

Assessing U.S. Aggregate Fluctuations Across Time and Frequencies*

Thomas A. Lubik
Federal Reserve Bank of Richmond[†]

Christian Matthes
Federal Reserve Bank of Richmond[‡]

Fabio Verona
Bank of Finland[§]

February 2019

Abstract

We study the behavior of key macroeconomic variables in the time and frequency domain. For this purpose, we decompose U.S. time series into various frequency components. This allows us to identify a set of stylized facts: GDP growth is largely a high-frequency phenomenon whereby inflation and nominal interest rates are characterized largely by low-frequency components. In contrast, unemployment is a medium-term phenomenon. We use these decompositions jointly in a structural VAR where we identify monetary policy shocks using a sign restriction approach. We find that monetary policy shocks affect these key variables in a broadly similar manner across all frequency bands. Finally, we assess the ability of standard DSGE models to replicate these findings. While the models generally capture low-frequency movements via stochastic trends and business cycle fluctuations through various frictions they fail at capturing the medium-term cycle.

JEL CLASSIFICATION: C32, C51, E32

KEY WORDS: Wavelets, bandpass filter, SVAR, sign restrictions, DSGE model

*We wish to thank Bob King, Annika Lindblad, and Denis Tkachenko for useful comments. We are also grateful to participants at the 2018 CEF Conference in Milan, the University of Adelaide, Carleton University, the 2019 Meeting of the Finnish Economic Association and the Bank of Finland. James Geary provided exceptional research assistance. The views expressed in this paper are those of the authors and should not be interpreted as those of the Federal Reserve Bank of Richmond, the Federal Reserve System, or the Bank of Finland.

[†]Research Department, P.O. Box 27622, Richmond, VA 23261. Email: thomas.lubik@rich.frb.org.

[‡]Research Department, P.O. Box 27622, Richmond, VA 23261. Email: christian.matthes@rich.frb.org

[§]Monetary Policy and Research Department. Snellmaninaukio, PO Box 160, 00101 Helsinki. Email: fabio.verona@bof.fi.

1 Introduction

Economists have often found it useful to separate long-run trends from business cycle fluctuations, which generally are considered those that occur with a cycle length of between two and eight years. On the statistical side, this approach is probably best characterized by the idea of a trend-cycle decomposition as in Beveridge and Nelson (1981), where the trend is associated with permanent movements in a time series as opposed to a business cycle being driven by transitory shocks. Conceptually, this idea is also inherent in filtering methods such as the Hodrick-Prescott (HP) filter, which has been the dominant approach in business cycle modeling to extract a trend from aggregate times and render them stationary. Such decompositions are convenient since they align with the idea of economic fluctuations as being driven by either permanent or temporary shocks that do not necessarily interact. In addition, monetary policy is often framed in terms of stabilizing the fluctuations of key variables around a trend that is unaffected by policy.

However, there is a growing awareness in the macroeconomics literature that this common view of economic fluctuations is no longer adequate to characterize the behavior of economic activity over time. For instance, Comin and Gertler (2006) argue that a substantial part of economic fluctuations is located in what they label a ‘medium-term cycle’, that is, fluctuations beyond a length of eight years, but falling short of a trend. Moreover, these medium-term fluctuations cannot be thought of in isolation of other frequency bands. Using a theoretical model, Comin and Gertler (2006) show that business cycles and medium-term cycles are intimately connected since they are driven by the same underlying temporary shock. Specifically, a temporary innovation to, say, productivity or the policy rate can reverberate throughout several frequency bands as they get propagated over time.¹

Against this background, we aim to provide a somewhat more encompassing view of cyclical behavior across all frequencies. In particular, we study three issues. First, we compute a decomposition of key macroeconomic time series using wavelet-based filtering. That is, we decompose a time series into several time series components, each of them fluctuating within a specific frequency band. We find the use of wavelets advantageous for our purposes since this filtering approach is more flexible than standard Fourier analysis and more traditional bandpass filtering. In particular, it allows different frequency movements to be more pronounced in some parts of the sample than others and thereby reveals time

¹Cogley (2001) makes a similar point for trend specifications where he shows the effects of trend specification errors are not confined to low frequencies, but are spread across the entire frequency domain. Researchers therefore have to have a clear understanding of the inter-relatedness of frequency bands for which a wavelet approach offers a convenient tool.

variation in the importance of different frequency components. The second question looks at the effects of identified monetary policy shocks across different frequency bands to assess the plausibility of medium-term cycles as being generated by temporary shocks. The third question asks whether standard dynamic stochastic general equilibrium (DSGE) models that are used in monetary policy analysis can replicate the volatility of different cycles of each macroeconomic variable under consideration and are thereby useful in addressing the policy questions raised.

We establish three main findings. First, the wavelet decomposition of key macroeconomic variables shows that the bulk of fluctuations in GDP growth, unemployment, and inflation occurs across different frequency bands. More than half of real GDP growth is explained by short-term, high-frequency components with only a third of fluctuations attributable to business cycle frequencies between two and eight years. Unemployment is dominated by medium-term fluctuations between eight and 32 years, and to a lesser extent by low-frequency movements while close to three-quarter of inflation and short-term interest rate fluctuations fall into the slow-moving trend component. The corollary to these results is that business cycles play only a secondary role in explaining overall aggregate fluctuations as real GDP growth is very much a high-frequency phenomenon, while the behavior of inflation is all trend.

Since these variables are central to thinking about monetary policy, both in terms of target variables as well as their information content for the state of the economy, we next assess the effects of monetary policy shocks on the individual frequency components. Using identified structural VARs with sign restrictions we find that across all frequency bands the results from an aggregate VAR carry over to individual components and short-term, business-cycle, medium-term and long-term components. In a baseline specification that includes only the overall data series, a contractionary policy shock, that is, an increase in the federal funds rate, lowers inflation, raises the unemployment rate, and decreases real GDP growth. We find similar patterns across most frequency bands, but as we increase the cycle length, the peak response moves further out, while precision of the impulse response estimates worsens and the quantitative importance declines. We take this as somewhat tentative evidence that monetary policy has an impact across all frequency bands and that a mechanism in line with interaction of endogenous growth and cycles as in Comin and Gertler (2006) is at play. In addition, we find that in the long run the relationship between the nominal interest rate and the inflation rate is positive, whereas in the short run an interest-rate increase lowers inflation. This relationship weakens or is non-existent over the

medium term, which arguably reflects a contrast between a demand effect in the short run and the Fisher effect in the long term.

Our third finding shows that standard DSGE models are in principle capable of replicating the behavior of macroeconomic variables in different frequency bands. We simulate artificial time series from three canonical DSGE models (Smets and Wouters, 2007; del Negro et al., 2015; and Christiano et al., 2016) and apply our wavelet decomposition to the same set of variables as before. Generally, all three models perform reasonably well for business cycle frequencies and for long-term fluctuations. In a sense, this is perhaps not surprising in that the models are built as business-cycle models around the idea that such fluctuations are the outcome of stochastic shocks and endogenous propagation. These DSGE models also include elements such as habit formation, investment adjustment costs, and wage and price indexation to impart persistence on the variables which helps match behavior at business-cycle frequencies.² Long-run behavior is captured by stochastic trends and time-varying inflation targets, which have been introduced successively over the course of model development to capture trends. We show, however, that these models largely fail in capturing behavior at medium-term frequencies, which is particularly prevalent in the case of unemployment and a monetary DSGE model with search and matching frictions in the labor market. We interpret these findings as a challenge for modelers to develop frameworks capable of capturing medium-term cycles.

This paper touches upon various literatures in macroeconomics and time series analysis. There has been a long-standing debate as to whether a frequency-based view of economic fluctuations is useful for analyzing and understanding policy. Perhaps emblematic of a critical viewpoint is Watson (1993) who argues that policy analysis at different frequencies is not relevant for policymakers and that the close relationship between a time series representation of a variable and its counterpart in the frequency domain, such as the spectrogram, invalidates the need for a separate analysis of frequency-specific considerations. This viewpoint is implicitly questioned by Onatski and Williams (2003) who study the effects of uncertainty, broadly understood, on monetary policy decisions. They show that when uncertainty enters a policymaker's decision problem at different frequencies it may have substantially different effects on outcomes. This criticism of the Watson-critique is taken up by Brock et al. (2007) who analyze the differential effects of various policy rules on outcomes across frequencies. In a follow-up paper Brock et al. (2013) demonstrate how

²Tkachenko and Qu (2012) and Sala (2015) estimate medium-size DSGE models in the frequency domain with a focus on business-cycle frequencies. They report similar findings as to the ability of such models to replicate observed behavior over the cycle.

reductions of variance at some frequencies lead to increases in variance at others, which then creates a policy trade-off. Our paper informs this debate in showing empirically the contributions of different frequency bands to the overall volatility of key macroeconomic variables and how they are impacted by monetary policy shocks.

Our paper also continues and contributes to the debate about the use of detrending methods in macroeconomics. Many empirical methods require the underlying data series to be stationary and thereby necessitate the use of a filter to remove trending components. However, as Canova (1998) has demonstrated different detrending methods extract different information from the underlying data series. This implies that the thus derived stylized facts can differ substantially qualitatively and quantitatively across different filtering methods.³ This insight is extended by Gorodnichenko and Ng (2010) and to the estimation of DSGE models. When researchers apply standard data transformations this induces biases in structural estimates and distortions in the policy conclusions. In order to address this issue Canova (2014) proposes joint modelling of the cycle and the trend within the model and the raw data.

We add to this literature by establishing a set of stylized facts based on the time-frequency decomposition inherent in wavelet analysis that has certain advantages over more traditional methods. Thereby, we also highlight the importance of joint theoretical modelling of economic behavior across all frequency bands and especially the medium term as an important component of economic fluctuations. While the importance of the medium run has been on economists' minds for a long time (e.g., Blanchard, 1997), there has been a flurry of recent research recent in the wake of Comin and Gertler's (2006) contribution that study the origin and effects of medium-term cycles (e.g., Beaudry et al., 2017; Cao and Huillier, 2018).

