

# Data Improvement and Labor Economics

Kevin F. Hallock, *Cornell University and  
National Bureau of Economic Research*

The expansion of available data for research has transformed empirical labor economics over the past generation. This paper briefly highlights some of the changes and describes a few examples of papers that illustrate the advances. It also documents the changing ways data have been used in the *Journal of Labor Economics* over the past 30 years, including a trend toward a higher fraction of papers using any data and, among those papers using any data, a higher fraction using nonpublic data, a higher fraction using international data, and more frequent use of multiple data sources. Finally, this paper describes work that came out of the recent Princeton Data Improvement Initiative—a program that considers and furthers improved data collection.

Mounds of papers written in empirical labor economics in recent decades have used existing, publicly available data. Since the early 1970s, the advent of more modern computing and new techniques in applied econometrics (many of which were created precisely to solve problems in labor economics) have triggered an avalanche of new papers. But, along with advances in computing power and econometric techniques, newer and more diverse data have been created, contributing significantly to the probing of interesting questions in labor economics. This paper briefly highlights some of

I am grateful to Linda Barrington, Fran Blau, David Card, Robert Hutchens, George Jakubson, Alan Krueger, Francesca Molinari, Robert Smith, and Chris Taber for comments and suggestions. Weishi (Grace) Gu and Xin Jin provided excellent research assistance. The Princeton Data Improvement Initiative was led by Alan Krueger, Ed Freeland, and Bill Barron. Maggie Newman is, of course, awesome. Contact the author at [hallock@cornell.edu](mailto:hallock@cornell.edu).

[*Journal of Labor Economics*, 2013, vol. 31, no. 2, pt. 2]  
© 2013 by The University of Chicago. All rights reserved.  
0734-306X/2013/31S1-0012\$10.00

the changes and describes a few examples of papers that illustrate the data advances. It also documents the changing ways data have been used in the *Journal of Labor Economics* over the past 30 years. The increase in the diversity of new data on modern labor economics is striking. Finally, this paper describes work that came out of the recent Princeton Data Improvement Initiative—a program that considers and furthers improved data collection.

The rest of this paper is organized as follows. In Section I, I describe certain sorts of data sources and their relationship to empirical labor economics. This is by no means an exhaustive list, and I apologize to the authors of papers I have left out. Rather, this is a simple overview of some examples of the diversity of data sources and changes over time in labor economics research. In Section II, I describe changes in the use of data in the papers published in the *Journal of Labor Economics* since the journal began in 1983 and document these trends. For example, a shrinking fraction of the papers in the journal have incorporated no data. Further, among those using any kind of data, an increasing fraction have used non-public data, multiple data sources, and data from countries outside the United States. In Section III, I describe a set of papers using novel data sources that came out of the Princeton Data Improvement Initiative (PDII), an effort to advance the way we think about data in labor economics and improve the data employed in labor economics research. Section IV offers some concluding comments.

## I. Data and Empirical Labor Economics

Clearly a large number of papers in empirical labor economics in the past decades have used existing data sources such as the Current Population Survey (CPS), the Census of Population, the Survey of Income and Program Participation (SIPP), the American Community Survey (ACS), the National Longitudinal Surveys (NLS), and the Panel Study of Income Dynamics (PSID). In fact, of the empirical papers published in the *Journal of Labor Economics* in the last 30 years, just under one-half have used data from a public source (see more on this in Sec. II). Although we have seen expansions in the number of novel new data sources recently, it isn't as though changes happened overnight. Nor is it the case that earlier research in modern labor economics used only one modest sort of data, even when exploring early questions. Take, for example, Ashenfelter's (1972) study of racial discrimination and trade unionism (although the focus was on the 1967 survey of economic opportunity, that is not the only data set used) or Freeman's (1976) study of the cobweb model and starting compensation for engineers.

Given that such a large fraction of the papers in empirical labor economics continue to use public use data, the Princeton Data Improvement Initiative was created to improve public use data by doing the type of re-

search that statistical agencies could do but do not have the resources to do; namely, develop new questions that could be added to public use surveys, extract more information from existing surveys, and assess the reliability of public use data. Public use data are still the mainstay of empirical labor economics, and it would be great if we could improve public use data. This is precisely why the Princeton Data Improvement Initiative started.

