Real Wage Growth over the Business Cycle: Contractual versus Spot Markets

Andriana Bellou
Université de Montréal

Barış Kaymak‡
Université de Montréal

Abstract

We derive and test the implications of implicit insurance contracts for *wage growth within a job spell*. By focusing on job stayers, we eliminate the confounding effect of changes in the cyclical composition of unobserved job quality, thereby addressing a major concern for earlier studies. We find that workers hired during recessions, or those who experienced unfavorable economic conditions since they were hired, receive larger wage raises during expansions, and are subject to smaller wage cuts during downswings. The change in the contemporaneous conditions, on the other hand, is not a significant determinant of wage growth. Our findings are consistent with a model of implicit insurance contracts without commitment. Our estimates imply that cyclical variations in average job quality explain only *a fifth* of the time-of-entry effects in wages.

J.E.L. Codes: E32, J31, J41
Keywords: Business Cycles, Wage Rigidity, Implicit Contracts, Job Quality

‡We thank Mark Aguiar, Paul Beaudry, Mark Bils, Rui Castro, Iourii Manovskii, Lyndon Moore and seminar participants at the Society of Economic Dynamics (2010), SOLE (2011), CEA meeting (2010), University of Montreal, University of Rochester and UQAM for their comments.
Department of Economics, Université de Montréal, C.P. 6128 succursale Centre-ville, Montréal, QC H3C 3J7. e-mail: andriana.bellou@umontreal.ca (corresponding author) and baris.kaymak@umontreal.ca

The prominent models of the labor market in modern macroeconomics assume that the exchange of labor services with compensation takes place within a short time period as in a spot market (see, for instance, Kydland and Prescott (1982) or Long and Plosser (1983)). Therefore, in a competitive labor market, a worker's wage equals his marginal product at all times, and, hence, changes in the wage rate can be explained entirely by contemporaneous changes in economic conditions. A similar prediction comes out of models where workers and firms constantly renegotiate a production surplus (see Mortensen and Pissarides (1994)). Most employment relationships are, however, long in nature. This provides potential welfare gains to decoupling wages from productivity in the short-run. For instance, firms may shield workers against arbitrary movements in their marginal product by underpaying them during expansions, and overpaying them during downswings, thereby providing a more stable flow of income. This is the idea behind the models of implicit wage contracts (Azariadis, 1975; Baily, 1974).[1]

Understanding the structure of the labor market is crucial. For instance, standard models of the business cycle with spot markets fail to explain the high volatility of employment and labor hours relative to the small movements in the measured wage rate. Incorporating contractual markets into the standard model brings theory closer to the data (Boldrin, 1995).[2] This is because the temporal fluctuations in productivity, which are critical for the allocation of hours, are not reflected in the contract wage. In such an environment, the standard estimates of the intertemporal labor supply elasticity (see MaCurdy (1981) among others) can also be misleading. The distinction, therefore, is essential, not only for evaluating macro models, but also for calculating welfare or gauging labor supply responses in public policy debates.

In a contractual market wages carry information about the economic conditions when the contract was (re)negotiated. Consequently, wages are history-dependent unlike in a spot market where they depend only on current conditions. A growing body of evidence point to the relevance of past labor market conditions for wages.[3] Beaudry and DiNardo (1991), in particular, find that workers who experienced better economic conditions since they started their jobs have higher wages on average, and that contemporaneous conditions are irrelevant.[4] This empirical pattern, apparently inconsistent with a spot market

[1]See Rosen (1985) for a survey of the implicit contracts literature.

[2]See also Rudanko (2009, 2010) and Hall (2005) for more recent evaluations of contracts and wage rigidity on unemployment fluctuations.

[3]Oreopoulos, Wachter, and Heisz (2006), for instance, find a persistent negative effect on wages of being hired during a recession. See also Kahn (2010); Freeman (1981); Baker, Gibbs, and Holmstrom (1994) for similar findings on cohort-entry effects in wages.

[4]See also Grant (2003) and Kudlyak (2010) for similar findings for U.S., McDonald and Worswick (1999) for Canada, and Bellou and Kaymak (2010) for a set of European countries.

model of wages, was considered as evidence for contractual arrangements between workers and their employers.

Nevertheless, one may be too quick to dismiss the spot market model based on this evidence. It is plausible that jobs that start during recessions are of particularly low quality (Okun, 1973). This argument is somewhat backed up by the finding that the expected duration is shorter for jobs that start in recessions (Bowlus, 1995). Similarly, workers in weak matches may quit their jobs in pursuit of better matches, leading to the selection of more productive employee-employer pairs over time (Topel, 1991). If this selection were especially stringent during an upswing, when there are plenty of job vacancies, then the jobs that survive economic expansions would be particularly productive. In a recent contribution Hagedorn and Manovskii (2010) argue indeed that the findings in Beaudry and DiNardo (1991) simply reflect such selection. Building on a search model of job selection with spot markets, they devise various proxies for job quality, and find that wages do not display history-dependence once the cyclical variation in average match quality is properly accounted for, upon which they reject models of insurance contracts.

In this paper, we propose a novel test of the spot market model against the implicit contracts model. Our identification strategy relies on the implications of the two models for wage growth in response to a change in economic conditions for workers who *do not* switch jobs. By focusing exclusively on job stayers, we are able to control for the confounding effect of match quality under the assumption that the latter is time invariant for a given employer-employee pair. Our approach eliminates the need to rely on proxies for job quality, thereby avoids any potential measurement issues that might arise from using proxies.

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 rate than worker R over 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 level of productivity, he's the first to try to leave given a set of offers.[5] Consequently, to prevent severance, the employer is more likely to offer a raise to worker R, or, to offer him a larger wage raise relative to worker B. Thus R's expected wage gain is larger than B's. If, on the other hand, the spot market model better describes the wage behavior, both workers should be paid equally at all times since they are equally produc-

[5]We assume that some degree of worker mobility is allowed, and that workers cannot sign a contract that ties them to an employer regardless of their outside options.

tive. Hence, there is no reason to believe that the wage adjustments over the business cycle should depend on the economic conditions at the time of the hire.

We test several versions of this argument. Our benchmark null hypothesis is the spot market model with time-invariant heterogeneous match quality. Our aim is to test whether workers who are paid higher insurance premiums conditional on productivity, receive lower wage raises during upswings or larger wage cuts during downturns relative to workers with lower insurance premiums.[6] To this end, we first focus on self-enforcing wage contracts where the firm can credibly commit to future payments foreseen by the contract, but the worker cannot guarantee to stay with the firm (Harris and Holmstrom, 1982). In this case, the differences in the insurance premiums can be identified by the initial and the best economic conditions since the start of the job.

Our results show that the wage growth of job stayers within an employment spell is history dependent. Workers who were hired during expansions, or those who experienced better economic conditions on the job have lower wage growth on average. The contemporaneous change in the unemployment rate, on the other hand, is not a significant determinant of wage growth for job stayers. This is at odds with the spot market model independently of cyclical variations in job quality.

One could be concerned that job stayers are selected non-randomly over the business cycle. If jobs that start in recessions were more likely to be shed in response to a negative productivity shock, then a *faux* catch-up effect would be created by positive selection of workers at the bottom. Nevertheless, we find that this is not true. Although average wage growth is somewhat lower when corrected for non-random sampling, the bias does not depend on a worker's history, and, therefore, does not effect our findings.

We extend our results in three crucial dimensions. First, we relax the assumption that match quality is constant. In particular, we define each job by two components: an initial job-specific productivity and an associated wage growth. A job-specific component in wage growth could confound our results if the survival of job stayers over the business cycle is endogenous to anticipated wage growth. Our analysis shows, however, that our tests are biased towards the spot market model if such endogeneity in growth rates is ignored. Using workers with several wage observations on the same job, we estimate the extended version of our model, and find that the history dependence in wage growth is indeed stronger. The effect of past labor market conditions on wage growth is large, but not persistent. Time-of-entry effects fade out rapidly with a half-life of 2 years. If the

[6]Technically, an insured worker receives indemnity whenever his wage is higher than his productivity, and pays a premium otherwise. We use the term premium more generally to denote the gap between wage and productivity in contractual markets. It is understood that the premium can be negative.

cohort-entry effects in wages were driven entirely by permanent job-specific factors, such fading out would not occur. This reinforces the evidence on contractual markets without commitment.

Second, we study a contractual market where neither the worker nor the employer can credibly commit to the contract (Thomas and Worrall, 1988). In this case, the optimal contract cannot be summarized by extremum moments, such as the best or the worst unemployment rate. We propose a new framework for testing for more general forms of contractual arrangements, without relying on the typically used indicators. First, we capture the contractual variation in wages that is due to differences in the history of economic conditions since the job started. Then we project wage growth on the estimated insurance premium. We find that workers who receive lower wages conditional on productivity due to differences in the time of signing the contract, have higher subsequent wage growth as predicted by the implicit contracts model. If the history-dependence were otherwise driven by differences in job quality, there would be no reason to expect a lower wage growth for high quality jobs. On the contrary, our estimates imply a positive correlation between job quality and subsequent wage growth.

Finally, we explore whether our results are driven by cyclical variations in human capital by controlling for available training measures in the data. Our estimates show that training investment is an important component of wages, but it does not display a cyclical history-dependence. Therefore, it has no effect on our findings.

Overall, we view that our paper contributes to the literature on the cyclical behavior of wages by addressing a serious concern for the existing evidence on implicit contracts raised in Bowlus (1995) and recently revived by Hagedorn and Manovskii (2010): that the observed history-dependence in wages can potentially be explained by cyclical changes in the composition of job quality. We find, however, that the wage growth of job stayers displays significant history dependence, even after accounting for persistent differences in job quality. This finding is robust to a broad class of selection models where match quality is not necessarily constant over time, and it suggests that a model of wage contracts without commitment provides a better description of wages. We believe, therefore, that it would be yet impetuous to dismiss contractual markets altogether.