In this paper, we exploit the benefits of wavelet analysis as a complementary approach to classical time series and spectral analysis. We first use the univariate wavelet transform for exploratory data analysis of US macroeconomic variables. In addition, we use the wavelet power spectrum to analyze the evolution over time of the variance of the variable at different frequencies. We then use this approach to isolate specific frequency components from each variable and use those frequency components in a standard VAR regression setup. Our paper thus contributes to a growing literature on the use of alternative filtering methods in economics and finance, such as Aguiar-Conraria et al. (2012) and Bandi et al. (2019).

The remainder of the paper is structured as follows. In the next section we present

³This observation is also in line with the recent criticism in Hamilton (2018) on the use and application of the HP-filter in macroeconomic modelling.

our first set of results, namely new stylized facts based on a wavelet decomposition of aggregate data. In Section 3, we use the decomposition to assess the effects and importance of monetary policy shocks across different frequency bands in a structural VAR framework. Section 4 considers the question whether existing DSGE models are able to capture these regularities. Section 5 concludes.

2 A Frequency-Band Decomposition of Aggregate Time Series

We use the wavelet methodology to decompose standard US macroeconomic time series into different components that can be associated with the scale of the underlying cycles. We regard this time-frequency decomposition, that is, a decomposition of a variable into components in the time domain with precise counterparts in the frequency domain, as a useful and informative alternative to typical trend-cycle decompositions that provides a more encompassing view of the nature of economic fluctuations. In what follows, we briefly discuss the methodology and detail the data used in our empirical exercise. We then present our baseline results, followed by an extensive robustness analysis with respect to alternative filtering methods and choices.

2.1 Methodology and Data

The analysis in this paper is based on a time-frequency decomposition of key economic time series. Our basic objective is to decompose a time series into individual components that can be cleanly and clearly associated with fluctuations at different frequencies or different lengths of a cycle, but are represented in the time domain. For this purpose, we employ wavelet multiresolution analysis (MRA) which performs such decomposition in a way similar to the traditional time series trend-cycle decomposition approach (e.g., Beveridge and Nelson, 1981; Watson, 1986), or other filtering methods like the Hodrick and Prescott (1997) or the Baxter and King (1999) band-pass filter. However, a wavelet approach aims at a more fine-grained understanding of the different components of a time series that make up what is considered a ‘cycle’ as opposed to a ‘trend’.⁴ Specifically, we employ a particular version of a wavelet transformation of a time series called the Maximal Overlap Discrete Wavelet

⁴Conceptually, our line of reasoning is informed by the notion of medium-term cycles as advocated by Comin and Gertler (2006). There is a growing understanding that the neat trend-cycle view of economic fluctuations is inadequate to capture the nature of economic activity.

Transform (MODWT).⁵

As an example, by using the specific form of a Haar wavelet filter any time series X_t can be decomposed into a scale component $S_{J,t}$ and J detail components $D_{j,t}$:

$$X_t = \sum_{j=1}^J D_{j,t} + S_{J,t}, \quad (1)$$

where these coefficients are given by:

$$D_{j,t} = \frac{1}{2^j} \left(\sum_{i=0}^{2^{j-1}-1} X_{t-i} - \sum_{i=2^{j-1}}^{2^j-1} X_{t-i} \right), \quad (2)$$

$$S_{J,t} = \frac{1}{2^J} \sum_{i=0}^{2^J-1} X_{t-i}. \quad (3)$$

Intuitively, the wavelet filter separates the original series X_t , which is defined in the time domain, into different time series components. These represent the fluctuations of X_t in a specific frequency band, that is, a range of frequencies, or length of cycles, that are grouped together.⁶ In this example, the smooth scale component $S_{J,t}$ at time t is computed as the weighted average of lagged values of X_t at scale J , while the detail components $D_{j,t}$ are overlapping weighted moving averages up to scale J . The bands are associated with different details j such that for small j , the wavelet component $D_{j,t}$ captures the higher-frequency characteristics of the time series, that is, its short-term fluctuations. As j increases, the components represent lower frequency movements of the series. Finally, the smooth component $S_{J,t}$ captures the lowest frequency dynamics, that is, the long-term behavior.⁷

The key parameter for the economic interpretation of the wavelet decomposition is the scale J which determines how fine-grained or detailed the decomposition is. For J large enough, the scale component $S_{J,t}$ approximates the true underlying trend of the series. If J is small, then the scale component includes fluctuations of shorter duration, which one may not normally associate with a trend.⁸ An alternative interpretation is that $S_{J,t}$ is the

⁵The MODWT version of the wavelet filter has become the standard in the empirical finance and forecasting literature, e.g. Berger (2016) or Faria and Verona (2018).

⁶The individual components, or *wave-lets*, thus make up the overall *wave* in a prescribed manner.

⁷As in the Beveridge and Nelson (1981) time-series decomposition into stochastic trends and transitory components, the wavelet coefficients $D_{j,t}$ can be viewed as components with different levels of calendar-time persistence operating at different frequencies, whereas the scaling component $S_{J,t}$ can be seen as the low-frequency trend of the time series under analysis.

⁸The Appendix contains a simple example how the scale parameter J is related to the idea of taking various differences of time series.

underlying scale of the time series upon which fluctuations of higher frequencies and shorter cycle durations are built. In that sense, our analysis falls in line with a more standard trend-cycle decomposition. On a final note, the filter discussed above is one-sided since future values are not needed to compute the wavelet coefficients of the transform of X_t at time t . This implies that the $D_{j,t}$ and $S_{J,t}$ lag X_t . Moreover, since the length of the filters increases with j , so does the delay. Hence, the coarser the scale, the more the $D_{j,t}$ and $S_{J,t}$ are lagging X_t . We use this fact in our VAR analysis below.

What distinguishes the wavelet decomposition is that the choice of the scale allows the researcher to hone in on and isolate specific frequency bands that are the objects of interest. While other filtering methods, such as Fourier analysis, also allow a researcher to focus on specific frequencies, a wavelet approach has some key advantages. Traditional decomposition techniques, such as spectral analysis of a time series, tend to impose strong assumptions about the data-generating process. Specifically, they often require data to be stationary or pre-filtered. However, many economic and financial time series are hardly stationary as they exhibit trends and patterns such as structural breaks, volatility clustering and long memory which the wavelet approach can handle with ease.

Unlike Fourier analysis, wavelets are defined over a finite window in the time domain, which is automatically and optimally resized according to the frequency of interest and the choice of the scale J . Wavelets and standard Fourier analysis are essentially approximations with basis functions, but Fourier basis functions are non-zero almost everywhere, making it harder for them to capture local phenomena. Using a short time window isolates the high-frequency features of a time series, while treating the same signal with a large time window reveals its low-frequency features. By varying the size of the time window, we can therefore capture time-varying and frequency-varying features of the time series at the same time. Wavelets are, thus, very useful when dealing with non-stationary time series, irrespective of whether the non-stationarity comes from the level of the time series (that is, from a long-term trend or jumps) or from higher-order moments (that is, from changes in volatility).

Wavelet filtering methods are similar to filtering by a set of band-pass filters so as to capture the fluctuations of a time series in different frequency bands, e.g. Christiano and Fitzgerald (2003). The band-pass filter is a combination of a Fourier decomposition in the frequency domain with a moving average in the time domain. It applies optimal Fourier filtering to a sliding window in the time domain with constant length regardless of the frequency being isolated. Wavelet filtering, in contrast, provides better resolution in the

time domain as the wavelet basis functions are both time-localized and frequency-localized.

In this paper, we use the maximal overlap discrete wavelet transform (MODWT) to compute the decomposition. This version is not restricted to a particular sample size: if the data are discrete the standard wavelet decomposition requires a sample of length 2^J for the decomposition to be exact; that is, it imposes a tight restriction on which and how many frequency bands can be considered and might require dropping observations. The MODWT avoids this problem and is also translation-invariant, that is, it is not sensitive to the choice of a starting point for the examined time series. Finally, implementation of the wavelet decomposition requires choice of a specific functional form for the filter that maps the original series into its components. We follow the literature and choose as a benchmark the Haar filter, but also consider the Daubechies filter as an alternative. Specifically, we employ the filter to decompose our time series of interest into seven individual series, labeled D_1, \dots, D_6 for the detail components and S_6 for the scale component, that is, we choose $J = 6$. The individual components are such that they add up to the underlying series. Given the scale of the decomposition as powers of two we can associate the components with individual frequency bands. Specifically, D_1 captures fluctuations up to four quarters, D_2 between four and eight quarters, up to D_6 which covers the band between 64 and 128 quarters. The scale component S_6 is associated with movements above 128 quarters.

We collect quarterly data on US macroeconomic aggregates, interest rates, and prices. Specifically, we report results for real GDP, the unemployment rate, the inflation rate for the overall personal consumption price index (PCE), the federal funds rate (FFR) and a 3-month and 10-year interest rate.⁹ The data are described in more detail in the Appendix. The full range of our sample covers 1954Q3 to 2017Q3. We utilize data in levels and in growth rates, where growth rates are computed as quarter-over-quarter values. Although not required for the wavelet filtering, we report results for GDP growth as it is the focus of policymakers' decisions. For our baseline decomposition we use a one-sided Haar filter, which are then employed in the VAR analysis. In a sense, the different scale components are generated regressors where we do not want to impart information onto the econometrician running the VAR than he could not possibly possess; that is, knowledge of the data at the end of sample should not be used to produce a decomposition for periods in the middle. For informative purposes and as a robustness check we also provide results for two-sided

⁹The 3-month Treasury rate at constant maturity is only available from 1981Q4 on. We use the 3-month Treasury rate from secondary market instead since it is available from 1947Q1. Preliminary analysis for the two series shows that they co-move extremely closely and that there is at most a level difference of up to 50 basis points.

