cef. up
Centro de Economia e Finanças da UPorto

cef.up working paper
2014-04

Estimating the Taylor Rule in the Time-Frequency Domain

Luís Aguiar-Conraria
Manuel M. F. Martins
Maria Joana Soares

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

Estimating the Taylor Rule in the Time-Frequency Domain

Luís Aguiar-Conraria* Manuel M. F. Martins† Maria Joana Soares‡

April 20, 2016

Abstract

We assess U.S. monetary policy across time and frequencies in the framework of the Taylor Rule (TR). With that purpose, we derive a multivariate generalization of the wavelet gain — the partial wavelet gain — a new tool which allows us, for the first time, to estimate the TR coefficients in the time-frequency domain. By using this and other continuous wavelet tools, we reach a number of results regarding the evolution of the TR coefficients along time that also have a frequency-domain nature — for example, the inflation coefficient has violated the Taylor principle unevenly across frequencies, and the evidence of a modified TR with a unit slope on output since 2009 is also uneven along time and across frequencies.

Keywords: Monetary Policy; Taylor Rule; Partial Wavelet Gain; Time-Frequency Estimation; Continuous Wavelet Transform.

JEL codes: C49, E43, E52.

*NIPE and Department of Economics, University of Minho, E-mail address: lfaguiar@eeg.uminho.pt. Corresponding author.

†Cef.up and Faculty of Economics, University of Porto, E-mail address: mmfmartins@fep.up.pt

‡NIPE and Department of Mathematics and Applications, University of Minho, E-mail address: jsoares@math.uminho.pt

This is a vastly revised and extended version of a paper that circulated previously under the title "Analysing the Taylor Rule with Wavelet Lenses".

1 Introduction

This paper uses continuous wavelet tools to estimate the coefficients of the Taylor Rule implicit in U.S. monetary policy between 1965 and 2014. The simultaneous variation of coefficients along time and frequencies, and the thorough statistical analysis provided by our tools, allow for detecting new stylized facts about the last five decades of U.S. monetary policy.

Taylor (1993) showed that the policy in 1986-92 was very well described by the simple parametric relation between the policy interest rate, the output gap and inflation

$$\mathrm{FFR}_t = 2 + \pi_t + \frac{1}{2} y_t + \frac{1}{2} (\pi_t - 2), \tag{1}$$

in which FFR is the (effective) federal funds rate, π is the inflation rate over the previous four quarters, y is the percent deviation of output from its potential and both the real equilibrium interest rate and the inflation target are assumed to equal 2 percent.

A first worth of the Taylor Rule (TR) is positive, and consists of its ability to parsimoniously describe U.S. monetary policy. Indeed, subsequent empirical studies have shown that such broad empirical success extends to periods before 1986 and after 1992, which is particularly noticeable given that, as documented *inter alia* by Kahn (2012) and Taylor (2012), there were frequent references to Taylor-type rules in the Federal Open Market Committee meetings since 1993 but not before — when policy discussions and decisions were more discretionary and focused on fine-tuning real activity with no special focus on long-run price stability.

A second worth of the TR is normative, as it came to be considered a useful benchmark for monetary policy, highly valuable to inform and aid policymakers' decisions, even if not to be followed mechanically. In fact, being an approximation to the optimal control solution of the monetary policy-maker's problem, the TR has proved to be quasi-optimal and more robust than a wide array of strictly optimal policy rules derived in specific macroeconomic models — see e.g. Taylor and Williams (2010). Moreover, it has the advantage that its simplicity makes it very easy to communicate and understand. Consistently, monetary policy is systematically modeled with a Taylor-type Rule in the New-Keynesian dynamic stochastic general equilibrium models that are the current mainstream for monetary policy conduct and analysis. Recently, given arguments that under the TR policy is more predictable, systematic, and thus effective (Taylor, 2012), the suggestion by Taylor and Williams (2010) that it could become an accountability device has received attention in policy circles, with a Bill introduced to the U.S. Congress then fueling discussions in academic circles — see e.g. Bernanke (2015).[1]

To deal with the observation that, in spite of the overall very good fit of the interest rates prescribed by the original TR, there are several episodes of systematic deviations between the FFR and the implied TR rate, the literature has explored essentially three avenues. First, episodes of substantial deviations have been formally identified and described as eras of discretionary monetary policy, as opposed to rules-based eras, with the latter associated with higher macroeconomic stability — e.g. Taylor (2012) and Nikolsko-Rzhevskyy, Papell, and Prodan (2014). Second, the original TR has been extended with several additional and refined explanatory variables — e.g. Clarida, Galí and Gertler (2000), Sims (2013), Sack and Rigobon (2003), Lubik and Schorfheide (2007), and Christensen and Nielsen (2009).

A third approach has allowed the coefficients of the TR to vary along time. Following Clarida, Galí and Gertler's (2000) finding that the U.S. interest rate policy has been more sensitive to inflation after 1979, the stability of the U.S. TR has been assessed with several time-series methods, such as threshold models (Bunzel and Wenders, 2010), time-varying parameters models (Trecroci and Vassalli, 2010), Markov-switching models (Assenmacher-Wesche, 2006), smooth-transition models (Alcidi, Flamini and Fracasso, 2011), instrumental variables quantile regressions (Wolters, 2012), and Hamilton's (2001) flexible approach to nonlinear inference (Kim, Osborn and Sensier, 2005). Time variation in the TR coefficients — notably the increase in inflation's coefficient — has been associated with policy regime changes and, in turn, phenomena as the Great Moderation (Canova, 2009) and the decline in inflation persistence (Benati, 2008). Considering that the TR is

[1] H.R. 5018 (113th) Federal Reserve Accountability and Transparency Act of 2014, discussed in the House Financial Service Committee, according to which the FED should explain to the House any systematic deviations of the policy interest rates from a reference policy interest rate that would correspond precisely to that implied by Taylor's (1993) Rule presented in (1). For details, see https://beta.congress.gov/113/bills/hr5018/BILLS-113hr5018ih.pdf

an approximation to the optimal control solution of the policymaker's maximization problem, its functional form and coefficients depend on the policymaker's preferences, on the structure of the economy, as well as on the information considered by both policymakers and the public. Hence, the TR may change for many reasons, as the literature has thoroughly addressed in the recent years — e.g. Favero and Rovelli (2003), Owyang and Ramey (2004), Dennis (2006), Surico (2007), Dolado, Maria-Dolores and Naveira (2005), Alcidi, Flamini and Fracasso (2011), Tillmann (2011), Hamilton, Pruitt and Borger (2011).

While the literature documents very thoroughly how TR coefficients change along time, it does not assess whether that variation has different intensity for different cyclical oscillations. There are, however, ample reasons for the TR coefficients to change differently at distinct frequencies. First, as the main focus of monetary policy is cyclical stabilization, one key concern of policymakers should be to understand and control which specific cyclical oscillations they want to, can, and do control at each period of time. For example, policymakers should care about the impact of policy across cyclical frequencies because oscillations at different frequencies may have different impacts on social welfare; or because controlling oscillations at some frequencies may imply a trade-off with larger variability at other frequencies (Yu, 2013); moreover, different circumstances may recommend different choices regarding these frequency-domain trade-offs. Second, while it is arguable that during most of the time policymakers react more strongly to persistent than to short-lived fluctuations in the main macroeconomic variables, the relative importance of controlling low versus medium and versus high frequency oscillations may change with circumstances; indeed, the well-known discussion about which inflation rate to consider in the TR — whether headline inflation or core inflation, which features smaller high-frequency variation — is an example of the difficulty in finding a once and for all best indicator for policy, given the frequency-domain trade-offs faced by the policymaker (see e.g. Mehra and Sawhney, 2010). Third, changes in the monetary policy regime may be closely related with changes in the relative intensity of the policy reaction at different frequencies; for example, a policymaker trying to conquer credibility may have to react very strongly to transitory changes in inflation, but once credibility is established, he or she may increase the focus on fluctuations in inflation of a more permanent nature.

Hence the motivation for this paper: to thoroughly describe the changes in the U.S. TR along time and across frequencies. We use continuous wavelet tools, with an approach consisting of a sequential analysis of partial wavelet coherencies, phase-difference diagrams and gains. The partial coherencies and phase-diagrams determine, for each time and frequency, the significance, sign and synchronization (lags or leads) between the policy interest rate and each of the macroeconomic variables in the TR, controlling for the other variable; the partial gains provide estimates of the coefficients associated to each macro variable in the TR, along time and across frequencies.

At the methodological level, our main contribution to the literature is to provide a multivariable generalization of the wavelet gain that allows for estimating multivariate functions in the time-frequency domain. We also use the multiple coherency (see Aguiar-Conraria and Soares, 2014) jointly with the partial coherency (and partial phase-difference) to refine the interpretation of the estimates given by the partial gain.

Regarding the coefficient on inflation, we emphasize four findings. First, it has changed much more markedly for cycles of intermediate duration than for longer cycles and shorter cycles likewise. Second, rather than a change from a constant coefficient below 1.0 before 1979 to a coefficient above 1.0 after 1979, there has been a gradual decrease of the inflation coefficient until the mid-1970s, followed by an increase after 1979 that is essentially completed at the start of the 1985-2003 rules-based era, and is most marked at the core business cycle frequencies. Third, the Taylor principle has been violated unevenly across frequencies, with the estimate of the inflation slope below 1.0 for longer the lower the corresponding frequencies. Fourth, since the mid-1980s, for cycles of period above 4 years the coefficient on inflation in the TR is consistently above the baseline value and full sample estimate of 1.5.

As regards the coefficient on the output gap, we emphasize two main findings. First, the full sample OLS estimate of 0.5 seems to be an artifact resulting from different coefficients across frequencies and along time. Second, the time-series evidence and policy-makers' statements pointing to a modified TR with a slope of 1.0 on the output gap in the U.S. TR since 2009 is not evenly explained across frequencies, as it is associated to stronger reactions of policy to output at the short-end and long-end of cyclical oscillations, but not at the most standard business cycle frequencies.