Our results do not entirely refute cyclical variations in the composition of job quality as they allow for pro-cyclical changes in average match quality for newly hired workers. Nevertheless, our calculations indicate that about 20% of the history dependence in wages can be explained in this way. By contrast, we find no evidence for cyclical changes in average match quality *on-the-job* through endogenous separations as argued in Hagedorn and Manovskii (2010).

In the next section, we outline a generic model of implicit insurance contracts and contrast its implications for wages with a spot market model with heterogeneous job quality. We describe our empirical tests in section 2. Section 3 presents our results. In section 5, we provide a discussion of our results and investigate whether models of human capital and training could plausibly explain our findings. Section 6 concludes.

1 Market Structure and the Cyclical Behavior of Wages

The marginal product of a worker i in job j at time t is given by

$$p_{ijt} = \beta_{ij} T_{ijt} + \phi_i X_{it} + \epsilon_{ijt}, \tag{1}$$

where T_{ij} is job tenure and X_{it} is years of market experience. The return to job seniority is specific to the match, and the return to experience differs among workers. The unobserved random error, ϵ_{ijt}, can be decomposed into four:

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

where y_t is the cyclical component of productivity, a_i is a worker-specific component and μ_{ij} is a persistent match-specific component. We assume that ν_{ijt} is outside the model and is independent of other components of productivity.

1.1 Self-Enforcing Insurance Contracts and the Distribution of Wages

The implicit contracts literature deals with arrangements where a risk-neutral employer insures a risk-averse worker against temporary fluctuations in marginal product. We assume that workers do not have alternatives means of insurance. While we make these apparently restrictive assumptions to simplify our analysis, they are not necessary. It would suffice to assume that firms have better (or less costly) means of diversifying risk, and that, when alternative means of insurance are available to the worker, they do not completely crowd out the possibility of insurance through the employer.

We consider contracts that are conditional on long-term changes in productivity, for instance due to accumulation of experience or seniority.[7] There are two sources of risk in equation (1): the time-specific component y_t, and the idiosyncratic component ν_{ijt}.

[7] While the contracts below may be improved upon by arrangements that provide consumption smoothing over the life-cycle, the analysis of such contracts calls for a proper treatment of moral hazard and hold-up problems (Lazear, 1979; Hashimoto, 1981), which are beyond the scope of our analysis.

We assume that the idiosyncratic risk can be completely eliminated across workers, and focus on contracts that aim to insure the worker against the intertemporal fluctuations in his productivity.[8]

Let $h_t = \{y_t, y_{t-1}, ..., y_{t_0}\}$ denote the history of time-specific marginal products on a job that started at $t = t_0$. A contract is a sequence of functions $w_t(h_t; \mathbb{X}_{ijt})$ for $t \geq t_0$ that assigns a history-dependent stream of wage payments to the worker. $\mathbb{X}_{ijt}$ denotes the set of additional productivity traits in equation (1). Let $v(y_t; \mathbb{X}_{ijt})$ denote the worker's utility from quitting his job and pursuing his career elsewhere. Denote the corresponding outside option for the employer by $\pi(y_t; \mathbb{X}_{ijt})$. Both options are functions of current conditions, i.e., they are forward looking. Also, both outside options are allowed to be conditional on individual productivity traits. In what follows, we suppress this dependency to save on notation. It should be understood, however, that the arguments below are conditional on $\mathbb{X}_{ijt}$. Future flows of wages and productivity are discounted geometrically at rate $\delta < 1$, which may reflect the possibility of exogenous termination of the employment relationship as well as the interest rate and time preferences. The optimal *self-enforcing* contract solves the following problem:

$$\max \sum_{k=1}^{\infty} \delta^k E_{t_0}[U(w_{t_0+k}(h_{t_0+k}))] \tag{2}$$

subject to

$$\sum_{k=1}^{\infty} \delta^k E_{t_0}[y_{t_0+k} - w_{t_0+k}(h_{t_0+k})] = 0 \tag{3}$$

$$E_t\left[\sum_{k=1}^{\infty} \delta^k U(w_{t+k}(h_{t+k}))\right] \geq v(y_t) \qquad \forall h_t \text{ and } t \geq t_0 \tag{4}$$

$$E_t\left[\sum_{k=1}^{\infty} \delta^k [y_{t+k} - w_{t+k}(h_{t+k})]\right] \geq \pi(y_t) \qquad \forall h_t \text{ and } t \geq t_0. \tag{5}$$

The efficient contract maximizes the welfare of the worker subject to three constraints. Equation (3) is the zero profit condition implied by the free entry assumption for firms. The inequality (4) invokes the worker's incentive compatibility constraint: at any time, the optimal contract is such that the worker prefers to honor the contract, given his outside option $v(y_t)$. The inequality (5) states a similar condition for the firm, where $\pi(y_t)$

[8]This assumption is not far from reality. Using matched employee-employer data from Italy, Guiso, Pistaferri, and Schivardi (2005) find that, firms absorb the idiosyncratic risk in productivity fluctuations, but only partially insure against aggregate movements at the firm level.

denotes the firm's outside option. We assume that the constraint set is non-empty.[9]

The problem above describes a setting where neither the worker, nor the firm can commit to honoring the contract in the future. At the other extreme, when both can credibly commit to the conditions of the contract, for instance, when mobility is too costly, the optimal contract solves (2), subject to constraint (3) only. If firms can commit to future payments, but the workers cannot, then the optimal contract maximizes (2), subject to (3) and (4), and so on. The following proposition gives the characterization of the optimal contract under two-sided lack of commitment.

Proposition 1 (Thomas and Worrall, 1988) *For any history $h_t = (h_{t-1}, y_t)$, there exist functions $\overline{w}(y_t)$, and $\underline{w}(y_t)$ such that, the optimal contract wage, $w_t(y_t)$ is*

$$w_t(h_{t-1}, y_t) = \begin{cases} \overline{w}(y_t) & \text{if } w_t(h_{t-1}) > \overline{w}(y_t) \\ w_t(h_{t-1}) & \text{if } \overline{w}(y_t) \geq w_t(h_{t-1}) \geq \underline{w}(y_t) \\ \underline{w}(y_t) & \text{if } w_t(h_{t-1}) < \underline{w}(y_t). \end{cases}$$

The optimal contract keeps the wage constant from one period to another, as long as the participation constraints do not bind. If the outside option of an agent changes substantially, so much as to render employment unsustainable at the previously agreed wage rate, the optimal contract calls for a wage adjustment to prevent separation. If, for instance, the worker receives a better wage offer after a contract is signed, then the firm offers a raise just enough to retain the worker.

1.1.1 The cross-sectional distribution of wages

When both the worker and the firm can credibly commit, the optimal contract features full insurance, i.e. a constant wage rate. Consequently, by equation (3), the wage rate must be equal to the expected productivity of the worker at t_0. This leads to a cross-sectional dispersion in wages, based only on differences observed at t_0 across workers, even for workers with the same marginal product at time t. The economic conditions at the start of a job are sufficient to capture this variation, leaving current economic conditions statistically redundant.

If firms can commit to future payments, but workers cannot, then the optimal contract features a constant wage, that increases only if the worker's outside option has sufficiently improved. Since firms commit to payments, wages are never adjusted downward. A worker's wage at time t reflects the highest wage he could command since the start

[9] The reader is referred to Harris and Holmstrom (1982) or Thomas and Worrall (1988) for the conditions that ensure the existence of an optimal contract in this setting.

of the job. The cross-sectional variation in wages reflects not only the conditions at t_0, but also the best economic conditions since then. These two moments exhaust all the variation in wages.

In the more general case, when neither party can commit, the contracted wage rate moves with marginal productivity, only when the latter is altered substantially, and remains constant otherwise. Current wages are still history dependent, but the form of this dependence cannot be represented by simple moments as above. Nevertheless, an indicator for the start year and the current year pair, (t_0, t), is sufficient to summarize the entire history of economic conditions between t_0 and t, and, hence, the cross-sectional distribution of wages.[10]

In all of these set-ups, the current wage distribution displays history-dependence, and current economic conditions do not matter, in sharp contrast to the spot market model. If different components of productivity are not measured correctly, however, a superficial history dependence in wage levels could also be observed in a spot market model. Regardless, a crucial distinction remains in the predicted wage paths of individual workers, which we turn to next.

1.1.2 Implications for wage growth

In contractual markets without commitment, wages are adjusted in response to changes in outside options. The incidence and the extent of an adjustment depends on the wage of the worker relative to his productivity, and thereby, on the history of economic conditions. Taking differences in Proposition 1, the wage growth along the optimal contract described by:

$$\Delta w_t(h_{t-1}, y_t) = \begin{cases} \overline{w}(y_t) - w_t(h_{t-1}) & \text{if } w_t(h_{t-1}) > \overline{w}(y_t) \\ 0 & \text{if } \overline{w}(y_t) \geq w_t(h_{t-1}) \geq \underline{w}(y_t) \\ \underline{w}(y_t) - w_t(h_{t-1}) & \text{if } w_t(h_{t-1}) < \underline{w}(y_t). \end{cases}$$

Whenever one of the incentive constraints binds, the wage is adjusted to reflect the reservation wage of the agent, which is forward looking, and, hence, depends only on the current conditions. Differences in wages across workers that arise from differences in market history are annulled when the reservation wage binds. This is the memoryless

[10] A common attempt to capture the two-sided lack of commitment is to include the maximum unemployment rate since the start of the job in the regression, along with the initial and the minimum unemployment rates. While the maximum unemployment rate may capture some additional variation in wages, this specification does not exactly correspond to contracts with two-sided lack of commitment, and makes it hard to give a meaningful interpretation to the coefficients. It is only a sufficient statistic for the case where workers can be enslaved by employers who cannot credibly commit to keeping them.

nature of contracts without commitment (Kocherlakota, 1996). The wage growth, therefore, depends *negatively* on the last period's wage whenever there is an adjustment. This generates history dependence in wage differences, which is absent in the spot market model.