(smoothed) wavelet filters, for alternative kernels, and for alternative filters, such as the Hodrick-Prescott and the Christiano-Fitzgerald bandpass filters.

2.2 Baseline Results

We report two sets of results. For purposes of exposition, we group the seven series into four categories which we label ‘Short Term’ (D_1, D_2), ‘Business Cycle’ (D_3, D_4), ‘Medium Term’ (D_5, D_6), and ‘Long Term’ (S_6). The short-term category captures high-frequency fluctuations under two years, which in macroeconomic applications are often discarded as noise, but may contain useful information about the incidence of shocks. The business-cycle category covers fluctuations at frequencies between 8 and 32 quarters (2-8 years), which most macroeconomic research on the sources of aggregate movements focuses on. This frequency band is, for instance, designed to be isolated by the application of the Hodrick-Prescott filter with a smoothing parameter of $\lambda = 1,600$.

We maintain this terminology for clarity, although one aspect of our paper is to argue for less rigid classifications in the standard trend-cycle methodology. Components D_5 and D_6 are grouped under ‘Medium Term’ fluctuations and cover frequencies up to 128 quarters (32 years). We note that this scale is shorter than the medium-term cycle adopted in Comin and Gertler (2006), defined as movements between 8 and 50 years. Finally, we associate S_6 with the ‘Long Term’ or, loosely speaking, the trend. We report the grouped wavelet decompositions for real GDP growth, the unemployment rate, the inflation rate, the federal funds rate, the 3-month and 10-year rate, and the difference between the latter two series, namely the term spread, in Figures 1-6. The decompositions into the individual wavelets are collected in the Appendix. Table 1 reports the variance decompositions by frequency.

We find that more than 50% of overall fluctuations in real GDP growth are explained by the short-term components D_1 and D_2 , roughly one third by the business cycle components D_3 and D_4 , with the rest by medium to long-term components.¹⁰ This raises the question whether and to what extent macroeconomic stabilization policy can affect this short-term component, especially since it is likely to contain measurement error. At the same time, the low-frequency component S_6 declines from above 4% to below 2% over the course of the sample (see Figure 1). This is in line with the secular decline in trend growth that has been found in numerous studies. However, this is not the full picture behind the recent lower growth rates, as the two medium-term components D_5 and D_6 essentially offset each other since 2000 and thereby do not contribute to the underlying growth trend. This comes

¹⁰The medium-frequency components as defined by Comin and Gertler (2006) thus make up only 12.5% of the overall fluctuations, with half falling on the band between 32-64 quarters.

largely from the business cycle components during the recovery from the Great Recession. The Great Moderation is most visible in the short-term components and to a lesser extent in the business-cycle band.¹¹ The wavelet decomposition shows that it is more of a higher frequency phenomenon. This observation lends support to the argument that the Great Moderation came about because of an improvement in the way monetary stabilization policy was conducted rather than a change in, for instance, inventory management.

The unemployment rate decomposition in Figure 2 and Table 1 reveals a slightly different pattern. Roughly one third of unemployment fluctuations are due to short-term and business-cycle movements, while medium- and longer-term frequencies (D_5 - S_6) each explain around 20%. Fluctuations in the unemployment rate can therefore be described as a medium-term cycle. What dominates the *level* of the unemployment rate is its long-term component S_6 , which could be interpreted loosely as a natural rate of unemployment. A focus of the next section is the extent to which the trend components are affected by monetary policy. What is striking is that the different components do not seem to comove closely. For instance, the unemployment rate is at 5.4% in 1990, while the long-term component S_6 is at 7.2%, the difference being made up by components D_4 - D_6 . In other words, the business cycle peak produces a negative unemployment gap relative to a very high natural rate on account of strong medium-term components which might be tied to labor force participation peaking in the late 1990s. Finally, the Great Moderation is considerably less visible in the unemployment rate, if at all.

We now turn to the nominal side of the economy. Figure 3 contains the results from the decomposition of the PCE inflation rate. 40% of inflation movements can be traced back to the long-term component S_6 . The business cycle component explains around one fifth of the overall variability, while medium-term components cover 25%. About 15% of inflation variability can be traced back to very short-term or noise components. As in the case of the unemployment rate, the scale of the decomposition is dominated by the trend S_6 . The monetary policy literature often interprets this component as the inflation target or the perception thereof. It can also be seen as a measure of the extent to which inflation expectations are anchored. In our decomposition, it shows a gradual rise from almost zero in the late 1960s to a peak of 6.2% in the early 1980s followed by a gradual decline until reaching the 2% target in the 2000s.

A similar pattern in terms of the Volcker disinflation can be found in the medium-term

¹¹Figure C.1 in the Appendix shows that this is largely due to the D_3 component, indicating that the Great Moderation is essentially a high-frequency event. To this point, see also Aguiar-Conraria et al. (2012) and Pancrazi (2015).

components D_5 and D_6 . What is striking is the run-up in trend inflation over the course of the 1970s and the drawn-out, three-decade long struggle to return it to 2%. Since the Federal Reserve arguably did not change its implicit inflation target over that time, this component may therefore be better described as the public’s perceived target. Our results then depict a striking loss of central bank credibility.¹² In light of this aspect, it is perhaps surprising that there is not much of a Great Moderation visible when interpreted as a binary event, that is, a break in policy or a structural change before or after the early 1980s. Instead, in the graphs in Figure 3 it is possible to discern the high volatility of the 1970s, preceded and followed by the more stable 1960s and 1980s, respectively. Interestingly, inflation volatility seems to have gone up again in the 2000s, especially around the Great Recession.

We report decompositions for the FFR and the 10-year rate in Figures 4 and 5.¹³ They show similar patterns as the inflation decompositions, whereby volatility in the 10-year rate can be attributed to almost 70% to the long-term component S_6 , ten percentage points more than for the short rates. Presumably, this reflects that longer rates are less subject to the vagaries of higher-frequency fluctuations. Since the interest rates share common components, especially in the medium and longer run, it is therefore often instructive to consider the term spread, in our case the difference between the 10-year and the 3-month rate. The term spread decomposition in Figure 6 puts most weight, almost 45%, on the business-cycle components. This supports the idea that at frequencies commonly associated with the business cycle the spread is a useful indicator of economic and financial conditions. Interestingly, the long-term component has gone up considerably since the early 1980s to a level of above 2%, implying that the difference between the short and long rates has become more persistent.

As a final exercise, we produce the power spectra from the wavelet decomposition for real GDP growth, unemployment and the federal funds rate in Figures 7-9.¹⁴ The wavelet-based spectra are akin to classical spectra in that they decompose a time series into frequencies and measure the contribution of each frequency to the overall behavior of a time series.¹⁵ Ordered by frequency, this spectral density is depicted in the right column of each figure. In addition, we show the time-frequency decomposition in the left column. Using the wavelet

¹²This interpretation is consistent both with the learning and inherent inflation persistence story in Primiceri (2006) or Sargent, Williams, and Zha (2006), the inflation misperception argument in Lubik and Matthes (2016), as well as a number of recent papers on evolving private sector beliefs, for instance, in Bianchi (2013).

¹³The results for the 3-month rate are almost identical to those for the FFR. The respective decompositions are shown in the Appendix.

¹⁴We briefly describe and discuss the concept behind the wavelet power spectra in the Appendix.

¹⁵As discussed before, the wavelet power spectrum does not require stationarity.

filter, the graph reports the wavelet power spectrum (WPS), a decomposition by time (on the horizontal axis) and frequency (on the vertical axis). The graph is coded as a heat map such that areas of higher activity are depicted as redder on the color spectrum. The solid black lines delineate 95% confidence regions. Note that the wavelet-based spectral density can be obtained by integrating the WPS over time.

The spectral peak of GDP growth in Figure 7 is just above two years which confirms our prior finding that output fluctuations are highly concentrated among the highest frequencies. The WPS, however, shows that this observation is largely driven by the late 1950s and 1970s which show considerably higher concentration of activity in the short term than what occurred during later periods. In the same vein, the Great Moderation is quite visible from the graph.¹⁶ The unemployment rate in Figure 8 shows two local spectral peaks associated with cycles of around 8 and 32 years. This is in line with our previous findings, but sharpens the observation of unemployment being subject to medium-term cycles towards the edges of that frequency band. Notably, the Great Moderation is not apparent from the WPS, while deep recessions in the mid-1970s and the Great Recession impart some higher frequency components on the decomposition. Finally, Figure 9 reveals that the FFR has a spectral peak at a very low frequency which we associate with the presence of an explicit or implicit inflation target. However, the WPS shows a local peak at a frequency of 8 years which is largely driven by the period from 1968 until the early 2000s.

Overall, what emerges from these decompositions is a multifaceted picture of macroeconomic fluctuations. Across all variables, the business-cycle components $D_3 - D_4$, that is, cycles between two and eight years, explain about one third of overall fluctuations. There is considerable heterogeneity across variables as far as the other components are concerned: 50% of real GDP growth is captured by high frequency components (cycles of less than 2 years). Essentially, much of quarterly GDP movements occurs at very high frequencies.¹⁷ In turn, short-run fluctuations do not seem to play much of a role for the other variables. The behavior of unemployment is dominated by medium-term movements with a cycle length of between 8 and 32 years and to a lesser extent by longer-term movements of lower frequency than that. Inflation and interest rates have sizeable long-term components, too. These components can be interpreted as “trends” and natural or potential rates. Their behavior

¹⁶This is in line with related research by Pancrazi (2015) who argues that the reduction of volatility of GDP after the mid-1980s is mainly a high-frequency phenomenon of cycles up to 4 years and that it is much milder, or absent, for other frequencies.

