

MIT Graduate Labor Economics 14.662 Spring 2017
Lecture Note 6: Job Loss and Job Search at the Micro
and Macro Level

David H. Autor

MIT and NBER

March 17, 2017

1 INTRODUCTION

As we discussed early in the semester, jobs play a surprisingly small roll in labor economics. Much of contemporary labor economics focuses on *skills* not *jobs* as the primary determinant of wages. If wages are set according to the Law of One Price for skills (LoOP), then we can pretty much ignore industry, occupation, firm, union power, minimum wages, rents, and even employment, unemployment and job search when thinking about wage determination. Extreme as that sounds, it may not be bad an approximation to the truth over the long run: perhaps skills *are* the key determinant of earnings over the course of a career. But over shorter horizons, these ‘local’ (some would say ‘frictional’—though I don’t like this term) determinants of earnings matter, and potentially matter a lot. A number of insights and observations have reinvigorated interest in these non-LoOP factors over the last decade:

1. The literature on job ‘tasks’ as a determinant of earnings has reanimated the (old) idea that workers are not necessarily paid for the skills that they possess but rather the work they accomplish. If so, the assignment of skills to tasks is likely important to wage determination. This suggests focusing on occupations (or individual jobs) as well as skills as objects of theoretical and empirical interest.
2. The growing availability of matched firm-worker data has facilitated the discovery of the importance of firms in wage determination. The leading modern exemplar of this view is the 2013 *QJE* paper by Card, Heining and Kline, but there are many papers on your reading list (and elsewhere) using similar data. The leading U.S. example is Song et al. “Firming Up Inequality.”
3. The Great Recession has reminded economists of the extremely large direct costs of job loss, as well as the potential for *hysteresis*, whereby current job loss may have a direct effect on future employment and earnings potential.

This subsection of 14.662 studies job loss and job search. There are three strands of literature I’d like to discuss. One focuses on the consequences of job loss at the worker level. A second studies the job search process—how workers escape unemployment. A third literature, which is empirically quite recent, studies interactions (pecuniary externalities) among job searchers in the labor market.

2 CONSEQUENCES OF JOB LOSS

2.1 JLS 1993

The starting point for the modern empirical literature on job loss is Jacobson, LaLonde and Sullivan’s 1993 paper, “Earnings Losses of Displaced Workers.” This paper innovates on at least two fronts. First, it mixes worker-level microdata, specifically quarterly earnings history for a 5% sample of Pennsylvania workers, with firm-level data on employment levels, growth, geography, and detailed industry. Such research data were almost unheard of at the time. Second, the paper develops a credible and transparent event study framework for estimating the earnings consequences of job loss. Two decades after its publication, researchers are still closely replicating JLS using data from around the world, and largely confirming its original findings.

As with any causal inference problem, the challenge that JLS face is establishing a valid counterfactual, here for the earnings of job losers had they not lost employment. For a twenty-five year old paper, the statement of the inference problem is remarkably modern, “In this study, we define displaced workers’ earnings losses to be the difference between their actual and expected earnings had the events that led to their displacement not occurred.” They take as their empirical target

$$E[Y_{it}|D_{i,s} = 1, I_{i,s-p}] - E[Y_{it}|D_{i,s} = 0, I_{i,s-p}] \quad \text{for all } I_{i,s-p},$$

where s is the date of separation, p is some number of periods, I_{s-p} is all information available at time $s - p$ prior to separation that would be expected to affect the evolution of earnings (prior to whatever events lead to separation), and $D_i \in \{0, 1\}$ designates separation (i.e., displacement). This expression is not estimable since we do not observe expected earnings over future states. Instead, JLS estimate

$$y_{it} = \alpha_i + \gamma_t + \mathbf{x}_{it}\beta + \sum_{k \geq -m} D_{it}^k \delta_k + \varepsilon_{it},$$

where k represents the quarter relative to job displacement and $k - m$ is quarters relative to displacement. They apply a minimum value of $m = 20$, meaning that they must observe worker earnings at least twenty quarters (five years) prior to displacement. Further specifications also add a worker-specific trend

$$y_{it} = \alpha_i + \omega_i t + \gamma_t + \mathbf{x}_{it}\beta + \sum_{k \geq -m} D_{it}^k \delta_k + \varepsilon_{it}.$$

For workers separating from ongoing firms, JLS can also estimate models that include firm effects interacted with time dummies, meaning that the specification compares displaced from non-displaced workers of the same employer.¹ JLS parameterize further for econometric efficiency to allow the trajectory of displaced workers to vary along three dimensions: 1) a dip in pre-displacement earnings; 2) a drop at displacement; and 3) a trajectory of recovery. You'll find these specifications carefully explained in the paper.