But first I will highlight examples of the richness with which labor economists have used data and begin to show changes over time. I will briefly describe seven areas of interest (although, clearly there are many, many more): (i) adding institutional, historical, or contextual detail to existing large-scale surveys, (ii) using specifically collected experimental data, (iii) using personnel or firm-specific data, (iv) using employee-employer matched data, (v) using special survey and nonexperimental data, (vi) using administrative data, and (vii) taking advantage of advances in panel data. It may well be the case that papers that are in none of these seven categories but only use one of the large, publicly available, surveys account for the majority of empirical papers in the past generation or more. But this introduction is about data changes and data improvement, so I will focus for a bit on these areas.

Although there have been many interesting papers that have used existing large-scale surveys, a new generation of papers began to add institutional or historical context to those sorts of data sets, including some noteworthy ones in the 1990s, such as, for example, Card and Krueger (1992) and Angrist and Krueger (1991).

Card and Krueger (1992) not only use earnings data from the 1980 census of population but they also consider school quality in a historical context. Using these data, they find that men who were educated in states with higher quality schools have higher returns to additional years of education. They also find higher returns to education for those who were from states with larger fractions of female teachers and states with better educated teachers.

Another paper of this type is Angrist and Krueger's (1991) study of compulsory schooling and earnings. This paper uses an instrument for education (quarter of birth) to estimate the rate of return to schooling and shows it is close to the least-squares estimate. The novelty of the data highlights another historical or institutional feature. Here the authors show that the season of birth is related to school attainment since school start age and compulsory schooling laws vary.

A second innovation in data used by labor economists could be called specifically collected experimental data. In this I highlight just three: the negative income tax experiments, Job Training Partnership Act (JTPA) training programs, and the Moving to Opportunity (MTO) evaluations.

Ashenfelter's (1983) paper is an interesting study of determining participation in income-tested social programs. It is an excellent example of

the use of large specifically collected experimental data in labor economics, in this case the Seattle and Denver Income Maintenance Experiments. The paper describes how estimates of the expected number of participants in means-tested social programs can be made with only knowledge of the income distribution. However, because of incentive effects and the fact that some potential participants will not apply for benefits, these simple estimates can be biased. This paper uses the negative income tax experimental data to document that such bias was small in this case.

Of course, there is also a large literature on training programs, some of which is based on experimental data. For example, Heckman and Smith (2000) study a great deal of evidence from the JTPA, an employment and training program in the United States which at one time had a budget of around \$1 billion per year. The data studied by Heckman and Smith (2000) come from experiments in 16 centers. Earlier, Ashenfelter (1978) used data from the Continuous Work History Sample (CWHHS), maintained at the Social Security Administration, to study training programs through a novel use of experimental methods.

In a paper that considers the effects of training programs, LaLonde (1988) uses data from myriad sources, including the National Supported Work Demonstration (NSW) of the Manpower Demonstration Research Corporation (MDRC), along with data from the PSID and CPS. He argues that many of the results using nonexperimental data cannot be replicated with experimental data. He notes, therefore, that researchers using nonexperimental data should use caution.

Neighborhood effects have also been studied with experimental data. Kling, Liebman, and Katz (2007) use experimental data from the MTO social experiments whereby participants were offered housing vouchers in a lottery conducted in five cities in the United States: Baltimore, Boston, Chicago, Los Angeles, and New York. These data were used in conjunction with administrative data on earnings and welfare benefits to study neighborhood effects. The authors find that 4–7 years after random assignment, families offered vouchers had lower poverty rates and lived in safer neighborhoods than those not offered vouchers. They also find no significant effects on adult economic self-sufficiency or on physical health; however, mental health benefits were large for adults and female youth. Measures of education, risky behavior, and health show benefits for female youths, while male youths did worse on those metrics.

A third innovation that has had important implications in empirical work is the use of personnel or firm-specific data. Medoff and Abraham (1980) discuss the relationships among experience, performance and earnings using personnel data from two firms. Later, Baker, Gibbs, and Holstrom (1994a) use 20 years of personnel records from a single firm to consider internal labor markets, “fast tracks,” ports of entry, exit rates, salaries, and a host of other issues. Baker, Gibbs, and Holmstrom (1994b) use the

same data to study the wage policy of the firm. Many papers have followed suit.

