\pi_t = c + \alpha_1 \, \pi^e_{t+1} + \alpha_2 \, ugap_t + \alpha_3 \, ss_t + \varepsilon_t \quad , \tag{1}$$

where c is a constant and ε_t the error term.

The NKPC and the policymakers' view of inflation dynamics are close, but not equivalent. For example, agents' expectations of next-period inflation (π^e_{t+1} in the NKPC) are not long-run expectations, and not necessarily shaped by the central bank's inflation target. Energy inflation may affect expectations and may, on occasions, have persistent impacts. The unemployment gap may result from various types of shocks and its fluctuations may exhibit varying persistence over time. Hasenzagl, Pellegrino, Reichlin, and Ricco (2021) have reconciled these two views through a multivariate semi-structural time series model with endogenous trend inflation, natural unemployment, and potential output.

In this paper, we take a different approach by using a frequency-domain decomposition of the time series of inflation and its NKPC determinants. To understand our approach, consider the decomposition of the time series of inflation (π_t) and of its NKPC determinants into four cyclical components – high frequency (HF, less than 2 years), business cycle frequencies (BCF, 2 to 8 years), medium frequencies (MF, 8 to 16 years), and low frequencies (LF, more than 16 years) – so that $\pi_t = \pi_t^{HF} + \pi_t^{BCF} + \pi_t^{MF} + \pi_t^{LF}$. The low-frequency (long-run) component of inflation should be closely related with inflation expectations, especially with their low-frequency component, and probably less so with resource utilization and energy shocks. At the other end of the frequency spectrum, the high-frequency (short-run) movements of inflation should be closely related with energy prices shocks, especially their dominant high-frequency component, and less so with cycles of unemployment and expectations. At intermediate frequencies (BCF and MF), unemployment is expected to play a key role, with the relevance of inflation expectations and supply shocks possibly changing during specific episodes.

We begin by decomposing the time series of inflation and its NKPC determinants with discrete wavelet tools. We then use standard linear regression and forecast methods to study the in-sample dynamics of inflation, and to forecast inflation out-of-sample in a consistent framework based on the NKPC.

Our first main contribution is to provide new stylized facts on inflation dynamics. Using a frequency domain approach is important in light of the controversies surrounding the Phillips curve's ability to explain inflation dynamics. The instability and flattening of the Phillips curve in the 1980s or 1990s (see e.g. Ball, Mazumder, Dynan, and Stock, 2011, Coibion and Gorodnichenko, 2015, and Del Negro, Lenza, Primiceri, and Tambalotti, 2020) has been often associated with phenomena related to specific frequencies such as the low-frequency anchoring of inflation expectations. Relatedly, some authors (e.g. Cogley and Sbordone, 2008) have analyzed the interaction between low-frequency and high-frequency variations in inflation dynamics. Moreover, the theoretical literature using Dynamic Stochastic General Equilibrium (DSGE) models emphasizes the importance of properly modeling medium- and low-frequency movements of inflation (Del Negro, Giannoni, and Schorfheide, 2015), and capturing interactions between macro variables at different frequencies (see e.g. Comin and Gertler, 2006, Beaudry, Galizia, and Portier, 2020 and Angeletos, Collard, and Dellas, 2020). Against this background, we analyze the dynamics of inflation in the frequency domain and provide answers to the following questions: Is the magnitude and (in)stability of the NKPC different across cyclical frequencies? Is there any flattening of the NKPC? If so, at which cycles? What is the contribution of each NKPC determinant in explaining inflation across cycles and over time?

We emphasize the following results about inflation dynamics. First, the estimates of the NKPC slopes for the three frequency bands comprising cycles up to 16 years (HF, BCF, and MF) are quite similar to the standard time series estimate, while the estimate for cycles longer than 16 years (LF) is three times larger. Second, the NKPC slope is remarkably stable for medium frequencies and increasingly stronger after the Great Recession for low frequencies, which suggests that the Phillips trade-off is highly relevant in the medium-to-long run. Third, there is some evidence of flattening after the Great Recession, and, more recently, during the Covid recession, but only distinctly for very short cycles (HF) and quite mildly for business cycles (BCF). Fourth, when assessing the contribution of each NKPC determinant in explaining inflation over time, we find that expectations dominate at cycles longer than 8 years, while energy inflation dominates at short cycles (HF). Moreover, energy inflation become highly relevant at the business and medium cycles once expectations became anchored around 2000. Most importantly, unemployment fails to account for a substantial part of the variation of inflation – even at business-cycle frequencies.

The second main contribution of this paper is the use of a frequency-domain approach in making out-of-sample forecasts of inflation. This contribution is important in light of the inability of standard NKPCs to beat simple time-series models in forecasting inflation out-of-sample (see e.g. Canova, 2007, Stock and Watson, 2007, Dotsey, Fujita, and Stark, 2018, and Berge, 2018), and implying that the Phillips curve is an unreliable model for inflation forecasting. The consensus in the literature is that good inflation forecasts must account for a slowly varying local mean for inflation (see e.g. Faust and Wright, 2013 and Chan, Clark, and Koop, 2018), which, of course, has clear frequency-domain implications. Moreover, the data indicate that a large part of the variance of inflation is due to its low-frequency fluctuations. This is not the case for all NKPC variables, implying that a flexible use of frequency-domain information across variables may improve the inflation forecast. In this context, we attempt to answer the following questions: Can frequency-domain NKPC-based forecasts not only beat those of the usually most successful benchmarks, but those of a standard time-series NKPC as well? Are all inflation frequencies equally important to forecast? Which NKPC determinants of inflation matter most? How does the forecast performance of the frequency-domain NKPC evolve over time?

Our main results for the out-of-sample inflation forecast can be summarized as follows. First, the best forecast of inflation (for different forecasting horizons) is one based solely on the forecasts of the low-frequency component of inflation. In other words, the key to a good inflation forecast is to forecast well its low frequency and ignore the other frequencies of inflation. This model consistently outperforms both the standard benchmarks and the forecasts with the time-series NKPC over the 21 years of our out-of-sample exercise. Its relative performance is particularly good for the later stages of the Great Recession and throughout the recovery. Second, it is also important to use the low frequency of the unemployment gap as a predictor. Adding other frequencies of the predictors further improves the forecasts of inflation, but only marginally.

Overall, our frequency-domain approach to inflation provides a refined analysis of inflation dynamics and an improved method for forecasting inflation. While unemployment may seem of little relevance in-sample, a key insight of this work is that it remains crucial in predicting inflation out-of-sample due to its role at cycles longer than typical business cycles. The Phillips curve is alive and well – as long as both time-domain and frequency-domain information is taken into account.

Our approach has a number of highly convenient features and a methodological innovation. First, the time series are decomposed into frequency-domain components that have clear economic meaning – high frequency (less than 2 years), business cycle frequency (2 to 8 years), medium frequencies (8 to 16 years), and long-run frequencies (longer than 16 years). Second, each time series is decomposed into frequency components that add up precisely to the original time series. This assures that we are neither ignoring nor overlapping information. Third, our approach to forecast is flexible as the forecast of each frequency component of inflation may depend on other frequency components of the predictors. Our method allows for forecasting a specific frequency of inflation using information of expectations, slack, and supply shocks from the same and possibly other frequencies. This methodological contribution should be quite useful in applications in which predictors have very different relative power across cyclical bands.

The rest of the paper is organized as follows. In section 2, we briefly discuss the studies that laid the groundwork for this work. In section 3, we present the data and the method. Sections 4 and 5 present the results for our in-sample analysis of inflation dynamics and for our out-of-sample forecasts. Section 6 concludes.

2 Related literature

This paper relates to three strands of literature, i.e. specification of the NKPC, frequency-domain analyses of Phillips curves, and the frequency-domain approach to forecasting.

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), as well as of strong specification and sampling uncertainty (Mavroeidis, Plagborg-Moller, and Stock, 2014).

Here, we use a state-of-the-art empirical NKPC (equation 1) in the spirit of Coibion and Gorodnichenko (2015), Fuhrer (2017), and Coibion, Gorodnichenko, and Kamdar (2018).[1] This NKPC

[1] Given our purpose of describing and forecasting inflation with an integrated framework that explores the frequency-domain information in the data, we follow the vast literature focusing on reduced-form Phillips curves, where identification depends upon the ability of expectations and energy inflation to control for changes in the position of the curve. Alternative and more sophisticated identification strategies that have recently gained popularity include the use of instrumental variables or cross-sectional (micro or regional) data to isolate the effect of

relates to a vast literature that gradually refined the micro-founded full information rational expectations NKPC (described by e.g. Woodford, 2003), replacing the labor share (e.g. Gali and Gertler, 1999) with the unemployment (or outputs) gap (e.g. Rudd and Whelan, 2007 and King and Watson, 2012) as proxy for marginal costs, adding controls for supply shocks (in the spirit of Gordon, 2011), and replacing rational expectations with survey expectations (e.g. Roberts, 1995).

Households survey inflation expectations, given by the Michigan survey of consumers, became the state-of-the-art in recent empirical NKPC for both theoretical and empirical reasons. Most notably, household expectations have been shown to be the closest possible to firm expectations, which while relevant according to micro foundations, are still unavailable (see Coibion and Gorodnichenko, 2015, Coibion, Gorodnichenko, and Kumar, 2018, and Pfajfar and Roberts, 2018). Moreover, the Michigan survey of household inflation expectations features intrinsic inertia due to the micro-founded inefficiency with which agents revise expectations, and thus dispenses of ad hoc inertial mechanisms (Fuhrer, 2017 and Coibion, Gorodnichenko, and Kamdar, 2018).[2]

While equation (1) has often been used to study the dynamics of inflation, it has rarely been used to forecast inflation out-of-sample. The literature of inflation forecasting based on Phillips curves mostly relies on some version of the Friedman-Phelps 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 US inflation with several models, including specifications of the NKPC with expectations of inflation taken from the Michigan survey of consumers. The key difference between our paper and Berge (2018) is that we run the forecasts in the frequency domain, which represents a novelty in the literature of Phillips curve-based inflation forecasting.[3]

---

demand shocks on real activity (see e.g. McLeay and Tenreyro, 2020, Hazell, Herreno, Nakamura, and Steinsson, 2020, Hooper, Mishkin, and Sufi, 2020, and Barnichon and Mesters, 2021).

