IIMB-WP N0. 681/2023

WORKING PAPER NO: 681

Wage Cyclicality Across Time and Frequencies

Srinivasan Murali
Economics Area
Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore – 560 076
srinim@iimb.ac.in

Shweta Sogani
Doctoral Student
Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore – 560 076
shweta.sogani22@iimb.ac.in

Year of Publication – March 2023

Wage Cyclicality Across Time and Frequencies

Srinivasan Murali*   Shweta Sogani†

IIM Bangalore   IIM Bangalore

March 2023

Abstract

Employing wavelet analysis, this paper documents how the cyclicality of real wages has evolved over time and at different business cycle frequencies in US. We use individual-level data from the CPS to construct the composition bias corrected quarterly wage series using Haefke et al. (2013)'s methodology. Utilizing continuous wavelet tools, we find that the cyclicality of wages for all the workers as well as new hires has increased over time. Additionally, we find that the increase in cyclicality is prevalent across all frequencies. This decline in wage rigidity over business cycles relates to the broader structural changes in the labour market and also has implications for labour search framework.

JEL codes: C49, E24, E32
Keywords: Wavelets, Wage cyclicality, Wage rigidity, Time-frequency analysis

*Indian Institute of Management Bangalore, Bengaluru, India; (Email: srinim@iimb.ac.in)

†Indian Institute of Management Bangalore, Bengaluru, India; (Email: shweta.sogani22@iimb.ac.in)

1. Introduction

This paper uses continuous wavelet tools to estimate the cyclicality of real wages across time and frequencies in US. Macroeconomists have always been interested in understanding the behaviour of real wages over the business cycle. A number of early empirical studies like Dunlop (1938), Tarshis (1939), Bodkin (1969), Neftci (1978), Sargent (1978), Geary and Kennan (1982), and Sumner and Silver (1989) used aggregate wage data and found the real wages to be acyclical or mildly procyclical. Following the arguments of Stockman (1983), Bils (1985) and Solon et al. (1994) showed that measuring wage cyclicality using aggregate wage data is misleading, as the cyclicality estimates suffer from composition bias. Aggregate data, by its way of construction, gives more weight to low-skilled workers during expansions than during recessions, thus introducing a countercyclical bias in the wage cyclicality estimates. Following this, a large number of studies like Shin (1994), Devereux (2001), Devereux and Hart (2006), Hart (2006), Carneiro et al. (2012), Haefke et al. (2013) using longitudinal data to keep the composition of workers fixed over the business cycle, found that the real wages are substantially procyclical.

There has been a renewed interest in understanding the elasticity of wages over business cycle, as the literature of search and matching models has resorted to wage rigidity as one of the ways to generate empirically consistent unemployment fluctuations. Studies like Shimer (2005) and Costain and Reiter (2008) showed that a standard labour search model generates too low a volatility in both unemployment and vacancies compared to the data. Labeled as the unemployment volatility puzzle, Hall (2005) showed that introducing rigid wages in search models can help in increasing the volatilities thus making them empirically consistent. In addition, Shimer (2004) and Pissarides (2009) argued that the model's behaviour depends only on the wage rigidity of newly hired workers and not on the continuing workers. Following this, a number of studies like Menzio (2005), Farmer and Hollenhorst (2006), Blanchard and Galí (2007), Hall and Milgrom (2008), Gertler and Trigari (2009), Shimer (2010), Michaillat (2012), and Christiano et al. (2016) introduced some form of wage rigidity to make their models more consistent with the data.

In this paper, we employ continuous wavelet tools, namely, wavelet coherency, wavelet phase-difference, and wavelet gain to analyse how the cyclicality of wages has changed over time and at different business cycle frequencies. Wavelet coherency, the time-frequency analog of correlation, measures the magnitude of co-movement of wages with the business cycle at every time period and frequency, while the phase-difference gives us the direction of this relationship over time and frequency along with lead/lag of wages over business cycles. In order to estimate our primary object of interest, i.e., wage elasticity over time and frequency, we need a regression setup with both time-varying and frequency-varying regression coefficients. Wavelet gain provides us with such a framework, thus enabling us to estimate wage cyclicality across time and frequencies.

As pointed out by the earlier literature, it would be misleading to estimate wage cyclicality using aggregate data. Hence following Haefke et al. (2013), we use individual level data from Current Population Survey (CPS) - Outgoing Rotation Groups (ORG) over the period 1979 – 2019 to construct the quarterly wage series free of composition bias. Our initial regression results replicates the overall findings of the literature and Haefke et al. (2013) in particular, i.e., wages are procyclical with respect to aggregate labour productivity, and wages of new hires are more elastic to aggregate business cycle conditions than that of continuing workers.

By employing wavelet techniques, this paper provides new insights on wage cyclicality over time and across frequencies. We find that wages are more elastic in the lower frequencies compared to higher frequencies. This suggests that wages take time to adjust, and the procyclical relationship strengthens over longer time horizons. Importantly, we find that the wage cyclicality of all workers as well as new hires has increased over time across all the frequency intervals. This finding relates to broader structural changes observed in the labour market including declining unionization and bargaining power of workers. Additionally, this finding also has implications for the calibration of labour search and matching models.

The rest of the paper is organized as follows. Section 2 describes the dataset and presents the regression estimates of wage cyclicality. Section 3 explains the wavelet tools used in our analysis while section 4 discusses the time-frequency results of wage cyclicality. Section 5 talks about the implications of our findings and section 6 concludes.

2. Data

It is well known from the studies like Bils (1985) and Solon et al. (1994) that estimates of wage cyclicality obtained from aggregate data suffers from composition bias. During a recession, more low-skilled workers earning lower wages end up losing their jobs compared to high-skilled workers. Similarly, the opposite happens during an expansion. Thus, the composition of the workforce varies over the business cycle. This implies that, during a recession, the aggregate wages are constructed over the workforce with more high-skilled workers compared to an expansion, thus introducing a countercyclical bias in the estimates. In order to overcome this, we need to use individual level worker data to keep the composition fixed over the business cycle.

2.1 Individual Level Wage Data

Majority of studies in this literature use household level panel data to estimate the wage cyclicality, thus ensuring the composition of workers is kept fixed over time. A number of studies like Bils (1985) and Shin (1994) use individual wage data from the National Longitudinal Survey of Youth (NLSY) while other studies like Solon et al. (1994) and Devereux (2001) use Panel Study of Income Dynamics (PSID) to estimate wage cyclicality. Even though both NLSY and PSID keep the composition of workers fixed over the business cycle, they provide information on individual wages only at a yearly frequency. Since we use wavelets in our analysis and wavelets are quite data demanding, it would be better if we have data at a higher frequency.

Hence, following Haefke et al. (2013), we use Current Population Survey (CPS) - Outgoing Rotation Groups (ORG) to obtain the individual level wage data. The major advantage of using CPS is, we will be able to construct wage series at a quarterly frequency, thus more suitable for the wavelet analysis. CPS-ORG is a monthly survey of US households that has been administered since 1979. Haefke et al. (2013) uses the CPS microdata to construct quarterly wage series for the period of 1979-2006. We closely follow the methodology of Haefke et al. (2013) in constructing the quarterly wage data. In addition, we extend the original data till 2019.[1]

The sample consists of both male and female non-supervisory workers in the private non-farm business sector, who are aged between 25 and 60. We measure wages as hourly earnings, obtained by dividing earnings by the usual hours worked. Further, the wages are deflated using the implicit deflator for private non-farm business sector. The panel structure of the CPS data is exploited to identify if the workers were newly hired. We match the workers with the preceding three monthly data files to identify new hires as those who were not working for at least one of the preceding three months.[2,3] In addition, the data also contains information on the demographic characteristics of the workers, and the industry and occupation of their work. Finally, the data on aggregate labour productivity, measured as the output per hour in the non-farm business sector, is obtained from the BLS productivity and cost program.

2.2 Constructing the Wage Series

A number of papers in the literature measuring wage cyclicality including Bils (1985), Solon et al. (1994), and Devereux (2001) show that, controlling for the composition of workers is critical to get an unbiased estimate of cyclicality. Hence, we need to control for the observed and unobserved characteristics of individual workers over the business cycle. Let w_{it} be the wage earned by worker i at time t. Then, following Haefke et al. (2013), the individual wages can be modeled as

[1]We thank Thijs van Rens for providing us with some of the missing data files needed for generating our estimates.

[2]There is a discontinuity in matching in the third and fourth quarters of 1985 and 1995 due to changes in sample design. Therefore, we have missing values for those quarters.

[3]There is a possibility of misreporting of employment status in our data. When an employed worker reports being unemployed at any point in the survey, that worker would be identified as a new hire, resulting in our estimates of wage cyclicality of new hires being biased towards zero. Since we find, in line with the literature, that the wages of new hires are more procyclical, this bias will only work against our result.

$$\log w_{it} = x_i'\beta + \log \hat{w_{it}}, \tag{1}$$

where x_i is the vector of individual characteristics – education, gender, marital status, race, and a fourth order polynomial in experience, while $\hat{w_{it}}$ captures the residual wage controlling for these factors. Even though these variables capture observed individual heterogeneity, they don't account for the individual fixed effects. Hence, majority of the papers in this literature take first difference of the wages to drop the fixed effects. However, doing this in our case will drop the wages of all the new hires from our analysis. Hence, just like Haefke et al. (2013), we work with wage levels controlling only for the observable factors, and not explicitly controlling for individual specific fixed effects. Haefke et al. (2013) shows that, controlling just for the observable factors works quite well in accounting for both observed and unobserved heterogeneity, taking care of the composition bias in the estimates of wage cyclicality. Thus, the residual $\hat{w_{it}}$ denotes the worker wages corrected for the composition bias.