Lazear (2000) is a case study of a firm that switched from paying its windshield installers from salaries to “piece rates.” Such sources of data are more common in other social sciences but are being increasingly used in labor economics. Craig Olson and I (Hallock and Olson 2010) used data on stock option exercise behavior inside a particular firm to estimate the value of options to employees, relative to the Black-Scholes value. We used data from another firm in Hallock and Olson (2012) to examine the relative value employees place on certain forms of compensation. The firm set the total level of compensation for each employee and then let the employees choose their pay “mix” (i.e., relative weights on salary, at-risk bonus, and stock options). We find, among other results, that women were much less likely to choose the at-risk forms of pay (bonus and stock) than men.

Employer-employee matched data is a fourth area of new data in labor economics. The best-known study in this area, Abowd, Kramarz, and Morgolis (1999), uses data on over 1 million employees from more than 500,000 employers in France. The authors decompose total annual compensation into separate parts; those related to employee-measured characteristics, individual heterogeneity, firm heterogeneity, and residual variation. These sorts of data and similar data created in the United States and elsewhere have led to a new research “industry” to study a wide variety of questions in labor economics and related areas. Of course, related, is the important work on job creation and destruction such as that so carefully done by Davis, Haltiwanger, and Schuh (1998).

The fifth area of data used in modern labor economics is that coming from surveys and is nonexperimental but takes advantage of experimental situations. Well-known examples of these are Card and Krueger (1994) and Ashenfelter and Krueger (1994). Card and Krueger (1994) examine the 1992 change in New Jersey’s minimum wage by surveying over 400 fast-food restaurants along the border of New Jersey (where the minimum was raised) and Pennsylvania (where the minimum wage stayed the same). The authors found that the increase in the minimum wage did not lead to decreases in employment. There are also special archival studies done, such as that by Goldin and Rouse (2000), to study blind auditions at eight major symphony orchestras. They find evidence that blind auditions increased the proportion of women in major symphony orchestras.

Ashenfelter and Krueger (1994) survey sets of identical twins at the Annual Twinsburg Twins Festival in Twinsburg, Ohio. The paper aimed to compare the wages of genetically identical twins with different levels of schooling. A novelty of the paper was that they asked both twins questions about their own and their siblings’ education and that of their parents so they could solve the measurement error problem created by dif-

ferencing. The paper found that adjusting for ability bias and for measurement error led to larger empirical estimates of returns to schooling than estimated using standard cross-sectional data.

Of course there are also many sorts of administrative data (the sixth category of data I discuss here) that have been used to creative ends. Well-known studies providing examples of such include Angrist (1990), Dyke et al. (2006), and Card, Chetty, and Weber (2007).

Angrist (1990) estimates that earnings of white Vietnam veterans were about 15% lower than their nonveteran counterparts in the 1980s, long after the war ended. He uses Social Security earnings data and the randomly assigned risk of induction into the Vietnam draft to create estimates of veteran status on civilian earnings.

Administrative data on welfare recipients from Missouri and North Carolina are used in Dyke et al. (2006) to study welfare-to-work programs. They find that more intensive training is initially associated with larger earnings losses that are made up with longer-run earnings gains.

Administrative data from Austria are used in studying severance pay and unemployment insurance in Card et al. (2007). The paper explores the permanent income hypotheses and other models of behavior using labor market data on over 500,000 people. Among the findings are that lump sum severance (on the order of 60 days of pay) reduces job-finding rates by about 10%, increasing the length of time for unemployment insurance (UI) benefits from 20 to 30 weeks lowers job finding rates in the first 20 weeks by 5%–9%, and increases in search behavior created by the program have little effect on the quality of the subsequent job match. Another excellent use of administrative data is in Jacobson, LaLonde, and Sullivan (1993), who use data on displaced workers in Pennsylvania from 1974 to 1986 to estimate the impacts of job displacement on earnings.

The seventh category is advances in panel data. Mincer and Polachek (1974) use the 1967 National Longitudinal Survey of Work Experience (NLS) where they discuss the problems of “potential experience” and the beauty of data like the NLS that were the first to use longitudinal structure to consider the earnings of women in a way that nicely foreshadows Blau and Kahn (2013). Subsequent papers that make great use of related data include Altonji and Pierret (2001), who use the National Longitudinal Survey of Youth (NLSY) 1979 to study employer learning and statistical discrimination, and Farber and Gibbons (1996), who study learning and wage dynamics.

