10-1-2004

Some Reflections on the Use of Administrative Data to Estimate the Net Impacts of Workforce Programs in Washington State

Kevin M. Hollenbeck

W.E. Upjohn Institute for Employment Research, hollenbeck@upjohn.org

Upjohn Institute Working Paper No. 04-109

Follow this and additional works at: https://research.upjohn.org/up_workingpapers

Part of the [Labor Economics Commons](#)

Citation

Hollenbeck, Kevin M. 2004. "Some Reflections on the Use of Administrative Data to Estimate the Net Impacts of Workforce Programs in Washington State." Upjohn Institute Working Paper No. 04-109. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. <https://doi.org/10.17848/wp04-109>

This title is brought to you by the Upjohn Institute. For more information, please contact repository@upjohn.org.

***Some Reflections on the Use of Administrative Data to Estimate the
Net Impacts of Workforce Programs in Washington State***

Upjohn Institute Staff Working Paper No. 04-109

Kevin Hollenbeck
W.E. Upjohn Institute for Employment Research
300 S. Westnedge Ave.
Kalamazoo, MI 49007
hollenbeck@upjohninstitute.org

October 29, 2004

JEL Codes: J24, J31, I28, C93

Paper presented at the APPAM meetings in Atlanta, GA on October 29. This paper builds on work that was done under contract to the Washington Workforce Training and Education Coordinating Board, under contract to the U.S. Department of Labor Employment and Training Administration for the National Workforce Investment Research Colloquium, and under subcontract to the University of Baltimore (prime contractor to the U.S. Department of Labor Employment and Training Administration) for work on the ADARE Consortium. The contractual support of these agencies as well as the resources and support of the Upjohn Institute are gratefully acknowledged. The usual caveat applies.

***Some Reflections on the Use of Administrative Data to Estimate the
Net Impacts of Workforce Programs in Washington State***

Abstract

The purpose of this paper is to reflect on the results, methodology, and processes used in a series of net labor market impact studies done for the State of Washington over the past six years. All of the studies relied on administrative data and used a technique referred to as quasi-experimental evaluation. The program interventions were the federal- and state-funded workforce development programs. The paper sets out eight “reflections” for analysts and policy makers to consider. These reflections identify lessons learned and uncertainties or issues that need more consideration and scrutiny.

Over the past six years, I have been involved in several studies in which estimates of the net labor market impacts of federally-funded workforce development program services for adults and dislocated workers in the state of Washington (JTPA and WIA) were computed (Hollenbeck and Huang [2003]; King, Schroeder, and Hollenbeck [2003]; Hollenbeck, King, and Schroeder [2003]; Hollenbeck [2004]; and Hollenbeck, Schroeder, King, and Huang [2004]). In all of these studies, the data came from administrative sources, and the estimation technique was “quasi-experimental.” The “treatment” group was individuals who received services funded by JTPA or WIA, and the “comparison” group was, most often, individuals who used the public labor exchange (called JobNet in Washington), but who did not receive JTPA or WIA services.¹ The purpose of this paper is to reflect on the results, methodology, and processes used in those studies to identify some “lessons learned” for analysts or policy makers and to point out uncertainties or issues that need more consideration and scrutiny.

RESULTS

Each of the studies that were conducted examined several labor market outcomes. The outcomes were, for the most part, alternative measures of employment and earnings. The studies rely on earnings records from the Unemployment Insurance wage record system as a primary data source, so the measures of employment and earnings use a calendar quarter for an accounting period. Washington is somewhat unique in that employers report quarterly hours of work in addition to earnings, so employment can be measured as having earnings in a quarter, or it can be measured as hours of employment in a quarter.

¹ In some instances, the “treatment” group was individuals who received certain types of services through JTPA or WIA and the “comparison” group was individuals who received other, less intensive, services through JTPA or WIA.

