Employment Duration over the Business Cycles: Quits vs. Firings*

Ismail Baydur†
University of Virginia

April 21, 2014

Abstract

This paper studies the cyclical behavior of employment duration. Models with on-the-job search predict that jobs created during a recession have shorter spells, because workers are more likely to accept low-quality jobs during a recession. In contrast, models with endogenous separation predict that jobs created during a recession endure longer, because firms contact a larger group of applicants and are able to hire high-quality workers. I test these competing predictions using data from the National Longitudinal Survey of Youth 1979 cohort. I estimate a proportional hazard model under the assumption that job terminations due to different reasons are competing risks. My results support the predictions of both models with on-the-job search and models with endogenous separation. A higher unemployment rate at the start of an employment relationship increases the probability that the worker quits to take or look for another job, but it decreases the probability the firm fires the worker. The net effect of these opposing forces on the overall duration of the employment is negative, but small, implying that match quality is weakly pro-cyclical. Furthermore, an increase in the current unemployment rate reduces the probability that the job spell end by the worker's quit decision, but it increases the probability that the firm fires the worker. These findings are consistent with pro-cyclical quits and counter-cyclical firings.

Keywords: Business cycles, employment, quits, firing, match quality, job duration
JEL Classifications: E24, E32, J22, J63, J64.

---

*I would like to thank Toshihiko Mukoyama, Sophie Osotimehin, and Huzeyfe Torun for their valuable comments. All errors are my own.

†Contact information: ib2s@virginia.edu

1 Introduction

The duration of an employment relationship is a signal about its match quality and can be used to explore how business cycles affect match quality. High quality matches are likely to endure longer, while low quality matches dissolve relatively quickly. Therefore, match quality is pro-cyclical if jobs created during recessions have shorter spells and counter-cyclical otherwise.

However, the theoretical predictions about the direction of the cyclical behavior of match quality is unclear, because job seekers and hiring firms respond to labor market conditions in opposite directions. On one hand, the unemployment rate is high during a recession, and job seekers compete for a relatively small number of job openings. Models with on-the-job-search imply that increased competition among the job seekers causes unemployed workers to accept low-quality jobs and quit later to take or look for another job.[1] Hence, jobs ending due to workers' quit decisions tend to be shorter. On the other hand, labor market conditions are favorable for hiring firms during a recession, because there are a relatively small number of hiring firms. Models with endogenous separation imply that favorable labor market conditions during a recession are likely to cause firms to hire high quality workers that remain at the hiring firm longer.[2] As a result, the aggregate labor market conditions at the start of a job have an ambiguous effect on employment duration.

In this paper, I empirically study the effects of aggregate labor market conditions on employment duration using data from the National Longitudinal Survey of Youth (NLYS) 1979 cohort. I use a proportional hazard model for job terminations, where I treat job terminations due to different reasons as competing risks. I am particularly interested in estimating the effects of unemployment rate separately for job terminations ended by a worker's quit decision to take or look for another job and those ended by the firm's firing decision.

I find that a high unemployment rate at the start of a job increases the probability that a worker quits his current job to take or look for another job, but it reduces the probability that a job spell ends by firm's firing decision. The overall effect of the unemployment rate at the start of a job on employment duration is negative, implying match quality is pro-cyclical. However, this effect is rather small. The median duration of a non-union job held by a 29 year-old white male with a high school degree falls from 44 weeks to only 42 weeks if the unemployment rate at the start of the job is one standard deviation above its sample mean.

Furthermore, aggregate data show that quits are pro-cyclical but firings are counter-

[1] See Mortensen (1994) for an application to U.S. data.

[2] See Mortensen and Pissarides (1994).

cyclical.[3] These regularities in the aggregate data suggest that workers and firms with ongoing employment relationships also respond to current labor market conditions in opposite directions. Using the estimates from the proportional hazard model, I find that an increase in the current unemployment rate reduces the probability that an already employed worker quits to take or look for another job, but it increases the probability that a job spell ends by firm's firing decision. These results are consistent with the aggregate behavior of quits and firings.

The main contribution of this paper is to utilize information about the reason for job terminations. The NLSY 1979 cohort includes detailed information about the reason why a job spell has ended. In particular, I observe job terminations due to a worker's quit decision specifically to take or look for another job and due to the firm's firing and lay-off decision. This feature of the data enables me to distinguish worker-initiated job terminations from firm-initiated job terminations.

Previous studies used the data from the NLSY 1979 cohort to estimate the cyclical behavior of job match quality, but these studies do not make a distinction among the cause-specific job terminations. For example, Bowlus (1995) and Mustre-del-Rio (2012) also find that match quality is pro-cyclical, but their estimates suggest a much stronger pro-cyclical behavior of match quality.