There are a few other newer areas or trends I will mention here. An additional category for data that may prove useful to labor economists is emerging—that of data that are being scraped from the web. An interesting example is Choi and Varian (2011) who use data from an index of real-time searches on Google to help make forecasts of economic measures such as unemployment claims. Another data set that is getting increased attention

is the Job Openings and Labor Turnover Survey (JOLTS) from the US Department of Labor. The sample size is approximately 16,000 business establishments. See, for example, Davis et al. (2010). Finally, one area that is not covered in detail in this paper is the dramatic shift away from time-series data in labor economics research toward micro data as documented in Stafford (1986) and Angrist and Krueger (1999).

## II. Changes in Data in Papers in the *Journal of Labor Economics* over Time

A data set on all papers published in the *Journal of Labor Economics* from the years 1983 to 2012 was created. Over that time 876 papers were published in the *Journal of Labor Economics*. Of those, 211 (24%) contained no data. The remaining 665 contain at least one data source. The first panel of figure 1 plots the fraction of papers in a given year that use any data at all. This fraction has steadily increased over time from just over 50% at the start of the journal to over 80% today.



The rest of the panels in figure 1 represent statistics from the 665 papers that contain at least one data source. The top right panel shows that, among papers that contain some data, the fraction that contain only publicly available data has declined from roughly three-quarters to just over 40%. This is consistent with the rise in internal personnel data, own-surveys, and other proprietary data over time.

The bottom left panel of the figure plots the fraction of papers that use more than one data source (e.g., both the CPS and the PSID). Finally, and perhaps most dramatic, is the fraction of papers published in the *Journal of Labor Economics* that use data from a country other than the United States (including those that also use US data). In the first 9 years of the journal, this fraction was under 20% (except for the year 1985, which included a special issue on “Trends in Women’s Work, Education and Family Building,” that included papers using data from around the world).


## III. The Princeton Data Improvement Initiative

The Princeton Data Improvement Initiative (PDII) began in 2007 and was funded by the Industrial Relations Section at Princeton University. The idea of doing something novel to explore empirical questions is obviously not new to the Industrial Relations Section at Princeton. Take, for example, the excellent report by Lester (1948) that explores company wage policies and considers new ways of gathering information to consider wages. The idea of the PDII was to bring together a set of researchers to improve data that are already being collected. The first meeting was on Amelia Island in Florida in 2007 and brought together labor economists, applied psychologists, directors, and other staff from the Census, the Bureau of Economic Analysis and the Bureau of Labor Statistics, and other

Fraction of Papers Using Any Data      Fraction using Only Publicly-Available Data



Fraction using Multiple Data Sources



Fraction using any Non-US Data

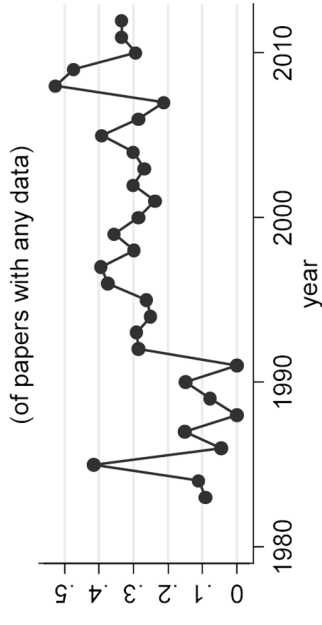


FIG. 1.—Data in papers in the *Journal of Labor Economics*, 1983–2012



areas of government, data consultants, and other experts. The second meeting was held at the Educational Testing Service in New Jersey in October of 2008. The papers in this special issue emerged from those meetings: Abraham et al. (2013), Autor and Handel (2013), Blau and Kahn (2013), Blinder and Krueger (2013), and Kleiner and Krueger (2013).

The PDII was created by Alan Krueger, Ed Freeland, and Bill Barron. Alan Krueger came up with the idea, Ed Freeland led the design of the survey, and Bill Barron led the outreach to the government statistical agencies. The survey was conducted by Westat between June 5 and July 20, 2008. The survey was administered to 2,513 individuals aged 18 or over who were in the labor force. The questions were patterned after the CPS, with a number of additional questions added for the purpose of the initiative. More details can be seen in Kleiner and Krueger (2013).