An obvious problem with studying worker separations is that these separations may be endogenous, and it's plausible that the worst workers (and, more troublingly, the workers with the worst expected earnings trajectories) disproportionately comprise the separators. The approach that JLS pursue to blunt this concern is to construct a "mass-layoff" sample that includes separators whose firms' employment in the year following their departure was 30-percent or more below their maximum levels during the late 1970's. This definition encompasses firms that closed around the time of workers' separations, as well as others that had large employment declines. This definition of "mass layoffs" has been remarkably persistent since its inception (perhaps unjustifiably so). A modern alternative to conditioning on the endogenous outcome of a worker separating during a mass layoff event is to instrument a worker's separation with the firm-level separation event.

2.1.1 SOME KEY FINDINGS FROM JLS

- High tenure, prime-age workers endure substantial and persistent earnings losses when displaced during or following mass layoffs. Even six years after their separations, their quarterly earnings remain \$1,600 below their expected levels. This loss represents 25 percent of their pre-displacement earnings. The estimated loss is even larger controlling for worker-specific time trends, suggesting that employers do not systematically displace workers whose earnings were already growing slowly. JLS observe that "because the estimated losses do not decline significantly after the third year following their separations, there is little evidence that displaced workers' earnings will ever return to their expected levels."
- The quarterly employment rates of to-be displaced workers in the sample depart only

¹This empirical setup makes computational demands that were also unheard of in its day. For at least a decade after the paper was published, Dan Sullivan kept running in his office the NeXT Unix computer on which he did the empirical work. For interested geeks: NeXT was a hardware and software company founded by Steve Jobs after his departure from Apple in 1985. NeXT was not commercially successful but its technology was legendary. Dan Sullivan's NeXT box, for example, continued working flawlessly without a reboot through at least 2003 (that's when I last saw it). Incidentally, Apple acquired NeXT in 1996 and Steve Jobs rejoined the company in 1997. The NeXTSTEP operating system became the foundation for MacOS, which is the current Macintosh operating system. MacOS is now a certified variant of the Unix 03 Standard.

slightly from their expected levels, except for the year following separation. Thus, the substantial earnings losses observed in the data are largely due to lower earnings for those who work, rather than an increase in the number of workers without quarterly earnings. This is somewhat at odds with the working paper by Gregor Jarosch (2016), “Searching for Security and the Consequences of Job Loss,” which finds serial correlation in non-employment for displaced workers.

- The events that lead to worker separations cause earnings to depart from their expected levels even before they leave their firms. Figure 2 shows that to-be displaced workers’ quarterly earnings begin to diverge meaningfully from their expected levels approximately three years prior to separation. That divergence accelerates during the quarters immediately prior to separation, so that by the quarter prior to displacement, these workers’ earnings are approximately \$1,000 below their expected levels.
- A different earnings loss pattern emerges for workers from the non-mass-layoff sample. Prior to separations, their earnings depart only slightly from their expected levels. Following separation, their earnings drop by only one-half as much as workers in the mass layoff sample. This group’s earnings fully recover 3-5 years following their separations. (This subgroup probably includes larger fractions of workers who quit their jobs or who had fewer firm-specific skills.) This set of results suggests that there’s nothing mechanical in the empirical setup that necessarily generates large loss estimates.
- As shown by Figure 4, when JLS use non-displaced workers in displaced workers’ former firms as a comparison group, the estimated earnings losses are smaller by about 20 percent. For example, five years after separation, displaced workers’ quarterly earnings are \$1,200 below (compared to \$1,500 below) their expected levels. JLS take this as evidence that the choice of comparison group matters little, but that’s not my read. At a minimum, it suggests that the mass-layoff ‘treatment’ is not clean; it’s correlated with the time-varying firm effect, even for stayers. It should be possible to use the reported treatment effects for the full sample, the mass-layoff sample, and the mass-layoff within-firm comparison, to triangulate what the 2SLS estimate would find.
- There is considerable heterogeneity in the earnings losses by demographic group. Losses do not differ much by birth cohort, but they are much larger for workers employed by very large firms and by primary metal industries. This pattern likely reflects the collapse of the U.S. steel sector in the early 1980s. The steel industry was a major Pennsylvania employer, and it’s the largest single contributor of observations to the separator sample.

