

Measuring What Employers Really Do about Entry Wages over the Business Cycle

Pedro Martins, Queen Mary, University of London and CEG-IST

Gary Solon, Michigan State University

Jonathan Thomas, University of Edinburgh

June 2009

Abstract

In models recently published by several influential macroeconomic theorists, rigidity in the real wages that firms pay newly hired workers plays a crucial role in generating realistically large cyclical fluctuations in unemployment. There is remarkably little evidence, however, on whether employers' hiring wages really are invariant to business cycle conditions. We review the small empirical literature and show that the methods used thus far are poorly suited for identifying employers' wage practices. We propose a simpler and more relevant approach – use matched employer/employee longitudinal data to identify entry jobs and then directly track the cyclical variation in the real wages paid to workers newly hired into those jobs. We illustrate the methodology by applying it to data from an annual census of employers in Portugal over the period 1982-2007. We find that real entry wages in Portugal over this period tend to be about one-and-a-half percent higher when the unemployment rate is one percentage point lower.

The authors thank Michael Elsby and Pedro Portugal for helpful discussions. Martins and Thomas are grateful for the support of the ESRC (RES-062-23-0546).

Measuring What Employers Really Do about Entry Wages over the Business Cycle

I. Introduction

Recent theoretical efforts to explain the amplitude of cyclical unemployment fluctuations have emphasized the potential importance of rigidity in the real wages that firms pay to their newly hired workers. But are firms' hiring wages really acyclical?

In section II, we review the recent theoretical literature. In section III, we argue that the empirical literature has not succeeded yet in identifying what employers really do about hiring wages, and then we propose a more direct and transparent empirical approach – using matched employer/employee longitudinal data to identify firms' entry jobs and then tracking the cyclical variation in the real wages paid to workers newly hired into those jobs. In section IV, we illustrate our approach with an application to data from the annual census of employers in Portugal. In section V, we briefly summarize our findings, and we discuss their macroeconomic implications.

II. The Macroeconomic Importance of Entry-Wage Cyclicity

Recent interest in the cyclical behavior of hiring wages has grown out of a debate about the ability of the canonical Mortensen-Pissarides (1994) model to generate realistically large cyclical fluctuations in unemployment. In the canonical model, wages within a worker/employer match are determined by Nash bargaining, and this yields a solution where the wage depends on the worker's productivity in the match and the worker's outside option, which depends in turn on the value of search and non-market activities. Because the value of non-market activities is less cyclical than labor productivity, the wage does not vary as much as productivity, and employment is

correspondingly more volatile. However, as pointed out by Shimer (2005), under standard parameter values, the outside option is almost as cyclical as productivity, with the implication that wages in the canonical model are likewise too cyclical to generate realistic levels of unemployment variability.

Consequently, several authors – such as Shimer (2004), Hall (2005), and Hall and Milgrom (2008) – have suggested wage stickiness as one way of generating more variability of unemployment within the model. Whether this is the right way to go, however, should ultimately be subject to empirical investigation. How much *do* wages actually vary over the business cycle?

Of course, there already is a great deal of evidence on certain aspects of that question. Numerous studies using longitudinal microdata have found that real wages are substantially procyclical.¹ As will be discussed in the next section, many of those studies find particularly large wage cyclicity for job changers, with more moderate cyclicity appearing for the wages of workers in ongoing matches.

This apparent difference in wage cyclicity between incumbent and newly hired workers seems at odds with the canonical model, in which continuous Nash bargaining would cause wages in ongoing matches to be as procyclical as those in new matches. This is easily accommodated, though, by a sensible modification of the model. All that matters for the job creation incentive (what determines how many new jobs there are) is the expected (discounted) value of wages in the match, not how they are distributed over time, a point that has long been recognized (Becker, 1962; Barro, 1977; Hall, 1980; and Boldrin and Horvath, 1995) and has been reemphasized in the context of the current

¹ See Solon, Barsky, and Parker (1994) and the literature review in Solon, Barsky, and Parker (1992). For more recent studies focusing on wage cyclicity in ongoing matches, see Devereux (2001) and Shin and Solon (2007).

debate (Shimer, 2004; Haefke, Sonntag, and van Rens, 2008; and Pissarides, forthcoming). Hence, it is argued, all that is needed to square the model with the data is to incorporate a labor contract that makes ongoing wages relatively insensitive to future productivity changes, for example, by making workers slightly risk averse while maintaining the canonical model's assumption that firms are risk neutral. In such an environment, if both firm and worker could commit to a future wage path, ongoing wages would be constant but at a level influenced by productivity at the point of match initiation, so wages of new hires could vary procyclically while those of "stayers" would be acyclical. An approach along these lines then suggests that the wage in a new match would be persistent and the total cost of the worker – and hence the profitability of the match – would vary with the start wage.

In fact this goes too far to the other extreme as now stayers' wages are *too* rigid relative to the data. Rudanko (2009) incorporates Thomas and Worrall's (1988) model of "limited commitment" into a search-matching model, so that ongoing wages respond to compete with outside options and therefore display a degree of procyclicality. For example, Rudanko finds that, if neither firm nor worker can commit, then the model can match the empirical aggregate wage volatility (and while she does not examine this, the continuing match wage will also be procyclical). However, the effect on magnifying unemployment and vacancy volatility is small – essentially confirming the point that the distribution of wages over time in an ongoing match is unimportant.

