THE STATA JOURNAL

The *Stata Journal* publishes reviewed papers together with shorter notes or comments, regular columns, book reviews, and other material of interest to Stata users. Examples of the types of papers include 1) expository papers that link the use of Stata commands or programs to associated principles, such as those that will serve as tutorials for users first encountering a new field of statistics or a major new technique; 2) papers that go "beyond the Stata manual" in explaining key features or uses of Stata that are of interest to intermediate or advanced users of Stata; 3) papers that discuss new commands or Stata programs of interest either to a wide spectrum of users (e.g., in data management or graphics) or to some large segment of Stata users (e.g., in survey statistics, survival analysis, panel analysis, or limited dependent variable modeling); 4) papers analyzing the statistical properties of new or existing estimators and tests in Stata; 5) papers that could be of interest or usefulness to researchers, especially in fields that are of practical importance but are not often included in texts or other journals, such as the use of Stata in managing datasets, especially large datasets, with advice from hard-won experience; and 6) papers of interest to those who teach, including Stata with topics such as extended examples of techniques and interpretation of results, simulations of statistical concepts, and overviews of subject areas.

The *Stata Journal* is indexed and abstracted by *CompuMath Citation Index*, *Current Contents/Social and Behavioral Sciences*, *RePEc: Research Papers in Economics*, *Science Citation Index Expanded* (also known as *SciSearch*), *Scopus*, and *Social Sciences Citation Index*.

For more information on the *Stata Journal*, including information for authors, see the webpage

http://www.stata-journal.com

Lee (2009) treatment-effect bounds for nonrandom sample selection

Harald Tauchmann
University of Erlangen-Nuremberg,
Rheinisch-Westfälisches Institut für Wirtschaftsforschung (RWI),
and Centre of Health Economics Research (CINCH)
Nürnberg, Germany
harald.tauchmann@fau.de

Abstract. Nonrandom sample selection may render estimated treatment effects biased even if assignment of treatment is purely random. Lee (2009, *Review of Economic Studies*, 76: 1071–1102) proposes an estimator for treatment-effect bounds that limit the possible range of the treatment effect. In this approach, the lower and upper bound correspond to extreme assumptions about the missing information that are consistent with the observed data. In contrast to conventional parametric approaches to correcting for sample-selection bias, Lee's bounds estimator rests on very few assumptions. I introduce the new command `leebounds`, which implements the estimator in Stata. The command allows for several options, such as tightening bounds by using covariates.

Keywords: st0364, leebounds, nonparametric, randomized trial, sample selection, attrition, bounds, treatment effect

1 Introduction

Random assignment of treatment provides an ideal setting for identifying treatment effects. Most prominent randomized trials are designed to generate a situation where randomness of treatment is guaranteed, ruling out any potential endogeneity bias. However, this ideal setting can easily be distorted by nonrandom sample attrition. For example, attrition may happen when participants dropout from a program, when a researcher is denied information on the outcome variable, or when death occurs during a clinical trial. While treatment is purely random in the original population, this does not hold for the actual estimation sample if attrition is linked to the treatment status, which potentially leads to attrition bias with perhaps unknown direction.

Parametrically correcting for attrition and selection bias has become a standard procedure in applied empirical research, rendering the seminal method by Heckman (1976, 1979) a workhorse of applied econometrics. This procedure is implemented in Stata by the `heckman` command. However, this parametric approach has been criticized for relying on restrictive assumptions, joint normality in particular, and for being vulnerable to misspecification (Puhani 2000; Grasdal 2001), which has led to the development of semiparametric approaches (Ichimura and Lee 1991; Ahn and Powell 1993). Though these estimators rely on less restrictive distributional assumptions, valid exclusion restrictions

are even more essential. More recently, researchers have proposed bound estimators that require very few assumptions and do not rely on valid exclusion restrictions. Rather than correcting point estimates for potential bias, these estimators determine an interval for the true treatment effect. The interval is based on extreme assumptions about the impact of selection on the estimated effect that are consistent with the data. One such estimator is Horowitz and Manski (2000). This approach does not involve any assumption about the selection mechanism; however, it is applicable only to outcome variables that are bounded to a certain interval because missing information is imputed on the basis of minimal and maximal possible values. This impedes its application to numerous problems and regularly yields very wide bounds.