Starting as far back as Becker (1962) and Mincer (1962), economists have had an appreciation for measures of experience in earnings functions. One major problem has been the lack of good measures of experience, especially in cross-sectional data. The lack of good measures of true labor market experience is particularly troublesome for some demographic groups who are less well attached to the labor force. Obviously estimates such as “age – education – 6” as a proxy for years of experience works relatively better for those individuals who have a strong attachment to the labor force. As noted above, volumes of papers have been written that use cross-sectional data such as the Current Population Survey, the American Community Survey, and the Decennial Census of Population. But none of these has perfectly satisfactory measures of experience.

Using data from the PSID as well as the PDII, Blau and Kahn (2013) carefully explore the possibility of adding retrospective questions about actual work experience to data sets like the March CPS. The paper uses multiple strategies. In one part, the authors use the PSID’s measure of actual (not potential) experience and find that those work history variables are important in a number of contexts.

They also use the PSID to compare respondents’ memory of their work history to measures accumulated over time in the regular panel surveys. Obviously only the former can be done in cross-section surveys (like the CPS, Census, and ACS), so making this comparison to actual experience (collected in each wave for the most recent year) to retrospective measures over longer time periods is important. Interestingly, Blau and Kahn (2013) document that the retrospective data match quite well with the annual survey data in the PSID. This suggests that there is promise in adding retrospective experience data to cross-sectional surveys.

In a next step Blau and Kahn (2013) include two important questions for their work in the PDII questionnaire to essentially ask retrospective questions in a CPS-type cross-sectional sample. The appendix to Blau and Kahn (2013) displays the actual questions included in the PDII for this

purpose. Blau and Kahn (2013) show that having actual experience data is particularly important for exploring a number of important areas in labor economics such as the gender wage gap, residual wage inequality, and human capital accumulation after one has left formal schooling.

In another paper that also considers issues of human capital, Autor and Handel (2013) use among other ideas the “task framework” of Autor, Levy, and Murnane (2003) as a way to consider what employers demand. The paper fundamentally creates a structure to consider the causal relationships among human capital endowments, occupational assignments, job tasks, and wages that is related to a Roy (1951) model of allocation of workers to occupations.

The authors note two important problems when studying such questions. First is the issue of linking tasks (which are characteristics of jobs) and human capital (which are characteristics of workers).

The second challenge is that of data. Autor and Handel (2013) note that most data sets used by social scientists have rich information on demographic and income measures but lack detail on tasks of jobs. Several in this area have turned to the Occupational Information Network (O\*NET) and the previous incarnation of those data, the Dictionary of Occupational Titles (DOT), for the job task data. These sources have very rich detail by occupation (not individual) on job requirements but are limited somewhat on task data. This is primarily due to the fact that the data are collected for occupations and not individuals.

Autor and Handel (2013) note that tasks are different from concepts such as experience and education and explicitly make the distinction between “returns to education” and “returns to tasks.” They go on to note that tasks are “not fixed worker attributes.” Further they note that workers can select into certain jobs and move to jobs when the value of tasks changes, thus motivating the Roy (1951) model.

An additional feature of the paper is to consider tasks at the person level (recall that the O\*NET data are at the level of the occupation). This is where the PDII comes in with this paper. Since there are data at the individual level on cognitive, interpersonal, and physical tasks, the authors can consider whether tasks vary not only between occupations (which they could do with other data) but also within occupations.

Blinder and Krueger (2013) use PDII data to add to the growing literature on “offshorability.” Note that this is not outsourcing. Outsourcing refers to movement of jobs outside of a company whether or not those jobs leave the country. Offshoring, the subject of their paper, is the movement of jobs out of the country whether or not they go outside of the company. So if, for example, Cornell moves its back-office alumni affairs staff from Ithaca to Toronto it has offshored jobs but not outsourced them. If, on the other hand, Cornell eliminates its dining staff in Ithaca and hires an outside contractor to do those jobs, it clearly outsourced but

has not offshored any jobs. Somewhat related to Autor and Handel (2013), another important distinction is relevant here: offshorability is a feature of one's job, not one's self.