This issue is considered further by Kudlyak (2007), who compares the cyclicity of the new hire wage with that of the "user cost of labor." The latter is defined to measure the cost of hiring a worker today relative to postponing this decision one period,

and is the primary determinant of the current employment decision (in the equilibrium of her matching model, the marginal product of labor will equal the user cost). If a worker hired next period will be paid differently from a hire today, the user cost must take into account this future difference in addition to the current wage. Kudylak emphasizes that the user cost may be even more procyclical than the new hire wage. To understand this, consider a worker hired in a downturn, and suppose full commitment by both worker and firm is possible. The equilibrium contract will pay a relatively low and constant wage. However, because of mean reversion, productivity is expected to improve, so a worker hired next period is likely to receive a higher constant wage. Thus the cost of filling a position today rather than tomorrow is the current wage *less* the discounted future expected savings on the lower wage contract, and the user cost falls more than the new hire wage in a recession (and vice versa in a boom). Even when commitment is imperfect, a similar argument may apply if wages in ongoing matches are sufficiently inertial relative to productivity.²

To summarize, even if wages in ongoing employment relationships are sticky, the wages of new hires may be highly procyclical, and the allocative importance of start wages is all the greater the more persistent those wages are as the match continues. As recent theoretical analyses have shown, with sufficiently procyclical start wages, the problem of explaining the volatility of unemployment remains. But *are* entry wages really very procyclical? In the next section, we explain why existing evidence is not

² Kudlyak (2009) makes a heroic attempt to measure the cyclical in the user cost of newly hired labor. The true value of that user cost as perceived by the employer depends on how long the employer believes the match will last, what the employer believes will be the future path of both the wage rate and the productivity of the match, and the employer's understanding of the interrelationships among these factors. For a data analyst to impute the employer's perception of the user cost therefore requires very strong assumptions. Based on the assumptions she makes, Kudlyak concludes that the user cost is more procyclical than the start wage.

dispositive, and we introduce a more direct approach for measuring entry-wage cyclicality.

III. Methods of Measuring Entry-Wage Cyclicality

A. Previous Empirical Studies

Pissarides (forthcoming) dismisses theories based on cyclically rigid wages on the ground that, empirically, hiring wages are quite procyclical. His main basis for the claim is a collection of microeconomic studies reporting that workers who change employers during expansions tend to experience much better wage growth than workers who change employers during recessions. Examples of such studies include Bils (1985) and Shin (1994) for the United States; Devereux and Hart (2006) for the United Kingdom; and Martins (2007) and Carneiro, Guimarães, and Portugal (2009) for Portugal. Similarly, Haefke, Sonntag, and van Rens (2008) find that individuals moving from non-employment to employment during expansions tend to receive higher wages than individuals moving from non-employment to employment during recessions.

Gertler and Trigari (2009), however, explain why such evidence does not rule out acyclical wage setting by firms. Even if every firm maintains cyclically rigid wages in each job title, workers taking new jobs will show procyclical wages if, during expansions, they face better opportunities to move into higher-paying industries, higher-paying firms within industries, or higher-paying jobs within firms. In contrast, during recessions, a larger share of job changers are displaced workers, who often suffer wage reductions relative to the jobs they have lost. As Gertler and Trigari (p. 73) put it, “Suppose, for example, that a highly skilled machinist takes a job as a low-paid cab driver in a

recession and then is re-employed as a high-paid machinist in a boom. In this case there is a cyclical movement in job match quality for the individual” even if the wages in both the machinist job and the cab driver job are acyclical.³ As Gertler and Trigari emphasize, this process of so-called “cyclical upgrading” in worker/employer matches (and cyclical *downgrading* in recessions) has long been discussed and documented. Early references include Reynolds (1951), Reder (1955), and Okun (1973). Recent analyses include Solon, Whatley, and Stevens (1997), McLaughlin and Bilts (2001), and Devereux (2004).

We should note that, from the workers’ perspective, their improved opportunities for moving to better jobs in an expansion are a true form of wage procyclicality. But that is not what theoretical analyses such as Hall (2005) are about. Rather, those analyses are attempting to explain the amplitude of cyclical fluctuations in unemployment by assuming rigidity in the hiring wages that particular firms pay in particular jobs. Gertler and Trigari therefore are correct in criticizing the existing empirical literature for rejecting such theories on the basis of evidence confounded by cyclical upgrading/downgrading.

Gertler and Trigari claim to get better estimates of cyclicity of hiring wages within particular jobs by using longitudinal U.S. worker data from 1990-1996 in the Survey of Income and Program Participation. In particular, they perform least squares estimation of a regression of the worker’s log real wage on the unemployment rate, the interaction of the unemployment rate with a dummy variable that equals 1 if the worker has been with the current employer for less than a year, dummy variables for every worker/employer match, and other control variables including the current tenure of the

³ Bewley (1999, p. 151) earlier made the exact same point: “the findings could result from the effects of recession on the level of jobs accepted during recessions rather than from the effects on hiring pay for particular jobs.”

match. The estimated coefficient of the interaction term is statistically insignificant, and Gertler and Trigari (pp.73-4) conclude that, “after we control for compositional effects, new hire wages appear no more cyclically sensitive than existing workers’ wages.”

