Manufacturing Earnings and Cycles: New Evidence*

Robert A. Hart†
Department of Economics, University of Stirling
James R. Malley
Department of Economics, University of Glasgow
Ulrich Woitek
Department of Economics, University of Munich

February 24, 2003

Abstract

In the time domain, the observed cyclical behavior of the real wage hides a range of economic influences that give rise to cycles of differing lengths and strengths. This may serve to produce a distorted picture of wage cyclicality. Here, we employ and develop frequency domain methods that allow us to assess the relative contribution of cyclical frequency bands on real wage earnings. Earnings are decomposed into standard and overtime components. We also distinguish between consumption and production wages. Frequency domain analysis is carried out in relation to wages alone (the univariate case) and to wages in relation to a selected range of cyclical economic indicators (multivariate). We establish that all key components of real wages are strongly pro-cyclical but display significant co-variations with more than one frequency band. Moreover, components are by no means uniformly associated with each of the chosen proxies for the cycle.

JEL Classification: E32, J31

*We are grateful to Michael Burda, Christian Dustmann, Campbell Leith, Jacques Melitz, Albrecht Ritschl, Harald Uhlig, and seminar participants at Humboldt University and the Glasgow-Strathclyde Joint Seminar in Macroeconomics for helpful comments and suggestions on this and earlier versions of the paper. We would also like to thank John Stinson and Sharon Gibson of the Bureau of Labor Statistics, U.S. Department of Labor for providing unpublished data relating to hours of work and consumer prices respectively.

†Corresponding Author: Robert Hart, Department of Economics, University of Stirling, Stirling FK9 4LA, UK, email: r.a.hart@stir.ac.uk.

Observed real wages are not constant over the cycle, but neither do they exhibit consistent pro- or counter-cyclical movements. This suggests that any attempt to assign systematic real wage movements a central role in an explanation of business cycles is doomed to failure. (Lucas, 1977)

1 Introduction

Recent years have witnessed a marked shift in economists' views of the behavior of real wages over the business cycle. The prevailing wisdom, emanating largely from aggregate time series investigations, is that wages are at most weakly pro-cyclical. Abraham and Haltiwanger (1995) conclude, "*Correcting for all of the measurement problems, estimation problems, and composition problems does not lead to a finding of systematically procyclical or counter-cyclical real wages.*" In contrast, evidence based individual-level longitudinal surveys (e.g. Bils, 1985; Solon *et al.*, 1994) supports the notion that wages are strongly pro-cyclical. This paper shows that another type of data disaggregation adds significant new insights into aggregate wage cyclicality. This concerns observing wage behavior in the frequency domain. We are interested in real hourly earnings in U.S. manufacturing where earnings are separated into those deflated by consumer prices and those by producer prices. We find that our frequency approach provides particularly valuable insights if we break down earnings into constituent parts. We provide a method of earnings decomposition into the standard hourly wage, the overtime mark-up, and the proportion of overtime workers in the total workforce.[1]

To appreciate the potential value-added of employing frequency methods, consider three cycles of relatively short, medium and long periods. Although by no means hard and fast or exhaustive representations, these might consist of (respectively) a wage contract cycle, a business cycle and a product cycle. Each type may associate systematically with the real wage. The relative strength and direction of the associations may differ, however. The start of a three-year United States wage contract, for example, may coincide with wage

[1]To obtain this breakdown we employ unpublished data relating to hours of work provided by the U.S. Bureau of Labour Statistics. These data are provided on an annual frequency and are available from 1959 to 1997. Further details on data sources and methods can be found in Appendix A.3.

adjustments designed to correct for unforeseen economic events at the previous negotiation time point. This process may be expected to generate a mix of pro- and countercyclical wage effects through time depending on the direction of deviations from expected outcomes. Additionally, the wage may respond positively to the business cycle. For instance, where compensation relates to marginal product, human capital investment may produce procyclical wages stemming from the fixity of the labor input. The wage may also associate positively with the product cycle. Top quality workers earning relatively high pay may be matched with new and innovative products with strong growth potential. As these products are eventually superseded by new innovations, wages may subsequently be associated with the hire of relatively poor quality and less well-remunerated workers.

Yet, all three cyclical effects will serve to condition a long time series of the real wage. This gives rise to a series of critical questions. Which, if any, is the frequency band dominating the cyclical behavior of the wage? If a given frequency dominates, what direction and strength of cyclicality does it exhibit? Does an association between cycles in a certain frequency band and the wage pattern represent a contemporaneous association or involve leads or lags? Pursuing such lines of enquiry leads to a more general question. Is the observed wage cyclicality in the frequency domain supportive of the general view arising from aggregate time series analysis or does it serve to modify that view? A seeming low correlation between the wage and a measure of the cycle may simply reflect the fact that the underlying time series is composed of a number of frequency bands between the two variables that are of different amplitudes and timing. Separately, one or more bands may display strong evidence of a systematic cyclical relationship. Taken together, countervailing influences may serve to mask underlying patterns.

The analysis of the wage's spectral representation allows us to tackle directly these issues since it can be decomposed into cyclical components defined over multiple economic cycle frequencies. The starting point is univariate analysis. A stationary time series can be broken down into superimposed harmonic waves of varying phases and amplitudes. To determine the length of the dominant cycle it is necessary to search the spectrum between the endpoints of the entire frequency interval and select the cycle that explains the greatest portion of the total variance of the wage. An apparent 5-year cycle may simply reflect the fact that the time series possesses a dominant cycle of that length. Alternatively, it may disguise the fact that there are two other underlying cycles - one longer and one shorter than 5 years - that combine to give the appearance of a prevailing 5-year cycle. Or there may be more than two underlying cycles. Investigating the frequency domain not only allows us to identify the number of cycles a series possesses but also to determine the contribution of each cycle to explaining the total variance of the wage and whether at this frequency the explained variance is significant.

However, the main economic interest behind the study of the frequency domain lies in an extended framework. We cannot discern whether observed cycles in the wage reflect underlying economic conditions unless we relate them to other variables that are both reflective of the ups and downs of economic activity and exhibit strong associations with the wage itself. Output, employment, unemployment, fixed capital formation as well as inventory and building investments have been variously adopted for this latter role. These economic indicators represent a range of cycles of differing phases and amplitudes. At one end of the spectrum, inventory investment is typically found to follow a three- to four-year cycle (Kitchin cycle). At the other, building investment cycles are around 20 years (Kuznetz cycle). In the middle, output often cycles in the five- to seven-year range (business cycle) while fixed capital formation is more likely to be in the seven- to ten-year range (Juglar cycle). Moving from univariate to multivariate analysis in the frequency domain necessarily involves describing associations of the wage with one or more of these economic cycles.

In this paper, we concentrate on average hourly real earnings.[2] We decompose real earning into three components: the standard wage, the premium mark-up, and the proportion of overtime workers. We apply spectral methods to each component part. This provides us with richer detail of wages over the cycle than previous studies that have tended to compare standard hourly wages with average hourly earnings. We show that not only do cyclical effects differ across earnings' components but that even a single component may co-vary with more than one representative measure of the economic cycle. We

[2]See the longitudinal microdata contribution of Devereux (2001). This argues the case for analyzing wage behavior over the cycle in this broader earnings context.

should expect this latter finding on a priori grounds. For instance, the level of the standard hourly wage rate typically reflects an agreement between (at least) two bargaining parties. The bargaining agenda is comprised of a variety of economic issues of concern to each side. The uppermost interest of workers in a given firm might be for wage changes to cover cost of living increases while management might place greatest weight on product demand. As a result, wage outcomes may be conditioned, to a greater or lesser extent, by the cyclical characteristics of representative proxies for each of these economic influences.