[2] 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 thorough review of the specification of the state-of-the-art empirical NKPC.

[3] Another (less relevant) difference with Berge (2018) is that we include a proxy for supply shocks.

2.2 The Phillips curve in the frequency domain

Our paper is by no means the first to explore the idea that the Phillips curve differs across frequencies.

Some authors have used spectral analysis, ultimately based on Engle (1974), to study wage and price inflation Phillips curves. For instance, Reinbold and Wen (2020) identify demand shocks and improve the identification of a static Phillips curve across frequencies using spectral analysis. Other papers use band-pass filters and focus on the Phillips relation at business-cycle frequencies (e.g. King and Watson, 1994). Overall, this literature has found a significant relationship between inflation and unemployment at business-cycle frequencies, but not at shorter or longer cycles in the context of static or accelerationist Phillips curves.

Another group of papers use frequency-dependent regression models (see Ashley and Verbrugge, 2009, 2020) to assess the (non)linearity of the Phillips curve. They find that accelerationist Phillips curves exhibit different slopes over the business cycle, with inflation reacting differently to persistent and non-persistent fluctuations of the unemployment gap.

Our paper most closely relates to studies that use wavelet tools to study the dynamics of inflation in the framework of Phillips curves. Using the continuous wavelet transform (CWT), Aguiar-Conraria, Martins, and Soares (2019) show that there is considerable variation of the coefficients of a US NKPC identical to (1), both across frequencies and over time. Using the discrete wavelet transform (DWT), Gallegati, Gallegati, Ramsey, and Semmler, 2011 not only find a structural break in the mid-1990s and a significant slope in the US accelerationist wage Phillips curve at business cycle frequencies, but also at cycles longer than 8 years. Combining DWT and CWT approaches, Fratianni, Gallegati, and Giri (2019) find that estimates of the UK wage Phillips curve have been more stable and consistently significant over time at medium-run cycles than at business cycle frequencies.

Our paper offers several contributions to this field. First, we innovate in studying inflation with the Maximal Overlap Discrete Wavelet Transform (MODWT) using a state-of-the-art empirical NKPC. Second, we innovate by using such framework to consistently study both the in-sample dynamics and the out-of-sample forecast of inflation. Third, we are the first to build inflation forecasts from its frequency-specific forecasts.[4]

2.3 Frequency-domain forecasts

Our approach extends the use of discrete wavelet methods to forecast out-of-sample economic and financial time series. This area of research includes Rua (2011) and Rua (2017), who forecasts GDP growth and inflation using a factor-augmented wavelets approach; Zhang, Gençay, and Yazgan (2017), Faria and Verona (2018, 2020b, 2021), who focus on forecasting stock market returns; Caraiani (2017), who forecasts exchange rates; and Faria and Verona (2020a), who forecast the bond risk premium and the equity risk premium.

Our paper relates closely to the work of Faria and Verona (2018, 2021) on improving forecasts of a variable of interest by summing the forecasts of its frequency components rather than directly forecasting the aggregate. In those papers, each frequency component of the variable of interest is forecasted using only information from the same frequency component of the predictors. We contribute to this literature by allowing – but not imposing – that each frequency component of inflation may depend on other frequency components of the predictors. For instance, the business-cycle frequencies or medium-term frequencies of, say, the unemployment gap, is allowed to affect low-frequency fluctuations of inflation. Our generalization of the wavelet-based approach to forecasting is potentially useful for many economics problems, and seems particularly valuable in the case of inflation in light of the different patterns of variance across frequencies of inflation, expectations, unemployment, and energy prices.

[4] Rather than decomposing the aggregate time series into frequencies, a recent alternative strategy involves splitting the price index into its components and empirically identifying the specific price categories sensitive to the degree of resource utilization. This cyclically sensitive inflation is then used in a Phillips curve framework to analyze inflation dynamics and forecast inflation (see Zaman, 2019 and Stock and Watson, 2020).

3 Data and method

Our data are US quarterly time series for 1978Q1–2020Q4 of inflation, the unemployment gap, expectations of inflation, and energy inflation. Let P_t be the consumer price index (CPI) provided by the US Bureau of Labor Statistics in quarter t. The annualized quarter-on-quarter inflation rate is computed as $\pi_t = 400 \ln (P_t/P_{t-1})$. 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 the difference between the quarterly average of the civilian unemployment rate provided by the US Bureau of Labor Statistics and a linear trend.

To analyze the Phillips curve in the frequency domain, we use wavelet filtering methods to decompose our time series into individual components that can be associated with fluctuations at different frequencies. The wavelet method used in this paper allows for decomposing any time series into a trend (or permanent) component and cyclical (or transitory) movements in a manner similar to the traditional time series trend-cycle decomposition approach (e.g. Beveridge and Nelson, 1981), or filtering methods as the Baxter and King (1999) bandpass filter or the Hodrick and Prescott (1997) filter.[5]

Specifically, we use the MODWT with the Haar filter. Hence, a time series y_t can be decomposed as:

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

where

$$y_t^{D_j} = \frac{1}{2^j}\left(\sum_{i=0}^{2^{j-1}-1} y_{t-i} - \sum_{i=2^{j-1}}^{2^j-1} y_{t-i} \right) \tag{3}$$

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

are the J wavelet detail components and

$$y_t^{S_J} = \frac{1}{2^J} \sum_{i=0}^{2^J-1} y_{t-i} \tag{4}$$

is the wavelet scaling component.

Expression (2) shows that the original time series y_t can be decomposed in different time series components, each representing the fluctuations of the original variable within a specific frequency band. For small j, the j wavelet detail components (expression (3)) 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 (expression (4)) captures the lowest frequency fluctuations (i.e. its trend).

The wavelet components resulting from the MODWT with Haar filter are easy to interpret as they are simply differences of moving averages. When $J = 1$, a time series y_t is decomposed into a transitory component ($y_t^{D_1}$) and a persistent scale component ($y_t^{S_1}$) as:

$$y_t = \underbrace{\frac{y_t - y_{t-1}}{2}}_{y_t^{D_1}} + \underbrace{\frac{y_t + y_{t-1}}{2}}_{y_t^{S_1}} .$$

When $J = 2$, the decomposition results in two detail components ($y_t^{D_1}$ and $y_t^{D_2}$) and a scale component ($y_t^{S_2}$) such that:

$$y_t = \underbrace{\frac{y_t - y_{t-1}}{2}}_{y_t^{D_1}} + \underbrace{\frac{y_t + y_{t-1} - (y_{t-2} + y_{t-3})}{4}}_{y_t^{D_2}} + \underbrace{\frac{y_t + y_{t-1} + y_{t-2} + y_{t-3}}{4}}_{y_t^{S_2}} .$$

The first detail component ($y_t^{D_1}$) remains unchanged, while the prior persistent component ($y_t^{S_1}$) is divided into an additional transitory component ($y_t^{D_2}$) and a new persistent one ($y_t^{S_2}$).

Here, we compute a J=5 level decomposition of our time series. As we use quarterly data, the first component (D_1) captures fluctuations with a period between 2 and 4 quarters, while the components D_2, D_3, D_4 and D_5 capture fluctuations with periods of 1–2, 2–4, 4–8, and 8–16 years, respectively. Finally, the smooth component S_5 captures fluctuations with a period longer than 16 years.[6] Subsequently, the high-frequency (HF) component of each variable (e.g. π_t) is computed as $\pi_t^{HF} = \pi_t^{D_1} + \pi_t^{D_2}$, the business-cycle-frequency (BCF) component (π_t^{BCF}) is computed as $\pi_t^{BCF} = \pi_t^{D_3} + \pi_t^{D_4}$, whereas its medium frequency (MF) and the low frequency (LF) components correspond to $\pi_t^{D_5}$ and $\pi_t^{S_5}$, respectively.

The original time series of our variables and their frequency components are reported in Figures 1 to 4. These figures highlight the fact that the original time series are the result of the aggregation of several underlying frequency components that exhibit quite different dynamics. Table 1 reports the variance decomposition by frequency of our variables. About half of the volatility of inflation occurs at the LF band, while a quarter of its volatility is due to HF fluctuations. Almost two-thirds of the variance of inflation expectations occurs at the LF band, and almost all their variance occurs at the three lower frequency bands (BCF, MF, and LF). The unemployment gap exhibits a more even distribution of its variance across the three lower frequency bands. Two-thirds of energy inflation variance, in contrast, is concentrated in the HF band.

4 Inflation dynamics

In this section we assess the in-sample dynamics of inflation through the lenses of the NKPC across cyclical frequencies and over time. In the first sub-section, we focus on the estimates and statistical significance of the NKPC coefficients. In the second sub-section, we turn our attention to the economic relevance of each NKPC determinant of inflation.

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

4.1 Time- and frequency-varying NKPC coefficients

We start by analyzing how different are the estimates of the NKPC across cyclical frequencies.

Table 2 shows linear regressions of the NKPC for our US quarterly data for 1978Q1–2020Q4. The first row displays the standard time series estimates, which are generally consistent with the literature (see e.g. Coibion and Gorodnichenko (2015). The second to fifth rows present estimates of the NKPC for the four frequency bands, i.e. estimates of OLS regressions

$$\pi_t^f = c^f + \alpha_1^f \pi_{t+1}^{e,f} + \alpha_2^f ugap_t^f + \alpha_3^f en_t^f + \varepsilon_t^f \quad , \tag{5}$$

where f=HF, BCF, MF, LF.

The estimate of the NKPC slope (α_2^f) is quite similar for high frequency (HF), business-cycle frequencies (BCF) and the medium-term frequencies (MF), and all are identical to the time series estimate. Yet, the slope is much more precisely estimated for business cycles and medium-run cycles. For the long run (LF), the estimate of the slope is three times that of the time series and other frequencies. Hence, the Phillips tradeoff is not merely a business cycle relationship – it is highly relevant in the medium-run and even longer cycles. This result comports with the findings of e.g. Del Negro, Giannoni, and Schorfheide (2015), Beaudry, Galizia, and Portier (2020), and Angeletos, Collard, and Dellas (2020).