Gertler and Trigari, however, misunderstand the properties of the sort of “fixed effects estimation” they implement by controlling for worker/employer matches. As is well-known from the literature on longitudinal estimation, controlling for match-specific fixed effects is the same thing as using only the within-match variation in the other regressors to identify the coefficients of those regressors.⁴ As a result, what Gertler and Trigari’s insignificant estimate for the coefficient of the interaction term really means is merely this – when one contrasts how a worker’s wage moves with changes in the unemployment rate during the first year with an employer (say, between the worker’s second and sixth months with the employer) with how the same worker’s wage moves with changes in the unemployment rate after the first year with the employer (say, between the fourteenth and eighteenth months of the match), the estimate of this contrast is statistically insignificant. As such, the estimated coefficient of the interaction term speaks only to an arcane aspect of how the worker’s wage evolves *after* the worker already has been hired. It says nothing whatsoever about cyclicity in the initial hiring wage.⁵

⁴ For example, see Wooldridge (2002, p. 273).

⁵ Where Gertler and Trigari control for match fixed effects, Carneiro, Guimarães, and Portugal (2009) (using the same Portuguese data base that we use) instead control for linearly separable worker fixed effects and firm fixed effects. Like the fixed effects estimation of Gertler and Trigari, the fixed effects estimation of Carneiro et al. does not seem to identify what the authors intend it to. First, controlling for linearly additive firm effects does not account for within-firm heterogeneity in match quality. Taking Gertler and Trigari’s skilled machinist as an example, controlling for firm fixed effects successfully adjusts for that worker’s cyclical upgrading only if the only reason the machinist is paid more by the firm that uses him as a machinist than the firm that uses him as a cab driver is that the first firm tends to pay *all* its workers more than the latter firm pays *all* its workers. Obviously, that need not be the case. Second, the particular regression model estimated by Carneiro et al. includes a dummy variable for matches in their first

To summarize, we agree with Gertler and Trigari’s point that empirical assessment of recent theories of hiring-wage rigidity requires an approach that identifies wage variation within particular jobs. Their approach, however, does not achieve this. In the next sub-section, we propose an alternative approach that does track hiring wages in particular jobs.

B. Our Methodology

To get at the question that recent researchers intended to address, we propose a more direct and pertinent methodology – identify specific entry jobs within particular firms, track the wages paid to newly hired workers in those jobs, and measure how those entry wages vary over the business cycle. A precedent for this approach was applied by Solon, Whatley, and Stevens (1997, p. 412), who found that “the entry wage for ‘laborers,’ the most common occupation for newly hired workers” in the Byers iron tubing and pipe manufacturing company, was procyclical over the 1919-1932 period. That result, however, applied to only one company observed over 70 years ago. In the present paper, we take a similar approach with a large number of firms observed between 1982 and 2007 in an annual census of employers in Portugal. Because we hope that our approach eventually will be applied to other countries besides Portugal, here we describe the general methodology. Idiosyncrasies that arise in the application to the Portuguese data set will be addressed in section IV.

year (which switches off after the first year of the match) and an interaction of that dummy with the unemployment rate. The sum of the coefficient of the interaction term and the coefficient of the unemployment rate is taken to represent the overall cyclical effect of hiring wages. Carneiro et al., however, do not justify that interpretation and, in certain cases, it clearly is incorrect. As a particularly simple example, consider the “full commitment” case in which a worker hired at a low wage during a recession is thereby locked into a low wage throughout the rest of the match. With identification coming off of the contrast between the first year and subsequent years of the match, this observation would bias the estimation away from finding a recession effect on the hiring wage.

Let w_{jt} denote the typical real wage paid in period t to workers newly hired into job j , which is a particular job into which a particular firm does a substantial amount of its hiring of new employees. A convenient statistical model for w_{jt} is

$$(1) \quad \log w_{jt} = \alpha_j + \beta_t + \varepsilon_{jt}$$

where α_j is a job fixed effect, β_t is a period fixed effect common to all entry jobs, and the zero-mean error term ε_{jt} represents temporary job-specific departures from the general period effect. The object of our analysis is to estimate a time series for β_t and relate it to business cycle conditions.

If the data set were a “balanced panel” with every one of the N sampled jobs observed in all periods, the estimation of β_t would be particularly simple. With the mean of α_j normalized to zero, the least squares estimator of each β_t would be simply

$$(2) \quad \hat{\beta}_t = \sum_{j=1}^N \log w_{jt} / N.$$

A figure that plotted the $\hat{\beta}_t$ time series along with a cycle indicator such as the unemployment rate would give an eyeball impression of the cyclical behavior of entry wages. A quantitative summary could be obtained by applying least squares to the regression of $\hat{\beta}_t$ on the unemployment rate with controls for secular time trends. If non-stationarity were a concern, the regression could be estimated in first differences instead of levels.

If the panel is unbalanced because observations of w_{jt} are missing for some jobs in some years, it is still straightforward to estimate β_t in equation (1) by applying least

squares to the regression of the available observations of $\log w_{jt}$ on a vector of period dummy variables (for which the estimated coefficients are the $\hat{\beta}_t$ series) and a vector of job dummies. Alternatively, exactly the same $\hat{\beta}_t$ are obtained by applying least squares to the “de-meanned” regression of $\log w_{jt}$ on a vector of period dummies, in which each variable is expressed as a deviation from the job’s time mean of the variable over the periods in which that job is observed. Either way of calculating $\hat{\beta}_t$ accounts for period-to-period changes in which jobs are in the sample by controlling for job fixed effects. This $\hat{\beta}_t$ series is the same as the sample means in equation (2) except that it is regression-adjusted for changes in sample composition. Again, the $\hat{\beta}_t$ series can be plotted along with the unemployment rate, and it can be regressed on the unemployment rate and time trend variables.

Although complications inevitably arise in practice (as will be seen in the next section), the basic methodology is almost absurdly simple. That is what we like about it. Our method looks in a direct and transparent way at what employers actually do about entry wages over the business cycle.

