

Iatrogenic Specification Error: A Cautionary Tale of Cleaning Data

Christopher R. Bollinger, *University of Kentucky*

Amitabh Chandra, *Dartmouth College, Institute for the Study of Labor (Bonn), and National Bureau of Economic Research*

It is common practice to use sensible rules of thumb for cleaning data. Measurement error is often the justification for removing (trimming) or recoding (winsorizing) observations where the dependent variable has values that lie outside a specified range. We consider a general measurement error process that nests many plausible models. Analytic results demonstrate that winsorizing and trimming are solutions for a narrow class of error processes. Indeed such procedures can induce or exacerbate bias. Monte Carlo simulations and empirical results demonstrate the fragility of cleaning. Even on root mean square error criteria, we cannot find generalizable justifications for these procedures.

I. Introduction

Empirical researchers frequently use simple rules of thumb to clean data on the basis of the dependent variable. As an example, researchers analyzing survey reports of wages and salaries often remove observations

We are grateful to Mark Berger, Dan Black, Charles Brown, Joel Horowitz, Chuck Manski, Douglas Staiger, Bo Honoré, participants at the Joint Statistical Meetings, seminar participants at the University of Kentucky, Princeton University, and Northwestern University for many useful comments. Chandra acknowledges the support of NIA P01 AG19783-02. Address comments to the authors at crboll@uky.edu or amitabh.chandra@dartmouth.edu.

[*Journal of Labor Economics*, 2005, vol. 23, no. 2]
© 2005 by The University of Chicago. All rights reserved.
0734-306X/2005/2302-0002\$10.00

whose value for the hourly wage is below the minimum wage or above some prespecified cutoff: sample exclusions based on wages can be found in Bound and Freeman (1992), Card and Krueger (1992), Katz and Murphy (1992), Juhn, Murphy, and Pierce (1993), and Buchinsky (1994). We cite these authors to illustrate the endorsement of this practice by leading scholars in the field. As we demonstrate in this article, the intuitively appealing strategy of discarding certain observations is not costless and can introduce specification error in cases where no error previously existed. Given the fact that the inconsistency is exacerbated by the analyst's actions, we borrow a term from the medical literature and term this form of bias "iatrogenic" specification error. In the medical literature an iatrogenic event is an adverse reaction to a well-intentioned treatment initiated by a physician, and we believe that parameter inconsistency that is caused by the analyst's well-intentioned actions shares the same features of physician-induced complications.

Given the widespread acceptance of this practice, the topic of "robust" estimation has received the attention of both economists and statisticians. In one of the earliest formal examinations, Stigler (1977) poses an interesting question: how much have methods such as trimming, winsorizing, the Edgeworth average, or Tukey's Biweight reduced the bias in the laboratory estimation of physical constants such as the speed of light or the density of the earth? Stigler concludes that the 10% trimmed mean, the smallest trimming amount considered in his study, is the most reliable estimator. In this he echoes the famous mathematician Legendre who recommended deleting those observations with errors "too large to be admissible" (1070). Stigler looks at the role of measurement error in the physical sciences; the error process may be vastly different in the social sciences where economic agents may have strategic or cultural incentives to inflate or deflate their reports. In the econometrics literature, Angrist and Krueger (2000) apply trimming and winsorizing techniques to the matched employer-employee data from Mellow and Sider (1983). When they trim both the employer and employee wage data, they find that the correlation between the two measures improves. Interestingly, this result does not hold for reports of hours worked. On the basis of this finding they conclude that "a small amount of trimming could be beneficial" (Angrist and Krueger [2000], 1349). Their prescription, which summarizes the intuition and current practice of most analysts, may be summarized as: "Loosely speaking, winsorizing the data is desirable if the extreme values are exaggerated versions of the true values, but the true values still lie in the tails. Truncating the sample is more desirable if the extremes are mistakes that bear no resemblance to the true values" (1349).

We examine this practice in detail here, using wages and earnings as a motivating example, though the results are likely to apply to other errors of measurement in survey data. We posit a general model of response

error in the dependent variable of a linear regression model and characterize the effect of different cleaning techniques on the estimated coefficients. We demonstrate, both analytically as well as through the use of simulations, that in general there is no reason to believe that removing what are referred to as obvious errors in the dependent variable reduces bias. This is similar in spirit to the findings of Hyslop and Imbens (2001), who examine instrumental variables approaches to solving the measurement error problem and find that they only apply to very specific measurement error processes. Our work is most closely related to that of MacDonald and Robinson (1985), who consider Bayesian estimation of an error components model in panel data when one of the error components is measurement error. They explicitly show that trimming can be thought of as an extremely dogmatic prior belief. Our article differs in that we explicitly consider a general error process and we do not consider a panel setting. Moreover we discuss an optimal trimming approach and other classical approaches to estimation. We demonstrate that the results in Stigler (1977) do not necessarily carry over in a regression framework. Indeed, trimming or winsorizing can bias coefficient estimates by as much as 10%–30%, and in many cases either induces bias that did not previously exist, or exacerbates the bias due to measurement error. The intuition for our result is simple: assuming that the researcher trims or winsorizes the data based on a lower bound of c or an upper bound of C , it will be shown that cleaning creates selection bias, and this is generally worse than the effects of measurement error in the dependent variable.

Our article is organized as follows: Section II describes identification with general measurement error in the dependent variable. We use the linear projection of the mismeasured dependent variable onto the covariates to derive analytical results. In Section III, we generalize the use of this projection to consider three specific models of measurement error (additive white noise, linear transformation, and the contaminated data process) that are found in social science data. Section IV examines the theoretical implications of trimming the data on bias as well as the asymptotic variances of the coefficients. We prove that only in highly specialized cases, unlikely to be found in social science data, does cleaning reduce bias. In these cases, we demonstrate that the information necessary to reduce bias leads to a simpler correction that requires fewer assumptions. This section also demonstrates that trimming will not necessarily reduce standard errors. Section V presents simulation results for the cases considered analytically. We generate quasi-simulated data from the 1990 U.S. Decennial Census to study the properties of winsorizing and trimming in a multivariate context. Finally, we present an empirical example from the March 2001 Current Population Survey (CPS). These simulations and examples support the results of the earlier two sections. Section VI

provides concluding comments. The appendix to this article provides detailed mathematical proofs and also considers the effects of winsorizing on bias and efficiency.

II. General Measurement Error in the Dependent Variable

To evaluate the widespread practice of “cleaning” data as described above we consider a general model for measurement error processes in the dependent variable. To keep the analysis simple, we focus on a linear regression model as the underlying structural model of interest to the researcher. Assume that the relationship between the “true” dependent variable and the covariate is described by

$$y_i^* = x_i'\beta + u_i. \quad (1)$$

Our maintained assumption is that the analyst is interested in estimates of β .¹ We assume a general process that relates the true value y_i^* to the observed value y_i :

$$y_i = h(y_i^*, \varepsilon_i). \quad (2)$$

There are six assumptions that are made for identification of the vector β and its associated covariance matrix:²

ASSUMPTION 1. $E(u_i|x_i) = 0$.

ASSUMPTION 2. x_i is a vector random variable with mean 0 and full rank second moment matrix V_x .

ASSUMPTION 3. Random sampling is employed.

ASSUMPTION 4. $h(.,.)$ has finitely many discontinuities.

ASSUMPTION 5. ε_i is independent of (y_i^*, x_i, u_i) .

ASSUMPTION 6. $\text{Cov}(y_i, y_i^*) > 0$.