With the sole exception of the long run, the estimates of the energy price coefficients are remarkably identical to the time series estimates across all frequency bands.

The coefficient associated with expectations differs the most across frequencies. Inflation reacts essentially one-to-one to changes in expectations at business-cycle frequency and medium run, but varies much less with expectations at high frequencies and much more at low frequencies. Overall, the standard time series estimate appears to be an artifact, i.e. it averages out the heterogeneous estimates across different cycles – estimates that essentially indicate that the reaction of inflation to expectations increases with the length of the cycle.

We now check whether the estimates of the NKPC coefficients for our cyclical frequencies are more or less stable over time than those of the standard time-series NKPC.

We do so by inspecting estimates obtained from expanding windows that start with the sample 1978Q1–1999Q4, and recursively add one quarter through 2020Q4.[7] Figure 5 shows these estimates and their statistical (in)significance. The charts in the first column plot estimates of α_1, the second column those of the NKPC slope (α_2), and the third column the coefficients associated with energy prices (α_3). The graphs in the upper row report the results for the standard time-series NKPC, and each subsequent row shows time-varying estimates for each frequency band, from HF through LF.

We find no evidence of flattening of the time-series NKPC after 2000Q1. The estimate of the slope features an abrupt change in the latter stages of the Great Recession, and becomes statistically significant after 2010 when it is larger (in absolute value) than in the previous ten years. While our result is not strictly comparable with most literature on the flattening of the Phillips curve (which dates it sometime between the mid-1980s and the early 1990s), it contrasts with papers that point to a flatter curve after the Fed announced its 2 % inflation target in 2012 (e.g. Bundick and Smith, 2020). The sensitivity of inflation to expected inflation and energy prices changes abruptly at the end of the Great Recession, but remains statistically significant throughout. While the former is fairly stable, the latter is stronger after the Great Recession.

There are considerable differences in the evolution over time of the estimates of the NKPC coefficients for the frequency bands. Focusing on the slope, we first note that it is more precisely estimated for most frequency bands (HF, BCF, MR) than for the aggregate time series. There is evidence of flattening after the Great Recession and during the Covid recession, but, as noted,

[7] We set the minimum length of our samples at 22 years, and thus focus on estimates from 2000Q1 onwards as estimates for frequency bands comprising medium- and long-run cycles would not be robust with shorter samples. The choice of the initial sample period is also consistent with the forecast period used in the out-of-sample forecast part of the paper.

these are only distinct for short cycles and quite mild for business cycles. The estimate of the NKPC slope is remarkably stable for medium frequencies and increasingly stronger after the Great Recession for low frequencies. Overall, these results suggest that the Phillips trade-off remains a significant business-cycle phenomenon and is increasingly a medium-to-long-run phenomenon in the post-Great Recession era.

The estimates of the coefficient associated to energy prices increase abruptly at the end of the Great Recession for both the HF and BCF, suggesting that its variation for the time-series NKPC is determined by these frequencies.

The estimates of the coefficient associated to expectations drop somewhat at the end of the Great Recession for the BCF and HF (recovering to much higher values during the Covid recession in the latter case). Their variation is much smaller for MF and, similar to BCF, overall indicative of a one-to-one change of inflation with expectations. In contrast, the sensitivity of inflation to expectations increases steadily and substantially between 2009 and 2015 at LF.

Overall, the dynamics of inflation is explained by the NKPC rather differently for the four frequency bands, which implies that it is useful to assess it separately for the different cycles.

4.2 Time- and frequency- varying NKPC determinants of inflation

We now move from statistical significance to economic relevance, conducting a full-sample decomposition of inflation using the coefficient estimates of the NKPC in Table 2 and the actual data. Figure 6 shows the contribution of each NKPC determinant to explain inflation from 1978Q1 to 2020Q4.

The top graph relates to the time-series NKPC. It shows that the good fit of that model comes essentially from expected inflation (blue line) and somewhat less from energy prices (green line). Notably, the explanatory contribution of unemployment (red line) to inflation is apparently quite limited throughout the entire sample.

The results for each of the four frequency components are reported in the graphs in the second and third row. Clearly, over the very short run, the NKPC explains inflation almost entirely through energy prices. Expectations and unemployment do not capture much of the high-frequency movements of inflation.

Over the long run, inflation co-moves almost perfectly with expectations. At low frequencies, energy prices and unemployment do not contribute visibly to a NKPC explanation of inflation. There is also a prominent role of expectations apparent over the medium run. Since the late 1990s, however, there are several episodes in which expectations do not explain medium-run inflation as well as energy prices. This shift may be related to the anchoring of expectations and their consequent limitation in explaining medium-run fluctuations of inflation.

Results for business cycles are similar to those of medium frequency cycles in spite of the higher volatility of all variables. Expectations and energy prices contribute most to explaining inflation, while the contribution of unemployment is marginal overall. Moreover, since the late 1990s, the contribution of expectations to explain inflation along the business cycle has often fallen considerably short of the contribution of energy prices. Again, we view such result as consistent with the anchoring of expectations, which is consistently dated by the literature at around 1999 (see e.g. Jorgensen and Lansing, 2019 and Carvalho, Eusepi, Moench, and Preston, 2021).

Overall, our in-sample results confirm the key role of expectations in explaining the low-frequency fluctuations of inflation. They attribute a somewhat larger than expected role to energy prices. They are highly relevant at HF, BCF, and even at MF, once expectations become anchored. In contrast, our results strongly suggest that unemployment has a very limited role in explaining inflation, even in the context of the NKPC, of which it should be a key determinant of inflation. This is apparent in the aggregate time-series NKPC and when we apply the NKPC to the frequency components of the data at which the role of unemployment should be particularly relevant (business cycles and medium-run cycles).

Our results in this section motivate our next steps in the paper: is the NKPC an useful model to forecast inflation out-of-sample? Is there any role for unemployment in the forecast of inflation?

5 Inflation forecast

In this section, we assess the performance of our frequency-domain approach to the NKPC to forecast inflation out-of-sample (OOS). This contrasts with the purely time-series benchmarks that dominate the literature and the standard time-series NKPC. The first sub-section is methodological. After briefly reviewing how forecasts are computed with benchmark models, we describe our method. In the second sub-section, we present our results in three steps. First, we compare the overall precision of our forecasts with those of the benchmarks. Second, we describe how the relative performance of our forecasting method evolves along time. Finally, we show which cyclical frequencies and which NKPC determinants of inflation actually improve forecasts.

Our OOS forecasting exercise targets the annualized h-period-ahead average inflation rate: $\pi^h_{t+h} = \frac{1}{h}\sum_{i=1}^{h}\pi_{t+i} = \frac{400}{h}\ln\left(P_{t+h}/P_t\right)$. We focus on the 1-quarter-, 4-quarters- and 8-quarters-ahead (h=1,4,8) horizons as they are the most relevant for policymakers.

Our OOS forecasts are direct forecasts produced with a sequence of expanding windows. We start by obtaining the first OOS forecasts with the sample 1978Q1–1999Q4. The sample is then increased by one observation and a new set of OOS forecasts is produced. This procedure is repeated until the end of the sample. Hence, the full OOS period runs from 2000Q1 to 2020Q4.

5.1 Forecasting models

5.1.1 Time series benchmarks

As is common in the literature on inflation forecasting, we compare the forecasting performance of the NKPC against two time series models. First, the random walk model of Atkeson and Ohanian

(2001) (henceforth denoted AO), for which the h-period-ahead forecast is given by $\hat{\pi}^h_{t+h} = \pi^h_t$. Second, the unobserved components model with stochastic volatility of Stock and Watson (2007) (henceforth denoted UCSV), which is given by $\pi_t = \tau_t + \varepsilon_t$ and $\tau_t = \psi_{t-1} + \eta_t$, where ε_t and η_t feature stochastic volatility. The forecasts of the UCSV model are computed as $\hat{\pi}^h_{t+h} = \tau_t$.[8]

5.1.2 Time-series NKPC

Ultimately, we want to assess the merit of our wavelet-based NKPC in forecasting US inflation relative to the time-series NKPC. Forecasts with the NKPC are obtained as follows.

At each step of the recursive sample and OOS period, for each h we first estimate a regression:

$$\pi^h_t = c^h + \alpha^h_1 \pi^e_{t+1} + \alpha^h_2 ugap_t + \alpha^h_3 en_t + \varepsilon^h_{t+h} \ , \tag{6}$$

and then compute the forecasts as:

$$\hat{\pi}^h_{t+h} = \hat{c}^h + \hat{\alpha}^h_1 \pi^e_{t+1} + \hat{\alpha}^h_2 ugap_t + \hat{\alpha}^h_3 en_t \ . \tag{7}$$

While inflation, expectations, and energy inflation are observed, the unemployment gap needs to be computed. At each step of our recursive forecasting procedure, we compute the unemployment gap by fitting a linear trend to the available unemployment data. Such procedure avoids any "look-ahead" bias in the predictive regression forecast based on the NKPC. This model is denoted henceforth as NKPC_TS.

[8] In the UCSV model, we use non-centered parameterization as in Chan (2018).

5.1.3 Wavelet-based NKPC

Our wavelet-based NKPC forecast method builds on the NKPC and on the filtered data obtained with the MODWT decomposition.[9]

For each frequency component f, f=HF, BCF, MF, LF, we specify a Phillips curve such as:

$$\pi_{t+h}^{h,f} = c^{h,f} + \alpha_1^{h,f}\pi_{t+1}^{e,f} + \alpha_2^{h,f}ugap_t^f + \alpha_3^{h,f}en_t^f + \varepsilon_{t+h}^{h,f} \quad . \tag{8}$$

In this baseline specification, each frequency component of inflation $\pi_{t+h}^{h,f}$ depends only on the same frequency component f of the NKPC determinants of inflation. We then generalize this specification by allowing other frequency components of the predictors into the NKPC of inflation at frequency f, $\pi_{t+h}^{h,f}$. This methodological contribution should prove highly useful, given that inflation and its determinants have different relative powers across cyclical bands. Moreover, theory and practice both suggest that they may interact with each other across frequencies.

