AER: Insights 2021, 3(1): 83–96
https://doi.org/10.1257/aeri.20190553

The Cyclical Behavior of Job Quality and Real Wage Growth†

By Andriana Bellou and Bariş Kaymak*

We study the empirical relevance of implicit insurance contracts for wage setting while accounting for cyclical fluctuations in average job quality. Using proxy measures, we find the latter to be acyclical, if not countercyclical, due to the cleansing effects of layoffs during recessions versus quits during expansions. Then, we study the cyclical behavior of wage growth among job stayers to test for contracts, circumventing differences in job quality altogether. Both methods strongly corroborate the prevalence of wage contracts in the labor market and imply a highly procyclical price for labor. (*JEL* E24, E32, J31, J41, J63)

Equilibrium models of the labor market foresee a tight link between employment and labor costs. The cyclical nature of this link, in particular, is key to determining the sources of business cycles and assessing the roles of productivity shocks, wage stickiness, or monetary policy in employment fluctuations. Yet, empirical work using aggregate time series data found wages to be predominantly unresponsive to employment fluctuations.[1] However, since most employment is long term, and often contractual, economists have long suspected that contemporaneous pay does not reflect the marginal cost of labor (Becker 1962). In contractual markets, wages are history dependent, reflecting the economic conditions when the contract was signed. This was heralded by Bils (1985), who found wages to be especially procyclical when workers change employment. In a seminal paper, Beaudry and DiNardo (1991) (henceforth BD), showed that the cyclical behavior of wages is best explained by a model of dynamic insurance contracts with costless worker mobility. In such a setting, wages are rigid downward, to insure the worker against drops in productivity, but flexible upward to prevent quits if the worker receives a better offer elsewhere (Harris and Holmstrom 1982, Thomas and Worrall 1988). This explains not only the apparent apathy in average wages but also the large body of empirical work documenting a strong correlation between wages and past labor

* Bellou: Université de Montréal (email: andriana.bellou@umontreal.ca); Kaymak: Université de Montréal (email: baris.kaymak@umontreal.ca). Pete Klenow was coeditor for this article. We thank Mark Aguiar, Paul Beaudry, Mark Bils, Rui Castro, and Lyndon Moore for their comments. Authors acknowledge support from Chaire de la Fondation J.W. McConnell en études Américaines, Social Sciences and Humanities Research Council of Canada, and Fonds de Recherche du Québec—Société et Culture.

† Go to https://doi.org/10.1257/aeri.20190553 to visit the article page for additional materials and author disclosure statement(s).

[1] See Basu and House (2016) for a recent survey.

market conditions.[2] Allowing for contractual patterns in wages proves the price for labor to be highly procyclical (Kudlyak 2014).

A well-recognized challenge in measuring wage cyclicality is that workers and the jobs they perform are not necessarily comparable along the business cycle. It was argued by Stockman (1983) and subsequently shown by Bils (1985) and Solon, Barsky, and Parker (1994) that employment of low-wage workers is more cyclical, pushing average wages up during recessions. These findings, along with the countercyclical nature of job destruction documented by Davis and Haltiwanger (1992), depicted recessions as episodes of job cleansing, whereby labor is reallocated to better use (Caballero and Hammour 1994, Mortensen and Pissarides 1994). A similar reallocation activity is also thought to take place during expansions, however, when workers frequently change jobs (Barlevy 2002). Recently, Hagedorn and Manovskii (2013) argue that job-to-job movements associated with quit behavior lead to procyclical movements in average job quality. Furthermore, they find that accounting for job quality overturns the earlier evidence for the contractual nature of the labor market.

In this paper, we investigate the effect of job destruction on average job quality and test its implications for the empirical relevance of contracts for wage setting. To gauge the reallocative effects of separations, we develop proxies for job quality that distinguish between cleansing effects of quits, mostly during expansions, and of layoffs, mostly during recessions. The relative strength of each determines the cyclicality of average quality of sustained jobs.

To test for implicit insurance contracts under heterogeneity in job quality, we study wage growth among job stayers. This allows us to control for time-invariant differences in job quality for a given employer-employee pair. To see the essence of our argument, consider two identical workers: worker B, who was hired during a boom, and worker R, who was hired in the *subsequent* recession. If the employment relationship is characterized by insurance contracts, worker B enjoys a higher wage during the recession because he was insured against a possible downturn prior to the recession. Nonetheless, his advantage is temporary. As the economy recovers from the recession, outside opportunities improve. Since R is paid less for the same work, he is more likely to quit given a set of offers. Consequently, to prevent severance, the employer has to offer a raise to worker R, or offer him a larger raise relative to B. Thus R's expected wage gain is larger than B's. If, on the other hand, the labor market is characterized by spot transactions, both workers should be paid equally at all times since they are identically productive. Hence, in a spot market there is no reason to expect cyclical wage adjustments to depend on past economic conditions.

