A'Hearn, Brian und Komlos, John:

Improvements in Maximum Likelihood Estimators of Truncated Normal Samples with Prior Knowledge of

Munich Discussion Paper No. 2003-8

Department of Economics
University of Munich

Volkswirtschaftliche Fakultät
Ludwig-Maximilians-Universität München

Online at https://doi.org/10.5282/ubm/epub.51

Improvements in Maximum Likelihood Estimators of Truncated Normal Samples with Prior Knowledge of σ
A Simulation Based Study with Application to Historical Height Samples

Brian A'Hearn
Dept. of Economics
Franklin & Marshall College
Box 3003
Lancaster PA 17604-3003
ph: (717) 291 3926
fax: (717) 291 4369
B_Ahearn@ACAD.FANDM.EDU
USA

John Komlos
Department of Economics
University of Munich
Ludwigstraße 33/IV
D-80539 Munich, Germany
Telephone: +49-89-2180-5824
+ 49-89-8983-9700
Voice Mail: +49-89-2180-3169
Fax: +49-89-33-92-33
email: jk@econhist.de

Abstract

Researchers analyzing historical data on human stature have long sought an estimator that performs well in truncated-normal samples. This paper reviews that search, focusing on two currently widespread procedures: truncated least squares (TLS) and truncated maximum likelihood (TML). The first suffers from bias. The second suffers in practical application from excessive variability. A simple procedure is developed to convert TLS truncated means into estimates of the underlying population means, assuming the contemporary population standard deviation. This procedure is shown to be equivalent to restricted TML estimation. Simulation methods are used to establish the mean squared error performance characteristics of the restricted and unconstrained TML estimators in

relation to several population and sample parameters. The results provide general insight into the bias-precision tradeoff in restricted estimation and a specific practical guide to optimal estimator choice for researchers in anthropometrics.

Introduction

Human height is a widely used synthetic indicator of biological living standards in many different settings, including but not limited to underdeveloped economies, historical contexts, and circumstances in which economic indicators are either unreliable or completely lacking, as, for example, for the children of Soweto (Cameron 2003). Physical stature is positively correlated with net nutrition - the balance between the quantity and quality of nutrient intake and the demands on those resources by the human organism for growth, metabolic maintenance, work, and for resistance to diseases. Of course, individual heights depend as much on genetic potential as on nutrition, but at the population level environmental factors play a very substantial role in determining height at a particular age (Bogin 1999).

Statistical analysis of height data is facilitated considerably by the biological law that height is approximately normally distributed within a population, and that its standard deviation is practically constant, varying over a range of only about one cm, while mean heights can easily vary by more than 15 cm within a population over time (Cole 2000; 2003; Baur and Komlos 2003).[1] Consequently, variations in a population's nutritional status affect primarily the mean of the distribution of heights, and not its form or dispersion.

The Problem of Truncation

Historical height data, though drawn from a normal population, are generally not random samples. In particular, height distributions drawn from military records prior to the introduction of universal conscription are often truncated from below, insofar as most armies imposed a minimum height requirement.[2] Often, the requirement was not strictly

enforced, and varied across units or over time, leaving the distribution with a deficient but not clearly-truncated left tail: the problem of shortfall (Komlos 2003). In the presence of truncation or shortfall, least squares (LS) estimators are biased. The nature of this bias is readily visualized as in Figure 1, which depicts the probability density function (pdf) of a truncated standard normal random variable. Without the lower tail, the mean of the distribution shifts to the right: in Figure 1, truncation at -1σ changes the mean of the distribution from zero to 0.29, clearly biased upward.

Figure 1. Density and mean of a truncated standard normal random variable.

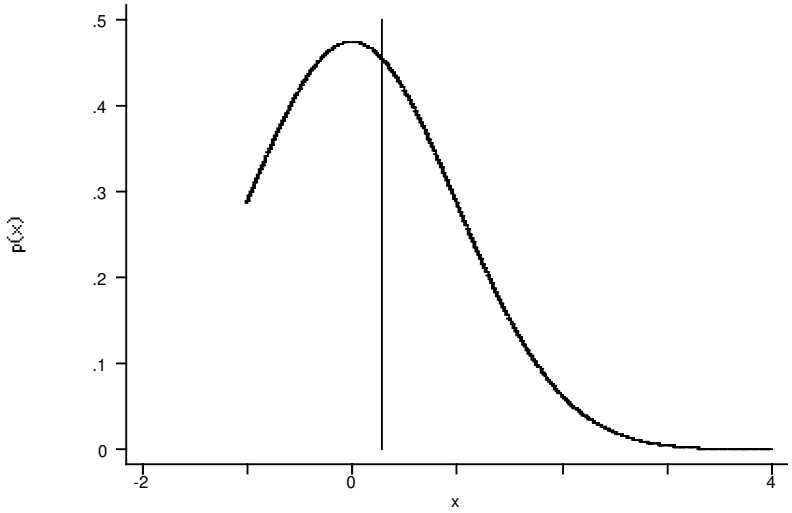

More precisely, suppose that observations on a latent normal random variable Y^*: $y^* = \mu + \varepsilon;\ \ \varepsilon \sim N(0,\sigma^2)$, are selected only if they exceed truncation point τ, that is, when $\mu + \varepsilon \geq \tau$, or $\varepsilon \geq \tau - \mu$. Conditional on being selected into the sample, then, $E[\varepsilon \mid y^* \geq \tau] > 0$ and $\text{var}(\varepsilon \mid y^* \geq \tau) < \sigma^2$. The sample mean of the observed variable Y (the LS estimator of μ), has expected value:

$$E\left(\frac{1}{n}\sum_{i=1}^{n} y_i\right) = E\left(\frac{1}{n}\sum_{i=1}^{n}(\mu+\varepsilon_i)\right) = \mu + \frac{1}{n}\sum_{i=1}^{n} E(\varepsilon_i) \qquad \text{(Eq. 1)}$$