The authors make another important distinction between "offshoring" and "offshorability." The former is what they call an "observable action" (although actually observing the action is likely more difficult than the idea of observing it) and the latter is a characteristics of a job. It is useful to think of jobs that are either easily offshorable (e.g., call center worker) and those that are clearly not (e.g., hair stylist).

Previous work (including Blinder 2009) has also examined offshorability by using O\*NET job content information to essentially assign a subjective index to the hundreds of occupations in the data. But note from the discussion of Autor and Handel (2013) that the O\*NET data are only at the level of occupation. So in order to pick up the within-occupation variation in offshorability, again the PDII comes to the rescue. A series of questions was developed and asked of each of the PDII respondents (more than 2,500 of them) about whether their job could be done from a "remote location" or whether it requires the respondent to be "physically present."

The results are interesting. According to their various measures, Blinder and Krueger (2013) estimate that around one in four jobs (when the survey was conducted in 2008) are potentially offshorable, and when asked whether a job is either offshorable or not, three of their measures agreed between 70% and 80% of the time. The authors indicate that the Census Bureau could generate historical data on offshorability without adding questions to either the CPS or the ACS since those who coded surveys used the standard questions supplemented only with detailed instructions on measuring offshorability.

There are many instances where statistics reported from one source might reveal information that appears inconsistent with that reported from another source. One recent example is the apparent inconsistency where employer costs for compensation are increasing at the same time worker wages remain flat. Since the wage data may come from the CPS and the employer costs data from Employer Costs for Employee Compensation it may be that different surveys are the reason. Or it may be that employer costs include information beyond wages and salaries (as they obviously do; take health insurance, for example). Abraham et al. (2013) consider differences in employment between household and establishment data. Changes resulting from the recent Patient Protection and Affordable Care Act make consideration of bettering the data collected on employer costs for employer compensation versus employee compensation received a very timely issue for labor economists and public policy researchers. A paper in the first year of the *Journal of Labor Economics* by Mellow and Sider (1983) on the accuracy of responses in labor market surveys is in the same spirit as Abraham et al. (2013).

Abraham et al. (2013) start with an interesting basic framework where they consider a two-by-two matrix for whether a wage and salary job is reported by a household (yes or no) and whether a wage and salary job is reported by the employer. Of course, things get considerably more complicated very quickly. Imagine just the case where an individual reports more than one job and what sorts of problems that alone would create.

The empirical part of the paper relies heavily on CPS records linked to unemployment insurance records linked to the same people. The authors note that the CPS-UI data they use were constructed by the Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) program as described in Abowd et al. (2009).

Abraham et al. (2013) find very large differences in both employment status and the number of jobs held when using the CPS-UI matched data. For example, 17.6% of workers in the CPS who are in jobs that should be covered by UI are not found in the UI data and 6.4% of UI workers do not show up in the CPS. When those with more than one job are considered, the problems are even more substantial. They go on to show that those in the "off diagonal" cells of the matrix mentioned above are not random and are linked to both observable person and job characteristics. Importantly those characteristics are associated with what they call marginal workers and marginal or nonstandard jobs.

The final paper, Kleiner and Krueger (2013), is about occupational licensing. Specifically, it documents the prevalence and nature of licensing and the relationship between licensing and certification on wages. The authors highlight three major types of occupational regulation. The first is registration, where individuals essentially file their information with some agency before performing in their occupation. More restrictive is certification, where some group (government, nonprofit) uses some form of test to certify some level of achievement. The strictest form of occupational regulation is licensure or what Kleiner and Krueger (2013) call the "right to practice." They note that, in 2003, 800 occupations were licensed in at least one state in the United States.

As part of the PDII survey data, several questions were included that had to do with licensure. These included three yes/no questions: (i) "Do you have a license or certification that is required by a federal, state or local government agency to do your job?" (ii) "Would someone who does not have a license or certificate be legally allowed to do your job?" and (iii) "Is everyone who does your job eventually required to have a license or certification by a federal, state or local government agency?" Thirty-five percent answered yes to the first question, and 6% noted in the second question that those without a license could do the work in the first question. Another 3% noted in question iii that they would eventually need a certification, so that makes the total that are or will be certified 38%.

