

Dissertation Awards

11-11-2020

Jobs Loss: Consequences and Labor Market Policy

Jonas Cederlöf
Uppsala University

Job Loss: Consequences and Labor Market Policy

Jonas Cederlöf

In March 2020, the U.S. economy saw an unprecedented increase in the number of layoffs and discharges due to the outbreak of the COVID-19 virus. Over 11 million American workers were displaced—a 670 percent increase compared to the same month the previous year. The unemployment rate rose quickly and was recorded at an all-time high of 14.7 percent in April 2020. Not only is displacement a traumatic event in and by itself, but it also has been shown to lead to significant and even permanent losses for workers in terms of future earnings, employment, wages, health, and mortality (Sullivan and von Wachter 2009). Furthermore, the large and potentially permanent negative consequences from job loss obviously puts a massive strain on public expenditures in terms of providing unemployment benefits, health care, and social assistance.

This thesis consists of four self-contained essays devoted to the topic of job loss, its consequences for individual workers, and how related labor market policies could be used and affect subsequent labor market outcomes for workers. The first essay studies the short- and long-run consequences of job loss for individual workers and explores why and under what circumstances the cost of displacement are most persistent. In the second essay, which is joint work with Peter Fredriksson, Arash Nekoei, and David Seim, we examine the effects of advance notice of job loss and how such notice affects the labor market prospects for workers. The third essay, coauthored with Martin Söderström and Johan Vikström, empirically analyzes the importance of caseworkers at the Public Employment Service. Here we examine not only caseworker value-added but also which features of a caseworker are important for job seeker outcomes and to what extent job seeker–caseworker matching matter. The fourth and final essay of this thesis examines a 30-week unemployment benefit extension and studies how the extension affects both benefit and unemployment duration, as well as the hazard to employment.

Essay 1

Saved by Seniority—Effects of Displacement for Workers at the Margin of Layoff

A large literature documents that displaced workers suffer significant and even permanent losses in terms of their future earnings, employment, and wages. Since the seminal study by Jacobson, Lalonde, and Sullivan (1993), the literature has relied on comparisons of displaced vis-à-vis nondisplaced

high-tenured workers, using mass layoffs or plant closures as an exogenous source of variation. However, most job loss occurs because of more marginal adjustments to employment. Strikingly, between 2003 and 2012, only about 8 percent of all involuntary separations in the United States were due to mass layoffs, many of which also involved lower-tenured workers.¹ Meanwhile, little is known to what extent the results from mass layoffs generalize to less drastic and more regular adjustments to employment, or if extraordinary events render extraordinary consequences for workers?

The first essay of this thesis studies the short- and long-run consequences of job loss for workers and explores why and under what circumstances the cost of displacement is most persistent. This is done using a novel source of identification, exploiting the use of a last-in-first-out (LIFO) rule used in layoffs in Sweden. The LIFO rule mandates that workers should be laid off in inverse order of seniority, whereby more recent hires ought to be let go before workers with higher tenure. Using detailed matched employer-employee data, containing information on job start and end dates, I rank workers according to their relative seniority (tenure) within an establishment which, by the LIFO rule, renders variation in the probability of displacement. Combining these data with a unique individual register dataset containing all layoff notifications involving at least five workers during 2005–2015, I identify occupation-specific cutoffs in downsizing establishments where the probability of displacement jumps discontinuously. This lends itself to a (fuzzy) regression discontinuity (RD) design, where the discontinuities enable me to go beyond standard mass layoff estimates and study earnings, employment, and wage losses for a broader and more representative population of workers who are laid off because of less drastic and more regular adjustments to employment. Exploiting the large heterogeneity across layoffs, I also pool different thresholds and separately estimate earnings losses and its persistence by worker characteristics and the size of layoff to investigate the key drivers of permanent earnings losses.

The main finding is that both the composition of workers and the size of the layoff have important consequences for how workers are affected by job loss, particularly in the long run. My estimates indicate that displaced workers on average suffer initial earnings losses of about 38 percent compared to their nondisplaced coworkers. While not being fully comparable, the size of these initial losses is close to what has been observed for displaced workers in the United States who are laid off during recessions (Davis and von Wachter 2011). As time progresses, however, the earnings gap between displaced and nondisplaced workers shrinks and is fully closed seven years after displacement. Crucially, this is not driven by the nondisplaced workers getting laid off at a later point in time. I also decompose cumulative earnings losses into different margins of adjustment and show that these losses

are primarily driven by lower wages and less employment, whereas the hours responses are of lesser importance.