In case of a truncated variable, the second term in Eq. 1 does not equal zero, since the conditional expectation of ε is positive. Nor does the sum of the ε's go to zero as $n \to \infty$. Hence, the sample mean is also inconsistent. This intuition generalizes to situations in which mean height μ_i varies across individuals *i* in relation to observable characteristics. LS estimators of the coefficients of that relationship are biased, and suffer the further problem of heteroskedasticity, the variance of ε_i depending positively on μ_i.[3]

Estimators for Truncated Samples, Part I

Two early approaches to the problem of bias in samples suffering shortfall or truncation took quite different approaches. The Quantile Bend Estimator (QBE) sought to fill in the missing lower tail observations, creating an artificial complete distribution as close as possible to a normal. The mean and standard deviation of this artificial distribution provided unbiased estimators of the underlying population parameters. Komlos and Kim criticized QBE estimates of time trends as displaying excessive short-term variability – fluctuations that were implausible, given biological limits to the variability in the physical stature of a population in the short run. Subsequent simulation studies by Heintel confirmed that the QBE suffers much higher sampling variability than alternative estimators, and its popularity has waned.[4]

In contrast to the QBE, Komlos and Kim (K&K) proposed to exclude those observations that should not have been present, excluding from analysis data below the highest truncation point known to have been applied by military authorities in any period or to any group in the data (Komlos and Kim 1990).[5] The means of the K&K truncated samples by group $\bar{y}_j^{TR}$, j=1…k, of course, remain biased, but the *differences* over time or across groups, i.e. $\bar{y}_j^{TR} - \bar{y}_{j+n}^{TR}$, have the same sign as the true differences in the underlying

population means. This is the case, because μ_j^{TR} is a monotonic function of μ_j, and $\bar{y}_j^{TR}$ is an unbiased estimator of μ_j^{TR}: if $\mu_1^{TR} < \mu_2^{TR}$ then it follows that $\mu_1 < \mu_2$.

A shortcoming of the K&K method is the difficulty of estimating the impact of multiple covariates on height, as the sample must be divided into mutually exclusive subgroups for which means can be calculated. However, it can be extended to the multivariate case by employing least squares regression with multiple dummy variables, again truncating at the highest known τ (truncated LS, or TLS). Though biased, the TLS estimator preserves the signs and relative sizes of coefficients. Chung and Goldberger (1984) showed that the regression slope coefficients relating two variables in a truncated population are proportional to those for the entire population. As TLS is an unbiased estimator of those truncated population slopes, if follows that TLS-estimated coefficients will on average have the same signs and relative sizes as the (complete) population parameters of interest, though not the correct absolute magnitudes.[6] Both K&K and TLS have been used extensively in the literature on historical heights, and offer greater robustness and stability than available alternatives. However, the Chung and Goldberger result implies that we would *expect* truncated means to vary less than true, suggesting that some of the apparently reduced sampling variability of the TLS estimator is illusory. We provide evidence on this below.

Estimators for Truncated Samples, Part II

Given the biologically determined normality of height distributions, an alternative approach to unbiased estimation is maximum likelihood (ML), based on the normal pdf. In the case of truncation, we require a likelihood function that conditions on the probability of *y* having been selected into the sample, i.e. on the chances that $Y^* > \tau$ or

$\varepsilon \geq \tau - \mu$. Intuitively, it is clear that the area under the normal pdf will not integrate to unity without the lower tail, and the function must be normalized accordingly. The pdf of the truncated normal random variable Y is thus:

$$f(y) = \frac{\frac{1}{\sigma}\phi\left(\frac{y-\mu}{\sigma}\right)}{1-\Phi\left(\frac{\tau-\mu}{\sigma}\right)} \quad \text{if } y \geq \tau; \ f(y) = 0 \text{ if } y < \tau. \qquad \text{(Eq. 2)}$$

ϕ denotes the standard normal pdf and Φ denotes its cumulative distribution function (cdf).[7] Note that τ should be adjusted for rounding.[8]

Equation 2 provides the basis for truncated maximum likelihood (TML) estimation of μ and σ. The TML estimator has the usual ML properties of unbiasedness, consistency, and asymptotic efficiency, yields a direct estimate of σ, is amenable to handling multiple truncation points, and permits all forms of hypothesis testing. While these desirable characteristics make TML appear the ideal estimator, experience has demonstrated that TML estimates frequently display implausible magnitudes of variation over time or across groups.[9]

Note that Equation 2 also provides a basis for inferring estimates of μ from K&K- or TLS-estimated truncated sample means $\bar{y}^{TR}$. Using $f(y)$ from Equation 2, for given values of μ, τ, and σ, $\mu^{TR} = \int_{\tau}^{\infty} y f(y)\, dy$.[10] While this integral cannot be solved analytically for μ as a function of μ^{TR} and the other parameters, it is possible to generate μ^{TR} for successive values of μ (given τ and σ) until a μ^{TR} is found that equals $\bar{y}^{TR}$. The associated value of μ is taken as an estimator for the population mean. As $\bar{y}^{TR}$ is an unbiased estimator of μ^{TR}, this converted-TLS (CTLS) method yields unbiased estimates

of μ. This method requires an assumption about σ - an assumption that will rarely be exactly correct and will therefore induce bias; it discards all observations below the highest known τ; it does not address the difficulty of hypothesis testing in the TLS framework; and it is quite cumbersome computationally. However, if the greater stability of the TLS estimator is real, CTLS should share that characteristic. Both issues are addressed in the next section.

Restricted and Unconstrained TML Estimation

In the simple linear regression model with normal disturbances, the LS and ML estimators are equivalent. The use of the truncated-normal pdf in CTLS – the same pdf used in TML estimation – suggests the intuition that the two procedures may be equivalent here too. The assumption of a particular value for σ would then be a form of restricted TML (RTML). This intuition can be tested by simulation, in which repeated samples are generated, $\hat{\mu}_{CTOLS}$ and $\hat{\mu}_{RTML}$ estimated, and their squared difference calculated. Table 1 presents the root mean squared difference for various values of τ and σ, based on 1,000 trials in each case.[11] The simulation results demonstrate that the two estimators are, in fact, essentially identical. The root mean squared difference between the two estimates is on the order of a hundredth of a millimeter with μ=165.