2 Econometric Method

2.1 Univariate Measure

In empirical research on economic cycles, the predominant paradigm has been to examine auto- or cross-covariances in the time domain.[3] The information on the cyclical structure contained in the autocovariance function can be transformed into frequency domain, revealing a more detailed picture. The *spectrum* of a process is defined as the Fourier transform of the autocovariance function $\gamma_x(\tau)$, $\tau = 0, \pm 1, \pm 2, ...$:

$$f_x(\omega) = \frac{1}{2\pi} \sum_{\tau=-\infty}^{\infty} \gamma_x(\tau) e^{-i\omega\tau}; \ \omega \in [-\pi, \pi]. \tag{1}$$

Figure 1 illustrates the plot of a spectrum.[4] The interpretation is like that of a probability density function; $f_x(\omega)d\omega$ is the part of the overall variance of X_t which is due to the component with frequencies over the interval $[\omega, \omega + d\omega]$. The total area under the spectrum equals the process variance:

$$\gamma_x(0) = \int_{-\pi}^{\pi} f_x(\omega) d\omega. \tag{2}$$

[3] Widely cited examples are Kydland and Prescott (1990) and Backus and Kehoe (1992). For a recent exception, working in the frequency domain, see A'Hearn and Woitek (2001).

[4] We estimate parametric spectra, i.e. we start in the time domain by fitting autoregressive models to the data. A detailed explanation of the estimation procedure can be found in the Appendix, Section A.1.1.

In other words, we can look at it as the plot of a decomposition of the variance against frequencies in the interval $[0, \pi]$.[5] After normalising the spectrum using the variance $\gamma_x(0)$, the area under the curve from ω_1 to ω_2 in Figure 1 is (half) the share of total variance of X_t which can be attributed to the composite of the waves in this range. Spectral analysis thus permits a natural decomposition of a series into cyclical components defined over frequency bands we are interested in. In terms of real wages, these may relate to 3- year wage contract cycles, 5-7 year business cycles and still longer cycles generated, for instance, by product and process innovations.

Figure 1: Spectrum

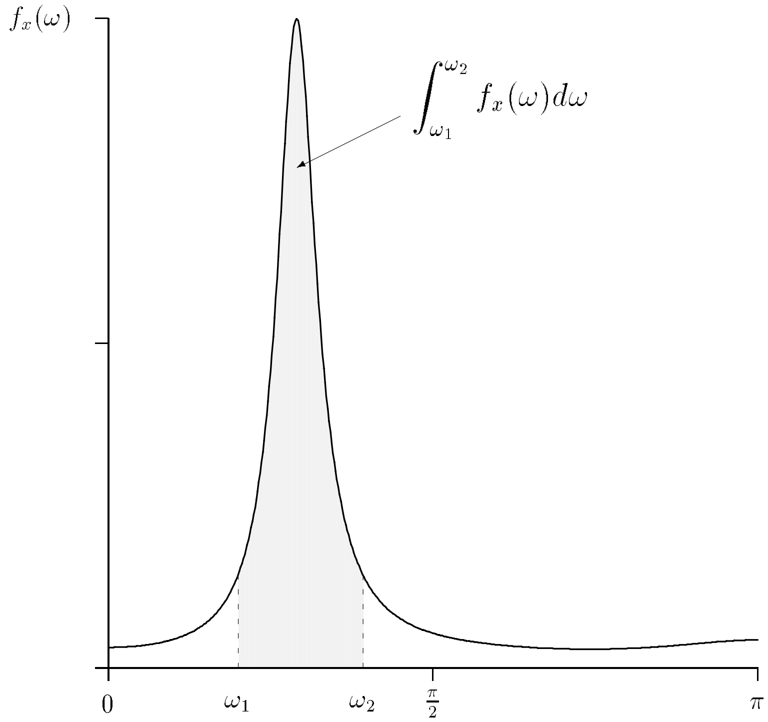

2.2 Multivariate Measures

Identifying each of the multiple wage cycles that combine to produce the observed wage time series does not in itself contain the most interesting information from an economist's viewpoint. To achieve this, we would need to establish what each of the wages represents, if anything, from an economic

[5] Since the spectrum is an even function, it is not necessary to plot it in the entire range $[-\pi, \pi]$.

perspective. The common procedure is to test whether or not a limited number of mainline cyclical indicators can explain significant degrees of the wage variation. One such indicator may be output deviations from trend, selected to capture business cycle activity. If output deviations are strongly correlated, say, with a 5-7 year wage cycle, and if that frequency range represents a dominant explanation of wage movements, we would be inclined to lean towards a business cycle explanation of wage cyclicality. Of course, we have a choice of proxies for the business cycle and each would be expected to offer greater or lesser co-variations with different wage cycles. Finding statistical support for business cycle effects may not complete the story, however. Shorter and longer frequency ranges, exhibiting strong correlation with other economic phenomena, may also add significantly to explained wage variation.

Suppose that the peaks and troughs of an influential constituent cycle of the wage time series coincide with the respective turning points of the selected business cycle measure. Then we would conclude that the wage is both procyclical and in phase with the cycle. But the two series may be highly procyclical and out of phase. For example, in common with time series analyses, adjustment impediments associated with bargaining may lead to consistent phase lags of the wage to the cycle. Or the two series may be partly in phase and partly out of phase.

We apply and develop frequency domain techniques to offer detailed insights into these aspects of wage cyclicality. We first consider 'explained variance' from a frequency domain perspective. This is achieved via the *squared coherency* measure (sc) which assesses the degree of linear relationship between cyclical components of two series X_t and Y_t, frequency by frequency. The sc is defined as

$$sc(\omega) = \frac{|f_{yx}(\omega)|^2}{f_x(\omega) f_y(\omega)}; \quad 0 \leq sc(\omega) \leq 1, \tag{3}$$

where $f_x(\omega)$ is the spectrum of the series X_t, and $f_{yx}(\omega)$ is the cross-spectrum for Y_t and X_t.[6] Using this expression, we can decompose $f_y(\omega)$ into and ex-

[6]The spectrum is the Fourier transform of the autocovariance function, and the cross-

plained and an unexplained part,

$$f_y(\omega) = \frac{|f_{yx}(\omega)|^2}{f_x(\omega)f_y(\omega)} f_y(\omega) + f_u(\omega) = \\ = sc(\omega)f_y(\omega) + f_u(\omega). \tag{4}$$

Integrating equation (4) over the frequency band $[-\pi, \pi]$ gives

$$\underbrace{\int_{-\pi}^{\pi} f_y(\omega)d\omega}_{\gamma_y(0)} = \underbrace{\int_{-\pi}^{\pi} sc(\omega)f_y(\omega)d\omega}_{\text{“explained” variance}} + \underbrace{\int_{-\pi}^{\pi} f_u(\omega)d\omega}_{\tilde{\sigma}^2}. \tag{4'}$$

The first term on the right in equation (4′) is the product of squared coherency between X_t and Y_t and the spectrum of Y_t; the second term is white noise. This equality holds for every frequency band $[\omega_1, \omega_2]$. We can plot total variance (the area under the spectrum) and explained variance as shown in Figure 2. Comparing the area under the spectrum of the explained component to the area under Y's spectrum in a frequency interval $[\omega_1, \omega_2]$ yields a measure of the explanatory power of X, analogous to an R^2 in the time domain. In contrast to sc however, R^2 is constant across all frequencies.