The residual wages are averaged over quarters for each subgroup, i.e., all workers and new hires. The wage index for subgroup j, $\hat{w}_{jt}$ is defined as follows:

$$\log \hat{w}_{jt} = \log w_{jt} - (x_{jt} - \bar{x}_j)'\beta, \tag{2}$$

where w_{jt} and x_{jt} refer to the average of the wages and the observables respectively for that subgroup of workers in quarter t, and $\bar{x}_j$ is the overall subgroup average of the characteristics.

2.3 Wage Cyclicality

Wage cyclicality captures the response of wage to change in aggregate labour productivity, and is measured as the coefficient of regression of log real wage index on log real labour productivity. Wages that are flexible will have the regression coefficient close to one while rigid wages will have the coefficient closer to zero. We estimate this regression in first differences in order to avoid spurious correlation.

$$\Delta \log \hat{w}_{jt} = \alpha_j + \eta_j \Delta \log y_t + \epsilon_{jt}, \tag{3}$$

where $\hat{w}_{jt}$ is the composition bias corrected real wage index and y_t real labour productivity. In addition, we also include quarter dummies to control for the seasonality while estimating this regression. The estimates of wage cyclicality obtained from this regression are shown in table 1.

Table 1: Wage Cyclicality

	All Workers		New Hires	
	1979-2019	1984-2019	1979-2019	1984-2019
Wage cyclicality	0.14	0.21	0.52	0.79
Standard error	0.12	0.15	0.40	0.47
Quarters	157	138	157	138

Note: Wage cyclicality is measured as the coefficient of regression of log real wage index on log real labour productivity as shown in equation (3). In addition, the regression also includes quarter dummies to control for seasonality.

Our results are consistent with the broad literature on wage cyclicality and with Haefke et al. (2013) in particular. To be specific, we are able to generate the three major findings of Haefke et al. (2013). First, wages are procyclical with respect to aggregate labour productivity. Second, the wages of new hires respond much more to productivity compared to the wages of all workers. Finally, comparing the estimates between 1979-2019 and 1984-2019, wages of both new hires and all workers are less elastic prior to 1984. This result provides an indication that wage rigidity might have reduced since Great Moderation starting 1984. Even though this regression provides preliminary evidence, we next employ wavelet tools in order to carefully examine the evolution of wage rigidity over the entire time period.

3. Wavelets

Continuous wavelet transform is a powerful tool that uses frequency domain analyses to dig into changes in data both over time and across frequencies. Thus, wavelets combines the power of time-varying regressions and spectral analysis in a single integrated framework. This makes wavelets an ideal tool for studying the changes in wage rigidity both over time and across frequencies. Aguiar-Conraria and Soares (2014) provides a detailed introduction and survey of the continuous wavelet transform and the various wavelet tools that can be used for analysing the data. We start by using wavelet coherency and wavelet phase-difference to study the relationship between wages and productivity. A number of papers like Aguiar-Conraria et al. (2012a), Aguiar-Conraria et al. (2012b), and Aguiar-Conraria et al. (2013) have also used these techniques under different contexts. In addition, we make use of wavelet gain, which is an analog of regression in a time frequency domain, to estimate the elasticity of wages across time and frequencies. Aguiar-Conraria et al. (2018) employs wavelet gain to estimate time-varying coefficients of the Taylor rule while Aguiar-Conraria et al. (2020) uses this tool to discuss the changes in Okun's law over time. We now provide a brief introduction of the various wavelet tools we use in our analysis.

3.1 Continuous Wavelet Transform

Fourier transform is a well known method to perform frequency domain analysis, where a time series is written as a combination of sine and cosine base functions. These sinosoid base functions do not change over time and hence Fourier transforms cannot capture the change in spectral characteristics of the signal over time. In order to estimate the changes in the time series across both time and frequencies, we need a base function that changes over time. A wavelet $\psi(t)$, as the name signifies, represents a small wave, oscillating around the time axis and loses its strength as it moves away from the centre. A wavelet transform decomposes the data in terms of time localized wavelets and hence enables us to capture the evolution of data in both time and frequency domains. The continuous wavelet transform of a series $x(t)$ with respect to a given wavelet $\psi(t)$ is given by the convolution

$$W_x(\tau, \mathrm{s}) = \frac{1}{\sqrt{|\mathrm{s}|}} \int_{-\infty}^{\infty} x(t)\psi^*\left(\frac{t-\tau}{\mathrm{s}}\right) dt \tag{4}$$

where $*$ denotes the complex conjugate. τ is the translation parameter controlling the time location while s is the dilation parameter capturing the width of the wavelet $\psi\left(\frac{t-\tau}{\mathrm{s}}\right)$. For $|s| > 1$, the function becomes broader thus leading to a lower frequency, and for $|s| < 1$, it becomes narrower corresponding to a higher frequency. In line with the literature, we use Morlet wavelet as our choice of the wavelet function and is given by $\psi(t) = \pi^{-\frac{1}{4}} e^{6it} e^{-\frac{t^2}{2}}$.

Wavelet power spectrum measures the variance distribution of the time series x across time and frequencies and is given by

$$(WPS)_x = W_x W_x^* = |W_x|^2. \tag{5}$$

In a time series analysis, we use covariance and correlation to study relationships between two variables. Similarly, in the context of wavelets, we define cross-wavelet power and coherency to study these relationships both across time and frequency. The cross-wavelet power between two series $y(t)$ and $x(t)$ is defined as the absolute value of the cross-wavelet transform, given by $W_{yx} = W_y W_x^*$. Similar to correlation, complex wavelet coherency between y and x is obtained by normalizing the cross-wavelet power with the square root of the wavelet powers of x and y, and is given by

$$\varrho_{yx} = \frac{S(W_{yx})}{[S(|W_y|^2)S(|W_x|^2)]^{1/2}}, \tag{6}$$

where S is a smoothing function across time and frequency. Denoting the smoothed cross-wavelet transform as S_{yx} and the square root of the smoothed wavelet power of x as $\sigma_x = \sqrt{S(|W_x|^2)} = \sqrt{S_{xx}}$, the complex wavelet coherency can be written as

$$\varrho_{yx} = \frac{S_{yx}}{\sigma_y \sigma_x} \tag{7}$$

Complex valued coherency can be represented in a polar form as $\varrho_{yx} = |\varrho_{yx}| e^{i\phi_{yx}}$. The absolute value of the complex wavelet coherency, denoted by $R_{yx} = |\varrho_{yx}|$, captures the magnitude of the relationship and is referred to as the wavelet coherency. Similarly, the angle of the complex coherency, ϕ_{yx}, represents the phase-difference between the two series. The wavelet phase-difference ϕ_{yx} provides information on the direction of the relationship across time and frequency, and also the relative leads and lags of the two series.

If the computed phase-difference is zero, then both the series exactly coincide with each other at the given frequency. If $\phi_{yx} \in (0, \frac{\pi}{2})$, then both the series move in the same direction (in phase), but y leads x. If $\phi_{yx} \in (-\frac{\pi}{2}, 0)$, then x leads y. Similarly, a phase-difference of π or $-\pi$ indicates an anti-phase relationship, with x leading if $\phi_{yx} \in (\frac{\pi}{2}, \pi)$ and y leading if $\phi_{yx} \in (-\pi, -\frac{\pi}{2})$.

3.2 Wavelet Gain

In order to answer how the elasticity of wage has changed over time and frequency, we need a regression setup with its coefficients depending both on time and frequency. Wavelet gain provides an analog of the regression framework across time and frequencies. With this tool, we will be able to estimate the wage cyclicality that is both time-varying and frequency-varying. The complex wavelet gain of y on x, denoted by $\mathcal{G}_{yx}$, is given by $\mathcal{G}_{yx} = \frac{S_{yx}}{S_{xx}} = \varrho_{yx}\frac{\sigma_y}{\sigma_x}$. Wavelet gain, G_{yx} is defined as the modulus of the complex wavelet gain

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

The wavelet gain can be interpreted as the absolute value of the regression coefficient of y on x at a given moment in time and a specific frequency. Thus, wavelet gain gives only the magnitude of the regression coefficient, while the sign of the coefficient can be obtained from the phase-difference ϕ_{yx}.

4. Results

We now use the wavelet tools discussed so far – wavelet coherency, phase-difference, and wavelet gain, to analyse the relationship between wages and productivity, and how the elasticity of wages behaves across time and at different frequencies.[4] Wavelet coherency gives the magnitude of correlaton between real wages and productivity over time and at different frequencies, while the phase-difference gives the direction of the relationship. Wavelet gain captures the absolute value of the coefficient of regression between wages and productivity, and hence measures the cyclicality of wages at different time and frequencies. We analyse the wage cyclicality of all the workers and new hires separately as the wages of new hires behave differently over the business cycles.

4.1 All Workers

The results from the wavelet analysis for all the workers are summarized in figure 1. Following Aguiar-Conraria et al. (2018), we analyse the relationship over a wide range of frequencies. We present results for coherency, phase-difference, and wavelet gain for three frequency intervals, $1.5 \sim 4$ years (shorter end of business cycles), $4 \sim 8$ years (longer end of business cycles), and $8 \sim 20$ years (long-run cycles).