My finding that earnings losses are transitory rather than persistent stands in contrast to the previous literature, which finds long-run earnings losses ranging between 10 and 20 percent of previous earnings. I make use of the large heterogeneity across layoffs to understand the main drivers of long-run earnings losses. I begin by estimating earnings losses of displaced workers using mass layoffs following the standard event study approach. I find large and highly persistent effects of displacement, thus ruling out that the transitory pattern observed in the RD analysis is context or time specific. I then proceed by producing separate RD estimates for each layoff. I correlate the short- and the long-run losses with characteristics of the workers, occupation, and establishment involved in the layoff, as well as economic conditions at the time of notification. While I find that older workers are more negatively affected by job loss, the key driver of persistent earnings losses turns out to be the relative size of the layoff. In fact, significant persistence can only be found among establishments executing mass layoffs, that is, those that displace more than 30 percent of their workforces. This pattern remains even when controlling for worker characteristics and economic conditions. Going further, I exploit the fact that there is variation in the size of the layoff relative to the local labor market, holding constant the layoff size in relation to the establishment. These estimates indicate that the key determinant of persistent earnings losses is the size of the layoff in relation to the local labor market, suggesting that negative spillovers and labor congestion play an important role for workers' future labor market outcomes.

To my knowledge, this is the first study to exploit seniority rules as an exogenous source of variation to involuntary job loss. Relative to the previous literature estimating earnings losses upon displacement, the findings shed new light on the question of how workers are affected by job loss and whether displacement creates lasting scars or merely temporary blemishes. Whereas previous literature suggests the former, my results indicate that while this may be true for high-tenured workers displaced under a mass layoff, earnings losses are temporary and transitory when considering more regular and less drastic adjustments to employment. These results may have important implications for both targeting and shaping of public policy.

The findings in this paper also speak to the theoretical literature explaining the observed earnings losses of displaced workers in models featuring search frictions, unemployment fluctuations, and job ladders (see, e.g., Ljungqvist and Sargent [1998]; Davis and von Wachter [2011]; Krolikowski [2017]; Kuhn and Jung [2019]). Many of the standard models of the labor market have trouble generating the magnitude and persistence of empirically observed losses or disagree on its sources (Davis and von Wachter 2011; Carrington and Fallick 2017). My results suggest that the way for a mecha-

nism generating long-run persistence at the individual level has been overemphasized in the theoretical literature.

Essay 2

How Does Advance Layoff Notice Affect the Labor Market Prospects for Workers?

(with Peter Fredriksson, Arash Nekoei, and David Seim)

Layoff rules are sometimes criticized for hampering the speed of adjustment after adverse shocks and creating inefficiencies in the allocation of resources. However, layoff rules also provide insurance for workers in part by mandating that firms share information on future reductions in labor demand, and they may force firms to share the costs associated with layoff.

The second essay in this thesis examines how advance layoff notice affects the labor market prospects for workers. More precisely, we utilize quasi-random variation in the length of notice periods to estimate the effects of advance notice on exposure to nonemployment, job mobility, subsequent wages, and earnings. The quasi-random variation comes from collective bargaining agreements, which stipulate that individuals above a certain age get longer notice periods. We estimate the causal effects of longer advance notice in a regression discontinuity design.

Employment protection legislation in most countries features advance layoff notice. The length of the notice period typically increases with tenure and tends to be longer for white-collar workers than for blue-collar workers. Notice periods are also longer in Northern and Continental Europe than in Anglo-Saxon countries (Organisation for Economic Co-operation and Development [OECD] 2013). In the United States, labor laws do not feature notice periods in the case of individual dismissals, but in the case of plant closures and mass layoffs, legislation sometimes mandates that workers are given two months of notice.

Contracts may also provide employment protection and advance notice. In the United States, for example, collective bargaining agreements often include notice periods and severance pay. Such agreements provide workers with extra protection relative to the law. In the Swedish context, the law provides a set of default rules, and the provisos in the collective agreements provide workers with additional protection.