They go on to show that the licensing is associated with roughly 18% higher wages.

#### IV. Concluding Comments

The last 30 years have seen extraordinary improvements in data for labor economists. This paper briefly highlights some of the changes and describes a few examples of papers that feature some of these changes. It also describes papers that came out of the even more recent Princeton Data Improvement Initiative—a program that considers additional advances in data collection. This paper also documents the changing ways data have been used in the *Journal of Labor Economics* over the past 30 years, including a trend toward a higher fraction of papers using any data and, among those papers using any data, an increasing fraction using nonpublic data and international data and more frequent use of multiple data sources.

It will be interesting to see what the next 30 years bring. I expect more advances in data and data availability—good news for labor economists.

#### References

- Abowd, John M., Francis Kramarz, and David Margolis. 1999. High wage workers and high wage firms. *Econometrica* 67 (March): 251–333.
- Abowd, John, Bryce E. Stephens, Lars Vilhuber, Fredrik Andersson, Kevin J. McKinney, Marc Roemer, and Simon Woodcock. 2009. The LEHD infrastructure files and the creation of the quarterly workforce indicators. In *Producer dynamics*, ed. Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts. Chicago: University of Chicago Press.
- Abraham, Katharine G., John Haltiwanger, Kristin Sandusky, and James R. Spletzer. 2013. Exploring differences in employment between household and establishment data. *Journal of Labor Economics* 31, no. 2, pt. 2: S129–S172.
- Altonji, Joseph G., and Charles R. Pierret. 2001. Employer learning and statistical discrimination. *Quarterly Journal of Economics* 116 (February): 313–50.
- Angrist, Joshua. 1990. Lifetime earning and the Vietnam era draft lottery: Evidence from Social Security earnings data. *American Economic Review* 80 (June): 313–36.
- Angrist, Joshua D., and Alan B. Krueger. 1991. Does compulsory school attendance affect schooling and earnings? *Quarterly Journal of Economics* 106, no. 4:979–1014.
- . 1999. Empirical strategies in labor economics. In *Handbook of labor economics*, vol. 3, ed. Orley C. Ashenfelter and David Card. Amsterdam: Elsevier.

- Ashenfelter, Orley. 1972. Racial discrimination and trade unionism. *Journal of Political Economy* 80, no. 3, pt. 1:435–64.
- . 1978. Estimating the effect of training programs on earnings. *Review of Economics and Statistics* 60 (February): 47–57.
- . 1983. Determining participation in income-tested social programs. *Journal of the American Statistical Association* 78 (September): 517–25.
- Ashenfelter, Orley, and Alan B. Krueger. 1994. Estimates of the economic return to schooling from a new sample of twins. *American Economic Review* 84 (December): 1157–73.
- Autor, David H., and Michael J. Handel. 2013. Putting tasks to the test: Human capital, job tasks, and wages. *Journal of Labor Economics* 31, no. 2, pt. 2: S59–S96.
- Autor, David H., Frank Levy, and Richard J. Murnane. 2003. The skill content of recent technological change: An empirical investigation. *Quarterly Journal of Economics* 118 (November): 1279–1333.
- Baker, George, Michael Gibbs, and Bengt Holmstrom. 1994a. The internal economics of the firm. *Quarterly Journal of Economics* 109 (November): 881–919.
- . 1994b. The wage policy of a firm. *Quarterly Journal of Economics* 109, no. 4:921–55.
- Becker, Gary. 1962. Investment in human capital: A theoretical analysis. *Journal of Political Economy* 70, no. 5, pt. 2:9–49.
- Blau, Francine D., and Lawrence M. Kahn. 2013. The feasibility and importance of adding measures of actual experience to cross-sectional data collection. *Journal of Labor Economics* 31, no. 2, pt. 2: S17–S58.
- Blinder, Alan S. 2009. How many U.S. jobs might be offshorable? *World Economics* 10 (April–June): 41–78.
- Blinder, Alan S., and Alan B. Krueger. 2013. Alternative measures of offshorability: A survey approach. *Journal of Labor Economics* 31, no. 2, pt. 2: S97–S128.
- Card, David, Raj Chetty, and Andrea Webber. 2007. Cash-on-hand and competing models of intertemporal behavior: New evidence from the labor market. *Quarterly Journal of Economics* 122 (November): 1511–60.
- Card, David, and Alan B. Krueger. 1992. Does school quality matter? Returns to education and school quality in the United States. *Journal of Political Economy* 100 (February): 1–40.
- . 1994. Minimum wage and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *American Economic Review* 84, no. 4:772–93.
- Choi, Hyunyoung, and Hal Varian. 2011. Predicting the present with Google trends. Working paper, Google.