The wavelet coherency shows that the regions of high coherency are sparsely distributed at the shorter end of the business cycle. Among the long-run cycles, in particular the upper end of the spectrum, we find that the coherency is consistently high throughout the entire sample period. Interestingly, at the business cycle frequencies of $4 \sim 8$ years, we do not find much correlation in the earlier part of our sample. But post 2000, the coherency has become very strong as seen by the emergence of dark red regions in the coherency plot. This shows that the relationship between real wages and productivity changes has strengthened over time.

Even though the coherency gives us a measure of the magnitude of the correlation, we need to pair it up with the phase-difference to understand the direction of

[4] The wavelet results are generated using ASToolbox v.2018 available at https://sites.google.com/site/aguiarconraria/wavelets-and-economics/the-astoolbox. Aguiar-Conraria and Soares (2014) provides a detailed description of this toolbox.

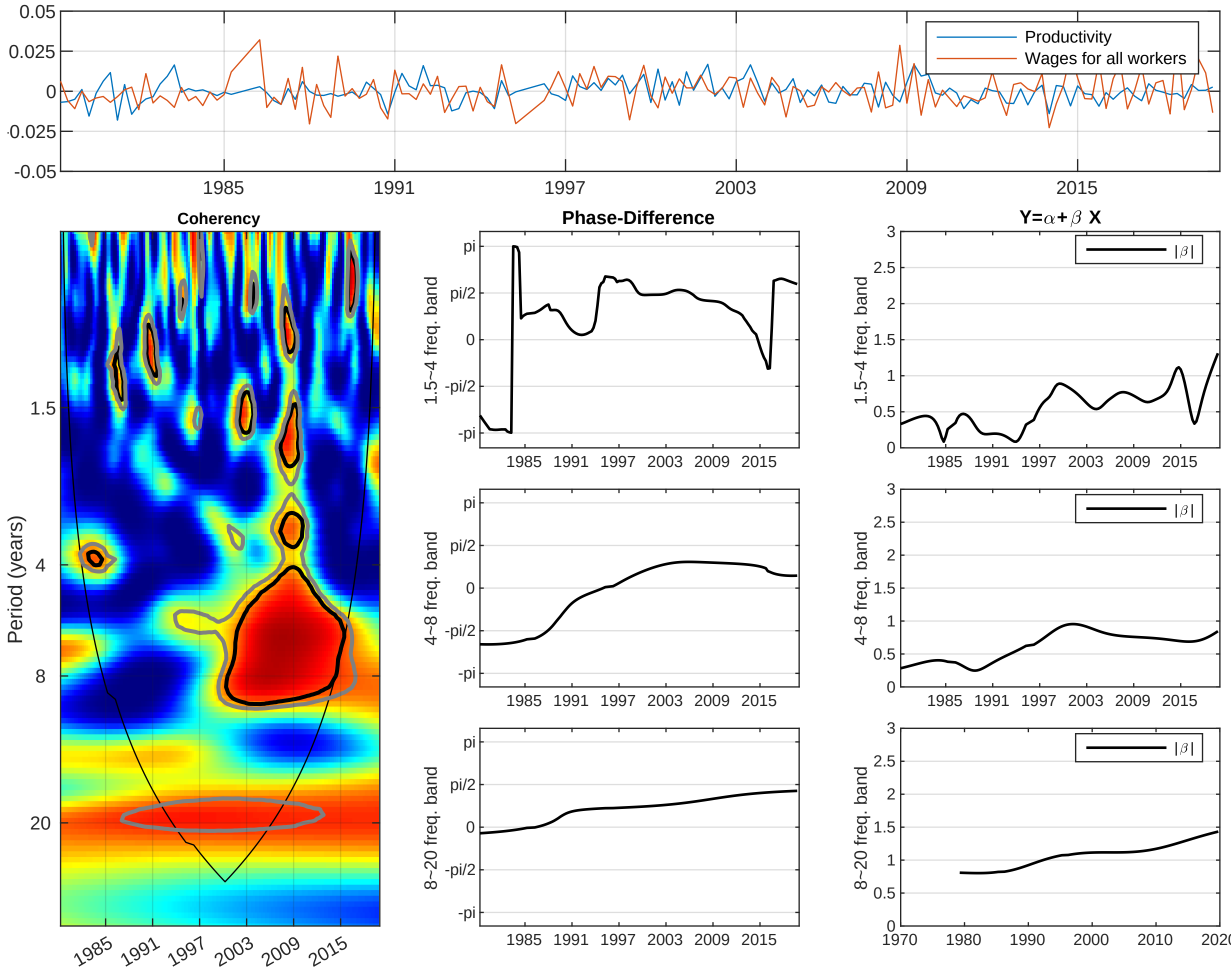

Figure 1: **All workers:** Change in log real wages and log real labour productivity between 1979-2019 (top). Wavelet Coherency between wages and productivity (left panel). Warm colours indicate high coherency, and cold colours indicate low coherency. The black parabolic line is the cone of influence that highlights the regions affected by edge effects. In the centre are the phase differences, and on the right are the wavelet gains for the three frequency intervals.

the correlation and also the lead/lag structure between the variables. The phase difference at the short end of the business cycle varies considerably over time. Despite that, wages and productivity are positively related for most of the sample, except for occasional negative relationship at the start and end of the sample. At lower frequencies, the phase differences are considerably more stable, with wages moving positively with productivity most of the time at both business cycle frequency and in the long run. Additionally, we find that, predominantly wages lead productivity in our sample, except during the early part of the sample where productivity leads wages.

We next turn towards the main focus of our paper, wage cyclicality, as measured by the wavelet gain. We find that there is considerable variation in our estimates at high frequencies, while the trend becomes smoother at lower frequencies. Comparing the average cyclicality across frequency bands, we find that the magnitude of cyclicality is higher at longer frequencies compared to shorter ones. This indicates that wages take time to adjust, and the procyclical relationship gets stronger when we look at longer time horizons. Importantly, we find that the wage elasticity shows an increasing trend over time across all the frequency intervals. Thus, wages have become more flexible over time with respect to changes in productivity, and this pattern holds across all the frequencies.

4.2 New Hires

Figure 2 summarizes the results for new hires. Even for the new hires, the regions of high coherency at shorter frequencies are scarcely distributed. Just like for all the workers, the coherency is consistently high at the long-run frequencies throughout the sample period. And post 2000, the correlation becomes stronger with increase in the spread of high coherency regions as seen in the coherency plot. Similar to our previous results, the phase-differences have significant variations at the short end of the business cycle. But the wages are predominantly procyclical at both business cycle and long-run frequencies, with wages leading productivity most of the time.

We earlier established that wages of new hires respond more to productivity changes compared to all the workers. Consistent with this finding, we find that the wavelet gain of new hires is larger than that of all workers across all the frequency intervals. Similar to the previous case, we find that there is substantial variation at higher frequencies. And more importantly, just like the case of all workers, we find that the wage cyclicality of new hires have also increased over time and this increase is seen across all frequencies. This analysis shows that the wages of all workers as well as new hires have become