In this article, I introduce the new command `leebounds`, which facilitates the estimation of alternative bounds proposed by Lee (2009). These alternative bounds impose more structure on the assumed selection mechanism and allow for outcome variables with unbounded support while often yielding more narrow bounds. Thereby, `leebounds` complements the contributions of Beresteanu and Manski (2000) and Palmer et al. (2011), who have already made other bounds estimators available to Stata users. Beresteanu and Manski (2000) provide Stata code for the bounds estimators introduced by Manski (1990) and add further refinements (Manski 1994, 1995, 1997; Manski and Pepper 2000) to the original approach. Unlike Lee's estimator, Manski's bounds are meant to obtain treatment-effect bounds under (nonrandom) treatment selection. Palmer et al. (2011) introduce a Stata command for the bounds estimator developed by Balke and Pearl (1997), which is closely related to Manski's estimators. Here the focus is on estimating treatment-effect bounds under imperfect compliance with a randomly assigned treatment.

In the following section, I summarize Lee's bounds estimator. In section 3, I describe the syntax of `leebounds`. In section 4, I illustrate the application of `leebounds`. In section 5, I conclude the article.

2 The Lee (2009) bounds estimator

2.1 The intuition behind the estimator

Lee (2009) proposes a bounds estimator that estimates an interval for the true value of the treatment effect in the presence of nonrandom sample selection. The estimator rests on only two assumptions: random assignment of treatment and monotonicity. The latter implies that assignment to the treatment group can affect attrition in only one direction. That means that besides observations for which the outcome variable is observed irrespective of the assigned treatment status, the actual estimation sample includes either observations where the outcome is observed because of receiving the treatment or observations where the outcome is observed because of not receiving the treatment, but not both simultaneously.

The bounds estimator trims either the treated or the nontreated observations so that the share of observations with observed outcome is equal for both groups. Trimming is either from above or from below. This corresponds to two extreme assumptions about missing information that are consistent with the observed data and a one-sided selection mechanism. That is, in the group that suffers less from attrition, either the largest or the smallest values of the outcome are regarded as "excess observations" and are excluded from the analysis. This implies that the treatment effect on those that never suffer from attrition is subject to estimation. In this article, I focus on the practical issue of how estimates for the bounds are calculated and, in particular, on how this procedure is implemented in Stata; for more theory, readers can refer to Lee (2009).

2.2 Estimation

Estimating treatment-effect bounds as suggested by Lee (2009) is computationally straightforward. Only a raw group mean and two trimmed group means of the outcome variable need to be calculated. Here Y_i denotes the outcome, T_i is a binary treatment indicator, and S_i is a binary selection indicator, with $S_i = 0$ indicating attriters for which Y_i is not observed. As usual, i indexes observations. The shares of observations with observed outcome in the treatment group, q_T, and its counterpart for the control group, q_C, can then be written as

$$q_T = \frac{\sum_i 1\,(T_i = 1, S_i = 1)}{\sum_i 1\,(T_i = 1)}$$
$$q_C = \frac{\sum_i 1\,(T_i = 0, S_i = 1)}{\sum_i 1\,(T_i = 0)}$$

Here $1(\cdot)$ denotes the indicator function. To simplify notation, we will consider the case $q_T > q_C$; that is, the treatment group suffers less from attrition.[1] Then

$$q = \frac{q_T - q_C}{q_T} \tag{1}$$

and $1 - q$ determines the quantiles at which the distribution of Y in the treatment group are trimmed to exclude extreme values of Y from the analysis. Hence,

$$y_q^T = G_{Y|T=1,S=1}^{-1}(q)$$
$$y_{1-q}^T = G_{Y|T=1,S=1}^{-1}(1-q)$$

determine the marginal values y_q^T and y_{1-q}^T of the outcome that enter the trimmed means, with G_Y^{-1} denoting the inverse empirical distribution function of Y.