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

$$\pi_t^{h,f} = c^{h,f} + \boldsymbol{\alpha_1^h}\pi_{t+1}^{e,f} + \boldsymbol{\alpha_2^h}ugap_t^f + \boldsymbol{\alpha_3^h}en_t^f + \varepsilon_t^{h,f} \quad . \tag{9}$$

where $\pi_t^{h,f}$, $\pi_{t+1}^{e,f}$, $ugap_t^f$ and en_t^f are 4x1 vectors of observables, $c^{h,f}$ is a 4x1 vector of intercepts and $\varepsilon_t^{h,f}$ is a 4x1 vector of residuals, while $\boldsymbol{\alpha_1^h}$, $\boldsymbol{\alpha_2^h}$ and $\boldsymbol{\alpha_3^h}$ are 4x4 matrices of coefficients.[10] We consider two cases of this wavelet-based model.

In a first case, matrices $\boldsymbol{\alpha_m^h}$, with $m = 1, 2, 3$, are diagonal. In this baseline specification, we

[9] We use a two-sided version of the Haar filter in the OOS exercise. To avoid a "look-ahead" bias, we compute the filtered time series of inflation and its determinants recursively at each iteration of the OOS forecasting procedure by using data from the start of the sample through the quarter at which we build the forecast. Therefore, our forecasts are made with current and past information only. To deal with boundary effects, we use a reflection rule, i.e. we extend the time series symmetrically at the boundaries before computing the filtered series.

[10] Regarding 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.

assume that only the components of the predictors at frequency f are used to forecast the frequency component of inflation at frequency f. We denote this model as NKPC_WAV_diag.

In a second case, we allow for interactions between inflation and its determinants across frequencies. In this model, all the coefficients in matrices $\boldsymbol{\alpha_m^h}$, $m = 1, 2, 3$, 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 a generalization of NKPC_WAV_diag, including 12 predictors for each frequency f of inflation, rather than 3, and it should lead to better in-sample fit. However, it remains to be seen whether relations between inflation and its determinants across frequencies are empirically so relevant that the improved in-sample fit does not harm the OOS performance. NKPC_WAV_diag is more parsimonious and may achieve better OOS performance. Finding which specification predicts inflation most accurately and robustly is a key contribution of our approach.

Formally, after running regression equation (9), the forecasts of each frequency component of inflation are computed as:

$$\hat{\pi}_{t+h}^{h,f} = \hat{c}^{h,f} + \boldsymbol{\hat{\alpha}_1^h} \pi_{t+1}^{e,f} + \boldsymbol{\hat{\alpha}_2^h} ugap_t^f + \boldsymbol{\hat{\alpha}_3^h} en_t^f \quad . \tag{10}$$

Finally, given that $\pi_{t+h}^{h} = \pi_{t+h}^{h,HF} + \pi_{t+h}^{h,BCF} + \pi_{t+f}^{h,MF} + \pi_{t+h}^{h,LF}$, then the h-quarter-ahead inflation forecast is given by the sum of the h-quarter ahead forecasts given by each frequency component of the NKPC:

$$\hat{\pi}_{t+h}^{h} = \hat{\pi}_{t+h}^{h,HF} + \hat{\pi}_{t+h}^{h,BCF} + \hat{\pi}_{t+f}^{h,MF} + \hat{\pi}_{t+h}^{h,LF} \quad . \tag{11}$$

In the context of forecasting stock market returns, Faria and Verona (2018) show that forecasts of the variable of interest may be improved by disregarding the forecasts of some of its frequencies.

In our context, this would mean that the h-quarter-ahead inflation forecast given by

$$\hat{\pi}^{h}_{t+h} = \kappa_{HF}\hat{\pi}^{h,HF}_{t+h} + \kappa_{BCF}\hat{\pi}^{h,BCF}_{t+h} + \kappa_{MF}\hat{\pi}^{h,MF}_{t+f} + \kappa_{LF}\hat{\pi}^{h,LF}_{t+h} \quad , \tag{12}$$

where κ_f are set to 0 (instead of being 1) for some frequencies *f*, could be better than that obtained by summing all frequencies forecasts of inflation (equation 11).

In principle, one could optimize the forecast by grid-searching the weights for each frequency that minimize the root mean squared error (RMSFE) for the entire OOS period. Given that such a procedure is not implementable in real-time, a simple method that considers only forecasts for the low-frequency component has strong motivations in the literature of wavelet-based forecasts (see Faria and Verona, 2018, 2021), as well as some motivation in the literature on inflation forecasting (e.g. Faust and Wright, 2013 and Chan, Clark, and Koop, 2018). Moreover, it finds support in the data – about half of the variance of inflation occurs at the low frequency (Table 1).

Accordingly, we assess the forecasting performance of the model $\hat{\pi}^{h}_{t+h} = \hat{\pi}^{h,LF}_{t+h}$, *i.e.* $\kappa_f = 0 \, \forall f = HF, BCF, MF$ in (12). That is, we assess whether the forecast of inflation may be improved by using only the forecast of its low-frequency component. As in the more general case, we consider a model strictly comprising the forecasts from the low frequencies of the predictors, which we denote NKPC_WAV_diag (LF), and a model in which the whole elements of the last row of matrices $\boldsymbol{\alpha^h_m}$, with $m = 1, 2, 3$, are allowed to be non null, which we denote NKPC_WAV_all (LF). This is the generalized version of our low-frequency forecast model. It allows for influence from the high-, business-cycle- and medium-frequency fluctuations of the NKPC predictors into the low-frequency of inflation, which, as Table 1 suggests, 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.

5.2 Results

5.2.1 Forecast precision

We now compare the ability of the wavelet-based NKPC to forecast inflation out-of-sample (OOS) in the period 2000Q1–2020Q4. Following common practice, we assess the forecast accuracy by computing the root mean squared error (RMSFE) for each model and computing the statistical significance of differences in RMSFE across methods. Table 3 reports our results.

Panel (a) shows the RMSFE of the AO random walk model and the UCSV model, the standard benchmarks for inflation forecasts. Panels (b) and (c) report the RMSFE of the time-series- and wavelet-based NKPCs, respectively, relative to that of the AO model. A value below 1 indicates the model outperforms the AO benchmark. Panel (d) reports the RMSFEs of the wavelet-based NKPC models relative to those of the time-series NKPC. Asterisks indicate statistical significance according to the Diebold and Mariano (1995) test of relative predictive accuracy at the 10 % (*), 5 % (**), and 1 % (***) levels.

The UCSV model yields inflation forecasts that are noticeably better than those of the AO model only at the 1 quarter-ahead horizon. Despite being much more sophisticated than the AO model, the rather limited performance of the UCSV model at longer horizons is in line with the literature (see e.g. Faust and Wright, 2013 and Jarocinski and Lenza, 2018).[11]

Panel (b) confirms that the traditional time-series NKPC fails to outperform the AO model.

The first lines of panels (c) and (d) indicate that the forecasts of the wavelet-based NKPC model that does not allow for interaction across frequencies (NKPC_WAV_diag) usually outperform those of the benchmark, as well as those of the time-series NKPC. However, the difference is statistically significant only relative to the time-series NKPC at the 8-quarter horizon. The second lines

[11] The UCSV model typically outperforms the AO model only in samples of highly volatile inflation (in this case, for h=1). In samples with relatively smooth time series, the estimated trend of the UCSV model is close to actual inflation. Thus, its inflation forecasts are close to the previous period.

of panels (c) and (d) indicate that the generalized wavelet-based NKPC model (NKPC_WAV_all, allowing for interactions between all frequencies) does not perform significantly better than the AO and the NKPC_TS models.

The third and fourth lines of panels (c) and (d) report the key results of this sub-section. Consistent with our conjecture, focusing solely on the low-frequency component of inflation (LF) produces forecasts for inflation that are substantially and statistically more accurate than those of both the AO and the NKPC_TS model.

The RMSFEs of the NKPC_WAV_diag (LF) model are 80 %, 77 % and 79 % of those of the AO benchmark, at the 1-, 4-, and 8-quarter horizons respectively, all statistically significant at the 5 % level. In turn, they amount to 83 %, 75 % and 55 % of the RMSFEs of the NKPC_TS at h=1, h=4 and h=8, respectively, all significant at the 10 % level and the latter significant at 5 %.

The generalized low-frequency model, NKPC_WAV_all (LF), produces even better forecasts than the NKPC_WAV_diag (LF) model: allowing influences from higher frequencies (HF, BCF and MF) of the NKPC inflation determinants produces wavelet-based NKPC forecasts of the low-frequency of inflation that are slightly even closer to actual inflation. The RMSFEs from the NKPC_WAV_all (LF) model are 79 %, 76 % and 77 % of those of the AO benchmark (h=1, h=4, h=8, respectively), which are all significant at the 5 % level; and they are 83 %, 74 % and 54 % of those of the time-series NKPC (significant at 10 % for h=1 and h=4, and significant at the 5 % level for h=8).

Overall, the wavelet-based NKPC forecasts of the low-frequency of inflation are substantially and significantly better forecasts of inflation than time series benchmarks and the standard time-series NKPC. Our frequency-domain approach to the NKPC effectively resurrects the forecast ability of the Phillips curve. The gain in forecast accuracy is above 20 % with regard to the naive time series benchmark. It is as high as 25 % at the 4-quarter horizon and 45 % at the 8-quarter horizon, regarding the time-series NKPC forecasts.

Figure 7 shows actual inflation, the forecasts obtained with our preferred model NKPC_WAV_all (LF), and those from the time-series NKPC. The performance of the time-series NKPC deteriorates markedly as the forecast horizon increases (to h=4 and h=8) and after the Great Recession, when forecasts are consistently overshooting actual inflation. Our in-sample results suggest that such poor forecasting performance comes from a combination of two factors: the steadiness of inflation expectations and its determinant role in that model, and the minor relevance of the unemployment gap, indicated by the low estimate of α_2.