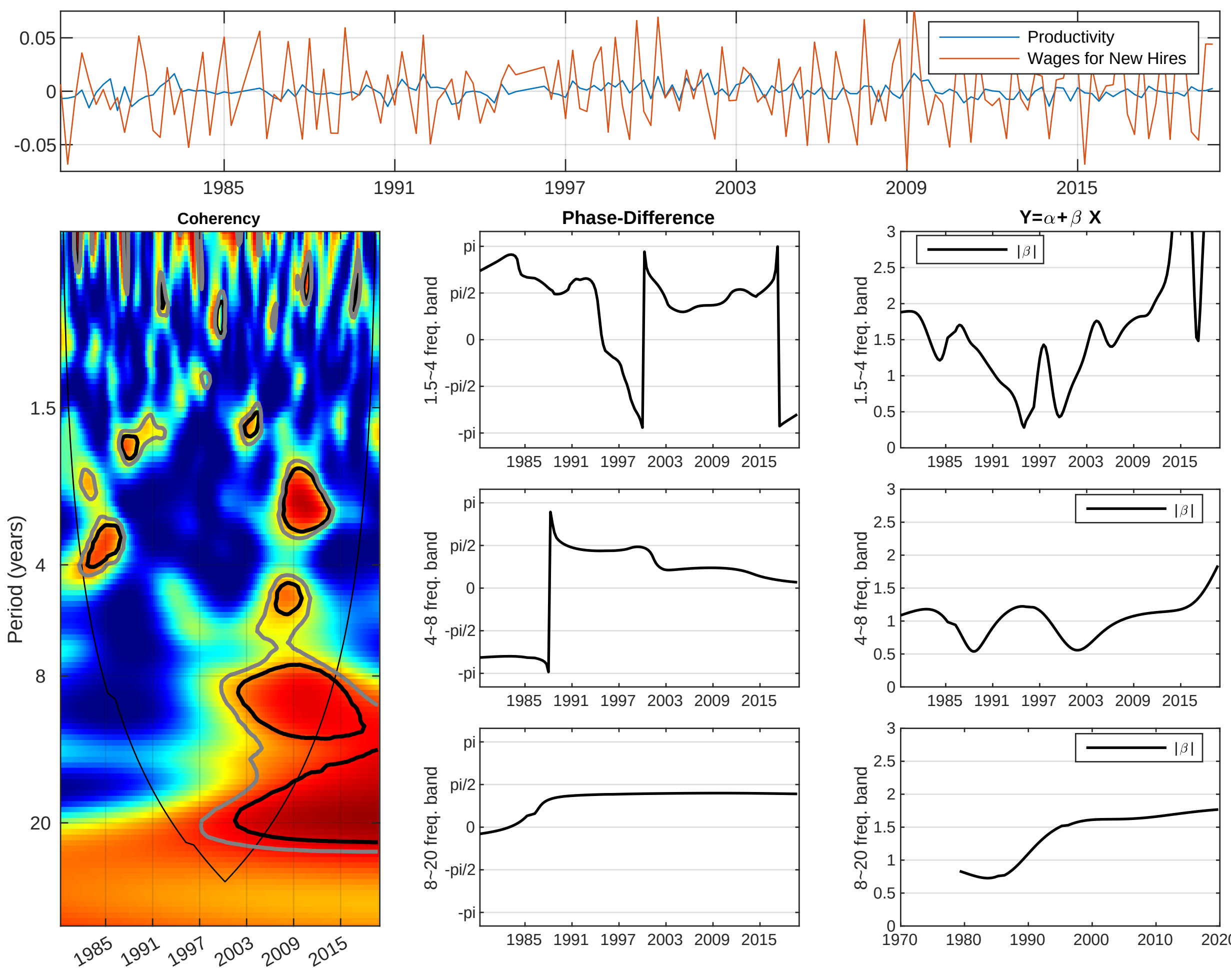

Figure 2: **New hires:** Log real wages and log real labour productivity between 1979-2019 (top). Wavelet Coherency between wages and productivity (left panel). Warm colours indicate high coherency, and cold colours indicate low coherency. The black parabolic line is the cone of influence that highlights the regions affected by edge effects. In the centre are the phase differences, and on the right are the wavelet gains for the three frequency intervals.

more procyclical over time.

4.3 Cyclicality across Gender and Skills

We now document the evolution of wage elasticity separately across gender and skills. Following the majority of literature, we classify workers with a college degree or more as 'high-skilled', while those with less than a college degree as 'low-skilled'. The results of the wavelet analysis for the different groups are provided in appendix A. The broad results continue to hold, i.e., wages are procyclical across the various subgroups of workers, and the wages are more flexible in the longer frequencies compared to the shorter ones. Interestingly, we find that, the increase in wage elasticity over time is predominantly driven by the low-skilled workers. Comparing figures A1 and A2, we can clearly see that there is an increasing trend in elasticity of wages among low-skilled men, while the elasticity doesn't increase among high-skilled men. And, we find a similar pattern among female workers as can be seen in figures A3 and A4, with the increase in wage cyclicality more prominent among low-skilled women. Appendix A also contains the results for new hires separately, where again we find that the increase in cyclicality is more pronounced among low-skilled workers.

5. Implications

We now discuss how our findings relates to the broader changes in the labour market and also its implications for labour search and matching models.

5.1 Great Moderation and Structural Changes

Using wavelet analysis, we find that the cyclicality of wages has increased over time for both new hires and all workers, and this has implications for understanding Great Moderation. Galí and Gambetti (2009) and Stiroh (2009) showed that the volatility and correlations of various labour market variables showed a marked decline post 1984. Studies like Champagne and Kurmann (2013) and Haefke et al. (2013) show that, in contrast to a number of macroeconomic variables, the volatility of real wages increased during this time. Using our analysis, we also find that the wage rigidity has reduced over time for all the workers as well as for new hires. This finding thus strengthens the argument proposed by Champagne and Kurmann (2013), Nucci and Riggi (2013), and others, that more flexible wages predominantly due to the emergence of performance-pay contracts led to Great Moderation.

This increase in wage elasticity that we document is also consistent with the declining worker power hypothesis put forth by Stansbury and Summers (2020). The bargaining power that enabled workers to bargain for long-term wage contracts, which helped them smooth their wages over the business cycles has been declining considerably. This in turn could lead to higher variability in wages over the business cycle. Additionally, we find that the increase in wage elasticity is concentrated among the low-skilled workers. This adds strength to the previous argument, as the decline in unionization rates was significantly larger among the non-college educated workers as shown in Stansbury and Summers (2020), and other studies relating the decline in unionization with the increase in income inequality, such as, Card (1996), Card et al. (2004), Farber et al. (2021), and Fortin et al. (2021).

5.2 Labour Search and Unemployment Volatility Puzzle

The estimates of wage elasticity also has implications for the model fit of canonical search and matching models. Shimer (2005) and Costain and Reiter (2008) documented that the benchmark search and matching model generates much lower volatility in unemployment and vacancies compared to the data. Hall (2005) by introducing equilibrium wage stickiness in the model in place of Nash bargaining showed that, this significantly increases the volatility in both vacancies and unemployment. Relatedly, Shimer (2004) and Pissarides (2009) showed that, job creation and unemployment in the model is influenced by the wage behaviour of newly hired workers, and not affected by that of existing employment relationships. This implies, in order to improve the fit of the model, the wages of new hires should not respond much over business cycles. Due to these findings, a large number of studies followed suit and introduced some form of wage stickiness in their models to match the data (Menzio (2005); Farmer and Hollenhorst (2006); Moen and Rosen (2006); Braun et al. (2006); Blanchard and Galí

(2007); Hall and Milgrom (2008); Gertler and Trigari (2009); Kennan (2010); Shimer (2010); Michaillat (2012); Christiano et al. (2016)).

Even though wage stickiness was one of the widely used solutions to improve the model fit, a number of studies like Bils (1985), Solon et al. (1994), Devereux (2001), and Haefke et al. (2013) document that real wages of all workers, controlling for composition bias, are quite procyclical and wages of new hires respond even more over the business cycles compared to the wages of existing workers. These studies show that, assuming rigid wages in a search model, particularly that of new hires, is not empirically valid. In this paper, we further document that the cyclicality of wages for all the workers and new hires has actually increased over time. Thus, with our finding, the assumption of wage rigidity to solve the unemployment volatility puzzle has become even more untenable.

6. Conclusion

We employ continuous wavelet tools to analyse how the wage cyclicality in US has evolved over time and across different frequencies. Using individual level wage data from CPS, we find that, (i) wages are procyclical, with wages of new hires more cyclical compared to continuing workers, (ii) wages are more elastic over longer time horizons compared to the shorter end of the business cycle, and (iii) wage cyclicality has increased over time across all the frequency intervals. This finding is consistent with the broader structural changes in the labour market and also has implications for labour search and matching models.

References

Aguiar-Conraria, Luís and Maria Joana Soares, "The continuous wavelet transform: Moving beyond uni-and bivariate analysis," *Journal of economic surveys*, 2014, *28* (2), 344–375.

_ , Manuel MF Martins, and Maria Joana Soares, "The yield curve and the macro-economy across time and frequencies," *Journal of Economic Dynamics and Control*, 2012, *36* (12), 1950–1970.

_ , _ , and _ , "Estimating the Taylor rule in the time-frequency domain," *Journal of Macroeconomics*, 2018, *57*, 122–137.

_ , _ , and _ , "Okun's law across time and frequencies," *Journal of Economic Dynamics and Control*, 2020, *116*, 103897.

_ , Pedro C Magalhães, and Maria Joana Soares, "Cycles in politics: wavelet analysis of political time series," *American Journal of Political Science*, 2012, *56* (2), 500–518.

_ , Pedro C Magalhaes, and Maria Joana Soares, "The nationalization of electoral cycles in the United States: a wavelet analysis," *Public Choice*, 2013, *156* (3), 387–408.

Bils, Mark J, "Real wages over the business cycle: evidence from panel data," *Journal of Political economy*, 1985, *93* (4), 666–689.

Blanchard, Olivier and Jordi Galí, "Real wage rigidities and the New Keynesian model," *Journal of money, credit and banking*, 2007, *39*, 35–65.

Bodkin, Ronald G, "Real wages and cyclical variations in employment: a re-examination of the evidence," *The Canadian Journal of Economics/Revue canadienne d'Economique*, 1969, *2* (3), 353–374.

Braun, Helge et al., "(Un) Employment Dynamics: The Case of Monetary Policy Shocks," in "2006 Meeting Papers" number 87 Society for Economic Dynamics 2006.

Card, David, "The effect of unions on the structure of wages: A longitudinal analysis," *Econometrica: Journal of the Econometric Society*, 1996, pp. 957–979.

— , Thomas Lemieux, and W Craig Riddell, "Unions and wage inequality," *Journal of Labor research*, 2004, *25*, 519–559.

Carneiro, Anabela, Paulo Guimarães, and Pedro Portugal, "Real wages and the business cycle: Accounting for worker, firm, and job title heterogeneity," *American Economic Journal: Macroeconomics*, 2012, *4* (2), 133–52.

Champagne, Julien and André Kurmann, "The great increase in relative wage volatility in the United States," *Journal of Monetary Economics*, 2013, *60* (2), 166–183.

Christiano, Lawrence J, Martin S Eichenbaum, and Mathias Trabandt, "Unemployment and business cycles," *Econometrica*, 2016, *84* (4), 1523–1569.

Costain, James S and Michael Reiter, "Business cycles, unemployment insurance, and the calibration of matching models," *Journal of Economic Dynamics and control*, 2008, *32* (4), 1120–1155.

Devereux, Paul J, "The cyclicality of real wages within employer-employee matches," *ILR Review*, 2001, *54* (4), 835–850.

— and Robert A Hart, "Real wage cyclicality of job stayers, within-company job movers, and between-company job movers," *ILR Review*, 2006, *60* (1), 105–119.