Using data from the 1979 cohort of the National Longitudinal Survey of Youth (NLSY), we find that poor matches dissolve not only during expansions when workers quit for better jobs but also during recessions when they are laid off. The cleansing effects of recessions are at least as important, resulting in countercyclical movements in average quality of remaining jobs. When both selection channels are accounted for, the cyclical wage behavior is consistent with implicit insurance contracts. This result is reinforced by the behavior of wage growth among job stayers

[2] The literature finds significant cohort effects (see, for instance, Baker, Gibbs, and Holmstrom 1994) and persistent negative effects of being hired in a recession (e.g., Kahn 2010).

in response to changes in economic conditions. We find that wage growth on the job depends only on past economic conditions and is insensitive to the contamporaneous state of the economy. This finding is robust to nonrandom selection of stayers over the business cycle, variations in the role of training and human capital, and situations with permanent differences in wage growth that are specific to a worker or job.

Next, we review the implications of job cleansing for the empirical tests of wage contracts. We show how cyclical changes in job quality can be either captured by appropriate proxy measures or sidestepped by studying wage growth among job stayers. Section II presents the data. Section III tests for contractual patterns in the data, and Section IV presents results on the cyclicality of cleansing effects in job quality. Section V verifies the robustness of the benchmark results. Section VI concludes.

I. Market Structure and the Cyclical Behavior of Wages

In this section, we discuss the implications of implicit contracts for wages when there is heterogeneity in job quality. We draw on the empirical literature on self-enforcing insurance contracts, which mainly focuses on arrangements that are binding on the employer but not the worker. We review the challenges that the current empirical tests face when there is cyclical variation in average job quality. Then we introduce our proxy measures for job quality and subsequently propose an alternative test that is robust to cyclical changes in the composition of jobs.

A. *Self-Enforcing Insurance Contracts and the Distribution of Wages*

Implicit insurance contracts are arrangements in which employers insure workers against temporary fluctuations in productivity. It is assumed that firms have better means of diversifying risk and that alternative means of insurance do not completely crowd out employer-provided insurance. The extent of the insurance and the implications for the empirical distribution of wages depend on the levels of commitment by the worker and employer. Full commitment by both parties implies that both the worker and employer are required to honor the contract irrespective of the economic conditions. Then, the optimal contract features full insurance, that is, a constant wage regardless of productivity. When there is free entry, the wage reflects the expected productivity of the worker at the time the contract is signed, and the economic conditions at the start of the job are sufficient to capture the cross-sectional variation in wages.

When the contract is binding *only* on the employer, the optimal contract features a downward rigid wage, which increases only if the worker's outside option sufficiently improves. The employer partially insures the worker by committing to never lower their wage in the future. Therefore, at any time, a worker's current wage reflects the highest wage he could command on the job, which is captured by the best economic conditions since the job started.

BD test these predictions with the following wage regression, where i denotes the worker and j denotes the job:

$$w_{ijt} = \theta_1 U_t + \theta_2 U_{ij0} + \theta_3 U_{ijt}^{\min} + X_{ijt}\Lambda + \epsilon_{ijt}. \tag{1}$$

The economic conditions are approximated by the unemployment rate: U_t is the current unemployment rate, U_{ij0} is the unemployment rate when the job started, and $U_{ijt}^{\min}$ is the minimum unemployment rate since then. If the labor market is spot, then wages are explained only by U_t. If the labor market is described by full-commitment contracts, then U_{ij0} is the only relevant variable. If, on the other hand, contracts are binding on the employer but not on the worker, then $U_{ijt}^{\min}$ is a sufficient statistic for wages. BD reject the spot market hypothesis in favor of contractual markets with one-sided commitment.

B. *Endogenous Separations and Cyclical Selection of Job Quality*

A superficial history dependence can arise in spot markets if unobserved components of productivity are correlated with past economic conditions. If improvements in job quality are led primarily by quits as opposed to layoffs, then bad matches dissolve, especially during expansions since quits are procyclical. By inference, jobs that survived expansions must be of higher quality and therefore pay higher wages. Such procyclical selection of job quality could also explain the empirical relevance of the lowest unemployment rate on average wages. This is the essence of Hagedorn and Manovskii's (2013) criticism of BD.

One way to address changes in the composition of job quality associated with quits as well as layoffs is to include proxies for match quality in the regression. To that end, assume a spot market ($\theta_2 = \theta_3 = 0$), set $\Lambda = 0$ in (1) for simplicity, and decompose the error as follows:

$$\epsilon_{ijt} = a_i + m_{ij} + \nu_{ijt},$$

where a_i is time-invariant worker quality, m_{ij} reflects job quality, and ν_{ijt} is random error. Matches dissolve either when workers quit for better matches or when firms lay off workers in response to adverse economic conditions. Either may alter the average quality of continuing matches.