The favorable performance of our model comes from the smoothness of its forecasts, avoiding excessively sharp fluctuations while still capturing the essence of the evolution of inflation over time. The figure shows that the forecast gain, compared to the time-series NKPC, is less evident for the 1-quarter ahead horizon, but remains notable for the 4-quarter and 8-quarter horizons, especially after the Great Recession. In the following sub-sections, we explore these results in detail. In the next sub-section, we assess whether our model outperforms the time-series NKPC systematically or only in specific episodes. In the sub-section after that, we assess which NKPC determinants of inflation drive the forecast success of our model.

5.2.2 Forecast timing

To explore the timing of the outperformance of the NKPC_WAV_all (LF) model versus both the AO model and the time-series NKPC model, in Figure 8 we show the cumulative differences between the squared forecast errors (SFE) of the NKPC_WAV_all (LF) model and those of the AO model (left-hand side charts) and the NKPC_TS (right-hand side charts) through the OOS period. The top charts show cumulative differences in SFE for 1-quarter-ahead forecasts, the middle charts relate to the 4-quarters-ahead forecasts, and the bottom charts relate to the 8-quarters-ahead forecasts.

In the plots in Figure 8, a rising line indicates the predictive regression of the NKPC_WAV_all

(LF) model outperforms the alternative model at the relevant forecast horizon.

A broad conclusion drawn from Figure 8 is that, along the 21 years of the OOS forecasting exercise, our NKPC_WAV_all (LF) model rarely underperforms any of the alternative models. Indeed, the cumulative difference of SFEs rarely decreases for all forecast horizons and relative to both models. Furthermore, the temporal pattern of forecasting outperformance of our model is rather similar for all forecast horizons, irrespective of the alternative model.

For 1-quarter-ahead forecasts, the relative performance of our model is particularly favorable in the late stages of the Great Recession and at the height of the 2020 Covid recession. At the 4-quarter-ahead horizon, our model beats both alternative models in the early stages of the recovery from the Great Recession until about 2012 in the case of the AO and until about 2017 for the time-series NKPC. At the 8-quarter-horizon, our NKPC_WAV_all (LF) model forecasts inflation much better than both benchmarks during the entire recovery from the Great Recession up to 2015 in the case of the AO model and up to 2017 for the time-series NKPC.

Overall, it is particularly noteworthy that a method based on long-run movements of the NKPC variables markedly improves 1-quarter-ahead forecasts of inflation in periods of high instability and uncertainty – the Great Recession and Covid recession – with no costs in its performance for longer forecast horizons following those episodes.

5.2.3 Forecast determinants

Having established that our low-frequency NKPC significantly and systematically outperforms the relevant benchmarks in forecasting inflation, we now ask whether these results resurrect the NKPC. Recall that, in sub-section 4.2, we found that the low-frequency of inflation is essentially explained, in-sample, by the low-frequency of expectations. If this was the only driver of our out-of-sample forecasting results, then the Phillips tradeoff would remain irrelevant.

In Table 4 we present in further detail the RMSFE of our low-frequency NKPC models relative to those of the AO model. The last column shows the relative RMSFE of our NKPC_WAV_all (LF) model, and the penultimate column shows the relative RMSFE of the NKPC_WAV_diag (LF) model. Given that the results are similar, we focus on the latter for simplicity of presentation. In the first three columns, we show the relative RMSFE that would be obtained if we restricted the NKPC determinants to the low frequency of expectations (first column), of expectations and unemployment (second column), or of expectations and energy prices (third column).

The comparison of the first and the second columns of Table 4 allows us to assess the contribution of the low-frequency component of the unemployment gap above that of the low-frequency component of expectations. At the 1-quarter-ahead forecast horizon, adding the unemployment gap to a NKPC_WAV_diag (LF) model restricted to expectations reduces its relative RMSFE from 83 % to 80 %. At the 4-quarter-ahead horizon, it decreases the relative RMSFE from 89 % to 78 %. At the 8-quarter-ahead forecast horizon, it cuts the relative RMSFE from 109 % to 78 %. Furthermore, the forecast gain given by the low-frequency of the unemployment gap is not just economically relevant, but statistically significant at the 5 % level all forecasting horizons. The low-frequency of energy prices, on the other hand, adds little, if anything, to the forecasts.

To see whether these findings apply to the entire OOS period, in Figure 9 we present, for the three forecast horizons, actual inflation (black lines) and the forecasts from the NKPC_WAV_diag (LF) model (blue lines), from that model restricted to $\pi^{e,LF}$ (red lines), from that model restricted to $\pi^{e,LF}$ and en_t^{LF} (yellow lines), and from that model restricted to π^{e,D_6} and $ugap_t^{LF}$ (green lines).

Figure 9 indicates that the gain in forecast accuracy given by the inclusion of the low-frequency of the unemployment gap essentially starts at 2009 for h=1, at 2010 for h=4, and at 2011 for h=8. There is also some gain in forecast accuracy at the 4-quarter-ahead and 8-quarter-ahead forecast horizons in the first three to four years of the OOS period.

Overall, we conclude that the unemployment gap is of crucial importance to forecast inflation with a frequency-domain NKPC. The key role of the low-frequency of unemployment may be clearly understood by looking with further detail at the wake of the Great Recession. Inflation falls substantially between 2011Q3 and 2016Q3 from 2.6 % to 1.2 %. In the same period, the low-frequency of unemployment increases from -1.1 % to 1.7 %, and the estimates of α_2^{LF} (reported in Figure 5) become increasingly negative. As a result, the low-frequency of unemployment decisively reduces the inflation forecasts to levels closer to actual inflation.

5.3 Robustness

We submitted our OOS forecast procedure to several robustness checks. First, we considered alternative proxies for slack: the unemployment rate (in levels) and the output gap (computed as the difference between real GDP and its linear trend). Second, we forecast personal consumption expenditure (PCE) inflation using the corresponding component for energy price inflation. Third, we experimented with alternative wavelet filters such as Daubechies and Coiflets of different lengths. Fourth, we computed the forecasts using rolling window estimates rather than expanding window estimates (with a window size of 88 quarters, the same as for our initial in-sample period). Fifth, we considered expanding windows starting in 1985Q1, to avoid the period of disinflation and skip the potential structural break at the start of the Great Moderation.[12]

Results are reported in Table 5. Overall, the results of these robustness analyses indicate that our conclusions are qualitatively – and often quantitatively – robust to all these changes.[13]

[12] Given the reduction of the sample size, we use a J=4 level MODWT decomposition in this simulation. Therefore, the results reported in the last two lines of panels c) and d) of Table 5 relate to forecasts of a smooth component of inflation that captures fluctuations with a period longer than 8 years. These are not strictly comparable to those reported in the paper and in the other robustness checks, where the low-frequency component of inflation corresponds to cycles with a period longer than 16 years.

[13] We ran several additional checks that we do not report here for the sake of brevity. The results were quite similar. In particular, we ran forecasts i) for other measures of slack (quadratic detrended unemployment, CBO detrended unemployment, quadratic detrended real GDP, CBO detrended real GDP), and ii) with the original MODWT frequency decomposition rather than using the sum of D_1 and D_2 for HF, and of D_3 and D_4 for BCF. The results are available upon request.

6 Concluding remarks

In this paper, we used a frequency-domain approach to an empirical New Keynesian Phillips Curve (NKPC) to study the in-sample dynamics of inflation and forecast inflation out-of-sample in a consistent framework.

With regard to the dynamics of inflation, we find that the unemployment gap is statistically significant at most cyclical frequencies. There are no signs of a flattening of the NKPC in the period 2000Q1–2020Q4 at business-cycle, medium-run, or long-run frequencies. Unemployment is, however, not economically relevant and the dynamics of inflation is explained essentially by expectations and energy inflation, with inflation expectations dominating in long cycles and energy inflation in short cycles. Furthermore, at business and medium-run cycles, energy inflation becomes dominant after 2000, when expectations become anchored.

We show that the key to obtaining a good forecast of inflation is to forecast the low frequency of inflation well and ignore its other frequencies. Furthermore, unlike for the in-sample analysis, the low-frequency component of the unemployment gap turns out to be crucial to realizing these forecasting gains.

So is the Phillips curve alive and well? We argue that it depends on the cyclical frequency and whether the NKPC is used to describe the dynamics of inflation or to forecast inflation.

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.

Angeletos, G.-M., F. Collard, and H. Dellas (2020): "Business-Cycle Anatomy," *American Economic Review*, 110(10), 3030–3070.

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.

——— (2020): "Finding a Stable Phillips Curve Relationship: A Persistence-Dependent Regression Model," Working Papers 201909R, Federal Reserve Bank of Cleveland.

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.

Ball, L., S. Mazumder, K. Dynan, and J. H. Stock (2011): "Inflation Dynamics and the Great Recession/Comments and Discussion," *Brookings Papers on Economic Activity*, p. 337.

Barnichon, R., and G. Mesters (2021): "The Phillips multiplier," *Journal of Monetary Economics*, 117, 689–705.

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.

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.

Bundick, B., and A. L. Smith (2020): "Did the Federal Reserve Break the Phillips Curve? Theory and Evidence of Anchoring Inflation Expectations," *Federal Reserve Bank of Kansas City Working Paper*, (20-11).

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.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2021): "Anchored inflation expectations," *Available at SSRN 3018198*.

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.

Chan, J. C. C. (2018): "Specification tests for time-varying parameter models with stochastic volatility," *Econometric Reviews*, 37(8), 807–823.

Cogley, T., and A. M. Sbordone (2008): "Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve," *American Economic Review*, 98(5), 2101–2126.

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.

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.

Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti (2020): “What's up with the Phillips Curve?,” *Brookings Papers on Economic Activity*, 1(Spring), 301–357.

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.