The paper proceeds as follows. In Section 2, we intuitively describe our methodology and in

Section 3 we describe the data. In Section 4 we apply wavelet tools to the data and provide a continuous time-frequency assessment of the U.S. TR. Section 5 concludes. In the appendix, we provide a self contained summary of our methodology, with an emphasis on our main methodological contribution: the partial wavelet gain that we use to estimate the coefficients of the Taylor Rule in the time-frequency domain.

2 Methodology

The continuous wavelet transform is an increasingly popular tool in econometric analysis. The most common argument to justify its use is the possibility of tracing transitional changes across time and frequencies — see Aguiar-Conraria and Soares (2014) for a review. So far, the analysis in the time-frequency domain with the continuous wavelet transform has been mostly limited to the use of the wavelet power spectrum, the wavelet coherency and the wavelet phase-difference. Aguiar-Conraria and Soares (2014) already extended these tools to allow for multivariate analyses. These multivariate tools are sufficient to assess the strength of the relation between several variables, but they are insufficient to estimate the magnitude of the relation. Just like (partial) correlation coefficients do not provide the same information as the regression coefficients.

Mandler and Scharnagl (2014) use the concept of the wavelet gain as a regression coefficient in the regression of y on x. In this paper, and, to our knowledge, for the first time, we will estimate an equation relating more than two variables (just like a regression of y on x and z) in the time-frequency domain. To do so, we generalize the concept of wavelet gain and define the *partial wavelet gain*, which can be interpreted as a regression coefficient in the regression of y on x after controlling for other variables. We will proceed in a non-conventional fashion and leave to the appendix all the technical details and the mathematical derivation of the relevant formulas. In the main text, we simply provide the formulas for the particular case of three variables and a constructed example that will illustrate our claim that by estimating the partial gain one is essentially estimating an equation in the time-frequency domain.

2.1 Partial wavelet gain with three variables

We illustrate the use of the formulas derived in the appendix, in terms of bivariate coherencies, for the case where we just have three series x_1, x_2 and x_3. In this case, it can easily be shown (see Aguiar-Conraria and Soares 2014) that the multiple wavelet coherency is given by:

$$R^2_{1(2\,3)} = \frac{R^2_{12} + R^2_{13} - 2\Re\left(\varrho_{12}\,\varrho_{23}\,\overline{\varrho_{13}}\right)}{1 - R^2_{23}}, \tag{2}$$

where ϱ_{ij} is the complex wavelet coherency and R_{ij} the wavelet coherency between x_i and x_j, $\Re$ means that we are collecting the real part, and the upperbar is used to denote complex conjugation. The complex partial wavelet coherency between x_1, and x_2, after controlling for x_3, is given by:

$$\varrho_{1\,2.3} = \frac{\varrho_{12} - \varrho_{13}\overline{\varrho_{23}}}{\sqrt{(1 - R^2_{13})(1 - R^2_{23})}}. \tag{3}$$

On the other hand, applying formula (A.13), it is easy to show that the partial wavelet gain $G_{1\,2.3}$ is given by:

$$G_{1\,2.3} = \frac{\left|\varrho_{12} - \varrho_{13}\overline{\varrho_{23}}\right|}{(1 - R^2_{23})}\,\frac{\sigma_1}{\sigma_2}. \tag{4}$$

Note that each of the above quantities is a function of time and frequency. In the case of the last formula, for example, we have the partial gain between x_1, x_2 (after controlling for x_3) for different times and frequencies. Therefore, we do not have one number. We have a matrix, whose information must be summarized for tractability.

2.2 Example: Partial gain, coherency and phase-difference

We now give a constructed example illustrating the application the wavelet gain and partial wavelet gain, proposed in this paper. Given the full control of the data generating processes, our example makes it clear that the partial wavelet gain may be interpreted as a regression coefficient in the time-frequency domain. The example also highlights that, because the (partial) wavelet gain is an absolute value, its interpretation must be associated with that of the wavelet (partial) phase-difference, which will tell us if the relation is positive or negative and will also tell us which variable is leading or lagging.

Figure 1 is a guide to interpret the (partial) phase-difference, between the two series x and y (controlling for the third variable): a partial phase-difference with value zero indicates that the time-series move together at the specified frequencies; if the partial phase-difference lies in the interval $(0, \frac{\pi}{2})$, then the series move in phase, but the time-series x leads y; if the partial phase-difference is in $(-\frac{\pi}{2}, 0)$, then it is y that is leading; a partial phase-difference of π indicates an anti-phase relation; if the partial phase-difference is in $(\pi/2, \pi)$, then y is leading; time-series x is leading if the partial phase-difference lies in $(-\pi, -\frac{\pi}{2})$.

Imagine that we have monthly data and that the data generating processes for X and Z are given by

$$X_t = \sin\left(2\pi \tfrac{t}{3}\right) + \sin\left(2\pi \tfrac{t}{8}\right) + \varepsilon_{x,t},$$
$$Z_t = \sin\left(2\pi \tfrac{t}{9}\right) + \varepsilon_{z,t},$$

while for Y is given by

$$Y_t = \begin{cases} 2\sin\left(2\pi \frac{t+3/12}{3}\right) + 1\sin\left(2\pi \frac{t-1}{8}\right) + Z_t + \varepsilon_{y,t}, & \text{for } t \leq 100 \\ 2\sin\left(2\pi \frac{t+3/12}{3}\right) - 3\sin\left(2\pi \frac{t-1}{8}\right) + Z_t + \varepsilon_{y,t}, & \text{for } t > 100 \end{cases}.$$

Suppose that we are interested in regressing Y against X in the time-frequency domain. What should we expect?

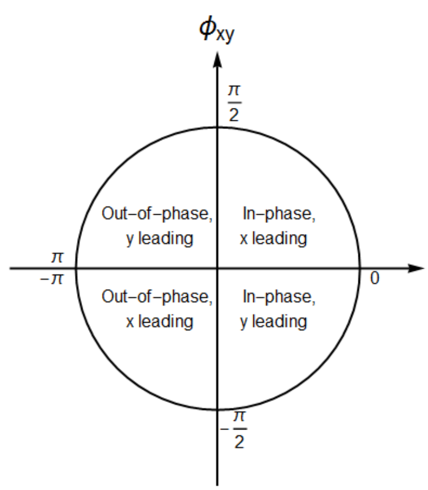

Figure 1: Phase-difference circle.

At frequencies that correspond to a period of 3 years, the estimated coefficient should be 2 throughout the sample, implying that the wavelet gain should be 2 also. The phase-difference should also indicate that Y slightly leads (by 3 months) X, meaning that the phase-difference between Y and X should be between 0 and $\pi/2$.

At the frequency corresponding to a 8 year period, the coefficient should be $+1$ in the first half of the sample and -3 in the second half. However, given that the wavelet gain is an absolute value, it would yield an estimate of $+3$ for the coefficient in the second half of the sample. To capture the negative sign of the relation, one has to use the information given by the phase-difference. In the first half of the sample, at this frequency, Y lags X (by 1 year) and the variables are in-phase. Therefore, the phase difference should be between $-\pi/2$ and 0. In the second half, Y lags X (by 1 year) and the variables are out-of-phase. Therefore, the phase-difference should be between $\pi/2$ and π.

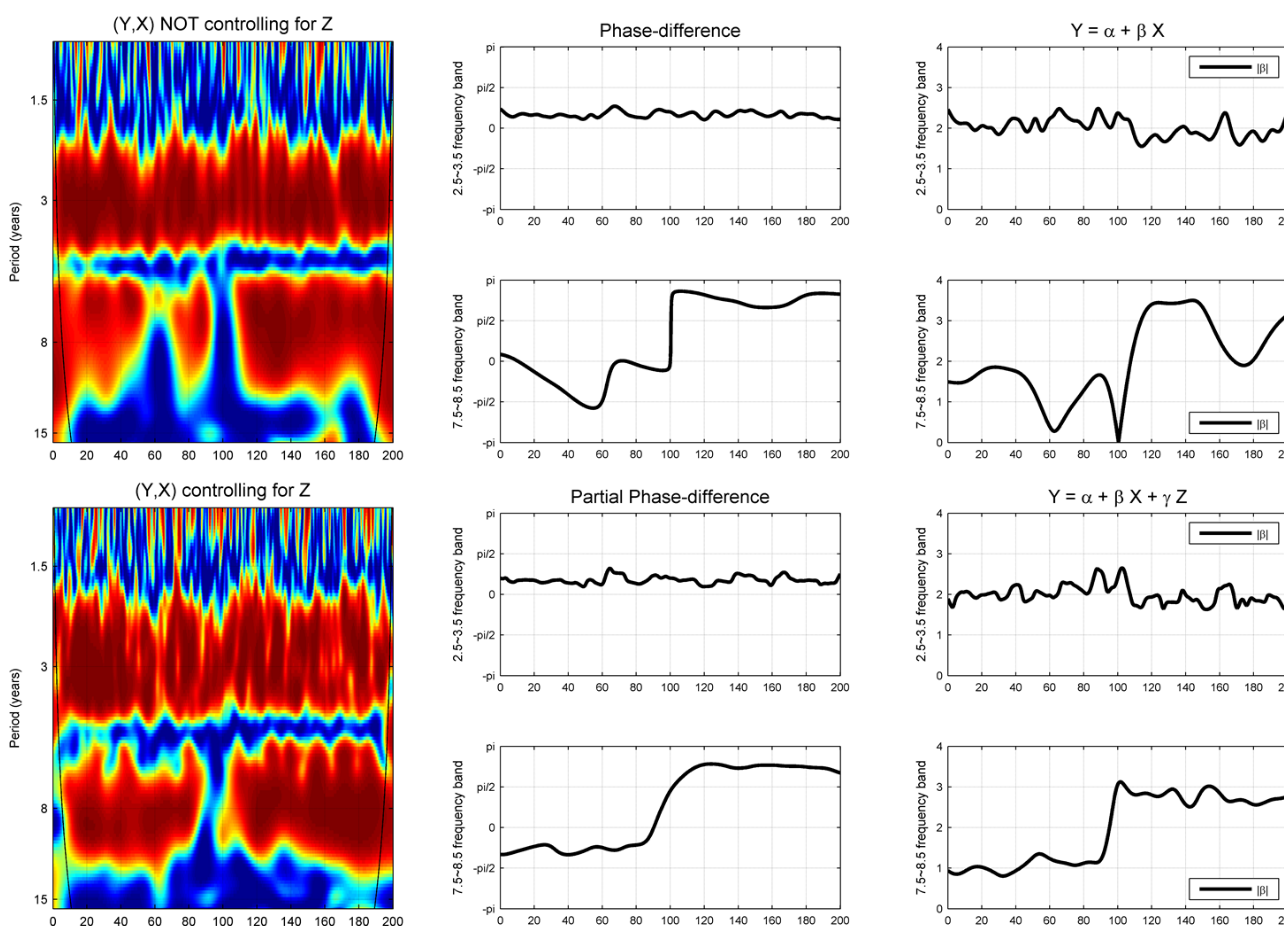

Figure 2: on the left — wavelet coherency between Y and X (top) and partial wavelet coherency between Y and X, after controlling for Z (bottom). The color code for coherency ranges from blue (low coherency –close to zero) to red (high coherency –close to one). In the middle –phase-differences (top) and partial phase-differences (bottom) between Y and X. On the right –wavelet gain (top) and partial wavelet gain of Y over X, after controlling for Z (bottom).