Let us begin with quit behavior. Each period, workers draw offers from a stationary distribution before ν_{ijt} is realized. Let $\tilde{w}_{it}$ denote the best wage offer that a worker, employed or unemployed, can obtain in the market. This offer depends only on the current economic conditions. A worker quits if $\tilde{w}_{it} > w_{ijt}$. Since the cyclical and worker-specific components of productivity are equally valuable in all jobs, and since $E[\nu_{ijt}] = 0$, a better wage offer must come from a better match. Quits therefore improve matches over time.

To capture the procyclicality of quits, let $\tilde{m}_{it} = f(y_t)$ represent the job quality associated with the outside offer, where y_t denotes aggregate productivity, and assume that $f' > 0$. Then, the average quality of jobs lasting T periods depends on the best offer ever received:

$$E\left[m_{ij} | m_{ij} > \max\left\{\tilde{m}_{it_0}, \ldots, \tilde{m}_{it_0+T}\right\}\right],$$

where $\max\{\tilde{m}_{it_0}, \ldots, \tilde{m}_{it_0+T}\} = f(\max\{y_{t_0}, \ldots, y_{t_0+T}\})$ by monotonicity. Therefore, the best conditions during the *entire job spell* are a natural proxy for match quality. Note that this proxy is different than $U^{\min}$ in equation (1), which approximates the best conditions since the beginning, t_0, until the time when wage is observed, $t \in \{t_0, t_0 + T\}$. When both are included in a wage regression, the identification of the proxy for match quality comes from the changes in aggregate conditions between t and $t_0 + T$.

Let us now turn to layoffs. Consider a fall in aggregate productivity that deems certain jobs unprofitable. Let $\underline{m}(y_t)$ be the minimum sustainable match quality with $\underline{m}' < 0$. Job survival increases with match surplus and declines during recessions. The average quality of jobs that last T periods is $E[m_{ij} | m_{ij} > \max\{\underline{m}(y_{t_0}), \ldots, \underline{m}(y_{t_0+T})\}]$. The condition for the expectation can be approximated by $\min\{y_{t_0}, \ldots, y_{t_0+T}\}$ since $\underline{m}$ is strictly decreasing in y_t. Therefore, the *worst* economic conditions during the entire job spell are a natural proxy for changes in job quality induced by layoffs during downturns.

The proxies above are more general than the log-cumulative market tightness over a job spell: $\log \sum_{\tau=t_0}^{t_0+T} \theta_\tau$, proposed by Hagedorn and Manovskii (2013), along two dimensions. First, the latter does not distinguish between selection effects of quits and layoffs. High values of this measure are associated with past expansions and, hence, higher job quality due to the procyclical nature of quits. But so are the low values of it, since jobs that survive recessions are also of higher quality due to increased layoffs. This makes it a poor measure of job quality.

Second, it combines cyclical and noncyclical components of selection. To see this, notice that cumulative tightness is the sum of job duration and average tightness during the job spell: $\log T + \log \bar{\theta}_{t_0,T}$. Job duration captures the notion that good matches last longer, regardless of the business cycle (Abraham and Farber 1987). Hence, the correlation between job duration and wages is positive anyhow. The correlation between average tightness and wages, conditional on duration, could be positive or negative depending on the relative strengths of selection during recessions versus expansions. To measure the cyclicality of job selection, job duration and average conditions must be included *separately* in wage regressions. If the cleansing effect of quits dominates that of layoffs, then the coefficient on average tightness should be positive.[3]

We estimate wage regressions with different combinations of the proxies described previously to control for unobserved match quality. Next, we turn to another test of implicit insurance contracts that is robust to changes in job composition.

C. *Using Wage Growth of Job Stayers to Distinguish between Models*

In spot markets, wages do not depend on past conditions, but changes in the composition of jobs render them seemingly history dependent. By contrast, in a contractual market, each worker's wage is adjusted upward whenever $U^{\min}$ falls. The two models can therefore be distinguished by studying wage growth for job stayers.

To see the key idea, take the difference between two consecutive wage observations for a job stayer:

$$\Delta w_{ijt} = \theta_1 \Delta u_t + \theta_3 \Delta u_{ijt}^{\min} + \Delta X_{ijt} \Lambda + \Delta \nu_{ijt}. \tag{2}$$

[3] In recent work, Galindo da Fonseca, Gallipoli, and Yedid-Levi (2020) apply a similar decomposition to study differences in wage setting across occupations.

Since worker quality and match quality are time invariant, they disappear in (2). Therefore, they can no longer cause a selection bias. Equation (2) encompasses all three models: the spot labor market ($\theta_3 = 0, \theta_1 < 0$), a contractual market with one-sided commitment binding on the employer ($\theta_1 = 0, \theta_3 < 0$), and a contractual market with full commitment ($\theta_1 = \theta_3 = 0$).