Engle, R. F. (1974): “Band spectrum regression,” *International Economic Review*, 15, 1–11.

Faria, G., and F. Verona (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*, 50(C).

——— (2021): “Out-of-sample time-frequency predictability of the equity risk premium,” *Quantitative Finance*, forthcoming.

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.

Fratianni, M., M. Gallegati, and F. Giri (2019): "Mr Phillips and the medium-run: temporal instability vs. frequency stability," *Mo.Fi.R. Working Papers*, 155.

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.

Gali, J., and M. Gertler (1999): "Inflation dynamics: A structural econometric analysis," *Journal of Monetary Economics*, 44(2), 195–222.

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.

Gordon, R. J. (2011): "The History of the Phillips Curve: Consensus and Bifurcation," *Economica*, 78(309), 10–50.

Hasenzagl, T., F. Pellegrino, L. Reichlin, and G. Ricco (2021): "A Model of the Fed's View on Inflation," *Review of Economics and Statistics*, forthcoming.

Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson (2020): "The Slope of the Phillips Curve: Evidence from U.S. States," NBER Working Papers 28005, National Bureau of Economic Research.

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.

Hooper, P., F. S. Mishkin, and A. Sufi (2020): "Prospects for inflation in a high pressure economy: Is the Phillips curve dead or is it just hibernating?," *Research in Economics*, 74(1), 26–62.

Jarocinski, M., and M. Lenza (2018): "An Inflation-Predicting Measure of the Output Gap in the Euro Area," *Journal of Money, Credit and Banking*, 50(6), 1189–1224.

Jorgensen, P., and K. J. Lansing (2019): "Anchored Inflation Expectations and the Flatter Phillips Curve," Working Paper Series 2019-27, Federal Reserve Bank of San Francisco.

King, R. G., and M. W. Watson (1994): "The post-war US Phillips curve: a revisionist econometric history," in *Carnegie-Rochester Conference Series on Public Policy*, vol. 41, pp. 157–219.

——— (2012): "Inflation and unit labor cost," *Journal of Money, credit and Banking*, 44, 111–149.

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.

McLeay, M., and S. Tenreyro (2020): "Optimal inflation and the identification of the Phillips curve," *NBER Macroeconomics Annual*, 34(1), 199–255.

Pfajfar, 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.

Reinbold, B., and Y. Wen (2020): "Is the Phillips curve still alive?," *Review-Federal Reserve Bank of St. Louis*, 102(2), 121.

Roberts, J. M. (1995): "New Keynesian Economics and the Phillips Curve," *Journal of Money, Credit and Banking*, 27(4), 975–984.

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.

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

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.

——— (2010): "Modeling inflation after the crisis," in *Proceedings-Economic Policy Symposium-Jackson Hole*, pp. 173–220. Federal Reserve Bank of Kansas City.

——— (2020): "Slack and cyclically sensitive inflation," *Journal of Money, Credit and Banking*, 52(S2), 393–428.

Verona, F. (2020): "Investment, Tobin's Q, and Cash Flow Across Time and Frequencies," *Oxford Bulletin of Economics and Statistics*, 82(2), 331–346.

Woodford, M. (2003): *Interest and prices: Foundations of a theory of monetary policy*. Princeton University Press.

Yellen, J. (2016): "Macroeconomic Research After the Crisis: speech at the 60th annual economic conference sponsored by the Federal Reserve Bank of Boston, October 14, 2016," .

Zaman, S. (2019): "Cyclical versus Acyclical Inflation: A Deeper Dive," *Economic Commentary*, (2019-13).

Zhang, K., R. Gençay, and M. E. Yazgan (2017): "Application of wavelet decomposition in time-series forecasting," *Economics Letters*, 158, 41–46.

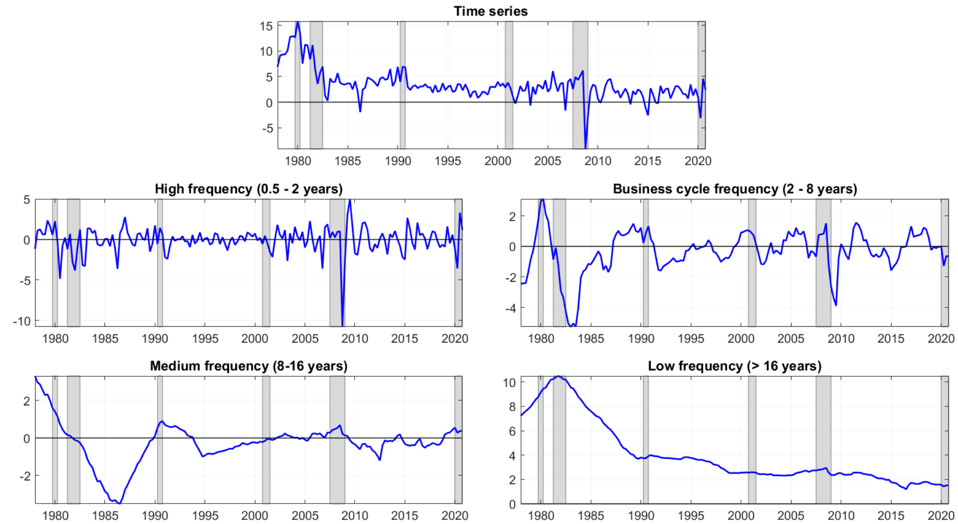

Figure 1: Inflation: time series and frequency decomposition

Top graph: time series of inflation, US 1978Q1–2020Q4. Remaining graphs: four frequency-specific time series resulting from the frequency decomposition of the inflation time series. Gray bars denote NBER-dated recessions.

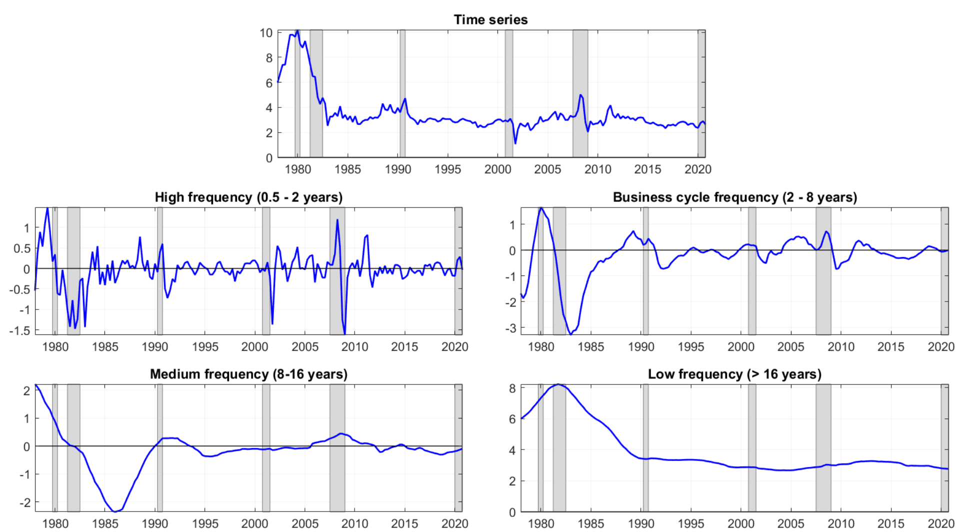

Figure 2: Inflation expectations: time series and frequency decomposition

Top graph: time series of inflation expectations, US 1978Q1–2020Q4. Remaining graphs: four frequency-specific time series resulting from the frequency decomposition of the inflation expectations time series. Gray bars denote NBER-dated recessions.

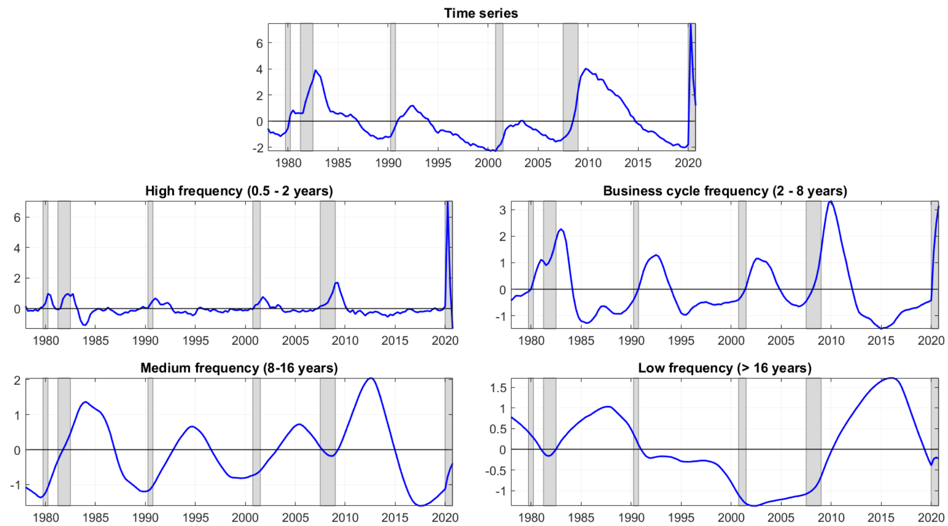

Figure 3: Unemployment gap: time series and frequency decomposition

Top graph: time series of unemployment gap, US 1978Q1–2020Q4. Remaining graphs: four frequency-specific time series resulting from the frequency decomposition of the unemployment gap time series. Gray bars denote NBER-dated recessions.