Table 1. Root mean squared difference between RTML and CTLS estimates (cm)

σ:	6.00	6.86	7.50
τ:			
150	.0012	.0012	.0013
163	.0012	.0012	.0013
166	.0012	.0013	.0012

1000 replications; sample size 1000; true mean 165 cm in all cases.

While yielding identical estimates of μ, RTML has the advantages of permitting hypothesis testing and allowing for multiple truncation points, and can be considered together with unconstrained TML in a unified framework. In the simulations summarized in Table 1, the true value of σ was used in estimating $\hat{\mu}_{RTML}$, i.e. 6.00 cm in the column 1 simulations, 6.86 in column 2, etc. In practice, of course, σ is never known. And whenever an incorrect restriction is imposed some degree of bias is induced. The figure of 6.86 cm has been suggested as a plausible figure for males based on data for modern populations, but any such rule of thumb will never be *exactly* correct (Frisancho 1990; Cole 2000). However, a general property of restricted ML estimators is their greater sampling precision in comparison with unconstrained estimation, and this is true *regardless of any bias* introduced by incorrect restrictions.[12]

The optimal choice of estimator thus depends on how we evaluate the tradeoff between bias and precision. A common criterion for balancing these risks is mean squared error (MSE):

$$MSE(\hat{\mu}) = E\left[(\hat{\mu} - \mu)^2\right] = \text{bias}(\hat{\mu})^2 + \text{var}(\hat{\mu}) \qquad \text{(Eq. 3)}$$

If the restrictions imposed are "close" to the true value, the bias induced will be small, and the reduction in variance substantial. So the choice between restricted and unrestricted estimation depends on the researcher's degree of confidence in the restrictions. The extent of the trade-off between bias and variance is an empirical question, depending on sample size (n), τ, μ and σ. In the remainder of this paper, we use simulation methods to explore this trade-off and provide a practical guide for which

estimator, $\hat{\mu}_{TML}$ or $\hat{\mu}_{RTML}$, is to be preferred in the range of circumstances typical of applied work with historical heights.

MSE Performance of Restricted and Unconstrained TML Estimators

The simulation results reported here are based on 2,000 replications for a range of n, σ, and τ. The ranges for each parameter were chosen as representative of values likely to be encountered in research on historical heights: μ=165 cm throughout; σ varies between 6 and 8 cm by increments of 0.5 cm, with special consideration give to 6.86 cm; τ varies from 150 cm to 167 cm; n = 250, 500, and 1000. The RTML procedure imposes the restriction that σ = 6.86 cm in all cases.

RTML results are summarized in Appendix Table 1, and illustrated in Figure 2. They indicate that a) MSE and both of its constituent elements increase with increasing τ; b) MSE is very high if $\tau > \mu$; c) if $\sigma \neq$ 6.86 cm then $E(\hat{\mu}) \neq \mu$, with bias increasing very rapidly as the σ restriction error exceeds about 5 mm; d) the variance of the estimates decreases as n increases, while bias is roughly constant.

Figure 2: MSE of the Restricted TML Estimator, n = 500.[13]

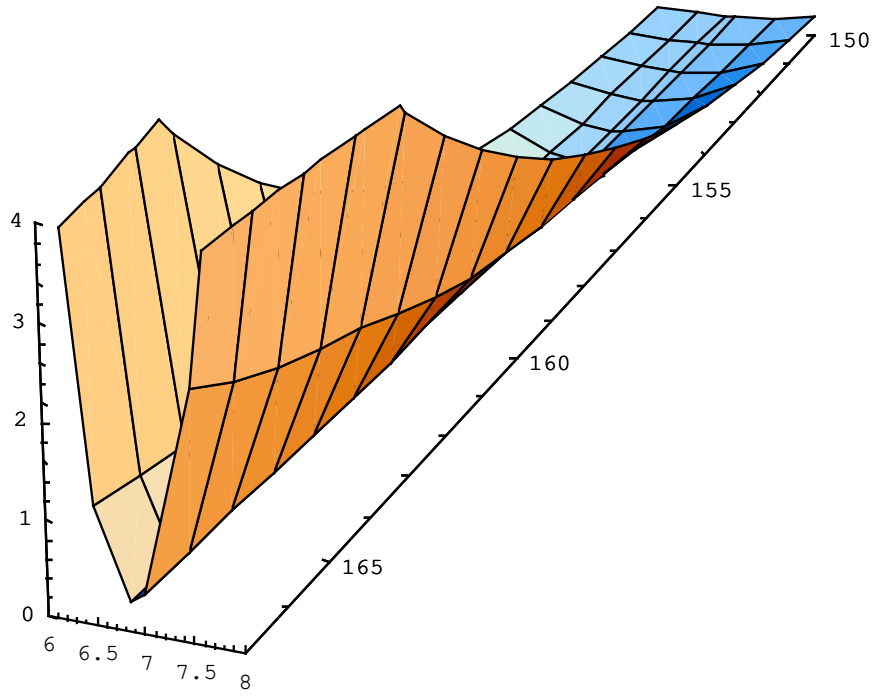

Note: τ is measured along the long axis, σ along the short, both in centimeters. MSE is measured vertically in cm^2 units.