The inclusion of the causes for job terminations distinguishes my paper from the earlier studies for four reasons. First, it allows me to estimate the effects of labor market conditions separately on worker-initiated and firm-initiated job terminations. In models with on-the-job search, workers accept low-quality jobs during recessions, but they later quit to take or look for another job. The mechanism in the models with on-the-job search imply that job match quality is pro-cyclical and the cause of a job termination is more likely to be due to a worker's quit decision. In contrast, models with endogenous separation imply that only high quality matches are created during recessions. Therefore, match quality is counter-cyclical and the cause of job termination is more likely to be due to the firm's firing decision. Earlier studies do not make a distinction between quits and firings and try to estimate the *net* effect of these opposing forces described in these theoretical models. Instead, I estimate the effects of unemployment rate on quits and firings separately and find supporting evidence for *both* of these class of models. However, these opposing forces are of similar magnitude and cancel each other's effect.[4]

Second, earlier studies discarded job spells for females, because they are likely to quit

[3]See Akerlof, Rose, and Yellen (1989) for a discussion. For an analysis of more recent data, see the Job Openings and Labor Turnover Survey Highlights (January 2014) report from The Bureau of Labor Statistics.

[4]Kahn (2008) finds a similar result using a matched employer-employee data set and after controlling for firm fixed effects.

their jobs due to personal concerns, e.g. pregnancy and child care, rather than professional concerns. Since I observe the specific reason for job terminations, I do not need to impose such a restriction to my sample.

Third, earlier studies also acknowledge that the current unemployment rate affects the duration of a job spell and controls for that to isolate the effect of the unemployment rate at the start of the job. However, as supported by the aggregate data, the effect of the unemployment rate on job termination decisions for workers and firms are in opposite directions. Without making a distinction among the job terminations according to their causes, the inclusion of the current unemployment rate on the right hand side causes the model to be misspecified.[5] Since I treat cause-specific job terminations as competing risks, my results do not suffer from this misspecification bias.

Finally, the inclusion of information about the reasons for job terminations necessitates a change in my estimation strategy. While the termination of a job is still the failure event in the proportional hazard model, it can occur due to several mutually exclusive reasons. For example, if a job spell ends by a worker's quit decision in the data, then a termination for the same job spell due to the firm's firing decision is never observed. In other words, each cause of job terminations is a competing risk for the other causes.

I estimate the hazard functions for job terminations under two different specifications. In the first specification, I estimate a proportional hazard model for each cause-specific event as in Cox (1972). This specification is similar to the standard application of Cox's proportional hazard model except that job terminations due to other reasons are treated as right-censored and handled as in the standard Cox proportional hazard model. Despite the similarities in implementation, the interpretation of the coefficient estimates are completely different when there are competing risks. While the proportionality assumption in Cox hazard model allows for inference solely based on the coefficient estimates, this is generally not possible when there are competing risks for the same failure event. For example, the probability of a job termination due to a worker's quit decision before a specific point in time depends on the survival probability of the job up to that point. However, the survival probability depends on all the cause-specific hazard functions, which makes the cause-specific job terminations interdependent. To facilitate inference from the cause-specific regressions, I calculate the probability of observing each cause-specific event, called the cumulative incidence function, using the estimates from all the cause-specific hazard estimations.

The second specification is based on an alternative method for cause-specific hazard

[5] To address this problem, Bowlus (1995) adds the square of the current unemployment rate as an explanatory variable. The coefficient estimate for the current unemployment in Mustre-del-Rio (2012) is insignificant.

regressions proposed by Fine and Grey (1999). They directly estimate a subhazard function, which counts job terminations due to other reasons in the risk set of the failure event rather than treating them as right-censored. An advantage of subhazard regressions is that, as in the standard Cox regression, inference can be made solely based on the coefficient estimates. Furthermore, potential bias in the coefficient estimates due to unobserved heterogeneity is less of a concern. I include only one randomly selected job spell for each individual in my sample. This sampling scheme produces unbiased coefficient estimates, although the baseline hazard is biased downward due to overrepresentation of longer job spells. While the potential bias due to unobserved heterogeneity is not an issue for inference from the subhazard regressions, it may potentially bias the estimates of the cumulative incidence functions from the cause-specific regressions, which uses the baseline hazard estimates in addition to coefficient estimates.

I provide estimation results from the subhazard regressions to address issues concerning unobserved heterogeneity in the cause-specific hazard regressions. The estimation results from both of the specifications are consistent with each other. However, cumulative incidence functions for the subhazard regressions are calculated from separately-run regressions and can potentially produce inconsistent results. For example, the sum of the probabilities of all possible cause-specific job terminations can potentially exceed one. An advantage of the cause-specific regressions over the subhazard regressions is that cumulative incidence functions are estimated simultaneously and robust to such inconsistencies.

In the next section, I describe the data set I use in this study. Section 3 describes the cause-specific and subhazard regressions. I present my estimation results from both of the specifications in Section 4. The last section concludes.

2 Data Description

I use data from the NLSY 1979 cohort in this study. A total of 12,686 individuals that were born between 1957 and 1964 participated in this survey. These individuals were interviewed annually from 1979 through 1994 and biennially thereafter until the survey ended in 2010. The data set I use in this study covers all the survey years.

The survey collects detailed information about each job a respondent holds or previously held. The structure of the survey enables me to create employment histories for all the individuals participating in the survey. I construct the data set for the employment duration analysis by linking each job across different survey years.[6] I measure the duration of a job

[6]I obtain some of the job-specific characteristics, e.g. job start and stop dates, from the Employer History Roster (Beta Version). This roster alleviates the more involved linking process across different survey years.

spell in weeks. Some of the job spells are right-censored due to the finite horizon of the survey and loss of follow-up.