The net impact studies derive their estimates from individual-level longitudinal data that cover three time frames. *Pre-registration* data cover the time period before the observed individuals received program services for the “treatment” group cases or received counterfactual services for the “comparison” group cases. *Program services* data record the experiences and services that were received by the “treatment” group cases or “comparison” group cases while they were participating in the program being analyzed or in the counterfactual program. *Post-exit* data provide information on the individuals after they have received program services and have been determined officially to have exited from the program.²

Much of the variation in outcome measures of the net impact studies reflects different lengths of time periods for the post-exit data. So, for example, the studies might measure employment and quarterly earnings at the most recently-available quarter. In one study, for example, that was the first calendar quarter of 2004 (denoted as 2004:Q1). So that study examined the following outcomes:

Employment (latest):	=	1 if earnings > \$100 in 2004:Q1; ³ else 0
Hours (latest):	=	hours of employment reported in 2004:Q1
Earnings (latest):	=	earnings in 2004:Q1
Conditional Earnings (latest):	=	earnings in 2004:Q1 if non-zero; else missing

² Whether an individual has exited from a program is sometimes difficult to determine and to date. For example, an individual in a training program may simply stop attending (perhaps because they became employed) without notifying anyone in the program. Many of the analyses being discussed use receipt of services from the Employment Service as the counterfactual. Again an individual may simply never return to the ES office. In these cases, Washington uses a six-month “soft exit” rule that indicates that no activity over a six month period is called an exit, and the “exit date” is set at the last date of service.

³ A small percentage of wage records have quarterly earnings of less than \$100 (2000 \$), and we typically delete those records as not reflecting substantial employment. We also typically “edit” extremely high levels of quarterly earnings.

However, individuals in the “treatment” group or “comparison” group may have exited in different calendar quarters, so that they may have quarterly time series of labor market experience data that have different lengths. So net impact studies may use a second concept of employment and earnings, namely, these concepts measured at the same number of quarters after exit for everyone. That is, the following concepts are used:

Employment + n: = 1 if earnings > \$100 in exit quarter + n; else 0

Hours + n: = hours of employment in exit quarter + n

Earnings + n: = earnings in exit quarter + n

Conditional Earnings + n: = non-zero earnings in exit quarter + n; else missing

Finally, to smooth out variations that might arise simply by examining data from a single quarter, employment and earnings outcomes may be averaged over several post-exit quarters. It is often the case that the first quarter after exit is omitted because, for technical reasons, it is only a partial quarter of earnings for some individuals. Averaged outcomes then, are defined as follows:

Employment (n – m): = arithmetic average of employment + n through employment + m

Hours (n – m): = arithmetic average of hours + n through hours + m

Earnings (n – m): = arithmetic average of earnings + n through earnings + m

Conditional Earnings (n – m): = arithmetic average of number of non-missing conditional earnings + n through conditional earnings + m

A striking finding from the different Washington studies is the relative consistency of the results despite different methods, different years of data, different post-exit time periods, and other differences. Tables 1 and 2 provide summaries of results from the various studies for

Table 1

Summary of Net Impact Results Across Studies, Adults

Study	Treatment	Comparison	Outcomes	Impact	Treatment mean ^a
I	JTPA II-A; exited in 97/98 n = 2,772	ES services; exited in 97/98; no other program n = 76,762	Employment + 3	10.9***	55.6%
			Earnings + 3	\$302***	\$1,603
			Employment (8–11)	7.4***	55.6%
			Earnings (8–11)	\$568***	\$1,603
	JTPA II-A; exited in 99/00 n = 2,463	ES services; exited in 97/98; no other program n = 157,568	Employment + 3	3.6***	58.4%
			Earnings + 3	\$179*	\$1,846
II	WIA adults; exited in PY00 n = 866	ES services; exited in PY00 n = 164,477	Employment + 4	5.9***	73.6%
			Earnings + 4	\$489***	\$2,418
III	WIA (exc. Core Only) adults; female; exited in PY00; 75% random sample n = 391	ES services; female; exited in PY00; 50% random sample n = 28,733	Employment (2–4)	10.9*** -- 15.5***	74.4%
			Earnings (2–4)	\$391 -- \$604***	\$2,009
	WIA (exc. Core Only) adults; male; exited in PY00; 75% random sample n = 292	ES services; male; exited in PY00; 50% random sample n = 39,241	Employment (2–4)	6.8 -- 8.0**	73.1%
			Earnings (2–4)	\$227 -- \$551	\$2,909
IV	WIA (exc. Core Only) adults; female; exited in PY00–PY01 n = 1,076	ES services; female; exited in PY00–PY01 n = 133,884	Employment (2–7)	10.0*** -- 10.2***	74.1%
			Earnings (2–7)	\$402*** -- \$455***	\$2,083
	WIA (exc. Core Only) adults; male; exited in PY00–PY01 n = 2,062	ES services; male; exited in PY00–PY01; n = 190,715	Employment (2–7)	11.9*** -- 12.7***	75.6%
			Earnings (2–7)	\$203*** -- \$294***	\$2,810

NOTE: Dollar figures in constant 2000\$. *** significant at the 0.01 level; ** significant at the 0.05 level; * significant at the 0.10 level.