Finally, note the influence of Z on variable Y: given that its influence occurs at the frequency corresponding to a 9 year period, excluding this variable, and therefore incurring in an omitted variable bias, should contaminate the relation between Y and X at the frequency corresponding to a 8 year period.

Figure 2 displays the results obtained with the use of the referred wavelet tools. In the left panel of the figure, we plot the wavelet coherency between Y and X (top) and partial wavelet coherency between Y and X after controlling for Z (bottom). In the middle part of the figure we present the (circular) means of the phase-differences (two top figures) and of the partial phase-differences (two bottom figures) corresponding to two different frequency bands (one for periods of 2.5 to 3.5 years and the other for periods of 7.5 to 8.5 years). On the right of the figure, we display the means of the wavelet gain of Y over X (two top figures) and of the partial wavelet gain of Y over X controlling for Z (two bottom figures), corresponding to the same frequency bands.

All the results we were expecting are confirmed in Figure 2. In particular, note how the relations between Y and X around the 8 year period are much more accurately estimated when we use the partial wavelet tool (which corresponds to control for the effects of variable Z).

3 The Data

Our data are quarterly time-series of the federal funds rate (FFR), inflation and the output gap, for the U.S. 1965:IV-2014:IV and correspond to the data used by Nikolsko-Rzhevskyy, Papell, and Prodan (2014) updated through the end of 2014. These are real-time data that were available to policymakers when interest rate decisions were made, consistently with the sort of data used in the vast majority of empirical research on monetary policy rules since Orphanides (2001). The source for output and inflation is the Real-Time Data Set for Macroeconomists created by Croushore and Stark (2011) and available at the Philadelphia Federal Reserve website, which provides vintages of data available since 1965:IV with the data in each vintage starting in 1947:I.[2]

Inflation is the year-over-year rate of change of the real-time GDP deflator. The output gap is the percent difference between real GDP and a real-time quadratic trend, i.e. a trend obtained

[2] http://www.philadelphiafed.org/research-and-data/real-time-center/real-time-data/data-files/.

fitting a quadratic function of time to the real GDP data from 1947:I through the vintage date (see Nikolsko-Rzhevskyy, Papell, and Prodan 2014 for further details, namely on the choice of the functional form for the trend and on timing issues). It should be emphasized that the output gap is detrended output, rather than filtered output, meaning that only its long-run variation has been removed, and thus analyses of oscillations at all other frequencies are warranted.

The source for the FFR is the FRED (Federal Reserve Economic Data) available at the website of the Federal Reserve of St. Louis, until 2008:IV.[3] From 2009:I onwards, when the policy interest rate has been constrained by the zero lower bound, we use the shadow FFR of Wu and Xia (2014) which is computed from a nonlinear term structure model and captures the overall monetary policy stance, including the effects of unconventional policies.[4]

In Figure 3, we plot the three time-series, on the left-hand side charts, and their wavelet power spectra, on the right-hand side, which measure the variance of the series at each time-frequency locus and provide a first time-frequency description of the data.

A first overall conclusion is that, with the exception of the output and inflation instability of the 1970s, the variability of the three time-series occurs at frequencies corresponding to periods larger than 4 years.

The chart of inflation shows its well-known gradual rise between the mid-1960s and the 1970s, the disinflation between 1980 and 1986, and the ensuing period of low and stable inflation, with particularly low rates following the recent financial and economic crisis. The wavelet power spectrum of inflation shows that during the inflationary period it has oscillated most specially at business-cycle frequencies ($4 \sim 8$ years). After that, during the disinflation period, the areas of statistically significant power spectrum become gradually thinner — which illustrates the subsequent anchoring of inflation (and its expectations) and the prolonged period of very low inflation variance during the Great Moderation.

The chart of the output gap shows the strong recession associated with the first oil shock in the mid-1970s, as well as the recession in the early 1980s associated with disinflation; it then shows the Great Moderation between 1984 and 2007, and the Great Recession starting in 2008. The

[3] http://research.stlouisfed.org/fred2/.

[4] http://faculty.chicagobooth.edu/jing.wu/research/data/WX.html.

wavelet power spectrum indicates a prevalence of cyclical oscillations with periods below 12 years until 1985 — namely 6 $\sim$ 8 years and 10 $\sim$ 12 years —, and then that shorter cycles gradually lost importance to longer cycles along the sample, with cyclical variability concentrated at cycles with a rather long period at the final part of the sample.

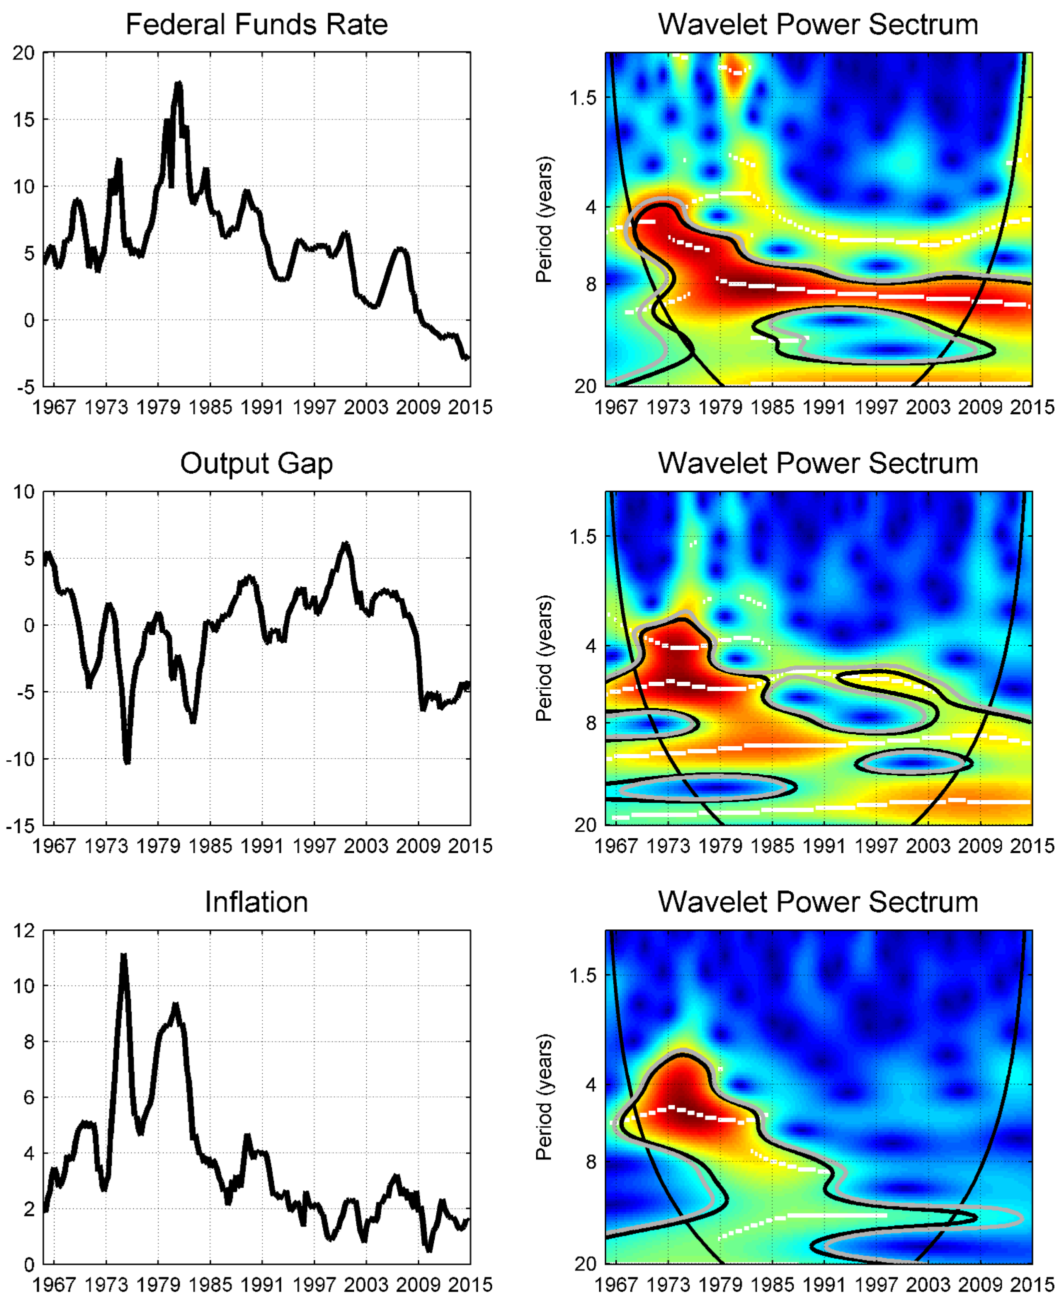

Figure 3: On the left: Plot of each time-series. On the right: The corresponding wavelet power spectrum. The black/gray contour designates the 5%/10% significance level. The cone of influence, which indicates the region affected by edge effects, is shown with a parabola-like black line. The color code for power ranges from blue (low power) to red (high power). The white lines show the local maxima of the wavelet power spectrum.