---

spectrum is the Fourier transform of the cross-covariance function:

$$f_{yx}(\omega) = \frac{1}{2\pi} \sum_{\tau=-\infty}^{\infty} \gamma_{yx}(\tau)e^{-i\omega\tau}; \ \omega \in [-\pi, \pi].$$

Figure 2: Explained Variance

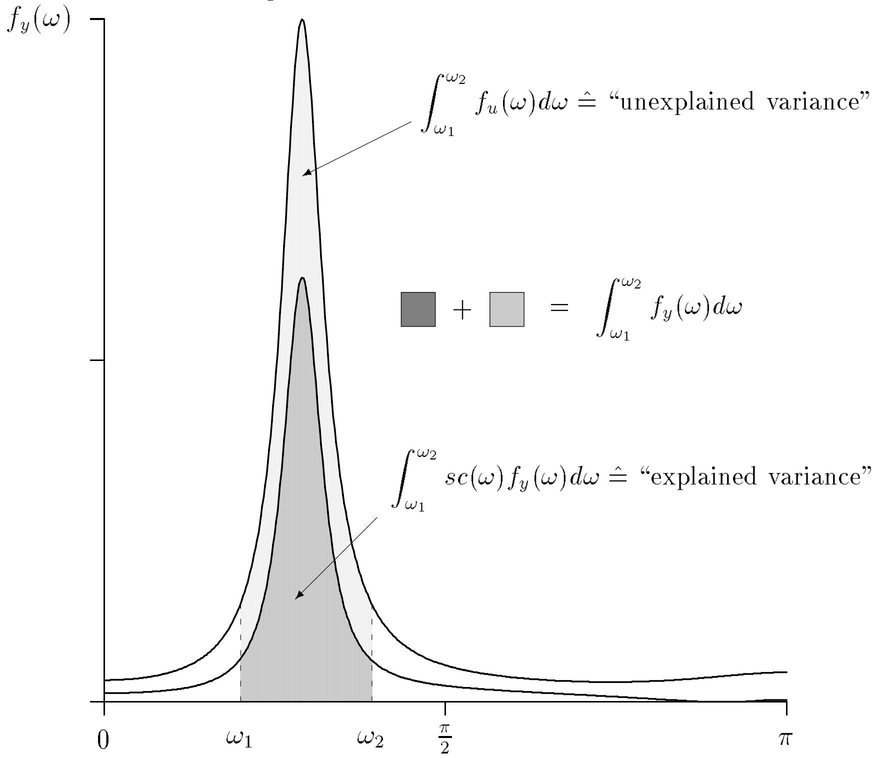

We now turn to the concepts of phase shift and dynamic correlation. Identifying lead-lag relationship between the series Y_t and X_t in time domain is carried out using the cross correlations at lags $\tau = \pm 1, \pm 2, \ldots ,$. In contrast, in the frequency domain, this can be achieved frequency by frequency using the *cross spectrum*. The cross spectrum, which is the Fourier transform of the covariance function of Y_t and X_t, is given by

$$f_{yx}(\omega) = c_{yx}(\omega) - iq_{yx}(\omega), \tag{5}$$

where $c_{yx}(\omega)$ is the *cospectrum* and $q_{yx}(\omega)$ is the *quadrature spectrum*. It can be used to derive the *phase spectrum* defined as

$$\phi_{yx}(\omega) = -\arctan(q_{yx}(\omega)/c_{yx}(\omega)). \tag{6}$$

The phase spectrum at frequency ω measures the lead of the cyclical component of Y_t at this frequency over the corresponding component of X_t. It can be interpreted as the negative of the angle which would transform the component in X_t into the best linear approximation of Y_t. To facilitate an intuitive interpretation, in Section 4, we present the phase shift relative to the relevant cycle range.

As pointed out by Croux *et al.* (2001), a measure like the squared coherency presented above is not suited for analysing the comovement of time series, because it does not contain information about possible phase shift between cycles in the series X_t and Y_t. In this sense, the correlation coefficient in time domain is more informative, since it is calculated lag by lag, providing both information on the lead-lag structure and the degree of linear relationship between the two series. Croux *et al.* (2001) propose an alternative measure, the so-called dynamic correlation $\rho(\omega)$, which measures the correlation between the "in-phase" components of the two series at a frequency ω:

$$\rho(\omega) = \frac{c_{xy}(\omega)}{\sqrt{f_x(\omega) f_y(\omega)}}; \quad -1 \leq \rho(\omega) \leq 1. \tag{7}$$

3 Real hourly earnings decomposition and the cycle

In the univariate analysis of real earnings within the frequency domain, we can ascertain the lengths of cycles that are most closely associated with component parts of the earnings measure. This gives vital pointers to the types of economic indicator to work within multivariate applications. The latter involve exploring the degrees to which each wage component co-varies and is in phase with selected cyclical indicators. In this section we concentrate on earnings decomposition itself and suggest the range of economic indicators that might apply on a priori grounds. We deal first with nominal earnings decomposition and then we consider the choice of price deflator in the real series.

3.1 Nominal earnings decomposition

A critical feature of earnings decomposition concerns the distinction between the hourly standard wage rate and overtime payments. Overtime has been an important recent phenomenon in the United States. During the 1980s and 1990s, the proportion of overtime workers in manufacturing grew to 40 percent of the workforce. From the early 1990s trough to early 1997, average weekly overtime in manufacturing increased by 1.6 hours to reach 4.9 hours, the highest since the Bureau of Labor Statistics first recorded these data in 1956 (Hetrick, 2000).

Average nominal hourly earnings, A, can be defined as a geometric average; that is

$$A_t = E_t^{\lambda_t} W_t^{1-\lambda_t} \tag{8}$$

where, E is average hourly earnings of overtime workers, W is the average standard hourly wage rate, and λ is the proportion of the total workforce working overtime. Additionally, given Fair Labor Standard Act regulations that set maximum weekly standard hours at 40 and the minimum overtime premium at 1.5, E can be expressed as

$$E_t = \mu_t W_t, \tag{9}$$

where μ is the mark-up required to convert the average standard wage to average wage earnings. Specifically, the mark-up is given by

$$\mu_t = (40 + 1.5V_t) \, / \, (40 + V_t) \tag{10}$$

where V is average weekly overtime hours of overtime workers.

Substituting (9) into (8) and taking logs gives

$$\ln A_t = \ln W_t + \lambda_t \ln \mu_t. \tag{11}$$

If $\lambda = 0$, all workers receive the average standard wage rate and $A = W$ in (8). If $\lambda > 0$, then $A > W$ due to the fact that a proportion of weekly hours for a proportion of workers are compensated at the hourly premium rate.[7]

The extent to which A diverges from W depends in (11) on (a) the size of the average premium mark-up, μ and (b) the proportion of overtime to total workers, λ. Essentially, changes in μ are solely dependent, as shown in (10), on changes in average overtime hours of overtime workers, V. For the great majority of workers, the maximum number of weekly hours before premium

[7]In Appendix A, we outline how we derive and construct (8) and the subsequent wage expressions from Bureau of Labor Statistics data. In our data set, movements in geometric and arithmetic average wage earnings correspond almost identically, e.g. the simple correlation is 0.99.

rates apply together with the level of the premium are fixed by legislation. By contrast, changes in λ in (11) are employment driven, dependent on the proportion of workers choosing to undertake overtime. Together, these variables recognise an essential overtime breakdown underlined in the work of Trejo (1993) dealing with union effects on overtime working.