Dunlop, John T, "The movement of real and money wage rates," *The Economic Journal*, 1938, *48* (191), 413–434.

Farber, Henry S, Daniel Herbst, Ilyana Kuziemko, and Suresh Naidu, "Unions and inequality over the twentieth century: New evidence from survey data," *The Quarterly Journal of Economics*, 2021, *136* (3), 1325–1385.

Farmer, Roger and Andrew Hollenhorst, "Shooting the auctioneer," 2006.

Fortin, Nicole M, Thomas Lemieux, and Neil Lloyd, "Labor market institutions and the distribution of wages: The role of spillover effects," *Journal of Labor Economics*, 2021, *39* (S2), S369–S412.

Galí, Jordi and Luca Gambetti, "On the sources of the great moderation," *American Economic Journal: Macroeconomics*, 2009, *1* (1), 26–57.

Geary, Patrick T and John Kennan, "The employment-real wage relationship: an international study," *Journal of Political Economy*, 1982, *90* (4), 854–871.

Gertler, Mark and Antonella Trigari, "Unemployment fluctuations with staggered Nash wage bargaining," *Journal of political Economy*, 2009, *117* (1), 38–86.

Haefke, Christian, Marcus Sonntag, and Thijs Van Rens, "Wage rigidity and job creation," *Journal of monetary economics*, 2013, *60* (8), 887–899.

Hall, Robert E, "Employment fluctuations with equilibrium wage stickiness," *American economic review*, 2005, *95* (1), 50–65.

_ and Paul R Milgrom, "The limited influence of unemployment on the wage bargain," *American economic review*, 2008, *98* (4), 1653–1674.

Hart, Robert A, "Worker–job matches, job mobility and real wage cyclicality," *Economica*, 2006, *73* (290), 287–298.

Kennan, John, "Private information, wage bargaining and employment fluctuations," *The Review of Economic Studies*, 2010, *77* (2), 633–664.

Menzio, Guido, "High frequency wage rigidity," *Manuscript. Univ. Pennsylvania*, 2005.

Michaillat, Pascal, "Do matching frictions explain unemployment? Not in bad times," *American Economic Review*, 2012, *102* (4), 1721–1750.

Moen, Espen R and Asa Rosen, "Incentives in competitive search equilibrium and wage rigidity," 2006.

Neftci, Salih N, "A time-series analysis of the real wages-employment relationship," *Journal of Political Economy*, 1978, *86* (2, Part 1), 281–291.

Nucci, Francesco and Marianna Riggi, "Performance pay and changes in US labor market dynamics," *Journal of Economic Dynamics and Control*, 2013, *37* (12), 2796–2813.

Pissarides, Christopher A, “The unemployment volatility puzzle: Is wage stickiness the answer?,” *Econometrica*, 2009, *77* (5), 1339–1369.

Sargent, Thomas J, “Estimation of dynamic labor demand schedules under rational expectations,” *Journal of political Economy*, 1978, *86* (6), 1009–1044.

Shimer, Robert, “The consequences of rigid wages in search models,” *Journal of the European Economic Association*, 2004, *2* (2-3), 469–479.

_ , “The cyclical behavior of equilibrium unemployment and vacancies,” *American economic review*, 2005, *95* (1), 25–49.

_ , “Labor markets and business cycles,” 2010.

Shin, Donggyun, “Cyclicality of real wages among young men,” *Economics Letters*, 1994, *46* (2), 137–142.

Solon, Gary, Robert Barsky, and Jonathan A Parker, “Measuring the cyclicality of real wages: how important is composition bias?,” *The quarterly journal of economics*, 1994, *109* (1), 1–25.

Stansbury, Anna and Lawrence H Summers, “Declining worker power and American economic Performance,” *Brookings Papers on Economic Activity*, 2020, *156*.

Stiroh, Kevin J, “Volatility accounting: A production perspective on increased economic stability,” *Journal of the European Economic Association*, 2009, *7* (4), 671–696.

Stockman, Alan C, “Aggregation bias and the cyclical behavior of real wages,” *Unpublished manuscript*, 1983.

Sumner, Scott and Stephen Silver, “Real wages, employment, and the Phillips curve,” *Journal of Political Economy*, 1989, *97* (3), 706–720.

Tarshis, Lorie, “Changes in real and money wages,” *The Economic Journal*, 1939, *49* (193), 150–154.

Online Appendix

A. Cyclicality across Gender and Skills

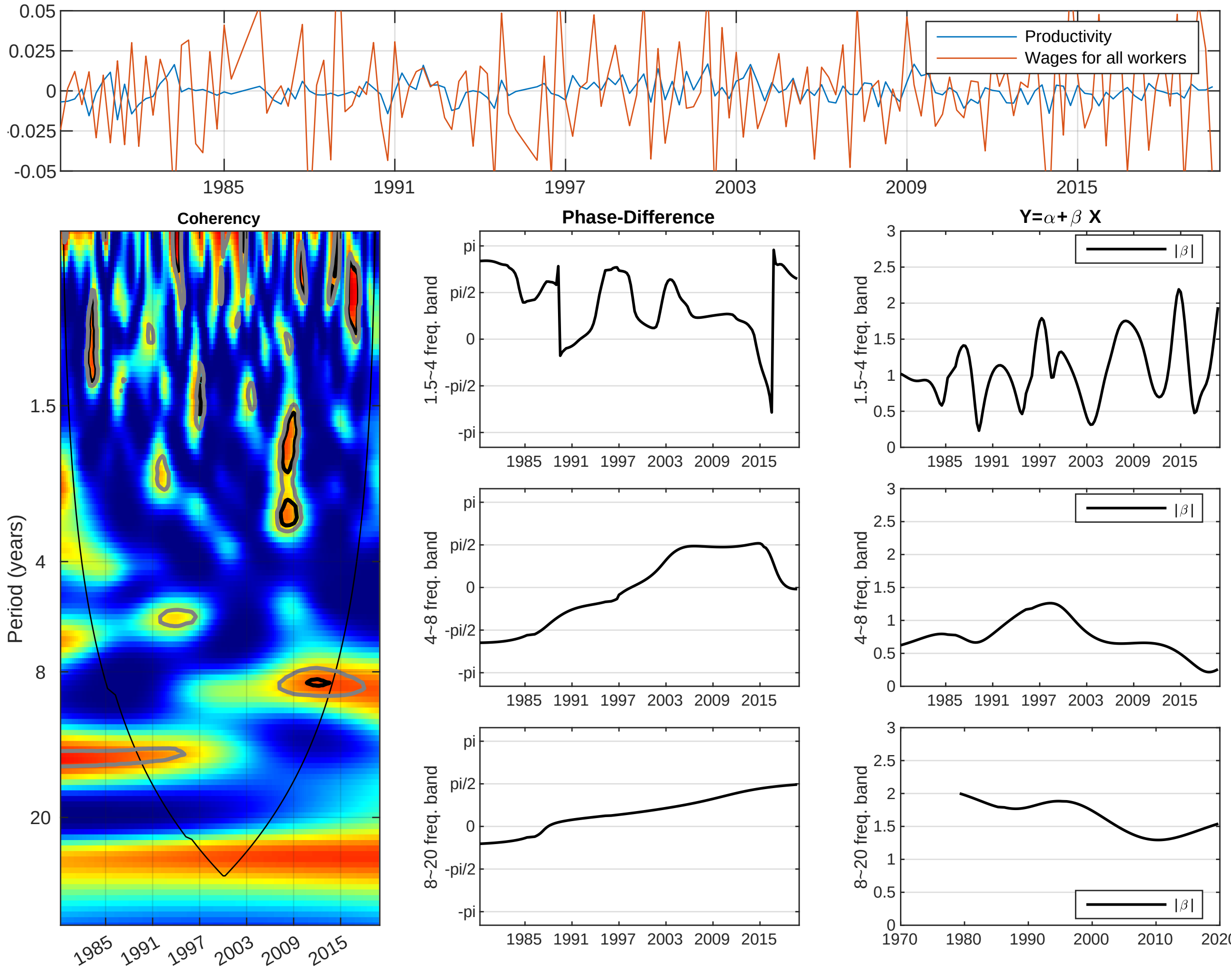

Figure A1: **High-skilled male, All workers:** Change in log real wages and log real labour productivity between 1979-2019 (top). Wavelet Coherency between wages and productivity (left panel). Warm colours indicate high coherency, and cold colours indicate low coherency. The black parabolic line is the cone of influence that highlights the regions affected by edge effects. In the centre are the phase differences, and on the right are the wavelet gains for the three frequency intervals.