IV. An Application to Firm-Level Data from Portugal

A. Description of the data

In this section, we illustrate our method for measuring entry-wage cyclicality with an application to Portuguese data from 1982-2007. As shown in figure 1, Portugal’s annual unemployment rate varied widely over that period, with peaks in the mid 1980’s, mid 1990’s, and mid 2000’s. Another relevant feature of the Portuguese labor market

over that period is its system of collective wage bargaining, which is explained in detail in Cardoso and Portugal (2005). The main message of that paper is that the bargained wage rates function as wage *floors*, with employers commonly paying wages above those floors. In the authors' words (p. 899), "the wage cushion works as a mechanism to overcome the constraints imposed by collective bargaining, allowing firms wide scope for action in their wage-setting policy.... This operation of institutional forces and market forces in the Portuguese economy may help explain why a typically European institutional framework is compatible with high wage flexibility and low unemployment."

Our data come from Quadros de Pessoal, an annual mandatory census of all employers in Portugal (except that most of the public sector is excluded). The census was conducted in March in each year through 1993 and in October from 1994 on. Employer identification numbers enable longitudinal matching of employers. Information on individual employees also is available for every year between 1982 and 2007 except 1990 and 2001. The employee information includes wage rate, detailed occupation code, and tenure with the firm. Employee identification numbers enable longitudinal matching of individual workers, even when they change employers, from 1986 on.

The data base suits our purposes very well. The combination of occupation and tenure information enables us to identify jobs into which each firm commonly recruits new employees, and we also observe the wages that the new employees in those jobs are paid. And by tracking those entry-job wages longitudinally, we can observe how they vary over the business cycle.

To identify a set of jobs into which, year after year, employers are observed to hire new workers, we apply a series of sample selection criteria. First, we consider only firms that employed at least 50 workers in at least five years of the 1982-2007 period. Second, from 1995 on, we define jobs within firms in terms of six-digit occupation codes (such as supermarket shelf stocker, bank teller, clothing machine operator). When the same occupation code appears for different firms, we treat each occurrence as a different occupation/firm job. We use a Ministry of Employment document to link the six-digit codes from 1995 on with the five-digit codes used through 1994. In the process, we lose some jobs for which the pre-1995 codes branch out into multiple post-1995 codes.⁶ Third, defining newly hired workers as those with no more than three months of tenure, we include a job in our main sample of entry jobs if, in at least half the years the firm is present in the data base, the particular job accounted for at least three new hires and at least 10 percent of the firm's new hires in that year.

We therefore are focusing on what Doeringer and Piore (1971) called “port-of-entry” jobs. We do not mean, however, to subscribe to their stark description in which firms hire into only a limited number of such jobs, with other jobs filled almost exclusively by internal promotions and reassignments. Recent studies, such as Baker, Gibbs, and Holmstrom (1994), have suggested that hiring from outside a firm can take place in a broad spectrum of jobs. Our focus on jobs that recurrently show new hires in the annual employer census is driven entirely by a pragmatic concern – to identify

⁶ We also check that all workers in the job are at the same “job level,” an eight-category variable coded as (1) apprentices, interns, trainees; (2) non-skilled professionals; (3) semi-skilled professionals; ... (8) top executives (top management). In most cases, this check is superfluous because it is unusual for the workers listed in the same occupation code within a firm to span more than one job level.

cyclical variation in hiring wages by job, we need those wages to be observed in multiple years spanning different business cycle conditions.

The main sample resulting from our selection criteria consists of 660 jobs in 598 firms. Because new hires are not observed in all jobs in all years, our panel of wages in entry jobs is “unbalanced,” with a total of 5,727 job/years. The unbalancedness occurs because some firms begin or end between 1982 and 2007 and because, in existing firms, it may happen that the workers observed in the job as of a particular census month include no one hired within the last three months. The number of entry-job wages observed per year ranges from a low of 55 in 1983 to a high of 415 in 2005.

Obviously, our analysis pertains to only a tiny fraction of all the jobs and firms encompassed in the Portuguese census of employers. It does, however, focus on those jobs that are most clearly recognizable as perennial entry jobs. Perhaps it is best to view our study as a sort of summary of 660 case studies of particular entry jobs.

As explained in the previous section, the dependent variable in our first-stage regressions is the typical logarithmic real hiring wage in job j in period t . The wage measure we start with is each newly hired worker’s monthly base pay (“corresponding to the normal hours of work”) divided by the worker’s normal monthly hours. We use Statistics Portugal’s monthly consumer price index (for March through 1993, for October afterwards) to convert this wage to real terms. Our main measure of the *typical* log hiring wage in a job/year is the modal value (with log wages measured to two decimal places). Typically, the majority of newly hired workers in a job/year are paid at the modal rate.

B. Empirical Analysis

As explained in section III.B, we begin by estimating the year effects β_t in equation (1) by applying least squares to the regression of the logarithm of job-specific real entry wages w_{jt} on year dummy variables with controls for job fixed effects. As discussed in section III.B, if the panel of entry jobs were balanced with entry wages observed for all jobs in all periods, the estimated year effects would be simply $\hat{\beta}_t = \sum_{j=1}^N \log w_{jt} / N$. With our highly unbalanced panel, our estimates essentially regression-adjust that simple statistic for year-to-year changes in sample composition.