Why is it important to differentiate between W, μ and λ in (11)? Standard and overtime components in the wage may well differ with respect to the cyclical indicator with which they most strongly co-vary as well as the degree to which they are in- phase with their dominant cyclical influence. Suppose at the trough of an output or employment cycle firms hold underutilized labor stock. The initial phase of the upturn may be dominated by labor dishoarding which involves effective hours being bought into line with paid-for hours. Planned increases in the stock of employment will depend on the anticipated time required to restore normal work intensity together with the expected degree and length of the ensuing growth period. In the later phases of the upturn, employment adjustment lags associated with search, hiring, and training may require firms to resort to temporary hours increases as an employment buffer (Nadiri and Rosen, 1973).

Overtime hours would be expected to feature prominently in this latter adjustment process. The degree to which the firm will extend overtime working will depend on the relative costs of alternative buffers, such as a greater than planned run-down of inventories (Topel, 1982). To the extent that overtime is used to offset shortfalls in planned employment growth - especially at times of exceptional demand peaks - overtime cycles are likely to be relatively short and, given potential substitution, may well correlate strongly with changes in business inventories. Moreover, since premium rates apply automatically to changes in overtime hours - with no pay negotiation involved - we would expect that, to a large degree, overtime pay would be in-phase with the cycle (see the discussion in Hall and Lilien, 1979).

Overtime adjustments occur not only through hours adjustments, μ, but also through employment adjustments, λ. Whether or not these two variables are close substitutes is an empirical question. Speedier responses may be achieved through changing the hours of workers already committed to overtime rather than through persuading marginal workers to move in and out of overtime completely. In this case, μ rather than λ is more likely to represent the in-phase buffer response to high demand discussed above. Changes in λ, by contrast, may represent a longer term restructuring of work organisation.

The view taken in the bulk of the existing literature is that the standard nominal wage, W, is likely to be influenced by economic conditions occurring over the entire spectrum of the business cycle. A firm's wage contract negotiations are undertaken, typically, at regular and pre-determined intervals. Even allowing for consideration of anticipated and historic economic extrapolations, nominal wage changes are likely partially to reflect prevailing short-run economic conditions at all stages of business activity. For this reason, economists have tended to use cyclical changes in output or employment or unemployment to represent such business cycle effects. But the degree of contemporaneous association with the cycle may be limited. Economic forecast errors, communications problems involving asymmetric information and time delays between contract negotiation and implementation may serve to produce pro-cyclical wage changes that are out of phase with the actual business cycle. Added to this, any given cycle indicator may itself respond less than immediately to actual business activity. Employment- and unemployment- related indicators are well known to be particularly prone to this type of problem.

3.2 Real earnings and the choice of price deflator

Time series analysts have found that the choice of price deflator has a strong bearing on the observed degree of cyclical wage responsiveness (for example Abraham and Haltiwanger, 1995). Generally, wages deflated by consumption prices, Cp, are found to be more pro-cyclical than wages deflated by production prices, Pp. We label these, respectively, consumption and production wages. This comparative observation appears to be confined to real wage responsiveness to the business cycle. However, it is not necessarily the case that business-related cycles are the only significant determinant of changes in prices, and especially producer prices.

Producer prices refer to the entire marketed output of firms. They include goods and services purchased by other firms as production inputs or capital investments. In these latter respects, and following Caballero and Hammour (1994), suppose that some firms treat a given recessionary period as a time to scrap outdated capital and to invest in product and process innovations. The costs of such cleansing activities will be offset by expected returns accruing to improved growth performance as the new goods and operational methods provide a competitive edge. For at least two reasons, these investment periods may mark the start of wage cycles that are linked to the evolution of new products and processes. Relatively high wages will be associated with the entry phases of product cycles. First, firms will wish to match high quality workers with their newly introduced and relatively high value added activities. Second, they will endeavor to ensure investment returns by introducing payment structures designed to encourage work application and effort. Through time, innovations will be superseded during later investment periods when new products and processes are established within the same firms and/or competing firms. Productive and competitive advantages will be eroded and relative wage growth will decline.

Product and process innovations involve investments in buildings and equipment and a key cyclical indicator for the latter is changes in fixed capital formation. To the extent that producer wages follow cycles of capital formation then this may tend to detract from the significant influences of conventional business cycle measures - such as output and employment changes - which are, on average, of shorter duration.

Another potentially important issue concerning the cyclical behaviors of Cp and Pp is the length of adjustment lag between actual and desired prices. For example, delayed price changes on the producers' side may involve component parts supplied to assembly plants by subcontractors that are based on fixed prices over specified contract periods. As for consumers, make-to-order companies will supply products at a future date, but at currently specified prices. It is difficult, a priori, to form expectations about the relative lag-lengths of Cp and Pp but it is important that the methodology allows us to test for systematic response differences between consumption and production wages.

4 Findings

Our empirical research strategy is to build out from the foundation of the univariate analysis of wave decomposition. As discussed in section 2, univariate methodology in the frequency domain allows us to establish (a) the number of different cycles that significantly influence earnings components and (b) the average lengths of these cycles. This helps us to determine the choice of cyclical economic indicators to be used in the multivariate analysis. We are then in a position to examine the degree to which the frequency of each component of wage earnings coheres with each of the selected indicators. We then proceed to the dynamic correlations discussed in Section 2.2. These allow us not only to measure if a wage component and a cyclical indicator exhibit pro-, counter-, or a-cyclical relationships but also the time lengths involved if significant lags or leads are present.

Before we can analyse the cyclical structure, we need to ensure that the series are stationary. Since all filtering methods in the literature potentially distort the cyclical structure dependent on the true data generating process,[8] our method of choice is the strategy proposed by Canova (1998): we compare the results for different filtering methods, and judge the robustness of the outcome. The filters are the Hodrick-Prescott (HP) filter (Hodrick and Prescott, 1997),[9] the Baxter-King filter (Baxter and King, 1999) in a modified version proposed by Woitek (1998) (BKM),[10] and the difference filter (D). In the following, we report the results for the BKM filter.[11]

As discussed in relation to Figure 1, we can decompose a series into cyclical components, defined over multiple frequency bands. The variables included in the first column of Table 1 are the constituent parts of (8) expressed in real terms. We show results using both Cp and Pp deflators. The next three columns show the share of the total variance of each wage series that is explained by the composite of waves in the respective frequency range. Taking the premium markup, μ, as an example, 0.01, 0.07 and 0.59 of the total vari-

[8]Recently it was demonstrated by Cogley and Nason (1995), King and Rebelo (1993) and Harvey and Jaeger (1993), that the widely used Hodrick-Prescott filter (Hodrick and Prescott 1997) is likely to generate spurious cyclical structure at business cycle frequencies if applied to difference stationary series. Similar points can be made with respect to the Baxter-King Filter (Guay and St-Amant 1997), and to moving-average filters in general (Osborn 1995). Moreover, there is the danger of spurious correlation between Hodrick-Prescott filtered series (Harvey and Jaeger 1993).

[9]With the usual smoothing weight of 100 (annual data) and for a series which is I(0), the HP filter leaves cycles with period length up to about 11 years almost undistorted in the data.