The chart of the federal funds rate (FFR) shows that nominal interest rates tended to increase with inflation since the mid-1960s, peaked at very high levels at the beginning of the 1980s and then gradually decreased until the end of the sample. The power spectrum of the FFR indicates that throughout the whole sample the variability of the policy rate has been systematically strong at cyclical frequencies of $8 \sim 10$ years, even though with particular strength during the disinflation, but has also been strong at shorter cycles ($4 \sim 8$ years) during the 1970s.

In Figure 4, we plot the Federal Funds Rate (FFR) and the Reference Policy Rule (RPR), i.e. the interest rate computed with equation (1) using our real-time output gap and inflation data. The figure conveys two main messages.

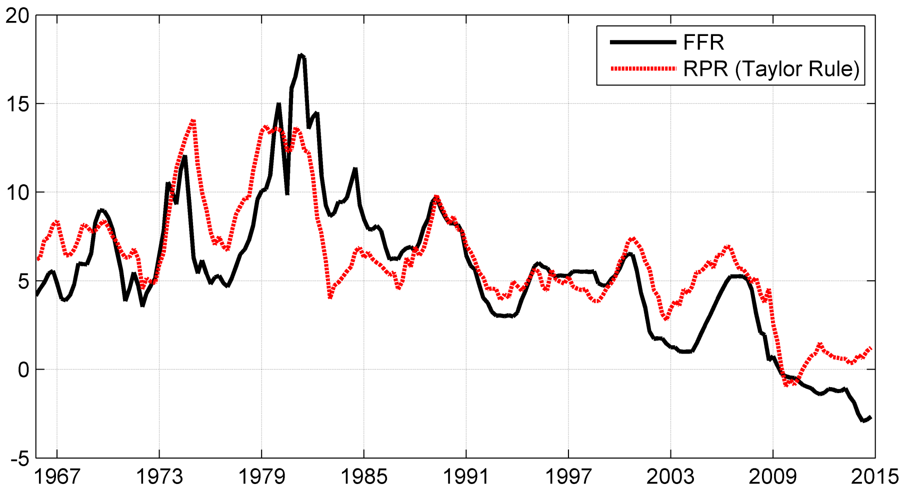

Figure 4: The Classic Taylor Rule — the proposed Reference Policy Rule, computed with our real-time output gap and inflation — and the Federal Funds Rate since 1965 (effective for 1965:IV-2008:IV, shadow for 2009:I-2014:IV).

First, it is remarkable how the original TR broadly mimics the overall path of the policy interest rate, given its simplicity, given that there are no references to interest rate rules in the Federal Open Market Committee discussions before 1993 (Kahn, 2012) — when policy was much more discretionary and policymakers' discussions focused on fine-tuning real activity and not on maintaining long-run price stability (Taylor, 2012) — and given that policymakers have never committed to a specific TR. The overall compliance of U.S. monetary policy to the TR gains support from the results of an OLS regression of the TR with our real-time data for 1965:IV-2014:IV (standard errors in parenthesis): $\mathrm{FFT}_t = \underset{(0.30)}{0.17} + \underset{(0.07)}{1.54}\pi_t + \underset{(0.05)}{0.51}y_t$. The estimates for the coefficients on inflation and the output gap are almost the same as the original formulation of

Taylor's Rule (the smaller intercept may be due to a higher inflation target or to a lower equilibrium real interest rate).

Second a closer look at the figure reveals that in many periods the FFR looks persistently close to the RPR, while in others it deviates systematically from the RPR. One approach in the literature has considered the former as episodes of rules-based policy, and the latter as ones of discretionary policy, typically associating better macroeconomic outcomes with the former. Notably, Taylor (2012) identifies a period of discretionary policy before 1985, with the FFR below the RPR during an era of fine-tuning until 1979, and above the RPR in the disinflation after 1979; in 1985-2003, policy has been rules-based, with its predictable systematic approach arguably key to the Great Moderation; in 2003-06, he identifies a policy of FFR substantially and persistently below the RPR, which he terms the Great Deviation (Taylor, 2011, 2012) and relates with the boom that led to the 2008 bust and the Great Recession; since then, policy has essentially been ad-hoc, as the FFR has been consistently below the RPR, including negative interest (shadow) interest rates that would not be prescribed by the original Taylor Rule. In the same vein, but using formal structural breaks tests, Nikolsko-Rzhevskyy, Papell, and Prodan (2014) find that the FFR followed quite closely the original Taylor Rule in 1965:IV-1974:III and in 1985:II-2001:I, deviating substantially from the rule in 1974:IV-1985:I and in 2001:II-2013:IV, with the former period split into one of too low interest rates (until 1979:IV) and another of too high interest rates (from 1980:I to 1985:I).

Interestingly, when they use a modified Taylor Rule with a coefficient of 1 on the output gap, Nikolsko-Rzhevskyy, Papell, and Prodan (2014) detect a further break and identify a period of rules-based policy in 2006:IV-2013:IV. Such modified rule is consistent with statements by the Governor of the FED pointing out that the implied interest rates are closer to those given by the optimal control solution of the FRB/US model than the interest rates implied by the original TR — see Bernanke (2011) and Yellen (2012). Indeed, it prescribes negative interest rates since 2009 — in line with the shadow FFR depicted in Figure 2 — which the original Taylor Rule does not — as also shown in the picture.

In this paper, rather than comparing the U.S. policy interest rate with the one given by the original TR, and rather than seeking for alternative coefficients that could improve the fit of a modified TR to the policy interest rate data, we give due consideration to the arguments (surveyed in section 1) that the policy rule coefficients may change both along time and across frequencies. Hence the worth of the continuous wavelet transform tools that we use in the next section to estimate the U.S. Taylor Rule in the time-frequency domain.

4 Results: The Taylor rule in the Time-Frequency domain

We now assess the relation between the FFR and the macroeconomic variables of the Taylor Rule in the time-frequency domain, using multivariate continuous wavelet tools, in particular a generalization of the wavelet gain for the case of functions with more than one explanatory variable.

We start with the multiple coherency,[5] which is the time-frequency analog of the R^2 in the typical regression. Then, we present and discuss the partial coherency, the partial phase-difference, and the partial gain between the FFR and each of the macroeconomic variables in the Taylor Rule, controlling for the effects of the other. The latter corresponds to estimating the coefficients associated to each macro variable in the TR allowing for their variation along time and across frequencies – i.e. estimating the Taylor Rule coefficients in the time-frequency domain.

While the interpretation of our econometric results proceeds along the standard approach in similar literature for the coherency and phase-differences (see e.g. Aguiar-Conraria, Martins and Soares, 2012), it is substantially extended to consider the parametric estimation provided by the partial gain.

Figure 5 summarizes our results. To facilitate the presentation, we give partial phase-difference and gain diagrams displaying mean values corresponding to three frequency intervals, namely for cycles of period $1.5 \sim 4$ years (the short end of business cycles), cycles of period $4 \sim 8$ years (the bulk of business cycles fluctuations) and cycles of period $8 \sim 20$ years (capturing long run relations). For the partial phase-differences, which are measured on a circular scale, the mean is computed as a circular mean, which is the appropriate notion of mean in this case; see, e.g. Zar (1998, pp. 598-599). The mean gain in a given frequency band is obtained by computing the absolute value of the mean of the corresponding complex gains.[6]

[5] In what follows, since we always deal with wavelet based measures, for simplicity, we will avoid using the word wavelet and simply write multiple coherency for multiple wavelet coherency, partial gain for partial wavelet gain, etc.

[6] To assess significance of multiple and partial coherencies (the latter ones also coincide with significance of the