Previous empirical analyses of employment protection originate from the seminal work of Lazear (1990). Using a panel of OECD countries, Lazear (1990) finds that severance pay increases unemployment. In a different panel of countries, Heckman and Pagés (2004) document a negative association between job security provisions and employment rates. Another strand of the literature exploits policy reforms to obtain quasi-experimental effects of employment pro-

tection on various outcomes. Kugler and Pica (2008) show that mandated severance pay reduces employment in Italy, while Autor, Kerr, and Kugler (2007) find similar effects on employment from increased dismissal costs.

We contribute to this literature along two margins. First, the employment outcomes in previous studies originate from both hiring and separations margins. By focusing on separations, we specifically investigate the role of advance notice for a laid-off worker, thereby uncovering the insurance role of such employer protection legislation. Second, our setting provides the ideal testing ground from an identification perspective. We exploit exogenous variation in the length of advance notice within establishments and displacement events, across individuals. This permits a compelling analysis of the causal effects of advance notice.

We also contribute to an older literature, which investigates the role of advance notice for laid-off workers using the Displaced Worker Survey (see Ruhm [1992, 1994] and the survey in Addison and Blackburn [1994]). That literature exploits cross-sectional variation in advance notice periods instigated by the Worker Assistance and Retraining Notification Act and finds that joblessness falls upon notification of job loss. Our paper breaks new ground by providing quasi-experimental variation applied to administrative data on long-term outcomes.

Advance notice policies are clearly related to severance pay and unemployment insurance policies. Pissarides (2001) analyzes a model where firms can offer severance pay and advance notice as part of an optimal contract, but unemployment insurance (UI) is exogenously set by the government. His analysis implies that optimal contracts are more likely to involve severance pay or advance notice when UI replacement rates are low. In his model, advance notice plays a role over and above severance payments if and only if the insurance properties are sufficiently superior.²

Our empirical work relates to the model in Pissarides (2001) in the sense that we think of (voluntary) severance pay as an outcome of other policy parameters, in our case mandatory advance notice. We examine whether firms are willing to make—and whether workers accept—an up-front severance payment in order to avoid the notice period. To our knowledge, this is the first paper that provides such an analysis.

The major results of the paper are the following. Longer notice periods cause prolonged periods of adjustment. For workers who are eligible for a longer notification period (the treatment group), the probability of remaining in the displacing firm increases during the first two years after notification, and the probability of moving to another firm falls during the same time period. As a result of an extension of the notice period, workers are less exposed to unemployment and nonemployment and spend less time outside the labor force. After two years, all employment responses have subsided and there are no differential effects on employment outcomes

for the treatment and control groups. We also show that the treatment group experiences smaller wage losses when finding a new job than the control group: wages in new jobs are 3 percent higher for the treatment group than for the control group. Moreover, firms make severance payments to workers in order to avoid the notice period: the extra payment accruing to the treatment group amounts to almost 60 percent of the monthly wage. Finally, we show that workers who are eligible for higher UI get lower severance payments. This is consistent with the view that they are more willing to accept lower severance pay since they have a better outside option should they leave the firm for unemployment.

The extra severance payment is part of the overall effect on earnings for treated workers. When we decompose the earnings effect accruing to workers during the first two years after notification, we find that around half of the earnings effects has to do with less exposure to nonemployment and a third is due to the increase in severance pay. The wage effect amounts to a fifth of the overall earnings effect. Over time, the wage effect becomes less important as individuals with short notification periods—who initially find a lower-quality job—move on to better-paying jobs at a greater rate than individuals with long notification periods.

Essay 3

What Makes a Good Caseworker?

(with Martin Söderström and Johan Vikström)

Countries around the world use job search assistance, monitoring schemes, and labor market programs to try to bring unemployed workers back to work. By now, there is extensive evidence on these policies (see, e.g., Card, Kluve, and Weber [2010, 2017]). However, much less is known about the caseworkers who provide the job search assistance, carry out the monitoring, and assign job seekers to programs (see the literature review below and McCall, Smith, and Wunsch [2016]). This is unfortunate, since a comprehensive picture of labor market policies requires that we understand the role of the human resources used to provide the services. It is, for example, important to know who becomes a caseworker, why some caseworkers perform better than others, and for whom caseworkers matter the most. While these are important questions, the evidence is scarce for two important reasons. First, in most cases, there is nonrandom sorting of job seekers to caseworkers, often because the most productive caseworkers are assigned the most disadvantaged job seekers. Second, high-quality data on caseworkers is often lacking. Usually, data do not link caseworkers to job seekers, and in the rare cases when such information is available, typically little is known about the caseworkers.