[10]The modified Baxter-King Filter uses Lanczos' σ factors to deal with the problem of spurious side lobes, which invariably arises with finite length filters. In contrast to the original filter proposed by Baxter and King (1999), our cut-off period is 15 years, allowing us to analyse cycles as long as the Juglar cycle.

[11]The results for all the other filters are comparable and will be made available on request.

ance earnings is explained by the composite waves in the (1) 7-10, (2) 5-7 and (3) 3-5 year frequency ranges, respectively. Only the third band is found to yield a significant share of the total variance. The final column of Table 1 gives the cycle length at which the explained variance is maximised across the entire spectrum for each series, e.g. μ has a cycle length of 4.18 years.

Table 1: Earnings: share of total variance

	(1)	(2)	(3)	Cycle
μ	0.01	0.07	0.59^{***}	4.18
λ	0.08	0.64^{***}	0.22	5.94
W/Cp	0.14	0.43^{***}	0.33	5.82
W/Pp	0.24^{**}	0.41^{***}	0.25	6.74

Notes:
(i) W: standard hourly wage; μ: premium markup; λ: proportion of workers working overtime; Pp: producer price index; Cp: consumer price index; (ii) share of total variance: share of variance which can be attributed to the composite of waves in the respective frequency range; (iii) (1): 7-10 years (Juglar cycle), (2): 5-7 years, (3): 3-5 years (Kitchin cycle); (iv) */**/***: share of total variance is significant at the 10/5/1 per cent level.

As alluded to above, statistically determining the cyclical length of underlying cycles in a series amounts to testing the null hypothesis of no cyclical structure or in other words that the series is white noise.[12] Accordingly, if the explained variances reported in columns 1-3 are significant then the series can

[12]To establish significance we follow Reiter and Woitek (1999) and simulate white-noise processes to assess whether the share of total variance in the frequency intervals of interest is significantly different from the result we would obtain if the data generating process was white noise. For example, we fit an AR model of order 5 to a white noise process, which has the same variance as the series under analysis, and repeat this 2000 times. We then use the univariate spectral measures from this experiment to derive the empirical distribution under the null hypothesis (i.e. no cyclical structure). Note that we employ empirical distributions since asymptotic distributions are extremely difficult to derive in this context. In any event, since we employ relatively short time series, the asymptotic properties would most certainly not be correct. In addition to the white-noise processes, we also used filtered random walks as null models, to ensure that the cyclical structure is not created by an inappropriately chosen filter. This procedure basically produced the same results.

be said to have cyclical structure of the length specified.

Our findings are as follows. First, the cycle of the premium mark-up, μ, is predominantly explained within the shortest (3-5 year) range. Second, the distribution of explained variance with respect to the proportion of overtime workers, λ, is similar to the Cp-deflated wage cycle, although the 5-7 year range for λ, is relatively more dominant than W/Cp. Third, the production wage has two significant cycles that fall into the 5-7 and 7-10 year ranges, with the dominant cycle being in the 5-7 year range. Fifth, the production wage cycle length is the longest followed by the proportion of overtime workers, the consumption wage cycle and the premium markup.

One or more of the frequency wave ranges shown in Table 1 displays a significant share of total variance association with each component of the wage. We selected four economic indicators to represent these ranges. First, gross fixed capital formation ($GFCF$) was chosen to capture the longer end of the time spectrum.[13] Moreover, as we argue in Section 3.2, it is a potentially appealing cyclical indicator in relation to the production wage. Second, output (Y) and employment (N) represent the middle time band. Output- and employment-based indicators are typically used to represent the business cycle. Finally, the change in business inventories, labeled inventory investment (II), is taken to represent cycles of relatively short duration. We argued in Section 3.1 that this may well provide an ideal indicator of overtime hours cyclicality, the driving variable in the overtime premium mark-up. Using the same methodology as Table 1, the results in Table 2 for the four indicators show the following. With regard to significant contributions to total variance, the cycle lengths in ascending order are II, N, Y and $GFCF$ (see last column, Table 2). II has one significant cycle in the 3-5 year range. N, Y and $GFCF$ each have one significant cycle in the 5-7 year range. Additionally, $GFCF$ displays a second significant cycle in the 7-10 year range.

[13]For the G7 countries, $GFCF$ is dominated by a long cycle in the 7-9 years range, with the exception of a slightly shorter cycle in the US series (Woitek, 1996), which is in line with the results in Table 2. This phenomenon can be explained by the existence of built-in-stabilisers in the US economy which are absent in European countries (Romer, 1999).

Table 2: Economic indicators: share of total variance

	(1)	(2)	(3)	Cycle
$GFCF$	0.17*	0.56***	0.20	6.58
Y	0.08	0.49***	0.33	5.94
N	0.09	0.43***	0.36	5.75
II	0.03	0.08	0.54***	3.68

Notes:
(i) $GFCF$: gross fixed captial formation, Y: output, N: employment, II: inventory investment; (ii) share of total variance: share of variance which can be attributed to the composite of waves in the respective frequency range; (iii) (1): 7-10 years (Juglar cycle), (2): 5-7 years, (3): 3-5 years (Kitchin cycle); (iv) $\star/\star\star/\star\star\star$: share of total variance is significant at the 10/5/1 per cent level.

Analogous to the univariate approach, we require a multivariate method that allows us to achieve two objectives. First, we want to find out at which business cycle frequency the ratio of explained to unexplained variance is at a maximum. Second, we need to determine whether the share of variance explained by our various indicators of the cycle, in a specific frequency band, is significant.[14] In other words, we would like to test the null hypothesis that the real wage component(s) and the cyclical indicators are unrelated in a specific frequency band. As in the univariate case the data may reveal multiple cycles between any two series. The cells of Table 3 refer to the proportion of total variance in the respective frequency range for each component of the real wage explained by the variance of II, N, Y and $GFCF$ respectively. Consider μ in Table 3 using the $GFCF$ cycle in conjunction with Table 1. This reveals that

[14]To determine whether the explained variance, ρ_{XY} between two series Y and X, in the relevant frequency band $[\omega_1, \omega_2]$, is significantly different from zero we implement the following procedure. First, we fit AR models to Y and X and, second, we conduct a parametric bootstrap to simulate the model under the null hypothesis (i.e. no interaction between the series). This produces a simulated series $\left(Y_t^S X_t^S\right)$ that has the univariate characteristics of the underlying data, but without interaction. Third, we fit a VAR of fixed order to $\left(Y_t^S X_t^S\right)$ and calculate ρ_{XY}^S. Fourth, these steps are then repeated for $s = 2000$ so that we can obtain an empirical distribution of ρ_{XY} under the null conditional on the series we are examining. Note that Priestley (1981, p705-706) develops a similar test of zero coherency for the classical spectral estimate, the periodogram.

0.00, 0.04 and 0.32 of the total variance of μ is explained by the total variance in $GFCF$ in, respectively, the (1) 7-10, (2) 5-7 and (3) 3-5 year frequency ranges. Although statistical significance is established in the 7-10 and 5-7 year ranges, the magnitudes are not economically meaningful.

More generally, the results in Table 3 suggest the following. The overtime mark-up, μ is the only wage component that associates significantly with changes in inventories, II. Without any doubt, most of its explained variation occurs within the shortest 3-5 year range. This variable is dominated by changes in overtime hours and we speculated earlier that the inventory variable offers a suitable choice of indicator. While other ranges are found to be significant, they account for very little of the total variation of the premium mark-up.