The nature of this dependence is simple. As an example, Figure 1 shows wages over a boom cycle in a contractual market without commitment. The economy goes from an average (a) state to a boom (b) followed by a recession (r). Since there is no commitment, wages lie between $\underline{w}(y_t)$ and $\overline{w}(y_t)$ at all times. The cross-sectional variation of wages conditional on productivity reflects different insurance premiums received by workers with different labor market histories. When the economy moves into an expansion from t to $t+1$, the participation constraint becomes binding for workers at the low end of the premium distribution, leading to wage raise. Meanwhile, the wages of workers with higher insurance premiums remain unchanged. Similarly, when the economy enters a recession at $t+1$, the employer's participation constraint binds, especially for high wage workers. These workers receive wage cuts, while those with lower wages are spared.

In both cases, the boom and the recession, the optimal contract calls for a smaller wage increase for workers who receive larger insurance premiums over their productivity. Each worker's position in the cross-sectional distribution of premiums in turn depends on the past labor market conditions. Workers who were hired during expansions, or those who experienced more favorable conditions since they were hired, find themselves receiving higher wages conditional on productivity, and are, therefore, subject to lower wage increases in general. This process mitigates the initial time-of-entry effects in wages over time.

That the insurance premiums converge with tenure is implicit in the formulation of the contracting problem. The outside options in inequalities (4) and (5) are forward looking, and, thus, depend only on the current economic conditions. The only difference between wage adjustments come from the existing wage payments, which are predetermined by the contract, and are dependent on the history. Unlike the predictions of contracts model for wage levels, this is robust to selection of jobs by quality over the business cycle as we demonstrate next.

1.2 A Spot Market Model of Wages with Cyclical Selection of Jobs

In a spot market model wages equal marginal product at all times. Given equation (1), wages can be captured by the following equation:

$$w_{ijt} = \beta_{ij} T_{ij} + \phi_i X_{it} + y_t + a_i + m_{ij} + \nu_{ijt} \ , \tag{6}$$

where $m_{ij} < \mu_{ij}$ denotes the part of the match surplus that is retained by the worker. Structurally, wages in a spot market depend only on the contemporaneous conditions. Still, history dependence can arise in a spot market if the econometrician fails to control for all components of productivity, and if these components are correlated with past economic conditions. Since tenure and experience are observed, y_t is captured by time effects, and the fixed worker effect can be uncovered using panel data, we focus on the omission of m_{ij}. The following example illustrates how the cyclical changes in the composition of jobs can lead to a superficial history-dependence in a spot market with endogenous quits.

1.2.1 Endogenous quits and cyclical selection

Suppose that $\beta_{ij} = \phi_i = 0$ for simplicity, which implies that $w_{ijt} = y_t + a_i + m_{ij} + \nu_{ijt}$.[11] Every period workers draw offers from a stationary distribution. We assume that offers are drawn 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, as in the previous section. A worker decides to quit his existing job if $\tilde{w}_{it} > w_{ijt}$. Since the cyclical and the worker-specific components are equally valuable in all jobs, and since $E[\nu_{ijt}] = 0$, a better wage offer must come from a better match.[12] The endogenous quit decision leads to destruction of poor matches over time, leading to the survival of only the best ones. Denoting the match quality corresponding to the best wage offer by $\tilde{m}_{it}$, the average match quality conditional on $T = t - t_0$ years of seniority is

$$E[m_{ij} | m_{ij} > \max\{\tilde{m}(y_{t_0}), \ ..., \ \tilde{m}(y_t)\}]. \tag{7}$$

The number of arguments in the max operator above increases with tenure, raising the average quality of surviving matches. When the match quality is not observed, this confounds the estimates of the return to seniority (Topel, 1991).

If the wage offer $\tilde{w}_{it}$ displays cyclical variation, a similar selection argument could also generate a *faux* history-dependence. For instance, if the wage offer $\tilde{w}_{it}$ is pro-cyclical, the selection in (7) applies more stringently to those workers who experienced better economic conditions. This creates a negative relation between the minimum unemployment rate experienced over a worker's tenure and the expected match quality.

Furthermore, the pro-cyclicality of the wage offer, $\tilde{w}_{it}$, transcends that of the match

[11] This version of the model is similar to the one considered in Hagedorn and Manovskii (2010).

[12] Otherwise the current employer could retain the worker by outbidding the outside offer.

quality $\tilde{m}_{it}$, implying that the average quality of new matches that are formed during expansions are higher. Therefore, jobs that start during expansions command higher wages, which makes the initial unemployment rate at the start of the job a statistically important determinant of wages, even in a spot market!

Is average match quality pro-cyclical? On the one hand, unemployed workers probably sample more offers during expansions, which raises the quality of matches. On the other hand, during recessions, employers sample more applicants, which also may increase the average quality of new matches. It is plausible, however, that a larger portion of the match surplus is captured by workers during expansions, leading to a pro-cyclical match-specific component in wages. In such a case, the implications of the spot market model are similar to those of the contractual model discussed in Section 1.1.1, yet the cyclical nature of match quality remains more of an empirical matter.

1.2.2 Using wage growth of job stayers to distinguish between models

In the spot market model with selection, it is the composition of jobs that is responsive to past conditions, not individual wages. As a result, it can be distinguished from contractual models by studying the wage differences of workers who do not switch jobs. To see this, let $\phi_i = \phi$ and $\beta_{ij} = \beta$ for now, and denote the first difference operator by $\Delta x_t = x_t - x_{t-1}$. Using (6) to calculate the change in the wages of job stayers:

$$\Delta w_{ijt} = \phi + \beta + \Delta y_t + \Delta \nu_{ijt}. \tag{8}$$

Since the unobserved worker and match quality effects are time invariant, they disappear in the wage growth equation. Contrary to the contractual markets, the wage growth of job stayers does not display any history dependence if the labor market is, in fact, characterized only by a spot market model.

1.2.3 Anticipated wage growth and cyclical selection of job stayers

It may be too restrictive to define job quality simply by a time-invariant match quality. Guvenen (2007) and Haider (2001), for instance, document significant and persistent variation in wage growth across workers. Differences in the training contents of various career paths, or differences in the ability to accumulate human capital (Becker, 1964) could lead to such variation. Similarly, differences in the role of firm specific human capital, or varying degrees of moral hazard problems could lead to variation in seniority premiums.

Using equation (6), the wage growth in the general model is:

$$\Delta w_{ijt} = \phi_i + \beta_{ij} + \Delta y_t + \Delta \nu_{ijt}. \tag{9}$$

Wage growth could display history-dependence if ϕ_i and β_{ij} are correlated with past economic conditions. In particular, differences in wage growth, if anticipated, could affect workers' career choices, leading to a potential endogeneity issue. In this section we show that the decision to switch jobs is independent of ϕ_i but not of β_{ij}. However, while expected match quality, m_{ij}, is *positively* related to having experienced favorable economic conditions in the past, anticipated wage growth β_{ij} is *negatively* correlated with past conditions, contrary to what one would expect in contractual markets. This, in fact, makes it easier to distinguish between the two models as outlined below.

Suppose that each job is characterized by a match quality level, m_{ij}, and an anticipated return to tenure β_{ij}. Let $\tilde{W}_{it}$ be the annualized present discounted value of the best offer a worker obtains at the beginning of year t. The corresponding expected flow value from the existing job for a worker with T years of tenure is:

$$(1-\delta)\mathbb{E}_t \sum_{k=t}^{\infty} \delta^{k-t} p_{ij} = a_i + m_{ij} + \beta_{ij}(T + \frac{\delta}{1-\delta}) + \phi_i(X_{it} + \frac{\delta}{1-\delta}). \tag{10}$$

A worker quits his job if the value in (10) is less than $\tilde{W}_{it}$. Since all firms equally reward general skills, a_i and X_{it}, in a competitive market, the decision boils down to the relative match qualities and the anticipated wage growths in the two jobs. Let $\tilde{w}^a_{it} = m_{ij'} + \beta_{ij'}\frac{\delta}{1-\delta}$ be the match-specific component of the best competing offer by firm j'. A worker quits his job if and only if

$$\tilde{w}^a_{it} > m_{ij} + \beta_{ij}(T + \frac{\delta}{1-\delta}).$$

Two points are worth noting. First, the inequality above is more likely to hold if either the current pay or the anticipated wage growth at the existing job are low. Second, it does not depend on ϕ_i, the return to experience.

Figure 2 contrasts the switching decisions during a recession and an expansion. The circles are isodensity curves for the distribution of match quality m_{ij} and the return to tenure β_{ij}. We assume, for the moment, that these are initially independent. The solid line shows the indifference line for switching jobs during a recession. The workers with (m_{ij}, β_{ij}) combinations that are above the line prefer to stay with their current jobs. Therefore, conditional on staying with the current job, both the match level, m_{ij}, and the expected wage growth are higher than their unconditional means. When the economy is in an expansion, workers sample more offers, leading to more stringent selection (dashed

line) during booms.

That job switches are endogenous implies that the longer the worker has stayed with his job, the more offers he must have sampled, and rejected. Therefore, jobs are selected positively with respect to β_{ij} over years. In addition, conditional on tenure, workers that experienced more favorable economic conditions during their tenure faced more stringent selection constraints, and therefore, must be working at jobs with higher β_{ij}'s on average. This is contrary to the predictions of the implicit contracts model: if a worker experienced more favorable conditions, he already receives a larger insurance premium and, therefore, experiences a *lower* wage growth during expansions.

Equation (9) could be estimated directly. A fixed effects panel estimation would identify ϕ_i at the worker level and β_{ij} at the job level. Since this estimation requires at least three wage observations per job, it necessarily disregards jobs with very short durations and the workers who have just started their career. Nevertheless, we estimate (9) and address the issues with sample selection afterwards. Next, we lay out the details of our empirical strategy.