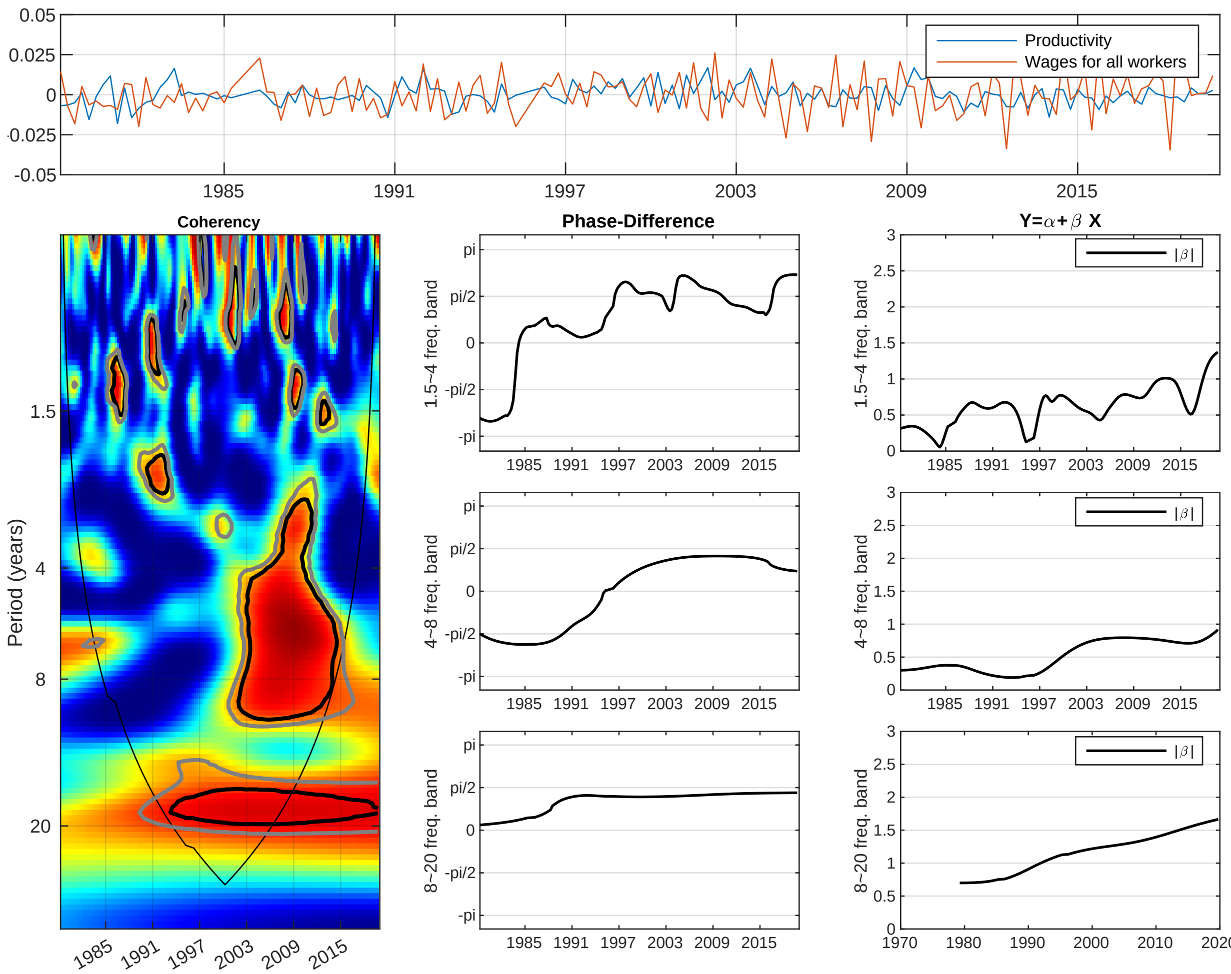

Figure A2: **Low-skilled male, All workers:** Change in log real wages and log real labour productivity between 1979-2019 (top). Wavelet Coherency between wages and productivity (left panel). Warm colours indicate high coherency, and cold colours indicate low coherency. The black parabolic line is the cone of influence that highlights the regions affected by edge effects. In the centre are the phase differences, and on the right are the wavelet gains for the three frequency intervals.

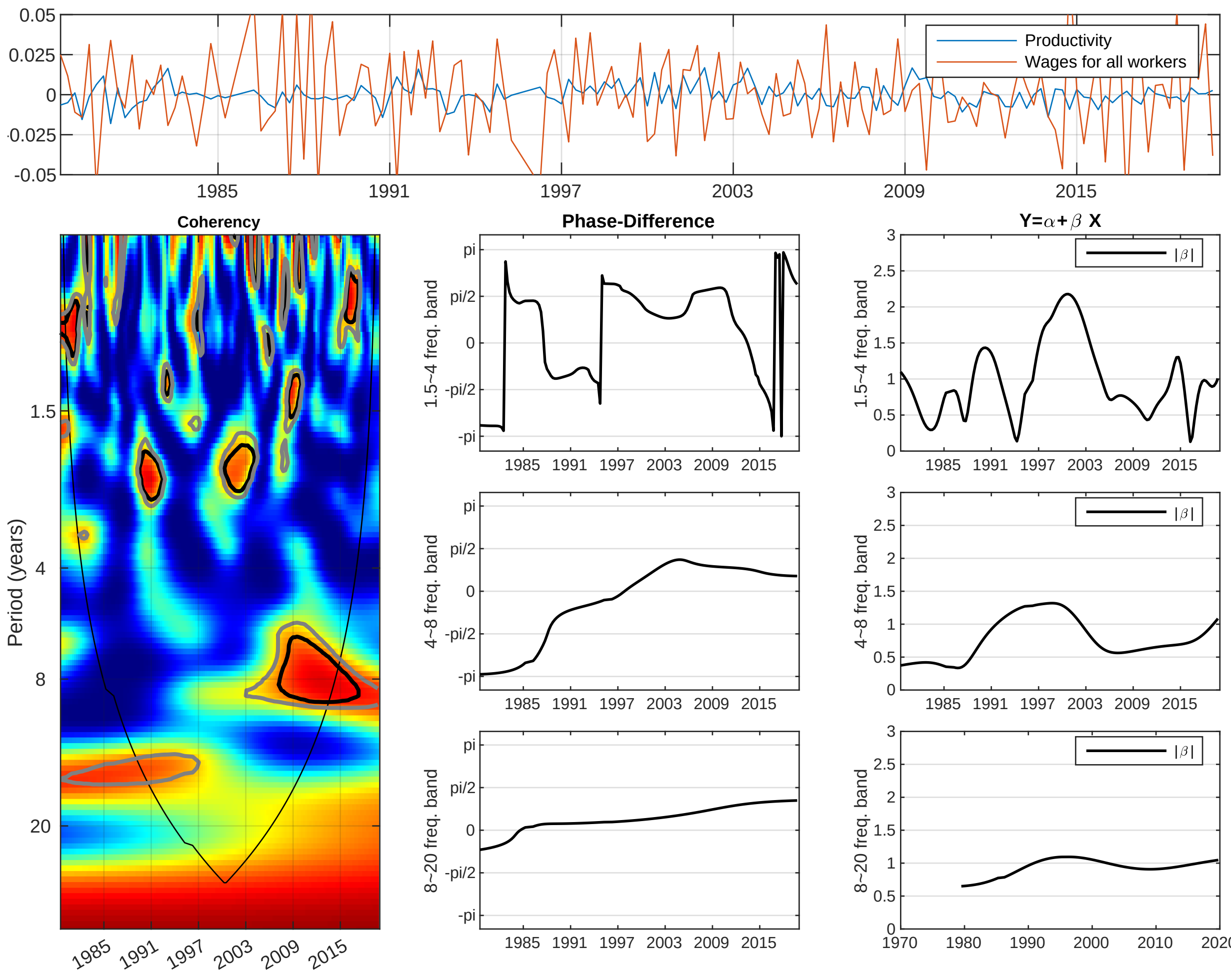

Figure A3: **High-skilled female, All workers:** Change in log real wages and log real labour productivity between 1979-2019 (top). Wavelet Coherency between wages and productivity (left panel). Warm colours indicate high coherency, and cold colours indicate low coherency. The black parabolic line is the cone of influence that highlights the regions affected by edge effects. In the centre are the phase differences, and on the right are the wavelet gains for the three frequency intervals.