^aAverage for pre-registration period.

adults served through Title II-A of JTPA and its successor program in WIA and for dislocated workers served through Title III of JTPA and its successor program in WIA.

The different studies all point to rather positive results for publicly-funded workforce development programs. The programs for the adult workers (not dislocated workers) seem to raise employment and earnings by about 20 percent. When disaggregated by sex as was done for

Table 2

Summary of Net Impact Results Across Studies, Dislocated Workers

Study	Treatment	Comparison	Outcomes	Impact	Treatment mean ^a
I	JTPA III; exited in 97/98 n = 4,475	ES services; exited in 97/98; no other program n = 84,106	Employment + 3	7.5***	86.4%
			Earnings + 3	\$239***	\$6,314
			Employment (8–11)	4.2***	86.4%
			Earnings (8–11)	\$871***	\$6,314
	JTPA III; exited in 99/00 n = 3,964	ES services; exited in 97/98; no other program n = 179,151	Employment + 3	2.2***	85.6%
			Earnings +3	–\$239***	\$5,911
II	WIA dislocated workers; exited in PY00 n = 14442,048	ES services; exited in PY00 n = 164,477	Employment + 4	10.8***	93.0%
			Earnings + 4	\$312**	\$6,844
IV	WIA (exc. Core Only) dislocated workers; female; exited in PY00–PY01 n = 745	ES services; female; exited in PY00–PY01 n = 133,884	Employment (2–7)	10.6*** -- 11.2***	92.6%
			Earnings (2–7)	\$247* -- \$339***	\$6,067
	WIA (exc. Core Only) dislocated workers; male; exited in PY00–PY01 n = 2,791	ES services; male; exited in PY00–PY01 n = 190,715	Employment (2–7)	9.9*** -- 11.1***	93.5%
			Earnings (2–7)	\$364*** -- \$406***	\$7,759

NOTE: Dollar figures in constant 2000\$. *** significant at the 0.01 level; ** significant at the 0.05 level; * significant at the 0.10 level.

^aAverage for pre-registration period.

the last studies shown in table 1, the results suggest that the net impacts for women are somewhat larger than they are for men. For men, the positive impact of JTPA and then WIA seemed to operate through employment (the percentage increase in earnings is approximately the same as the percentage increase in employment). For women, the percentage increase in earnings exceeds the percentage increase in employment, so women must have experienced a positive net impact in hours of work or wage rates in addition to the net impact on employment.

For dislocated workers, the employment and earnings net impacts are positive and significant. But in this case, for both men and women, the percentage increase in employment is greater than the percentage increase in earnings, suggesting that impacts on hours or wage rates

were not positive. Overall average earnings increased, but the strongly positive employment impacts were offset partially by negative wages or hours impacts.

The consistency of the results leads me to the following reflection:

Reflection #1: Consistent evidence suggests that the federal job training programs as administered in Washington are effective, especially in increasing employment rates, but also in generating higher earnings. For (non-dislocated worker) adults, the employment impact is on the order of 15–20 percent, and the earnings impact is on the order of 10–20 percent for men and 20–40 percent for women. For dislocated workers, the employment impact is on the order of 10–15 percent. The earnings impact is on the order of 5–10 percent for both males and females.

Because the studies have been conducted over a period of years, and because the underlying data for each study cover a period of a few years, the studies can provide some evidence on the sensitivity of the net impacts to the business cycle. In particular, the first study, which estimated the net impacts of individuals served by JTPA, followed one cohort of program exiters (97/98) at an upward trending point in the cycle (between trough and peak). The results, in the first row of Study I in the tables, show substantially positive impacts in the short-run (three quarters after exit). However, the labor market peaked in 99/00, and the impacts for the 97/98 cohort became dampened (see tables for impacts averaged over 8 to 11 quarters after exit), and the impacts for the 99/00 cohort at +3 were considerably smaller than the earlier cohort of exiters.

The labor market plummeted after the program year 99/00, so the results shown in the bottom panels of the table pertain to a downward trending business cycle. The results show fairly substantial employment impacts even though the labor market was quite soft, with slightly smaller impacts when more quarters of outcome data are considered. In short, the more recent

studies do not seem to be driven by cyclical pressures. This leads to a second reflection about the study results.