Regardless of the process in equation (2), one summary of the joint dis-

¹ If the analyst is not interested in β per se, but other features of the joint distribution between y and x such as $\text{Cov}(y, x)$ or $\text{Var}(y|x)$, then it is possible that our results do not apply. Further analysis is required to understand the applicability of trimming, winsorizing, or even rescaling for this class of problems. We thank an anonymous referee for suggesting this caveat.

² The mean independence assumption is stronger than necessary for identification of the vector β , but it allows for a simpler analysis below. The zero mean for x_i is the usual normalization. The fourth assumption is necessary for moments to be well defined, and assumption 6 simply ensures that the measurement error process is not so perverse that y_i is uninformative about y_i^* (covariance of zero) or that y_i and y_i^* are negatively related. Indeed, the necessary condition would simply be that the sign of the covariance were known and that the covariance is not zero. The fifth assumption is the strongest one. It implies that the measurement error process is independent of x_i and u_i except through y_i^* and insures that $f(y_i|y_i^*) = f(y_i|y_i^*, x_i, u_i)$.

tribution of y_i and y_i^* is the population linear projection of y_i on y_i^* :

$$y_i = \delta + \gamma y_i^* + e_i. \quad (3)$$

Here, $\delta = E(y_i)$, $\gamma = [\text{Cov}(y_i, y_i^*)]/[V(y_i^*)]$, and $E(e_i) = E(e_i y_i^*) = 0$. The linear projection is not a statement about the data generating process, but rather a summary measure of the joint distribution of (y_i, y_i^*) . The actual measurement process, as defined by $h(y_i^*, \varepsilon)$, may be substantially more complicated. Assumption 6 insures that $\gamma > 0$.

The researcher is only able to observe (y_i, x_i) . Substituting equation (1) into equation (3) yields

$$y_i = \delta + x_i' \beta \gamma + \gamma u_i + e_i. \quad (4)$$

Assumption 5 insures that $\text{Cov}(x_i, \gamma u_i + e_i) = 0$ and $E(\gamma u_i + e_i) = 0$. This defines the population linear projection of y_i on x_i ,

$$y_i = a + x_i' b + \eta_i, \quad (5)$$

where $b = \gamma \beta$, $a = \delta$, and $\eta_i = \gamma u_i + e_i$.

Therefore, the ordinary least squares (OLS) regression of y_i on x_i yields a consistent estimate of b which is proportional to β . The parameters of interest are identified up to an unknown scaling constant. This would imply that estimates of ratios of the parameters are consistent. In some settings, identification up to scale is considered sufficient. For example, in wage regressions the coefficients on years of education and years of labor market experience can be combined to consistently identify the relative return of experience to education, a fact that might be sufficient for estimation of schooling choices. In general, however, we assume the researcher is interested in recovering the parameters β . This suggests two important identification approaches: obtain information about the scaling constant γ or obtain information about one of the elements in β . While it may be possible to obtain some consistent estimate of one element in β from auxiliary regressions or economic theory, the use of validation data may permit estimation of γ . Bound and Krueger (1991) and Bollinger (1998) have examined the structure of response error when y is the natural log of annual labor market earnings using Social Security income data matched to the Current Population Survey. They find a point estimate for γ is 0.90. This estimate could be used in log wage models to rescale slope coefficients to account for measurement error.³

³ The results in Bound and Krueger (1991) and Bollinger (1998) rely on estimates from the the 1977 and 1978 CPS-SSA matched files. It is possible that the structure of measurement error has changed over time, thereby reducing the applicability of the rescaling option since it hinges critically on knowledge of the correct γ . Estimation of γ is further complicated by the fact that low-earning respondents in the SSA data may be reporting their CPS earnings correctly. Examining these hypotheses is an important avenue for future research.

III. Specific Measurement Error Models

The above analysis holds for general examples of measurement error. In this section, we present three special cases of the above model. These cases are chosen because they are commonly examined or supported in the literature or lead to results in the context of this article that are of interest. Assumptions 1–6 are maintained, and additional assumptions are also imposed.

A. Additive White Noise

The classical measurement error process is often assumed: $y_i = y_i^* + \varepsilon_i$, and $E(\varepsilon_i) = 0$. Indeed, the error term in regression models is often motivated as measurement error. The parameters of the linear projection of y on y^* are $\gamma = 1$, $\delta = a = 0$, and the least squares estimates are consistent for the parameters of interest β . In this model, if y_i^* were hourly wages, it would be possible to have observations less than the minimum wage (or for that matter even negative observations) and observations above whatever threshold is deemed as a maximum. While it may be true that observations outside the acceptable region are measured with error, observations within the acceptable region are also measured with error. However, as is well known, classical measurement error does not lead to any bias—the estimated standard errors are inflated but all statistical tests remain valid. Researchers will often point out that “standard errors are too large” because of the additional measurement error. Standard errors are meant to capture the variation in estimates due to differences across samples. As long as the data generating process does not change, the sampling variation of the estimating coefficients will depend on the variation in both the structural model and the variation in the error model. Hence, estimates of the standard error are not biased, but rather reflect the variation across samples for this data generating process.

B. Linear Measurement Error

A second case is where the data generating process is linear: $y_i = d + \gamma y_i^* + \varepsilon_i$. Here, the parameters in the linear projection of y on y^* are $\gamma = g$ and $\delta = d$, and the model can either have $\gamma > 1$ or $\gamma < 1$. The data generating process can lead to observations outside the “acceptable” range. Because of the values of δ and the distribution of ε_i , even if $\gamma < 1$, it is quite possible to have both observations that are “too high” and observations that are “too low.” Empirical work by Bollinger (1998) and Bound and Krueger (1991) supports the possibility that $\gamma < 1$. For example, using nonparametric regression on the 1978 CPS-SSA matched data, Bollinger (1998) estimates that γ is equal to 0.91 for men and 0.97 for women. He estimates the intercepts δ to be \$1,364 and \$211, respectively. Cognitive psychologists have noted that this model, with $\gamma < 1$, will arise when

respondents exhibit “regression to the mean.” If survey respondents give answers that try to make them appear “average,” then those below the mean report higher values, on average, while those above the mean report lower values, on average. Similarly, the hot deck procedure used by the census to impute earnings can also lead to a regression to the mean (Hirsch and Schumacher 2001). To our knowledge, no study has found any variable with $\gamma > 1$.

C. Contaminated Data

A third example is a simple contaminated sample: $y_i = (y_i^*) \times 1(\varepsilon_{1i} > \kappa) + (d + \varepsilon_{2i}) \times 1(\varepsilon_{1i} < \kappa)$. The term $1(\cdot)$ is the indicator function and $(\varepsilon_{1i}, \varepsilon_{2i})$ are mean zero and mutually independent. This model produces a mixture: with some probability $p = \Pr(\varepsilon_{1i} > \kappa)$, we observe the true variable y_i^* , while with probability $(1 - p)$ we observe only noise: $d + \varepsilon_{2i}$. This leads to a model where we have some correctly measured observations and some observations where the observed y has no relationship to the actual y^* . In this model $\gamma = p$ and $\delta = d(1 - p)$. Again, some observations may fall outside a given range, depending on the distribution of ε_{2i} and the value of d . An important implication of this model is that estimates of the slope parameter β can be obtained if an estimate of p is available. Horowitz and Manski (1995) note that the expectation of y^* given x cannot be bounded unless information about d is available. Our analysis does not contradict this, but rather points out that in a linear model, the slopes can be identified up to the contamination rate. In many cases researchers have a priori bounds for the contamination rate. The bounds on the contamination rate will yield trivial bounds for the slope coefficients. If $\underline{p} < p < \bar{p}$, then the elements of β , β_j , are bounded by $[\underline{b}_j/\underline{p}, \bar{b}_j/\bar{p}]$.