In the third essay of this thesis, we address both of these issues. First, we break the caseworker–job seeker sorting by

exploiting that many local employment offices in Sweden use date-of-birth-rules to allocate job seekers to caseworkers, creating as-if random allocation. Second, we have access to uniquely fine-grained information on caseworkers, such as labor market experiences and cognitive ability, and we can link job seekers to caseworkers. The quasi-random allocation and the fine-grained data allow us to provide new and credible evidence on the importance of caseworkers and what makes a good caseworker.

In Sweden, employment services are provided by caseworkers at local public employment offices. These offices have extensive discretion to design the rules for allocating job seekers to caseworkers. It turns out that many offices use job seekers' date of birth (day in the month) to allocate them to caseworkers. We exploit that the day in the month you are born (1st to 31st) is uncorrelated with individual characteristics. Thus, if job seekers are allocated to caseworkers using a date-of-birth rule, this creates as-if random allocation, since all caseworkers within a local office will have job seekers with similar observed and unobserved characteristics.

Our analysis shows that even within the offices that use a date-of-birth allocation rule, some job seekers are not allocated using the rule. Some offices make exemptions for special groups, such as youths, disabled workers, or immigrants, and allocate these groups to caseworkers who are believed to be able to provide the best support to them. Since these exemptions may introduce sorting, we use an IV-framework exploiting that we can identify the caseworker each job seeker would have had if they had been allocated using the date-of-birth rule. This rule-predicted caseworker is then used as an instrument for the caseworker assigned to the job seeker. This identification strategy adds to the existing literature on caseworkers, which mainly includes studies based on conditional independence assumptions, assuming that the allocation of job seekers to caseworkers is random conditional on observed job seeker characteristics (see, e.g., Lechner and Smith [2007]; Behncke, Frölich, and Lechner [2010a,b]; Arni, van den Berg, and Lalive [2017]; Arni and Schiprowski [2019]). One recent exception, however, is Schiprowski (2020), who exploits unplanned absences to study the effects of a meeting with a caseworker and to study productivity differences across caseworkers. Our paper uses detailed data on caseworkers and the date-of-birth allocation to provide more comprehensive evidence on caseworker performance.

Our unique administrative data on caseworkers include rich measures of labor market history, such as information on previous occupations and personal experience with unemployment. For most male caseworkers, we also have information on cognitive and noncognitive ability from enlistment tests. Staff records provide information on experience (tenure at the public employment service) and wages for each caseworker. To this, we add information on demographics such as gender, level and type of education, and country of origin.

Using the fine-grained caseworker data, we initially study who becomes a caseworker. One conclusion is that caseworkers in Sweden are a heterogeneous group that includes former blue-collar workers; individuals with university degrees in social work, business economics, and human relations; and both natives and nonnatives. Interestingly, caseworkers have, on average, lower cognitive skills, substantially more experience of unemployment, but similar noncognitive skills as other public sector employees with similar types of occupations.

We then study caseworker performance in three different parts. In the first part, we analyze how different observed caseworker characteristics are related to caseworker performance as measured by the reemployment rate among their job seekers. Even though we are able to study a heterogeneous group of caseworkers, few observed caseworker characteristics predict caseworker performance. The most important characteristic is the gender of the caseworker: job seekers with female caseworkers have 3.1 percent shorter unemployment durations than those with a male caseworker. There is also some evidence that caseworkers with higher wages perform better, but this may, of course, reflect both that high-performing caseworkers are rewarded with higher wages and that higher wages motivate caseworkers to perform better. However, many other caseworker characteristics, such as type of education, level of education, experience from previous occupations, and personal experience with unemployment are not related to caseworker performance. These results are consistent with results from the teacher literature, which finds little evidence of a relationship between teacher quality and observed teacher characteristics (Rockoff 2004; Rivkin, Hanushek, and Kain 2005; Rockoff et al. 2011). Moreover, caseworkers with higher cognitive ability do not perform better than low-ability caseworkers. There is, however, some suggestive evidence indicating that caseworkers with higher noncognitive ability may have a positive impact on job seekers' job-finding rates early on in the unemployment spell.