¹⁷We use final data in our empirical study, that is, the last data vintage available. In contrast, policymakers operate in a real-time environment where initial data releases are subject to sometimes large measurement error. Lubik and Matthes (2016) show that this can lead to what looks like policy mistakes *ex post*.

arguably conforms to conventional wisdom, that is, inflation seems to be all trend, driven by the Federal Reserve’s implicit and then later explicit inflation target.

This naturally raises the question whether stabilization policy aimed at the business cycle is misdirected or misses important aspects that policymakers should focus on.¹⁸ An immediate follow-up question is whether models that are being used to describe and analyze monetary policy are consistent with the heterogeneity in fluctuations. We address these two questions in turn in the following two sections. First, we investigate whether identified monetary policy shocks have differential effects on key variables for different frequencies; and second, we study whether some standard DSGE models are capable of replicating the wavelet-based variance decompositions in this section.

2.3 A Comparison of Alternative Filters

We assess the robustness of our baseline findings for the one-sided Haar filter along several dimensions. First, we consider a two-sided version of the Haar filter. The second exercise considers an alternative kernel for the wavelet decomposition, namely the Daubechies filter. The third robustness check uses filters that are more common in the macroeconomics literature, specifically the Christiano-Fitzgerald bandpass filter and the Hodrick-Prescott filter. As before we focus on four broad frequency bands for exposition. The decompositions are reported in Figures 10-12.

Figures 10 and 11 contain the decompositions of, respectively, real GDP growth and unemployment for the one- and two-sided Haar filter and the Daubechies filter. By definition the two-sided filter is smoother than a one-sided filter since it uses all available information over the whole span of the sample and not just up to the data point at which the filter is applied. This is evident by comparing the one-sided Haar filter with its two-side counterpart in the figures. Generally, there are no large differences in terms of the overall direction and volatility for both unemployment and real GDP growth, but the one-sided filter imparts more volatility to the short-term and business-cycle components than the other filters. Moreover, the one-sided Haar filter lags the other filters in picking up general directional movements. This is especially visible in the medium- and long-term components of unemployment in Figure 11.¹⁹ The fact that the one-sided Haar is slow in picking up the rise and subsequent fall in trend unemployment in the 1970s and 1980s is simply a feature of how it is constructed. As discussed before, we prefer a one-sided filter since we use the

¹⁸This argument has been made most succinctly by Brock et al. (2008, 2013).

¹⁹We find similar patterns in the decompositions for inflation and the interest rates. These results are included in the Appendix.

individual components as variables in a VAR which rests on the idea that the innovations are one-step ahead forecast errors and thereby do not reflect the full information in the sample.

The figures also report results for the Debauchies filter as an alternative to the two-sided Haar filter. The Haar filter produces less volatile components than the Daubechies, but the difference seems minor. There are a few episodes where the two filters do not overlap each other. For instance, the medium-term components of inflation in the mid-1970s differ noticeably, but these occurrences are the exception. We prefer the Haar over the Debauchies implementation of the wavelet decomposition since the former has a more intuitive interpretation (see the discussion in the Appendix). The differences between the various implementations of the decomposition are small enough, however, not to affect the conclusions drawn in the next two sections.²⁰

In contrast, the decompositions based on two widely used filters in macroeconomic analyses are materially different. Figure 12 compares our baseline filter with the corresponding bandpass filter of Christiano and Fitzgerald (2003) (CF) and the canonical Hodrick-Prescott (HP) filter. In a sense, the CF filter and our Haar filter are conceptually similar in that they explicitly isolate specific frequency bands and represent them in the time domain. This is evident from comparing the two filters in the figure for unemployment rate decompositions as an illustrative example. At business-cycle frequencies the CF filter extracts more volatile components, but is arguably not that different from the wavelet-based filter. The exception are the longer-term components, especially D_6 and S_6 , where the two filters pick out different peaks and are generally not that well aligned.²¹

In contrast, the HP filter produces quite different series. For a smoothing parameter of $\lambda = 1,600$ it extracts the business-cycle frequencies corresponding to our components D_3 and D_4 . It is considerably more volatile than the wavelet decomposition. More striking is the pattern for lower frequencies. The figure reports the HP trend which is computed as the difference between the original series and the business cycle component obtained with $\lambda = 1,600$. It is akin to the S_6 component with wavelets, which is the “residual” part of the series. This slow-moving component is quite different from the other series and thus raises concerns as to whether the HP filter introduces spurious dynamics (see Hamilton, 2017).

²⁰We performed the empirical exercises in Sections 3 and 4 using alternative wavelet decompositions. The results are available on request.

²¹However, recall that the CF filter is optimized to extract business-cycle frequencies and not low frequencies. At the same time, this is another argument in favor of wavelet filters since it treats all frequencies the same way.

3 The Frequency-Specific Effects of Monetary Policy Shocks

We now study whether and to what extent monetary policy shocks affect key aggregate variables across different frequencies. The previous section demonstrated that the behavior of GDP growth, unemployment, and the inflation rate differs in terms of the contribution of various frequency bands to overall volatility. Whereas the majority of fluctuations in GDP growth are located among the highest frequencies, that is, the short-term components, the unemployment rate is more evenly split between a large medium-term and lower-frequency components. In turn, most of the movements in the inflation rate are driven by the long term which we might associate with the inflation target. As we think of monetary policy as trying to stabilize movements in GDP growth and unemployment against a background of stable prices or constant inflation, the question is whether policy is successful in affecting these variable at frequencies that are the main drivers of their overall volatility.

Our approach is as follows. We assess the effects of monetary policy shocks on individual frequency bands by using the filtered series as explanatory variables in a VAR. Given a plausible identification of policy shocks, we then compute impulse response functions to these shocks for the various decompositions. We begin by assessing the plausibility of our preferred identification scheme in a standard model. To this end, we estimate a three-variable VAR in an activity variable, that is, either the unemployment rate or real GDP growth, inflation, and the federal funds rate. We then identify a structural monetary policy shock using a sign restriction approach where we assume that all restrictions are imposed only on impact. Specifically, we assume that a contractionary monetary policy shock - one that raises the federal funds rate on impact - lowers output, increases unemployment and lowers inflation.

Figure 13 reports impulse responses to an identified policy shock from the two VARs. The left column shows the responses in the model with unemployment, the right those of the model with GDP growth. In this baseline specification, a rise in the interest rate by 25 basis points increases unemployment by 10bp with a hump-shaped peak after 3-4 quarters of 20bp. It lowers inflation by 60bp on impact before gradually returning to its long-run level. Similarly, a contractionary monetary policy shock lowers GDP growth by almost 1.5 percentage points and inflation shows a similar decline as in the other specification.²²

In the next step, we add the frequency components to the baseline specification, either in

²²Both the unemployment rate and GDP overshoot their long run level in their adjustment path after the shock in line with the interest rate responses. That is, monetary policy responds endogenously to the worsening economic conditions due to the unanticipated contraction by loosening policy.

terms of unemployment or real GDP growth, as an activity variable. For each specification we identify the policy shock separately, which allows for the possibility that there could be differences across models. We consider two alternative specifications. First, we add the seven frequency bands, D_1 - S_6 , of each variable included in the VAR one by one to the baseline specification. This results in a six-variable VAR, estimated separately for each band. We report selected impulse responses for GDP growth in Figures 14-16, where the left column shows the responses of the aggregate variables and the left column the corresponding responses for a frequency band. The respective responses with unemployment as the activity variable can be found in the Appendix.

We find that the responses of the high-frequency components D_2 are significant, and are in line with the baseline results and what theory would suggest; however, the responses are not large quantitatively and economically small. Nevertheless, this indicates that the monetary transmission mechanism works as theoretical reasoning and practical experience would indicate. The response of the business-cycle component D_4 is not significant on impact but becomes more sharply estimated a few quarters out. As before, the direction of the responses is consistent with the identification scheme on the overall series. Contractionary shocks are thus likely to have their strongest impact in a few quarters which is in line with the idea that monetary policy stabilizes business cycles with a lag.

Finally, the response for the long-term component S_6 is drawn out and not significant over the business cycle horizon, but exhibits comovement between the federal funds rate response and inflation. In other words, at longer horizons and cycles, the Fisher effect, namely that interest rates and inflation rates are positively correlated, comes through; whereas at higher frequencies this correlation moves in the opposite direction as the demand-constricting effect of higher rates reduces inflation. It is clear from these findings that in the transition between high frequency and low frequency movements the comovement patterns for these two variables switch.

The second VAR specification adds the filtered series in groups that represent broader frequency bands. Since the wavelet decomposition is fully additive we cannot include all individual series. We therefore focus on a specification that looks at the business-cycle components ($D_3 + D_4$), the medium-term cycles ($D_5 + D_6$), and the long term (S_6). This results in a twelve-variable VAR, where we identify the policy shock by imposing sign restrictions on impact on the aggregate variables only. Figure 17 contains the respective responses for the model with GDP growth, while the corresponding responses for unemployment are reported in Figure 18. In each graph, the top row contains the aggregate responses, followed

by the short-term, medium-term, and long-term components in separate rows.

A contractionary monetary policy shock has a negative impact on real activity in each frequency band whereby the largest response is for the business cycle component. If we just look at the short-term frequencies (not reported), the impact effect is larger. This possibly reflects the dominant role of high-frequency movements in GDP growth (see Table 1). The responses of the business-cycle and medium-term components return to zero after 20 and 30 quarters, respectively. The response of the long-term component on the other hand remains negative over the full projection horizon of 10 years. This indicates that monetary policy shocks can have long-lasting effects even on GDP *growth*. We find a similar pattern for the unemployment rate, with oscillating behavior of the higher-frequency components and a more drawn out response of the trend. In terms of the size of the policy-induced movements, the business-cycle and medium-term components are roughly similar, in contrast to the variance decompositions in Table 1. This suggests that there are other shocks that drive movements in the unemployment rate in these frequency bands.