- A key puzzle raised by this paper is why earnings losses from job displacement are so *large*. Remarking on this, JLS write, “Our figures indicate that workers displaced in the best of circumstances have losses that are at most only one-third less severe than the average losses depicted in Figure 2. This finding suggests that some valuable attribute of the employment relationship itself is lost when high-tenure workers are displaced.” What precisely is lost?
- A noteworthy pattern is that losses are twice as large for workers leave the original sector, but are still substantial for those who find new employment in the same sector (same 4-digit industry or back in manufacturing or non-manufacturing). JLS write, “The findings for both displaced manufacturing and non-manufacturing workers indicate that a substantial portion of their earnings losses result from the loss of some highly firm-specific component of earnings. Even those who found new employment in the same industry experienced large and persistent losses.”
- One interesting observation in rereading this twenty year old paper is how *short* it is relative to contemporary top journal papers. This brevity does not reflect a lack of substance; what *is* absent is the innumerable extensions, specification checks, and non-illuminating models that one would typically find in a contemporary paper. It’s hard to feel nostalgic for the ‘good old days’ of empirical work because most of it was mostly terrible. JLS, however, is an early example from the ‘credibility revolution,’ showing the potential of transparent empirical work applied to (a 1990s version of) ‘big data’ to produce self-evidently compelling results. Contemporary top journal papers are generally more exhaustive (and *exhausting*) than this 1993 article, but this is not obviously all to the good. (Indeed, the 2009 follow-up paper by Sullivan and Von Wachter, which is substantively important but not methodologically innovative, is nearly twice as long at 42 vs. 25 pages.) The American Economic Association will soon launch a new short-format journal, one objective of which is to counter the trend towards logorrhea in economics publishing.

2.2 NON-PECUNIARY COSTS OF JOB LOSS

The 2009 paper by Jacobson and Von Wachter is 15 year follow-up to JLS, studying the effect of job loss on worker mortality by matching Pennsylvania earnings records to Social Security Administration death records covering 1980–2006. JvW’s UI records are identical to those used by JLS except that they have been extended by four years (JLS have 1974 to 1986, JvW have 1974 to 1990). JvW include workers born between 1920 through 1959 whereas JLS limited the sample to those born after 1930 to 1959; including workers who are

a decade older increases power for analysis of mortality outcomes. The results of this paper are at once unsurprising and stunning. They suggest that involuntarily displaced workers lose approximately 1.3 to 1.6 years of life, with the largest losses suffered by the workers who are youngest when displaced. This finding underscores the fact that job loss is frequently a traumatic event, costing workers not only earnings but a also sense of identity and purpose.²

2.3 RECESSIONS AND THE COST OF JOB LOSS

The 2011 *Brookings* paper by Davis and Von Wachter offers the most comprehensive and up-to-date follow-up to the original JLS paper. It addresses three questions:

1. Are the original JLS findings broadly upheld outside of the specific context of job loss in Pennsylvania during the 1980-1983? This recession wiped out the steel sector and a large swath of U.S. manufacturing, so there's good reason to believe that job loss had particularly scarring effects over this period.
2. How do the costs of job loss vary with economic conditions at the time of job loss?
3. Can these high costs be rationalized by contemporary models of job search (building on Mortensen and Pissarides, 1994)?

Using national Social Security Longitudinal earnings records for 1974 through 2008, Davis and Von Wachter estimate that cumulative earnings losses associated with job displacement average 1.4 years of pre-displacement earnings for males who are displaced in mass-layoffs when the national unemployment rate is below 6 percent. These estimated losses are doubled—2.8 years of pre-displacement earnings—if displacement occurs when the unemployment rate exceeds 8 percent. These findings echo Couch and Placzek's 2010 *AER* paper, which replicates JLS using Connecticut data and finds impacts that are large and robust but nevertheless 25 to 40 percent smaller than those in JLS. Couch and Placzek argue that this discrepancy reflects the contrast between the severe Pennsylvania recession studied by JLS and the comparatively placid labor conditions prevailing in Connecticut during 1993 through 2004.

Davis-Von Wachter provide an expansive and thoughtful characterization of the key facts (in the data and in the world) on the earnings consequences of job loss. If you plan to read only one paper from the JLS-inspired literature, this is the one to read. (I've excerpted many key results in my lecture slides.)

²A 2015 article in the *Annual Review of Sociology* by Jennie E. Brand reports that job displacement is associated with "...declines in psychological and physical well-being; loss of psychosocial assets; social withdrawal; family disruption; and lower levels of children's attainment and well-being. Although reemployment mitigates some of the negative effects of job loss, it does not eliminate them."

Related but not identical is the widely cited 2012 *AEJ: Applied* paper by Oreopoulos, Von Wachter and Hsiez, “The Short- and Long-Term Career Effects of Graduating in a Recession.” This paper documents durable scarring effects of graduating college during a recession that reduce earnings and raise the frequency of job change over a decade. Figure 1a of this paper is brilliant. The counterpart to the adverse earnings outcomes experienced by workers graduating in recessions is the *improvement* in the quality of individuals entering some lower-ranked occupations in slack labor markets. A 2015 *AEJ: Policy* paper by Stevens, Miller, Page, and Filipksi finds that staffing levels in nursing homes moves counter-cyclically, which may partly explain why most additional deaths that occur when the labor market is tight are among the elderly, particularly elderly women and those residing in nursing homes. The 2015 NBER working paper by Markus Nagler (who is auditing 14.662 this spring), Marc Piopiunik, and Martin West finds that individuals entering the teaching profession during recessions are significantly more effective in raising student test scores than average teachers.