The intuition is simple. Consider two workers with identical productivity: worker B, who was hired when the unemployment rate was 4 percent, and worker R, who was hired in the *subsequent* recession when the unemployment rate was 8 percent. In a contractual market, worker B maintains his wage over the recession, because he was insured against a possible downturn prior to the recession. As the economy recovers, the unemployment rate decreases to 6 percent. This reduces $U^{\min}$ for R but not for B. Consequently, R's wage increases while B's stays the same. If the labor market were spot instead, B and R would be paid the same wage, decreasing during the recession and increasing during the recovery.

A potential issue with estimating wage growth regressions for job stayers is that the error term $\Delta \nu_{ijt}$ may be correlated with the variables of interest conditional on staying on the job. If matches with lower realizations of ν_{ijt} were discontinued, then observed wage growth of stayers would be biased upward—presumably more so for low-wage workers, whose employment is more cyclical. Another issue may arise if there are persistent differences in wage growth across workers or jobs and the anticipation of these differences affects the quit behavior. We address these concerns in Section V.

II. Data

The data come from the 1979 cohort of the NLSY for the years 1979–2008 (Bureau of Labor Statistics 2009). The NLSY is a panel that closely tracks workers' jobs and their start and end dates, making it ideal for our purposes. We use the nationally representative cross-sectional sample and focus on respondents over 21 working at least 15 hours a week in the private sector.[4] The resulting sample is fairly broad—including part-time jobs, workers with multiple jobs, etc.—and constitutes a conservative choice for testing for implicit contracts. In the online Appendix, we show that the results are stronger for subsamples where contracts are more likely to be prevalent, such as workers with a single, full-time job.

To measure the cyclical fluctuations in workers' outside options in the labor market, Bils (1985) and BD rely on the unemployment rate. Hagedorn and Manovskii (2013) follow, except in constructing their proxies for match quality, where they use the more recently available labor market tightness measures. To facilitate comparison, we too use the unemployment rate to construct variables pertaining to implicit contracts but use the labor market tightness to construct our proxies for match quality. This choice has no bearing on our results (see the online Appendix).

[4] The online Appendix gives details on variable construction and sample selection.

Table 1—Real Wages and Unemployment History

	log$w$ (1)	log$w$ (2)	log$w$ (3)	log$w$ (4)	Δlog$w$ (5)	Δlog$w$ (6)
U_t	−1.63 (0.24)	−0.86 (0.19)	0.54 (0.22)	0.39 (0.23)		
ΔU_t					−0.18 (0.20)	0.31 (0.24)
U_{ij0}		−2.29 (0.20)		−0.57 (0.33)		
$U_{ijt}^{\min}$			−4.20 (0.37)	−3.54 (0.51)		
$\Delta U_{ijt}^{\min}$						−2.64 (0.63)
Observations	58,967	58,967	58,967	58,967	30,869	30,869

Notes: All specifications control for cubic polynomials in experience and tenure, a quadratic time trend, and indicators for industry and region. Columns 1–4 also control for worker fixed effects and include all observations. Coefficients and standard errors are multiplied by 100. Standard errors are clustered by start year and current year interactions and are shown in parentheses. Columns 5 and 6 include job stayers only.
Source: Data come from the 1979 cohort of the NLSY.

III. Cyclical Behavior of Wage Growth

We first document the history dependence in wages by estimating (1). The control variables are worker fixed effects, cubic polynomials in tenure and experience, a quadratic time trend, and indicators for region and industry. Table 1 shows the results. When regressed only on U_t, wages appear procyclical. A 1 percentage point increase in U_t is associated with a 1.63 percent drop in wages. When U_0 is introduced, however, the coefficient on U_t rises and eventually becomes statistically insignificant in column 4, which contrasts all three models at once. On average, wages decline by 0.57 percent in response to U_0 and by 3.54 percent in response to $U^{\min}$. Assuming a 6 percentage point difference in the unemployment rate between peak and trough, this implies a 24.7 percent $(= 6 \times (-0.57 - 3.54))$ gap in starting wages with a standard error of 5 percent.

The estimates confirm the earlier studies. The coefficient on $U^{\min}$ was estimated as −2.9 percent in BD and −2.5 percent in Grant (2003). The coefficients on other variables are also similar, with the exception that Grant (2003) finds a stronger effect for U_t.[5] Our findings are also consistent with Bils (1985), who finds wages to be more procyclical among new hires.

Next, we turn to job stayers. In the benchmark test, wage growth is regressed on ΔU_t and $\Delta U^{\min}$ (see equation (2)). The control variables are differences in cubic polynomials in tenure and experience, in a quadratic time trend, and indicators for industry and region. The last two columns in Table 1 show the results. Wage growth does not respond to ΔU_t for job stayers. If the labor market were spot, then wage growth would be procyclical. It is instead acyclical and displays notable history

[5] The estimates in BD come from the PSID (1976–1984), and those in Grant (2003) use the NLSY (1979–1998).

dependence. The coefficient of $\Delta U^{\min}$ remains significant at -2.64 percent.[6] This is at odds with a spot labor market but consistent with a contractual market, where wages are adjusted whenever the worker's outside option binds.