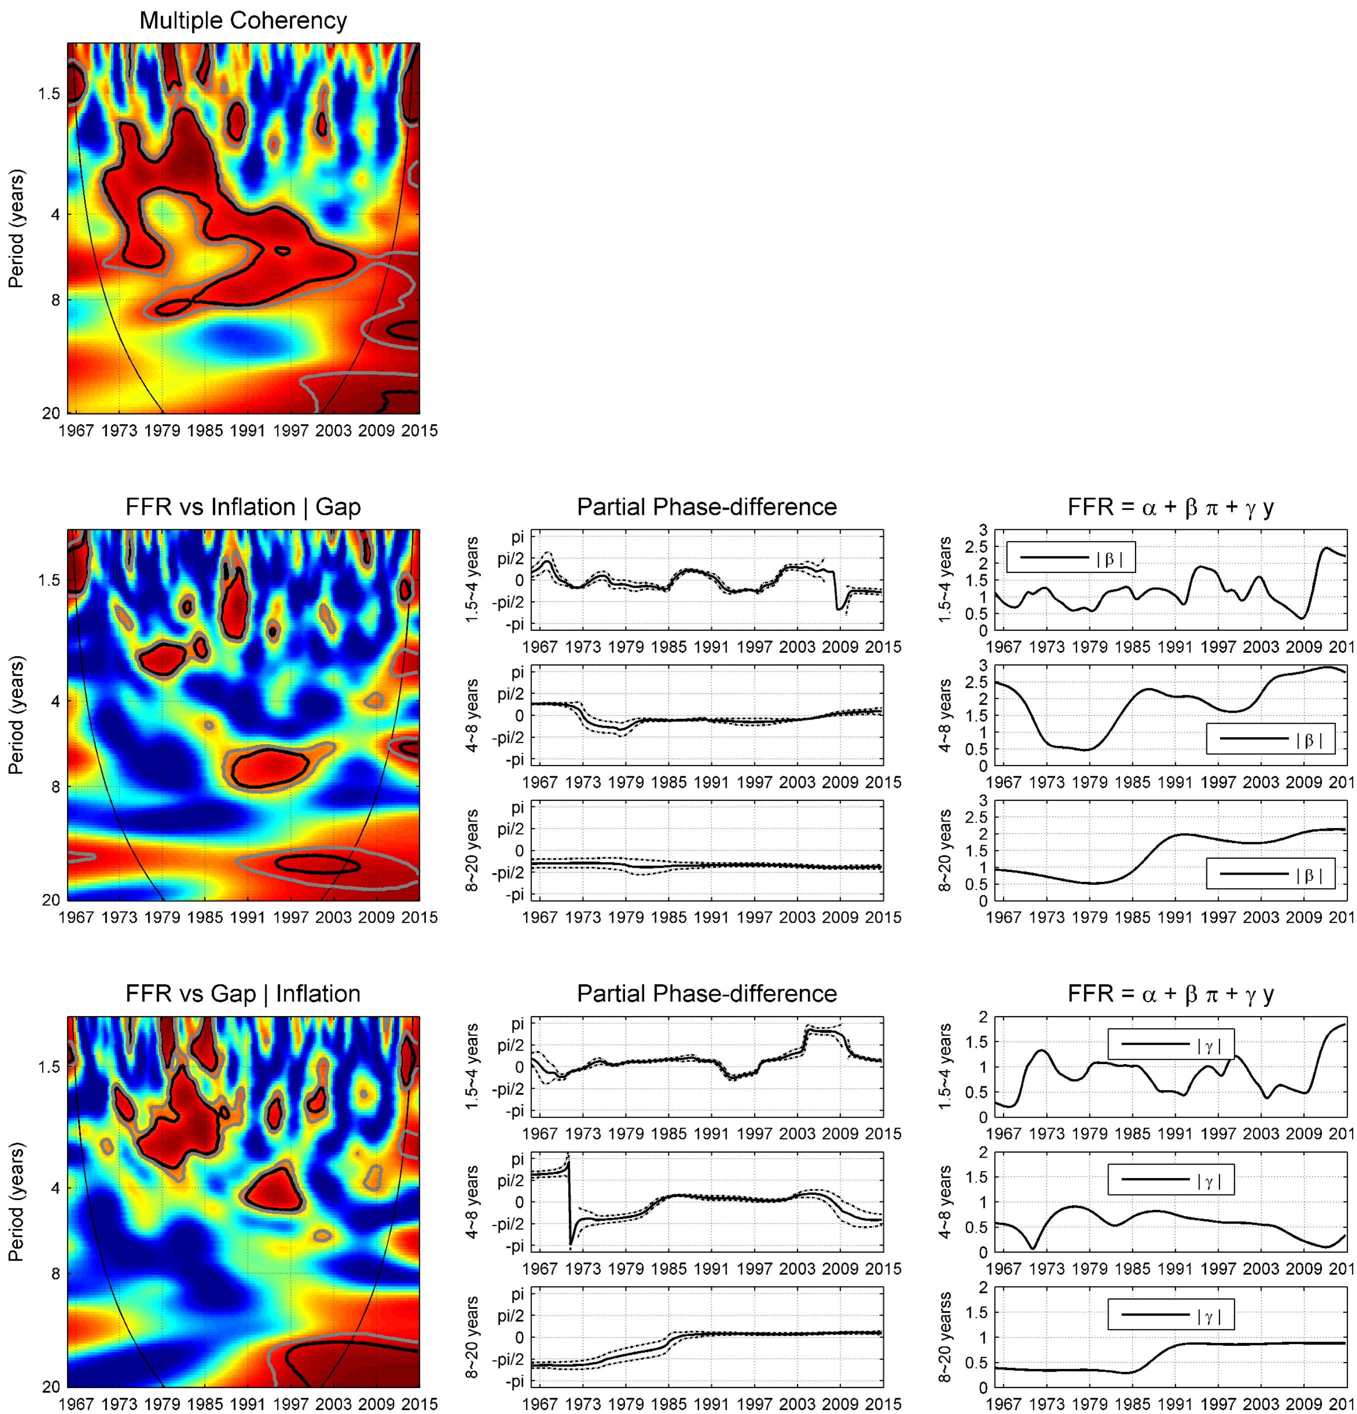

Figure 5: On the left – multiple wavelet coherency (top) and partial wavelet coherency between interest rate and inflation (middle) and between interest rate and the output gap (bottom). The black/gray contour designates the 5%/10% significance level. The color code for coherency ranges from blue (low coherency – close to zero) to red (high coherency – close to one). In the middle – partial phase-differences. On the right – partial wavelet gain.

partial gains — see the methodological appendix) we rely on Monte-Carlo simulations (with 5000 replications) after fitting an ARMA model. Confidence intervals for the circular mean at each point in time were also computed — we used the formulas proposed in Zar (1996), p. 604; see also Berens (2009) — and the interpretation of the mean phase at each point is done considering values as extreme as the two-end points of the corresponding interval. The limits of the confidence intervals for the mean phases are indicated in the pictures with black dashed-lines.

FFR and the Taylor Rule The multiple coherency indicates, for each time-frequency location, the proportion of the variation in the FFR that is jointly explained by the corresponding variations of inflation and the output gap. Hence, it measures the overall fit of the TR in the time-frequency domain: regions with a significant multiple coherency mean that inflation and the output gap are jointly significant explanatory variables of the FFR at those time-frequency locations. The first chart of Figure 5 confirms that the TR is overall a very good model for the FFR, as shown by the prevalence of regions depicted in red and yellow, as well as by the large regions within the gray and dark contours of statistical significance. The further time-frequency details given by the multiple coherency suggest that the overall fit of the TR has gradually shifted towards cycles of longer length. At higher frequencies ($1.5 \sim 4$ year cycles) it is high during the 1970s and 1980s, but hardly after 1991. At typical business cycles frequencies ($4 \sim 8$ years) it is strong and significant specially between 1985 and 2005. At longer cycles ($8 \sim 20$ years) multiple coherency starts increasing in the 1990s and becomes statistically significant since the beginning of the 2000s. Such pattern is consistent with the high intensity of shocks and resulting macro and policy volatility of the 1970s, the change to a more systematic monetary policy regime and the moderation of macroeconomic volatility since 1985, the gradual conquer of credibility during the 1990s, and then the persistence of the slowdown and of the policy reaction since the financial crisis of the late 2000s.

The multiple coherency is of assistance in the interpretation of the results given by the partial coherencies, especially when the explanatory variables are highly related, as is the case in the TR. Our partial coherencies — to be analyzed in the next sub-sections — capture the co-movement between each explanatory variable (inflation and output) and the FFR, filtering out the effect of the other. Yet, there is typically a strong co-movement between inflation and the output gap, the Phillips Curve — indeed, the predictive power of the gap over inflation is often invoked to motivate its inclusion in the TR. In such circumstances, while the overall significance of the model is high, the significance of individual co-movements for both explanatory variables may appear mistakenly low. It is therefore important that the partial coherencies are interpreted together with the multiple coherency. A notable example is the time-frequency region between 1973 and 1980 for frequencies of $4 \sim 8$ years: while both partial coherencies are mostly blue, the multiple coherency is mostly red and statistically significant; hence, in spite of the apparent lack of statistical significance of the partial coherencies, we are able to interpret the evolution of the coefficients on inflation and on the output gap in that time-frequency region.

FFR and inflation The partial coherency between the FFR and inflation exhibits different patterns across our three ranges of frequency-bands. At the short-run frequencies (period of $1.5 \sim 4$ years) the partial coherency is strong and significant from the second half of the 1970s until the beginning of the 1990s. At the typical business cycles frequencies ($4 \sim 8$ years), the coherency is strong and significant between 1985 and the end of the rules-based era, 2003. At the lower frequencies (period $8 \sim 20$ years) the coherency is consistently strong throughout the whole period, but is only significant after the beginning of the 1990s.

When the coherencies are significant, the phase-differences for both the $4 \sim 8$ years and the $8 \sim 20$ years frequency bands are stable and consistently located in the interval $(-\pi/2, 0)$, indicating a positive co-movement — as expected in the TR — with inflation leading the FFR. Note, however, that at the $8 \sim 20$ frequencies the lag between the interest rates and inflation is larger (phase difference closer to $-\pi/2$), which suggests that U.S. monetary policy has reacted more timely to changes in inflation at business cycle frequencies than at longer cycles. The phase-differences vary more in the frequency band of $1.5 \sim 4$ years, but, when the coherency is significant, they overall indicate a positive co-movement, with the FFR lagging inflation until 1986 and leading inflation in 1987-91.

We now focus on the time-frequency partial gain from FFR to inflation, displayed in the upper three charts of the right-hand-side of Figure 5. We have seen above that the full sample OLS estimate of the slope on inflation in the TR is essentially Taylor's baseline value of 1.5. Looking at the time-frequency estimates, we now see that they exhibit considerable variation around that value, with important differences across our frequency bands, which indicates that the TR implicit in U.S. monetary policy has changed along both dimensions — time and frequency.

The most interesting result, which is common to all frequency bands, is that the gain is below 1.0 — violating the Taylor principle — between around early 1970s and early 1980s. At the short-run frequencies, when significant (from the second half of the 1970s until the beginning of the

1990s), the gain fluctuates between 0.5 and 1.5. It falls below the Taylor principle threshold of 1.0 between 1974 and 1981, and then fluctuates within the range of 1.0 to 1.5 until the beginning of the 1990s. It is interesting to note that since the rules-based era of 1985-2003, the inflation coefficient is sistematically above the baseline vale of 1.5 both at the core of the business cycle frequencies ($4 \sim 8$ years) and at long run frequencies ($8 \sim 20$ years). To be more precise, at the core business cycle frequencies, the gain increases from 0.5 in 1979 to 2.5 in 1987. After that — when statistically significant — remains at values between 1.5 and 2.5. At the long run frequencies, it fluctuates around 2, starting in 1990 and throughout until the end of the sample.

Our results thus add important information to the studies of U.S. monetary policy that, following Clarida et al. (2000), document that U.S. policy reacted more to inflation after 1979, showing that the timing and size of the change in reaction differs across cyclical frequencies.

FFR and the output gap The partial coherency between the FFR and output exhibits different patterns across the frequency-bands, which to some extent resemble the patterns of the multiple coherency and of the partial FFR-inflation coherency — a gradual shift of co-movement towards cycles of longer length. Yet, two differences are noteworthy. First, coherency at short-run frequencies (period of $1.5 \sim 4$ years) are much more pervasive and indeed significant from the early 1970s until the beginning of the 2000s, and then in a further episode in 2011-2014; second, coherency is much more limited at the typical business cycles frequencies ($4 \sim 8$ years), at which it is only significant during the 1990s. At the lower frequencies (period $8 \sim 20$ years) the coherency is strong only since the beginning of the 1990s and significant merely since the mid-1990s.