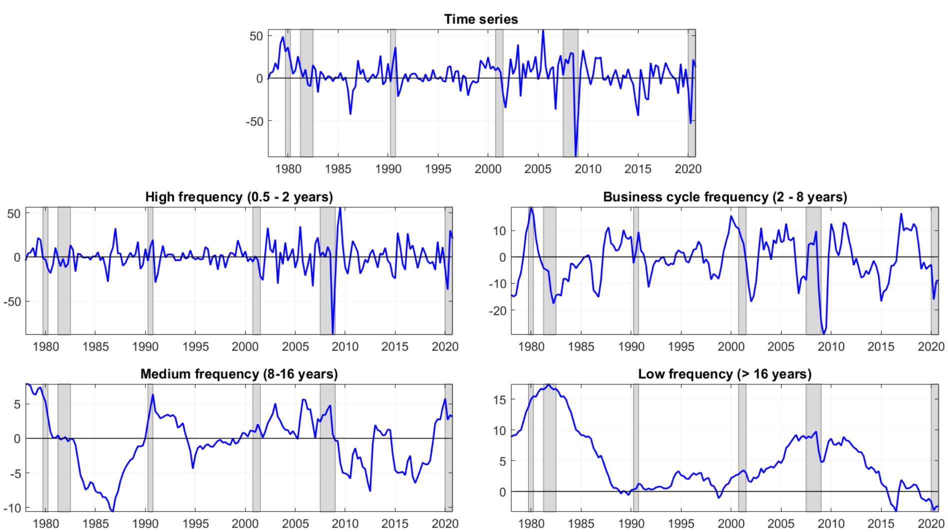

Figure 4: Energy inflation: time series and frequency decomposition

Top graph: time series of energy inflation, US 1978Q1–2020Q4. Remaining graphs: four frequency-specific time series resulting from the frequency decomposition of the energy inflation time series. Gray bars denote NBER-dated recessions.

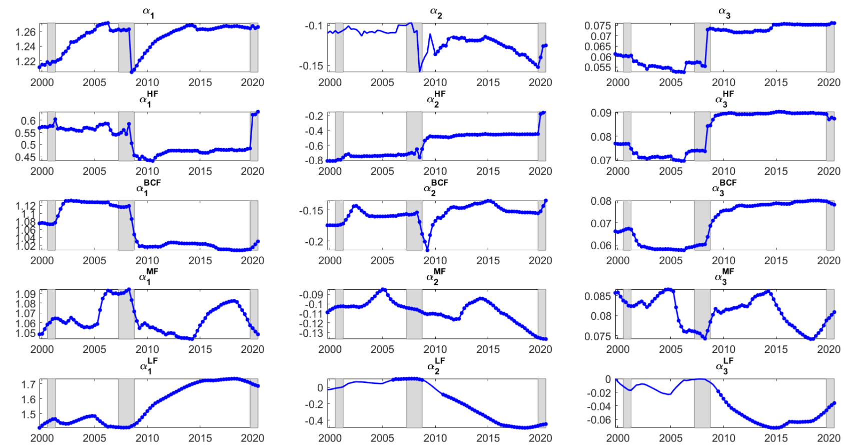

Figure 5: Estimated coefficients of the New Keynesian Phillips Curve over time and across frequencies

Recursive estimates of the coefficients of the Phillips curve ((1)) for the original time series data (upper row) and filtered data for different frequency bands (remaining rows). HF: high frequency, cycles with periods between 2 and 8 quarters; BCF: business cycle frequency, cycles with periods between 2 and 8 years; MF: medium frequency, cycles with periods between 8 and 16 years; LF: low frequency, cycles longer than 16 years. Sample periods are expanding windows starting in 1978Q1–1999Q4, recursively including one additional quarter through 2020Q4. Gray bars denote NBER-dated recessions. Statistically significant coefficients (at 5 %) are reported with a circled marker.

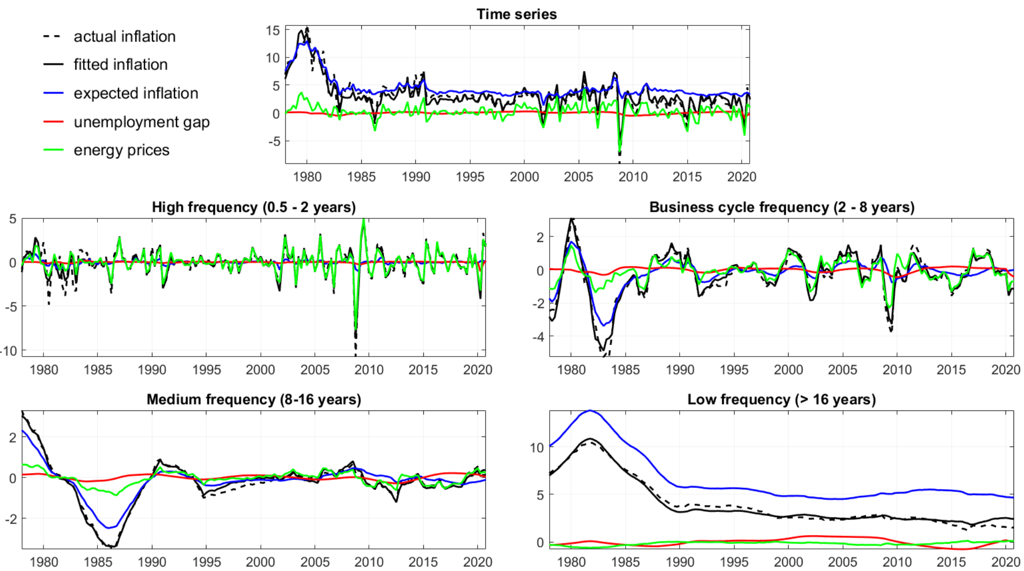

Figure 6: Phillips curve-based decomposition of inflation

Notes: The lines show the contributions of each variable for the period 1978Q1–2020Q4. Constant terms omitted. Top graph: decomposition of the time series of inflation. Remaining graphs: decomposition of four frequency-specific time series of inflation.

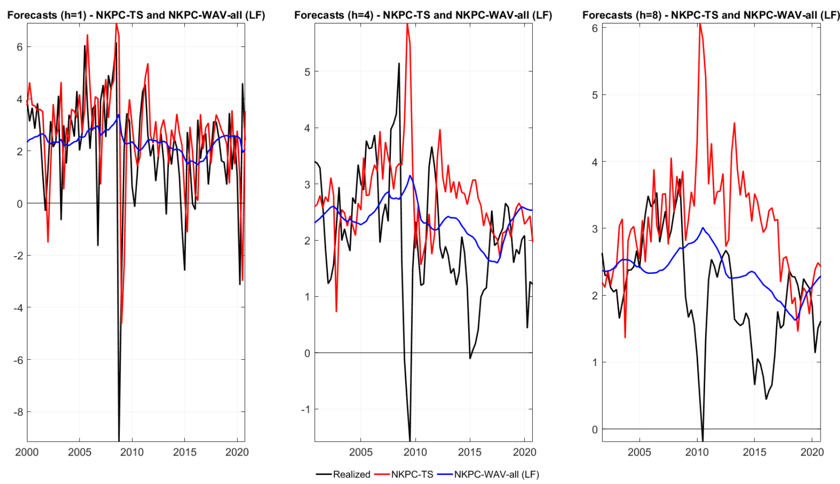

Figure 7: New Keynesian Phillips Curve forecasts for different forecasting horizons: time series vs wavelet

Forecasts for h=1 (left) , h=4 (middle), and h=8 (right). Black lines: realized inflation. Red lines: forecasts with the NKPC_TS model. Blue lines: forecasts with the NKPC_WAV_all (LF) model. The out-of-sample period is 2000Q1–2020Q4.

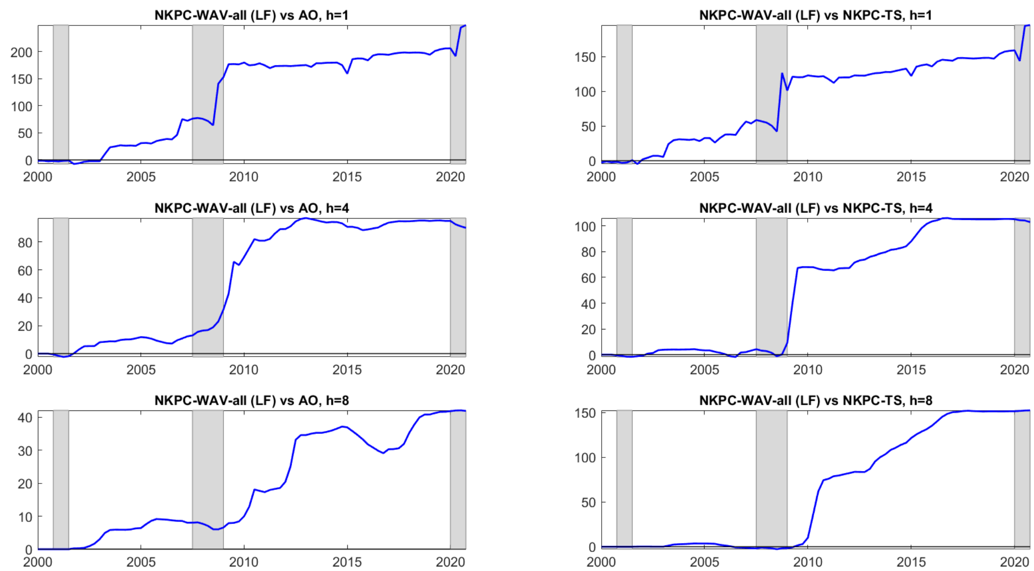

Figure 8: Cumulative differences in squared forecast errors

Left side graphs: cumulative difference between the squared forecast errors of the NKPC_WAV_all (LF) model and those of the AO model, for h=1 (top) , h=4 (middle), and h=8 (bottom). Right side graphs: cumulative difference between the squared forecast errors of the NKPC_WAV_all (LF) model and those of the NKPC_TS model, for h=1 (top) , h=4 (middle), and h=8 (bottom). Gray bars denote NBER-dated recessions.