2 Testing for Contractual Markets: Empirical Implementation

Our purpose is to empirically distinguish between the two models by studying the wage growth of job stayers. A natural way to test the spot markets model is to estimate (8) along with measures of past labor market performance and see if wage growth displays history dependence in a way that is consistent with the implicit contracts model. This constitutes our benchmark test. Following Bils (1985), we measure economic conditions by the unemployment rate. As in Beaudry and DiNardo (1991), we include the unemployment rate at the beginning of a job spell, and the minimum unemployment rates since the start of the job as measures of past economic conditions. The estimated equation is

$$\Delta_k w_{ijt} = \theta \Delta_k u_t + \Delta_k X_{ijt} \Lambda + \gamma_1 u_{t_0} + \gamma_2 u^{min}_{t_0,t-k} + \Delta_k \nu_{ijt}. \tag{11}$$

where $u^{min}_{t_0,t-k} = \min\{u_{t_0}, u_{t_0+1}, ..., u_{t-k}\}$, and X_{ijt} denotes a worker's productivity characteristics such as tenure, experience etc. Δ_k is the k-period difference operator, where k denotes the number of time periods between two consecutive observations of a worker - job pair.[13] If $\Delta_k \nu_{ijt}$ is orthogonal to u_{t_0} and $u^{min}_{t_0,t-k}$, then γ_1 and γ_2 can consistently be es-

[13] In the data, k varies across workers, jobs and time. To save on notation, this variation is suppressed throughout the paper.

timated with an OLS regression. Under the null hypothesis of spot markets, they should be close to zero.

A symmetric approach is to take the wage difference equation from a contractual model, and to test if contemporaneous conditions matter. We use a contractual market with one-sided commitment, where wage growth depends on the *change* in the minimum unemployment rate since the start of the job. We run the following specification in this case:

$$\Delta_k w_{ijt} = \theta \Delta_k u_t + \Delta_k X_{ijt} \Lambda + \gamma_3 \Delta_k u^{min}_{t_0,t-k} + \Delta_k \nu_{ijt} \tag{12}$$

In a contractual market, $\gamma_3 < 0$ and $\theta = 0$, whereas in a spot market, $\gamma_3 = 0$ and $\theta < 0$. The two tests above are similar, with the former testing for more general history dependence and the latter specifically for one-sided commitment contracts. Empirically, the first test is more robust, since it is harder to identify $\Delta_k u_t$ from $\Delta_k u^{min}_{t_0,t-k}$ in the data, especially for workers that are hired during recessions. This is a problem particularly in short panels, where most of the available identification is cross-sectional.

The use of the initial and the minimum unemployment rate, as in the literature, assumes that employers commit to future employment and pay, but workers switch employers at no cost. In more general arrangements where neither side can fully commit, such extremum moments do not capture a sensible variation in wages. We, therefore, develop a second test that is more robust to different contracting schemes.

Our second test draws on the observation that the wage adjustments in a contractual market depend negatively on the insurance premium received by the worker. If a worker already earns a higher premium because he experienced favorable economic conditions while he was on the job, he is subject to more severe wage cuts during recessions and smaller pay raises during booms relative to a worker of equal productivity. To see whether wage differences are consistent with this prediction, we estimate the following equation:

$$\Delta_k w_{ijt} = I(t \times t-k)\Theta + \Delta_k X_{ijt} \Lambda + \gamma W^p_{t-k} + \Delta_k \nu_{ijt} \tag{13}$$

where W^p_{t-k} denotes the insurance component of wages. To identify this component, we first regress wages on the interactions of job-start year and current year indicators controlling for other variables in (13). Given the contemporaneous changes in economic conditions, the combination of start year and current year captures the entire history of economic conditions, and, hence, all possible contracting arrangements defined over the history. This procedure is identical to estimating (13) by replacing W^p_{t-k} with actual wage, W_{t-k}, and using a full set of indicators for all possible $\{t_0, t-k\}$ pairs as instruments for

W_{t-k}. We use the TSLS estimate since it has the added advantage of efficiency.

Essentially, W^p_{t-k} is the average wage in year $t-k$ of all workers hired in t_0. If the wages were described by a spot market model, W^p_{t-k} would capture the cyclical differences in match-specific productivity. There is, however, no reason to believe that workers with higher match qualities should be subject to larger cuts in downswings, and small raises during upswings. If anything, one would expect a positive correlation between match quality and subsequent wage growth, as indicated by observed measures of productivity, such as education.

Note that this strategy is more subtle than simply testing for a negative relationship between wage growth and wage levels. While there are other models that predict such reversion in wages,[14] they do not invalidate the test above unless the nature of this relationship displays a particular cyclical variation. We discuss such possibilities in Section 5.

$I(t \times t-k)$ contains the full set of interaction indicators for the current year and the previous year, and captures the changes in the economic conditions between two consecutive interviews. This also ensures that W^p_{t-k} is identified by indicators pertaining to history strictly prior to year $t-k$.

A potential problem with estimating wage growth regressions with job stayers is that the error term $\Delta_k \nu_{ijt}$ 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 upwards, presumably more so for low wage workers. If, for instance, there is a threshold productivity level, below which the match is dissolved, then workers with low productivity are more likely to hit this threshold. Given the correlation between starting wages and the unemployment rate, this sort of non-random selection could lead to a spurious history dependence in wage growth even in a spot market model.

One way to get around this problem is to make a distributional assumption about $\Delta_k \nu_{ijt}$, and apply the Heckman correction procedure. This requires an exclusion restriction for robust identification. We use three. The first is whether the worker has a working spouse. Those with a working spouse may have limited geographical mobility, and be more likely to stay on the job, conditional on productivity. Second, we use the total duration of the job. Those with a low match quality are more likely to switch jobs, therefore, have a shorter job duration, which, by construction, predicts the likelihood of staying on

[14] For instance, a model of human capital accumulation with on-the-job training, as in (Ben-Porath, 1967), and heterogeneity in the initial capital endowments, would predict that workers with lower wages invest more in their training and enjoy faster wage growth.

the same job. Under our assumptions, match quality does not enter equations (11) and (12), and hence, acts as a valid exclusion restriction.[15] Third, we use the total number of jobs that the worker held per year *after* his current job. A worker who switched jobs multiple times is more likely to have sampled from the unconditional wage distribution on his *current* job. The probability of staying on the current job is lower for such workers. At the same time, there is no reason to expect any dependence between the number of jobs held in the future, and the shocks to wage growth in the current job.

Another potential problem could arise if wage growth contains a permanent match-specific effect as described in section 1.2.3. We address this concern by estimating more general cases of (11) and (13), where we allow for fixed effects in wage growth by worker and then by job. Ignoring the job-specific growth rates leads to under-rejection of our null hypothesis. This is because only jobs with higher growth rates survive favorable economic conditions, whereas the insurance contracts foresee lower growth rates for workers who experienced favorable conditions.

The same is not necessarily true, however, for our second test. Although jobs are selected positively with respect to wage growth, conditional on the past wage rate, the growth rate could be selected negatively. Since workers compare discounted future wages, a workers's decision to stay when he has a low wage at his current job indicates anticipation of fast wage growth. This could lead to a negative correlation between wage levels and wage growth among surviving matches if m_{ij} and β_{ij} are initially independent as depicted in the figure. If this correlation is sufficiently positive, i.e. workers with higher wages, also have high wage growths, then a potential selection bias goes in the opposite direction. Besides, even when this type of selection predicts lower wage growth conditional on current wage, it is not clear, whether this effect systematically depends on the past economic conditions. If not, then this has no consequence for our test. The effect of cyclical selection in growth rates on the estimate of γ is therefore an empirical question. We test the robustness of our results to cyclical selection in growth rates by explicitly allowing for job-specific fixed effects in wage growth.

3 Data and Estimation Results

3.1 Data

The data come from the 1979 cohort of the National Longitudinal Survey of Youth (NLSY) for the years 1979 - 2008. NLSY is a panel that closely tracks workers' jobs, their start dates

[15] This exclusion restriction may fail if $\Delta_k \nu_{ijt}$ is correlated with m_{ij}, a possibility that we address shortly.

and end dates, making it ideal for our purposes. [16] We use the nationally representative cross-sectional sample for our estimations, and restrict the sample to males of 21 years of age or older, working full-time (35+ hours) at the time of the interview in the private sector. We exclude workers with multiple jobs, and those that are enrolled in school. We drop jobs that started before 1976, and before the respondent was 16 years old. Appendix A provides details on the construction of variables.

Table 1 shows the summary statistics. Workers in the NLSY are slightly younger than an average US worker, which also explains the slightly lower job tenure. After our restrictions, there are 2,437 individual workers with 22,329 valid observations. Figure 3 shows the national unemployment rate that we use in our regressions. The sample period starts at the onset of the 1981 recession, and includes three full cycles, providing substantial variation for our estimations.

3.2 Wages and past labor market conditions

We begin our analysis by documenting the history dependence in wage levels. The variables of interest are the contemporaneous unemployment rate, the unemployment rate when the worker started his current job, and the minimum unemployment rate since the beginning of his current job. We include controls for individual fixed effects, cubic polynomials in tenure and experience, and indicators for region and industry. We also include a quadratic time trend to capture any long term relation between average wages and the unemployment rate.

Table 2 shows the effect of past labor market conditions on wage levels. When regressed only on the contemporaneous unemployment rate, wages appear to be strongly procyclical. An increase in the unemployment rate by 1 percentage point is associated with a 1.35% decline in wages. When the initial unemployment rate is introduced, however, the coefficient on the contemporaneous unemployment rate substantially declines, and eventually becomes statistically insignificant in column 3 when the minimum unemployment rate is included in the regression. The specification in column 4 contrasts all three models at once. The past unemployment rates are not only statistically but also quantitatively important. On average, wages decline by 0.80% in response to the unemployment rate at the time of hire, and by a remarkable 2.85% in response to the minimum unemployment rate since then. Assuming a 4 percentage point difference in the initial unemployment rate between peak and trough, this implies a 14.6% (=4*(0.80+2.85) gap in

[16] Another panel used in the literature has been the PSID, however it is much harder to identify the job switches in the PSID. See Kambourov and Manovskii (2008) for a detailed discussion on this issue.

starting wages. The effect of contemporaneous unemployment rate, on the other hand, is virtually zero.