The response of the FFR and inflation components for both VAR specifications is very similar. At higher frequencies, the FFR rises and inflation falls, where especially the business-cycle components move together closely. The response of the respective trend components is different, however. Inflation and the FFR do not react much on impact and in the near term, but move together positively over the longer horizon. A contractionary policy shock thus has a long-lasting negative effect on the long-term component of the FFR and inflation. These results confirm the existence of a Fisher effect in the long-term component, whereas in the short term the demand-constricting effect of an interest rate hike dominates as in standard monetary policy models. Moreover, the results also show that contractionary policy lowers the long-term component persistently presumably through an expectations effect: tightening policy gains credibility, anchors inflation expectations, and lowers inflation overall.

4 Assessing DSGE Models

We now investigate whether several medium-scale DSGE models can replicate the stylized facts identified above. Such models have been developed explicitly with an eye on replicating the performance for business cycle movements and the long run. This raises the question whether they can, in fact, capture behavior along all frequency bands identified by our wavelet decomposition. In a preview of the results, we find that the models generally do well for business-cycle frequencies and in the long term as these are frequency bands

which the models are designed to replicate. However, the models generally fail at capturing medium-term frequencies.

4.1 DSGE Models and Simulation

In the DSGE literature it is well known that various modeling devices are useful in matching persistence in the data, at least over the business cycle (see the programmatic papers by Christiano et al., 2005, and Christiano et al., 2010, and also the seminal DSGE models by Smets and Wouters, 2003, 2007). Examples are modifications to utility, such as habits in consumption, production, such as investment adjustment costs, and highly persistent shock processes. At the same time, stochastic trends have proved to be a flexible modeling component to capture drifting behavior over time. This section studies whether these modeling elements are useful across all frequencies.

We select models based on their widespread use in monetary policy analysis and their consistency with the specific data that we have considered so far. Moreover, we want to give the chosen models a fair chance at capturing the patterns found in the wavelet decomposition. We therefore require that one of the underlying drivers of business cycles is a stochastic trend in productivity which can smoothly vary over time. This specification is well known to match the movements in the GDP trend. We thus focus on three canonical models in the literature: Smets and Wouters (2007), del Negro et al. (2015), and Christiano et al. (2016).²³

Smets and Wouters (2007) (SW) is a further development of the canonical Smets and Wouters (2003) New Keynesian DSGE model. It is the prototype of a medium-sized, optimization-based model designed to jointly capture the evolution of output and inflation and the monetary policy process. To this end, the model contains a variety of shocks and frictions that have come to be accepted as central to understanding aggregate fluctuations. The basic setup involves a representative household that makes consumption choices and supplies labor to a competitive labor market. On the production side there are monopolistically competitive firms that employ labor and capital to generate output, make investment decisions and set prices. The third type of agent in this model is a policymaker who sets interest rates based on given feedback rules.

The model features nominal price stickiness and sticky wages with backward inflation indexation to capture slow-moving aspects of these variables. On the real side, there is habit formation in consumption and investment adjustment costs designed to create hump-

²³We use computer codes for these models available at Volker Wieland's Macroeconomic Model Data Base (MMB): <https://www.macromodelbase.com/> and from the journal websites of the published articles.

shaped responses of these aggregate demand components. The model is driven by seven structural shocks including a monetary policy disturbance. One key distinguishing feature of Smets and Wouters (2007) as opposed to Smets and Wouters (2003) is that the former does not have a time-varying inflation target. The model is estimated using Bayesian methods over the period 1966-2004 for seven key aggregate variables, but the set of observables does not include the unemployment rate. We can therefore not compare their model with our decomposition along this margin.²⁴ We take their parameters estimates as given and simulate the model under this specification.

The second model that we consider, del Negro et al. (2015) (dNGS), is an extension of the SW model. It introduces a time-varying target inflation rate and incorporates financial frictions in the vein of Christiano et al. (2014). The model is estimated for a slightly larger dataset than the SW model and over the period 1964-2008. The key finding of the paper is that the model is compatible with Great Recession outcomes in that it successfully predicts a sharp contraction in economic activity along with a drawn-out but modest decline in inflation. The third model is Christiano et al. (2016) (CET). While it is built around the same nominal structure as SW, CET introduce a much richer labor market setting governed by search and matching frictions and various wage determination mechanism. We report results both for a benchmark specification with Nash bargaining and an alternative, namely alternative offer bargaining. What is important for our purposes is that the framework models the unemployment rate in contrast to the previous two DSGE models. Christiano et al. (2016) estimate the model over the sample period 1951-2008, with the same end date as del Negro et al. (2015).

Our simulation procedure is as follows. We take the estimated models as given and fix the parameter values at the reported posterior medians. The models are then simulated by drawing from the innovation distributions over 10,000 periods. This is repeated 1,000 times, whereby we record the last observations to coincide with the length of our sample. From this sampling distribution we then compute the variance decomposition from the wavelet filter as in section 2 and report 90%-confidence regions for the mean estimate of the variance decomposition. We group the individual wavelet decompositions into the categories ‘Short Term’ (D_1 - D_2), ‘Business Cycle’ (D_3 - D_4), ‘Medium Term’ (D_5 - D_6), and ‘Long Term’ (S_6).

²⁴They also use the GDP deflator to measure inflation whereas we report results for PCE inflation. The differences between these two inflation measures are minor.

4.2 DSGE Models and Frequency-Band Decompositions

We report the results of the simulation exercise in Table 2. Our focus is on the results for four key variables, namely real GDP growth, inflation, the federal funds rate, and the unemployment rate. We compare the simulation results to two different sets of underlying data: first, our original sample which covers 1954Q3 to 2017Q3; and second, the actual sample period over which the respective model was originally estimated. For all three models this excludes the Great Recession period and its aftermath. The latter results are reported separately in Table 3.

The SW model is remarkably successful in replicating the overall volatility of real GDP components across all frequency groups, essentially matching the data exactly: around 60% is attributed to the short-term component, 30% to the business cycle component and a much smaller percentage to the medium and long term. The same pattern is found for the dNGS model and with some minor differences for the Nash-bargaining specification of the CET model. A key driver for this finding is the specification of the exogenous productivity process as a stochastic trend which is now standard modeling device in DSGE models. The wavelet decomposition thus confirms the importance of this assumption.

Turning to the nominal variables, inflation and interest rates, the performance of the SW model notably deteriorates. While the model is consistent with the short-term and medium-term components in inflation, the contribution of the business cycle and long-term fluctuations is essentially flipped. The SW model attributes only 15% to the long term and more than one third to the business cycle. In contrast, the dNGS model comes much closer to the patterns in the data, although it underpredicts the contribution of the long-term component by almost 10 percentage points. The key difference between the two models is that del Negro et al. (2015) incorporate a time-varying inflation target which is stationary but highly persistent. Over the sample period it effectively pins down the trend movements in the inflation rate. As discussed before, trend inflation in the data might simply reflect the changing implicit or explicit inflation target, which in the DSGE modelling sense can be captured by such an exogenous process.

Interestingly enough, the Nash-bargaining specification of Christiano et al. (2016) has difficulty with this pattern as it attributes considerable variability to the short-term and business-cycle components and not enough to the long-term component. Notably, their model does not feature a time-varying inflation target which reinforces the point raised above. However, the alternative specification of the CET model with wage determination based on an alternative offer bargaining mechanism is on point in capturing the behavior

of inflation across all frequency bands. It is well known, e.g., Krause et al. (2008), that a Nash-bargaining nominal wage mechanism does not impart enough inflation persistence in a New Keynesian search and matching framework which these results confirm. On the other hand, the alternative offer bargaining mechanism implies endogenous wage inertia which then translates into inertial prices (see the discussion in Christiano et al., 2016).

The frequency-specific patterns of the FFR in the data do not differ much from that of the inflation rate, although the wavelet decomposition attributes 80% of its movements to medium-term and long-term components, as opposed to two thirds in the case of inflation. A similar pattern is discernible for the three DSGE models in that they cannot replicate the importance of the long-term component and the relative lack thereof in the business-cycle frequencies. Most strikingly, the trend in the interest rate is associated with almost 60% of movements in the data, only half of which the dNGS model can capture. As before, the alternative specification of the CET model does remarkably well for the behavior of the FFR.

We finally consider the decompositions of the unemployment rate which of our three DSGE models only Christiano et al. (2016) can address. In the data, half of the fluctuations in unemployment are captured by the medium-term component with the remainder roughly equally attributed to the business cycle and the long term. Under the Nash-bargaining specification the CET model attributes half of unemployment fluctuations to the long-term component, one third to the medium term component and the remainder to the business cycle. The model gets the broad pattern of fluctuations at different frequencies right: what matters for explaining the unemployment rate are the medium to long-term components, but not those that are arguably more directly shaped by monetary policy.

At the same time, the alternative bargaining specification of CET results in a considerably worse performance for the unemployment rate. It also has problems with the decompositions of GDP growth where it attributes too much volatility to the long-term and medium-term components. Yet, its performance for the two nominal variables, the inflation rate and the interest rate is spot on, where the baseline specification with Nash bargaining put too much weight on the business cycle components. Comparing the two approaches to modeling wage determination these findings indicate that alternative offer bargaining generates more persistence in the model than Nash bargaining does.²⁵ The flip

²⁵This is, of course, related to the Shimer (2005) puzzle who argues that the standard search and matching model cannot replicate the observed volatility and also persistence of the unemployment rate and vacancies, that is, open positions. Alternative offer bargaining therefore presents an attractive solution to the Shimer puzzle which does not have to rely on exogenous wage stickiness.

side of this finding is that the former imparts too much persistence which hurts the model's performance with respect to the medium and long-term components of GDP growth and unemployment.