The DvW paper ‘confronts’ leading models of unemployment fluctuations with evidence on the present-value earnings losses associated with job displacement. Their chief conclusion is that, “the 1994 model of Dale Mortensen and Christopher Pissarides, extended to include search on the job, generates present-value losses that are only one-fourth as large as observed losses. Moreover, present-value losses in the model vary little with aggregate conditions at the time of displacement, unlike the pattern in the data.” Bob Hall’s published comments on the DvW explains the puzzle quite succinctly (one of the lovely features of *Brookings* papers is that discussant and audience comments are published with the papers):

“Evidence on the cost of recruiting suggests that, right after a new hire, an employer has about \$1,000 invested in the worker. The bargaining structure of the DMP model interprets this amount as the employer’s capitalized share of the surplus the job generates. If the bargain splits the surplus roughly equally, the worker has a similar stake in the job. The worker’s loss from a layoff that occurs immediately after the hire would thus be about \$1,000, which is far below the figure that the paper calculates for the typical layoff occurring 3 or more years after the hire.

The DMP model as normally developed is focused on unemployment and is exceedingly stripped down with respect to how the typical employment relationship evolves after the hire. All that matters for the analysis of unemployment is the present value of the expected margin the employer will earn from the relationship from the difference between the worker’s productivity and the worker’s wage. Given the objective of the model, it is no shortcoming that the model cannot generate realistically big figures for the consequences of job loss.

In their conclusion, the authors lay out some of the ideas from labor economics that would belong in a master model of the employment relationship that deals both with the issues that gave rise to the DMP model and with many issues of governance of the ongoing relationship. I think the paper performs an important service in making it clear that the master model faces an important challenge in explaining how workers move from having, on average, only a roughly \$1,000 stake in a brand-new job to having around \$100,000 at stake after more than 3 years of tenure. The implied gradient of accumulation of the worker’s share of job- specific capital is remarkably steep, and thus a real challenge to empirical model builders.”

Gregor Jarosch’s 2016 working paper (“Searching for Job Security and the Consequences of Job Loss,” which was Gregor’s job paper) takes up this challenge.³ Working with German Social Security records, Jarosch offers a novel observation about why earnings losses from job loss appear so persistent: job separations appear to lead to subsequent separations; that is, they are serially correlated. This finding is somewhat at odds with DvW, who write, “the findings from administrative data pertain to annual or quarterly earnings. Hence, the earnings losses potentially arise from reductions in both employment and wages. However, the earnings loss for the median worker in the sample is about as large as, and more persistent than, the mean loss (von Wachter and others 2011, Schoeni and Dardia 2003). This result and survey-based evidence that most job losers return to employment (for example, Farber 1999) suggest that the bulk of earnings losses after job displacement reflects reductions in wage rates or hours worked.” Nevertheless, the German data are well suited to exploring this possibility because they record daily wages, something lacking in quarterly U.S. unemployment earnings records.

Proceeding on the idea that job loss catalyzes further job losses, Jarosch writes a quality ladder model where workers move upward into more productive and stable jobs through the process of on-the-job search. When workers are involuntarily displaced, they have to start again at the bottom of the ladder, where job security is lower. This leads to serial correlation in job separations—that is, one involuntary separation causes (in expectation) subsequent separations. Jarosch pairs this idea with the assumption of a stochastic human capital acquisition process: workers gain human capital while working and lose it while unemployed. This amplifies the power of the job quality ladder for explaining the data: because job loss leads to both further job loss and a deceleration or reversal in the human capital acquisition process, these two forces can jointly explain why job stability following

³Notably, Jarosch was the RA for the DvW paper, which underscores the value of RA’ing for your human capital acquisition!

displacement eventually recovers but earnings do not fully rebound. Jarosch’s paper offers a fresh answer to a longstanding puzzle that, in my view, the literature has not satisfactorily grappled with. At present, I find it difficult to assess the degree to which Jarosch’s model offers novel insight into the data versus merely rationalizing the facts that that the model was built to explain.⁴

2.4 THE CAUSAL EFFECT OF UNEMPLOYMENT DURATION ON OFFER WAGES

The sophisticated paper by Schneider, Von Wachter and Bender (AER 2016), “The Effect of Unemployment Duration on Wages: Evidence from Unemployment Insurance Extensions,” presents direct evidence on the causal effect of unemployment duration on offer wages (where offer wages can be thought of as a measure of the market’s perception of a worker’s skill/productivity/human capital). The empirical challenge that the paper carefully outlines is that even with exogenous increases in unemployment durations (induced by age discontinuities in the German UI system), duration may affect reemployment wages through two distinct channels: a causal effect on offer wages (this effect will generally be negative) and a causal effect on reservation wages (which could in theory be positive or negative). The fairly in-depth methodological discussion in SVWB leads to the following straightforward empirical approach. First, test whether there is any effect of UI extensions on reemployment wages *conditional on unemployment duration*. If not, this implies that the reservation wage channel is not binding – meaning that, in practice, offered wages are generally higher than reservation wages – and hence any causal effect of UI extension on wages will operate through changes in *offer wages* as a function of unemployment duration. Next, estimate the effect of extended UI on reemployment wages without controlling for realized duration: given the conclusion above, this can then be interpreted as the causal impact on offer wages. (This topic is covered formally in your second problem set.)