Overall, the estimates confirm the marked history-dependence in wages documented in earlier studies. The coefficient on the minimum unemployment rate was estimated as -2.9% in Beaudry and DiNardo (1991), and -2.5% in Grant (2003). The coefficients on other variables are also similar, with the exception that Grant (2003) finds a stronger effect for the contemporaneous unemployment rate in his sample.[17] Our findings are also consistent with Bils (1985) who finds the wages of job stayers to be acyclical, while the wages of newly hired workers to be highly procyclical.

3.3 Wage growth and labor market history

Is the history-dependence an indication of self-enforcing contractual arrangements in the labor market, or simply an artifact of unobserved match quality? Next, we turn to the wage growth of job stayers to disentangle the two interpretations.

We begin with our benchmark specification where we assume that the wage contracts are characterized by full commitment by the employer, but not by the worker. Table 3 shows the effect of past and contemporaneous unemployment rates on the wage growth of workers who do not change jobs between two consecutive interviews. We start by regressing the wage growth on the initial unemployment rate, and the minimum unemployment rate since the worker started his job until the last wage observation. If wages are determined on the spot, these variables should be insignificant. We control for differences in cubic polynomials in tenure and experience, a quadratic time trend, and indicators for industry and region.

The results indicate that wages of job stayers are mildly procyclical; a one percent increase in the current unemployment rate leads to a 0.26% decline in wages. Nevertheless, the estimate is not statistically significant. By contrast, columns 2 and 3 include the initial and the minimum unemployment rates. Both variables are strong predictors of wage growth. A worker who started his job when the unemployment rate was one percentage point higher, experiences, on average, 0.51% higher wage growth. Similarly, a one percentage point higher minimum unemployment rate is associated with 1.61% additional wage growth. The change in the current unemployment rate, on the other hand, is statistically irrelevant. When we include both the initial and the minimum unemployment rate in column 4, the initial unemployment rate becomes insignificant, indicating

[17] The estimates in Beaudry and DiNardo (1991) come from the PSID (1976 - 1984), and those in Grant (2003) use NLSY (1979 - 1998).

that self-enforcing contracts with one-sided commitment are better descriptions of reality than full-commitment contracts.

If the contractual market with one-sided commitment is the true model, then the *change* in the minimum unemployment rate is a sufficient statistic for wage adjustments. Now, we set our null hypothesis to be the implicit contracts model with one-sided commitment, and we test it by including the contemporaneous changes in the minimum unemployment rates. The results in column 5 show that a one percent increase in the minimum unemployment rate leads to a 3.39% decline in wages of job stayers, whereas the change in the current unemployment rate does not matter. We cannot reject the model of contracts with one-sided commitment.

Overall, our findings are at odds with a pure spot market model of wage determination as contemporaneous changes in the economic conditions do not have a significant effect on the wages of job stayers in any of our models in Table 3. Instead, our results are consistent with a contractual market, where wages are adjusted whenever the worker's outside option binds. Furthermore, workers with higher insurance premiums, because they experienced favorable conditions since they were hired, enjoy smaller wage raises.

Our results do not entirely rule out potential changes in the composition of job quality over the business cycle. In fact, that the initial unemployment rate is significant in column 4 of Table 2, but not of Table 3 suggests that the observed dependence on initial conditions in BD type regressions is, in part, due to lower quality jobs created during recessions. This is in line with the finding that jobs that start in recessions have shorter durations (Bowlus, 1995).

To gauge the extent of cyclical variations in job quality, we computed the predicted initial wage for each job using the estimates in the last column of Table 2. We, then, added the predicted on-the-job wage growth from the one-sided commitment model estimated in column 5 of Table 3.[18] Figure 4 shows the variance of predicted wages by job tenure. The time-of-entry effects vanish by a remarkable 80% after 7 years of tenure. If the time-of-entry effects were driven by changes in job quality, the dispersion would have persisted. The figure shows that the variance converges to 0.005, indicating that only about a fifth of the time-of-entry effects are, in fact, due to variations in average job quality.

Our results also suggest that the cyclical movements in average match quality occurs mostly through newly hired workers, and not through on-the-job selection of match quality via endogenous quits, as modeled in Hagedorn and Manovskii (2010). The coefficient of the minimum unemployment rate in Table 2, -2.85%, remains significant at -3.39% in

[18]Given the estimates, predicted initial wage is: $\hat{w}_{t_0} = -(0.8 + 2.85) * U_{t0,i}$ as initially $U_{t0} = U_{t0,t}^{min}$. Consequently, the predicted wage in year t is: $\hat{w}_{t_0} = \hat{w}_{t_0} - 3.39 * (U_{t0,t}^{min} - U_{t0})$.

Table 3 for job stayers. This implies that the dependence of wages on the minimum unemployment rate was entirely due to contracts. This distinction between compositional changes in job quality through new hires versus through on-the-job selection is not apparent in Hagedorn and Manovskii (2010) since their proxies of match quality therein combine the two components by construction. It is possible that the proxies are picking up on variations in the quality of newly created jobs.[19]

3.4 Are job stayers special?

One might be worried that restricting our tests to job stayers exposes our test to a potential sample selection bias. 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. Same may not apply to high wage workers if they are further from their quit threshold. We address this concern by correcting our estimates for sample selection using a two-step Heckman selection model, where two measures of match quality, job duration and total number future job switches after the current job are used as exclusion restrictions along with the presence of a working spouse.

The last two columns in Table 3 show the results. All three variables are significant determinants of staying on the job. The probability of staying on the job increases with the quality of the job. Those with a working spouse are also more likely to stay on the job, perhaps due to limited geographical mobility of the worker.

The inverse mills ratio has the expected negative sign, implying that job stayers have a slightly higher wage growth. Nevertheless, the coefficient is not statistically significant when the minimum unemployment rate is included in the regression. The effect of past economic conditions on wage growth is virtually unchanged, showing that our results are not driven by sample selection.

3.5 Anticipated wage growth

Our benchmark results abstract from potential match-specific variation in the growth rate of wages. Ignoring this variation could lead to an endogeneity problem if survival of jobs is dependent on anticipated wage growth. We address this concern by including worker and job fixed effects in our regressions. Table 4 shows the results.

[19] One of the proxies, for instance, is the sum of all labor market tightness values for the duration of a job, which is expected to correlate negatively with match quality. This measure could be decomposed into two: duration of a job (as employed by Bowlus (1995)) and the average labor market tightness during the job, which would enable a separate test of on-the-job selection of match quality.

When fixed worker effects are included, the coefficient of the minimum unemployment rate is 1.58%, compared to 1.51% in our benchmark specification. The coefficient on the change in the minimum unemployment rate remains effectively the same. This is consistent with our conclusion in section 1.2.3 on the effect of experience on selection: since market experience is rewarded equally at all jobs, endogenous survival of jobs does not lead to a selection effect in worker-specific wage growth.

In the following columns of Table 4, where we control for job fixed effects, the coefficient on the minimum unemployment rate increases further to 2.01%. This is consistent with the hypothesis that jobs with higher anticipated growth rates are more likely to survive expansions, and suggests that our benchmark estimates were biased towards zero. The coefficient on the change in the minimum unemployment rate, on the other hand, remains similar at -2.96%, indicating that such selection is absent.

To reconcile the seemingly contradictory results, note that the selection through endogenous quits is especially operational for workers that were hired during recessions, and, therefore, are looking to improve their job quality. For such workers, however, the minimum unemployment rate coincides exactly with the current unemployment rate during economic recovery. Therefore, the coefficient on the change in the minimum unemployment is identified mostly during downturns, when quits are especially low.[20] Since in one-sided commitment models, wages are rigid downwards, the coefficient of the *change* in the minimum unemployment rate is unaffected by such selection. The level of the minimum unemployment rate, however, still captures larger real wage reductions for workers with high insurance premium. This distinction comes from the lack of commitment on part of the firms, and becomes more apparent when we consider contracts with two-sided lack of commitment in Section 4

The contemporaneous change in the unemployment rate is close to zero and insignificant in all of the specifications in Table 4.

4 Contracts with Two-Sided Lack of Commitment

We now generalize our test to contracts where neither the worker nor the employer can fully commit to the contract. We examine whether the subsequent wage growth depends negatively on the insurance premium that the worker receives (or pays). This is essentially a test of mean-reversion in the insurance premium. Since, in contractual markets, insurance premiums across workers are dispensable when the participation constraints

[20]Some independent variation of the two variables is also available for workers who experienced at least one complete business cycle on the same job.

bind, workers who start in a disadvantaged position catch up with other workers during expansions. Similarly, those with initial wage advantages lose them in severe downswings.

To implement our idea, we regress wage growth on the lagged wage rate and use the full set of history indicators as instruments for the lagged wage rate. To control for changes in the economic conditions between two consecutive interviews, we control for a full set of interactions of indicators for the current year, t, and the last interview year $t-k$. This ensures that the identification of the insurance premium comes entirely from different histories of economic conditions across job spells. The main advantage of this method is that it is robust to any contractual arrangement that can be defined over the history of economic conditions. It therefore encompasses the case where neither the employer nor the worker can commit to honoring the contract.[21] Another advantage is that we do not need to rely on proxies, such as the unemployment rate or the market tightness, to capture economic conditions.