Our final exercise considers the importance of the sampling period for the assessment of the models. We subject each of the four model specifications to the same test, namely whether they could replicate the behavior of the wavelet decompositions for the full length of our empirical sample from 1954Q3 to 2017Q3. However, in their published versions the estimation periods of the three model frameworks differ. Specifically, the estimation period for the SW model is 1966-2004, for the dNGS model it is 1964Q1-2008Q3, and for the CET framework the sample period is 1954Q1-2008Q4. The former two periods are similar, they both miss 10 years at the beginning of our sample and then the Great Recession and its aftermath; whereas the CET sample differs from ours in that it ends at the onset of the Great Recession. Although the underlying idea of structural DSGE modeling is that structural parameters are generally invariant over these sample periods it is also well known that sample size and sample period can affect structural parameter estimates (e.g., Canova and Ferroni, 2012).

In Table 3 we therefore contrast select decompositions for the actual estimation sample with the simulated sample for the same number of observations. The decompositions for the SW and dNGS sample periods are very similar to each other for real GDP growth and inflation. The biggest difference is the long-term inflation component in the dGNS sample which includes an additional four years before the onset of the Great Recession. In our full sample, this component explains 41% of inflation movements, in the shorter sample only 34%. For both models, the biggest discrepancy can be found in the behavior of the FFR. The shorter sample attributes much more volatility to the business-cycle and medium-term components for the policy rate, that is, 60% compared to 40% in the full sample. The long-term component explains about one third of the volatility in the SW sample, but almost 60% in the full sample. This discrepancy is arguably due to differences in policy across the two periods, either at the beginning of the full sample period between 1954 and 1964 or during the Great Recession. In any case, the long-term FFR component is more pronounced over the longer period. The results in Table 3 do show, however, that the SW model struggles to match these facts, while the dNGS model is closer to the data.

This pattern is also evident from the CET model, where we find that the long-term FFR component in the full sample explains more of the overall volatility. This suggests that the Great Recession period has had a noticeable effect on the behavior of the FFR -

which may not be surprising since during this period the Federal Reserve held its policy rate essentially fixed at its effective lower bound of zero. This, by itself, imparts persistence onto the FFR. Interestingly, such differences are not visible in any of the other variables, GDP growth and inflation, with the exception of the unemployment rate. It therefore seems that the behavior of the policy rate is largely disentangled from that of other macroeconomic aggregates. Table 3 also shows that the behavior of the unemployment rate is different across the samples and that the CET model under alternative offer bargaining cannot capture the behavior in the different sample either.

In summary, we conclude that the three canonical DSGE models are able to replicate the wavelet decomposition we found in the data. We identify a stochastic trend in productivity and a time-varying inflation target as the key modeling elements. The random walk component in the former and a highly persistent inflation target capture the long-term components in real GDP and inflation exceptionally well. Replicating the frequency-specific components of the unemployment rate proves to be more difficult. While the decompositions from the simulated data go broadly in the same direction, the challenge is that the variance decomposition is more evenly distributed among frequency bands than for the other variables. Based on these findings we advocate wavelet decompositions as a straightforward tool to assess the validity of a DSGE model as a data-generating process, especially with respect to the contribution of individual modelling elements.²⁶

5 Conclusion

This paper advances three main findings. First, we show that more than two thirds of inflation and unemployment fluctuations in the US occur at low frequencies, whereas at most a quarter are attributable to business cycle frequencies. However, it is mainly these latter fluctuations that are the focus of monetary policymakers and researchers: policy objectives are normally phrased in terms of stabilizing fluctuations around trends or potential. This dichotomy is generally reflected in the DSGE models that are used to study monetary policy and its effects. Frequency-specific decompositions such as the one we performed using wavelet methodology thus produce information relevant for policymakers.

Our second finding shows that several standard DSGE models do a credible job of

²⁶In a sense, we are simply confirming the results of Sala (2015) who estimates the SW-model in the frequency domain using likelihood-based methods and working off the counterpart of the time-series representation of a state-space model. He finds that this DSGE model broadly performs well and matches the data at various frequencies, but fails at capturing labor market data and the interactions between real and nominal variables. However, he uses stationary, thus pre-filtered data, and can therefore not speak to the overall decomposition into the several frequency bands.

replicating behavior at business-cycle frequencies that we identified in the data. However, the models need to be suitably modified to account for long-term movements via stochastic trends or time-varying inflation targets. They generally fail in capturing behavior at medium-term cycles of between 8 and 32 years. We demonstrate in a third set of results that monetary policy shocks exert influence over all frequency bands and in a broadly similar manner with the exception of the relationship between short-term interest rates and inflation where the Fisher effect prevails in the long run.

Our paper thus contributes to a growing area of research that suggests that the notion of a cycle relevant for stabilization policy should be extended to include at least the medium term. Specifically, the analysis in the paper indicates that temporary shocks can have long-lasting effects that traditional business cycle modelling largely abstracts from. Future work could therefore study time-frequency decompositions in models with such a transmission mechanism as in, for instance, Comin and Gertler (2006). Similarly, the findings in this paper also support the idea that what matters for monetary policy is less the short-term response of policy rates to deviations of economic activity from some target, but rather the credible anchoring of expectations.²⁷ Typical analyses of optimal monetary policy focus on weighted averages of the unconditional variances of policy targets. It is common to compare policies by considering, for example, a weighted average of the unconditional variances of inflation and unemployment. However, such computations mask the effects of policies on the variance of fluctuations at different frequencies. Frequency-based optimal policy in the vein of Brock et al. (2013) would thus be an interesting extension based on the analysis in this paper.

References

- [1] Aguiar-Conraria, Luis, Manuel M.F. Martins, and Maria Joana Soares (2012): “The Yield Curve and the Macro-Economy Across Time and Frequencies”. *Journal of Economic Dynamics and Control*, 36, pp. 1950-1970.
- [2] Bandi, Federico M., Benoit Perron, Andrea Tamoni, and Claudio Tebaldi (2019): “The Scale of Predictability”. *Journal of Econometrics*, 208(1), pp. 120-140.

²⁷Inflation expectations can be anchored by the execution of tough anti-inflationary policies. In that sense, short-term stabilization policies and commitment to a long-term target are essentially two sides of the same coin since the former helps ensure the latter. However, the joint determination of policy is rarely modelled in DSGE models where the inflation target is often assumed rather than chosen.

- [3] Baxter, Marianne, and Robert G. King (1999): “Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series”. *The Review of Economics and Statistics*, 81(4), pp. 575-593.
- [4] Beaudry, Paul, Dana Galizia, and Franck Portier (2017): “Putting the Cycle Back into Business Cycle Analysis”. Manuscript.
- [5] Berger, Theo (2016): “Forecasting Based on Decomposed Financial Return Series: A Wavelet Analysis”. *Journal of Forecasting*, 35(5), pp. 419-433.
- [6] Beveridge, Stephen, and Charles 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), pp. 151-174.
- [7] Bianchi, Francesco (2013): “Regime Switches, Agents’ Beliefs, and Post-World War II U.S. Macroeconomic Dynamics”. *Review of Economic Studies*, 80(2), pp. 463-490.
- [8] Blanchard, Olivier J. (1997): “The Medium Run”. *Brookings Papers on Economic Activity*, 2, pp. 89-158.
- [9] Brock, William A., Steven N. Durlauf, and Giacomo Rondina (2008): “Frequency-Specific Effects of Stabilization Policies”. *American Economic Review*, 98(2), pp. 241-245.
- [10] Brock, William A., Steven N. Durlauf, and Giacomo Rondina (2013): “Design Limits and Dynamic Policy Analysis”. *Journal of Economic Dynamics and Control*, 37, pp. 2710-2728.
- [11] Brock, William A., James Nason, and Giacomo Rondina (2007): “Simple versus Optimal Rules as Guides to Policy”. *Journal of Monetary Economics*, 54(5), pp. 1372-1396.
- [12] Canova, Fabio (1998): “Detrending and Business Cycle Facts”. *Journal of Monetary Economics*, 41(3), pp. 475-512.
- [13] Canova, Fabio (2014): “Bridging DSGE Models and the Raw Data”. *Journal of Monetary Economics*, 67, pp. 1-15.
- [14] Canova, Fabio, and Filippo Ferroni (2012): “The Dynamics of US Inflation: Can Monetary Policy Explain the Changes?” *Journal of Econometrics*, 167, 47-60.

- [15] Cao, Dan, and Jean-Paul L’Huillier (2018): “Technological Revolutions and the Three Great Slumps: A Medium-run Analysis”. *Journal of Monetary Economics*, 96, pp. 93-108.
- [16] Christiano, Lawrence J., Mathias Trabandt, and Karl Walentin (2010): “DSGE Models for Monetary Policy Analysis”. In: Benjamin M. Friedman and Michael Woodford (eds.): *Handbook of Monetary Economics*, Volume 3a, Chapter 7, pp. 285-367.
- [17] Christiano, Lawrence J., Martin S. Eichenbaum, and Charles L. Evans (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy”. *Journal of Political Economy*, 113(1), pp. 1-45.
- [18] Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt (2016): “Unemployment and Business Cycles”. *Econometrica*, 84(4), pp. 1523-1569.
- [19] Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno (2014): “Risk Shocks”. *American Economic Review*, 104(1), pp. 27-65.
- [20] Christiano, Lawrence J., and Terry J. Fitzgerald (2003): “The Band Pass Filter”. *International Economic Review*, 44(2), pp. 435-465.
- [21] Cogley, Timothy (2001): “Estimating and Testing Rational Expectations Models When the Trend Specification is Uncertain”. *Journal of Economic Dynamics and Control*, 25, pp. 1485-1525.
- [22] Comin, Diego, and Mark Gertler (2006): “Medium-Term Business Cycles”. *American Economic Review*, 96(3), pp. 523-551.
- [23] Del Negro, Marco, Marc P. Giannoni, and Frank Schorfheide (2015): “Inflation in the Great Recession and New Keynesian Models”. *American Economic Journal: Macroeconomics*, 7(1): pp. 168-96.
- [24] Faria, Gonçalo, and Fabio Verona (2018): “Forecasting Stock Market Returns by Summing the Frequency-Decomposed Parts”. *Journal of Empirical Finance*, 45, pp. 228-242.
- [25] Gorodnichenko, Yuriy, and Serena Ng (2010): “Estimation of DSGE Models when the Data are Persistent”. *Journal of Monetary Economics*, 57, pp. 325-340.
- [26] Hamilton, James D. (2017): “Why You Should Never Use the Hodrick-Prescott Filter”. Forthcoming, *The Review of Economics and Statistics*.