Reflection #2: The evidence is weak, at best, but it appears as though the impacts of federal job training programs are weakly influenced by the business cycle. Researchers should examine findings from similar studies done over substantial periods of time in other states as well as continuing these types of studies in Washington in order to gain confidence about the relationship between the business cycle and the effectiveness of job training programs.

METHOD

To set the stage for a discussion about methodology, compare and contrast a random assignment experiment to the “quasi-experimental” approach used in the Washington studies. Consider these two approaches in a different arena—clinical testing of a pharmaceutical. In the clinical testing of a pharmaceutical, trials are set up by identifying an eligible population (for example, individuals with a certain health condition), and then randomly assigning individuals into a treatment group that receives the drug and a control group that receives a placebo. During and after the drug regimen, the clinicians monitor particular outcomes, most especially the condition for which the pharmaceutical was developed. The efficacy of the drug is easily measured as the difference in outcomes between the treatment and control groups.

The quasi-experimental analog to the clinical trial would involve identifying individuals who happened to have been eligible for the drug and happened to have taken it. This is the “treatment” group. The individuals in this group are compared to individuals who also were eligible, but who did not take the drug. We generally assume that the actions of these individuals represent a good counterfactual for the “treatment” group; that is, they represent what might have happened to the individuals in the “treatment” group if the drug were not available. Their

actions represent the next best alternative. The group of individuals who did not take the drug is the “comparison” group.

There are, of course, substantial differences between the experiment and the quasi-experiment. With an experimental design, data are generally collected before the treatment as well as during and after. With the quasi-experimental design, the pre-program and service period data must be recalled or come from administrative sources that happen to exist. But most importantly, an experimental design guarantees that the treatment is not correlated with any characteristics of the individuals who receive the treatment. Consequently, the analyst can have high statistical confidence that outcomes are attributable to the treatment. With the quasi-experiment, the analyst is never quite sure why the comparison group members did not receive the treatment, and therefore outcomes may be explained in part by the same characteristics that explain why some individuals received the treatment and others didn't.

As Imbens (2004) points out, whether data from a quasi-experimental design can be used to identify net impacts is an empirical question. Basically, if the data conform to two assumptions, then net impact estimates from a quasi-experimental design gain credence. The two assumptions are “conditional independence of the treatment” and “overlap” (see Hollenbeck 2004). The first assumption essentially suggests that membership in the treatment is random once all of the characteristics observed by the econometrician are controlled. The second assumption essentially implies that there are no observable variables that perfectly explain participation in the treatment. For any set of characteristics, there is a non-null probability that the individual may be in the treatment or in the comparison group.

With the Washington studies, we have consistently defined the treatment and comparison groups by program exit dates. For example, a treatment group would be defined as individuals

who exited in PY 2000 (July 2000 – June 2001) and the comparison group would be individuals who existed from JobNet during that same period. However, this approach allows for substantial variation in the extent of program services between the treatment and comparison groups. An individual in the treatment group may have registered for JTPA several years prior to exiting from WIA (there was a substantial administrative exiting from JTPA and enrollment into WIA in late 1999 or early 2000 when WIA was started), and thus received services for several years. On the other hand, an individual in the comparison group may have registered for the ES in June 2001 and exited immediately. Or vice versa—the WIA individual may have been in the program for a very short amount of time relative to a lengthy period of services received by the ES individual.

Up to this point, there has been no attempt to analyze the sensitivity of the results to this definition of the treatment compared to a definition based on quarter of registration. In other words, the treatment group could be defined by individuals who encountered the job training program during a particular period of time, and the comparison group would be individuals who registered for the ES during the same period of time. In this scenario, the length of the treatment would be (relatively) constant. Thus, the following reflection:

Reflection #3: Sensitivity tests of the net impact results should be conducted in which the treatment and comparison sets of individuals are defined relative to registration date rather than exit date.

In practice, the quasi-experimental approach identifies a relatively large data set of observations, from which the comparison data set may be chosen. In the pharmaceutical testing analog, it is usually the case that the treatment cases can be identified, but there is a large pool of observations about individuals who did not take the drug. Some of the individuals were eligible (i.e., had the same health condition as the treatment cases) and some were not eligible. Two

types of net impact estimates are possible in this case. First, estimates can be derived using the whole set of observations in the comparison set (call these whole sample estimators); second, estimates can be derived by using a subset of observations that “match” the treatment group (called matched sample estimators; see Hollenbeck [2004]). The rationale for the second approach is obvious—estimates should only be derived from matching (eligible) observations. However, the rationale from the first approach suggests that information about the effect of the treatment can be gained from cases that are “nearly” eligible.