Based on the actions taken by the caseworkers, we also examine caseworker traits. Inspired by Arni, van den Berg, and Lalive (2017), we define "supportive" caseworkers as those who more often use supportive actions, such as sending their job seekers to labor market training, whereas "restrictive" caseworkers are those who more often use restrictive policies such as workfare. Furthermore, we define "active" caseworkers as those who more frequently meet with their job seekers. Our results show that active caseworkers perform better than other caseworkers. This adds to the rather few existing studies: Arni, van den Berg, and Lalive (2017) find that caseworkers who emphasize support have better outcomes, while Behncke, Frölich, and Lechner (2010b) show that tougher caseworkers are more successful than supportive ones.

The second part of the paper examines caseworker–job seeker matching and focuses on caseworker–job seeker similarity, since it has been argued that sharing the same social background can enhance communication and trust. This is also what we find. Being assigned a caseworker with the same gender leads to a higher job-finding rate, but it is not the case that immigrant caseworkers provide better support to immigrants. Using our fine-grained data, we are also the first to show that matching job seekers to caseworkers with similar labor market experiences and/or similar educational background leads to substantially shorter unemployment durations. Besides improved communication and trust, this may also reflect that experience from working in the same sector as the job seeker enables caseworkers to understand the individual-specific labor market opportunities, and that caseworkers can use their social networks to help job seekers with similar labor market experiences.

The third and final part of the paper examines the overall importance of caseworkers by estimating caseworker fixed effects. This takes into account differences due to both observed and unobserved caseworker characteristics. The overall conclusion is that there are economically important differences between caseworkers. A one standard deviation increase in the distribution of caseworker fixed effects not only increases the job-finding rate among the job seekers by around 0.1 standard deviation but also renders about 5 percent higher earnings after three years. This confirms that caseworkers indeed can affect how quickly job seekers get back to work, a result consistent with Schiprowski (2020). It is also in line with the results from other economic contexts; a large literature has documented substantial differences in teacher quality (Rockoff 2004; Rivkin, Hanushek, and Kain 2005; Rothstein 2010; Chetty, Friedman, and Rockoff 2014a,b), and several studies have shown that managers matter for firm policies and firm performance (Bertrand and Schoar 2002; Bloom et al. 2014; Lazear, Shaw, and Stanton 2015).

Essay 4

Extended Unemployment Benefits and the Hazard to Employment

How does the generosity of UI affect job search behavior? While providing a safety net for unexpected job loss, the provision of UI creates disincentives for job search by lowering the alternative cost to working. The question of how benefit levels and their overall generosity affects time in, and the hazard out of, unemployment has a long tradition in labor economics and has been subject to extensive research. The “spike” in the hazard rate out of unemployment coinciding with UI exhaustion is a widely established empirical result since the seminal work by Katz and Meyer

(1990a,b). This result generally has been attributed to shirking behavior among job seekers, holding off finding a new job until approaching benefit exhaustion. However, later work by Card, Chetty, and Weber (2007) challenges this view by attributing the lion’s share of such spikes to flight out of the labor force. They argue that “spikes are generally smaller when the spell length is measured by the time to next job than when it is defined by the time spent on the unemployment system” (p. 1). Hence, “. . . the size of the spike in re-employment rates at exhaustion in the current U.S. labor market (and many other labor markets) remains an open question. Further work on estimating these hazards using administrative measures of time to next job would be particularly valuable” (p. 16). Indeed, if benefit exhaustion renders job seekers to leave the labor force, the expected cost of extending UI benefits could be exaggerated if transition to work is higher from unemployment than nonemployment.

This paper contributes to the debate about the timing of reemployment and UI exhaustion while adding to the large literature on the effects of UI on job search behavior. In particular, I examine the effect on unemployment duration, and exit to employment, of an exogenous 30-week UI benefit extension in Sweden. For identification, I take advantage of a feature in the Swedish UI system that entitles individuals with a child below the age of 18 to 90 weeks of unemployment benefits instead of the statutory 60 weeks. As assignment to the extended UI benefit is determined by the age of a job seekers’ youngest child at the time of regular UI exhaustion (60 weeks), I exploit the quasi-experimental variation generated around the age threshold using a regression discontinuity design. This allows me to estimate the causal effect of increasing potential duration of UI on actual benefit duration, unemployment duration, and hazard to employment. Further, I allow the effects to vary with duration on UI and in unemployment to test whether job seekers time employment to benefit extension.