IV. Effect of Cleaning

We assume that the researcher truncates above and below the mean. The cleaning approaches we consider are defined by

$$\{y_i, x_i | c \leq y_i \leq C\} \quad (6)$$

for known constants (c, C) , such that $c < E(y) < C$. We compare the slopes obtained from a least squares projection of the cleaned y_i on x_i to those obtained from the uncleaned data regressed on the covariate (i.e., relative to b , the biased estimate of β from the uncensored data). Since the choice for the researcher is to clean or not clean, this is the relevant comparison. As noted in the section above, the slope b may be larger or smaller in magnitude than the true slope β .

A. Analytic Results: Trimmed Data

We first derive analytic results under the additional assumption of joint normality.

ASSUMPTION 7. (y_i, x_i) are jointly normal.

As Goldberger (1981) demonstrates, the slope vector from the least squares projection of y_i on x_i in the truncated sample (where observations above C and below c are discarded) is given by:

PROPOSITION 1. Under assumptions 1–7, the truncated slope b^* is attenuated relative to the slope b from the least squares projection in the full sample:

$$b^* = \left[\frac{\theta}{1 - (1 - \theta)\rho^2} \right] b \quad (7)$$

with

$$\theta = \frac{V(y_i | c \leq y_i \leq C)}{V(y_i)} \quad (8)$$

and

$$\rho^2 = \frac{b^2 \sigma_x^2}{V(y_i)}. \quad (9)$$

Proof. As Goldberger (1981) notes, $0 \leq \{\theta/[1 - (1 - \theta)\rho^2]\} \leq 1$.⁴ Clearly, if $\gamma \leq 1$, then the attenuation bias of the measurement error is exacerbated by the attenuation bias of the sample truncation. Hence, only if the researcher is certain that $\gamma > 1$ can truncation alleviate bias from measurement error. For this case, the optimal level is determined by finding (c, C) such that $\{\theta/[1 - (1 - \theta)\rho^2]\}\gamma = 1$. With two unknown terms and only one restriction, there are many solutions. Too little truncation will fail to fully correct for the bias, while too much will overcorrect. Therefore, selecting the trimming bounds on the basis of a priori values for the supports of y^* (as proxied by the 1 and 99 percentiles of the wage distribution, or trimming at the minimum wage) will not necessarily correspond to the optimal trimming rule. Other solutions exist as well. For example, if the analyst chooses $c = E(y) - c^*$ and $C = E(y) + c^*$, only

⁴ Since y_i is normally distributed, the variance of the doubly truncated distribution can be expressed (see Maddalla 1983) as:

$$V(y_i | c \leq y_i \leq C) = V(y) \left[1 + \frac{\left[\frac{\phi[(c - E(y))/V(y)] \phi[(c - E(y))/V(y)] - \phi[(C - E(y))/V(y)] \phi[(C - E(y))/V(y)]}{\Phi[(C - E(y))/V(y)] - \Phi[(c - E(y))/V(y)]} \right]^2}{\left[\frac{\phi[(c - E(y))/V(y)] - \phi[(C - E(y))/V(y)]}{\Phi[(C - E(y))/V(y)] - \Phi[(c - E(y))/V(y)]} \right]^2} \right]$$

the term c^* needs to be found in order to devise an optimal trimming rule. It is implicitly described in the next proposition.

PROPOSITION 2. Under assumptions 1–7 and $\gamma > 1$, an optimal trimming rule of the form $\{y_i, x_i | c \leq y_i \leq C\}$ with $c = E(y) - c^*$ and $C = E(y) + c^*$ may be derived implicitly as:

$$2 \left(\frac{c^*}{V(y)} \right) \left\{ \frac{\phi[c^*/V(y)]}{\Phi[c^*/V(y)] - \Phi[-c^*/V(y)]} \right\} = \frac{\gamma - 1}{\gamma - \rho^2}. \quad (10)$$

Proof. See the appendix.

Because the solution involves the cumulative distribution function of the standard normal distribution, there is no closed form expression. The optimal cleaning depends on the variance of the observed y , the correlation between y and x , and γ . The right hand side of (10) is increasing in γ and the left hand side of (10) is decreasing in c^* . Therefore, as γ increases, the truncation points must move closer to the mean and the data must be truncated more heavily. In order to use this approach a number of highly restrictive assumptions must be met. First, the data must be jointly normally distributed. Any discrete variables in x_i will violate this assumption. Second, the measurement error process must result in a projection equation for y_i on y_i^* where $\gamma > 1$. Finally, specific information on γ must be obtained in order to arrive at a truncation rule. Ad hoc cleaning approaches that ignore the strong nature of these assumptions may be of little value in reducing the bias in b .

The variance of the measurement error is often used as a measure of the severity of the error. Since $b^* = [\theta/1 - (1 - \theta)\rho^2]b$ under trimming, the derivative of $\theta/[1 - (1 - \theta)\rho^2]$ with respect to σ_e^2 reveals how the truncation bias is affected by the measurement error. This result motivates the next proposition.

PROPOSITION 3. Under assumptions 1–7, the absolute value of the difference between elements of b and b^* becomes larger, since $(\partial/\partial\sigma_e^2)\{\theta/[1 - (1 - \theta)\rho^2]\} < 0$.

The above proposition may appear to be prima facie counterintuitive, but the intuition behind it is simple. As the measurement error becomes more severe (as measured by its variance), trimming is more likely to result in deleting observations based on the regression error instead of the measurement error, thereby causing sample-selection bias.

B. Does Trimming Reduce Standard Errors?

A second reason sometimes cited for trimming is the reduction of standard errors. To examine this procedure more rigorously we begin by noting that the truncation will introduce heteroskedasticity by reducing the variance of errors in the tails of the distribution. Therefore, the asymptotic variance of the estimated slope from the truncated data can be

derived from the expression $AV(\hat{b}^*) = Q^{-1}E[(y_i - x_i'b^*)^2 x_i x_i' | c \leq y_i \leq C]Q^{-1}$, where $Q = E(x_i x_i' | c \leq y_i \leq C)$.

In the appendix we demonstrate that this expression for asymptotic variance may be written as the sum of two terms:

$$AV(\hat{b}^*) = \theta V(y_i) \left[1 - \left(\frac{\theta}{1 - (1 - \theta)\rho^2} \right) \rho^2 \right] Q^{-1} + Q^{-1} E\{[x_i'b^* - m(x_i)]^2 x_i x_i' | c \leq y_i \leq C\} Q^{-1}. \quad (11)$$

The size of the leading term is indeterminate relative to its full-sample OLS counterpart.⁵ This fact is in contrast to similar comparisons for the mean (and consequently the results in Stigler [1977]). For the trimmed mean, the term $\theta V(y_i) < V(y_i)$ and, therefore, trimming necessarily reduce the variance in the leading term. Here, the comparison is not so straightforward. The second term is due to heteroskedasticity from the trimming and is necessarily positive definite. Hence, even if the leading term is smaller than the variance of the OLS estimate on the full sample (in a positive-definite sense), the second term may reverse or at least mitigate that difference.