The resulting $\hat{\beta}_t$ series is plotted in figure 2, with the value for 1982 normalized to zero. Recall from section IV.A that the series pertains to March for the observations from 1982 through 1993, it pertains to October for the observations from 1994 through 2007, and it is missing the years 1990 and 2001. The figure also displays the seasonally adjusted unemployment rates for the corresponding months.⁷ Eyeballing the figure in a way that takes account of the upward secular trend in wages gives the impression that entry wages and the unemployment rate are inversely related, but that the relationship is very loose.

As discussed in section III.B, we can quantify the cyclicity of entry wages by estimating regressions of the $\hat{\beta}_t$ time series on the unemployment rate and secular time trend controls. The first row of table 1 shows the estimated coefficient of the unemployment rate when this regression controls for a linear time trend and is estimated

⁷ These are the Eurostat unemployment rates available at http://sdw.ecb.europa.eu/quickview.do?SERIES_KEY=132.STS.M.PT.S.UNEH.RTT000.4.000. As these go back only to 1983, for March 1982 we use the first-quarter unemployment rate from Statistics Portugal.

by weighted least squares (weighting by the number of entry jobs observed per year, in an effort to correct for the heteroskedasticity resulting from the wide variation in the per-year sample size). The procyclical coefficient estimate of -1.59 (with estimated standard error 0.42) implies that, when the unemployment rate is one point higher (say, .07 instead of .06), real entry wages tend to be about one-and-a-half percent lower.

Heteroskedasticity diagnostics applied to the residuals suggest that the weighting procedure is reasonable, but in the next row of table 1 we use ordinary least squares instead of weighted least squares. The resulting coefficient estimate, -1.50 (with estimated standard error 0.52), is very close to the weighted result. In the third row, we apply ordinary least squares to the first-differenced version of the regression. A linear time trend in levels implies that, in first differences, we need to control for the varying time interval between adjacent observations, which is usually 12 months but is 24 months for the 1989-1991 and 2000-2002 changes and 19 months for the change from March 1993 to October 1994. The resulting estimate, -1.89 (0.94), is a little more procyclical than the previous estimates.

In the fourth row, we revert to the analysis in the first row except that, as a simple check for dynamics, we include the 12-month lag of the unemployment rate as an additional regressor.⁸ The estimated coefficient of the lagged term, 0.50 (0.83), is insignificantly positive, and the estimated coefficient of the contemporaneous term, -2.17 (0.78), is significantly procyclical. Combining the estimates, the implied wage response to a one-point increase in the unemployment rate in *both* periods is -1.67.

⁸ The analysis of Portuguese wage cyclicality in Carneiro, Guimarães, and Portugal (2009) uses *only* a one-year lag of the unemployment rate and does not report results for the contemporaneous unemployment rate.

In the fifth row, we redo the regression from row 1 except that we control for a quadratic time trend. The estimated coefficient of the squared trend variable is statistically insignificant, and the estimated coefficient of the unemployment rate, -1.22 (0.58), becomes a little less procyclical.

The next two rows repeat the analysis from row 1 but with alternative approaches to the first-stage estimation of the $\hat{\beta}_t$ series. In the sixth row, instead of representing the “typical” job/year wage w_{jt} with the modal wage of new hires, we use their average wage. In the second-stage regression, this results in a slightly more procyclical estimate of -1.74 (0.58). In the seventh row, we use the modal wage, but we use a more inclusive sample of entry jobs in the first stage. Instead of requiring that the firms employ at least 50 workers in at least five years, we require that they employ at least 25. And instead of requiring that the job involves at least three new hires and at least 10 percent of the firm’s new hires in at least half the years the firm is in the data base, we require at least two new hires and at least 10 percent of the firm’s new hires in at least a quarter of the years. Again, in the second stage, this results in a slightly more procyclical estimate of -1.82 (0.37).

Finally, because some analyses later in this section will involve only the years 1986-2006, in the last row we redo the analysis of row 1 with only those years. The resulting coefficient estimate, -1.21 (0.39), is a little less procyclical.

In general, throughout table 1, the estimated coefficients of the unemployment rate stay fairly close to -1.5. Recognizing that, with modestly procyclical labor force participation, the change in the unemployment rate probably is a slightly attenuated version of proportional changes in employment, these estimates imply that the cyclical

elasticity of entry wages and the cyclical elasticity of employment are of similar magnitudes. In that sense, our estimates of the procyclicality of entry wages are substantial. The result becomes more striking when one considers that the “cyclical upgrading” discussed in section III.A suggests that we may be underestimating the true procyclicality of entry wages. If, in a recession, employers are able to recruit a higher quality of workers at any given wage, the effective wage they pay per efficiency unit of labor is that much lower.⁹ In section V, we will discuss further the macroeconomic implications of the magnitude of our estimates.

Another perspective on the estimates is that they are similar to the estimates of the overall cyclicity of workers’ real wages in the U.S. literature that uses longitudinal data to relate workers’ year-to-year changes in log wages to changes in the unemployment rate.¹⁰ This similarity is surprising for two reasons. First, the estimates are for different countries. Second, they are measuring different facets of wage cyclicity. As discussed in section II, wage smoothing in long-term employment relationships could cause overall wages to be less procyclical than the hiring wages in entry jobs. On the other hand, as discussed in section III.A, workers’ procyclical opportunities to move to higher-paying jobs (both within and across firms) contribute to estimates of workers’ overall wage procyclicality, while our estimates of entry-wage cyclicity are designed to hold jobs constant.

⁹ So far, our only empirical exploration of this possibility with the Portuguese data has been to re-estimate some of the regressions in table 1 with average years of education of newly hired workers, instead of their log hiring wage, as the dependent variable. These analyses have not shown any significant evidence of cyclical upgrading with respect to measured education of new hires.