Table A2 and Figure 3 present the results for unconstrained ML estimation. They indicate that a) the bias of the TML estimator is zero, so that MSE is driven by variance alone; b) MSE rises with τ, but much more rapidly than was the case for the restricted estimator: MSE at $\tau = 167$ is more than ten times its value at $\tau = 159$; c) variance decreases with sample size, yielding a sharper decrease in MSE than for the restricted estimator since bias plays no role in its determination; d) MSE rises with the underlying population standard deviation σ. It is clear that for small n and $\tau \geq \mu$, $\hat{\mu}_{TR}$ is unreliable. (In Figure 3, MSE is off the scale for $\tau \geq 166$.[14]) At the extremes of $\tau = 167$, $\sigma = 8$, and $n = 250$, MSE is 12 cm^2. This implies a root MSE of 3.5 cm, which is very large relative to the differences typically observed over time or across groups of historical data. Note that MSE values are off the scale (greater than 4) for $\tau \geq 166$ at the higher values of σ.

Figure 3. MSE of the Unrestricted TML Estimator, n = 500.

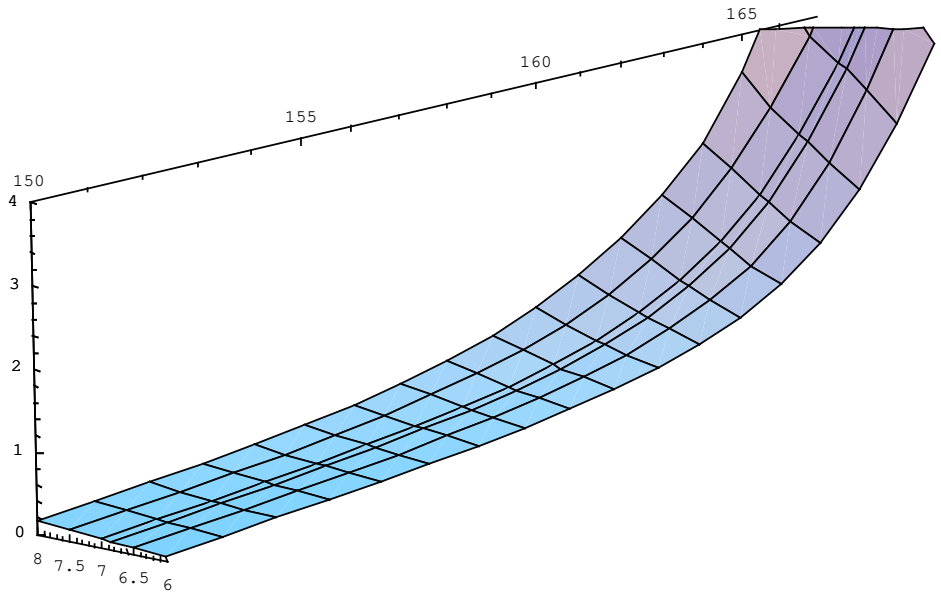

Note: τ is measured along the long axis, σ along the short, both in centimeters. MSE is measured vertically in cm^2 units.

The practical question is: over what range of parameter values of μ and σ is the restricted estimator preferred to unconstrained TML? Table A3 reports differences in MSE: $\hat{\mu}_{TML} - \hat{\mu}_{RTML}$. Positive numbers thus indicate superior performance of the restricted estimator. Figure 4 makes it clear that if the σ restriction is approximately true (say, within half a centimeter) and the truncation point exceeds about 160 cm, the restricted estimator offers substantial advantages over unconstrained TML. Considering Table A3, the choice is clearest at the extremes (the upper right and lower left corners). At truncation points well above the mean in small samples, the restricted estimator offers dramatically better precision than $\hat{\mu}_{TML}$, which far outweighs its bias. However, at truncation points well below the mean in large samples, unconstrained estimation is

generally to be preferred; it performs less well only in the immediate neighborhood of σ = 6.86 cm and then only by a very slight amount.

Figure 4. MSE Difference $\hat{\mu}_{TML} - \hat{\mu}_{RTML}$, n = 500.

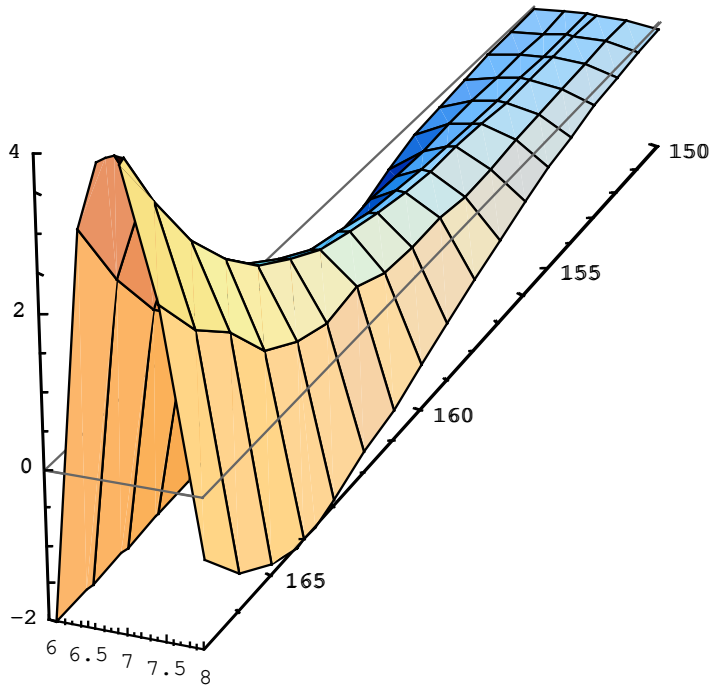

Note: τ is measured along the long axis, σ along the short, both in centimeters.
The MSE difference is measured vertically in cm^2 units.
The front lower left corner of the surface is off the scale, i.e. less than –2.

In historical datasets one typically encounters truncation points less than but not far from the mean, and sample sizes that are relatively small on account of the fact that they have to be drawn manually from archival records. Thus, the center row of Table 3A, for τ = 163 cm, and the columns for n = 250 or 500 can be considered typical of such situations. Figure 5 depicts the MSE graphically for τ = 163 and n = 250 and 500. (Appendix Figure 1 shows the same information for the more extreme case of τ = 165.) It

becomes clear that the optimal choice of estimator depends on the degree of confidence in the σ restriction. If one is confident that σ is within a half-centimeter or so of 6.86, $\hat{\mu}_{\text{CTLS}}$ or $\hat{\mu}_{\text{RTR}}$ is clearly preferred over $\hat{\mu}_{\text{TR}}$. The reduction in risk can be substantial in small samples: on the order of 2 cm^2 when the restriction is approximately correct. That implies a reduction of about 1.4 cm in the expected error (i.e. root MSE). This is large, relative to likely true differences over time and across groups.