The proportion of overtime workers, λ, displays strong associations with the middle bands of the $GFCF$, Y and N indicators. More than any other wage component, λ reveals the usefulness of testing associations over a range of cycle ranges. Both long- and, especially, short- cycle ranges also significantly co-vary with this wage component.

The dominant cyclical influence on the consumption wage occurs within the middle (5-7 year) time band. This is true for three of the cyclical indicators - i.e. the $GFCF$, Y and N. These findings are consistent with the univariate results. Note that, while explaining considerably less of the total variance, the longest (7-10) range is also significant for the $GFCF$ and Y indicators. Again, this establishes that more than one length of cycle may significantly influence wage earnings components.

Finally, the production wage relates exclusively to the 5-7 range and is confined only to the $GFCF$ indicator. Few earlier studies have been able to link producer wages to the business cycle. These results point to the view - as discussed in Section 3.2 - that such wages may be more appropriately explained by different types of economic cycle. In particular, we highlighted a potential connection with movements in fixed capital formation.

Table 3: Real earnings and economic cycles: explained variance

	GFCF			Y			*N*			*II*		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
μ	0.00*	0.04**	0.32	0.00	0.05***	0.47**	0.00	0.04**	0.47**	0.01*	0.04	0.34**
λ	0.05**	0.58**	0.17***	0.07***	0.62***	0.18***	0.07***	0.62***	0.19***	0.06	0.62	0.14
W/Cp	0.10**	0.36**	0.08	0.09*	0.37**	0.13	0.06	0.34**	0.10	0.09	0.34	0.06
W/Pp	0.12	0.31*	0.06	0.09	0.30	0.13	0.06	0.25	0.11	0.08	0.25	0.03

Notes:
(i) W: standard hourly wage; μ: premium markup; λ: proportion of workers working overtime; Pp: producer price index; Cp: consumer price index, $GFCF$: gross fixed captial formation, Y: output, N: employment, II: inventory investment; (ii) explained variance: share of variance explained by variance of cycle measure in the respective frequency range; (iii) (1): 7-10 years (Juglar cycle), (2): 5-7 years, (3): 3-5 years (Kitchin cycle); (iv) */**/***: share of total variance is significant at the 10/5/1 per cent level.

Table 4: Dynamic correlation at real earnings cycle frequencies

		$GFCF$		Y		N		II	
	period	phase	ρ_c	phase	ρ_c	phase	ρ_c	phase	ρ_c
μ	4.18	0.01	0.82^{***}	0.01	0.94^{***}	0.05	0.86^{***}	0.06	0.74^{***}
λ	5.94	0.02	0.49^{**}	0.03	0.86^{***}	0.04	0.56^{***}		
W/Cp	5.82	0.04	0.71^{***}	0.12	0.60^{**}	0.18	0.25		
W/Pp	6.74	−0.09	0.41^{*}						

Notes:
(i) W: standard hourly wage; μ: premium markup; λ: proportion of workers working overtime; Pp: producer price index; Cp: consumer price index, $GFCF$: gross fixed captial formation, Y: output, N: employment, II: inventory investment; (ii) ρ_c: in-phase correlation, phase shift: relative to cycle length; (iii) $^{*}/^{**}/^{***}$: dynamic correlation is significant at the 10/5/1 per cent level.

While we have found that standard real wages as well as μ and λ are pro-cyclical, we have yet to ascertain whether or not they are in phase with the various cyclical indicators. In order to establish the degrees to which phase shifts are important, we estimated the dynamic correlations expressed in equation 7. Our procedure was as follows. Indicators were selected if they displayed significant associations - for one or more lengths of cycle - with a wage component in Table 3. Thus, for example, correlations with respect to all four indicators were undertaken in relation to μ while only one correlation (with respect to GFCF) was estimated for W/Pp. The estimated correlations are presented in Table 4 and, perhaps surprisingly, they indicate a predominance of significant leads between wage components on the various indicators. Taken alone, however, these results are misleading. We are able to translate phase shifts into months using the cycle lengths (labelled 'periods' in Table 4) at which the explained variances of the wage components are maximised (see last column of Table 1). This is achieved simply by multiplying period $\times$ relative phase $\times$ 12. In all cases, these calculations reveal phase shifts of less than one year. In effect, the variables are contemporaneously related at the annual frequencies.

5 Conclusions

Frequency domain techniques establish that U.S. wage earnings are markedly pro-cyclical. This is true for all wage components. Moreover, while the earnings components are not strictly in phase with the dominant cycles, we do not detect leads or lags that extend beyond one year's duration. Beyond these general conclusions, studying the frequency domain permits more detailed insights into the cyclical behavior of earnings.

In the first place, both univariate and multivariate findings indicate that each component measure of the wage may display significant co-variations with more than one cycle measure in different frequency ranges. It turns out that there is typically a dominant range in terms of explaining total wage variation within each component. But findings of two significant ranges are quite common and, occasionally, we find three significant ranges. For example, it is not misleading to claim that a cycle with a 5-7 year length associates most strongly with the cyclical movements of the real standard hourly wage. It is misleading, however, to treat it as the only significant range association.

Secondly, the fact that earnings components respond to a range of frequency ranges suggests that, in the multivariate analysis of wage cyclicality, use should be made of more than one economic indicator of the economic cycle. For longer cycles, fixed capital formation offers a useful cyclical proxy to investigate. This associates relatively strongly with producer wages. For shorter cycles, inventory investment is found to relate strongly to the hours-dominated measure of the wage premium. In between lie the more familiar output and employment measures and these associate particularly strongly with consumer wages and with the proportions of overtime workers. Reliance on one representative indicator certainly provides only a partial insight into cyclical forces acting on wage earnings.

A Appendix

A.1 Estimation of the Spectrum

A.1.1 Autoregressive Spectra

To estimate the spectra, we fit autoregressive models in the time domain, and calculate the spectra of the estimated models. This method is based on the seminal work by Burg (1967), who shows that the resulting spectrum is formally identical to a spectrum derived on the Maximum Entropy Principle. This is seen to be a more reasonable approach then the normally used periodogram estimator. The periodogram employs the assumption that all the covariances outside the sample period are zero. Given that economic time series are notoriously short, this seems to be a problematic assumption[15] Consider a univariate AR model of order p, with residual variance σ^2. The spectrum is given by

$$f(\omega) = \frac{1}{2\pi} \frac{\sigma^2}{\left|1 - \sum_{j=1}^{p} \alpha_j e^{-i\omega j}\right|^2}; \ \omega \in [-\pi, \pi]. \tag{A1}$$

Equation (A1) is the analogue to the univariate spectrum in equation (1). With a VAR model of order p, the spectral density matrix is given by

$$\mathbf{F}(\omega) = \frac{1}{2\pi} \mathbf{A}(\omega)^{-1} \mathbf{\Sigma} \mathbf{A}(\omega)^{-\star}; \ \omega \in [-\pi, \pi]. \tag{A2}$$

$\mathbf{\Sigma}$ is the error variance-covariance matrix of the model, and $\boldsymbol{A}(\omega)$ is the Fourier transform of the matrix lag polynomial $\mathbf{A}(L) = \mathbf{I} - \mathbf{A}_1 L - \cdots - \mathbf{A}_p L^p$.[16] The diagonal elements of this matrix are the analogue to the univariate spectrum in equation (1), and the off-diagonal elements are the cross-spectra defined in footnote 7.