Kroft, Lange and Notowidigdo (QJE 2013) assess the causal effect of time out of the labor force on interview callbacks using a resume audit study performed in 100 cities. Kroft et al. find that the likelihood of receiving a callback for an interview significantly decreases with the length of a worker’s unemployment spell (as implicitly reported on the worker’s CV). The majority of this decline occurs during the first eight months of search. Duration dependence is stronger when the local labor market is tighter, which is consistent with screening models where employers use a worker’s unemployment spell length as a signal of unobserved productivity and recognize that this signal is less informative in weak labor markets. Like all resume audit studies, this paper cannot draw a tight link between callbacks and ultimate

⁴Queue an interesting philosophical debate about what a model is *supposed* to do: rationalize facts, guide inquiry, predict untested facts (like the Higgs boson) or some combination of these things.

employment or wage outcomes.

The 2016 working paper by Jarosch and Pilossoph, “Statistical Discrimination and Duration Dependence in the Job Finding Rate,” writes down a model where these initial callback effects (in Kroft et al.) don’t actually matter much because employers were very unlikely to hire workers with longer unemployment durations conditional on interviewing them (which is why employers don’t bother to do the interviews). This is an interesting hypothesis, but one could easily write down models that reach a different conclusion (e.g., based on the model and evidence in Pallais’ 2014 *AER* paper, “Inefficient Hiring in Entry Level Labor Markets”). The question requires empirical testing. I’d say that Schmeider et al. provides reasonably compelling evidence that unemployment durations do affect employment outcomes.

3 MICRO-MACRO LINKAGES IN JOB SEARCH

As underscored by Davis and Von Wachter, Oreopoulos et al., and Kroft et al., the consequences of job loss (and the gains from job search) depend upon the business cycle. Workers losing and seeking jobs in relatively slack labor markets suffer larger losses and make smaller gains than workers losing and seeking jobs in tighter labor markets. One explanation for this pattern is that ‘worse’ jobs are offered during recessions. A second possibility is that labor market congestion is bad for job seekers: when there are comparatively more job seekers than job vacancies, job seekers fare worse. This latter possibility implies that even in normal (non-recessionary) job markets, there may be important (negative) pecuniary externalities among groups of job seekers. If so, this would have important implications for job search assistance policies, among other topics. Two important recent papers provide evidence on this point: Lalive, Landais, and Zweimüller (AER 2015), “Market Externalities of Large Unemployment Insurance Extension Programs,” and Crépon, Duflo, Gurgand, Rathelot, and Zamora (QJE 2013), “Do Labor Market Policies Have Displacement Effects? Evidence from a Clustered Randomized Experiment.” Both papers offer both theory and data on this question. The theoretical models are closely related and, more remarkably, the evidence is as well. It’s worth exploring at least one of the two models. I’ll focus on the model in Crépon et al.

3.1 MODEL SETUP

- There is one sector with one type of worker. The labor force is normalized to size $n + u = 1$, where n and u are the fraction employed and unemployed respectively.
- When workers are employed, jobs end randomly with probability s .
- Unemployed people search for jobs and firms open vacancies to hire them. Denote total

job search effort exercised by the unemployed as u_e and total opened vacancies as v .

- The number of matches resulting from the aggregated search effort and available vacancies is given by the standard Diamond-Mortensen-Pissarides matching function $m(u_e, v)$. This function is increasing and concave in both of its arguments and is homogeneous of degree one.
- The tightness of the labor market is $\theta = \frac{v}{u_e}$. This implies:
 - The probability that a vacancy is filled is $\frac{m(u_e, v)}{v} = m\left(\frac{u_e}{v}, 1\right) = m\left(\frac{1}{\theta}, 1\right) = q(\theta)$, which is decreasing in θ .
 - The probability that an unemployed worker exercising one unit of search effort finds a job is $\frac{m(u_e, v)}{u_e} = \frac{m(u_e, v)}{v} \times \frac{v}{u_e} = \theta q(\theta) = f(\theta)$, which is increasing and concave in θ (given the assumptions on the matching function).
- Assume initially that everyone exerts search effort 1, but when workers become unemployed, a fraction π is assigned to job search assistance, which increases their search effort (or search productivity) to $e > 1$. (These assumptions are relevant to the empirical setting in Crépon et al. For purposes of the Lalive et al. paper, one could assume that a fraction π of workers qualify for REBP, which reduces their search effort to $e < 0$.)
- In steady state, there are two types of workers, untreated u_0 and treated u_1 , and total search effort is $u_e = u_0 + eu_1$. Treated workers will have higher exit rates from unemployment. They account for a share eu_1/u_e of search effort and hence the same share of job offers: $\frac{eu_1 m(u_e, v)}{u_e} = eu_1 f(\theta)$. The exit rate of treated workers is therefore $eu_1 f(\theta)/u_1 = ef(\theta)$. Similarly, the exit rate of untreated workers is $f(\theta)$.
- Displacement effects will be observed if changes in search effort by a subset of workers leads to a change in the overall tightness of the labor market θ .

