Inequality Among the Wealthy

Frank A. Cowell

Contents

CASE/150
June 2011

Centre for Analysis of Social Exclusion
London School of Economics
Houghton Street
London WC2A 2AE
CASE enquiries – tel: 020 7955 6679

Centre for Analysis of Social Exclusion

The Centre for the Analysis of Social Exclusion (CASE) is a multi-disciplinary research centre based at the London School of Economics and Political Science (LSE), within the Suntory and Toyota International Centres for Economics and Related Disciplines (STICERD). Our focus is on exploration of different dimensions of social disadvantage, particularly from longitudinal and neighbourhood perspectives, and examination of the impact of public policy.

In addition to our discussion paper series (CASEpapers), we produce occasional summaries of our research in CASEbriefs, and reports from various conferences and activities in CASEreports. All these publications are available to download free from our website. Limited printed copies are available on request.

For further information on the work of the Centre, please contact the Centre Manager, Jane Dickson, on:

Telephone:	UK+20 7955 6679
Fax:	UK+20 7955 6951
Email:	j.dickson@lse.ac.uk
Web site:	http://sticerd.lse.ac.uk/case

Editorial Note and Acknowledgements

Frank Cowell (email: f.cowell@lse.ac.uk) is a Professor in the Economics Department at the London School of Economics and a CASE associate. He is grateful to the Nuffield Foundation for financial support and to John Hills for helpful comments on an earlier draft. Alexander Teytelboym, Maximillian Eber and Girija Girish provided valuable research assistance.

Abstract

Using the evidence from the Luxembourg Wealth Study it appears that the distribution of wealth in the UK is considerably less than in Canada, the US or Sweden. But does this result come from an underestimate of inequality among the wealthy and of the wealth differential between the rich and the rest? Using a Pareto model for the upper tail of the distribution we can see that the inequality of comparisons of the UK with the other countries is indeed robust.

Keywords: wealth distribution
JEL numbers: D31

1 Introduction

Is wealth inequality in the UK very high? It is an emotive question, one to which we are likely to react using a variety of suspicions and prejudices. But if we are to provide something other than an evasive answer we need to be clear about two other questions: High on what criteria? High relative to what?

In this paper we will interpret these two questions as follows. First we address the question of how to measure wealth inequality in principle and in practice, focusing on special problem areas that are characteristic of wealth distribution. Second, we examine whether wealth in the UK is more unequally distributed than in other comparable developed countries, using the best available data for making such comparisons and taking into account methods to deal with the special measurement issues for wealth.

The paper is structured as follows. Section 2 discusses the international data source used here. Section 3 discusses several important technical issues pertaining to wealth-inequality measurement and presents a first pass at the breakdown of inequality across countries, Sections 4 and 5 discuss the way in which a functional form may be used to model the upper tail of the distribution and to provide refined estimates of the wealth-inequality breakdown and Section 6 concludes.

2 The Data

As is well known, wealth data present special problems of empirical analysis in comparison with data on incomes or earnings. There are several issues in connection with the tails of the distribution: the data sometimes miss out the assets possessed by those with little wealth; the data may be sparse and possibly unreliable in the upper tail – precisely the part of the distribution where one would like detailed information in order to make useful inequality comparisons. To some extent the problems of the lower tail have been overcome by the recent availability of datasets with a broader coverage of assets and of individuals. However, this broader coverage does not offer the same improvements in analysing the distribution of wealth amongst the wealthy, the issue that will be treated here.

If the comparisons are to be made across countries then one obviously has to overcome further difficulties: wealth concepts and conventions for collecting or reporting data may differ between countries. However, this type of problem can now be addressed by using the Luxembourg Wealth Study (LWS) described in Sierminska et al. (2006), which provides a harmonised internationally comparable database for a small number of developed countries. Here we use this to focus on net worth in four countries around the turn of the millennium: Canada (1999), Sweden (2002), the UK (2000) and the US (2000).[1] Of course the fact that the wealth data have been carefully harmonised to ensure, as far as possible, international comparability does not mean that the data sources underlying LWS are going to be perfect in every respect: indeed it can be argued that it is in respect of "wealth amongst the wealthy" that some of the LWS data may be less than ideal.[2]

For the comparisons undertaken here the unit of analysis is the household. The wealth concept used is the LWS-defined Net Worth 1 which consists of the following components:

Total Financial Assets (TFA):	Sum of	deposit accounts, bonds, stocks and mutual funds
Total Nonfinancial Assets (TNA):	Sum of	value of principal residence and investment property
Net Worth 1:	TFA + TNA – total debt	

Again see Sierminska et al. (2006) for a detailed discussion of these wealth components.