When coherencies are significant, the phase-differences consistently indicate a positive co-movement between the FFR and the output gap, as expected, with the FFR slightly leading output. The only major exception occurs in the early 1990s'downturn – specifically between 1993 and 1998 – for short cycles ($1.5 \sim 4$ years), when phase-differences are located in the $(-\pi/2, 0)$ and are statistically different from 0, meaning that the output gap was leading the FFR. Overall (with the exception of 1993-1998 at short cycles), the partial coherencies and phase-differences indicate that U.S. monetary policy has attempted to preemptively stabilize the output gap for most of the time span and frequency bands at which the FFR and the output gap have co-moved significantly;

the positive co-movement in the data is consistent with such a forward-looking policy approach, given the lags in the transmission of interest rates policy to real activity (and inflation).

To obtain quantified results about the coefficient of output in the U.S. TR, we now assess the time-frequency partial gain from the FFR to the output gap, displayed in the three charts on the bottom right-hand-side of Figure 5.

At the short-run frequencies (period of $1.5 \sim 4$ years), during the extended period in which it is consistently significant (from the early 1970s until the beginning of the 2000s), the gain is close to 1.0 most of the time, with the only notable exception in 1989-1993 when it falls to the baseline value of 0.5. After 2009 – including the 2011-2014 period of significant gain – when policy was highly expansionary to fight the Great Recession, the gain increases markedly and reaches values close to 2.0 by the end of the sample

At the business cycles frequencies ($4 \sim 8$ years), when significant (1990s) the gain starts with a value close to 1.0 in 1990 and gradually falls until 1997, then maintaining the baseline value of 0.5 until the end of the 1990s. Previously, since the early 1970s – when the partial coherency is not significant but the multiple coherency is – the gain fluctuates between the full sample estimate of 0.5 and the value of 1.0 that it features at the beginning of the 1990s.

At frequencies corresponding to the $8 \sim 20$ years period, when statistically significant – i.e. after the mid-1990s – the gain is consistently close to 1.0, a level that it maintains until the end of the sample period. Previously, when coherency started increasing although not being statistically significant – in the second half of the 1980s – the gain increased from a value somewhat below 0.5 and reached the level of 1.0 at the beginning of the 1990s.

Our estimates for the gains of the output gap across frequencies hence suggest that the full sample OLS estimate of 0.5 is an artifact resulting from different coefficients across frequencies and along time. Before the beginning of the 1990s, values of the gain below 0.5 at the $8 \sim 20$ cycles are offset by values mostly above 0.5 at the $4 \sim 8$ years and $1.5 \sim 4$ years cycles; after 1991, values of the gain close to 1.0 at the $8 \sim 20$ cycles are compensated by values gradually smaller and close to 0.5 after 1997 at the $4 \sim 8$ cycles, and after 2003 at the $1.5 \sim 4$ years cycles. In the brief episode of the early-1990s in which the gain is above 0.5 at both the $4 \sim 8$ and $8 \sim 20$ years frequency bands, it is particularly low at the $1.5 \sim 4$ years cycles – actually around 0.5 between

1989 and 1993.

Our estimates for the gains in the latter part of the sample show that the finding elsewhere in the literature and statements by policy-makers pointing to a coefficient of 1.0 on the output gap in the U.S. TR since the Great Recession is not evenly explained across frequencies. After 2009, the estimate for the gain at the $8 \sim 20$ years cycles band is consistently 1.0, and the gain at the $1.5 \sim 4$ years frequency band sharply increases from 0.5 to more than 1.5, while the gain at the $4 \sim 8$ years cycles is close to 0.0 and not significant. With time-series tools, Nikolsko-Rzhevskyy, Papell and Prodan (2014) found that since 2007 U.S. monetary policy follows a modified Taylor Rule with a coefficient of 1.0 on the output gap; a coefficient twice as large as that in the original TR is consistent with the preferences for a balanced approach to stabilize output and prices stated by Federal Reserve Governors during the Great Recession — see Bernanke (2011) and Yellen (2012); moreover, it is consistent with negative policy interest rates since 2009, in line with the estimated shadow FFR for that period — which the original TR with a 0.5 coefficient on the output gap is not. Our framework shows that the prevalence of a modified TR with a slope of 1.0 on the output gap since 2009 is associated to policy actions focusing on the short-end ($1.5 \sim 4$ years) and long-end ($8 \sim 20$ years) of cyclical oscillations, and not on the most standard business cycle frequencies ($4 \sim 8$ years).

5 Conclusions

In this paper we assessed U.S. monetary policy in 1965:IV-2014:IV across time and frequencies in the framework of the Taylor Rule (TR). While variation in the TR coefficients along time has already been the subject of a vast literature, there was no study yet of variations of the TR coefficients simultaneously in the continuous time and frequency domains. Yet, there are compelling reasons to expect changes in the TR coefficients that differ across cyclical frequencies, and so documenting those changes enhances the understanding of U.S. monetary policy under the lens of the TR.

Following the most common and adequate practice in the literature, we use real-time data (on inflation and the output gap) available to policymakers when policy decisions were made. As regards the policy interest rate, we pursue a recently proposed approach and, in 2009-2014, replace the effective fed funds rate (FFR) with a shadow FFR able to capture the negative interest rates implied by recent unconventional quantitative monetary policy.

We use a set of continuous wavelet tools — the wavelet coherency, phase-difference, and gain — that allow for assessing the intensity, sign and synchronization (or lead/lag) of the co-movement between our time-series, as well as for estimating the respective regression coefficient in the time-frequency domain and providing statistical inference for all these measures. In particular, we employ partial wavelet tools to describe the co-movements along time and across frequencies between the policy interest rate and inflation (controlling for the output gap) and between the policy interest rate and the output gap (controlling for inflation). Methodologically, our main contribution to the literature is to provide a multi-variable generalization of the wavelet gain that allows for estimating multivariate functions in the time-frequency domain.

Regarding results, we provide a set of stylized facts on the U.S. TR in the last five decades that would not have been possible to detect with pure time- or frequency-domain tools, nor with the time-frequency domain tools available thus far. In particular, we provide estimates of the TR slopes that are allowed to vary both across frequencies and along time.

Regarding the relation between inflation and monetary policy, we uncovered a number of results that may be summarized as follows. First, within the framework of the TR, we document a gradual shift of the co-movement between the FFR and inflation towards cycles of longer length, along the last five decades of U.S. monetary policy. Second, we confirm that the co-movement between the FFR and inflation has been positive at all frequencies, but find synchronization for higher frequencies, while for lower the FFR has lagged inflation. Third, the inflation coefficient in the U.S. TR has changed much more markedly for cycles of intermediate duration ($4 \sim 8$ years) than for longer cycles ($8 \sim 20$) and shorter cycles likewise ($1.5 \sim 4$ years). Fourth, rather than a change from a constant coefficient below 1.0 before 1979 to a coefficient above 1.0 after 1979, there was a gradual decrease of the inflation coefficient until the mid-1970s, followed by an increase after 1979 that is essentially completed at the start of the 1985-2003 rules-based era, and is most marked at the $4 \sim 8$ year cycles. Fifth, our estimates suggest that the Taylor principle has been violated for longer at frequencies that correspond to cycles of longer period, although statistical uncertainty is higher the lower are frequencies, notably before 1990. Finally, we show since the mid-1980s — 1984 for cycles of period $4 \sim 8$ years and 1987 for cycles of period $8 \sim 20$ years — the coefficient on inflation in the TR is consistently above the baseline value and full sample estimate of 1.5.

As regards the output gap, we emphasize the following findings. First, within the framework of the TR, we document a gradual shift of the co-movement between the FFR and the output gap towards cycles of longer length, along the last five decades of U.S. monetary policy; with this respect, the co-movement between the FFR and output has been stronger at shorter oscillations ($1.5 \sim 4$ years frequency band) and weaker at business cycles frequencies ($4 \sim 8$ years) than the co-movement with inflation. Second, we confirm that the co-movement between the FFR and output has been positive at all frequencies, with the policy rate leading output for most of the time and frequencies, consistently with an anti-cyclical stance and with the lags of policy impact. Third, we document that the full sample OLS estimate of 0.5 is an artifact resulting from different coefficients across frequencies and along time: before the 1990s, estimates below 0.5 at the $8 \sim 20$ years cycles are offset by estimates mostly above 0.5 at the $4 \sim 8$ years and $1.5 \sim 4$ years cycles; after 1991, estimates close to 1.0 at the $8 \sim 20$ cycles are compensated by estimates gradually smaller and close to 0.5, after 1997 at the $4 \sim 8$ cycles and after 2003 at the $1.5 \sim 4$ years cycles. Fourth, we show that the time-series evidence and policy-makers' statements pointing to a modified TR with a slope of 1.0 on the output gap in the U.S. TR since the Great Recession is not evenly explained across frequencies, being associated to stronger reactions of policy to output at the short-end ($1.5 \sim 4$ years) and long-end ($8 \sim 20$ years) of cyclical oscillations, but not at the most standard business cycle frequencies ($4 \sim 8$ years).

Acknowledgments

Research at NIPE has been carried out within the funding with COMPETE reference number POCI-01-0145-FEDER-006683, with the FCT/MEC's (Fundação para a Ciência e a Tecnologia, I.P.) financial support through national funding and by the ERDF through the Operational Programme on "Competitiveness and Internationalization – COMPETE 2020" under the PT2020 Partnership Agreement.

Research at Cef.up has been financed by Portuguese Public Funds through FCT (Fundação para a Ciência e a Tecnologia) and by the European Regional Development Fund through COMPETE 2020 – Programa Operacional Competitividade e Internacionalização – in the framework of the project UID/ECO/04105/2013.

References

[1] Aguiar-Conraria, Luís, Manuel M. F. Martins, and Maria J. Soares (2012) "The Yield Curve and the Macro-Economy Across Time and Frequencies" *Journal of Economic Dynamics and Control*, 36, 1950–1970.

[2] Aguiar-Conraria, Luís, and Maria J. Soares (2014) "The Continuous Wavelet Transform: Moving beyond Uni-and Bivariate Analysis" *Journal of Economic Surveys*, 28, 344–375.