Finally, an often overlooked reason for why trimming may not reduce standard errors is the effect of truncation on sample size. The finite sample variance of \hat{b}^* is given by

$$V(\hat{b}^*) = \frac{AV(\hat{b}^*)}{N \times (\Phi\{[C - E(y_i)]/V(y_i)\} - \Phi\{[c - E(y_i)]/V(y_i)\})}. \quad (12)$$

The term $\Phi\{[C - E(y_i)]/V(y_i)\} - \Phi\{[c - E(y_i)]/V(y_i)\} < 1$ measures the proportion of the sample discarded from the truncation rule, and increases in this proportion will raise estimates of the finite sample variance.

In conclusion, we find that comparisons between the variance of the truncated estimates and the variance of the full sample estimates is complicated and depends on the underlying parameters of the joint distribution. It is not possible to sign this difference, even under normality. In fact, the simulations below show little or no effect of trimming on standard errors.

⁵ The leading term is comparable to the asymptotic variance expression for the OLS estimate in the full sample: $V(y_i)(1 - \rho^2)E(x_i x_i')^{-1}$. Furthermore, $\theta V(y_i)(1 - \{\theta/[1 - (1 - \theta)\rho^2]\}\rho^2) \leq V(y_i)(1 - \rho^2)$. However, the difference $E(x_i x_i')^{-1} - E(x_i x_i' | c \leq y_i \leq C)^{-1}$ is necessarily positive semidefinite, since the variance of x_i in the truncated sample cannot be larger than the variance of the full sample (a sufficient condition here is joint normality; see Goldberger [1981]). Hence, a comparison of the leading terms is indeterminate.

C. Analytic Results: Winsorized Data

Rather than truncation, winsorized data are censored at the points c and C . Here, no observations are removed, but values of y_i outside of the region (c, C) are transformed as:

$$y_i^w = \begin{cases} C & \text{if } y_i \geq C \\ y_i & \text{if } c < y_i < C \\ c & \text{if } y_i \leq c \end{cases} \quad (13)$$

In the appendix we show that the analytical results from winsorizing are similar to those obtained for trimming. The empirical results suggest winsorizing has less of an impact on the slope coefficients (relative to OLS) than truncation. Again, if $\gamma > 1$, an optimal choice of winsorizing points is available, but of course γ is unknown. Ascertaining the effect of winsorizing on the size of the standard errors is conceptually similar to the effect of trimming on standard errors, but algebraically more difficult. We provide a discussion of this point in the appendix. One advantage winsorizing has over trimming is that the penalty of lost data does not effect the expression for the finite sample variance. But overall, winsorizing and trimming have similar effects on both slope estimates and asymptotic variance.

D. Cleaning Data in the General Case

The results derived in the previous section rely heavily upon normality. However, as Goldberger (1981) demonstrates, without normality some coefficients may be attenuated by truncation, while others may be inflated. Clearly, as a theoretical matter, truncation or winsorizing cannot be relied upon to adjust slope coefficients for the bias in general. In contrast, the results from Section II were derived under much weaker conditions.

V. Results

To suggest more general results, we present a set of Monte Carlo simulations. We draw data from the 1990 PUMS and estimate the returns to schooling, treating estimates from the full PUMS file as population parameters. The results of different cleaning procedures are gauged against these parameters. This allows for a complicated measurement model, with relationships similar to those found in typical economic data.⁶

⁶ Our working paper (Bollinger and Chandra 2003) provides more detailed Monte Carlo simulations for the univariate and multivariate case, and data from the normal, uniform, and log-normal distributions. Of note are the results for the realistic case of $\gamma = 0.9$. We note that cleaning procedures are always dominated by not cleaning the data. Even though the bias from not cleaning the data is a little over 10%, the bias from trimming is uniformly greater. One percent winsorizing or the use of median regression are neutral rules with respect to point

A. Evidence from U.S. Decennial Census Data

We begin with evidence from quasi-simulated data drawn from the PUMS samples of the 1990 U.S. Decennial Census. The advantage of these data is that they provide a complex multivariate distribution for analysis. We study the problem of estimating the returns to schooling, using a standard “Mincerian specification” (i.e., $\ln Wage = \beta_0 + \beta_1 Schooling + \beta_2 Experience + \beta_3 Experience^2 + \beta_4 Black + \mu$) to describe the relationship between hourly wages, years of schooling, race, and potential experience. We first select prime-age men who are working full-time in a nonagricultural industry. We remove individuals who earn less than the minimum wage. The resulting 346,900 observations constitute a sample that closely simulates the ideal population distribution assumed by many researchers. Column 1 of table 1 presents mean and standard deviation parameters for this pseudopopulation. Black men comprise 8.3% of the population, the mean years of potential experience is 17.67, and the mean years of schooling is 13.37. The regression parameters generated by an OLS regression on all 346,900 observations are reported in column 2 of table 1. For our simulations, we randomly draw samples of size 1,000 from the pseudopopulation of 346,900 observations. Ordinary least squares performed on these data (without any cleaning) are reported in the third column of the table. In columns 4–8, we report the effect of alternative cleaning procedures when no measurement error has been added to the PUMS data.⁷ The idea of cleaning data with no error might strike the reader as a peculiar exercise. Our motivation for doing so is to demonstrate that cleaning procedures are not benign and can introduce significant bias when they are not required; alternatively, if the degree of contamination is low, the iatrogenic error from cleaning data may be substantial. In this situation, the cleaning procedures do not generally perform better than not cleaning the data. In general, the root mean square error (RMSE) from the cleaning procedures (including median regression) is greater than that from doing nothing. Whereas a 1% trimming rule

estimates, but both procedures are dominated by not cleaning the data on the basis of an RMSE criteria. When $\gamma > 1$, a 1% trimming rule clearly dominates not cleaning the data. Before this conclusion is embraced too quickly by practitioners, we raise two important caveats: First, even though 1% trimming works, 5% trimming is much worse than not cleaning the data; the optimal trimming rule is therefore not a known constant and small perturbations from the optimal truncation will generate large biases relative to not cleaning the data. In fact, the “best rule” for the case of $\gamma > 1$ would be to use 5% winsorizing. Second, we reiterate the difficulty in being able to justify a behavioral model for why γ would exceed one.

⁷ In the last column of the table, we report results from performing median regression. While they are not explicitly studied in our article, we include these estimates because several readers of our article argued that it may be viewed as an alternative to trimming and winsorizing.