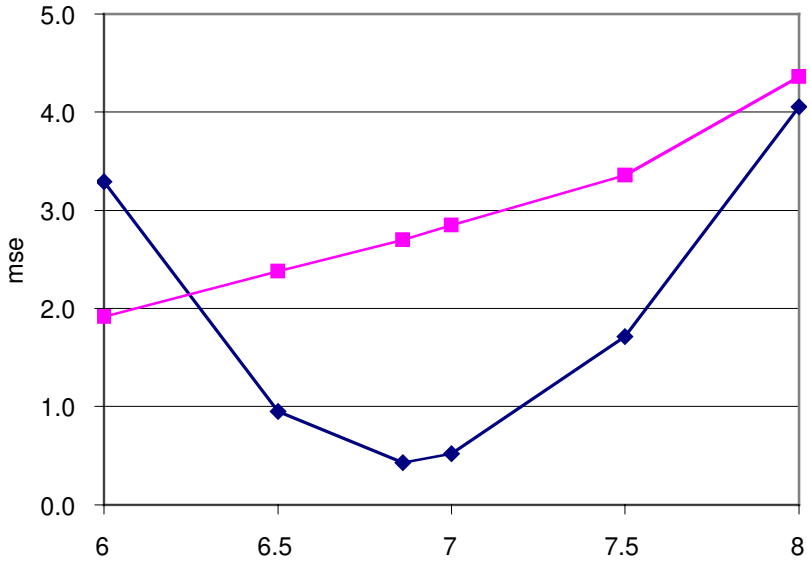

Figure 5a. MSE of TML and RTML Estimators, $\tau = 163$ cm, $n = 250$.

Figure 5b. MSE of TML and RTML Estimators, $\tau = 163$ cm, $n = 500$.

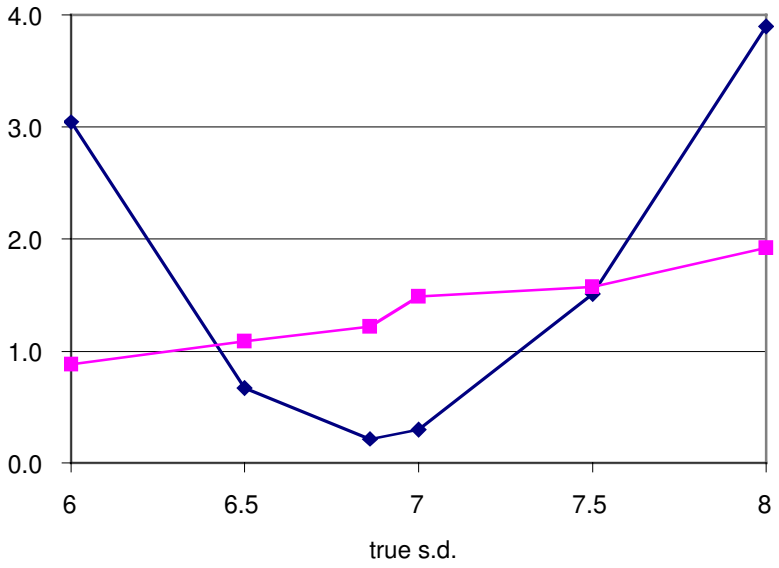

Summary

Much research with historical height data must confront the problem of samples truncated from below, a situation that induces bias in the standard LS estimator. Two currently popular approaches are the calculation of truncated sample means (the K&K or TLS estimators) and truncated maximum likelihood estimation (TML). The TLS method has the advantages of simplicity, accuracy in estimating the sign of differences over time or across groups, and robustness. TML has the advantages of unbiasedness, yielding an estimate of the population's standard deviation, the ability to handle multiple truncation points, and allowing hypothesis testing.

Wwe have described a method for inferring estimates of population parameters from TLS-estimated truncated sample means. We show by simulation that this Converted-TLS estimator is a restricted form of TML, in which the population standard deviation is constrained to take a value chosen by the researcher. The choice between a restricted and an unconstrained ML estimator boils down to the tradeoff between the bias induced by incorrect restrictions and the superior precision offered by the restricted estimator.

This tradeoff is evaluated in terms of the mean squared error criterion, which is calculated for a range of parameter values using simulation methods. The results show that the restricted estimator's precision is a decisive advantage if the truncation point is close to or above the population mean. This is particularly the case with small sample sizes. In situations most typical of research with historical height data, choice of estimation technique will depend on the degree of confidence in the restriction that the historical population's standard deviation equals a particular value such as the modern

figure of 6.86 cm. If that restriction is true within ± 0.5 cm, the restricted estimator is preferred in all situations except in large samples with a relatively low truncation point.

REFERENCES

A'Hearn, Brian. "Anthropometric Evidence on Living Standards in Northern Italy, 1730-1860," Journal of Economic History, forthcoming.

Baten, Jörg and Markus Heintel. "Zum Problem der Verteilungen mit Shortfall bei der Nutzung des Indikators Durchschnittliche Körpergröße," Historical Social Research, vol. 20, no. 4 (1995), pp. 135-154.

Mariluise Baur and John Komlos (2003) From the Fattest to the Tallest: the Fate of the American Population in the 20th Century," Unpublished Manuscript, University of Munich.

Bogin, Barry. 1999. Patterns of Human Growth. Cambridge: Cambridge University Press, 2nd edition.

Chung C. and A. Goldberger. 1984. „Proportional Projections in limited Dependent Variable Models," Econometrica 52: 531-34.