These findings do not refute on-the-job search. They simply indicate that separations are not a source of cyclical variation in average match quality, and they cannot explain the significance of $U^{\min}$ in Table 1. Cyclical movements in job quality may still arise from new jobs. Gertler, Huckfeldt, and Trigari (2016) recently argue that the procyclical movements in wages of newly hired workers reflect quality differences. The difference specification does not allow us to test this directly since U_0 is differenced away. Nonetheless, cyclicality of job quality for new hires can be gauged by examining how much and how fast the time-of-entry effects in wages fade out with tenure. If contracts were irrelevant, then time-of-entry effects would only reflect differences in match quality. As a result, wage dispersion at the time of entry would be permanent. With contracts, however, wages are updated as the market improves. Consequently, initial differences in wages disappear, as all workers eventually experience favorable conditions on the job.[7] This can be implemented by first computing the differences in wages at the time of hire using the estimates in the fourth column of Table 1 and then simulating an entire wage sequence for each worker by adding the predicted on-the-job wage growth estimated in the last column. Given the estimates, predicted initial wage is $\hat{w}_{i0} = -(0.57 + 3.54) \times U_{i0}$, as initially $U_{i0} = U_i^{\min}$. Consequently, the predicted wage after $\tau = t - t_0$ years on the job is $\hat{w}_{i\tau} = \hat{w}_{i0} - 2.64 \times (U_{i,t0,t}^{\min} - U_{i0})$. Denote the variance of the predicted wages for tenure year τ by ν_τ^{bm}, where *bm* refers to the benchmark estimate. The hypothetical scenario without contracts can be computed by setting $\theta_3 = 0$ instead. Denote this hypothetical variance series by ν_τ^0. If the time-of-entry effects were driven *only* by changes in job quality so that contractual variations did not matter, then the relative dispersion of the simulated wage series ν_τ^{bm}/ν_τ^0 would remain at 1. If, instead, the variation in starting wages entirely reflects differences in contractual terms, then it should fade away completely with tenure.

The solid line in Figure 1 shows the relative variance of predicted wages by tenure. The dashed series correspond to predicted variances using estimates that are two standard deviations below and above the benchmark: -2.64 ± 0.63. The variance of the predicted wage series declines by over 80 percent over time, indicating that at most a fifth of the time-of-entry effects is, in fact, due to variations in average job quality. Even with an elasticity that is two standard deviations below the benchmark estimate, over half the variation dies out.[8] We conclude that the time-of-entry effects in wages largely reflect variations in contractual terms rather than permanent differences in match quality.

[6] One might be concerned that these estimates are based on different samples. When we estimate the fourth column in Table 1 among job stayers, the coefficients are -3.84 (standard error 0.59) for $U^{\min}$ and -0.51 (standard error 0.39) for U_0. This is expected since job switchers do not help identify $U^{\min}$ in Table 1.

[7] Since the stochastic process for the unemployment rate is stationary and bounded from below, $U^{\min}$ converges for all workers.

[8] To address possible issues that may arise from attrition by tenure, we repeated the figure for jobs that lasted at least ten years. The time-of-entry effects for these jobs decline faster, converging similarly to about 20 percent of their initial value after 15 years.

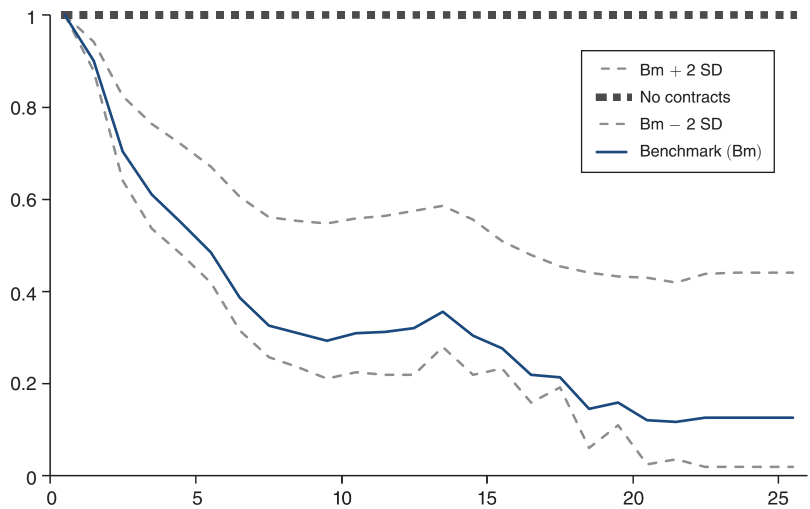

FIGURE 1. PERSISTENCE OF TIME-OF-ENTRY EFFECTS IN WAGES

Notes: Figure shows the relative variance of time-of-entry effects by years of tenure. The solid line is the model prediction, and the dashed lines show the error bounds based on the estimates in Table 1.