Table 1
Effect of Cleaning Procedures on Uncorrupted Data, Evidence from the Returns to Schooling in 1990 PUMS Data

	Mean	Population β	No Cleaning	Trim 1%	Trim 5%	Wins 1%	Wins 5%	Median
Schooling	13.3741	.092	.0918	.0878	.0704	.0910	.0856	.1001
SE (years of schooling)	2.21890078	.0072	.0068	.0075	.0070	.0086
RMSE0078	.0083	.0227	.0076	.0095	.0118
Potential experience	17.6700	.0374	.0375	.0361	.0294	.0372	.0353	.0402
SE (schooling)	8.59760084	.0079	.0069	.0082	.0076	.0095
RMSE0084	.0080	.0106	.0082	.0079	.0099
Potential experience ² /100	3.8639	-.0535	-.0538	-.0522	-.0419	-.0535	-.0508	-.0565
SE (potential experience)	3.36430218	.0202	.0175	.0211	.0196	.0246
RMSE0218	.0203	.0210	.0211	.0198	.0248
Black (1 = yes)	.0831	-.1419	-.1416	-.1374	-.1072	-.1417	-.1339	-.1640
SE (black)	.27620597	.0544	.0492	.0579	.0531	.0697
RMSE0597	.0545	.0602	.0579	.0538	.0731
Constant8608	.8639	.9293	1.2376	.8746	.9649	.7280
SE1275	.1195	.1112	.1229	.1150	.1430

NOTE.—Dependent variable is ln hourly wage. PUMS data are restricted to white (non-Hispanic) and black men in the 1990 PUMS files of the Decennial Census who are ages 25–55 during the census reference week. Nonworkers and respondents with hourly wages less than \$3.35 in 1989 (the nominal value of the minimum wage) are deleted from the analysis. Col. 1 reports means and standard deviations for this sample of 346,900 individuals, and col. 2 reports the parameters from estimating the model $\ln Wage = \beta_0 + \beta_1 Schooling + \beta_2 Experience + \beta_3 Experience^2 + \beta_4 Black + u$ on this sample. Reported estimates in other columns are empirical sample moments from 1,000 replications each with a sample size of 1,000.

improves the estimation of the coefficients on experience and experience squared, it is inferior to not cleaning the data as regards the estimation of the coefficients on schooling and race. Together these results confirm those from the univariate case: no cleaning procedure is neutral when applied to already clean data.

Measurement error is added to the data in table 2. We select two values for the variance of this error using the results of Bound and Krueger (1991), who note that $\text{Var}(\ln Y) = 0.458$ and 0.529 with corresponding error variances 0.083 and 0.116 . This implies that the variance of the error is 18% and 22% of the total variation in $\ln Y$. Rogers, Brown, and Duncan (1993) find even higher implied estimates of the variance of the measurement error. Therefore, to study empirically relevant cases we simulate measurement error whose variance is $0.1 \text{Var}(Wage)$ and $0.3 \text{Var}(Wage)$. In order to keep the reported number of results manageable we only report results from the latter simulation, but note that the results from the former are quantitatively similar (see Bollinger and Chandra [2003] for details). In the case of additive white noise we find that trimming is once again dominated by not cleaning the data. A case can be made for a 1% winsorizing rule over not cleaning the data, but it is important to note that significant bias is introduced when the censoring rate is increased to 5%. For this case, least squares is found to be superior to median regression. When $\gamma = 0.9$ there is no cleaning procedure that strictly dominates OLS. A 1% winsorizing rule provides superior estimates on an RMSE criterion for many coefficients but simultaneously raises the bias on others. For example, the coefficients on Exp , Exp^2 , and Black all have lower RMSE when a 1% winsorizing rule is applied, but the coefficient on schooling has a larger RMSE at the same time. When $\gamma = 1.1$ winsorizing at 1% and 5% are preferred to doing nothing. In this case, trimming procedures dominate not cleaning the data on an RMSE criteria but can be worse in terms of the bias component.

As noted in previous sections, another identification approach is to rescale the estimates. The optimal trimming rule derived in Section IV requires both information about γ and normality. Clearly, if $\gamma < 1$, trimming or winsorizing will be dominated by the rescaling approach. Even when $\gamma > 1$, the amount of trimming or winsorizing necessary depends upon the variance of the errors (see Bollinger and Chandra 2003). Examining the first column of table 2 shows that knowledge of γ alone will be sufficient to arrive at consistent estimates.

Even if γ is not known, the rescaling results can be used to perform sensitivity analysis. For example, researchers might ask: how sensitive to different values of γ are the conclusions we draw from our OLS estimates? The robustness of these conclusions may be examined either by placing bounds on γ , as suggested in Manski (1995), or alternatively by asking what values of γ support the conclusions typically drawn (an approach

Table 2
Effect of Cleaning Procedures on Corrupted Data,
Evidence from the 1990 PUMS

	No Cleaning	Trim 1%	Trim 5%	Wins 1%	Wins 5%	Median
Error model 1:*						
Schooling	.0925	.0883	.0704	.0916	.0859	.1001
SE	.0081	.0074	.0068	.0078	.0072	.0091
RMSE	.0082	.0084	.0227	.0078	.0095	.0121
Potential experience	.0377	.0363	.0294	.0374	.0354	.0402
SE	.0084	.0078	.0070	.0081	.0075	.0096
RMSE	.0084	.0079	.0107	.0081	.0078	.0100
Potential experience ² /100	-.0539	-.0524	-.0419	-.0537	-.0510	-.0566
SE	.0215	.0201	.0178	.0209	.0193	.0248
RMSE	.0215	.0202	.0213	.0209	.0195	.0250
Black (1 = yes)	-.1426	-.1394	-.1097	-.1426	-.1346	-.1616
SE	.0639	.0578	.0513	.0616	.0564	.0730
RMSE	.0639	.0578	.0606	.0616	.0569	.0756
Constant	.8535	.9222	1.2396	.8667	.9621	.7281
SE	.1345	.1237	.1147	.1289	.1203	.1502
Error model 2:†						
Schooling	.0833	.0795	.0634	.0825	.0773	.0901
SE	.0074	.0067	.0062	.0071	.0066	.0082
RMSE	.0115	.0142	.0294	.0119	.0161	.0084
Potential experience	.0339	.0327	.0264	.0337	.0318	.0360
SE	.0075	.0071	.0061	.0073	.0068	.0087
RMSE	.0083	.0086	.0127	.0083	.0088	.0088
Potential experience ² /100	-.0485	-.0471	-.0375	-.0483	-.0459	-.0502
SE	.0195	.0182	.0158	.0190	.0175	.0225
RMSE	.0201	.0193	.0225	.0197	.0191	.0227
Black (1 = yes)	-.1276	-.1250	-.0976	-.1276	-.1205	-.1439
SE	.0576	.0518	.0458	.0554	.0507	.0641
RMSE	.0594	.0545	.0637	.0572	.0551	.0641
Constant	.7672	.8295	1.1165	.7789	.8651	.6558
SE	.1208	.1108	.1020	.1158	.1075	.1351
Error model 3:‡						
Schooling	.1018	.0972	.0775	.1008	.0945	.1104
SE	.0090	.0082	.0075	.0087	.0080	.0100
RMSE	.0132	.0097	.0164	.0123	.0084	.0208
Potential experience	.0415	.0400	.0325	.0412	.0390	.0441
SE	.0092	.0087	.0075	.0090	.0083	.0106
RMSE	.0101	.0090	.0089	.0097	.0085	.0125
Potential experience ² /100	-.0594	-.0577	-.0466	-.0591	-.0562	-.0616
SE	.0239	.0224	.0192	.0232	.0214	.0275
RMSE	.0246	.0227	.0205	.0238	.0216	.0286
Black (1 = yes)	-.1560	-.1527	-.1204	-.1562	-.1478	-.1789
SE	.0702	.0635	.0554	.0675	.0618	.0803
RMSE	.0716	.0644	.0594	.0690	.0621	.0884
Constant	.9380	1.0126	1.3612	.9522	1.0567	.7982
SE	.1480	.1352	.1238	.1418	.1321	.1656

NOTE.—Dependent variable is ln hourly wage. PUMS data are restricted to white (non-Hispanic) and black men in the 1990 PUMS files of the Decennial Census who are ages 25–55 during the census reference week ($n = 346,900$). Nonworkers and respondents with hourly wages less than \$3.35 in 1989 (the nominal value of the minimum wage) are deleted from the analysis, and measurement error is added to observed ln wage using the specified error models. Reported estimates are empirical sample moments from 1,000 replications each with a sample size of 1,000. The variance of $\ln(Wage) = 0.3144$.