---

1. For the opposite case of $q_T < q_C$, all arguments hold symmetrically, with q being defined as $(q_C - q_T)/q_C$, the control group being trimmed at y_q^C and y_{1-q}^C, respectively, and the treatment group remaining untrimmed. For $q_T = q_C$, both the upper and the lower bound coincide with the difference in raw means.

Using this notation, we calculate estimates for the upper bound and the lower bound as

$$\widehat{\theta}^{\text{upper}} = \frac{\sum_i 1\left(T_i = 1, S_i = 1, Y_i \geqslant y_q^T\right) Y_i}{\sum_i 1\left(T_i = 1, S_i = 1, Y_i \geqslant y_q^T\right)} - \frac{\sum_i 1\left(T_i = 0, S_i = 1\right) Y_i}{\sum_i 1\left(T_i = 0, S_i = 1\right)} \quad (2)$$

$$\widehat{\theta}^{\text{lower}} = \frac{\sum_i 1\left(T_i = 1, S_i = 1, Y_i \leqslant y_{1-q}^T\right) Y_i}{\sum_i 1\left(T_i = 1, S_i = 1, Y_i \leqslant y_{1-q}^T\right)} - \frac{\sum_i 1\left(T_i = 0, S_i = 1\right) Y_i}{\sum_i 1\left(T_i = 0, S_i = 1\right)} \quad (3)$$

Lee (2009) considers a purely continuous outcome variable Y. Yet, especially in survey data, variables that are inherently continuous are often imprecisely reported, resulting in "ties" in the observed outcome data. Monthly disposable income can serve as an example, for individuals tend to report a round number, such as \$1,000 or \$1,500. Such ties may violate the intuition behind (2) and (3) if the marginal values y_q^T and y_{1-q}^T are frequent. For this reason, `leebounds` excludes the $q \cdot N_T^S$ (rounded down to the nearest integer) smallest—respectively, largest—values of Y when calculating the trimmed means. Here N_T^S denotes the number of observations in the treatment group for which the outcome variable is observed. This means that a certain fraction of the observations that exhibit the marginal values y_q^T and y_{1-q}^T enter the trimmed means. With no ties in Y, this procedure coincides with (2) and (3).

2.3 Tightening bounds

Estimating Lee (2009) bounds does not involve any covariates. This corresponds to the assumption of random assignment of treatment, under which the differences in conditional and unconditional expectations of Y coincide. Yet covariates that are determined before treatment can be used to tighten treatment-effect bounds. Covariates that have some explanatory power for attrition are used to split the sample into cells, and bounds are separately calculated for each cell. Finally, a weighted average of cells' bounds is computed. The appropriate weights are the probabilities of cell membership for those that never suffer from attrition (Lee 2009, 1094). These probabilities are unknown. However, because of random assignment of treatment and monotonicity, they can consistently be estimated by $\sum_i 1\left(J_i = 1, S_i = 1, T_i = 0\right) / \sum_i 1\left(S_i = 1, T_i = 0\right)$ for each cell J, where $J_i = 1$ indicates membership in J. Lee (2009) shows that such averaged bounds are tighter than those that do not use any covariates (Lee 2009, 1086).[2] Tightening bounds is offered by `leebounds` as an option.

Technically, only a limited number of discrete[3] variables can be used for tightening, because the number of observations and the joint distribution of treatment status and selection must allow for estimating the bounds for each cell. Thus estimation regularly fails if a large number of covariates are used. Tightening could also fail if the control group suffers relatively more from attrition for some cells, while attrition is more frequent

---

2. The proof is for the population parameters, not for their sample analogs. Hence, especially for ill-suited covariates, estimated bounds may fail in getting tighter with option `tight()`.
3. In practice, continuous variables (for example, `age`) must be transformed into categorical ones (`age classes`).

in the treatment group for other cells. Because of sampling error, this will frequently occur if the sample is split into too many cells.[4] `leebounds` checks for this, issues a warning if it detects a selection pattern that is heterogeneous across cells, and saves a macro that indicates the type of the selection pattern.