[3] Alcidi, Cinzia, Alessandro Flamini, and Andrea Fracasso (2011) "Policy Regime Changes, Judgment and Taylor Rules in the Greenspan Era" *Economica*, 78, 89–107.

[4] Assenmacher-Wesche, Katrin (2006). "Estimating Central Banks' Preferences from a Time-varying Empirical Reaction Function." *European Economic Review*, 50, 1951–1974.

[5] Benati, Luca (2008) "Investigating Inflation Persistence Across Monetary Regimes." *The Quarterly Journal of Economics*, 123, 1005–1060.

[6] Bernanke, Ben S. (2011) "Witness before the Committee on Banking, Housing, and Urban Affairs of the U. S. Senate", March 1, available at http://www.gpo.gov/fdsys/pkg/CHRG-112shrg65824/pdf/CHRG-112shrg65824.pdf.

[7] Bernanke, Ben S. (2015) "The Taylor Rule: A benchmark for monetary policy?", Ben Bernanke' s Blog, April 28, 2015 11:00am, available at http://www.brookings.edu/blogs/ben-bernanke/posts/2015/04/28-taylor-rule-monetary-policy.

[8] Berens, Philipp (2009) "CircStat: A MATLAB Toolbox for Circular Statistics." *Journal of Statistical Software*, 31 (10), 1–21.

[9] Bunzel, Helle, and Walter Enders (2010) "The Taylor Rule and "Opportunistic" Monetary Policy." *Journal of Money, Credit and Banking*, 42, 931–949.

[10] Canova, Fabio (2009) "What Explains the Great Moderation in the U.S.? A Structural Analysis." *Journal of the European Economic Association*, 7, 697–721.

[11] Christensen, Anders M., and Heino B. Nielsen (2009) "Monetary Policy in the Greenspan Era: A Time Series Analysis of Rules vs. Discretion." *Oxford Bulletin of Economics and Statistics*, 7(1), 69–89.

[12] Clarida, Richard, Jordi Galí, and Mark Gertler (2000) "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics*, 115, 147–80.

[13] Croushore, Dean and Tom Stark (2011) "A real-time data set for macroeconomists." *Journal of Econometrics*, 105 (1), 111–130.

[14] Dennis, Richard (2006) "The Policy Preferences of the US Federal Reserve." *Journal of Applied Econometrics*, 21, 55–77.

[15] Dolado, Juan J., Ramón Maria-Dolores, and Manuel Naveira (2005) "Are Monetary Policy Reaction Functions Asymmetric? The Role of Nonlinearity in the Phillips Curve." *European Economic Review*, 49, 485–503.

[16] Engle, Robert (1976) "Interpreting spectral analyses in terms of time-domain models." *Annals of Economics and Social Measurement*, 5(1), 89–109.

[17] Favero, Carlo A., and Riccardo Rovelli (2003) "Macroeconomic Stability and the Preferences of the Fed: A Formal Analysis, 1961-98." *Journal of Money, Credit and Banking*, 35, 545–56.

[18] Hamilton, James D. (2001). "A Parametric Approach to Flexible Nonlinear Inference." *Econometrica*, 69, 537–573.

[19] Hamilton, James D., Seth Pruitt, and Scott Borger (2011) "Estimating the Market-Perceived Monetary Policy Rule." *American Economic Journal: Macroeconomics*, 3, 1–28.

[20] Kahn, Geroge A. (2012) "The Taylor Rule and the Practice of Central Banking," in *The Taylor Rule and the Transformation of Monetary Policy*, edited by Evan F. Koenig, Robert Leeson, and George A. Kahn, Chapter 3, Hoover Institution, Stanford University.

[21] Kim, Dong H., Denise Osborn, and Marianne Sensier (2005) "Nonlinearity in the Fed's Monetary Policy Rule." *Journal of Applied Econometrics*, 20, 621–639.

[22] Lubik, Thomas A. and Frank Schorfheide (2007) "Do Central Banks Respond to Exchange Rate Movements? A Structural Investigation." *Journal of Monetary Economics*, 54, 1069–1087.

[23] Mandler, Martin, and Michael Scharnagl (2014) "Money growth and consumer price inflation in the euro area: a wavelet analysis." Discussion Paper, Deutsche Bundesbank No 33/2014.

[24] Mehra, Y. P., and Sawhney, B. (2010) "Inflation measure, Taylor rules, and the Greenspan-Bernanke years." *Federal Reserve Bank of Richmond Economic Quarterly*, 96 (2), 123–151.

[25] Nikolsko-Rzhevskyy, Alex, David H. Papell, and Ruxandra Prodan (2014). "Deviations from rules-based policy and their effects." Journal of Economic Dynamics and Control , 49, 4-17.

[26] Orphanides, Athanasios (2001) "Monetary Policy Rules Based on Real-time Data." *American Economic Review*, 91, 964–985.

[27] Owyang, Michael T., and Garey Ramey (2004) "Regime Switching and Monetary Policy Measurement." *Journal of Monetary Economics*, 51, 1577–1597.

[28] Sack, Brian, and Roberto Rigobon (2003) "Measuring the Reaction of Monetary Policy to the Stock Market." *Quarterly Journal of Economics*, 118, 639–669.

[29] Sims, Eric (2013) "What Measure of Output Should be Targeted in Interest Rate Rules?" Mimeo, March 7, 2013. Available at http://brassmonkey.tamu.edu/common/files/papers/Sims-Eric.pdf

[30] Surico, Paolo (2007) "The Fed's Monetary Policy Rule and US Inflation: The Case of Asymmetric Preferences." *Journal of Economic Dynamics and Control*, 31, 305–324.

[31] Taylor, John B. (1993) "Discretion Versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

[32] Taylor, John B. (2011) "Macroeconomic Lessons from the Great Deviation." In *NBER Macroeconomics Annual 2010*, edited by Daron Acemoglu and Michael Woodford, pp. 387–395. University of Chicago Press.

[33] Taylor, John B. (2012) "Monetary Policy Rules Work and Discretion Doesn't: A Tale of Two Eras." *Journal of Money, Credit and Banking*, 44, 1017–1032.

[34] Taylor, John B., and John C. Williams. (2010) "Simple and Robust Rules for Monetary Policy." In *Handbook of Monetary Economics*, Volume 3, edited by Benjamin M. Friedman and Michael Woodford, pp. 829–859. Elsevier.

[35] Torrence, Cristopher and Compo, Gilbert P. (1998). "A practical guide to wavelet analysis." *Bulletin of the American Meteorological Society* 79: 61–78.

[36] Tillmann, Peter (2011) "Parameter Uncertainty and Nonlinear Monetary Policy Rules." *Macroeconomic Dynamics*, 15, 184–200.

[37] Trecroci, Carmine, and Matilde Vassalli (2010) "Monetary Policy Regime Shifts: New Evidence from Time-Varying Interest Rate Rules." *Economic Inquiry*, 48,933–950.

[38] Wolters, Maik H. (2012) "Estimating Monetary Policy Reaction Functions Using Quantile Regressions." *Journal of Macroeconomics*, 34, 342–361.

[39] Wu, Jing Cynthia, and Fan Dora Xia (2014) "Measuring the macroeconomic impact of monetary policy at the zero lower bound." NBER working paper No. w20117, May.

[40] Yellen, Janet (2012) "The Economic Outlook and Monetary Policy", Remarks at the Money Marketeers of New York University, April 11, available at http://www.federalreserve.gov/newsevents/speech/yellen20120411a.pdf.

[41] Yu, Xiangrong. (2013) "Measurement Error and Policy Evaluation in the Frequency Domain." *Journal of Macroeconomics*, 38, 307–329.

[42] Zar, Jerrold H. (1996) Biostatistical Analysis (3rd Ed.), Prentice-Hall. Upper Saddle River, New Jersey.

A Methodological Appendix

The continuous wavelet transform is an increasingly popular tool in econometric analysis. The most common argument to justify its use is the possibility of tracing transitional changes across time and frequencies — see Aguiar-Conraria and Soares (2014) for a review. So far, the analysis in the time-frequency domain with the continuous wavelet transform has been mostly limited to the use of the wavelet power spectrum, the wavelet coherency and the wavelet phase-difference. Aguiar-Conraria and Soares (2014) already extended these tools to allow for multivariate analyses. These multivariate tools are sufficient to assess the strength of the relation between several variables, but they are insufficient to estimate the magnitude of the relation. Just like (partial) correlation coefficients do not provide the same information as the regression coefficients.

Mandler and Scharnagl (2014) use the concept of the wavelet gain as a regression coefficient in the regression of y on x. In this paper, and, to our knowledge, for the first time, we will estimate an equation relating more than two variables (just like a regression of y on x and z) in the time-frequency domain. To do so, we generalize the concept of wavelet gain and define the *partial wavelet gain*, which can be interpreted as a regression coefficient in the regression of y on x after controlling for other variables. At the end of this appendix, we provide a example to illustrate the application of this tool..

A.1 The Continuous Wavelet Transform

For all practical uses, a wavelet $\psi(t)$ is a function that oscillates around the t-axis and looses strength as it moves away from the center, behaving like a small wave. The specific wavelet we use in this paper is the complex-valued function (selected from the so-called *Morlet wavelet* family) defined by $\psi(t) = \pi^{-\frac{1}{4}} e^{6\,i\,t} e^{-\frac{t^2}{2}}$. Given a time-series $x(t)$, its *continuous wavelet transform* (CWT), with respect to a given wavelet ψ, is the function of two variables, $W_x(\tau, s)$, given by

$$W_x(\tau, s) = \frac{1}{\sqrt{|s|}} \int_{-\infty}^{\infty} x(t) \overline{\psi}\left(\frac{t-\tau}{s}\right) dt. \tag{A.1}$$

A.2 Uni and bivariate tools

All the quantities we are going to introduce are functions of time (τ) and scale (s). To simplify the notation, we will describe these quantities for a specific value of the argument, (τ, s), which will be omitted in the formulas.

A.2.1 Wavelet power spectrum and the phase angle

In analogy with the terminology used in the Fourier case, the (local) *wavelet power (spectrum)* is defined as

$$(WPS)_x = |W_x|^2 \,. \tag{A.2}$$

This gives us a measure of the variance distribution of the time-series in the time-frequency plane.