The explanatory variables include personal and job characteristics at the start of a job such as age, gender, race, education, and whether the job is protected by a union. I include unemployment rate at the start of the job, u_0, to account for the aggregate labor market conditions when the job is created. I also include the current unemployment rate, u_t, as a time-varying regressor to capture the on-going labor market conditions. I obtain data from the Bureau of Labor Statistics for the national unemployment rate. The time series is not seasonally adjusted so that it is consistent with the data from NLYS.

The NLSY 1979 also provides detailed information about the reason why a job spell ended. A detailed description of the reasons for job terminations is available in the Appendix. In particular, I observe whether the job ended due to the worker's quit decision to take or look for another job or due to the firm's firing decision. The theory predicts that workers and firms response to aggregate labor market conditions are different and the overall effect on the duration of a job spell is ambiguous. On one hand, workers are likely to accept low-quality jobs during recessions due to tight labor market conditions and quit later during booms to take or look for a better job. Workers' incentive to quit for a better job tends to reduce the duration of the job. On the other hand, firms hire among a larger applicant pool during recessions and can potentially form high-quality matches that endure longer. Hence, the duration of a job spell can be different for the jobs ending by the worker's quit decision and those ending by the firm's firing decision.

Using information about the reason why a job spell ended, I test these theoretical predictions in this paper. I expect both the starting and the current unemployment rate to affect quits and firings in opposite directions. While workers accept a low-quality job and are more likely to quit when the starting unemployment rate is high, they are less likely to quit if the current unemployment rate is high. In contrast, firms tend to form better matches when the starting unemployment rate is high, but they are more likely to fire a worker when the current unemployment rate is high.

3 Estimation Strategy

The Cox proportional hazard model is widely applied to duration data when time to a failure event is of interest. In the analysis of employment duration, the failure event of interest is the termination of a job. In this paper, the failure event of interest is still the termination of a job. However, there are multiple causes of job terminations, and only the first of these causes for job termination, if any, is observed. In other words, each reason for job termination is a competing risk for the other reasons. In this section, I describe two alternative approaches proposed in the literature when there are competing risks: cause-specific hazard regressions and regression on a subhazard function.[7]

3.1 Cause-Specific Hazard Functions

Let the hazard function for job terminations be:

$$h(t) = \lim_{\Delta t \to 0} \frac{P(t \leq T < t + \Delta t | T \geq t)}{\Delta t} \tag{1}$$

The hazard function is the instantaneous probability that a job is terminated at time T conditional on surviving up to time t. Cox (1972) further imposes that the hazard function for job terminations, conditional on a set of explanatory variables at time t, $X(t)$, takes the following proportional form:

$$h(t|X(t)) = h_0(t)\exp(X(t)\beta) \tag{2}$$

where $X(t)$ are time-varying explanatory variables, β is a vector of parameters common across all job spells, and $h_0(t)$ is the baseline hazard. The baseline hazard is also common across all job spells, and its form is left unspecified. Cox (1972) describes a semi-parametric approach for obtaining estimates of the model parameters, $\hat{\beta}$, through the maximization of the following partial likelihood function:

$$\mathscr{L} = \prod_{i:C_i=1} \frac{h(t_i|X_i(t_i))}{\sum_{j:t_j \geq t_i} h_(t_i|X_j(t_i))} = \prod_{i:C_i=1} \frac{\exp(X_i(t_i))}{\sum_{j:t_j \geq t_i} \exp(X_j(t_i))} \tag{3}$$

where $C_i = 0$ if the job spell is right-censored. Note that right-censored job spells enter the partial likelihood function only through the denominator. Further, the baseline hazard can be recovered non-parametrically after obtaining $\hat{\beta}$ even though it cancels from the estimating equation. The proportionality assumption implies that the hazard functions are strictly parallel and inference is possible solely based on $\hat{\beta}$. Specifically, a positive (negative) value of $\hat{\beta}$ implies that the probability of terminating a job increases (decreases) with an increase in the value of the explanatory variable.

A standard application of the Cox proportional hazard model can be misleading when there are competing events. The proportional hazard model in equation (2) assumes that the explanatory variables affect the probability of terminating a job in the same way regardless of its cause. However, the model is misspecified under such a restriction if the effect of one of

[7] Refer to Putter, Fiocco and Geskus (2006) for a general discussion.

the explanatory variables is different for each cause-specific job termination. In this study, a quit and a firing are competing events for job terminations, and the theory predicts that both the starting and current unemployment rate affect the probability of terminating a job due to quits or firings in opposite directions.

To empirically test for potentially different effects of the starting and current unemployment rates on the duration of a job, I define a separate hazard function for each cause-specific job termination. Formally, let k denote one of the K possible cause of job terminations. The hazard function for terminating a job due to reason k is:

$$h_k(t|X(t)) = h_{0,k}(t) \exp(X(t)\beta_k) \tag{4}$$

The specification in equation (4) is similar to the standard specification in equation (2) except that it is now separately defined for K different possible reasons for job terminations. Both the baseline hazard functions and the parameters are allowed to differ across different types of job terminations. β_k's can be estimated separately for each cause-specific hazard function by maximizing the partial likelihood function in equation (3). However, the occurrence of a competing event is treated as right-censored in each of these estimations.