2.4 Standard errors and inference

Estimates for the treatment-effect bounds are subject to sampling error. Lee (2009, 1088) provides analytic standard errors for them; we refer to the original paper for details about calculating standard errors. Analytical standard errors (or, alternatively, boot-strapped standard errors) are implemented in `leebounds`. Using these standard errors, one can determine "naive" confidence intervals that cover the interval $[\theta^{\text{lower}}, \theta^{\text{upper}}]$ with probability $1 - \alpha$. Interestingly, on the basis of Imbens and Manski (2004), Lee (2009, 1089) also derives a confidence interval for the treatment effect itself, that is, the scalar parameter of ultimate interest. This interval is tighter than the combined confidence interval for θ^{lower} and θ^{upper}. It captures both uncertainty about the selection bias and uncertainty about the sampling error. `leebounds` optionally estimates the confidence interval for the treatment effect.

3   The leebounds command

`leebounds` requires Stata 11 or higher. The prefix command `bootstrap` is allowed but is not recommended. `pweight`s (default), `fweight`s, and `iweight`s are allowed; see [U] **11.1.6 weight**. Observations with a negative weight are skipped for any type of weight.

3.1 Syntax

`leebounds` *depvar treatvar* $\big[$ *if* $\big]$ $\big[$ *in* $\big]$ $\big[$ *weight* $\big]$ $\big[$, <u>sel</u>ect(*varname*)
    <u>tight</u>(*varlist*) <u>cieffect</u> <u>vce</u>(<u>analytic</u>|<u>bootstrap</u>) <u>level</u>(#) $\big]$

depvar is a numeric outcome variable, and *treatvar* is a binary treatment indicator that can be either numeric or a string variable. The (alphanumerically) larger value of *treatvar* is assumed to indicate treatment.

3.2   Options

`select`(*varname*) specifies a binary selection indicator. *varname* can be either numeric or a string variable. The (alphanumerically) larger value of *varname* is assumed to indicate selection. If no selection indicator is specified, any observation with nonmissing information on *depvar* is assumed to be selected, and all observations with missing information on *depvar* are assumed to be not selected.

---

4. This may also indicate a violation of the monotonicity assumption.

tight(*varlist*) specifies a list of covariates for computing tightened bounds. With
tight() specified, the sample is split into cells defined by the covariates in *varlist*.
Continuous variables in *varlist* will cause the estimation procedure to fail.

cieffect requests calculation of a confidence interval for the treatment effect. This
interval captures both uncertainty about the selection bias and uncertainty about
the sampling error.

vce(analytic | bootstrap) specifies whether analytic or bootstrapped standard errors
are calculated for estimated bounds. analytic is the default. bootstrap allows for
the suboptions reps(#) and nodots. For vce(analytic), the covariance for the
estimated lower and upper bound is not computed. If this covariance is relevant, one
should choose vce(bootstrap). Instead of specifying vce(bootstrap), one can use
the prefix command bootstrap, which allows for numerous additional options. Yet
leebounds's internal bootstrapping routine is much faster than the prefix command,
allows for sampling weights by performing a weighted bootstrap, and makes the
option cieffect use bootstrapped standard errors.

level(#) sets confidence level. One can change the reported confidence level by retyp-
ing leebounds without arguments and specifying only the option level(#). This
affects the confidence interval for the bounds, but it does not affect the confidence
interval requested with cieffect.