- Davis, Steven J., R. Jason Faberman, John C. Haltiwanger, and Ian Rucker. 2010. Adjusted estimates of worker flows and job openings in JOLTS. In *Labor in the new economy*, ed. Katherine G. Abraham, James R. Spletzer, and Michael J. Harper. Chicago: National Bureau of Economic Research and University of Chicago Press.
- Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1998. *Job creation and destruction*. Cambridge, MA: MIT Press.
- Dyke, Andrew, Carolyn J. Heinrich, Peter R. Mueser, Kenneth R. Troske, and Kyung-Soeng Jeon. 2006. The effects of welfare-to-work program activities on labor market outcomes. *Journal of Labor Economics* 24, no. 3:567–607.
- Farber, Henry S., and Robert Gibbons. 1996. Learning and wage dynamics. *Quarterly Journal of Economics* 111, no. 4:1007–47.
- Freeman, Richard B. 1976. A cobweb model of the supply and starting salary of new engineers. *Industrial and Labor Relations Review* 29 (January): 236–48.
- Goldin, Claudia, and Cecilia Rouse. 2000. Orchestrating impartiality: The impact of “blind” auditions on female musicians. *American Economic Review* 90 (September): 715–41.
- Hallock, Kevin F., and Craig A. Olson. 2010. New data for answering old questions regarding employee stock options. In *Labor in the new economy*, ed. Katherine G. Abraham, James R. Spletzer, and Michael J. Harper. Chicago: National Bureau of Economic Research and University of Chicago Press.
- . 2012. Employees’ choice of method of pay. Working paper, Cornell University.
- Heckman, James J., and Jeffrey Smith. 2000. The sensitivity of experimental impact estimates (estimates from the national JTPA study). In *Youth employment and joblessness in advanced countries*, ed. David G. Blanchflower and Richard B. Freeman, 331–56. Chicago: University of Chicago Press.
- Jacobson, Louis S., Robert J. LaLonde, and Daniel Sullivan. 1993. Earning losses of displaced workers. *American Economic Review* 83 (September): 685–709.
- Kleiner, Morris M., and Alan B. Krueger. 2013. Analyzing the extent and influence of occupational licensing on the labor market. *Journal of Labor Economics* 31, no. 2, pt. 2: S173–S202.
- Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz. 2007. Experimental analysis of neighborhood effects. *Econometrica* 75 (January): 83–119.
- LaLonde, Robert J. 1988. Evaluating econometric evaluations of training programs using experimental data. *American Economic Review* 76, no. 4:604–20.

- Lazear, Edward P. 2000. Performance pay and productivity. *American Economic Review* 90 (December): 1346–61.
- Lester, Richard. 1948. Company wage policies: A survey of patterns and experience. Research Report no. 77, Industrial Relations Section, Princeton University.
- Medoff, James L., and Katharine G. Abraham. 1980. Experience, performance and earnings. *Quarterly Journal of Economics* 95 (December): 703–36.
- Mellow, Wesley, and Hal Sider. 1983. Accuracy of response in labor market surveys: Evidence and implications. *Journal of Labor Economics* 1, no. 4:331–44.
- Mincer, Jacob. 1962. On-the-job training: Costs, returns and some implications. *Journal of Political Economy* 70, no. 5, pt. 2:50–79.
- Mincer, Jacob, and Solomon Polachek. 1974. Family investments in human capital: Earnings of women. *Journal of Political Economy* 82, no. 2, pt. 2: S76–S108.
- Roy, A. D. 1951. Some thoughts on the distribution of earnings. *Oxford Economic Papers* 3, no. 2:135–46.
- Stafford, Frank. 1986. Forestalling the demise of empirical labor economics: The role of micro data in labor economics research. In *The handbook of labor economics*, ed. Orley Ashenfelter and Richard Layard, chap. 7. Amsterdam: North Holland.