The implied user cost of labor, defined by total expected outlays on a worker, is highly procyclical. The elasticity of starting wages to the unemployment rate is -3.21 percent $(= -2.64$ percent $- 0.57$ percent$)$. Figure 1 shows a half-life of five years in initial differences, which is consistent with an annual discount factor of 0.87. Therefore, the semi-elasticity of the total labor cost with respect to the unemployment rate is -10.5 percent over 3 years, -14.0 percent over 5 years, and -19.4 percent over 10 years, conditional on survival.[9] Accounting for the possibility of premature separation lowers these figures. A 25 percent annual separation rate, corresponding to 4 years of expected job tenure, implies elasticities of 7.6 percent, 8.5 percent, and 9.2 percent instead for a 3-, 5-, or 10-year contract. These elasticities are still much higher than 1.6 percent, estimated in column 1 of Table 1 ignoring contractual effects.

IV. Is There Cyclical Selection in Match Quality?

We have established that the observed history dependence in wages is largely explained by wage movements on the job rather than cyclical changes in the composition of jobs. In this section, we provide further evidence to this conclusion and investigate the effect of separations on average job quality by estimating (1) using proxies for job quality.

Table 2 shows the results. The proxy for job quality used in the first column is the cumulative market tightness as in Hagedorn and Manovskii (2013). The estimates confirm their result with an elasticity of 6.02 percent. The second column

[9] $-3.21 \times (1 - \delta^{T+1})/(1 - \delta)$, where T is the duration of the match and $\delta = 0.87$.

Table 2—Real Wages, Unemployment History, and Proxies for Match Quality

	log$w$ (1)	log$w$ (2)	log$w$ (3)	log$w$ (4)	log$w$ (5)	log$w$ (6)
U_t	−0.89	−1.00	−1.07	0.33	−0.14	−0.40
	(0.20)	(0.25)	(0.23)	(0.23)	(0.23)	(0.23)
U_{ij0}				−0.71	−0.77	−0.73
				(0.31)	(0.32)	(0.37)
$U_{ijt}^{\min}$				−2.00	−3.13	−2.75
				(0.43)	(0.49)	(0.50)
$\log(duration \times \bar{\theta}_j)$	6.02			5.61		
	(0.28)			(0.28)		
$\log(duration)$		5.36	3.12		5.10	3.99
		(0.26)	(0.33)		(0.25)	(0.34)
$\log \bar{\theta}_j$		5.05			−3.18	
		(1.34)			(1.55)	
$\log \theta_j^{\min}$			−7.09			−6.03
			(1.22)			(1.18)
$\log \theta_j^{\max}$			14.65			3.66
			(1.47)			(2.23)
Observations	58,967	58,967	58,967	58,967	58,967	58,967

Notes: All specifications control for individual fixed effects, cubic polynomials in experience and tenure, a quadratic time trend, and indicators for industry and region. U_t, U_{ij0}, $U_{ijt}^{\min}$ are contemporaneous, initial, and the minimum unemployment rate experienced since the start of the job until year t. $\bar{\theta}_j$, $\theta_j^{\min}$, $\theta_j^{\max}$ denote the average, minimum, and maximum labor market tightnesses ever observed during job spell j. Coefficients and standard errors are multiplied by 100. Standard errors are clustered by start year and current year interactions and are shown in parentheses.

Source: Data come from the 1979 cohort of the NLSY.

regresses wages on duration and average tightness separately. The elasticity of wages to job duration is 5.36 percent. The coefficient on average tightness is also significant at 5.05 percent, suggesting that separations lead to procyclical selection in job quality. Column 3 replaces average tightness by minimum and maximum tightness levels ever experienced on the job. This allows for a distinction between job selection through layoffs during downturns and through quits during upswings, respectively. The estimates suggest significant selection effects during both times, with a stronger degree of selection during upswings resulting in procyclical match quality on average.

The results are overturned when we include U_0 and $U^{\min}$ in the regression (columns 4–6). The unobserved match quality appears procyclical in column 4. A different result emerges when job duration and average tightness are included separately. While the coefficient on job duration remains significant at 5.10 percent, the coefficient on average tightness becomes significantly negative, indicating countercyclical selection in job quality in evidence of the cleansing effects of recessions. In the last column, we replace average tightness by minimum and maximum tightness during a job spell. This allows us to test for the effects of quits and layoffs separately. The coefficient on minimum tightness is negative and significant, in support of job cleansing during recessions, whereas the coefficient on maximum tightness is half as large and statistically insignificant, implying a limited role for job cleansing during expansions.