Chay, Kenneth Y. and James L. Powell, (2001) "Semiparametric Censored Regression Models," Journal of Economic Perspectives, 15, 4: 29-42.

Cole, T.J. (2000) "Galton's Midparent height revisited, " Annals of Human Biology 27, 4: 401-405.

Cole, T.J. (2003) "The secular trend in human physical growth: a biological view," Economics and Human Biology I, no. 2.

Cuff, Timothy. „Variation and Trends in the Stature of Pennsylvanians, 1820-1860," in John Komlos and Jörg Baten, eds*., The Biological Standard of Living in Comparative Perspective*, Stuttgart: Franz Steiner, 1998, pp. 208-235.

EViews 4 User's Guide Quantitative Irvine CA: Micro Software, 2000.

Frisancho, R.A. (1990) Anthropometric standards for the assessment of growth and nutritional status. Ann Arbor, MI: University of Michigan Press.

Greene, William. (1993) Econometric Analysis, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall.

Heintel, Markus. (1996) "Historical Height Samples with Shortfall: a Computational Approach," History and Computing, vol. 8, no. 1, pp. 24-37.

Heintel, Markus, Sandberg, L. and R. Steckel. (1998). "Swedish Historical Heights Revisited. New Estimation Techniques and Results," in John Komlos and Jörg Baten, eds., *The Biological Standard of Living in Comparative Perspective*, Stuttgart: Franz Steiner, 1998, pp. 449-458.

Judge, George et al. Introduction to the Theory and Practice of Econometrics, 2nd ed. New York: John Wiley and Sons, 1988.

Komlos, John. 2003. "How to (and How Not to) Analyze Deficient Height Samples," Unpublished Manuscript, University of Munich.

Komlos, John. (1999) "On the Nature of the Malthusian Threat in the Eighteenth Century," Economic History Review 52, no. 2 (Nov): 730-48.

Komlos, John, 1993. "The Secular Trend in the Biological Standard of Living in the United Kingdom," Economic History Review, 46: 115-44.

Komlos, John, 1989. Nutrition and Economic Development in the Eighteenth Century Habsburg Monarchy (Princeton: Princeton University Press).

Komlos, John, 1985. "Stature and Nutrition in the Habsburg Monarchy: The Standard of Living and Economic Development in the Eighteenth Century" American Historical Review 90: 1149-61.

Komlos, John and Joo Han Kim. (1990) "Estimating Trends in Historical Heights," Historical Methods, vol. 23, no. 3 (Summer), pp. 116-120.

Ruud, Paul. An Introduction to Classical Econometric Theory. Oxford: Oxford University Press, 2000.

Stata Reference Manual. Release 7, College Station, TX: Stata Press, 2001, vol. 4.

Twarog, S. 1997. Heights and Living Standards in Germany, 1850-1939: The Case of Württemberg. In Health and Welfare during Industrialization, edited by R Steckel and R. Floud. Chicago: The University of Chicago Press.

Table A1. Mean Squared Error of the Restricted TML Estimator, $\hat{\mu}_{RTML}$.

		n = 250			*n = 500*			*n = 1000*		
τ	σ	$bias^2$	*var.*	*MSE*	$bias^2$	*var.*	*MSE*	$bias^2$	*var.*	*MSE*
167	6.00	6.105	0.620	6.725	6.180	0.383	6.563	6.122	0.187	6.309
	6.50	1.023	0.692	1.714	0.952	0.301	1.253	0.932	0.153	1.085
	6.86	0.000	0.591	0.594	0.000	0.314	0.315	0.000	0.157	0.158
	7.00	0.082	0.613	0.695	0.123	0.300	0.423	0.112	0.151	0.263
	7.50	2.226	0.571	2.797	2.323	0.279	2.602	2.374	0.135	2.509
	8.00	6.759	0.555	7.314	6.918	0.271	7.190	6.855	0.135	6.990
165	6.00	4.350	0.574	4.923	4.334	0.282	4.616	4.293	0.148	4.440
	6.50	0.720	0.534	1.254	0.725	0.259	0.984	0.669	0.128	0.796
	6.86	0.000	0.526	0.528	0.000	0.254	0.255	0.000	0.131	0.131
	7.00	0.096	0.536	0.632	0.103	0.261	0.364	0.083	0.126	0.209
	7.50	1.731	0.537	2.268	1.530	0.278	1.808	1.759	0.123	1.882
	8.00	4.967	0.489	5.456	5.038	0.251	5.289	5.086	0.108	5.194
163	6.00	2.857	0.436	3.293	2.821	0.222	3.043	2.823	0.105	2.928
	6.50	0.498	0.449	0.947	0.462	0.213	0.674	0.462	0.114	0.575
	6.86	0.000	0.428	0.428	0.000	0.218	0.218	0.000	0.113	0.113
	7.00	0.040	0.483	0.523	0.052	0.249	0.301	0.059	0.116	0.175
	7.50	1.305	0.408	1.713	1.305	0.206	1.511	1.194	0.107	1.301
	8.00	3.648	0.405	4.053	3.699	0.197	3.896	3.651	0.105	3.757
161	6.00	1.830	0.353	2.183	1.750	0.166	1.916	1.760	0.091	1.851
	6.50	0.281	0.383	0.664	0.304	0.193	0.497	0.282	0.092	0.374
	6.86	0.000	0.405	0.406	0.000	0.199	0.199	0.000	0.098	0.098
	7.00	0.027	0.352	0.379	0.037	0.189	0.226	0.045	0.093	0.138
	7.50	0.768	0.406	1.174	0.678	0.188	0.866	0.844	0.094	0.938
	8.00	2.528	0.398	2.926	2.541	0.186	2.727	2.553	0.093	2.646
159	6.00	1.057	0.286	1.343	1.062	0.157	1.219	1.017	0.074	1.091
	6.50	0.184	0.325	0.509	0.240	0.162	0.402	0.172	0.076	0.247
	6.86	0.000	0.322	0.322	0.000	0.171	0.171	0.000	0.085	0.085
	7.00	0.020	0.317	0.337	0.022	0.157	0.180	0.025	0.079	0.104
	7.50	0.565	0.320	0.885	0.542	0.166	0.708	0.520	0.088	0.607
	8.00	1.657	0.348	2.005	1.712	0.161	1.873	1.682	0.084	1.767