* $\ln Wage = \ln Wage^* + e$; $\text{Var}(e) = 0.3 \times \text{Var}(Wage)$.

† $\ln Wage = 0.9 \ln Wage^* + e$; $\text{Var}(e) = 0.3 \times \text{Var}(Wage)$.

‡ $\ln Wage = 1.1 \ln Wage^* + e$; $\text{Var}(e) = 0.3 \times \text{Var}(Wage)$.

suggested in a similar context by Bollinger [2003]). Further, researchers may not have detailed information about γ but may have information about the likely range of γ . It is difficult to use that information for trimming and winsorizing, but it can be trivially used in a rescaling approach.

B. Empirical Example from the CPS

We also examine data-cleaning approaches using the March 2001 Current Population Survey. There are two measures of hourly wage that a researcher could exploit in these data. One is the hourly wage constructed from the reported annual earnings, weeks worked, and usual hours worked. The CPS also asks the actual hourly wage for workers who are paid hourly. Most researchers do not use this variable, as the resulting sample is smaller and only represents hourly wage workers. In this context, the two measures provide an interesting comparison to examine the implications of trimming as is typically practiced. Our sample consists of males, working full-time, year-round in nonagricultural positions who are not self-employed. In March 2001, we find 2,626 men who are full-time, year-round nonagricultural hourly wage workers.

The first column of table 3 presents the log wage regression on the reported hourly wage, while the second column uses the constructed hourly wage. We restricted the constructed hourly wage sample to contain only those workers who also reported an hourly wage, hence any difference between the two columns reflects only differences in the measurement of the hourly wage, rather than sample differences. One perspective with these results is that the first column represents the “true coefficients,” while the results in column 2 are biased due to response error. Interestingly, most of the coefficients in column 2 are larger in magnitude than their corresponding coefficients in column 1. The exception to this is the coefficient on bachelor’s degree. This is a case where trimming might be useful. However, the fact that the coefficient on bachelor’s is the exception demonstrates that it is difficult to find perfect cases.

If column 1 represents the “true coefficients,” then columns 3 and 4 represent attempts to correct column 2. Trimming at about half of the minimum wage is a common practice (see Angrist and Krueger 2000). Column 3 represents this approach. Another logical correction is to trim at the minimum wage; column 4 represents this approach. Comparing column 3 with columns 1 and 2, we find that the coefficients on experience and experience squared and black are largely unaffected by the trimming, and are still “too large.” The coefficient on less than high school has actually increased in magnitude, and made the bias worse. The coefficients on associate’s degree and graduate degree have declined in magnitude,

Table 3
Effect of Cleaning Procedures: Evidence from the March 2001 CPS

	Baseline Regressions		Constructed Hourly Wage		Reported Hourly Wage		Constructed Hourly Wage		Reported Hourly Wage	
	Reported Hourly Wage (1)	Constructed Hourly Wage (2)	Trim at Half Minimum Wage (3)	Trim at Minimum Wage (4)	Trim at Half Minimum Wage (5)	Trim at Minimum Wage (6)	Wins at Half Minimum Wage (7)	Wins at Minimum Wage (8)	Wins at Half Minimum Wage (9)	Wins at Minimum Wage (10)
Potential experience	.028 (.002)	.033 (.003)	.033 (.003)	.030 (.003)	.028 (.002)	.028 (.002)	.033 (.003)	.032 (.003)	.028 (.002)	.028 (.002)
Potential experience ²	-.045 (.005)	-.051 (.006)	-.051 (.006)	-.046 (.005)	-.046 (.005)	-.045 (.004)	-.052 (.006)	-.049 (.005)	-.045 (.005)	-.045 (.005)
Less than high school (1 = yes)	-.261 (.024)	-.287 (.031)	-.292 (.029)	-.280 (.028)	-.252 (.024)	-.253 (.023)	-.290 (.030)	-.285 (.028)	-.258 (.024)	-.256 (.023)
Associate's degree (1 = yes)	.170 (.028)	.191 (.037)	.187 (.034)	.166 (.032)	.168 (.028)	.168 (.027)	.191 (.035)	.184 (.034)	.170 (.028)	.169 (.028)
Bachelor's degree (1 = yes)	.254 (.028)	.242 (.037)	.235 (.034)	.211 (.032)	.253 (.028)	.270 (.027)	.238 (.035)	.232 (.034)	.254 (.028)	.258 (.028)
Graduate degree (1 = yes)	.377 (.058)	.590 (.076)	.575 (.070)	.575 (.066)	.375 (.057)	.450 (.056)	.584 (.072)	.577 (.069)	.7 (.058)	.398 (.056)
Black (1 = yes)	-.087 (.025)	-.135 (.033)	-.136 (.031)	-.113 (.030)	-.090 (.025)	-.093 (.024)	-.138 (.032)	-.130 (.030)	-.088 (.025)	-.090 (.025)
Constant	2.240 (.025)	2.197 (.032)	2.220 (.030)	2.282 (.029)	2.238 (.024)	2.241 (.024)	2.203 (.031)	2.230 (.029)	2.240 (.025)	2.240 (.024)
Observations	2,626	2,626	2,612	2,534	2,622	2,605	2,626	2,626	2,626	2,626

NOTE.—Sample is drawn from the March 2001 CPS. The sample consists of males who are not self-employed, working full-time, year-round in nonagricultural positions, who were paid hourly. Reported hourly wage refers to respondent's report of hourly wage (true wage); constructed hourly wage is constructed using annual earnings, hours, and weeks worked (and therefore constitutes a noisy measure of wage). The correlation between the two wage measures is 0.39 (and 0.54 in logs). The standard deviations for actual and constructed hourly pay are 7.4 and 12.5, respectively (.45 and 0.58 in logs). Reported hourly pay ranged from \$1 to \$99, whereas constructed hourly wages ranged from \$0.02 to \$184.13.

reducing the bias but not eliminating it. The coefficient on bachelor's degree has decreased in magnitude, increasing the bias in this coefficient.

Trimming at the minimum wage, represented by column 4, improves some coefficients but not others. The coefficients on experience and experience squared are now both smaller in magnitude and closer to the "ideal" column 1. The coefficient on less than high school has now declined in magnitude, but it is still somewhat larger than the coefficient in column 1. The coefficient on associate's degree has declined and is biased relative to the target in column 1; it now underestimates the magnitude. The coefficient on graduate degree has not changed any further and still overstates the target in column 1. The coefficient on black has declined in magnitude and is now closer to the coefficient in column 1 but is still larger in magnitude. The coefficient on bachelor's degree has declined further in magnitude, increasing the bias still further. The conclusion we take from this is that there is no clear advantage to trimming. While it certainly may reduce the bias for some estimates, it is simultaneously making other estimates worse. Since it is rare to have a target (as we do in this case), it is only through serendipity that one will pick the right trimming rule even if the researcher is only interested in one specific coefficient.