Table 5 shows our results. As before, we also control for differences in cubic polynomials of experience and tenure, a quadratic time trend, and indicators for industry and region. The estimates for the constant-match-quality model indicate that a worker who enjoys a 1% lower wage rate, for instance, because he was hired in a recession, enjoys a 0.11% larger wage growth on average. When the estimate is adjusted for non-random selection using the Heckman procedure, it remains similar at -0.15 in the second column.[22]

These findings confirm our earlier conclusion. If the wages were described by a spot market model, the variation in wages predicted by the history indicators would correspond to real, match-specific productivity differences between jobs that were selected differently over the business cycle. There is, however, no reason to believe that workers with higher match qualities should be subject to larger cuts in downswings, and small raises during upswings. On the contrary, one would expect larger wage raises in good matches, for instance due to increased investment in job specific capital (Becker, 1964). In addition, given that average wage growth is positively related to observed productivity characteristics, such as education, one would expect the unobserved match quality to be also positively related with wage growth.

The third column adds worker fixed effects in wage growth. The coefficient of lagged wages declines substantially to -31%. This is consistent with positive selection of job stayers with respect to anticipated wage growth. It implies that the labor force is composed

[21] A similar strategy was employed in Beaudry and DiNardo (1995) to estimate the intertemporal elasticity of labor supply in contractual markets.

[22] See Wooldridge (2002), pg. 567 for the estimation procedure used in this column.

of workers with particularly steep wage profiles during expansions. It could be that employment in jobs with flatter wage profiles is more cyclical in general, or that employers prefer to hire relatively inexperienced workers during expansions rather than recessions. The estimated history-dependence is more sensitive to selection compared to Table 4. This is because, the coefficient of the minimum unemployment rate is identified by workers for whom selection by anticipated wage growth is less likely to be important.[23]

When we add job fixed effects, the coefficient on lagged wage further declines to -0.50. Since the gap between two interviews varies in the data, this figure is not per annum. Given the distribution of time lags in our sample between two consecutive interviews, the annual change corresponds roughly to 80% of the estimated coefficient. A worker who was paid a one percent higher premium because of the history of economic conditions on the job, therefore, enjoys about 40% lower wage growth on average per year. This implies that 90% of the time-of-entry effects on starting wages fade out within 4 years, faster than our earlier estimate of 7 years when we ignored persistent variations in wage growth by worker and by job. This is consistent with the hypothesis that workers that are hired during expansions not only have better jobs, but also better prospects for wage growth.

These findings confirm our earlier conclusion that the time-of-entry effects in wages fade out relatively quickly. In our view, the theory of self-enforcing insurance contracts provides a natural interpretation of this result. Differences in the economic conditions when a contract is signed lead to job-cohort effects in wages. As the economic conditions fluctuate, wages are updated to reflect workers' and firms' outside options so as to prevent separation. Since the outside options are forward looking, the initial differences disappear.

5 Job Training and Human Capital Models

An important component of a worker's wage is his human capital. Perhaps, a systematic variation in human capital accumulation over the business cycle could explain our findings. Consider a generic on-the-job training model, à la Ben-Porath (1967) for instance, where skills are general and training activity takes time away from work. Since wages are procyclical, it is rational to invest in human capital during recessions, and work during booms. But then workers who are hired during booms, and those who experience favorable market conditions on-the-job would have accumulated less human capital, leading to lower wages. In addition, if there are diminishing returns to human capital investment,

[23] We also estimated our model using the initial and the minimum unemployment rate as instruments for lagged wage. The estimated coefficient indeed decreased by only 9 percentage points from -0.18 to -0.27.

these workers would also have steeper wage profiles relative to those hired in recessions. Both of these predictions are in contrast with our findings and the implications of the implicit contracts model.

Nonetheless, one could argue for a model with procyclical job training. If the employer bears the effective costs of training, if, for instance, the training activity is firm-specific, then potential liquidity problems during recessions may lead to lower training activity. In a related paper, Gibbons and Waldman (2006) develop a model of task specific human capital with learning to explain the presence of persistent cohort effects in wages. In this model, workers who enter the market in bad times are assigned to low-level tasks, which not only are associated with low entry wages, but also slower human capital accumulation, and hence lower wage growth. Therefore, although the model explains the job entry cohort effects in wage levels, it is at odds with our empirical findings on wage growth.

To empirically evaluate the implications of training and human capital for our findings, we directly control for training activity using the available measures in the NLSY. Although the training measures are imperfect, as probably most informal training activity goes unrecorded, we think that the available measures could give us an idea about the plausibility of a human capital explanation of our results.

The NLSY questions workers on the amount of time spent on training activities since the last time the worker was interviewed. Based on the responses, we constructed two variables: total hours of training activity between two wage observations, and the total cumulative amount of training since the worker first entered the labor market.[24].

The first column in Table 6 shows the BD regression with our training variables included. Overall, controlling for training does not effect the findings in the literature. The initial and the minimum unemployment rate remain significantly negative, while the current unemployment rate is not significant. The coefficient on the minimum unemployment rate declines from -2.85 to -2.14. The cumulative training variable has a positive and significant coefficient. We estimate the rate of return to a year of training (2000 hours) to be around 6.0%. This is somewhat lower than the return to a year of education.

The next two columns re-estimate our benchmark specifications, (11) and (12), including training variables among the covariates. Both estimations yield very similar coefficients as before. In fact, the coefficient on the minimum unemployment rate increases slightly from 1.37% in Table 3 to 1.71%, while the coefficients on U_0 and ΔU_t remain virtually the same. When we regress wage growth on the change in the minimum unemployment rate, the coefficient declines slightly to -2.82% with a standard error of 0.83,

[24]The appendix provides a detailed description of the variables.

compared to our benchmark estimate of -3.39%.[25]

Based on these results, we conclude that our empirical findings are not likely to be driven by cyclical fluctuations in human capital or training activity.

6 Conclusion

We study the wage growth of job stayers and show that wage adjustments over the business cycle show significant dependence on past economic conditions. In addition, changes in contemporaneous conditions do not have a significant effect on wage growth when past labor market conditions are controlled for. This is at odds with the spot market model of the labor market where wages equal marginal product at all times, and, hence, wage growth depends only on contemporaneous economic conditions.

We find that workers who were hired in booms, and those who experienced favorable economic conditions during their tenure on the job, have lower wage raises during expansions, and larger wage cuts during recessions. This pattern of wage adjustments is consistent with a contractual labor market, where employers and workers partake in an implicit agreement to shield wage payments from fluctuations in a worker's marginal product, without fully committing themselves to future payments and work.

Our results, therefore, indicate a decoupling of the marginal product from wage payments providing a potential explanation for the low elasticity of wages over the business cycle. If workers are paid below their marginal product during booms, and are overpaid during recessions, then the cyclicality of wages will be much lower than the underlying fluctuations in their marginal products.

Our results also draw attention to significant cohort entry effects in wages. Workers who are hired during recessions enjoy a lower wage rate in general. Nevertheless, our estimates indicate that these effects are relatively short-lived and they disappear in a few years.

References

Azariadis, C. 1975. "Implicit Contracts and Underemployment." *Journal of Political Economy* 83 (6):1183 – 1202.

[25] We also tried more general specifications that include polynomials of total cumulative training, lags of total training etc. The results are very similar to those reported in Table 6.

Baily, M. N. 1974. "Wages and Employment under Uncertain Demand." *Review of Economic Studies* 41 (1):37–50.

Baker, G., M. Gibbs, and B. Holmstrom. 1994. "The Wage Policy of a Firm." *Quarterly Journal of Economics* 109 (4):921–955.

Beaudry, P. and J. DiNardo. 1991. "Implicit Contracts and the Movements of Wages over the Business Cycle." *Journal of Political Economy* 99:665–688.

———. 1995. "Is the Behavior of Hours Worked Consistent with Implicit Contract Theory?" *Quarterly Journal of Economics* 110 (3):743–768.

Becker, Gary. 1964. *Human Capital*. New York: National Bureau of Economic Research.

Bellou, A. and B. Kaymak. 2010. "Wages, Implicit Contracts, and the Business Cyle: Evidence from a European Panel." Mimeo. University of Montreal.

Ben-Porath, Yoram. 1967. "The Production of Human Capital and Life Cycle of Earnings." *Journal of Political Economy* 75 (4).

Bils, Mark. 1985. "Real Wages over the Business Cycle: Evidence from Panel Data." *Journal of Political Economy* 93 (4):666 – 689.

Boldrin, M. 1995. "Labor Contracts and Business Cycles." *Journal of Political Economy* 103 (5):972–1004.

Bowlus, A. J. 1995. "Matching Workers and Jobs: Cyclical Fluctuations in Match Quality." *Journal of Labor Economics* 13 (2).

Freeman, R. 1981. "Career Patterns of Colleger Graduates in a Declining Job Market." *NBER Working Paper* 750.

Gibbons, R. and M. Waldman. 2006. "Enriching a Theory of Wage and Promotion Dynamics inside Firms." *Journal of Labor Economics* 24 (1):59 – 107.

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

Guiso, L., L. Pistaferri, and F. Schivardi. 2005. "Insruance within the Firm." *Journal of Political Economy* 113 (5):1054 – 1087.

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

Hagedorn, M. and I. Manovskii. 2010. "Job Selection and Wages over the Business Cycle." Mimeo. University of Pennsylvania.

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

Hall, R. E. 2005. "Employment Fluctuations with Equilibrium Wage Stickiness." *American Economic Review* 95 (1):50–65.

Harris, M. and B. Holmstrom. 1982. "A Theory of Wage Dynamics." *Review of Economic Studies* 49 (3):315 – 333.

Hashimoto, M. 1981. "Firm-Specific Human Capital as a Shared Investment." *Journal of Political Economy* 71 (3):475–482.

Kahn, L. 2010. "The Long-Term Labor Market Consequences of Graduating From College in a Bad Economy." *Labour Economics* 17 (2):303–316.

Kambourov, G. and I. Manovskii. 2008. "Rising Occupational and Industry Mobility in the United States." *International Economic Review* 49 (1):41–79.