Table A2. Mean squared error of the Unrestricted TML Estimator, $\hat{\mu}_{TML}$.

		n = 250			*n = 500*			*n = 1000*		
τ	σ	$bias^2$	*var.*	*MSE*	$bias^2$	*var.*	*MSE*	$bias^2$	*var.*	*MSE*
167	6.00	0.000	9.431	9.431	0.000	3.894	3.894	0.000	1.604	1.604
	6.50	0.000	9.559	9.559	0.000	4.450	4.450	0.000	1.875	1.875
	6.86	0.000	10.728	10.728	0.000	4.771	4.771	0.000	1.970	1.970
	7.00	0.000	12.476	12.476	0.000	4.704	4.704	0.000	2.083	2.083
	7.50	0.000	16.214	16.214	0.000	4.836	4.836	0.000	2.191	2.191
	8.00	0.000	12.233	12.233	0.000	6.515	6.515	0.000	2.468	2.468
165	6.00	0.000	4.474	4.474	0.000	1.898	1.898	0.000	0.821	0.821
	6.50	0.000	4.669	4.669	0.000	2.051	2.051	0.000	1.015	1.015
	6.86	0.000	5.575	5.575	0.000	2.470	2.470	0.000	1.123	1.123
	7.00	0.000	5.936	5.936	0.000	2.392	2.392	0.000	1.183	1.183
	7.50	0.000	5.913	5.913	0.000	2.819	2.819	0.000	1.336	1.336
	8.00	0.000	7.776	7.776	0.000	3.146	3.146	0.000	1.391	1.391
163	6.00	0.000	1.876	1.876	0.000	0.841	0.841	0.000	0.420	0.420
	6.50	0.000	2.374	2.374	0.000	1.078	1.078	0.000	0.507	0.507
	6.86	0.000	2.959	2.959	0.000	1.318	1.318	0.000	0.589	0.589
	7.00	0.000	2.907	2.907	0.000	1.289	1.289	0.000	0.644	0.644
	7.50	0.000	3.468	3.468	0.000	1.646	1.646	0.000	0.712	0.712
	8.00	0.000	4.594	4.594	0.000	1.988	1.988	0.000	0.864	0.864
161	6.00	0.000	0.909	0.909	0.000	0.440	0.440	0.000	0.217	0.217
	6.50	0.000	1.215	1.215	0.000	0.588	0.588	0.000	0.290	0.290
	6.86	0.000	1.537	1.537	0.000	0.691	0.691	0.000	0.348	0.348
	7.00	0.000	1.465	1.465	0.000	0.697	0.697	0.000	0.375	0.375
	7.50	0.000	1.907	1.907	0.000	0.927	0.927	0.000	0.422	0.422
	8.00	0.000	2.355	2.355	0.000	1.139	1.139	0.000	0.558	0.558
		0.000								
159	6.00	0.000	0.478	0.478	0.000	0.240	0.240	0.000	0.123	0.123
	6.50	0.000	0.664	0.664	0.000	0.337	0.337	0.000	0.158	0.158
	6.86	0.000	0.832	0.832	0.000	0.364	0.364	0.000	0.193	0.193
	7.00	0.000	0.896	0.896	0.000	0.393	0.393	0.000	0.205	0.205
	7.50	0.000	1.100	1.100	0.000	0.525	0.525	0.000	0.271	0.271
	8.00	0.000	1.422	1.422	0.000	0.667	0.667	0.000	0.332	0.332

Table A3. MSE Difference, Unconstrained – Restricted TML, $\hat{\mu}_{TML} - \hat{\mu}_{RTML}$.

		n = 250	*n = 500*	*n = 1000*
τ	σ	*MSE diff.*	*MSE diff.*	*MSE diff.*
167	6.00	2.706	-2.669	-4.705
	6.50	7.845	3.197	0.790
	6.86	10.134	4.456	1.812
	7.00	11.781	4.281	1.820
	7.50	13.417	2.234	-0.318
	8.00	4.919	-0.675	-4.522
165	6.00	-0.449	-2.718	-3.619
	6.50	3.415	1.067	0.219
	6.86	5.047	2.215	0.992
	7.00	5.304	2.028	0.974
	7.50	3.645	1.011	-0.546
	8.00	2.320	-2.143	-3.803
163	6.00	-1.417	-2.202	-2.508
	6.50	1.427	0.404	-0.068
	6.86	2.531	1.100	0.476
	7.00	2.384	0.988	0.469
	7.50	1.755	0.135	-0.589
	8.00	0.541	-1.908	-2.893
161	6.00	-1.274	-1.476	-1.634
	6.50	0.551	0.091	-0.084
	6.86	1.131	0.492	0.250
	7.00	1.086	0.471	0.237
	7.50	0.733	0.061	-0.516
	8.00	-0.571	-1.588	-2.088
159	6.00	-0.865	-0.979	-0.968
	6.50	0.155	-0.065	-0.089
	6.86	0.510	0.193	0.108
	7.00	0.559	0.213	0.101
	7.50	0.215	-0.183	-0.336
	8.00	-0.583	-1.206	-1.435

Note: positive numbers indicate superior MSE performance of the restricted estimator.

Appendix Figure 1a. MSE of TML and RTML Estimators, $\tau = 165$ cm, $n = 250$.