A second perspective on the estimates in columns 1 and 2 is that they both contain measurement error in the dependent variable. One would expect that trimmed versions of the two regressions would converge to some set of correct estimates. Columns 5 and 6 are trimmed versions of column 1. As one might expect, the reported wage has fewer observations below the trimming points than the constructed wage. In column 5 there is very little change in the coefficients. In column 6 there is little change the coefficients on experience, experience squared, or less than high school. The coefficient on bachelor's degree increases in magnitude. This is in sharp contrast to trimming the constructed wage, where the coefficient decreased in magnitude. The coefficients on graduate degree and black both increase in magnitude slightly. We conclude that it appears only serendipitous if any coefficients converge with trimming. If one is interested in the return to a college degree, trimming is likely to be undesirable. While if one is interested in the black-white wage gap, trimming at an even higher threshold may be desirable. The effects of trimming are unclear since we cannot even predict which direction the slope coefficients will change when we trim. Without a priori information, it is difficult or impossible to know if trimming has reduced bias, increased bias, or some of both.

VI. Conclusions

The common practice of cleaning data by removing observations where the dependent variable is larger or smaller than some threshold is used

in the hope of reducing the impact of measurement error. While this sounds sensible, it may make matters worse. Analytical results using normality demonstrate that cleaning strategies using truncation or winsorizing work only in the case where measurement error in the dependent variable results in an upward bias on the magnitude of coefficients. This case is not empirically supported by investigations of response error in earnings data. Under certain circumstances it may be possible to achieve an optimal cleaning strategy, but if the information necessary for that result were available, a simpler approach based on rescaling the estimated coefficients would work too.

Our empirical results demonstrate that a 1% winsorizing rule does not alter the results in a meaningful manner, but we note that this policy is generally dominated by doing nothing to the data. Still, winsorizing is clearly better than truncation. We caution, however, that small increments to this rule (e.g., to 5%) can dramatically increase bias and render the cleaning undesirable on mean squared error grounds. Two important extensions we hope to address in future work is cleaning based on covariates, and cleaning based on panel data where the researcher has multiple observations on the dependent variable. In both cases, it may be possible to develop cleaning procedures that exploit other information.

Appendix

Derivation of Equation (10)

To solve for the expression in (10), begin by noting that $E(y) = \delta + \gamma\alpha + \gamma\beta\mu_x$ and $V(y) = \gamma^2\beta^2\sigma_x^2 + \gamma^2\sigma_u^2 + \sigma_e^2$. Additionally, we simplify the analysis by only considering a symmetric truncation scheme, where $c = E(y) - c^*$ and $C = E(y) + c^*$, so that only c^* need be found. Consider first the expression for θ in this case:

$$\theta = 1 + \left(\frac{\{[c - E(y)]/V(y)\}\phi\{[c - E(y)]/V(y)\} - \{[C - E(y)]/V(y)\}\phi\{[C - E(y)]/V(y)\}}{\Phi[(C - \delta - \gamma\alpha - \gamma\beta\mu_x)/(\gamma^2\beta^2\sigma_x^2 + \gamma^2\sigma_u^2 + \sigma_e^2)] - \Phi[c - E(y)]/V(y)} \right) - \frac{\{\phi[c - E(y)]/V(y)\} - \{\phi[C - E(y)]/V(y)\}}{\{\Phi[C - E(y)]/V(y)\} - \{\Phi[c - E(y)]/V(y)\}}.$$

Substituting the symmetric expressions for c, C yields

$$= 1 + \frac{[-c^*/V(y)]\phi[-c^*/v(y)] - [c^*/V(y)]\phi[c^*/V(y)]}{\Phi[c^*/V(y)] - \Phi[-c^*/V(y)]} - \frac{\left[\frac{\phi[-c^*/V(y)] - \phi[-c^*/V(y)]}{\Phi[c^*/V(y)] - \Phi[-c^*/V(y)]} \right]^2}{\left[\frac{\phi[-c^*/V(y)] - \phi[-c^*/V(y)]}{\Phi[c^*/V(y)] - \Phi[-c^*/V(y)]} \right]}$$

$$= 1 - 2 \left(\frac{c^*}{V(y)} \right) \left(\frac{\phi[c^*/V(y)]}{\Phi[c^*/V(y)] - \Phi[-c^*/V(y)]} \right).$$

Next, noting that λ is observable and γ is assumed to be known, solve $\{\theta/[1 - (1 - \theta)\lambda]\}\gamma = 1$ for θ in terms of γ and λ :

$$\theta = \frac{1 - \lambda}{1 - \gamma}.$$

Substituting for θ and solving yields the implicit relationship expressed in equation (10):

$$2 \left(\frac{c^*}{V(y)} \right) \left(\frac{\phi[c^*/V(y)]}{\Phi[c^*/V(y)] - \Phi[-c^*/V(y)]} \right) = \frac{\gamma - 1}{\gamma - \lambda}.$$

Proof of Proposition 3

To show that the bias gets worse with more variance in ε , differentiate the bias term with respect to the variance of the measurement error:

$$\begin{aligned} & \frac{\partial}{\partial \sigma_\varepsilon^2} \left[\frac{\theta}{1 - (1 - \theta)\rho^2} \right] \\ &= \left\{ \frac{\partial \theta}{\partial \sigma_\varepsilon^2} [1 - (1 - \theta)\rho^2] - \theta \left[\rho^2 \frac{\partial \theta}{\partial \sigma_\varepsilon^2} - (1 - \theta) \frac{\partial \rho^2}{\partial \sigma_\varepsilon^2} \right] \right\} / [1 - (1 - \theta)\rho^2]^2. \end{aligned}$$

This term is negative if and only if the numerator is negative. Considering only the numerator and grouping similar derivatives yields

$$\frac{\partial \theta}{\partial \sigma_\varepsilon^2} (1 - \rho^2) + (1 - \theta) \theta \frac{\partial \rho^2}{\partial \sigma_\varepsilon^2}.$$

As noted in Goldberger (1981), both ρ^2 and θ are bounded in the unit interval. Inspection of the definition of ρ^2 clearly demonstrates that $\partial \rho^2 / \partial \sigma_\varepsilon^2 < 0$. Now, consider the definition of θ : inspection reveals that this has the truncation points standardized by the mean and variance of y . Hence, increasing σ_ε^2 is equivalent to increasing c and decreasing C for a truncated standard normal random variable. Since θ is also the ratio of the variance of the truncated standard normal to the variance of the untruncated standard normal (see Goldberger), increasing c and decreasing C will result in a lower variance for the truncated distribution and thus a lower θ . Hence, by inspection, $\partial \theta / \partial \sigma_\varepsilon^2 < 0$. Combined with the bounds on ρ^2 and θ , the result is established.

Does Trimming Reduce Standard Errors?

Under assumptions 1–7, the conditional distribution of $y_i|x_i$ is $N[x_i'b, V(y_i)(1 - \rho^2)]$. We can use the results in Goldberger (1981) to obtain an expression for the truncated second moment matrix as:

$$Q = E[x_i x_i' | c \leq y_i \leq C] = E[x_i x_i'] - (1 - \theta)E[x_i x_i'] b b' E[x_i x_i'].$$

The term $E[(y_i - x_i'b^*)^2 x_i x_i' | c \leq y_i \leq C]$ is considered by using the law of iterated expectations. The $E[(y_i - x_i'b^*)^2 | x_i, c \leq y_i \leq C]$ can be decomposed into the variation of y_i around the conditional mean in the truncated distribution, and the squared difference between the conditional mean and the linear projection term $x_i'b^*$. Doing so yields:

$$E[(y_i - x_i'b^*)^2 | x_i, c \leq y_i \leq C] = \theta V(y_i) \left[1 - \left(\frac{\theta}{1 - (1 - \theta)\rho^2} \right) \rho^2 \right] + [x_i'b^* - m(x_i)]^2,$$

where $m(x_i)$ is the conditional mean of y_i given x_i in the truncated sample. Combining these terms produces the equation in the article:

$$AV(\hat{b}^*) = \theta V(y_i) \left[1 - \left(\frac{\theta}{1 - (1 - \theta)\rho^2} \right) \rho^2 \right] Q^{-1} + Q^{-1} E\{[x_i'b^* - m(x_i)]^2 x_i x_i' | c \leq y_i \leq C\} Q^{-1}.$$