3.2 EQUILIBRIUM

- In steady state, inflows and outflows from each group must be equal

$$u_1 e f(\theta) = s \pi n,$$

$$u_0 f(\theta) = s(1 - \pi)n$$

- Using the fact that $1 - n = u = u_0 + u_1$, we can rearrange terms to obtain a relationship between the equilibrium number employed n and labor market tightness θ .

$$\begin{aligned}
1 - n &= \frac{s\pi n}{ef(\theta)} + \frac{s(1 - \pi)n}{f(\theta)} \\
&= \frac{s\pi n + es(1 - \pi)n}{ef(\theta)} \\
&= n \frac{s(\pi/e + (1 - \pi))}{f(\theta)}
\end{aligned}$$

$$\begin{aligned}
1 &= n \left[\frac{s(\pi/e + (1 - \pi))}{f(\theta)} + 1 \right] \\
1 &= n \left[\frac{s(\pi/e + (1 - \pi)) + f(\theta)}{sf(\theta)} \right]
\end{aligned}$$

$$n = \frac{f(\theta)}{s(\pi/e + (1 - \pi)) + f(\theta)} \tag{1}$$

- This can be thought of as a *supply* curve of workers. The tighter the labor market, θ , the faster workers are reemployed conditional on separation, the greater number of workers in status n rather u . We can write $\theta = \theta_B(n)$ as the required level of tightness to produce employment level n , where θ is increasing in n .
- This locus is plotted in Figure I of Crépon et al. It's crucial to note that the $\theta_B(n)$ curve is fairly flat for low levels of employment (low θ) and steep when employment is high. This follows because the function $f(\theta) = m(\theta, 1)$ is concave (increasing) due to the constant returns to scale assumption for the matching function (thus, it takes ever larger increments to θ to raise n by a given amount). This implies that the function $\theta_B(n)$ is convex.

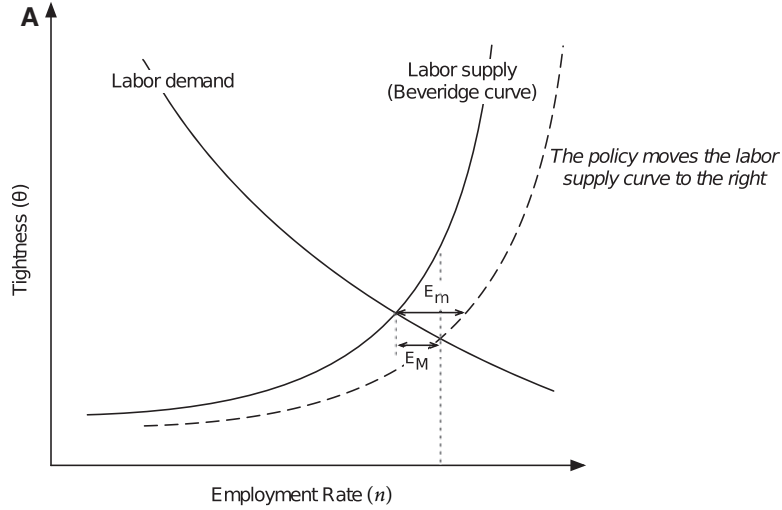


Figure 1a from Crépon et al. (2013)

- Now we need the labor demand side of the equation. Following Landais, Michaillat, and Saez (2012), the paper assumes that the aggregate production function has diminishing returns to scale. This is a non-standard assumption, but it's not crazy: in the short run, an increase in the number of available workers increases employment less than proportionally; that is, labor demand is not perfectly elastic. Specifically:

$$y = an^\alpha, \alpha \in (0, 1).$$

- An equally central assumption is that wages are *fixed* in the short-run at w_0 . (This possibility is something that Lalive et al. explore and test.) Let r equal the interest rate, c the flow cost of having a vacancy, J_V the asset value of a vacancy and J_E the asset value of a filled job. In equilibrium, a firm must make zero profit from opening a vacancy. The system of Bellman equations is:

$$\begin{aligned} rJ_V &= -c + q(\theta)(J_E - J_V), \\ rJ_E &= \alpha an^{\alpha-1} - w_0 + s(J_V - J_E). \end{aligned}$$