¹⁰ See Bils (1985), Solon, Barsky, and Parker (1994), and the literature review in Solon, Barsky, and Parker (1992).

To address the first issue of between-country differences, in table 2 we use longitudinal worker data from the Portuguese census of employers to replicate the analyses in the U.S. longitudinal literature. In the tradition of Bils (1985), we estimate the regression of workers' year-to-year change in log real wages on change in the unemployment rate. Following Solon, Barsky, and Parker (1994), we break the estimation into two steps. In the first stage, we apply ordinary least squares to the regression of the worker's change in log real wages on year dummy variables and the interaction of the worker's age with the time interval between the adjacent employer censuses. As shown by Shin (1994), inclusion of the latter variable in the differenced specification is appropriate if the log wage/age profile is quadratic in levels. In the second stage, much as in the third row of table 1, we apply ordinary least squares to the regression of the estimated year effects on change in the unemployment rate and controls for the time interval between the adjacent censuses.¹¹ The estimates are based on the period 1986-2006 because worker identifiers are not available before 1986 and we have not yet solved problems with longitudinally matching worker observations from the newly available 2007 data.

The resulting coefficient estimate for the unemployment rate, shown in the first row of table 2, is -1.64 (with estimated standard error 0.48). This estimate is more procyclical than the entry-wage counterpart for 1986-2006 shown in the last row of table 1. This suggests that, between the two factors that cause wage cyclicity from the

¹¹ In this instance, we control for a quadratic in the interval, corresponding to a quadratic in the log wage/age profile. We apply ordinary instead of weighted least squares in the second stage for reasons discussed in Dickens (1990). In our worker-level analysis, as opposed to the entry-job analysis in table 1, our first-stage estimation involves millions of observations. As a result, sampling error in the estimated year effects is inconsequential, and the error term in the second-stage regression is dominated by unobserved time effects common to all workers.

workers' perspective to differ from cyclical-ity of job-specific entry wages, the cyclical-upgrading factor dominates the wage-smoothing factor. The estimate of -1.64, which is based on both male and female workers, is similar to the estimates for male workers in the U.S. longitudinal literature and higher than those for female workers. Furthermore, it is based on the same measure of workers' wages used in the entry-wage analysis – the ratio of the worker's monthly base pay to the worker's normal monthly hours. This measure leaves out overtime pay and other such non-base payments, which typically are included in the wage measures used in the U.S. longitudinal literature. In the second row of table 2, we include such pay by measuring the wage as the ratio of the worker's total monthly pay to the worker's total monthly hours. The resulting coefficient estimate for the unemployment rate, -1.86 (0.46), is even more procyclical. The substantial procyclical-ity of real wages in Portugal accords with the conclusion of Cardoso and Portugal (2005), mentioned at the beginning of section IV.A, that "wage cushions" above the wage floors prescribed in collective bargaining enable flexible wages in Portugal.

As discussed in sections II and III.A, the wage cyclical-ity literature based on longitudinal worker data has estimated particularly strong wage procyclical-ity for workers who change employers and more moderate procyclical-ity for workers who stay with the same employer. In the third row of table 2, we redo the estimation in row 1 with a sample restricted to workers who change employers between censuses. As in the previous literature, we find particularly procyclical wages for changers, with an estimated unemployment rate coefficient of -2.53 (0.48).¹² Of course, the wage procyclical-ity for

¹² The closest counterpart to this estimate in the analysis of the 1986-2005 Portuguese data by Carneiro, Guimarães, and Portugal (2009) pertains to male "accessions," i.e., men with employer tenure of less than a year (who may or may not have been employed with another firm a year earlier). Carneiro et al. apply ordinary least squares in one stage to a levels regression of the worker's log wage (deflated by a lagged

workers who change employers is expected to exceed other measures of wage procyclicality because, for this group of workers, cyclical upgrading/downgrading is especially salient and wage smoothing is not a factor.

In the last row of table 2, we redo the row 1 analysis for a sample restricted to workers who stay with the same employer from one census to the next. In the first-stage regression, we control for the interactions of time between censuses with both age and tenure with employer. The second-stage estimate of the unemployment rate coefficient, -1.61 (0.48), is about as procyclical as both our estimate of overall wage cyclicality and most of our estimates of entry-wage cyclicality.¹³ The former comparison is somewhat surprising because, in the U.S. literature, the estimates of stayers' wage cyclicality usually are noticeably less than the estimates of overall wage cyclicality.¹⁴ With respect to the latter comparison, any wage-smoothing factor for stayers apparently is offset by cyclical upgrading of stayers across jobs within firms.

To summarize, longitudinal worker-based estimates of real wage cyclicality in Portugal are fairly similar to those in the U.S. literature, with the Portuguese estimates appearing a little more procyclical. In Portugal, however, unlike the United States at present, it is possible to use detailed employer/employee longitudinal data to conduct a direct and transparent analysis of the cyclical behavior of wages paid to newly hired workers in a large number of entry jobs. As discussed in section III, these data are better

version of the consumer price index) on the one-year lag of the unemployment rate, a quadratic in time, a quadratic in age, and controls for worker fixed effects. The estimated coefficient of their unemployment variable is -2.73. See our footnote 5 for a discussion of our concerns about the way in which Carneiro et al. pool their estimation of the cyclicality in accession wages with the cyclicality of other wages.

¹³ In the study described in the preceding footnote, Carneiro, Guimarães, and Portugal (2009) also report results for male stayers, for whom the estimated coefficient of the one-year lag of the unemployment rate is -1.50.