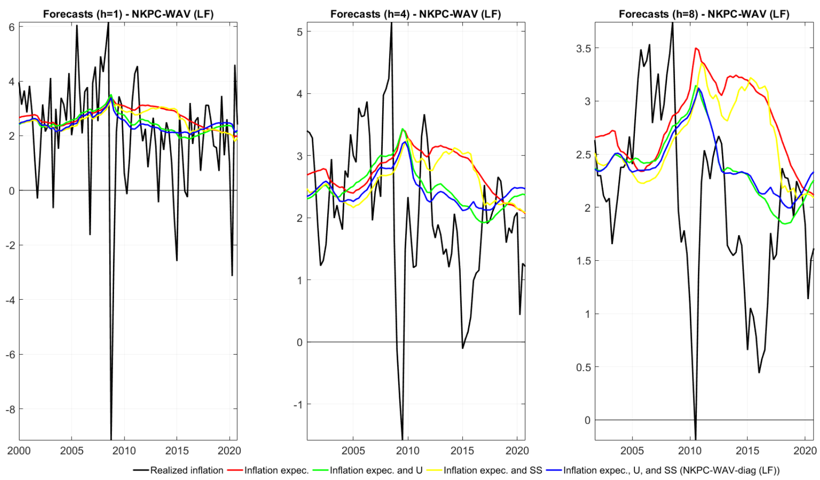

Figure 9: New Keynesian Phillips Curve forecasts for different forecasting horizons with wavelet: importance of different variables

Forecasts of π^{LF} for h=1 (left) , h=4 (middle), and h=8 (right). Black lines: realized inflation. Red lines: forecasts using $\pi^{e,LF}$. Green lines: forecasts using $\pi^{e,LF}$ and $ugap^{LF}$. Yellow lines: forecasts using $\pi^{e,LF}$ and en^{LF}. Blue lines: forecasts with the NKPC_WAV_diag (LF) model. The out-of-sample period is 2000Q1–2020Q4.

	HF	BCF	MF	LF
Inflation	27	17	10	47
Inflation expectations	5	16	16	63
Unemployment gap	14	32	28	26
Energy inflation	66	24	4	6

Table 1: 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. 1978Q1-2020Q4. HF: high frequency, cycles with periods between 2 and 8 quarters; BCF: business cycle frequency, cycles with periods between 2 and 8 years; MF: medium frequency, cycles with periods between 8 and 16 years; LF: low frequency, cycles longer than 16 years. Percentages may not add up to 100 due to rounding.

	c	α_1	α_2	α_3	$\bar{R}^2$
NKPC	-1.39 (0.19)	1.27 (0.05)	-0.13 (0.05)	0.08 (0.005)	0.89
NKPC_HF	-0.03 (0.05)	0.63 (0.14)	-0.15 (0.08)	0.09 (0.004)	0.82
NKPC_BCF	-0.05 (0.03)	1.03 (0.04)	-0.13 (0.02)	0.08 (0.003)	0.94
NKPC_MF	-0.07 (0.01)	1.05 (0.02)	-0.14 (0.01)	0.08 (0.005)	0.98
NKPC_LF	-2.43 (0.10)	1.69 (0.03)	-0.45 (0.04)	-0.04 (0.01)	0.98

Table 2: Estimates of the New Keynesian Phillips Curve

Estimates of equation (1); US data, sample period 1978Q1–2020Q4. 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 with periods between 2 and 8 quarters; BCF: business cycle frequency, cycles with periods between 2 and 8 years; MF: medium frequency, cycles with periods between 8 and 16 years; LF: low frequency, cycles longer than 16 years. Standard errors in parenthesis.

Model	Forecasting horizon h=1	h=4	h=8
a) Benchmarks			
AO random walk	2.81	1.63	1.16
UCSV	2.33***	1.62**	1.16
b) Time-series Phillips Curve (vs AO)			
NKPC_TS	0.96	1.03	1.44
c) Wavelet-based Phillips Curves (vs AO)			
NKPC_WAV_diag	0.95	0.94	0.98
NKPC_WAV_all	0.97	1.00	1.09
NKPC_WAV_diag (LF)	0.80**	0.77**	0.79**
NKPC_WAV_all (LF)	0.79**	0.76**	0.77**
d) Wavelet-based Phillips Curves (vs NKPC_TS)			
NKPC_WAV_diag	0.99	0.91	0.68**
NKPC_WAV_all	1.01	0.97	0.76
NKPC_WAV_diag (LF)	0.83*	0.75*	0.55**
NKPC_WAV_all (LF)	0.83*	0.74*	0.54**

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

Panel a): Root mean squared forecast errors (RMSFEs) at different forecasting horizons (h=1, h=4, and h=8) for the AO model and the UCSV model. Panels b) and c): RMSFEs relative to those of the AO model ($RMSFE_j$ / $RMSFE_{AO}$ for model j). Panel d): RMSFEs relative to those of the NKPC_TS model ($RMSFE_j$ / $RMSFE_{NKPC\text{-}TS}$ 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 2000Q1–2020Q4.

				WAV_diag (LF)	WAV_all (LF)
predictors	$\pi^{e,LF}$	$\pi^{e,LF}, ugap_t^{LF}$	$\pi^{e,LF}, en_t^{LF}$	$\pi^{e,LF}, ugap_t^{LF}, en_t^{LF}$	
forecasting horizon					
h=1	0.83*	0.80**	0.83*	0.80**	0.79**
h=4	0.89	0.78**	0.87	0.77**	0.76**
h=8	1.09	0.78**	1.03	0.79**	0.77**

Table 4: On the importance of each low frequencies of the predictors

RMSFEs relative to those of the AO model. Asterisks indicate statistical significance of the Diebold and Mariano (1995) test of comparative predictive accuracy at the 10 % (*) and 5 % (**) levels, relative to the AO model. The out-of-sample period is 2000Q1–2020Q4.

	Measure of slack						Measure of inflation			Wavelet filter		
	Unemployment rate			Output gap			PCE inflation			Daubechies length 4		
	Forecasting horizon			Forecasting horizon			Forecasting horizon			Forecasting horizon		
Model	h=1	h=4	h=8	h=1	h=4	h=8	h=1	h=4	h=8	h=1	h=4	h=8
a) Benchmarks												
AO random walk	2.81	1.63	1.16	2.81	1.63	1.16	1.93	1.23	0.91	2.81	1.63	1.16
UCSV	2.33***	1.62**	1.16	2.33***	1.62**	1.16	1.83	1.51**	1.12**	2.33***	1.62**	1.16
b) Time-series Phillips Curve (vs AO)												
NKPC_TS	0.96	1.02	1.34	0.94*	0.99	1.47	1.02	1.15	1.61*	0.96	1.03	1.44
c) Wavelet-based Phillips Curves (vs AO)												
NKPC_WAV_diag	0.95	0.93	0.96	0.95	0.94	1.05	0.98	0.98	1.01	0.96	0.99	1.00
NKPC_WAV_all	0.97	1.01	1.11	0.96	0.98	1.06	1.01	1.05	1.19	0.97	1.04	1.10
NKPC_WAV_diag (only LF)	0.80**	0.78*	0.81*	0.79**	0.75**	0.78**	0.86*	0.84	0.91	0.80**	0.79**	0.81*
NKPC_WAV_all (only LF)	0.79**	0.77**	0.79**	0.79**	0.74**	0.77**	0.84*	0.81	0.86	0.80**	0.79**	0.81*
d) Wavelet-based Phillips Curves (vs NKPC_TS)												
NKPC_WAV_diag	0.99	0.91*	0.72***	1.01	0.95	0.72**	0.96	0.85**	0.63**	1.00	0.96	0.70*
NKPC_WAV_all	1	0.99	0.83	1.03	0.99	0.72	0.99	0.91	0.74	1.02	1.01	0.77
NKPC_WAV_diag (only LF)	0.83*	0.77*	0.61**	0.84*	0.76	0.53**	0.84	0.73	0.57**	0.83*	0.76	0.57**
NKPC_WAV_all (only LF)	0.82*	0.75*	0.59**	0.84*	0.75	0.52**	0.82	0.71*	0.54**	0.84*	0.77	0.56**

Table 5: Relative out-of-sample root mean squared forecast errors - robustness checks

Panel a): Root mean squared forecast errors (RMSFEs) at different forecasting horizons (h=1, h=4 and h=8) for the AO model and the UCSV model. Panels b) and c): RMSFEs relative to those of the AO model ($RMSFE_j$ / $RMSFE_{AO}$ for model j). Panel d): RMSFEs relative to those of the NKPC_TS model ($RMSFE_j$ / $RMSFE_{NKPC\text{-}TS}$ 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 2000Q1-2020Q4.

	In-sample length			**In-sample start**		
	Rolling window			1985:Q1		
	Forecasting horizon			Forecasting horizon		
Model	h=1	h=4	h=8	h=1	h=4	h=8
a) Benchmarks						
AO random walk	2.81	1.63	1.16	2.81	1.63	1.16
UCSV	2.33***	1.62**	1.16	2.33***	1.62**	1.16
b) Time-series Phillips Curve (vs AO)						
NKPC_TS	0.98	1.00	1.05	0.98	1.05	1.17
c) Wavelet-based Phillips Curves (vs AO)						
NKPC_WAV_diag	0.96	0.98	0.97	0.95	1	1.04
NKPC_WAV_all	0.97	1.00	1.03	0.95*	0.96	0.99
NKPC_WAV_diag (only LF)	0.79**	0.74**	0.76**	0.84**	0.88	0.94
NKPC_WAV_all (only LF)	0.79**	0.74**	0.76**	0.82**	0.83**	0.86*
d) Wavelet-based Phillips Curves (vs NKPC_TS)						
NKPC_WAV_diag	0.99	0.98	0.92	0.98	0.96	0.89*
NKPC_WAV_all	1	1.01	0.98	0.97	0.92*	0.85
NKPC_WAV_diag (only LF)	0.81*	0.75*	0.73	0.86	0.84	0.81*
NKPC_WAV_all (only LF)	0.81*	0.74	0.73	0.84*	0.80*	0.73**

Table 5: continue

Panel a): Root mean squared forecast errors (RMSFEs) at different forecasting horizons (h=1, h=4 and h=8) for the AO model and the UCSV model. Panels b) and c): RMSFEs relative to those of the AO model ($\mathrm{RMSFE_j}$ / $\mathrm{RMSFE_{AO}}$ for model j). Panel d): RMSFEs relative to those of the NKPC_TS model ($\mathrm{RMSFE_j}$ / $\mathrm{RMSFE_{NKPC\text{-}TS}}$ 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 2000Q1-2020Q4.