Simplifying

$$J_E = \frac{\alpha an^{\alpha-1} - w_0 - sJ_V}{r + s},$$

and substituting into J_V gives:

$$\begin{aligned}
 J_V &= -c + q(\theta) \left(\frac{\alpha an^{\alpha-1} - w_0}{r+s} - J_V \right) \\
 J_V (1 + q(\theta)) &= -c + q(\theta) \left(\frac{\alpha an^{\alpha-1} - w_0}{r+s} \right) \\
 J_V &= \frac{-c}{(1 + q(\theta))} + \frac{q(\theta)}{(1 + q(\theta))} \left(\frac{\alpha an^{\alpha-1} - w_0}{r+s} \right) \\
 0 &= \frac{-c(r+s) + q(\theta)(\alpha an^{\alpha-1} - w_0)}{(1 + q(\theta))(r+s)} \\
 0 &= q(\theta)(\alpha an^{\alpha-1} - w_0) - c(r+s) \\
 0 &= \alpha an^{\alpha-1} - w_0 - \frac{c(r-s)}{q(\theta)} \tag{2}
 \end{aligned}$$

- This equation implies a decreasing relationship between θ and n . A higher θ leads to a lower fill rate, $q(\theta)$, meaning that the expected cost of a vacancy rises with the labor market tightness. Thus, demand is downward sloping in θ . Implicitly these two equations in θ lead to an equilibrium employment level n

$$\left(\frac{q(\theta^*) w_0 + c(r-s)}{\alpha a q(\theta^*)} \right)^{1/(\alpha-1)} = \frac{f(\theta^*)}{s(\pi/e + (1-\pi)) + f(\theta^*)}$$




Figure 1b from Crépon et al. (2013)

3.3 EMPIRICAL IMPLICATIONS: THE RIGID WAGE CASE

Crépon et al. assume throughout that the equilibrium wage is fixed at w_0 . This is a strong assumption that greatly simplifies the analysis.

1. Starting from an initial situation where $\pi = 0$ and $e = 1$, an increase in the fraction counseled ($\pi > 0 \rightarrow e > 1$), generates a rise in the denominator of the labor supply equation $\pi/e + 1 - \pi$, which implies a decrease in equilibrium tightness (the labor supply curve shifts rightward). Overall employment rises. But this means that employment of non-treated workers falls because their search effort has not risen but labor market tightness has fallen. This externality can be visualized by comparing the ‘micro’ and ‘macro’ responses of employment. The micro response is the rightward shift in the labor supply curve, not accounting for the slope of the demand curve. This rightward shift is the amount a single individual would gain in expected employment from a unilateral increase in search effort. The ‘macro’ effect accounts for the movement along the demand curve, which partly offsets the micro effect since the labor market slackens as search effort intensifies.
2. A second implication is the size of the externality on non-treated workers is increasing in the share treated, π .
3. A third implication comes from the slope of the labor supply curve. As noted above, this curve is shallow when the labor market is slack and steep when it is tight. This implies, as shown in Figure Ib, that a given size rightward shift in the labor supply curve has little net positive effect on aggregate employment when the labor market is slack and a larger net effect when the labor market is tight. Thus, externalities from enhanced search are predicted to be greater in a slack market.

3.4 WHAT IF WAGES ARE *not* RIGID?

- The model developed in the online Appendix of Lalive, Landais and Zweimüller (LLZ, *AER* 2015) is quite similar to Crépon et al. but it also considers the possibility that labor market interventions may affect equilibrium wages. The intervention that LLZ have in mind is the converse of Crépon et al.: a subset of workers is given extremely generous long-term unemployment insurance that is expected to reduce their job search effort. They are interested in analyzing direct and spillover impacts on employment and earnings of treated workers and untreated workers residing in the same local labor market.