3.3   Stored results

leebounds stores the following in e():

Scalars
    e(N)   number of observations
    e(Nsel)   number of selected observations
    e(cilower)   lower bound of treatment-effect confidence interval (if cieffect was specified)
    e(ciupper)   upper bound of treatment-effect confidence interval
                      (if option cieffect was specified)
    e(trim)   (overall) trimming proportion
    e(level)   confidence level
    e(cells)   number of cells (if option tight() was specified)
    e(N_reps)   number of bootstrap repetitions (if option vce(bootstrap) was specified)

Macros
    e(cmd)   leebounds
    e(cmdline)   command as typed
    e(title)   title in estimation output
    e(depvar)   name of dependent variable
    e(treatment)   binary treatment indicator
    e(wtype)   weight type
    e(wexp)   weight expression
    e(select)   *varname* (if option select() was specified)
    e(cellsel)   cell-specific selection pattern, homo or hetero (if option tight() was specified)
    e(covariates) *varlist* (if option tight() was specified)
    e(trimmed)   treatment or control
    e(vce)   *vcetype* specified in vce()
    e(vcetype)   title used to label Std. Err.
    e(properties) b V

Matrices
    `e(b)`   vector of estimated treatment-effect bounds
    `e(V)`   variance–covariance matrix of the estimates (covariance set to zero for
                      `vce(analytic)`)

Functions
    `e(sample)`   marks estimation sample

4   Examples

We use `cancer.dta`, which is shipped with Stata, for a simple illustrative application; for serious applications of Lee's bounds estimator besides that in Lee (2009), see Augurzky et al. (2012) or Cawley and Price (2013). We analyze how being treated with an active ingredient (`drug == 2 | drug == 3`) versus being treated with a placebo (`drug == 1`) affects survival time (`studytime`). We treat the data as if information on survival time were available for only those who died during the study (`died == 1`). This is not entirely correct for those who did not die (`died == 0`), because we know that they survived at least for the rest of the study period. Yet, in our illustration, we regard them as attriters without any (valid) information on the outcome `studytime`.

```
. sysuse cancer.dta, clear
(Patient Survival in Drug Trial)

. generate activedrug = (drug == 2 | drug == 3)

. leebounds studytime activedrug, select(died)

Lee (2009) treatment effect bounds

Number of obs.                   =     48
Number of selected obs.          =     31
Trimming porportion              =    0.5489
```

studytime	Coef.	Std. Err.	z	P>|z|	[95% Conf. Interval]	
activedrug						
lower	2.866667	3.909154	0.73	0.463	-4.795134	10.52847
upper	14.3	3.163771	4.52	0.000	8.099123	20.50088

The output displays that 48 individuals participated in the trial. Out of these, 31 died during the study and 17 survived. The latter are regarded as not selected because we have no precise information about survival time. The trimming proportion corresponds to q; see (1). The value 0.5489 indicates that the control group is trimmed by more than half, because the survival rate is much higher among individuals who were treated with an active drug. Correspondingly, the estimated treatment-effect bounds are pretty wide, ranging from a 2.87- to a 14.30-month gain in survival time. Taking standard errors into account, the lower bound does not significantly deviate from zero. To obtain a confidence interval for the treatment effect (see section 2.3), one can choose the `cieffect` option.

```
. leebounds studytime activedrug, select(died) cieffect

Lee (2009) treatment effect bounds

Number of obs.                    =     48
Number of selected obs.           =     31
Trimming porportion               =   0.5489
Effect 95% conf. interval         : [-3.5633   19.5039]
```

studytime	Coef.	Std. Err.	z	P>|z|	[95% Conf. Interval]
activedrug					
lower	2.866667	3.909154	0.73	0.463	-4.795134 10.52847
upper	14.3	3.163771	4.52	0.000	8.099123 20.50088

This interval is narrower than the combined confidence intervals for the bounds. One can allow for a less strict level of confidence by specifying `level(90)`. To illustrate the `vce()` option, we opt for bootstrapped rather than analytic standard errors.

```
. set seed 13052007

. leebounds studytime activedrug, sel(died) cie level(90) vce(boot, reps(250))

............................................... 50
............................................... 100
............................................... 150
............................................... 200
............................................... 250

Lee (2009) treatment effect bounds

Number of obs.                    =     48
Number of selected obs.           =     31
Trimming porportion               =   0.5489
Effect 90% conf. interval         : [-1.9390   18.1498]
```

studytime	Observed Coef.	Bootstrap Std. Err.	z	P>|z|	Normal-based [90% Conf. Interval]
activedrug					
lower	2.866667	3.749864	0.76	0.445	-3.301311 9.034644
upper	14.3	3.00403	4.76	0.000	9.358811 19.24119