A surprising finding from the Washington studies is that the whole sample estimators and matched sample estimators are quite similar. In particular, a technique called “propensity score blocking” produced net impact estimates that were quite similar to the matched sample estimators with lower standard errors. Mueser, Troske, and Gorislauskys (2003) find this same result using administrative data for Missouri.⁴ These results suggest to me that this technique may be the preferred technique for these kinds of estimates, as stated in the following:

Reflection #4: Whole (comparison) group estimation using propensity score blocking seems to result in reasonable estimates that have more precision than matched sample estimates.

The administrative data that have been used in the Washington net impact studies typically have considerable pre-registration information on employment and earnings from UI wage record data, and socio-demographic information about individuals based on applications for program services that are completed at the time of registration. So, for example, there might be 10 quarters of wage record data (earnings, hours, and industry) producing 30 variables that measure pre-registration labor experience. Whereas the only information about education may be one or two variables taken from the individual’s application. However, theory suggests that

⁴ Note that these authors call the technique, “matching by propensity score category.”

both prior labor market experience and education are important determinants of program success. So, it was necessary to generate a small number of variables that summarize pre-registration employment and earnings.

The variables that we constructed were as follows:

- Percent of quarters employed prior to registration
- Average prior quarterly earnings, zeros included in the calculation
- Coefficient of the fitted earnings trend regression for all the prior quarters
- Variance of the estimated earnings trend coefficient
- Percent of quarters with multiple employers during the prior quarters
- Had an earnings dip, defined as a 20 percent or more decline in quarter-to-quarter earnings
- Number of quarters between dip and registration
- Size of the dip in percentage terms

These summary variables seem to work well. They generally have the expected signs and are significant in the logit regressions that are estimated in order to calculate propensity scores. The policy implication from the fact that these variables seem to nicely summarize pre-registration employment and earnings is that these variables may be good indicators for programs to monitor when attempting to serve clients.

Reflection #5: The variables created to summarize pre-registration employment and earnings seem to “work” well. They have strong explanatory power in program participation equations. Nevertheless, it may be appropriate to experiment somewhat with these variables to see if a more parsimonious or statistically powerful set of variables could be identified.

As mentioned, two key empirical assumptions that are made in the quasi-experimental methodology are “conditional independence of the treatment” and “overlap.” Most of the studies

of the net impacts of federal training programs in Washington acknowledge the importance of these assumptions and attempt to test them empirically. However, there is not wide agreement on the particular specification tests that are most appropriate. Some of the tests in the literature seem quite arbitrary.

For example, many papers that use a quasi-experimental method use a graphical presentation to convince readers that there is substantial overlap in the comparison and treatment samples. These graphs typically display the probability density functions of the propensity scores, and show that the density function for the treatment group overlaps with the comparison group considerably. (See Dehejia and Wahba 1999, for example.) However, there does not seem to be a quantitative test of whether there is adequate “overlap” or “support.” An *ad hoc* test was suggested by a study done by Battelle Memorial Institute (n.d.). This study asserted that a reasonable assurance of overlap is that the propensity score that identifies the lowest quintile of *p*-scores for the treatment sample should approximate the 80th percentile of the *p*-scores for the matched comparison set. However, the Battelle study does not really justify this assertion.

Similarly, specification testing of the “conditional independence” assumption has been *ad hoc*. Hollenbeck, King, and Schroeder (2003) conducted an *ex post* regression approach to testing whether the variable identifying if an observation was in the treatment group or not had been purged of correlation from observable covariates. Mueser, Troske, and Gorislawsky (2003) use a “backcasting” technique for specification testing by examining how well the matched comparison group “backcasts” the employment history of the sample in the 5th through 8th quarter prior to program registration. Hollenbeck (2004) reserved a portion of the treatment sample for specification testing.

This leads to my next reflection:

Reflection #6: There appears to be little consensus on the appropriate way to test for “overlap” or for specification testing of the “conditional independence of the treatment.” Much theoretical and empirical progress needs to be made in these areas.