The main findings are threefold. First, while the increase in potential duration on UI increases actual duration on UI by about 2.7 weeks on average, I find no evidence of it prolonging duration in registered unemployment or negatively affecting the hazard to employment. This suggests that the 30-week benefit extension did not prolong average unemployment duration, as job seekers were unemployed just as long on average but with a somewhat higher replacement rate. The absence of negative effects on unemployment duration and future employment is believed to be driven by job seekers’ access to fairly generous post-UI programs, which weakens the disincentive effects of the benefit extension. Second, being eligible to 30 additional weeks of UI does not appear to have affected job search behavior prior to the actual extension period. That is, I find no evidence of job seekers lowering their search efforts due to the anticipation of extended benefits. Third, I find distinct spikes in the exit from UI at benefit exhaustion, but no such spikes are present

in the hazard to employment. This therefore speaks in favor of the interpretation made in Card, Chetty, and Weber (2007).

Notes

1. Calculations are based on data from the Bureau of Labor Statistics by combining data from the Mass Layoff Statistics program (which ended in March 2013) with the Job Openings and Labor Turnover Survey. That survey reports the total number of layoffs and discharges, which is made up of all involuntary separations initiated by the employer. Both these data sources can be accessed at <http://www.bls.gov>.
2. A severance payment is a pure transfer from the worker to the firm and thus does not affect the private surplus of the match. By contrast, advance notice may have a negative effect on the private surplus if the match is kept alive after it has become unproductive. These costs have two components: first, there may be variable costs associated with keeping an unproductive job alive; second, for the worker-firm pair, unemployment income is a pure subsidy that is forgone by keeping the match alive. In the framework of Pissarides (2001), the insurance value of advance notice must thus outweigh these losses relative to the severance pay.