Bootstrapped standard errors are similar to their analytical counterparts. Even the 90% confidence interval for the treatment effect overlaps the value of zero. Finally, we try to tighten the bounds by using a covariate. The only one available is `age`, which we have to transform into a categorical variable. Here we choose three age categories, with each category having roughly the same number of observations.

```
. _pctile age, percentiles(33 66 99)

. generate agecat = recode(age,r(r1),r(r2),r(r3))

. leebounds studytime activedrug, select(died) cieffect tight(agecat)

Tightened Lee (2009) treatment effect bounds

Number of obs.                    =     48
Number of selected obs.           =     31
Number of cells                   =      3
Overall trimming porportion       =   0.5489
Effect 95% conf. interval         : [ 0.1028   19.6897]
```

studytime	Coef.	Std. Err.	z	P>|z|	[95% Conf. Interval]
activedrug					
lower	7	4.155293	1.68	0.092	-1.144225 15.14423
upper	12.55556	4.29805	2.92	0.003	4.131531 20.97958

Tightening yields much narrower bounds for the treatment effect. Indeed, with the `tight()` option specified, the 95% treatment-effect confidence interval does not include the value zero.

```
. heckman studytime activedrug i.agecat, select(died = activedrug i.agecat)

Iteration 0:   log likelihood = -125.92466
Iteration 1:   log likelihood = -125.57366
Iteration 2:   log likelihood = -125.47902
Iteration 3:   log likelihood = -125.47786
Iteration 4:   log likelihood = -125.47786
```

Heckman selection model		Number of obs	=	48
(regression model with sample selection)		Censored obs	=	17
		Uncensored obs	=	31
		Wald chi2(3)	=	7.86
Log likelihood = -125.4779		Prob > chi2	=	0.0489

	Coef.	Std. Err.	z	P>|z|	[95% Conf. Interval]
studytime					
activedrug	9.84113	4.180875	2.35	0.019	1.646767 18.03549
agecat					
58	-3.662228	3.57097	-1.03	0.305	-10.6612 3.336745
67	-6.869451	3.478465	-1.97	0.048	-13.68712 -.0517839
_cons	12.1558	2.635776	4.61	0.000	6.989771 17.32182
died					
activedrug	-1.945951	.542586	-3.59	0.000	-3.0094 -.8825016
agecat					
58	.9610446	.5307829	1.81	0.070	-.0792707 2.00136
67	.8677531	.5754751	1.51	0.132	-.2601574 1.995664
_cons	1.1392	.5072488	2.25	0.025	.1450108 2.13339
/athrho	.0328594	.5928681	0.06	0.956	-1.129141 1.19486
/lnsigma	1.948617	.1272405	15.31	0.000	1.69923 2.198004
rho	.0328476	.5922284			-.810725 .8320799
sigma	7.018972	.8930975			5.469734 9.007014
lambda	.2305562	4.158752			-7.920448 8.381561

```
LR test of indep. eqns. (rho = 0):   chi2(1) =      0.00   Prob > chi2 = 0.9555
```

Finally, to compare `leebounds` with fitting a conventional sample selection model, we run Stata's `heckman` command using the same data. Here the variable `agecat` enters both equations of the Heckman model as a control. `heckman` yields a point estimate centered between the lower and the upper bound estimated by `leebounds` with `agecat`

used for tightening. However, the result from `heckman` is imprecise. The estimated confidence interval for the treatment effect is almost as wide as its counterpart from `leebounds`. Hence, in this particular example, the restrictive assumptions inherent to the Heckman selection model do not pay off in terms of substantially reduced uncertainty about the size of the treatment effect. One reason for this may be that the data lack variables that explain selection into the estimation sample while not being directly linked to the outcome variable `studytime`.