3.2 Cumulative Incidence Functions and Inference

While the estimation procedure with cause-specific hazard functions is the same as with the standard Cox proportional hazard model, the interpretations of the parameter estimates are different. Because the distributions of time to a job termination for each cause-specific event are potentially dependent, the sign of the parameter estimates alone cannot determine the effect of a covariate on the duration of employment. When the hazard functions are estimated separately for each cause-specific job termination, the effect of a change in the variable of interest on a cause-specific job termination depends nonlinearly on baseline hazard functions and parameter estimates of the other cause-specific hazard functions.

To illustrate this point, let the baseline cumulative cause-specific hazard function be:

$$H_k(t) = \int_0^t h_k(s)ds \tag{5}$$

Then, the probability of surviving from any event at time t is:

$$S(t) = \exp(-\sum_{k=1}^{K} H_k(t)) \tag{6}$$

The survival probability now depends on the baseline and parameter estimates not only from the hazard regression of the event of interest, but also from the hazard regressions of the other competing events. Further, the probability of failing from cause k before time t is:

$$I_k(t) = \int_0^t h_k(s)S(s)ds \tag{7}$$

The probability in equation (7) is called the cumulative incidence function. The effect of a change in the starting and current unemployment rates can now be examined for quits and firings by constructing cumulative incidence functions for each event using the baseline and parameter estimates from the cause-specific hazard regressions.

3.3 Regression on a Subhazard Function

As an alternative to cause-specific hazard regressions, Fine and Gray (1999) propose a methodology that allows inference on cumulative incidence functions solely based on estimates of β. They define a subhazard function for the competing risk k as follows:

$$\bar{h}_k(t) = \lim_{\Delta t \to 0} \frac{P(t \leq T < t + \Delta t | T \geq t \cup (T \leq t \cap K \neq k))}{\Delta t} \tag{8}$$

The subhazard function shows the instantaneous probability of a job ending due to reason k conditional on surviving up to time t *or* ending before time t due to a reason other than k. As in Cox's proportional hazard model, Fine and Grey (1999) also assume that the subhazard function is proportional to a baseline hazard function:

$$\bar{h}_k(t|X(t)) = \bar{h}_{0,k} \exp(X(t)\beta_k) \tag{9}$$

The subdistribution hazard function in equation (9) can be estimated in a way that is analogous to equation (3). The only difference in the estimation procedure is in the treatment of the risk set. According to equation (8), job spells that have already ended due to another cause are still considered to be in the risk set for the competing risk k. Since these observations can potentially become right-censored and drop from the risk set, Fine and Grey (1999) weight them using the Kaplan-Meier estimate of the survivor function for the censoring distribution.

One of the advantages of the estimation strategy proposed by Fine and Grey (1999) is that inference can now be made solely based on $\hat{\beta}$. Note that the baseline cumulative incidence function and subhazard function for the competing risk k are related as follows:

$$\mathrm{CIF}_k = 1 - \exp\left(-\int_0^t \bar{h}_k(s)ds\right) \tag{10}$$

The estimates of β have a similar interpretation to the standard Cox proportional hazard model. A positive (negative) value of $\hat{\beta}$ implies that the effect of increasing the value of the explanatory variable increases (decreases) the probability of terminating a job due to cause k.

In the next section, I implement both of these estimation methods using the duration data from NLSY 1979 cohort.

4 Results

4.1 Sample Restrictions

Following Bowlus (1995), I restrict the data set to include only private sector employment. Jobs that start before the individual completes all schooling or is younger than 16 years old are dropped from the sample. Further, jobs with missing job start and stop dates and those lasting less than two weeks are not included in the sample. Unlike Bowlus (1995), I still include females in my sample. Bowlus (1995) restricts the sample to only males on the grounds that females are likely to quit for reasons other than poor match quality, such as marriage, pregnancy, and childcare. The information about the reason for job terminations allows me to distinguish job terminations due to professional concerns from personal concerns. Therefore, I do not need to make such a restriction on my sample.

I define three reasons for job terminations for the empirical analysis: quits, firings, and other reasons. Models with on-the-job search imply that job spells are shorter for those jobs created during a recession because workers quit to take or look for a better job. Accordingly, quits due to reasons other than to take or look for another job are included in the other reasons category. Similarly, models with endogenous separation imply longer spells for jobs created during a recession, because these matches are expected to be high quality, and firms are less likely to fire these workers. Thus, firings include discharges and layoffs. Termination of temporary and seasonal jobs are included in the other reasons category, because these jobs are set for a fixed term regardless of match quality. Terminations due to closings are also included in the other reasons category since all jobs regardless of match quality are terminated with this type of job termination.[8]

[8]Inclusion of these type of job terminations in the firings category does not change the conclusions of this

The original data set consists of one observation per job and multiple spells for each individual. If there is an individual-specific unobserved component, the job spells for the same individual are potentially correlated and the estimates of β are biased. To address the concerns about unobserved heterogeneity, I randomly select one spell per individual. Bowlus (1995) points out that such a restriction on the sample produces unbiased estimates of β. However, the estimates for the baseline hazard functions are still biased, because longer spells are now overrepresented in the sample. Since cumulative incidence functions are constructed from the estimates of the baseline hazard functions, they are potentially biased too. While the bias in cumulative incidence function calculations is problematic for inference from the cause-specific hazard regressions, this is not a concern for the subhazard regressions because inference is possible solely based on the estimates of β.