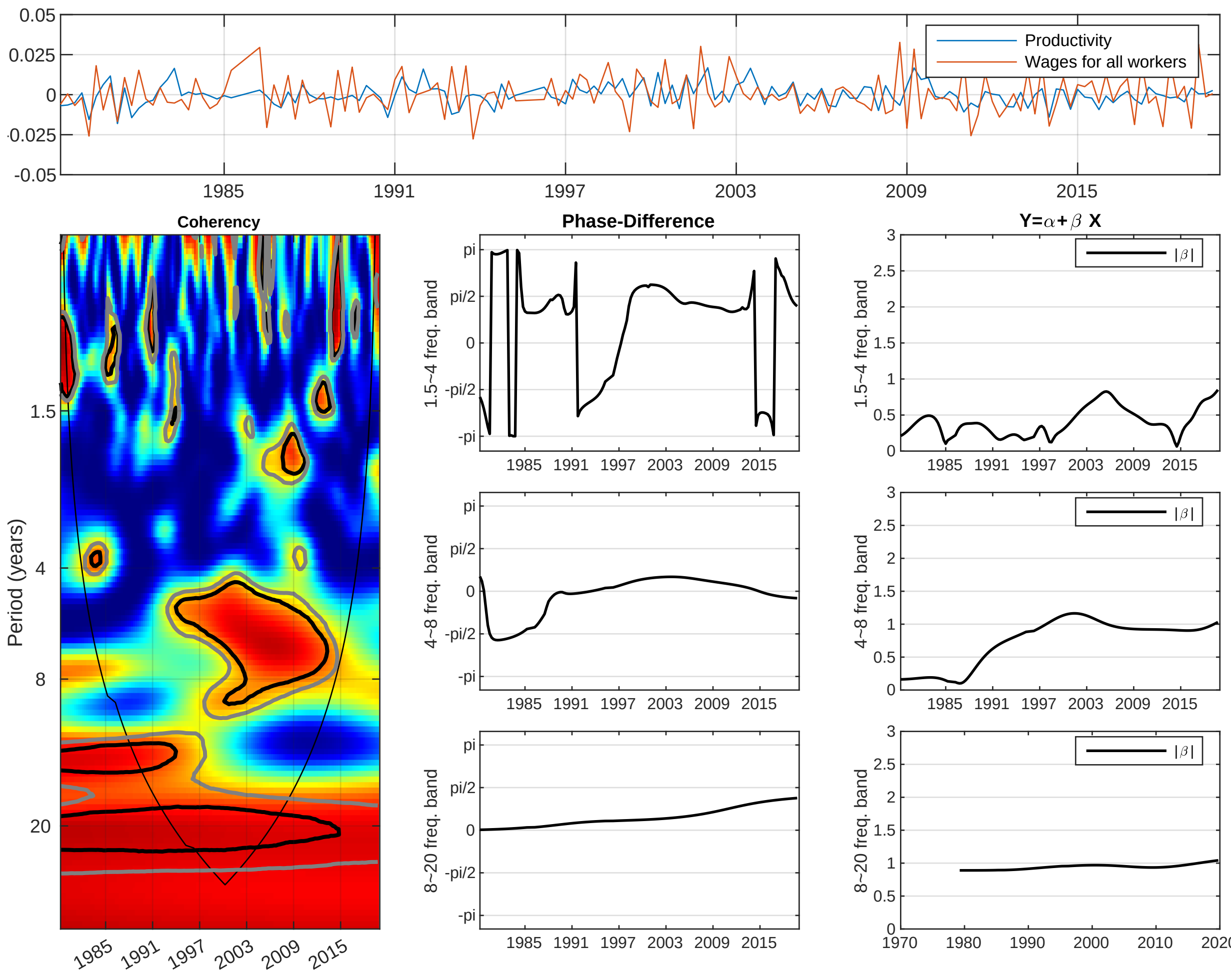

Figure A4: **Low-skilled female, All workers:** Change in log real wages and log real labour productivity between 1979-2019 (top). Wavelet Coherency between wages and productivity (left panel). Warm colours indicate high coherency, and cold colours indicate low coherency. The black parabolic line is the cone of influence that highlights the regions affected by edge effects. In the centre are the phase differences, and on the right are the wavelet gains for the three frequency intervals.

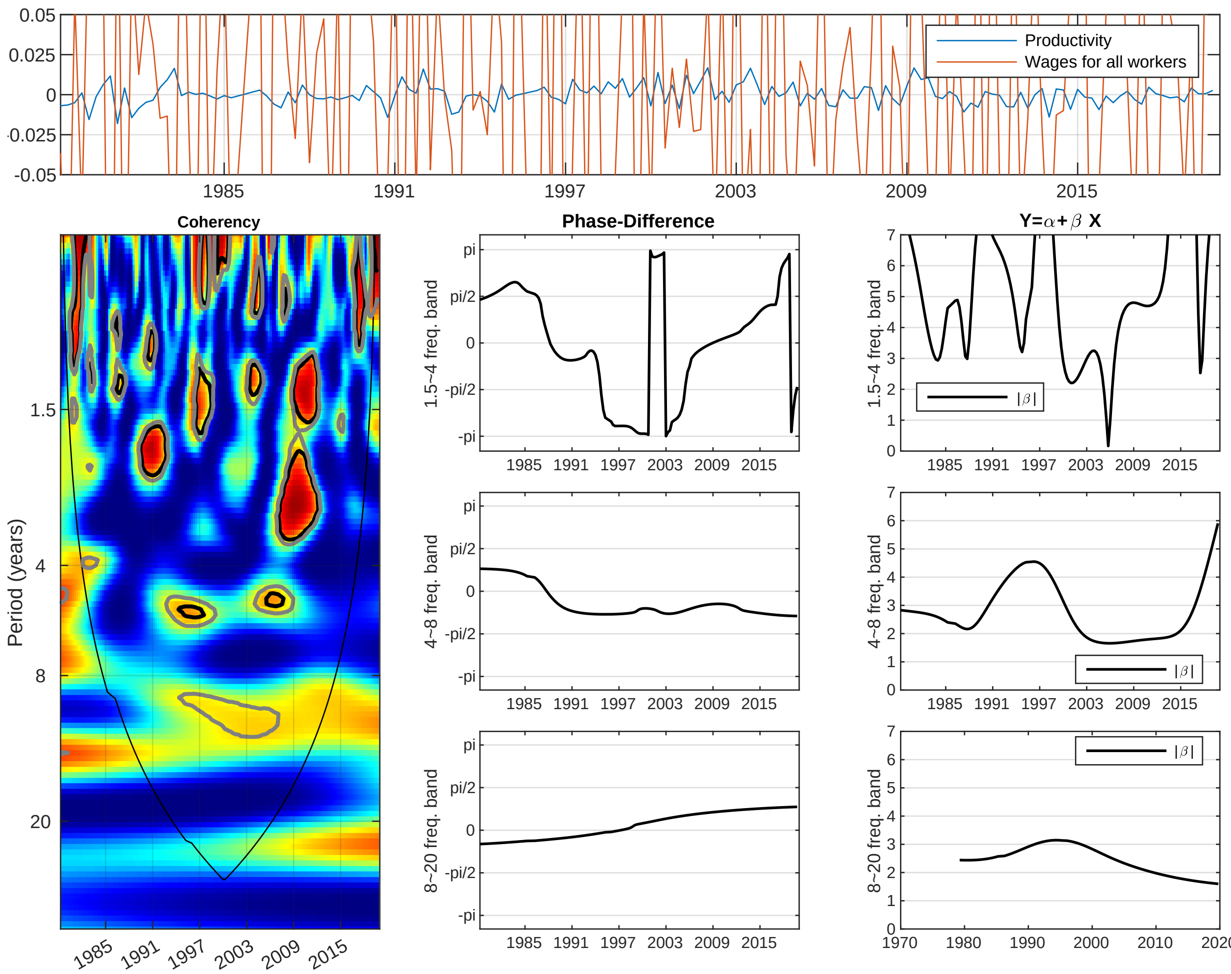

Figure A5: **High-skilled male, New hires:** Change in log real wages and log real labour productivity between 1979-2019 (top). Wavelet Coherency between wages and productivity (left panel). Warm colours indicate high coherency, and cold colours indicate low coherency. The black parabolic line is the cone of influence that highlights the regions affected by edge effects. In the centre are the phase differences, and on the right are the wavelet gains for the three frequency intervals.

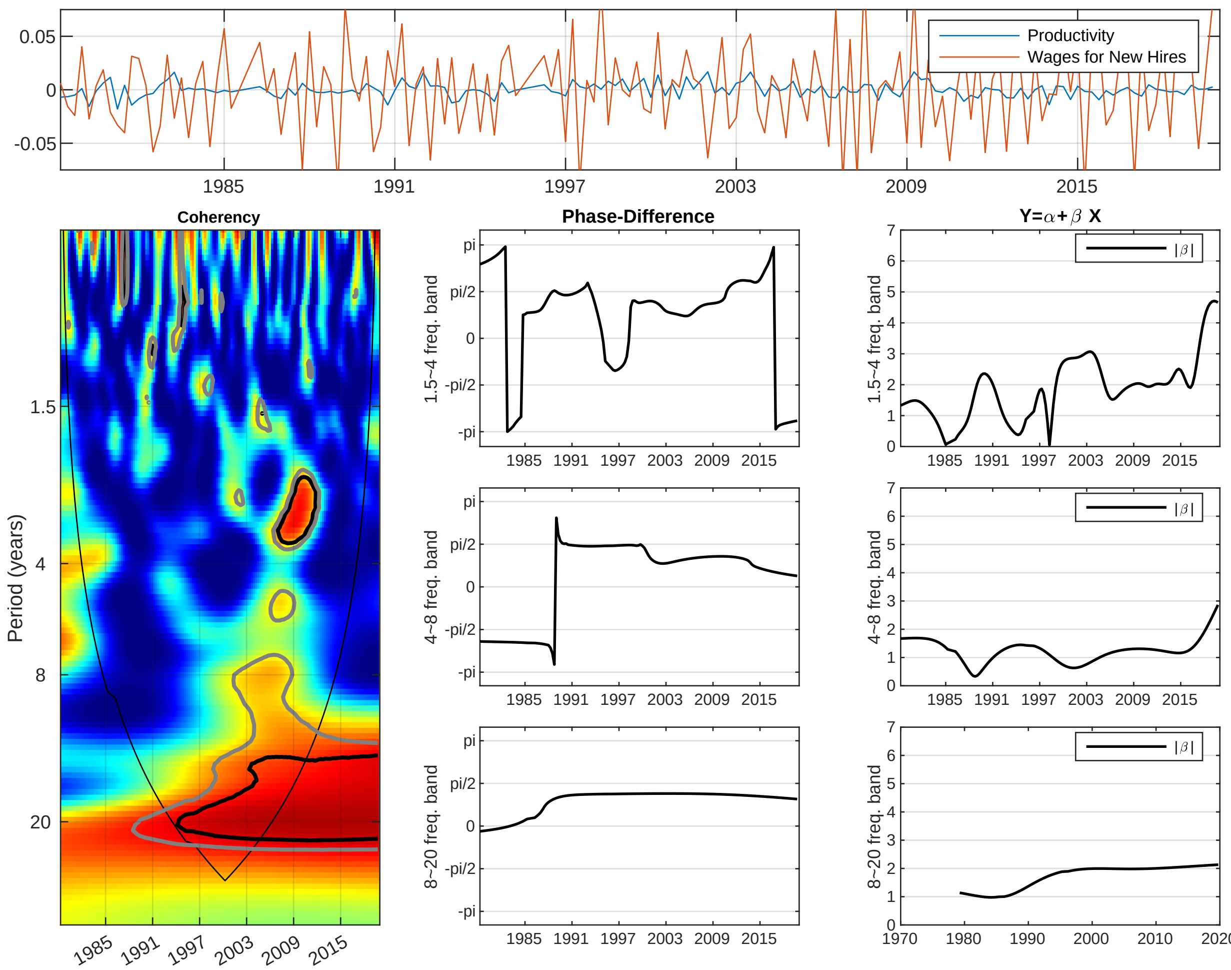

Figure A6: **Low-skilled male, New hires:** Change in log real wages and log real labour productivity between 1979-2019 (top). Wavelet Coherency between wages and productivity (left panel). Warm colours indicate high coherency, and cold colours indicate low coherency. The black parabolic line is the cone of influence that highlights the regions affected by edge effects. In the centre are the phase differences, and on the right are the wavelet gains for the three frequency intervals.