¹⁴ But some studies – such as Solon, Barsky, and Parker (1994) and Shin and Solon (2007) – report only moderate differences.

suites for assessing recent theories of firms' practices vis-à-vis setting hiring wages. Our analyses of the Portuguese data have produced a robust finding that real entry wages by job tend to be about one-and-half percent lower when the unemployment rate is one percentage point higher.

V. Summary and Discussion

Using longitudinally matched data from Portugal's annual census of employers, we have tracked the cyclical behavior of the real wages paid to newly hired employees in several hundred jobs. We hope that eventually our methodology will be applied to additional countries as the requisite data become available.

Our estimates suggest that real entry wages in Portugal tend to be about one-and-a-half percent lower when the unemployment rate is one percentage point higher. As noted in section IV.B, this procyclicality in entry wages is substantial in the sense that the cyclical elasticity of this "price variable" is of approximately the same magnitude as the cyclical elasticity of "quantity variables" like employment and unemployment. But are our estimates of entry-wage cyclicality large in other economically meaningful ways? We conclude with two remarks on this question.

First, some recent (as discussed in section II) and older attempts at theoretically explaining large cyclical changes in unemployment have emphasized rigidity of wages. Like many recent studies of other aspects of wage cyclicality, our study of real entry wages in Portugal has found that wages are not rigid, but rather respond considerably to business cycle conditions. Of course, this finding does not preclude that, in some sense, wages are "insufficiently" variable. Pursuing that idea, however, will require not a theory

of wage *rigidity*, but a theory of why wages are not even more variable than they are. Putting the point another way, the very existence of large cyclical fluctuations in employment and unemployment indicates that wages do not vary enough to prevent those fluctuations. The continuing challenge to macroeconomic theorists is to develop a logically coherent and empirically relevant explanation of why employment, unemployment, and wages vary over the business cycle in the ways that they actually do.

Second, returning to the question posed in section II, can the Mortensen-Pissarides model in particular account for the cyclical variability of unemployment in light of the magnitude of the entry-wage cyclicalities that we have found? As discussed in section II, the answer depends not only on the cyclicalities in hiring wages, but also on the persistence of those wages as the worker/employer match continues. At one extreme, if workers hired at a low wage during a recession are somehow locked into a permanent employment relationship at a perpetually low wage, the labor cost relevant to employers' recruiting decisions could be as cyclical as, or even more cyclical than, the initial hiring wage. At the other extreme, if the shortfall in hiring wages during a recession vaporizes almost immediately after the employment relationship is initiated, cyclicalities in the initial hiring wage might have only trivial impact on hiring decisions. Thus, assessing the allocative importance of initial hiring wages in the Mortensen-Pissarides model (or, for that matter, any model that recognizes the prevalence of long-term employment relationships) will require a deeper understanding of the durability of employment relationships, the evolution of wages and productivity in those relationships, and the dependence of both on current, past, and anticipated business cycle conditions. The empirical literature pioneered by Beaudry and DiNardo (1991), which is summarized in

Thomas and Worrall (2007), has begun to explore these issues, and we believe further research in that area could be highly productive.

References

- Baker, George, Michael Gibbs, and Bengt Holmstrom. 1994. "The Internal Economics of the Firm: Evidence from Personnel Data." *Quarterly Journal of Economics*, 109(4): 881-919.
- Barro, Robert J. 1977. "Long-Term Contracting, Sticky Prices, and Monetary Policy." *Journal of Monetary Economics*, 3(3): 305-16.
- Beaudry, Paul, and John DiNardo. 1991. "The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data." *Journal of Political Economy*, 99(4): 665-88.
- Becker, Gary S. 1962. "Investment in Human Capital: A Theoretical Analysis." *Journal of Political Economy*, 70(5): S9-49.
- Bewley, Truman F. 1999. *Why Wages Don't Fall during a Recession*. Cambridge, MA: Harvard University Press.
- Bils, Mark J. 1985. "Real Wages over the Business Cycle: Evidence from Panel Data." *Journal of Political Economy*, 93(4): 666-89.
- Boldrin, Michele, and Michael Horvath. 1995. "Labor Contracts and Business Cycles." *Journal of Political Economy*, 103(5): 972-1004.
- Cardoso, Ana Rute, and Pedro Portugal. 2005. "Contractual Wages and the Wage Cushion under Different Bargaining Settings." *Journal of Labor Economics*, 23(4): 875-902.
- Carneiro, Anabela, Paulo Guimarães, and Pedro Portugal. 2009. "Real Wages and the Business Cycle: Accounting for Worker and Firm Heterogeneity." Unpublished.
- Devereux, Paul J. 2001. "The Cyclicity of Real Wages within Employer-Employee