- [27] Hodrick, Robert, and Edward C. Prescott (1997): “Postwar U.S. Business Cycles: An Empirical Investigation”. *Journal of Money, Credit, and Banking*, 29(1), pp. 1-16.
- [28] Krause, Michael U., David Lopez-Salido, and Thomas A. Lubik (2008): “Inflation Dynamics with Search Frictions: A Structural Econometric Analysis”. *Journal of Monetary Economics*, 55, pp. 892-916.
- [29] Lubik, Thomas A., and Christian Matthes (2016): “Indeterminacy and Learning: An Analysis of Monetary Policy in the Great Inflation”. *Journal of Monetary Economics*, 82, pp. 85-106.
- [30] Onatski, Alexei, and Noah Williams (2003): “Modeling Model Uncertainty”. *Journal of the European Economic Association*, 1(5), pp. 1087-1122.
- [31] Pancrazi, Roberto (2015): “The Heterogeneous Great Moderation”. *European Economic Review*, 74, pp. 207-228.
- [32] Primiceri, Giorgio (2006): “Why Inflation Rose and Fell: Policymakers’ Beliefs and US Postwar Stabilization Policy”. *The Quarterly Journal of Economics*, 121, pp. 867-901.
- [33] Sala, Luca (2015): “DSGE Models in the Frequency Domain”. *Journal of Applied Econometrics*, 30, pp. 219-240.
- [34] Sargent, Thomas, Noah Williams, and Tao Zha (2006): “Shocks and Government Beliefs: The Rise and Fall of American Inflation”. *American Economic Review*, 96(4), pp. 1193-1224.
- [35] Shimer, Robert (2005): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies”. *American Economic Review*, 95(1), pp. 25–49.
- [36] Smets, Frank, and Raf Wouters (2003): “An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area”. *Journal of the European Economic Association*, 1(5), pp. 1123-1175.
- [37] Smets, Frank, and Raf Wouters (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach”. *American Economic Review*, 97(3), pp. 586-606.
- [38] Tkachenko, Denis, and Zhongjun Qu (2012): “Frequency Domain Analysis of Medium Scale DSGE Models with Application to Smets and Wouters (2007)”. *Advances in Econometrics : DSGE Models in Macroeconomics – Estimation, Evaluation and New Developments*, Volume 28, pp. 319-385.

- [39] Watson, Mark W. (1986): “Univariate Detrending Methods with Stochastic Trends”. *Journal of Monetary Economics*, 18, pp. 49-75.
- [40] Watson, Mark W. (1993): “Measures of Fit for Calibrated Models”. *Journal of Political Economy*, 101(6), pp. 1011-1041.

Table 1: Variance Decomposition US 1954 - 2017

	Short Term		Business Cycle		Medium Term		Long Term
	D1: 2-4Q	D2: 4-8Q	D3: 8-16Q	D4: 16-32Q	D5: 32-64Q	D6: 64-128Q	S6: >128Q
GDP Growth	32.7	22.8	19.0	12.6	6.1	2.5	4.3
Unemployment	1.3	3.8	9.6	18.5	25.4	20.3	21.1
Inflation	8.6	7.4	7.7	9.9	9.3	15.7	41.3
Federal Funds	1.5	2.8	5.6	10.2	10.5	12.8	56.6
3-Month Rate	1.3	2.4	4.8	8.9	10.1	13.1	59.3
10-Year Rate	0.6	1.3	2.5	4.0	7.5	15.5	68.5
Term Spread	5.5	9.7	17.1	26.7	19.1	5.4	16.5

Table 2: Variance Decomposition for Simulated Data

	Short Term	Business Cycle	Medium Term	Long Term
<u>ΔRGDP</u>				
Data	56	32	8	4
SW	59	29	10	2
	(48-71)	(21-38)	(5-18)	(1-5)
dNGS	60	28	10	2
	(50-70)	(20-35)	(5-17)	(1-3)
CET (Nash)	65	27	6	2
	(56-74)	(21-34)	(3-9)	(1-5)
CET (AOB)	27	40	23	10
	(20-35)	(29-50)	(14-37)	(3-22)
<u>Inflation</u>				
Data	16	18	25	41
SW	20	35	29	15
	(13-31)	(22-48)	(14-36)	(3-36)
dNGS	17	20	29	34
	(7-31)	(8-32)	(14-47)	(9-66)
CET (Nash)	27	39	18	16
	(18-37)	(24-50)	(9-29)	(3-42)
CET (AOB)	13	22	20	44
	(4-25)	(6-41)	(8-36)	(10-80)
<u>FFR</u>				
Data	4	16	24	57
SW	16	36	33	17
	(9-23)	(23-50)	(16-51)	(3-38)
dNGS	12	24	33	31
	(5-21)	(10-39)	(17-53)	(8-61)
CET (Nash)	25	38	19	16
	(16-34)	(25-50)	(11-31)	(3-41)
CET (AOB)	10	23	25	43
	(3-17)	(7-42)	(11-44)	(9-79)
<u>Unemployment</u>				
Data	5	29	45	21
CET (Nash)	6	18	30	46
	(2-11)	(5-33)	(12-52)	(13-79)
CET (AOB)	1	10	30	59
	(0-3)	(2-21)	(10-57)	(23-86)

Table 3: Variance Decomposition for Simulated Data - Alternative Sample

	Short Term	Business Cycle	Medium Term	Long Term
<hr/>				
Smets-Wouters				
<hr/>				
FFR				
Full Sample	4	16	24	57
SW Sample	8	25	35	32
Simulated	16	36	33	17
dNGS				
<hr/>				
FFR				
Full Sample	4	16	24	57
dNGS Sample	7	25	32	36
Simulated	12	24	33	31
CET (AOB)				
<hr/>				
FFR				
Full Sample	4	16	24	57
CET Sample	6	21	28	45
Simulated	10	23	25	43
Unemployment				
Full Sample	5	29	45	21
CET Sample	6	27	37	30
Simulated	1	10	30	59
<hr/>				