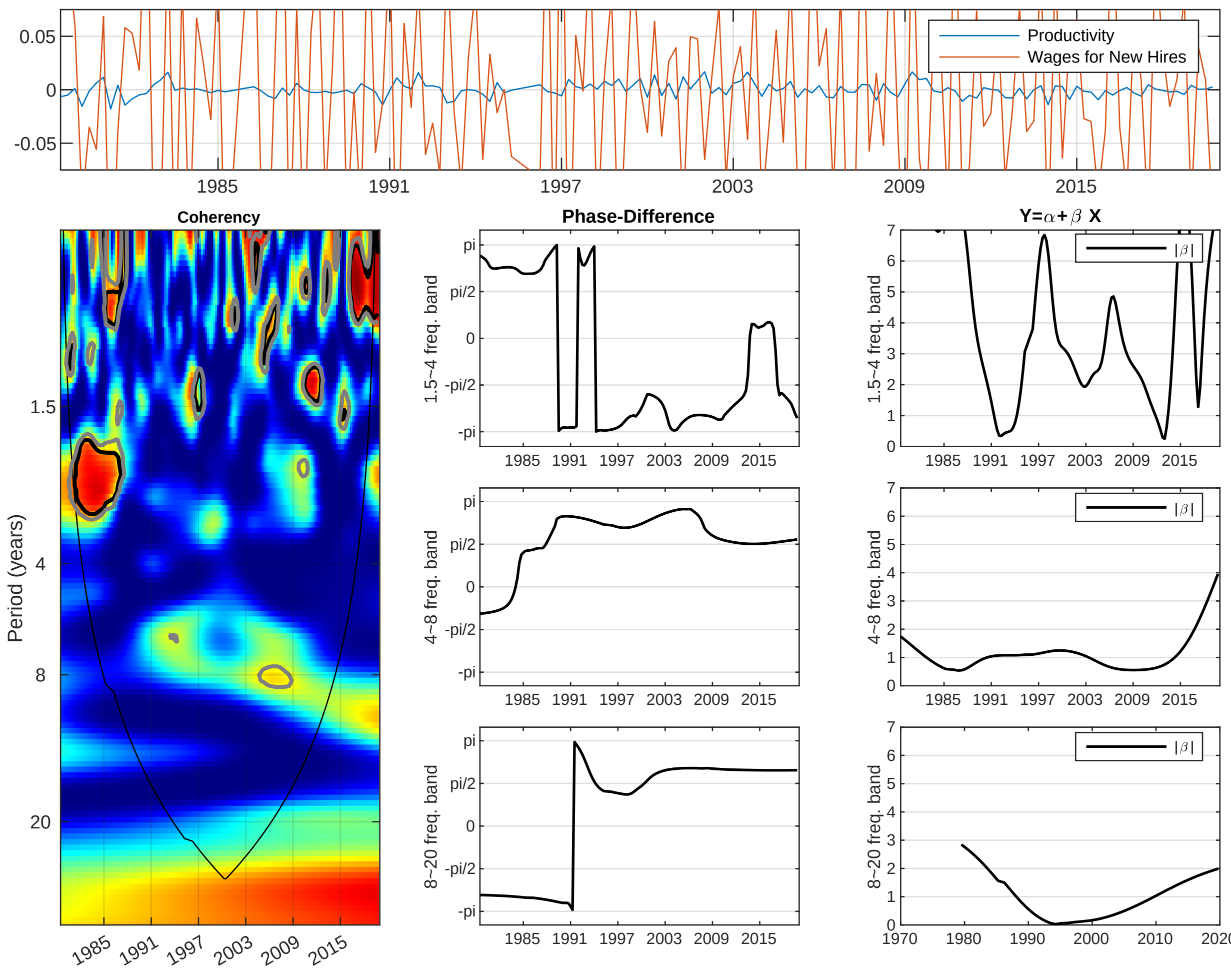

Figure A7: **High-skilled female, New hires:** Change in log real wages and log real labour productivity between 1979-2019 (top). Wavelet Coherency between wages and productivity (left panel). Warm colours indicate high coherency, and cold colours indicate low coherency. The black parabolic line is the cone of influence that highlights the regions affected by edge effects. In the centre are the phase differences, and on the right are the wavelet gains for the three frequency intervals.

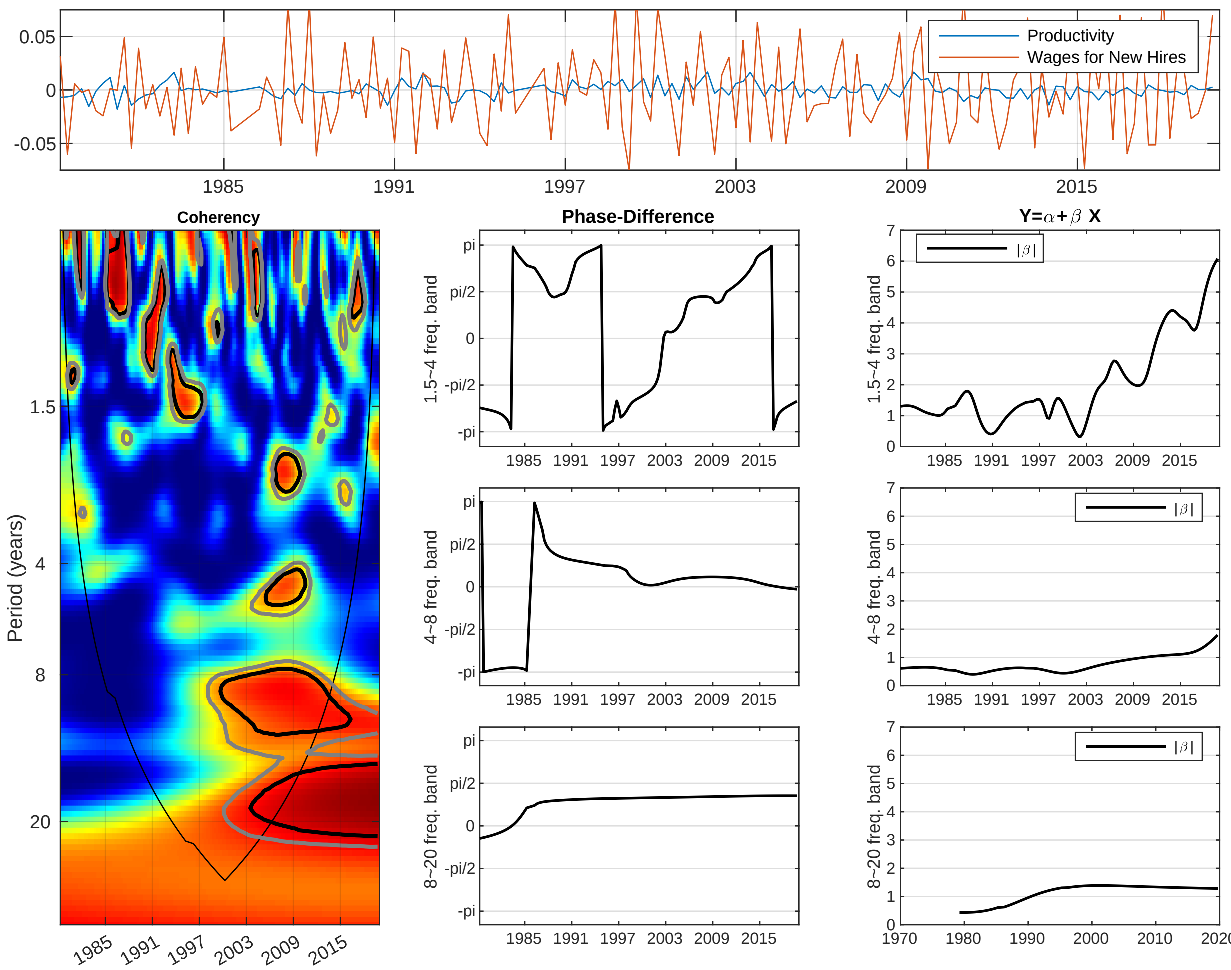

Figure A8: **Low-skilled female, New hires:** Change in log real wages and log real labour productivity between 1979-2019 (top). Wavelet Coherency between wages and productivity (left panel). Warm colours indicate high coherency, and cold colours indicate low coherency. The black parabolic line is the cone of influence that highlights the regions affected by edge effects. In the centre are the phase differences, and on the right are the wavelet gains for the three frequency intervals.