- LLZ are agnostic on how w_0 and w_1 (wages of treated and untreated) are affected by treatment and so consider multiple cases. If treatments directly affect wages (by shifting reservation wages as well as search behavior), then labor demand will also depend on the treatment (through the dependency of vacancies on workers' reservation wages, see equation (2)). If the treatment does not change reservation wages, then equilibrium will be affected only through the effect of search intensity on labor market tightness (see equation (1)).
- Figure Ia and Ib of LLZ (loosely) sketch these two cases. In the first case, wages are exogenously set. An increase in benefits given to the treatment group B_1 reduces search effort by treated workers and increases labor market tightness but does not change the locus of the demand curve; that is, for given tightness θ , there is no change in quantity demanded. Total employment falls. But holding job search effort constant, greater tightness raises employment rates of non-treated workers. Thus, if we assume that effort of *non-treated* workers is unaffected by B_1 , then their employment must rise. LLZ refer to this as a labor demand externality.
- Conversely, if reservation wages of treated workers rise ($\partial w_1/\partial B_1 > 0$), and firms cannot discriminate between treated and non-treated workers in hiring and wage setting (which is a bit of an awkward assumption), then the demand curve shifts inward. LLZ refer to this as a wage externality. Depending on which effect dominates (the inward labor supply shift, which tightens θ , or the inward labor demand shift, which slackens θ), employment of untreated workers could either rise or fall.
- Three summary points on this somewhat rough model:
 1. If wages are independent of benefits, and returns to labor are decreasing (that is, the demand curve is downward sloping), then $\partial\theta/\partial B_1 > 0$ and therefore employment externalities on non-treated workers should be positive (Figure 1a).
 2. If returns to labor are almost constant (labor demand is very elastic) and wages are strongly correlated with outside options of workers (B enters the wage bargain), then $\partial\theta/\partial B_1$ may be negative, in which case, employment of non-treated workers will be reduced by an increase in B_1 (though this case seems very unlikely to me, especially because we'd assume that employers can offer different wages to treated and untreated workers—or, equivalently, choose not to hire workers with artificially high reservation wages).
 3. The equilibrium adjustment of θ will partly depend on the fraction of unemployed who receive increased benefits. Thus, treatment intensity, defined as the frac-

tion of individuals eligible for extended benefits, also determines the magnitude of the externalities. In a model with rigid wages and diminishing returns, the positive externality on untreated unemployed should increase with the fraction of unemployed who are treated. [Again, bear in mind that the ‘treatment’ in LLZ is opposite in sign to the treatment in Crépon et al.]

3.5 RESULTS: LLZ

The Crépon paper fields and analyzes an extremely ambitious and well-designed RCT. It’s very difficult to develop experiments that are high-powered enough to credibly affect macroeconomic aggregates, of course, and the Crépon paper is somewhat underpowered as a result (though still impressive). The LLZ paper, by contrast, analyzes a startlingly high-powered quasi-experiment, and the results are remarkable. The findings of these two papers are highly consistent, which adds to their credibility.

To protect its assets after World War II from Soviet appropriation and to provide the capital needed for reconstruction, Austria nationalized its iron, steel, and oil industries, large segments of its heavy engineering and electrical industries, most coal mines, and the nonferrous metals industries. Firms in the steel sector were part of a large holding company, the Oesterreichische Industrie AG, OeIAG. By the mid-1970s this holding company was in bad shape due to shrinking markets, overstaffing, over-concentration in outmoded smokestack industries, insufficient research and development, and low productivity. Initially, the Austrian government covered the losses by subsidies. But in 1986, after the steel industry was hit by an oil speculation scandal and failure of a U.S. steel plant project, this protectionist policy was abolished. New management was appointed and a strict restructuring plan was implemented. This led to layoffs due to plant closures and downsizing, particularly in the steel industry.

To mitigate labor market impacts in heavily exposed regions, the Austrian government enacted a law that extended UI entitlement to 209 weeks for workers who at the beginning of their unemployment spells satisfied each of the following criteria: (i) age 50 or older; (ii) a continuous work history (780 employment weeks during the last 25 years prior to the current unemployment spell); (iii) residing in one of 28 selected labor market districts for at least 6 months prior to the claim; and (iv) commenced a new unemployment spell after June 1988 or had a spell in progress in June 1988. As LLZ remark, the treatment is “massive”—treated workers received an *extra three years* of covered unemployment with unchanged benefit levels.

The data are drawn from the universe of UI spells in Austria from 1980 through 2010. LLZ compute continuous work history for the prior 25 years for each individual at any point in time in order to determine eligibility status for REBP. They also use social security data to compute wages before and after each unemployment spell, as well as the total duration

of non-employment after the end of an employment spell. There is some fuzziness to this calculation.

The strategy is diff-in-diff. LLZ contrast eligible versus almost-eligible workers in treated (REBP) and non-treated (non-REBP) counties. To control for the potential endogeneity bias in the choice of REBP counties, they completely remove the steel sector from the analysis, eliminating all individuals who were employed in the steel sector immediately prior to becoming unemployed as well as unemployed whose subsequent employer is in the steel industry. The share of the steel sector in total employment is never larger than 15% in REBP counties, so this still leaves a very large sample. Nevertheless, it's not clear whether this approach mitigates or amplifies the endogeneity problem.

There are two key sets of results:

1. Employment effects: see slides.
2. Wage effects: This is a complex issue that is nicely expositied in Schmeider, Von Wachter and Bender (2014). An exogenous increase in a worker's unemployment duration could have two effects on re-employment wages. It may shift the wage offer distribution; we'd generally expect this effect to be negative. It may also change the worker's reservation wage, which could have a positive or negative effect on reemployment wages. It's likely that the first channel will dominate, but that's not a given. LLZ follow Schmeider et al.'s proposed workaround for this confound, which is to estimate the impact of extended unemployment benefits on reemployment earnings conditional on spell duration.