5   Conclusion

In this article, I introduced the new command `leebounds`, which implements Lee's (2009) treatment-effect bounds for data with random assignment of treatment that suffer from nonrandom sample selection. In addition to calculating point estimates for the bounds, the command accommodates the calculation of confidence intervals for the treatment effect and tightened bounds on the basis of covariates. `leebounds` complements the contributions of Beresteanu and Manski (2000) and Palmer et al. (2011), who have made other bounds estimators available to Stata users that, unlike Lee's estimator, deal with selection into treatment and imperfect compliance with a randomly assigned treatment.

6   Acknowledgments

7   References

Ahn, H., and J. L. Powell. 1993. Semiparametric estimation of censored selection models with a nonparametric selection mechanism. *Journal of Econometrics* 58: 3–29.

Augurzky, B., T. K. Bauer, A. R. Reichert, C. M. Schmidt, and H. Tauchmann. 2012. Does money burn fat? Evidence from a randomized experiment. Ruhr Economic Papers No. 368. http://hdl.handle.net/10419/67304.

Balke, A., and J. Pearl. 1997. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association* 92: 1171–1176.

Beresteanu, A., and C. F. Manski. 2000. Bounds for Stata: Draft version 1.0. Software at http://faculty.wcas.northwestern.edu/%7ecfm754/bounds_stata.zip, documentation at http://faculty.wcas.northwestern.edu/%7ecfm754/bounds_stata.pdf.

Cawley, J., and J. A. Price. 2013. A case study of a workplace wellness program that offers financial incentives for weight loss. *Journal of Health Economics* 32: 794–803.

Grasdal, A. 2001. The performance of sample selection estimators to control for attrition bias. *Health Economics* 10: 385–398.

Heckman, J. J. 1976. The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models. *Annals of Economic and Social Measurement* 5: 475–492.

———. 1979. Sample selection bias as a specification error. *Econometrica* 47: 153–161.

Horowitz, J. L., and C. F. Manski. 2000. Nonparametric analysis of randomized experiments with missing covariate and outcome data. *Journal of the American Statistical Association* 95: 77–84.

Ichimura, H., and L.-F. Lee. 1991. Semiparametric least squares estimation of multiple index models: Single equation estimation. In *Nonparametric and Semiparametric Methods in Econometrics and Statistics*, ed. W. A. Barnett, J. Powell, and G. Tauchen, 3–49. Cambridge: Cambridge University Press.

Imbens, G. W., and C. F. Manski. 2004. Confidence intervals for partially identified parameters. *Econometrica* 72: 1845–1857.

Lee, D. S. 2009. Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Review of Economic Studies* 76: 1071–1102.

Manski, C. F. 1990. Nonparametric bounds on treatment effects. *American Economic Review* 80: 319–323.

———. 1994. The selection problem. In *Advances in Econometrics, Sixth World Congress*, ed. C. Sims, vol. 1, 143–170. Cambridge: Cambridge University Press.

———. 1995. *Identification Problems in the Social Sciences*. Cambridge, MA: Harvard University Press.

———. 1997. Monotone treatment response. *Econometrica* 65: 1311–1334.

Manski, C. F., and J. V. Pepper. 2000. Monotone instrumental variables: With an application to the returns to schooling. *Econometrica* 68: 997–1010.

Palmer, T. M., R. R. Ramsahai, V. Didelez, and N. A. Sheehan. 2011. Nonparametric bounds for the causal effect in a binary instrumental-variable model. *Stata Journal* 11: 345–367.

Puhani, P. A. 2000. The Heckman correction for sample selection and its critique. *Journal of Economic Surveys* 14: 53–68.

About the author

Harald Tauchmann is a professor of health economics at the Friedrich-Alexander-Universität Erlangen-Nuremberg, a guest researcher at Rheinisch-Westfälisches Institut für Wirtschaftsforschung (RWI), Essen, Germany, and a research fellow at the Centre of Health Economics Research (CINCH), Essen, Germany. His research interests include health economics, applied econometrics, and statistical methods.