Analytic Results: Winsorized Data

When data are winsorized, no observations are removed, but values of y_i outside of the region (c, C) are transformed as follows:

$$y_i^{ws} = \begin{cases} C & \text{if } y_i \geq C \\ y_i & \text{if } c < y_i < C \\ c & \text{if } y_i \leq c \end{cases}$$

As in the section on trimming, let b represent the vector of full-sample (uncleaned) coefficients. Under bivariate normality, the winsorized coefficient vector, b^{**} is given by:

$$b^{**} = \left\{ \Phi \left[\frac{C - E(y_i)}{V(y_i)} \right] - \Phi \left[\frac{c - E(y_i)}{V(y_i)} \right] \right\} \times \left(1 - \frac{\left\{ \phi \left[\frac{c - E(y_i)}{V(y_i)} \right] - \phi \left[\frac{C - E(y_i)}{V(y_i)} \right] \right\}^2}{\left\{ \Phi \left[\frac{C - E(y_i)}{V(y_i)} \right] - \Phi \left[\frac{c - E(y_i)}{V(y_i)} \right] \right\}} \right) b. \quad (14)$$

Does Winsorizing Reduce Standard Errors?

The effects of winsorizing on the variance are derived similarly to the results for trimming. We start by noting that the $AV(\hat{b}^{**}) = E(x_i x_i')^{-1}$

$\times E[(y_i^w - x_i' b^{**})^2 x_i x_i'] E(x_i x_i')^{-1}$. Here, the term $E[(y_i^w - x_i' b^{**})^2 | x_i]$ can be broken into three terms:

$$\begin{aligned} E[(y_i^w - x_i' b^{**})^2 | x_i] &= \Phi\left[\frac{c - E(y_i)}{V(y_i)}\right] (c - x_i' b^{**})^2 + \left\{1 - \Phi\left[\frac{C - E(y_i)}{V(y_i)}\right]\right\} \\ &\quad \times (C - x_i' b^{**})^2 + \left\{\Phi\left[\frac{C - E(y_i)}{V(y_i)}\right] - \Phi\left[\frac{c - E(y_i)}{V(y_i)}\right]\right\} \\ &\quad \times E[(y_i - x_i' b^{**})^2 | x_i, c \leq y_i \leq C]. \end{aligned}$$

Combined with results from the previous section, we obtain

$$\begin{aligned} AV(\hat{b}^{**}) &= \left\{\Phi\left[\frac{C - E(y_i)}{V(y_i)}\right] - \Phi\left[\frac{c - E(y_i)}{V(y_i)}\right]\right\} \\ &\quad \times \left[\theta V(y_i) \left\{1 - \left[\frac{\theta}{1 - (1 - \theta)\rho^2}\right]\right\} E(x_i x_i')^{-1} Q E(x_i x_i')^{-1}\right] \\ &\quad + E(x_i x_i')^{-1} E\{[x_i' b^{**} - m^w(x_i)]^2 x_i x_i' | c < y_i < C\} E(x_i x_i')^{-1} \\ &\quad + \Phi\left[\frac{c - E(y_i)}{V(y_i)}\right] (E(x_i x_i')^{-1} E\{[(c - x_i' b^{**})^2 x_i x_i' | y_i \leq c\} E(x_i x_i')^{-1}) \\ &\quad + \left\{1 - \Phi\left[\frac{C - E(y_i)}{V(y_i)}\right]\right\} \\ &\quad \times (E[x_i x_i']^{-1} E\{[(C - x_i' b^{**})^2 x_i x_i' | y_i \geq C\} E(x_i x_i')^{-1}). \end{aligned}$$

Again, the comparison is difficult. Here, the first term will necessarily be smaller than the OLS expression. However, the second, third, and fourth terms are all positive definite. As in the trimming case, the impact on standard errors depends upon the parameters of the model.

References

- Angrist, Joshua D., and Alan B. Krueger. 2000. Empirical strategies in labor economics. In *Handbook of labor economics*, vol. 3A, ed. Orley Ashenfelter and David Card. Amsterdam: Elsevier Science.
- Bollinger, Christopher R. 1998. Measurement error in the current population survey: A nonparametric look. *Journal of Labor Economics* 16, no. 3:57–71.
- . 2003. Measurement error in human capital and the black-white wage differential. *Review of Economics and Statistics* 85:578–87.
- Bollinger, Christopher, and Amitabh Chandra. 2003. Iatrogenic specification error. Technical Working Paper no. 289, National Bureau of Economic Research, Cambridge, MA.
- Bound, John, and Richard Freeman. 1992. What went wrong? The erosion

- of relative earnings and employment among black men in the 1980s. *Quarterly Journal of Economics* 107:201–32.
- Bound, John, and Alan B. Krueger. 1991. The extent of measurement error in longitudinal earnings data: Do two wrongs make a right? *Journal of Labor Economics* 9:1–24.
- Buchinsky, Moche. 1994. Changes in the U.S. wage structure, 1963–1987: Application of quantile regression. *Econometrica* 62:405–58.
- Card, David, and Alan B. Krueger. 1992. School quality and black-white relative earnings: A direct assessment. *Quarterly Journal of Economics* 107:151–200.
- Goldberger, Arthur S. 1981. Linear regression after selection. *Journal of Econometrics* 15:357–66.
- Hirsch, Barry T., and Edward J. Schumacher. 2001. Match bias in wage gap estimates due to earnings imputations. Unpublished manuscript, Trinity College, San Antonio, TX.
- Horowitz, Joel L., and Charles F. Manski. 1995. Identification and robustness with contaminated and corrupted data. *Econometrica* 63: 281–302.
- Hyslop, Dean R., and Guido W. Imbens. 2001. Bias from classical and other forms of measurement error. *Journal of Business and Economic Statistics* 19:475–81.
- Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce. 1993. Wage inequality and the rise in the returns to skill. *Journal of Political Economy* 101:410–42.
- Katz, Lawrence, and Kevin M. Murphy. 1992. Changes in relative wages, 1963–1987. *Quarterly Journal of Economics* 107:35–78.
- MacDonald, Glenn M., and Chris Robinson. 1985. Cautionary tales about arbitrary deletion of observations; or, throwing the variance out with the bathwater. *Journal of Labor Economics* 3:124–52.
- Maddala, G. S. 1983. *Limited dependent and qualitative variables in econometrics*. Cambridge: Cambridge University Press.
- Manski, Charles F. 1995. *Identification problems in the social sciences*. Cambridge, MA: Harvard University Press.
- Mellow, Wesley, and Hal Sider. 1983. Accuracy of response in labor market surveys: Evidence and implications. *Journal of Labor Economics* 1: 331–44.
- Rodgers, Willard L., Charles C. Brown, and Greg J. Duncan. 1993. Errors in survey reports of earnings, hours worked and hourly wages. *Journal of the American Statistical Association* 88:1208–18.
- Stigler, Stephen M. 1977. Do robust estimators work with real data? *Annals of Statistics* 5:1055–98.

Copyright of Journal of Labor Economics is the property of University of Chicago Press and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.