References

- Addison, John T., and McKinley Blackburn. 1994. "Policy Watch: The Worker Adjustment and Retraining Notification Act." *Journal of Economic Perspectives* 8(1): 181–190.
- Arni, Patrick, and Amelie Schiprowski. 2019. "Job Search Requirements, Effort Provision and Labor Market Outcomes." *Journal of Public Economics* 169: 65–88.
- Arni, Patrick, Gerard J. van den Berg, and Rafael Lalive. 2017. "Treatment versus Regime Effects of Carrots and Sticks." IFAU Working Paper No. 2017:25. Uppsala, Sweden: Institute for Evaluation of Labour Market and Education Policy.
- Autor, David H., William R. Kerr, and Adriana D. Kugler. 2007. "Does Employment Protection Reduce Productivity? Evidence from US States." *Economic Journal* 117(521): 189–217.
- Behncke, Stefanie, Markus Frölich, and Michael Lechner. 2010a. "A Caseworker Like Me—Does the Similarity between the Unemployed and Their Caseworkers Increase Job Placements?" *Economic Journal* 120(549): 1430–1459.
- . 2010b. "Unemployed and Their Caseworkers: Should They Be Friends or Foes?" *Journal of the Royal Statistical Society: Series A* 173(1): 67–92.
- Bertrand, Marianne, and Antoinette Schoar. 2002. "Managing with Style: The Effect of Managers on Corporate Policy." *Quarterly Journal of Economics* 118(4): 1169–1208.
- Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. 2014. "Does Management Matter? Evidence from India." *Quarterly Journal of Economics* 128(1): 1–51.
- Card, David, Raj Chetty, and Andrea Weber. 2007. "The Spike at Benefit Exhaustion: Leaving the Unemployment System or Starting a New Job?" *American Economic Review* 97(2): 113–118.
- Card, David, Jochen Kluve, and Andrea Weber. 2010. "Active Labour Market Policy Evaluations: A Meta-Analysis." *Economic Journal* 120(548): 452–477.
- . 2017. "What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations." *Journal of the European Economic Association* 16(3): 894–931.
- Carrington, William J., and Bruce Fallick. 2017. "Why Do Earnings Fall with Job Displacement?" *Industrial Relations* 56(4): 688–722.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff. 2014a. "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates." *American Economic Review* 104(9): 2593–2632.
- . 2014b. "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes." *American Economic Review* 104(9): 2633–2679.
- Davis, Steven J., and Till von Wachter. 2011. "Recessions and the Costs of Job Loss." *Brookings Papers on Economic Activity* Fall(1993): 1–72.
- Heckman, James J., and Carmen Pagés. 2004. *Law and Employment: Lessons from Latin America and the Caribbean*. Chicago: University of Chicago Press.
- Jacobson, Louis S., Robert J. Lalonde, and Daniel G. Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review* 83(4): 685–709.
- Katz, Lawrence F., and Bruce D. Meyer. 1990a. "The Impact of the Potential Duration of Unemployment Benefits on the Duration of Unemployment." *Journal of Public Economics* 41(1): 45–72.
- . 1990b. "Unemployment Insurance, Recall Expectations, and Unemployment Outcomes." *Quarterly Journal of Economics* 105(4): 973–1002.
- Krolikowski, Pawel. 2017. "Job Ladders and Earnings of Displaced Workers." *American Economic Journal: Macroeconomics* 9(2): 1–31.
- Kugler, Adriana D., and Giovanni Pica. 2008. "Effects of Employment Protection on Worker and Job Flows: Evidence from the 1990 Italian Reform." *Labour Economics* 15(1): 78–95.
- Kuhn, Moritz, and Philip Jung. 2019. "Earnings Losses and Labor." *Journal of the European Economic Association* 17(3): 678–724.
- Lazear, Edward P. 1990. "Job Security Provisions and Employment." *Quarterly Journal of Economics* 105(3): 699–726.
- Lazear, Edward P., Kathryn L. Shaw, and Christopher T. Stanton. 2015. "The Value of Bosses." *Journal of Labor Economics* 33(4): 823–681.
- Lechner, Michael, and Jeffrey A. Smith. 2007. "What Is the Value Added by Caseworkers?" *Labour Economics* 14(2): 135–151.
- Ljungqvist, Lars, and Thomas J. Sargent. 1998. "The European Unemployment Dilemma." *Journal of Political Economy* 106(3): 514–550.
- McCall, Brian, Jeffrey Andrew Smith, and Conny Wunsch. 2016. "Government-Sponsored Vocational Education for Adults." In *Handbook of the Economics of Education*, Eric A. Hanushek, Stephen Machin, and Ludger Woessmann, eds. 479–652. Amsterdam: Elsevier.
- Organisation for Economic Co-operation and Development. 2013. *Employment Outlook*. Paris: OECD.
- Pissarides, Christopher A. 2001. "Employment Protection." *Labour Economics* 8(2): 131–159.
- Rivkin, Steven G., Eric A. Hanushek, and John F. Kain. 2005. "Teachers, Schools, and Academic Achievement." *Econometrica* 73(2): 417–458.
- Rockoff, Jonah E. 2004. "The Impact of Individual Teachers on Student Achievement: Evidence from Panel Data." *American Economic Review* 94(2): 247–252.
- Rockoff, Jonah E., Brian A. Jacob, Thomas J. Kane, and Douglas O. Staiger. 2011. "Can You Recognize an Effective Teacher When

- You Recruit One?" *Education Finance and Policy* 6(1): 43–74.
- Rothstein, Jesse. 2010. "Teacher Quality in Educational Production: Tracking, Decay, and Student Achievement." *Quarterly Journal of Economics* 125(1): 175–214.
- Ruhm, Christopher J. 1992. "Advance Notice and Postdisplacement Joblessness." *Journal of Labor Economics* 10(1): 1–32.
- . 1994. "Advance Notice, Job Search, and Postdisplacement Earnings." *Journal of Labor Economics* 12(1): 1–28.
- Schiprowski, Amelie. 2020. "The Role of Caseworkers in Unemployment Insurance: Evidence from Unplanned Absences." *Journal of Labor Economics* 38(4): 1189–1225.
- Sullivan, Daniel, and Till von Wachter. 2009. "Job Displacement and Mortality: An Analysis Using Administrative Data." *Quarterly Journal of Economics* 124(3): 1265–1306.