A.2 Modified Baxter-King Filter

Baxter and King (1999) construct a bandpass filter of finite order K which is optimal in the sense that it is an approximate bandpass filter with trend-reducing properties and symmetric weights, which ensure that there is no phase

[15] See the discussion in (Priestley, 1981, p. 432, 604-607). A recent applications to economic time series is A'Hearn and Woitek (2001).

[16] L is the backshift operator; the superscript '$\star$' denotes the complex conjugate transpose.

shift in the filter output. In time domain, the impact of the filter on an input series y_t is given by the finite moving average. In the frequency domain, the filter is characterised by its Fourier transform $\alpha(\omega)$.[17] To find the weights a_j, one solves the minimisation problem

$$\min_{a_j} Q = \int_{-\pi}^{\pi} |\beta(\omega) - \alpha(\omega)|^2 \, d\omega, \text{ s.t. } \alpha(0) = 0; \tag{A3}$$

where $|\beta(\omega)|$ is the "ideal" filter gain with cut-off frequencies ω_1 and ω_2.[18] The constraint ensures that the resulting filter has trend reducing properties.[19]

Solving the minimisation problem leads to the following results:[20]

$$\begin{aligned} a_j &= b_j + \theta; \ j = 0, \pm 1, \ldots, \pm K; \\ b_j &= \begin{cases} \dfrac{\omega_2 - \omega_1}{\pi} & \text{if } j = 0 \\ \frac{1}{\pi j}(\sin \omega_2 j - \sin \omega_1 j) & \text{if } j = \pm 1, \pm 2, \ldots \end{cases}; \\ \theta &= \frac{-\sum_{j=-K}^{K} b_j}{2K+1}; \end{aligned} \tag{A4}$$

The original Baxter-King filter has an undesireable property, which is known as *Gibb's phenomenon*, due to the fact that the ideal filter, which is a discontinuous function of ω, is approximated by a finite Fourier series. This approximation leads to side lobes in the gain function of the filter. (Priestley 1981, p. 561-3, Koopmans 1974, p. 187-9). While the relative contribution of some components for the overall variance of the series is exaggerated (i.e. they are multiplied by a gain greater than 1), other components are suppressed (i.e. they are multiplied by a gain less than 1).

An obvious solution to this problem is to increase the filter length. But since we are restricted by the limited availability of economic data, there is not much to be gained from changing the length of the filter. A more appropriate solution is to apply spectral windows. As an example, consider the so called

[17]See e.g. Koopmans (1974), p. 165 ff.

[18]The gain of a filter measures the change in the amplitude of the input components if transformed by the filter. The ideal bandpass filter gain $|\beta(\omega)|$ takes the value 1 in the frequency interval $[\omega_1, \omega_2]$ and 0 outside this interval.

[19]In order to remove the component with the frequency $\omega = 0$ from the series, the filter weights must sum to zero.

[20]The filter is symmetric (i.e. $a_j = a_{-j}$), and therefore does not impose a phase shift on the output.

Lanczos's σ factors (Bloomfield 1976, p. 129-137). We replace the truncated weights of the optimal filter b_j in equation (A4) by the modified weights $b_j^\star$, which are obtained from

$$b_j^\star = b_j \frac{\sin\left((2\pi j)/(2K+1)\right)}{(2\pi j)/(2K+1)}; |j| = 1, \ldots, K. \tag{A5}$$

After this step, the modified filter weights of the Baxter-King filter $a_j^\star$ can be calculated as demonstrated above (Woitek, 1998).

A.3 Derviation of earnings expression (7) from BLS data

The BLS calculate the *earnings rate*[21] (i.e. average hourly earnings), by dividing gross payrolls, GP, by total hours, thus

$$A_t = \frac{GP_t}{N_t H_t}. \tag{A6}$$

We can decompose equation (A6) by differentiating between overtime workers and workers working only standard hours. Under the FLSA, overtime is compensated at a premium rate for hours in excess of 40 per week. We assume that overtime workers are compensated for 40 weekly hours at the standard rate and then at the mandated premium rate for additional weekly hours.[22]

Accordingly, we may re-express equation (A6) in the form

$$A_t = \frac{N_t W_t H_t^s + N_t^o W_t (40 - H_t^s) + N_t^o W_t^o V_t}{N_t H_t^s + N_t^o (40 - H_t^s) + N_t^o V_t} \tag{A7}$$

where, $N^o \leq N$ are the number of employees working overtime, W and W^o are the average standard and average overtime hourly wage rates, H^s is average weekly standard hours of non-overtime workers, $V = (H^o - 40)$ is average overtime hours of overtime workers, and H^o is average total weekly hours worked by overtime workers. The numerator of (A7) comprises three parts.

[21] See, BLS Handbook of Methods, 1997, Ch. 2 Employment, Hours, and Earnings from the Establishment Survey.

[22] Actually, there is evidence (Trejo, 1993) that some overtime workers receive the premium before the 40 hour limit. Unfortunately, our data are such that we cannot accommodate this possibility.

The first term allows for all N workers to be paid at the standard rate, W for standard hours, H^s. These latter hours are averaged over non-overtime workers and we expect $H^s < 40$. Therefore, the second term allows for the fact that N^o overtime workers, assumed to be working 40 standard hours, are compensated at W for $(40 - H^s)$ hours. The final term shows that overtime workers are further compensated at the overtime rate, W^o for overtime hours, V. The maximum number of standard hours and the hourly overtime premium are fixed by legislation.

There are two problems with the definition of A in (A7): (i) the BLS Establishments Survey does not provide data on N^o and V; (ii) the arithmetic average used to calculate (A7) is additive and accordingly cannot be algebraically decomposed into its separate parts. We deal with each of these problems in turn.

The wage rate - i.e. W_t in (A7) - is approximated by BLS by adjusting average hourly earnings through the elimination of premium pay for overtime at a rate of time

$$W_t = \frac{GP_t}{N_t[\bar{H}_t^s + 1.5 \times \bar{H}_t^0]} \tag{A8}$$

where $\bar{H}^s(= H - \bar{H}^0)$ and $\bar{H}^0$ are, respectively, standard and overtime hours averaged over all workers (i.e. overtime and non-overtime workers). [23] No adjustment is made for other premium payment provisions such as holiday work, late shift work, and premium overtime rates other than those at time and one-half. W is calculated only for manufacturing industries because data on overtime hours are not calculated in other industries. This is the principal reason why we concentrate our attention on the manufacturing sector.

The BLS Current Population Survey does gather (unpublished) data pertaining to N^o.[24] Strictly, N^o defines the number of workers working in excess

[23] Note that (A6) follows the definition given by BLS (Handbook of Methods, 1997, Ch. 2, p 22): "[Average hourly earnings excluding overtime] ... are computed by dividing the total production payroll ... by the sum of the total production worker hours and one-half of the total overtime hours, which is equivalent to the payroll divided by standard hours."