Kocherlakota, N. 1996. "Implications of Efficient Risk Sharing without Commitment." *Review of Economic Studies* 63:595 – 609.

Kudlyak, M. 2010. "The Cyclical Price of Labor When Wages Are Smoothed." Federal Reserve Bank of Richmond Working Paper 10-13.

Kydland, F. E. and E. Prescott. 1982. "Time to Build Aggregate Fluctuations." *Econometrica* 50 (6):1345–1370.

Lazear, Edward P. 1979. "Why is There Mandatory Retirement?" *Journal of Political Economy* 87 (6):1261–84.

Long, J. B. and C. I. Plosser. 1983. "Real Business Cycles." *Journal of Political Economy* 91 (1):39 – 69.

MaCurdy, T. E. 1981. "An Empirical Model of Labor Supply in a Life-Cycle Setting." *Journal of Political Economy* 89 (6):1059 – 1085.

McDonald, J. T. and C. Worswick. 1999. "Wages, Implicit Contracts and the Business Cycle: Evidence from Canadian Macro Data." *Journal of Political Economy* 107 (4):884 – 892.

Mortensen, Dale and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61:397–415.

Okun, A. M. 1973. "Upward Mobility in a High-Pressure Economy." *Brooking Papers on Economic Activity* 1973 (1):207 – 261.

Oreopoulos, P., T. Von Wachter, and A. Heisz. 2006. "The Short- and Long-Term Career Effects of Graduating in a Recession: Hysteresis and Heterogeneity in the Market for College Graduates." NBER Working Paper No: W12159.

Oyer, P. 2006. "Initial Labor Market Conditions and Long-Term Outcomes for Economists." *Journal of Economic Perspectives* 20 (3):143–160.

Rosen, S. 1985. "Implicit Contracts: A Survey." *Journal of Economic Literature* 23 (3):1144 – 1175.

Rudanko, L. 2009. "Labor Market Dynamics under Long-Term Wage Contracting." *Journal of Monetary Economics* 56 (2):170–183.

———. 2010. "Aggregate and Idiosyncratic Risk in a Frictional Labor Market." Forthcoming in American Economic Review.

Thomas, J. and T. Worrall. 1988. "Self-enforcing Wage Contracts." *Review of Economic Studies* 55 (4):541 – 553.

Topel, R. 1991. "Specific Captal, Mobility, and Wages: Wages Rise with Job Seniority." *Journal of Political Economy* 99 (1):145 – 176.

Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross-Section and Panel Data*. MIT Press, Cambridge.

A Data

The analysis focuses on male respondents in the cross-sectional sample, who at the time of the interview were not enrolled in school and were employed.

Wages: The wage is the hourly rate of pay constructed by the NLSY. Nominal wages are deflated using the annual CPI index (All Urban Consumers, U.S City Average, All Items) from the Bureau of Labor Statistics (base period 1982-84). Wages were deflated using the CPI of the year when the worker last worked for the job as reported at the time of the interview. Observations with missing wage information or real wages below $1 and above $100 are dropped.

Hours: These are the usual weekly hours worked. Observations with missing information on hours were dropped. The sample includes only full-time workers (usual weekly hours of 35 or more).

Class of the job: The sample includes workers in the private sector only, thus dropping government employees, self-employed and those working without pay.

Industry Classification: The NLSY has employed the 3-digit 1970 and 1980 Census classification system in the 1979-2000 surveys in order to code all jobs into industry groups. Beginning 2002, the 3-digit 2000 Census codes were used to classify industries of all jobs reported by the respondents. To minimize potential inconsistencies or the effect of coding changes due to switching from the 1970/1980 to 2000 classification system for respondents who did not change jobs between consecutive interviews, 9 broader industry groups are defined based on the reported industry classification. The groups are: Agriculture, Forestry and Fisheries; Mining; Construction; Manufacturing; Utilities, Transportation and Warehousing; Wholesale and Retail Trade; Finance, Insurance, Real Estate, Rental and Leasing; Professional, Scientific, Technical Services, Management, Administrative and Waste Management Services, Educational Services, Health Services, Accommodation and Food Services, Arts, Entertainment and Recreation, Other Services; Public Administration

Job start date: The starting date of the job is identified by subtracting tenure (constructed by the NLSY and measured in weeks) from the date the worker last worked for the job as reported at the interview date. Jobs that started prior to 1976 are disregarded.

Current age: The current age corresponding to each job observation is constructed as the difference between the year the worker last worked at the job as reported at the time of the interview and the birth year. The age at the start of the job is calculated as the difference between the start year of the job and the birth year of the respondent. We only consider jobs that started when the respondent was 16 or older. Moreover, we restrict attention to workers with current age 21 years old and above.

Experience: This is actual experience measure in weeks constructed by adding for consecutive interviews the "total number of weeks the respondent worked since the last interview". This variable is constructed by the NLSY for all respondents of ages 16 years old

and above. The results are very similar to the usage of current age at each job observation as a measure of potential experience.

Unemployment rate: The unemployment rate is the quarterly, seasonally adjusted, civilian unemployment rate for ages 16+ obtained from the Bureau of Labor Statistics. The contemporaneous unemployment rate is the unemployment rate at the date (quarter, calendar year) when the respondent reported last working for the job. The initial unemployment rate corresponds to the unemployment rate at the date (quarter, calendar year) the job started. The minimum quarterly unemployment rate in the wage growth specifications is calculated as the historical minimum unemployment rate recorded between the date (quarter, calendar year) the job started and the last interview date (quarter, calendar year) before the contemporaneous year. All specifications are robust to the usage of annual instead of quarterly unemployment.

Training Variables: At every survey respondents were asked if they had participated in any training programs since the previous interview. Detailed information, then, were collected on the duration, intensity and the type of the training spells. The training data used in our estimations cover 1979 to 2004.The earlier surveys, 1979 to 1986, do not provide these details for training spells that lasted less than a month. For longer spells, the respondents reported the beginning and ending dates of each training spell (in month and year) and the average number of hours a week spent for training. This enables a construction of the total time investment in training in hours since the last interview. If the respondent was currently enrolled in a training program, an additional dummy variable was created. Until 1988, up to three training spells were recorded. Later this limit was raised to four. The respondents were however asked if they had fourth (fifth after 1986) training program to report. Based on this question, it is possible to calculate the number of workers for which this limit was binding. The limit was binding for a total of only 80 observations (about 0.2% of the sample) in all years.