PROCESS

The final area in which I address comments concerns the process through which the quasi-experimental impact studies have been conducted. These studies rely on administrative data bases that states help to procure and, in some cases, construct. Matching and merging data sets is complex and expensive, and the current fiscal crises occurring in most states are placing more and more demands on fewer and fewer state employees. At the same time, it seems as though more and more concerns are being expressed concerning data confidentiality. The future of the quasi-experimental net impact evaluations that are being conducted may depend on their value as perceived by state administrators. Thus it is important to attempt to bring resources to the effort, to be as efficient as possible in processing the data, and to demonstrate the benefits to state officials of the estimates that are derived.

Reflection #7: The fiscal conditions of states may be “squeezing out” the priority that states are putting on constructing administrative data sets such as those needed to conduct the quasi-experimental net impact evaluations that are the basis of this paper. Efforts should be made to reduce the marginal costs of these efforts by finding other resources (such as the ADARE consortium) or by reducing the processing or personnel costs imposed on states. Furthermore, efforts should be made to increase the marginal benefits of the efforts by marketing the power and value of the results for good program management at the state, as well as, national level.

The net impact estimates that have been produced provide consistent, rigorous evidence that workforce development programs in Washington seem to bring positive results to participants. However, it is important for program administrators to insist on multiple studies using many different methodologies in order to get a fully developed picture of the impact of

their programs. In particular, the quasi-experimental net impact estimates do not delve very deeply into the precise types of services that are provided to eligible participants. Thus it behooves program administrators to supplement these quantitative estimates with other sources of information about program impacts, such as qualitative studies of the impacts of programs on various subpopulations of clients. This would be my last reflection.

Reflection #8: Program administrators should never rely on a single study when making high stakes programmatic decisions. Quantitative evaluations should be supplemented with qualitative studies that delve into the operational components of programs at the local level.

This paper presents a few reflections based on several studies of the impact of federally-funded workforce development programs on employment and earnings of Washington participants. The results from several studies that use different data and different techniques are fairly consistent—these programs have a positive impact on clients. Less clear cut is how those impacts vary over the business cycle. In terms of methodology, the studies seem to suggest that “matching by propensity score groups” is perhaps the preferred methodology, and that the methods used to control for pre-registration earnings and employment are rather robust and powerful. More work needs to be done, however, in testing whether the results are stable when the treatment is defined based on registration date rather than exit date, and in developing reasonable specification tests and test of overlap/support. Finally, the paper warns of a possible retrenchment in states’ abilities to conduct these kinds of studies because of their severe fiscal problems.

REFERENCES

- Battelle Memorial Institute. N.d. "Net Impact Evaluation: Appendix A, Technical Appendix." Unpublished report. Battelle Memorial Institute, Seattle, WA.
- Dehejia, Rajeev H. and Sadek Wahba. 1999. "Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs." *Journal of the American Statistical Association* 94(448): 1053–1062.
- Hollenbeck, Kevin. 2004. "Using Administrative Data for Workforce Development Program Evaluation." Upjohn Institute staff working paper no. 04-103. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.
- Hollenbeck, Kevin M. and Wei-Jang Huang. 2003. "Net Impact and Benefit-Cost Estimates of the Workforce Development System in Washington State." Upjohn Institute technical report no. TR03-018. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.
- Hollenbeck, Kevin, Christopher T. King, and Daniel Schroeder. 2003. "Preliminary WIA Net Impact Estimates: Administrative Records Opportunities and Limitations." Paper presented at the Bureau of Labor Statistics and the Workforce Information Council symposium, "New Tools for a New Era!" held in Washington, DC, July 23-24.
- Hollenbeck, Kevin, Daniel Schroeder, Christopher T. King, and Wei-Jang Huang. 2004. "Net Impact Estimates for Training Services Provided under the Workforce Investment Act." Draft report submitted to the Office of Policy and Research, Employment and Training Administration, U.S. Department of Labor, September 30.
- Imbens, Guido W. 2004. "Nonparametric Estimation of Average Treatment Effects under Exogeneity: A Review." *The Review of Economics and Statistics* 86(1): 4–29.
- King, Christopher T., Daniel Schroeder, and Kevin Hollenbeck. 2003. "WIA Intensive and Training Service Net Impact Estimates." Unpublished manuscript. W.E. Upjohn Institute for Employment Research, Kalamazoo, MI.
- Mueser, Peter, Kenneth R. Troske, and Alexey Gorislavsky. 2003. "Using State Administrative Data to Measure Program Performance." Unpublished manuscript. University of Missouri, Department of Economics, Columbia, MO.