4.2 Estimation Results

Table 1 presents the estimation results. The first three columns show the estimation results from the cause-specific hazard regressions for job terminations due to quits, firings, and other reasons, respectively. The last three columns show the results from the subhazard resgressions. The coefficients of interest are those for the unemployment rate at the start of the job spell, u_0, and throughout the duration of the job, u_t, for quits and firings.

The effects of the explanatory variables can be directly inferred from the estimates of the subhazard regressions. The effect of u_0 is positive and statistically significant for the job spells ending by a worker's quit decision to take or look for another job. The positive sign implies that a high unemployment rate at the start of a job spell increases the probability that the worker is more likely to quit his current job. This result is consistent with the predictions of the models with on-the-job search. Regarding the effects of u_0 on job spells ending by firm's firing decision, the sign of the coefficient from the subhazard regressions supports the predictions of the models with endogenous separation. In contrast to the job terminations due to quits, the sign of the coefficient for u_0 is negative and statistically significant. The negative coefficient implies that a high employment rate at the start of a job reduces the probability that a firm will fire the worker in the future.

In the aggregate data, the cyclical behavior of aggregate quits and firings are qualitatively different. While quits are strongly pro-cyclical, firings are counter-cyclical. This macro-level observation suggests that workers and firms also respond to u_t in opposite directions. The negative coefficient for u_t from the subhazard regression for quits indicates that the probability that a job spell ends by a worker's quit decision is lower during recessions. In

paper.

	Cause-Specific Hazard			Subhazard		
Variable	Quits	Firings	Other	Quits	Firings	Other
u_0	.157*	-.052*	-.134*	.344*	-.101*	-.214*
	(0.017)	(.022)	(.018)	(.021)	(.030)	(.024)
u_t	-.104*	.138*	.090*	-.360*	.151*	.116*
	(.017)	(.020)	(.0166)	(.024)	(.029)	(.023)
HS	.086	-.269*	-.148*	.189*	-.227*	-.056
	(.063)	(.068)	(.054)	(.063)	(.067)	.055
COL	.063	-1.337*	-.627*	.379*	-1.063*	-.355*
	(.076)	(.129)	(.078)	(.076)	(.129)	(.078)
AGE	.661*	.071	.024	.541*	-.072*	-.101*
	(.046)	(.039)	(.024)	(.045)	(.033)	(.022)
SQAGE	-.012*	-.002*	-.001*	-.010*	.001*	.001*
	(.000)	(.001)	(.000)	(.000)	(.001)	(.000)
NWHITE	.082	.325*	.061	.026	.277*	-.000
	(.043)	(.055)	(.042)	(.043)	(.054)	(.042)
GEN	-.150*	-.299*	.347*	-.176*	-.330*	.408*
	(.041)	(.056)	(.040)	(.041)	(.055)	(.040)
UNION	-.473*	-.084	-.608*	-.187*	.267*	-.374*
	(.055)	(.061)	(.055)	(.053)	(.061)	(.054)
Occurrence:	2522	1416	2542	2522	1416	2542
# of observations: 7655						
# of right-censored observations: 1175						

Table 1: UNION=1 if the job is covered under a union contract or collective bargaining agreement; NWHITE=1 if the respondent is black or hispanic; SQAGE=age squared; HS=1 if the respondent is a high school graduate, and COL=1 if he completed 16 or more years of education. Standard errors are given in parentheses. * indicates significant at 5%

contrast, the coefficient estimate from the subhazard regression for firings is positive and statistically significant. The positive coefficient implies that the probability that a job spell ends by a firm's firing decision is higher during a recession. Both of these estimates are consistent with the cyclical behavior of quits and firings.

The results for the effect of u_t are crucial for isolating the effect of u_0, as the duration of a job spell is affected by the current cyclical fluctuations. Bowlus (1995) and Mustre-del-Rio (2012) both include the unemployment rate as an explanatory variable for the hazard regressions. However, the model suffers from misspecification bias if u_t has opposite effects on the decisions of workers and firms. In both of these papers, the coefficient estimate for u_t is statistically insignificant when it is added linearly to the model. To account for the cyclical patterns in quits and firings, Bowlus (1995) further adds the squared value of u_t to the right-hand side variables, and the estimates for the explanatory variables involving u_t becomes significant. By distinguishing job separations according to their causes, I separately identify the effects of current cyclical fluctuations on the duration of a job spell ended by quits and firings. The opposite signs for quits and firings support the discussion about the effect of u_t raised in Bowlus (1995).

4.3 Cumulative Incidence Functions

While the estimates from the subhazard regressions provide a direct inference on the effects of u_0 and u_t on quits and firings, using these estimates to make inferences about the overall duration of employment can be misleading. The subhazard functions for different reasons of job terminations are estimated separately, and the probability of job termination can potentially exceed one when the value of one of the explanatory variables is changed.[9]

To evaluate the overall behavior of employment duration, I use the coefficient estimates from the cause-specific hazard regressions. Note that the estimates of the coefficients from the cause-specific regressions alone are not informative about the effects of u_0 and u_t, although the signs agree with the estimates from the subhazard regressions. Therefore, I obtain the cumulative incidence functions for each job termination category using the coefficient and baseline hazard estimates from the cause-specific hazard regressions. By construction, the probability of job termination is less than unity at any point in time.