When the wavelet ψ is complex-valued, as in our case, the wavelet transform W_x is also complex-valued. In this case, the transform can be expressed in polar form as $W_x = |W_x| \, e^{i\phi_x}$, $\phi_x \in (-\pi, \pi]$. The angle ϕ_x is known as the *(wavelet) phase*.

A.2.2 Cross wavelet tools

The *cross-wavelet transform* of two time-series $x(t)$ and $y(t)$, denoted by W_{xy} is defined as

$$W_{xy} = W_x \overline{W}_y \,, \tag{A.3}$$

where W_x and W_y are the wavelet transforms of x and y, respectively. The absolute value of the cross-wavelet transform, $|W_{xy}|$, will be referred to as the *cross-wavelet power*. The cross-wavelet power of two time-series depicts the covariance between two time-series at each time and frequency.

We define the *complex wavelet coherency* of x and y, ϱ_{xy}, by

$$\varrho_{xy} = \frac{S\left(W_{xy}\right)}{\left[S\left(|W_x|^2\right) S\left(|W_y|^2\right)\right]^{1/2}} \,, \tag{A.4}$$

where S denotes a smoothing operator in both time and scale. For notational simplicity, we will denote by S_{xy} the smoothed cross-wavelet transform of two series x and y and also use σ_x and σ_y to denote, respectively, $\sqrt{S(|W_x|)^2} = \sqrt{S_{xx}}$ and $\sqrt{S(|W_y|)^2} = \sqrt{S_{yy}}$. With these notations, we will simply write the formula for the complex coherency as

$$\varrho_{xy} = \frac{S_{xy}}{\sigma_x \sigma_y}. \tag{A.5}$$

The wavelet coherency, which we will denote by R_{xy}, is defined simply as the absolute value of the complex wavelet coherency, i.e. is given by

$$R_{xy} = |\varrho_{xy}|. \tag{A.6}$$

With a complex-valued wavelet, we can compute the phase of the wavelet transform of each series and, by computing their difference, we can then obtain information about the possible delays of the oscillations of the two series, as a function of time and frequency. It follows immediately from (A.3) that the *phase-difference*, which we will denote by ϕ_{xy}, can also be computed as the angle of the cross-wavelet transform. Another slightly different way to define the phase-difference makes use of the angle of the complex wavelet coherency, instead of the angle of the cross-wavelet transform; this definition, although not strictly coinciding with the difference between the individual phases, due to the smoothing, has the advantage of allowing a more direct generalization for the multivariate case.

Finally, we define the *complex wavelet gain of y over x*, denoted by $\mathcal{G}_{yx}$, by

$$\mathcal{G}_{yx} = \frac{S_{yx}}{S_{xx}} = \varrho_{yx} \frac{\sigma_y}{\sigma_x} \tag{A.7}$$

and, following Mandler and Scharnagl (2014), we call *wavelet gain*, which we denote by G_{yx}, to the modulus of the complex wavelet gain, i.e.

$$G_{yx} = \frac{|S_{yx}|}{S_{xx}} = R_{yx} \frac{\sigma_y}{\sigma_x}. \tag{A.8}$$

Recalling the interpretation of the Fourier gain as the modulus of the regression coefficient of y on x at a given frequency (see, e.g. Engle 1976), it is perfectly natural to interpret the wavelet gain as the modulus of the regression coefficient in the regression of y on x, at each time and frequency.

A.3 Multivariate wavelet analysis

Let p ($p > 2$) time-series $x_1, x_2, \ldots, x_p$ be given. We first introduce a set of notations.

We will denote by W_i the wavelet spectrum corresponding to the time-series x_i and by W_{ij} the cross-wavelet spectrum of the two series x_i and x_j. Just as in the case of ordinary wavelet coherency, to compute partial wavelet coherencies it is necessary to perform a smoothing operation on the cross-spectra. We will denote by S_{ij} the smoothed version of W_{ij}, i.e. $S_{ij} = S\left(W_{ij}\right)$, where S is a certain smoothing operator. We will use $\mathscr{S}$ to denote the $p \times p$ matrix of all the smoothed cross-wavelet spectra S_{ij}, i.e. $\mathscr{S} = (S_{ij})_{i,j=1}^{p}$. [7] For a given matrix A, A_i^j denotes the sub-matrix obtained by deleting its i-th row and j-th column and A_{ij}^d denotes the co-factor of the element in position (i, j) of A, i.e. $A_{ij}^d = (-1)^{(i+j)} \det A_i^j$. For completeness, we use the notation $A^d = \det A$. Finally, for a given integer j such that $2 \leq j \leq p$, we denote by q_j the set of all the indexes from 2 to p with the exception of j, i.e. $q_j = \{2, \ldots, p\} \setminus \{j\}$.

A.3.1 Multiple and partial wavelet coherency and partial phase-difference

The *squared multiple wavelet coherency* between the series x_1 and all the other series $x_2, \ldots, x_p$ will be denoted by $R^2_{1(23\ldots p)}$ and is given by the formula

$$R^2_{1(23\ldots p)} = 1 - \frac{\mathscr{S}^d}{S_{11}\mathscr{S}_{11}^d}. \tag{A.9}$$

The *complex partial wavelet coherency* of x_1 and x_j ($2 \leq j \leq p$) allowing for all the other series will be denoted by $\varrho_{1\,j.q_j}$ and is given by

$$\varrho_{1\,j.q_j} = -\frac{\mathscr{S}_{j1}^d}{\sqrt{\mathscr{S}_{11}^d}\sqrt{\mathscr{S}_{jj}^d}}. \tag{A.10}$$

The *partial wavelet coherency* of x_1 and x_j allowing for all the other series, denoted by $R_{1\,j.q_j}$, is defined as the absolute value of the above quantity, i.e. $R_{1\,j.q_j} = \frac{|\mathscr{S}_{j1}^d|}{\sqrt{\mathscr{S}_{11}^d}\sqrt{\mathscr{S}_{jj}^d}}$, and the *squared partial wavelet coherency* of x_1 and x_j allowing for all the other series, is simply the square of

[7] To be more correct, $\mathscr{S}$ depends on the specific value (τ, s) at which the spectra are being computed, i.e. there is one such matrix for each (τ, s).

$R_{1\,j.q_j}$. Having defined the complex partial wavelet coherency $\varrho_{1\,j.q_j}$ of series x_1 and x_j controlling for all the other series, we simply define the *partial phase-difference* of x_1 and x_j given for all the other series, denoted by $\phi_{1\,j.q_j}$, as the angle of $\varrho_{1\,j.q_j}$.

A.3.2 Partial wavelet gain

We define the *complex partial wavelet gain* of series x_1 over series x_j allowing for all the other series, denoted by $\mathcal{G}_{1\,j.q_j}$, by the formula

$$\mathcal{G}_{1\,j.q_j} = \frac{\mathscr{S}^d_{j1}}{\mathscr{S}^d_{11}} \tag{A.11}$$

and the *partial wavelet gain*, denoted by $G_{1\,j.q_j}$, as the modulus of the above quantity, i.e.

$$G_{1\,j.q_j} = \frac{|\mathscr{S}^d_{j1}|}{\mathscr{S}^d_{11}}. \tag{A.12}$$

Naturally, the partial wavelet gain can also be computed using the partial wavelet coherency, as

$$G_{1\,j.q_j} = R_{1\,j.q_j} \frac{\sqrt{\mathscr{S}^d_{jj}}}{\sqrt{\mathscr{S}^d_{11}}}. \tag{A.13}$$

For $j = 2, \ldots, p$, the values $G_{1,j.q_j}$ can be interpreted as the coefficients (in modulus) in the multiple linear regression of x_1 in the explanatory variables $x_2, \ldots, x_p$, at each time and frequency.

A.3.3 Formulas in terms of coherencies

The above formulas for the partial wavelet coherency and for the partial wavelet gain were given in terms of the smoothed spectra S_{ij}. We can also define these quantities in terms of simple complex coherencies (i.e. wavelet complex coherencies between pairs of series).

Corresponding to the matrix $\mathscr{S}$, we now consider the matrix $\mathscr{C} = (\varrho_{ij})_{i,j=1}^{p}$ of all the complex wavelet coherencies ϱ_{ij}. Then, we can define the multiple wavelet coherencies by the following alternative formula

$$R^2_{1(23\ldots p)} = 1 - \frac{\mathscr{C}^d}{\mathscr{C}^d_{11}}, \tag{A.14}$$

the complex partial wavelet coherency by

$$\varrho_{1\,j.q_j} = -\frac{\mathscr{C}_{j1}^d}{\sqrt{\mathscr{C}_{11}^d}\sqrt{\mathscr{C}_{jj}^d}}, \tag{A.15}$$

and the partial wavelet gain by

$$G_{1\,j.q_j} = \frac{|\mathscr{C}_{j1}^d|}{\mathscr{C}_{11}^d}\frac{\sigma_1}{\sigma_j}. \tag{A.15}$$

The proof of the above results is a simple application of the properties of determinants; see Aguiar-Conraria and Soares (2014) for details concerning the multiple and partial coherencies.

A.4 Statistical Significance

To test significance of the wavelet power spectrum, one can rely on the results of Torrence and Compo (1998), which shows that the local wavelet power spectrum of a white noise or an AR(1) process, normalized by the variance of the time series, is very well approximated by a chi-squared distribution. Testing the wavelet power spectrum against a flat spectrum (white noise) is a good starting point. If one wants to consider more complicated null hypotheses, rather a white or red noise, one usually relies on Monte-Carlo simulations.

To test significance of coherency and partial coherency there are no good theoretical results. The ones that exist impose too stringent restrictions. Therefore, one usually relies on Monte-Carlo simulations. In our case, we fit an ARMA model to each of the series and construct new samples by drawing errors from a Gaussian distribution with a variance equal to that of the estimated error terms; for each set of time-series we perform the exercise several times, and then extract the critical values.

By comparing the formulas of the (partial) gain with the (partial) coherency, for example, comparing formula (A.6) with formula (A.8), it should be apparent that if one has the value zero, so does the other. Therefore, when we test the null hypothesis that the (partial) coherency is zero, we are simultaneously testing the null huypothesis that the (partial) gain is zero. Our Monte-Carlo simulations confirm this assertion.