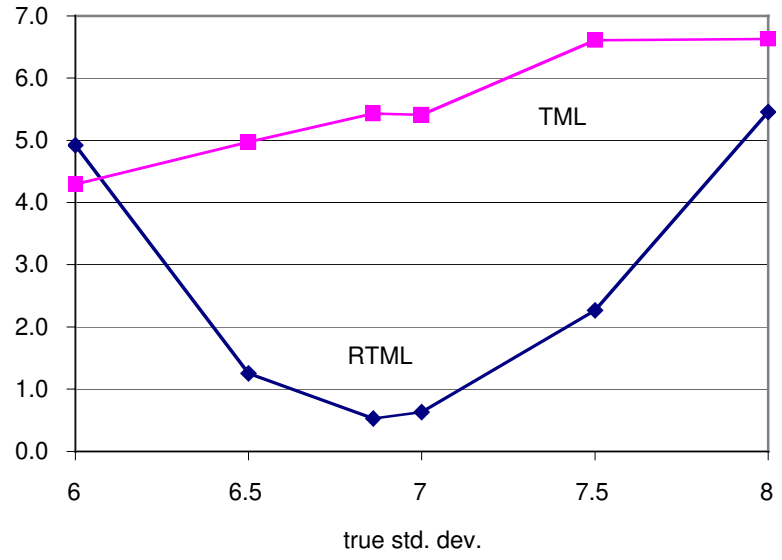

Appendix Figure 1b. MSE of TML and RTML Estimators, $\tau = 165$ cm, $n = 500$.

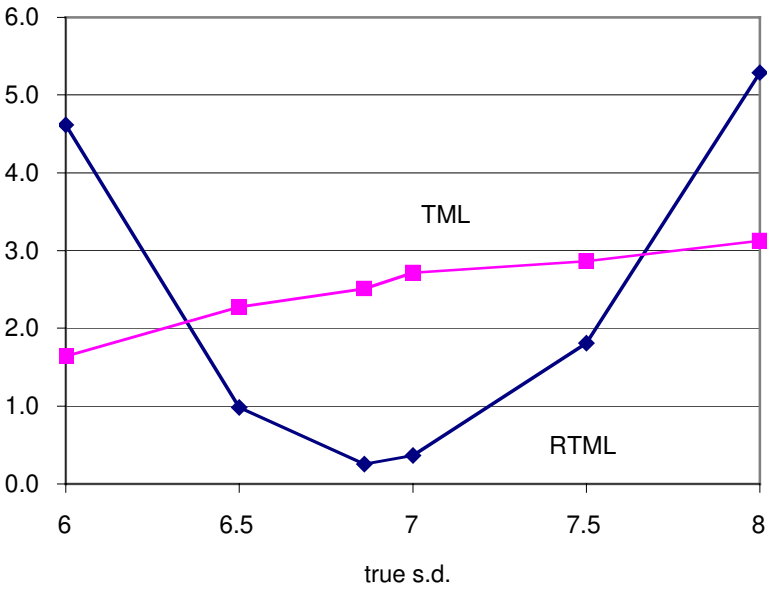

[1] More specifically, the standard deviation of heights appears to range between 6 and 7 cm among modern male populations, between 5.3 and 6.5 cm among females: Cole (2000) p. 402.

[2] Truncation from above is also observed on occasion, but will not be dealt with here.

[3] See Ruud (2000) Ch. 28 for a full discussion of truncated distributions.

[4] Heintel has shown that both the TLS and TML estimators defined below offer superior mean squared error (Heintel 1995, 1996). See, in addition, Komlos (1985, 1989) and Heintel, Sandberg and Steckel (1998) for empirical investigations confirming the unreliability of the QBE.

[5] By applying the highest-known τ to all groups, the artificial impact of potentially different shortfall patterns in the sub-samples is avoided by "equalizing" the bias over the complete sample (see Heintel and Baten, 1998, footnote 17, for an interesting study of artificial correlations if one fails to equalize the bias). When τ is unknown, Heintel's truncation point estimator can be used. This procedure first smoothes the histogram of sample heights using a kernel density estimator, then identifies $\hat{\tau}$ as the point where the estimated density's slope is maximal (i.e., where its first difference is greatest) (Heintel 1996).

[6] Chung and Goldberger 's results apply to a wide variety of selection and censoring rules, and require only minimal assumptions, not including normality. In the particular case of interest here, the truncated population slope coefficients can be shown to be biased toward zero, relative to the complete population slopes. The authors base their discussion on the

concept of projection. For the connection between projection and LS regression, see Ruud (2000) Part 1.

[7] Ruud (2000) Ch. 28 and Greene (1993) Ch. 22 provide more detail.

[8] If measurements were to the nearest half-centimeter, for example, the truncation point should be set at $\hat{\tau}$ - 0.25 cm. Aside from this, rounding has no significant effect on the performance of the estimators considered here, even if heaping, or the clustering of observations around particular values such as even integers, may pose more problems. See Komlos (1999) for a preliminary investigation.

[9] Examples include ML estimates by Twarog and A'Hearn of time trends in German and Italian data, respectively. In both cases, heights could fluctuate by as much as 2-3 cm across five-year birth cohorts, a figure as large as cumulative gains (losses) one would expect over about two decades of improving (deteriorating) living standards. Twarog (1997), A'Hearn (forthcoming). In the absence of famine, wars, or comparable events and their aftermaths, such dramatic and temporary fluctuations in height of a population are completely implausible.

[10] Ruud (2001, p. 804) gives an alternative expression for the truncated mean as the population mean plus a term involving the hazard rate (ratio of pdf to cdf) evaluated at the population mean, and shows how this equation can be the basis of a feasible weighted non-linear least squares (FWNLS) estimation procedure. It is further shown that FWNLS is not efficient, however.

[11] The simulations were run using code written by the authors for both the CTLS and RTML estimators in STATA6. STATA7 has a built-in TML command ("truncreg"), which can be used subject to any linear constraints.

[12] For a proof see Judge et al 1988, pp. 235-40.

[13] The data were smoothed before graphing in this and the other figures.

[14] The increase in MSE with σ is seen in the torque of the surface, which twists up in the back, down in the front from the viewer's perspective. This effect is pronounced for high values of τ.