[24] Basic information regarding the Survey and the published data can be found in BLS Handbook (Chapter 1) and the February 1994 issue of the BLS publication called Employment and Earning. There is a complication with these data. If, for example, a person worked 40 hours a week at a manufacturing job and then worked another 20 hours in the same week as a clerk in a store, that person would be shown as working 60 hours that week

of 40 weekly hours and, as mentioned previously, this is how we define overtime workers, that is

$$\lambda_t = \frac{N_t^o}{N_t} \tag{A9}$$

where λ_t is the proportion of workers working overtime at time t. Then, H^s in (A7), the number of standard hours worked by non overtime workers is given by

$$H_t^s = \frac{\bar{H}_t^s - \lambda_t \times 40}{1 - \lambda_t}. \tag{A10}$$

Further, V in (A7) is given by

$$V_t = \frac{\bar{H}_t^0}{\lambda_t}. \tag{A11}$$

The RHS of (A7) decomposes A into the contributions of total and overtime workers. To differentiate explicitly between workers working overtime and those working only standard hours while retaining A, we re-express real earnings as a geometric instead of an arithmetic average; thus our equation (8) in the text

$$A_t^\star = E_t^{\lambda_t} W_t^{1-\lambda_t} \tag{A12}$$

where, E is average hourly earnings of overtime workers and W and λ are defined in (A8) and (A9), respectively. Additionally, given FLSA regulations, E can be expressed in terms of (9) and (10) as in the main text, with V defined in (A11).

All employment, hours and earnings data are from the BLS Establishments Survey. The proportion of employees working greater than 40 hours per week are unpublished annual figures from the BLS Current Population Survey. The index of industrial production and price data are from the Federal Reserve Board and the BLS respectively.[25]

---

in their manufacturing job. We note that, while dual job holding is generally an important phenomenon, it is clear from the breakdowns of rates of dual job holding provided by Paxson and Sicherman (1996) that manufacturing occupations tend to exhibit below-average rates.

[25]The producer price index is for all commodities and the consumer price index is for

References

Abraham, K. G. and Haltiwanger, J. C. (1995), "Real Wages and the Business Cycle." *Journal of Economic Literature* **XXXIII**, 1215–1264.

A'Hearn, B. and Woitek, U. (2001), "More International Evidence on the Historical Properties of Business Cycles." *Journal of Monetary Economics* **47**, 299–319.

Backus, D. K. and Kehoe, P. J. (1992), "International Evidence on the Historical Properties of Business Cycles." *American Economic Review* **82**, 864–888.

Baxter, M. and King, R. G. (1999), "Measuring Business Cycles. Approximate Band-Pass Filters for Economic Time Series." *Review of Economics and Statistics* **81**, 575–593.

Bils, M. (1985), "Real Wages over the Business Cycle: Evidence from Panel Data." *Journal of Political Economy* **93**, 666–689.

Bloomfield, P. (1976), *Fourier Analysis of Time Series: An Introduction*. New York, London, Sidney, Toronto: John Wiley & Sons.

Burg, J. P. (1967), "Maximum Entropy Spectral Analysis." In: *Proceedings of the 37th Meeting of the Society of Exploration Geophysicists*, reprinted in Childers, D. G. (1978), *Modern Spectrum Analysis*. New York: IEEE Press.

Caballero, R. and Hammour, M. (1994), "The Cleansing Effect of Recession." *American Economic Review* **84**, 1075–1084.

Canova, F. (1998), "Detrending and Business Cycle Stylized Facts." *Journal of Monetary Economics* **41**, 475–512.

all urban consumers. Note that the latter (CPI-U-X1) is a BLS retrospectively calculated estimate of how the all items index would have moved (from 1967 to 1982) had the CPI been calculated using the flow of services approach to shelter. Prior to 1967 we use the official series reported by the BLS. In December 1982 the production CPI began to use flow of services approach. Hence as in other studies of wages (see Abraham and Haltiwanger, 1995) we employ the CPI, which is most consistent with the series currently reported by the BLS. We thank the BLS for providing the unpublished 1967-82 consumer price data as well as the Current Population Survey data.

Cogley, T. and Nason, J. M. (1995), “Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research.” *Journal of Economic Dynamics and Control* **19**, 253–278.

Croux, C., Forni, M., and Reichlin, L. (2001), “A Measure of Comovement for Economic Variables: Theory and Empirics.” *Review of Economics and Statistics* **83**, 232–241.

Guay, A. and St-Amant, P. (1997), “Do the Hodrick-Prescott and Baxter-King Filters Provide a Good Approximation of Business Cycles?”, Working Paper No. 53, Center for Research on Economic Fluctuations and Employment (CREFE).

Hall, R. E. and Lilien, D. L. (1979), “Efficient Wage Bargains under Uncertain Supply and Demand.” *American Economic Review* **69**, 868–879.

Harvey, A. C. and Jaeger, A. (1993), “Detrending, Stylized Facts and the Business Cycle.” *Journal of Econometrics* **8**, 231–247.

Hetrick, R. L. (2000), “Analyzing the Recent Upward Surge in Overtime Hours.” *Monthly Labor Review* **123**, 30–33.

Hodrick, R. and Prescott, E. (1997), “Postwar U.S. Business Cycles: An Empirical Investigation.” *Journal of Money, Credit and Banking* **29**, 1–16.

King, R. G. and Rebelo, S. T. (1993), “Low Frequency Filtering and Real Business Cycles.” *Journal of Economic Dynamics and Control* **17**, 207–231.

Koopmans, L. H. (1974), *The Spectral Analysis of Time Series*. New York, San Francisco, London: Academic Press.

Kydland, F. E. and Prescott, E. C. (1990), “Business Cycles: Real Facts and A Monetary Myth.” *Federal Reserve Bank of Minneapolis, Quarterly Review* **14**, 3–18.

Lucas, R. E. (1977), “Understanding Business Cycles.” In: K. Brunner and A. Metzler (Eds.), *Stabilization of the Domestic and International Economy, Carnegie-Rochester Series on Public Policy, Vol.5*, reprinted in Lucas, R. E. (1981), *Studies in Business-Cycle Theory*. Oxford: Basil Blackwell.

Nadiri, M. I. and Rosen, S. (1973), *A Disequilibrium Model of Demand for Factors of Production*. New York: National Bureau of Economic Research.

Osborn, D. R. (1995), "Moving Average Detrending and the Analysis of Business Cycles." *Oxford Bulletin of Economics and Statistics* **57**, 547–558.

Paxson, C. H. and Sicherman, N. (1996), "The Dynamics of Dual Job Holding and Job Mobility." *Journal of Labor Economics* **14**, 357 – 393.

Priestley, M. (1981), *Spectral Analysis and Time Series*. London: Academic Press.

Reiter, M. and Woitek, U. (1999), "Are there Classical Business Cycles?", University of Glasgow, Discussion Papers in Economics No. 9905.

Romer, C. D. (1999), "Changes in Business Cycles: Evidence and Explanations." *Journal of Economic Perspectives* **13**, 23–44.

Solon, G., Barsky, R., and Parker, J. A. (1994), "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?" *Quarterly Journal of Economics* **109**, 1–26.

Topel, R. H. (1982), "Inventories, Layoffs, and the Short-Run Demand for Labor." *American Economic Review* **72**, 769–787.

Trejo, S. J. (1993), "Overtime Pay, Overtime Hours, and Labor Unions." *Journal of Labor Economics* **11**, 253–278.

Woitek, U. (1996), "The G7-Countries: A Multivariate Description of Business Cycle Stylized Facts." In: W. Barnett, G. Gandolfo, and C. Hillinger (Eds.), *Dynamic Disequilibrium Modelling: Theory and Applications*, chap. 10, 283–309, Cambridge: Cambridge University Press.

Woitek, U. (1998), "A Note on the Baxter-King Filter.", University of Glasgow, Discussion Papers in Economics, No 9813.