Overall, Table 2 shows that average match quality in continuing jobs is somewhat countercyclical and that the evidence on implicit insurance contracts is robust to selection effects in match quality. Past unemployment rates are significantly negative in all three columns, whereas the current unemployment rate is statistically insignificant. In fact, the -2.75 percent estimate on $U^{\min}$ in the last column, obtained when our proxies for match quality are included in the regression, is very close to -2.64 percent, our estimate obtained from the sample of job stayers (Table 1, column 6). This shows that when a flexible specification is used, the proxy method yields an elasticity similar to the one implied by the wage growth of job stayers.

In the online Appendix, we show that the latter conclusion is robust to various sampling restrictions. However, countercyclicality of job cleansing is sensitive to the inclusion of part-time or secondary jobs. Among full-time workers, we find the selection effects in job quality to be acyclical.

V. Empirical Concerns and Extensions

We subjected the benchmark estimates based on job stayers reported in Table 1 to a series of sensitivity checks. A detailed analysis of robustness to sample selection and an analysis of cyclical changes in human capital accumulation can be found in the online Appendix. Here, we discuss potential concerns that might arise due to nonrandom selection of job stayers or permanent differences in wage growth. The upshot of our analysis below is that while these factors are generally at play, they do not correlate with the measures of past economic conditions and hence do not affect our findings.

A. *Are Job Stayers Special?*

One might be worried that focusing on job stayers biases our results. It is possible, for instance, that low-wage workers who experience negative productivity shocks quit their jobs, leading to higher observed wage growth among stayers. The same may not apply to high-wage workers if they are further from their reservation wage. Generally, this affects the estimated intercept in the wage growth equation but is not a concern for the estimated elasticity to $U^{\min}$, our main variable of interest, unless selection of stayers is somehow history dependent. In particular, it would be a concern only if, given the contemporaneous conditions, selection were more stringent for people who, for historical reasons, had a larger decline in their minimum unemployment rate.

This concern can be addressed by the two-step Heckman correction procedure, which requires an exclusion restriction for robust identification. Under our assumptions, match quality does not enter equation (2) and hence measures of match quality satisfy the exclusion restriction. We use the following proxies for match quality: total job duration, $\log \theta_j^{\min}$, and $\log \theta_j^{\max}$. Since these proxies are constant within a match, they do not affect wage growth under the null hypothesis, but they predict the probability of staying on the job by construction: those with a low match quality are more likely to switch jobs and therefore have a shorter job duration, for instance. The marginal effects of job duration, $\log \theta_j^{\min}$ and $\log \theta_j^{\max}$, in the first step are 32.4

Table 3—Real Wage Growth and Unemployment History: Job Stayers

	$\Delta \log w$ (1)	$\Delta \log w$ (2)	$\Delta \log w$ (3)
ΔU_t	0.25	0.42	0.29
	(0.24)	(0.25)	(0.23)
$\Delta U_{ijt}^{\min}$	−2.54	−2.62	−2.46
	(0.63)	(0.71)	(0.71)
Inverse Mills Ratio	−2.17		
	(0.52)		
Worker fixed effects	no	yes	yes
Job fixed effects	no	no	yes
Sample size	40,145	30,228	30,863

Notes: All specifications control for differences in cubic polynomials of experience and tenure, differences in a quadratic time trend, and indicators for industry and region. Estimates in column 1 are corrected for nonrandom sample selection using the Heckman correction procedure. Sample includes job stayers only. Coefficients and standard errors are multiplied by 100. Standard errors are clustered by start year and current year interactions.
Source: Data come from the 1979 cohort of the NLSY.

(standard error 2.9), −43.3 (standard error 4.8), and 64.5 (standard error 9.1) percent. Table 3 shows the estimation results from the second step in the first column. The coefficient on the inverse Mills ratio is −2.17 percent (standard error 0.52), implying that job stayers have a higher wage growth on average. However, the coefficient of $\Delta U^{\min}$ remains at −2.54, implying that such selection is not related to past economic conditions.

B. *Anticipated Wage Growth*

The benchmark specification assumes the predictable component of job quality to be time invariant. However, differences in career paths, firm-specific human capital, or varying degrees of moral hazard could lead to predictable variations in wage growth across jobs. Similarly, differences in the ability to accumulate human capital could lead to persistent differences in wage growth across workers (Haider 2001, Guvenen 2007). Ignoring this variation could lead to an endogeneity problem in the estimation of (2) if survival of jobs were dependent on anticipated wage growth. This can be addressed by including worker and job fixed effects in our regressions. This requires at least three wage observations per job. The estimation thus necessarily leaves out jobs with very short durations and workers who have just started their career.