Figure 1: Wages over a boom cycle in a contractual market without commitment

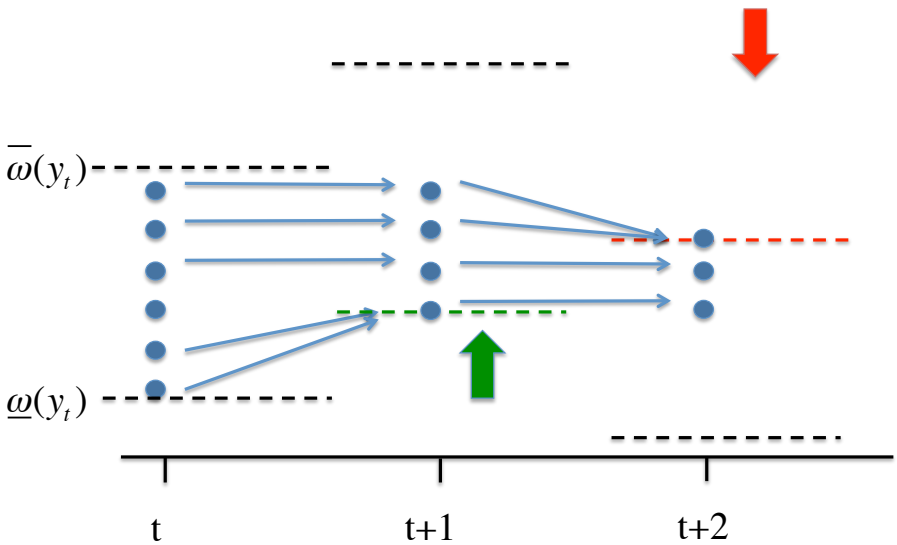

Figure 2: Endogenous survival and wage growth on-the-job

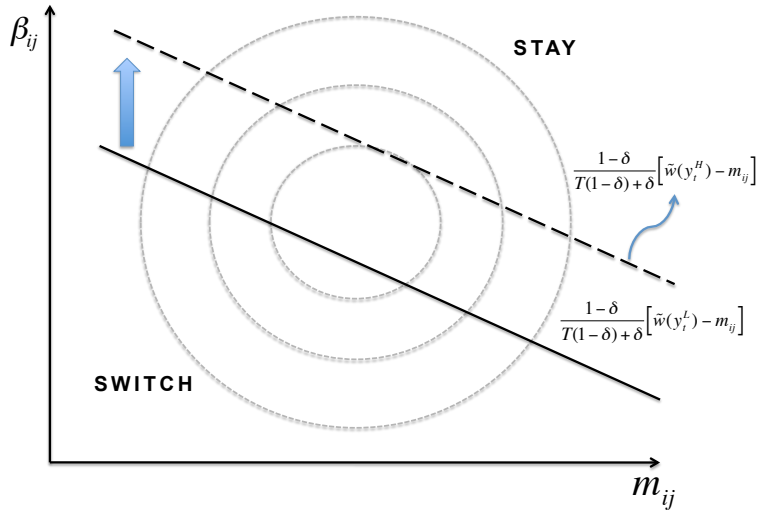

Figure 3: The unemployment rate: 1979 - 2008

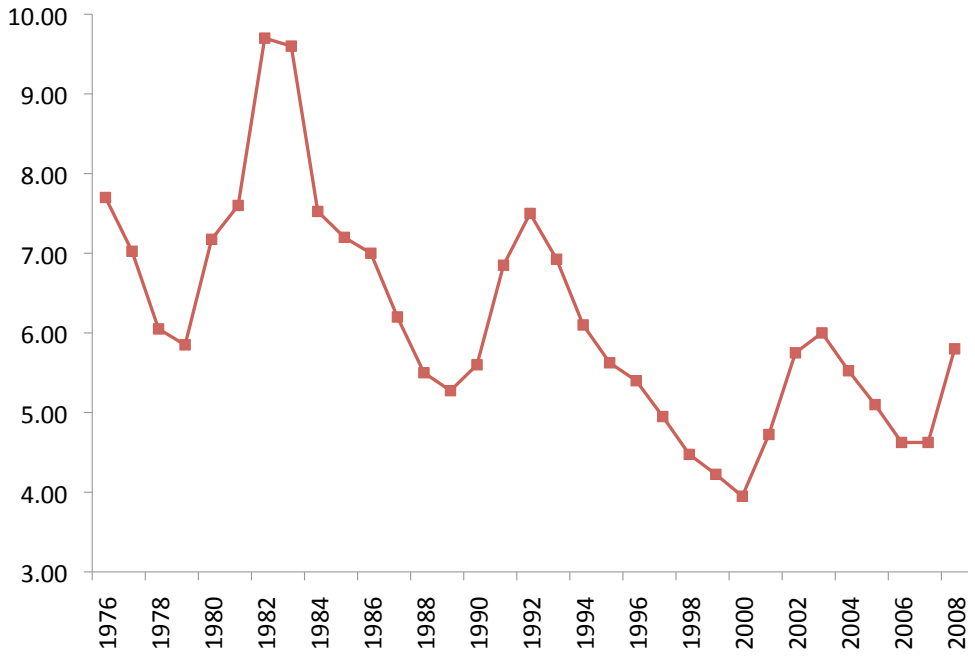

Table 1: Descriptive Statistics

	Mean	Standard Deviation
Age	31.6	7.4
Tenure	5.1	5.3
Income (log)	2.1	0.5
Weekly hours	45.6	8.6
Number of workers	2,437	
Number of observations	22,329	

Note.— Data comes from the 1979 cohort of the NLSY (1979 - 2008). Men of ages 21 and older who work full time in the private sector.

Figure 4: Time of entry effects in wages

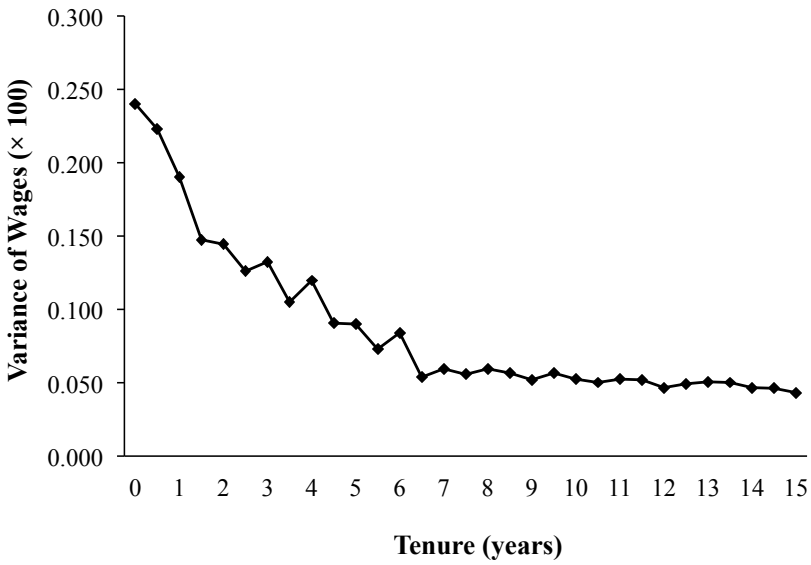

Table 2: Real Wages and Unemployment History

Dependent Variable: hourly wage rate (logs)				
U_t	-1.35**	-1.06**	-0.09	-0.25
	(0.26)	*(0.23)*	*(0.24)*	*(0.25)*
U_{t_0}		-2.03**		-0.80**
		(0.26)		*(0.33)*
$U^{min}_{t_0,t}$			-3.65**	-2.85**
			(0.42)	*(0.51)*
Sample size	22,329	22,329	22,329	22,329

Note .— All specifications control for individual fixed effects, cubic polynomials in experience and tenure, a quadratic time trend, and indicators for industry and region. Data comes from the 1979 cohort of the NLSY (1979 - 2008). Sample includes Men of ages 21 and older who work full time in the private sector. Coefficients and standard errors are multiplied by 100. Standard errors are clustered by start year and current year interactions.
$*$, $**$indicate statistical significance at 5%, 1%.

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

Dependent Variable: $\Delta_k \ln w_{ijt}$							
	I	II	III	IV	V	VI	VII
$\Delta_k U_t$	-0.26	-0.14	0.26	0.25	0.27	0.24	0.30
	(0.25)	*(0.25)*	*(0.28)*	*(0.28)*	*(0.29)*	*(0.28)*	*(0.29)*
$\Delta_k U^{min}$					-3.39**		-3.46**
					(0.84)		*(0.83)*
U_{t_0}		0.51**		0.09		0.09	
		(0.18)		*(0.20)*		*(0.2)*	
$U^{min}_{t_0,t-k}$			1.61**	1.51**		1.57**	
			(0.37)	*(0.43)*		*(0.43)*	
Inverse Mills Ratio						-1.24*	-0.91
						(0.62)	*(0.63)*
		First Step Probit Estimation					
Job Duration						0.12**	0.12**
						(0.004)	*(0.004)*
Working Spouse						1.43**	1.10*
						(0.45)	*(0.5)*
# future jobs/year						-75.27**	-82.65**
						(9.02)	*(10.87)*
Sample size	15,616	15,616	15,616	15,616	15,616	15,616	15,616

Note.– 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 VI. and VII. are adjusted for non-random sample selection using the Heckman correction procedure. Data comes from the 1979 cohort of the NLSY (1979 - 2008). Sample includes men of ages 21 and older who work full time in the private sector. Coefficients and standard errors are multiplied by 100. Standard errors are clustered by start year and current year interactions.
$*$, $**$indicate statistical significance at 5%, 1%.

Table 4: Real Wage Growth and Unemployment History with Profile Heterogeneity

Dependent Variable: $\Delta_k \ln w_{ijt}$							
$\Delta_k U_t$	-0.08	0.06	0.41	0.40	0.41	0.46	0.39
	(0.25)	*(0.25)*	*(0.28)*	*(0.28)*	*(0.30)*	*(0.29)*	*(0.30)*
$\Delta_k U^{min}$					-3.37**		-2.96**
					(0.87)		*(0.98)*
U_{t_0}		0.96**		0.28			
		(0.31)		*(0.37)*			
$U^{min}_{t_0,t-k}$			1.77**	1.58**		2.01**	
			(0.42)	*(0.51)*		*(0.57)*	
Worker Fixed Effects	Yes	Yes	Yes	Yes	Yes	–	–
Job Fixed Effects	No	No	No	No	No	Yes	Yes
Obs.	15,616	15,616	15,616	15,616	15,616	15,616	15,616

Note.— All specifications also control differences in cubic polynomials of experience and tenure, indicators for region and industry. Data comes from the 1979 cohort of the NLSY (1979 - 2008). Sample includes men of ages 21 and older who work full time in the private sector. Coefficients and standard errors are multiplied by 100. Standard errors are clustered by start year and current year interactions.
$*$, $**$ indicate statistical significance at 5%, 1% .

Table 5: Real Wage Growth and No-Commitment Contracts

Dependent Variable: $\Delta_k \ln w_{ijt}$				
	I	II	III	IV
$\log W^p_{t-k}$	-0.11**	-0.15**	-0.31**	-0.50**
	(0.02)	*(0.03)*	*(0.03)*	*(0.04)*
Worker Fixed Effects	No	No	Yes	No
Job Fixed Effects	No	No	No	Yes
Sample Size	15,616	15,616	15,616	15,616

Note.— TSLS estimates where the last observed wage, $\log W_{t-k}$, is instrumented by a full interaction of start year and last year indicators, $I(t_0 \times t-k)$. Estimate in column II is adjusted for non-random selection using the Heckman correction. All specifications also control differences in cubic polynomials of tenure and age, indicators for region and industry. Data comes from the 1979 cohort of the NLSY (1979 - 2008). Sample includes men of ages 21 and older who work full time in the private sector. Standard errors are clustered by start year and current year interactions.
$*$, $**$ indicate statistical significance at 5%, 1% .

Table 6: Training, Unemployment and Wage Growth

Dependent Var.	$\ln w$	$\Delta_k \ln w_{ijt}$	$\Delta_k \ln w_{ijt}$
U_t	-0.43		
	(0.24)		
U_{t_0}	-1.02**	0.04	
	(0.33)	*(0.20)*	
$U_{t_0,t}^{min}$	-2.14**	1.71**	
	(0.50)	*(0.47)*	
$\Delta_k U_t$		0.19	0.08
		(0.29)	*(0.30)*
$\Delta_k U^{min}$			-2.82**
			(0.83)
$\sum Tr_t/2000$	6.00**	-0.50	-0.47
	(0.95)	*(0.52)*	*(0.52)*
$Tr_t/2000$		3.14	2.94
		(2.50)	*(2.50)*
Sample size	20,464	14,171	14,171

Note.– All specifications control for differences in cubic polynomials of experience and tenure, differences in a quadratic time trend, and indicators for industry and region. Tr_t denotes the training activity between two consecutive wage observations, and $\sum Tr_t(/2000)$ denotes the total cumulative training over a worker's career. Coefficients are multiplied by 2000 to reflect a year of training, and then by 100 to report percentages. Data comes from the 1979 cohort of the NLSY (1979 - 2004). Sample includes men of ages 21 and older who work full time in the private sector. Coefficients and standard errors are multiplied by 100. Standard errors are clustered by start year and current year interactions.
$*$, $**$ indicate statistical significance at 5%, 1%.