Figure 1 shows the cumulative incidence functions for each cause-specific job terminations. The cumulative incidence functions are drawn for a 29 year-old high-school graduate white male whose job is not protected by a union. The unemployment rate is set equal to the average value of the unemployment rate for the survey years, 6.10%, and it is assumed to be equal to this value for all of the time periods from the start of the job. The plots for all three reasons are stacked so that the differences show the probability of observing the corresponding cause-specific job termination before time t. At any time t, the difference between the sum of cumulative incidence functions and one represents the survival probability. Thus, the median duration of a job is 44 weeks.

Figure 2 shows the effects of a change in u_0 on the cumulative incidence functions for quits, firings, and other reasons. In each plot, the solid curves show the cumulative incidence functions when u_0 is equal to its sample mean. The dashed and dotted curves correspond to the cumulative incidence functions when u_0 is one standard deviation, 1.46%, above or

[9] The probability of ending a job exceeds one after 150 weeks when the starting unemployment rate is one standard deviation above its sample mean.

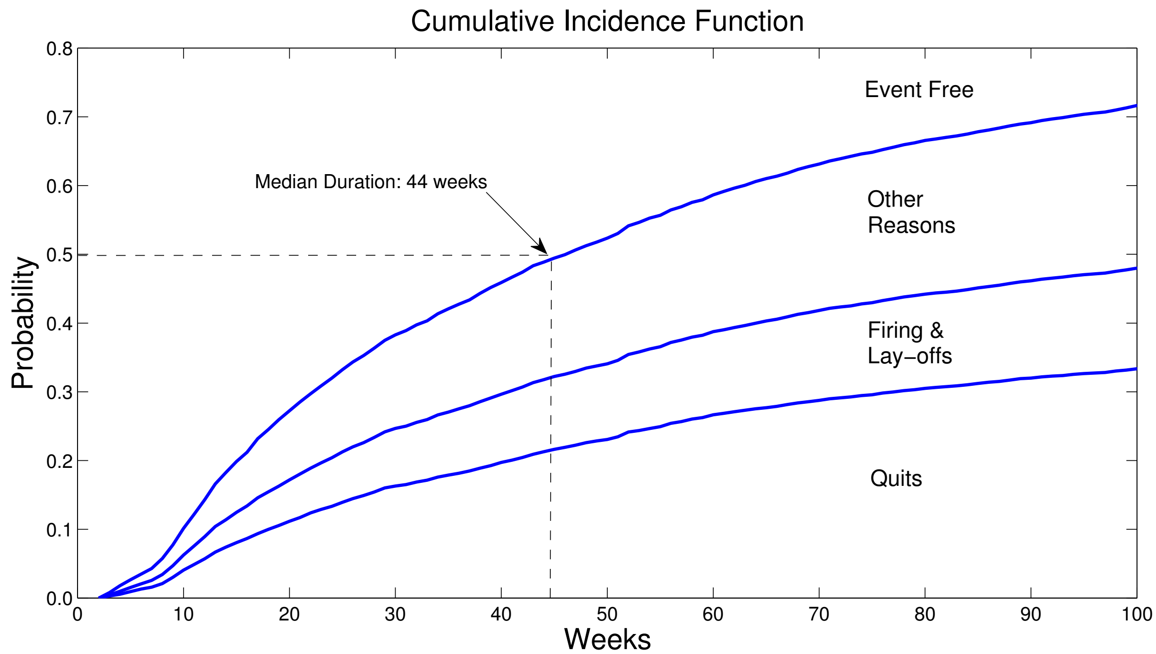

Figure 1: Cumulative incidence functions for quits, firings, and other reasons. The cumulative incidence functions are stacked so that the distance between two curves represents the probabilities of the different events.