The results are shown in the last two columns of Table 3. When fixed worker effects are included, the coefficient on $\Delta U^{\min}$ remains the same. This is not too surprising. Since market experience is rewarded equally at all jobs, the decision to switch jobs does not depend on worker-specific characteristics. In the last column, where we control for job fixed effects, the coefficient on $\Delta U^{\min}$ is −2.46, indicating that past economic conditions are not a major factor for interpreting the cyclical relation between separations and anticipated wage growth. Meanwhile, the coefficient on the contemporaneous change in the unemployment rate is close to zero and insignificant in all of the specifications in Table 3.

VI. Discussion

Both the results from wage regresssions with proxy measures for job quality and those from wage growth regressions using job stayers show the prevalence of wage contracts and are inconsistent with a spot labor market. The pattern of history dependence is consistent with implicit insurance contracts with labor mobility, where wages are rigid downward but flexible upward. Accounting for such contractual effects suggests that labor costs are highly procyclical.

The results also provide an insight into the cyclical workings of job reallocation. They imply that the relative job quality of stayers to leavers remains roughly stable over the business cycle. This suggests that the cyclical movements in average match quality occur mostly through newly hired workers. Our calculations indicate that about 20 percent of the history dependence in wages can be explained this way. By contrast, we find no evidence for procyclical movements in average match quality through separations. The results suggest that poor matches are equally, if not more, likely to dissolve during recessions as they are during expansions.

REFERENCES

Abraham, Katharine G., and Henry S. Farber. 1987. "Job Duration, Seniority, and Earnings." *American Economic Review* 77 (3): 278–97.

Baker, George, Michael Gibbs, and Bengt Holmstrom. 1994. "The Wage Policy of a Firm." *Quarterly Journal of Economics* 109 (4): 921–55.

Barlevy, Gadi. 2002. "The Sullying Effect of Recessions." *Review of Economic Studies* 69 (1): 65–96.

Basu, Susanto, and Christopher L. House. 2016. "Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models." In *Handbook of Macroeconomics*, Vol. 2, edited by John B. Taylor and Harald Uhlig, 297–354. Amsterdam: North-Holland.

Beaudry, Paul, and John DiNardo. 1991. "The Effect of Implicit Contracts on the Movements 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, Part 2): 9–49.

Bellou, Andriana, and Bariş Kaymak. 2021. "The Cyclical Behavior of Job Quality and Real Wage Growth: Dataset." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E119151V1.

Bils, Mark J. 1985. "Real Wages over the Business Cycle: Evidence from Panel Data." *Journal of Political Economy* 93 (4): 666–89.

Bureau of Labor Statistics. 2009. "National Longitudinal Survey of Youth 1979 Cohort, 1979–2016 (Rounds 1–22)." Produced and distributed by the Center for Human Resource Research (CHRR), The Ohio State University.

Caballero, Ricardo J., and Mohamad L. Hammour. 1994. "The Cleansing Effect of Recessions." *American Economic Review* 84 (5): 1350–68.

Davis, Steven J., and John Haltiwanger. 1992. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." *Quarterly Journal of Economics* 107 (3): 819–63.

Galindo da Fonseca, João, Giovanni Gallipoli, and Yaniv Yedid-Levi. 2020. "Match Quality and Contractual Sorting." *Labour Economics* 66: 101899.

Gertler, Mark, Christopher Huckfeldt, and Antonella Trigari. 2016. "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires." National Bureau of Economic Research Working Paper 22341.

Grant, Darren. 2003. "The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from the National Longitudinal Surveys." *ILR Review* 56 (3): 393–408.

Guvenen, Fatih. 2007. "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review* 97 (3): 687–712.

Hagedorn, Marcus, and Iourii Manovskii. 2013. "Job Selection and Wages over the Business Cycle." *American Economic Review* 103 (2): 771–803.

Haider, Steven J. 2001. "Earnings Instability and Earnings Inequality of Males in the United States: 1967–1991." *Journal of Labor Economics* 19 (4): 799–836.

Harris, Milton, and Bengt Holmstrom. 1982. "A Theory of Wage Dynamics." *Review of Economic Studies* 49 (3): 315–33.

Kahn, Lisa B. 2010. "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy." *Labour Economics* 17 (2): 303–16.

Kudlyak, Marianna. 2014. "The Cyclicality of the User Cost of Labor." *Journal of Monetary Economics* 68: 53–67.

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.

Solon, Gary, Robert Barsky, and Jonathan A. Parker. 1994. "Measuring the Cyclicality of Real Wages: How Important Is Composition Bias." *Quarterly Journal of Economics* 109 (1): 1–25.

Thomas, Jonathan, and Tim Worrall. 1988. "Self-Enforcing Wage Contracts." *Review of Economic Studies* 55 (4): 541–54.

This article has been cited by:

1. Mark Bils, Yongsung Chang, Sun-Bin Kim. 2022. How Sticky Wages in Existing Jobs Can Affect Hiring. *American Economic Journal: Macroeconomics* **14**:1, 1-37. [Abstract] [View PDF article] [PDF with links]