- Matches.” *Industrial and Labor Relations Review*, 54(4): 835-50.
- Devereux, Paul J. 2004. “Cyclical Quality Adjustment in the Labor Market.” *Southern Economic Journal*, 70(3): 600-15.
- Devereux, Paul J., and Robert A. Hart. 2006. “Real Wage Cyclicalities of Job Stayers, Within-Company Job Movers, and Between-Company Job Movers.” *Industrial and Labor Relations Review*, 60(1): 105-19.
- Dickens, William T. 1990. “Error Components in Grouped Data: Is It Ever Worth Weighting?” *Review of Economics and Statistics*, 72(2): 328-33.
- Doeringer, Peter B., and Michael J. Piore. 1971. *Internal Labor Markets and Manpower Analysis*. Lexington, MA: D. C. Heath.
- Gertler, Mark, and Antonella Trigari. 2009. “Unemployment Fluctuations with Staggered Nash Wage Bargaining.” *Journal of Political Economy*, 117(1): 38-86.
- Haefke, Christian, Marcus Sonntag, and Thijs van Rens. 2008. “Wage Rigidity and Job Creation.” Discussion Paper No. 3714, IZA.
- Hall, Robert E. 1980. “Employment Fluctuations and Wage Rigidity.” *Brookings Papers on Economic Activity*, 1980(1): 91-123.
- Hall, Robert E. 2005. “Employment Fluctuations with Equilibrium Wage Stickiness.” *American Economic Review*, 95(1): 50-65.
- Hall, Robert E., and Paul R. Milgrom. 2008. “The Limited Influence of Unemployment on the Wage Bargain.” *American Economic Review*, 98(4): 1653-74.
- Kudlyak, Marianna. 2007. “The Cyclicalities of the User Cost of Labor with Search and Matching.” Unpublished.
- Kudlyak, Marianna. 2009. “The Cyclical Price of Labor When Wages Are Smoothed.”

- Unpublished.
- Martins, Pedro S. 2007. "Heterogeneity in Real Wage Cyclicalities." *Scottish Journal of Political Economy*, 54(5): 684-98.
- McLaughlin, Kenneth J., and Mark Bils. 2001. "Interindustry Mobility and the Cyclical Upgrading of Labor." *Journal of Labor Economics*, 19(1): 94-135.
- Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies*, 61(3): 397-415.
- Okun, Arthur M. 1973. "Upward Mobility in a High-Pressure Economy." *Brookings Papers on Economic Activity*, 1973(1): 207-52.
- Pissarides, Christopher A. Forthcoming. "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?" *Econometrica*.
- Reder, Melvin W. 1955. "The Theory of Occupational Wage Differentials." *American Economic Review*, 45(5): 833-52.
- Reynolds, Lloyd G. 1951. *The Structure of Labor Markets*. New York: Harper & Brothers.
- Rudanko, Leena. 2009. "Labor Market Dynamics under Long-Term Wage Contracting." *Journal of Monetary Economics*, 56(2): 170-83.
- Shimer, Robert. 2004. "The Consequences of Rigid Wages in Search Models." *Journal of the European Economic Association*, 2(2-3): 469-79.
- Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review*, 95(1): 25-49.
- Shin, Donggyun. 1994. "Cyclicalities of Real Wages among Young Men." *Economics*

- Letters*, 46(2): 137-42.
- Shin, Donggyun, and Gary Solon. 2007. "New Evidence on Real Wage Cyclicity within Employer-Employee Matches." *Scottish Journal of Political Economy*, 54(5): 648-60.
- Solon, Gary, Robert Barsky, and Jonathan A. Parker. 1992. "Measuring the Cyclicity of Real Wages: How Important Is Composition Bias?" Working Paper No. 4202, National Bureau of Economic Research.
- Solon, Gary, Robert Barsky, and Jonathan A. Parker. 1994. "Measuring the Cyclicity of Real Wages: How Important Is Composition Bias?" *Quarterly Journal of Economics*, 109(1): 1-25.
- Solon, Gary, Warren Whatley, and Ann Huff Stevens. 1997. "Wage Changes and Intrafirm Job Mobility over the Business Cycle: Two Case Studies." *Industrial and Labor Relations Review*, 50(3): 402-15.
- Thomas, Jonathan P., and Tim Worrall. 1988. "Self-Enforcing Wage Contracts." *Review of Economic Studies*, 55(4): 541-54.
- Thomas, Jonathan P., and Tim Worrall. 2007. "Limited Commitment Models of the Labour Market." *Scottish Journal of Political Economy*, 54(5): 750-73.
- Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.

Figure 1. Annual Unemployment Rates in Portugal, 1982-2007




Figure 2. Estimated Period Effects for Log Entry Wages in Available Months and Corresponding Unemployment Rates, March 1982-October 2007



Table 1. Estimates of Cyclicalities of Log Entry Wages in Portugal, 1982-2007

Estimation Method	Estimated Coefficient of Unemployment Rate (and Estimated Standard Error)
1. Weighted least squares for regression of log modal entry wage year effects on unemployment rate and linear time trend	-1.59 (0.42)
2. Same as (1) but ordinary least squares	-1.50 (0.52)
3. Ordinary least squares for first differences	-1.89 (0.94)
4. Same as (1) with one-year lag of unemployment rate as additional regressor	Contemporaneous: -2.17 (0.78) Lagged: 0.50 (0.83)
5. Same as (1) with quadratic time trend	-1.22 (0.58)
6. Same as (1) with log average wage	-1.74 (0.58)
7. Same as (1) with more inclusive definition of entry jobs	-1.82 (0.37)
8. Same as (1) for 1986-2006	-1.21 (0.39)

Table 2. Estimates from Regressions of Portuguese Workers' Change in Log Wages on Change in Unemployment Rate, 1986-1987 to 2005-2006

Estimation Method	Estimated Coefficient of Change in Unemployment Rate (and Estimated Standard Error)
1. All longitudinally matched workers, with controls for age and time	-1.64 (0.48)
2. Same as (1) but with wage measure including "non-base" pay	-1.86 (0.46)
3. Same as (1) but for sample restricted to workers changing employers	-2.53 (0.48)
4. Same as (1) but for sample restricted to workers staying with same employer (with additional control for tenure with employer)	-1.61 (0.48)