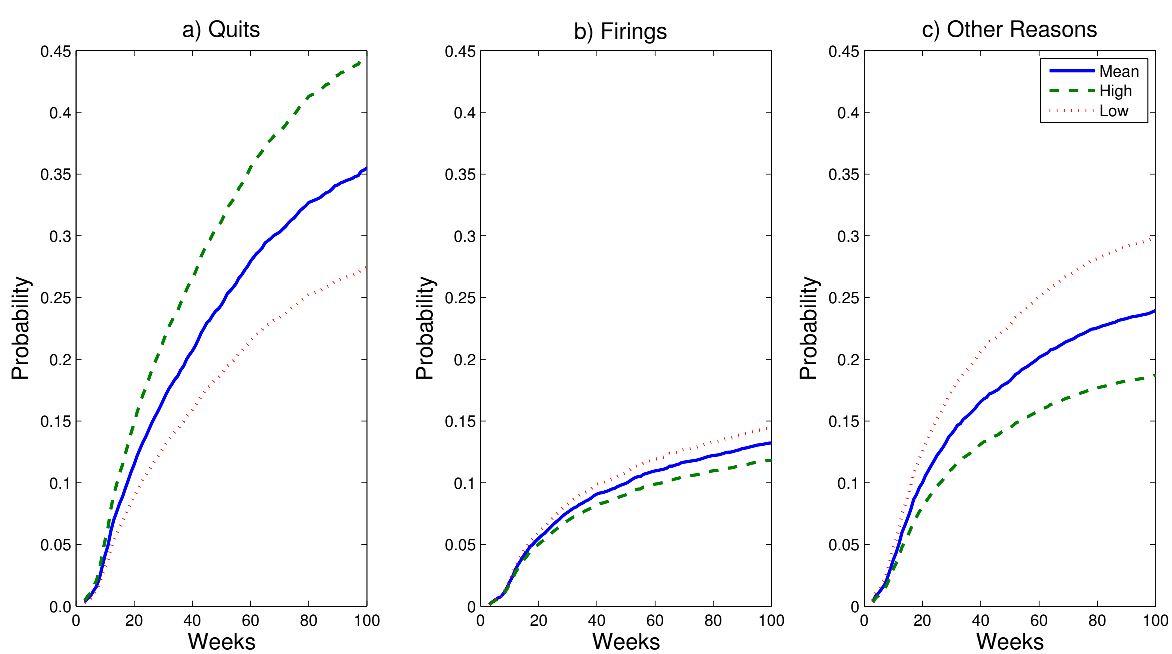

Figure 2: Changes in cumulative incidence functions in response to a change in u_0.

below its sample mean. The current unemployment rate is still kept at its average value for all of the remaining time periods.

The cumulative incidence functions constructed from the cause-specific hazard regressions are consistent with the results from the subhazard regressions. When u_0 is equal to its sample mean, the probability of quitting a job is equal to 0.225 at the median employment duration. This probability increases to 0.288 if u_0 is one standard deviation above its sample mean and decreases to 0.173 when u_0 is one standard deviation below. Firings respond to changes in u_0 in the opposite direction. At the median employment duration, the probability of firing a worker decreases from 0.094 to 0.085 when u_0 is one standard deviation above its sample mean. This probability increases to 0.102 when u_0 is one standard deviation about its sample mean. The behavior of job terminations are similar to those due to firings. At the median employment duration, the probability of terminating a job due to reasons other than quits and firings is equal to 0.173. This probability decreases to 0.137 when u_0 is above its sample mean and increases to 0.215 when u_0 is below its sample mean.

The overall effect of these opposing forces on the duration of a job spell is ambiguous. Taken together with job terminations due to other reasons, the overall effect of u_0 on the duration of employment is negative but small. The duration of employment decreases from 44 weeks to 42 weeks if u_0 is one standard deviation above its sample mean and it increases by only one week when u_0 is one standard deviation below its sample mean. Duration of employment is used as a proxy for match quality in the literature. Therefore, these findings suggest that match quality is weakly pro-cyclical.

Similar results hold for the effects of u_t. Figure 3 shows the change in the cumulative incidence functions for cause-specific job terminations after a change in u_t. In each plot, the solid curves show the cumulative incidence functions when u_t is equal to its sample mean. The dashed and dotted curves correspond to the cumulative incidence functions when u_t is permanently one standard deviation above or below its sample mean for all the periods after the job spell has started.

Changes in u_t affect the cumulative incidence functions constructed from the cause-specific hazard regressions in the same direction implied by the coefficient estimates from the subhazard regressions. At the median employment duration, the probability of quitting a job decreases from 0.225 to 0.187 if u_t is permanently increased by one standard deviation above its sample mean and increases to 0.268 when u_t is permanently one standard deviation below its sample mean. Unlike quits, the probability of being fired increases with u_t as implied by the subhazard regressions. At the median employment duration, the probability of firing a worker increases from 0.094 to 0.116 when u_t is permanently one standard deviation above its sample mean, but it decreases to 0.075 when u_t is permanently one standard deviation