Figure 1: Wavelet Decompositions: Real GDP Growth




Figure 2: Wavelet Decompositions: Unemployment




Figure 3: Wavelet Decompositions: Inflation




Figure 4: Wavelet Decompositions: Federal Funds Rate




Figure 5: Wavelet Decompositions: 10-Year Treasury Rate




Figure 6: Wavelet Decompositions: Term Spread




Figure 7: Wavelet Power Spectra: Real GDP Growth




Figure 8: Wavelet Power Spectra: Unemployment

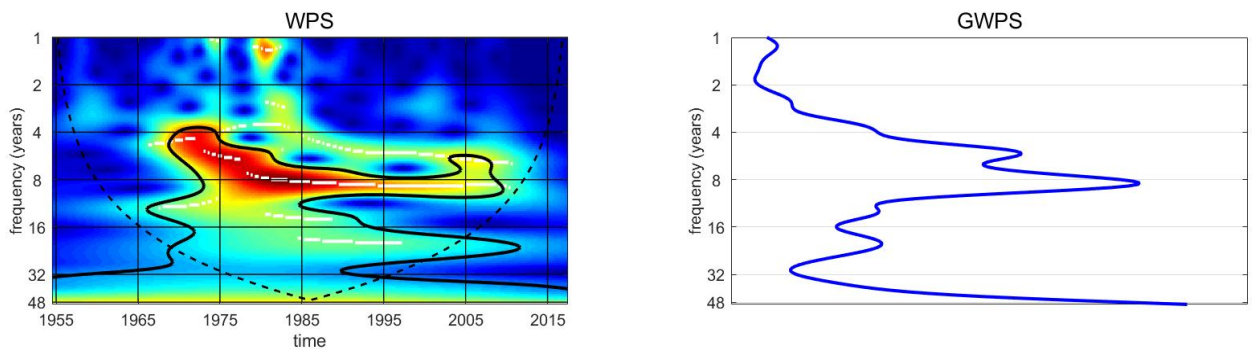


Figure 9: Wavelet Power Spectra: Federal Funds Rate

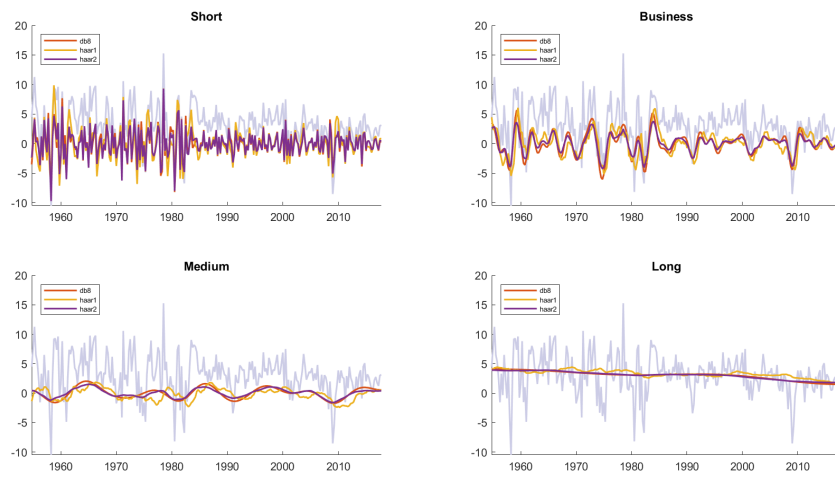


Figure 10: Wavelet Decompositions for Alternative Filters: Real GDP Growth



Figure 11: Wavelet Decompositions for Alternative Filters: Unemployment

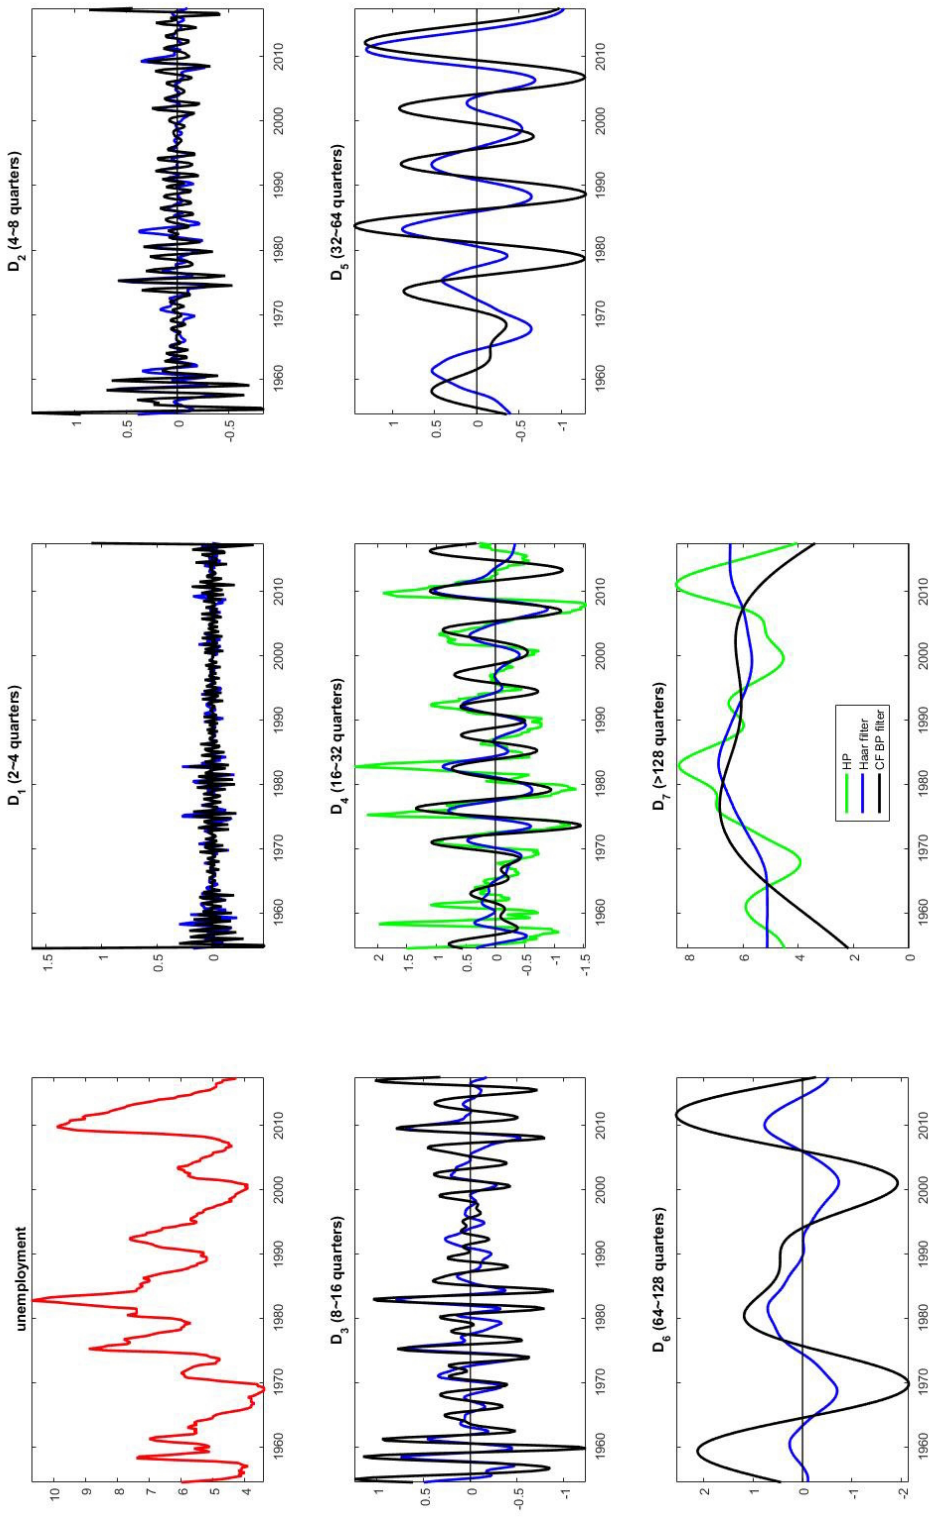


Figure 12: Decompositions with HP and CF Filters: Unemployment




Figure 13: Impulse Response Functions: 3-Variable Baseline VAR with Sign Restrictions.





Figure 14: Impulse Response Functions with D2 Components




Figure 15: Impulse Response Functions with D4 Components





Figure 16: Impulse Response Functions with D7 Components

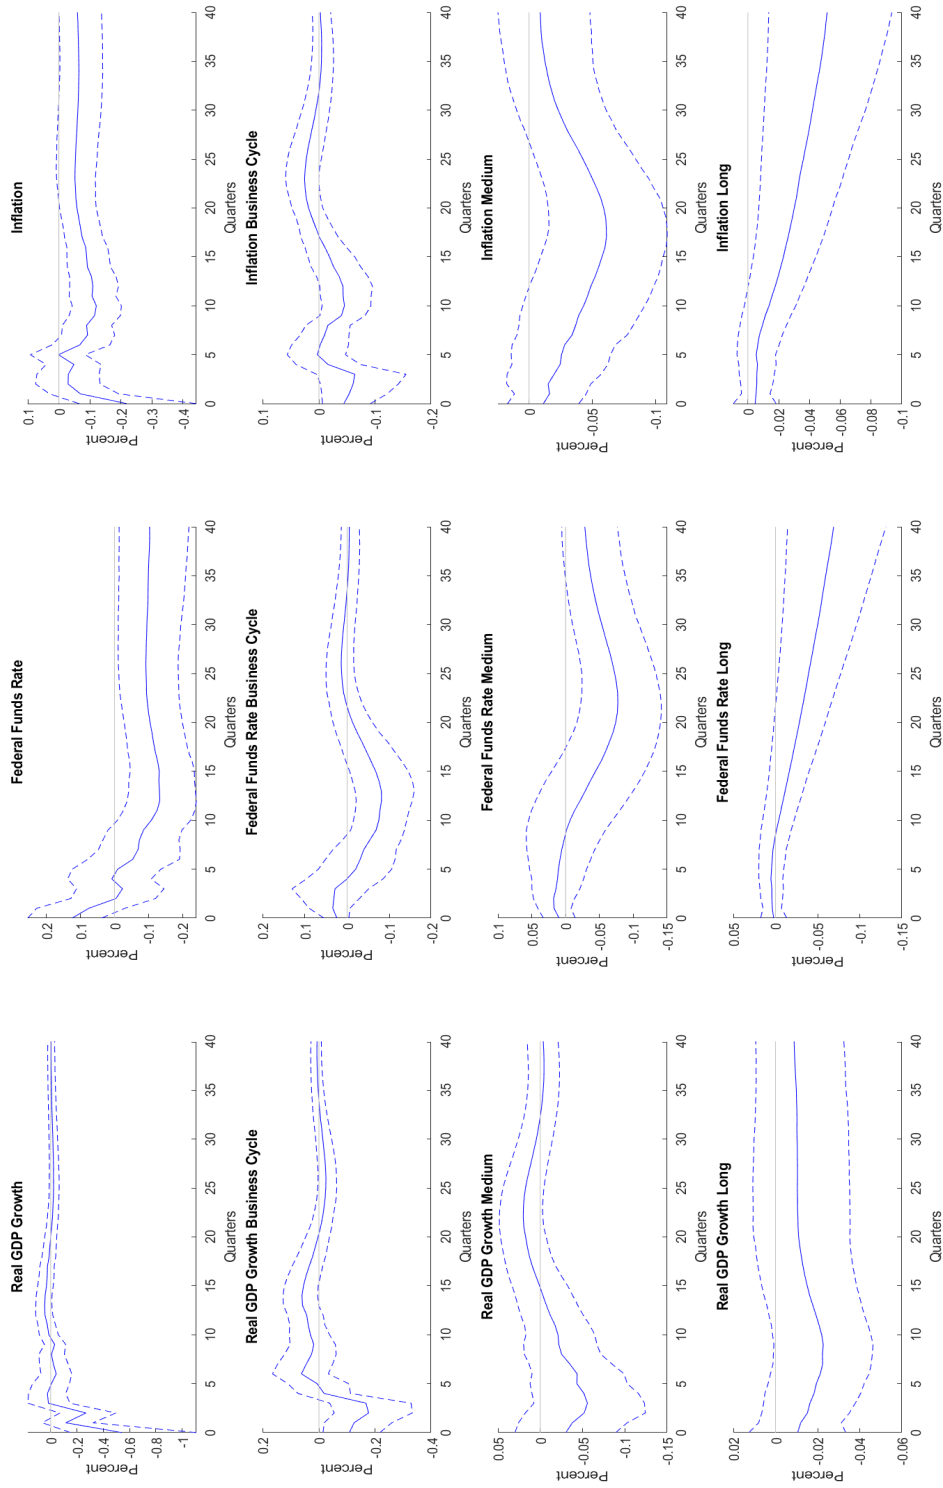


Figure 17: Impulse Response Functions with All Components: GDP Growth




Figure 18: Impulse Response Functions with All Components: Unemployment