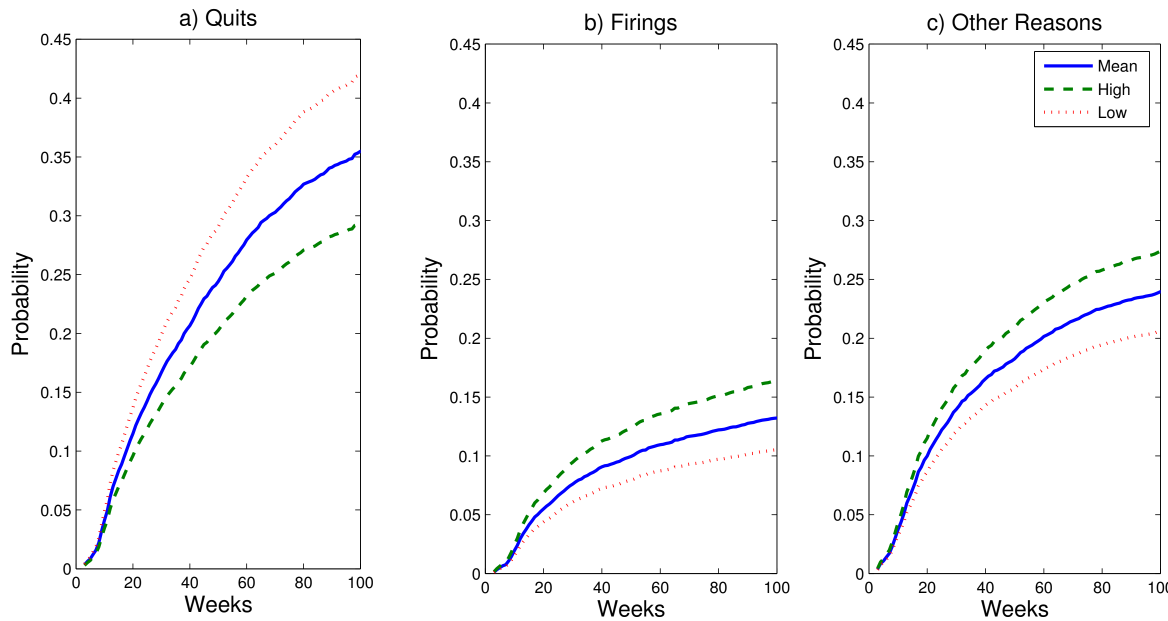

Figure 3: Changes in cumulative incidence functions in response to a change in u_t.

below its sample mean. The response of job terminations due to other reasons is similar to the response of firings. At the median employment duration, the probability of terminating a job due to a reason other than quits and firings increases from 0.173 to 0.198 when u_t is permanently above its sample mean, but it decreases to 0.149 when u_t is permanently below its sample mean.

5 Conclusion

Workers and firms respond to labor market conditions at the start of the employment relationship in opposite directions. In models with on-the-job search, job seekers have an incentive to take a low-quality job during a recession due to increased competition among the job seekers. The workers' incentive to take low-quality jobs implies that jobs created during a recession are likely to have shorter spells. In contrast, models with endogenous separation imply that hiring firms can potentially form better matches because they hire workers from a bigger applicant pool. Therefore, jobs created during a recession are likely to have longer spells. The net effect of these opposing forces on the duration of employment is a priori ambiguous.

In this paper, I empirically test for these theoretical predictions of the effects of labor market conditions on the duration of employment. Using data from NLYS 1979 cohort, I estimate a proportional hazard model under the assumption that different causes of job terminations are competing risks. I use information about the reason why a job spell has ended to distinguish job terminations due to a worker's quit decision from job terminations due to a firm's firing decision. Making a distinction between quits and firings is the main contribution of this paper, because it allows me to test separately for *both* of these opposing forces rather than estimating their *net* effect on the duration of employment.

Two methods have been widely used in the literature to estimate hazard models when there are competing risks. I apply both of these methods in this paper and they produce consistent results. I find that an increase in the unemployment rate at the start of an employment relation increases the probability that the worker quits his job to take or look for another job, but it reduces the probability that the firm fires the worker. These results support both the models with on-the-job search and the models with endogenous separation. Furthermore, the net effect of these opposing forces on the duration of employment is negative. Previous papers in the literature using the NLSY 1979 cohort also find a negative effect, but I find this effect to be much smaller. When the unemployment rate at the start of the employment relationship is one standard deviation above its sample mean, the median duration of a non-union job held by a 29 year-old white male with a high school degree decreases from 44 weeks to 42 weeks. These results suggest that match quality is weakly pro-cyclical.

Pro-cyclical quits and counter-cyclical firings imply that the responses of workers and firms in ongoing employment relationships to current labor market condition also move in opposite directions. My estimates from the proportional hazard model are also consistent with this aggregate behavior. The probability that a job spell ends with worker's quits decision decreases if the current unemployment rate is high. In contrast, the probability that a job spell ends with firm's firing decision is higher when the economy is in recession. The opposite response of quits and firings to the current unemployment rate supports the discussion in Bowlus (1995) about the non-linear response of the employment duration to the current unemployment rate.

My results provide useful evidence for developing theoretical models of labor markets. While jobs created during recessions are on average low quality, this result does not necessarily imply that the predictions of the models with endogenous separation are negligible. Instead, both mechanisms find empirical support and are equally important for theoretical models studying the cyclical behavior of match quality.

References

[1] Akerlof, G., Andrew Rose, and Janet Yellen (1988): "Job Switching and Job Satisfactionin the U.S. LaborMarket.", *Brookings Papers on Economic Activity*, No. 2, pp. 495-582.

[2] Bowlus, A. J. (1995): "Matching Workers and Jobs: Cyclical Fluctuations in Match Quality.", *Journal of Labor Economics*, Vol. 13, No. 2, pp. 335-350.

[3] Cox, D. R. (1972): "Regression Models and Life Tables.", *Journal of Royal Statistical Society*, 34(2): 187220.

[4] Mustre-del-Rio, J. (2012): "Job Duration and the Cleansing and Sullying Effects of Recessions.", *Federal Reserve Bank of Kansas City Research Working Papers*, 12-08.

[5] Kahn, L. B. (2011): "Job Durations, Match Quality and the Business Cycle: What We Can Learn from Firm Fixed Effects.", Unpublished manuscript, Yale School of Management, New Haven.

[6] Mortensen, D. T., and Christopher A. Pissarides (1994): "Job Creation and Job Destruction in the Theory of Unemployment.", *Review of Economic Studies*, 61(3): 397415.

[7] Mortensen, D. T. (1994): "The Cyclical Behavior of Job and Worker Flows.", *Journal of Economic Dynamics and Control*, 18: 1121-1142.

[8] Putter, H., M. Fiocco, and R. B. Geskus (2007): "Tutorial in Biostatistics: Competing Risks and Multi-State Models.", *Statistics in Medicine*, 26: 23892430.