[1]The sources used for the LIS harmonised database are as follows. *Canada*: *Survey of Financial Security*, an interview survey (with over-sampling of the wealthy) from Statistics Canada . *Sweden*: *Wealth Survey*, an interview survey combined with administrative records, provided by Statistics Sweden *United Kingdom: British Household Panel Survey*, a panel interview survey. *United States*, *Survey of Consumer Finances*, an interview survey (with over-sampling of the wealthy) from the Federal Reserve Board and U.S. Department of Treasury.

[2]For example, *w*hile the BHPS has the advantages compared to other UK sources of wealth data (HMRC does not provide effective coverage of wealth in the lower tail; the Wealth and Assets survey is only recently available and so cannot provide the run of years in BHPS) it is known to under-record financial assets, in the light of the evidence from these other UK sources. This under-recording this may affect the upper tail of the wealth distribution disproportionately.

3 Wealth inequality: A first look

How should we measure inequality of wealth? Clearly it would be helpful if we were to apply tools that are familiar and accepted in other contexts, such as income inequality. If one can just carry across some standard tools from the study of income and expenditure distributions, then it would be easier to compare different types of economic inequality and one could just carry across any required statistical techniques.

3.1 Inequality measures

In any study using inequality measures there are some standard caveats. The sparse data in the upper tail of the distribution of income or wealth may present problems for "top-sensitive" inequality measures. Likewise the lower tail of the distribution will typically present difficulties for "bottom-sensitive" inequality measures: for example measurement error concerning low values of wealth. These problems may affect how we can measure inequality: they will rule out the use of some indices and restrict range of application of others.

However, there is a further important practical difficulty. In the case of incomes it is often assumed that income is necessarily non-negative; in practice there may be negative incomes but usually the number of these is small and it is common practice just to ignore them. In the case of wealth the presumption that we are dealing with a non-negative quantity cannot be justified. It is a fact of life that many people enter a period of indebtedness at some point in their life. So, if we are interested in the inequality of net worth we have to accept that in principle this could be negative for some people at some point in their lives. Moreover the proportion of the population that has negative net worth at any given moment could be non-negligible (see below) and therefore a representative sample of population will inevitably contain a corresponding proportion of those with negative wealth. So the inequality index has to be defined for negative values. This precludes quite a large set of otherwise inequality indices; fortunately it does leaves a few practical and well known-indices including the coefficient of variation, the relative mean deviation and the Gini coefficient (Amiel et al. 1996). Here we will focus principally on the Gini.

We want to focus on the inequality of the rich. To make this precise we need to introduce some simple notation. Let the "rich" be defined as the

	Gini overall (G)	**Share rich** (s_{R})	**rich** (G_{R})	**Gini non-rich** (G_{N})	**between** (G_{B})
			($p_{\mathrm{R}} = 0.10$)		
UK	0.665	0.456	0.260	0.607	0.356
Sweden	0.893	0.582	0.314	1.045	0.482
Canada	0.747	0.532	0.293	0.710	0.432
US	0.836	0.705	0.349	0.779	0.605
			($p_{\mathrm{R}} = 0.05$)		
UK		0.301	0.223	0.618	0.251
Sweden		0.406	0.316	0.941	0.356
Canada		0.374	0.261	0.703	0.324
US		0.575	0.318	0.748	0.525
			($p_{\mathrm{R}} = 0.01$)		
UK		0.101	0.157	0.644	0.091
Sweden		0.175	0.326	0.891	0.165
Canada		0.151	0.132	0.721	0.141
US		0.329	0.198	0.777	0.319

Source: LWS

Table 1: Gini decomposition for the top 10 percent, top 5 percent, top 1 percent

top p_{R} proportion of the population and define the proportion of non-rich as $p_{\mathrm{N}} := 1 - p_{\mathrm{R}}$. Correspondingly let s_{R} and s_{N} be the wealth shares of the rich and non-rich groups and μ, μ_{R}, μ_{N} be the mean wealth overall and in the two groups. We have the following:

$$\mu = p_{\mathrm{R}}\mu_{\mathrm{R}} + p_{\mathrm{N}}\mu_{\mathrm{N}} \tag{1}$$

$$s_{\mathrm{R}} = \frac{p_{\mathrm{R}}\mu_{\mathrm{R}}}{\mu} \tag{2}$$

$$s_{\mathrm{N}} = \frac{p_{\mathrm{N}}\mu_{\mathrm{N}}}{\mu}. \tag{3}$$

We can then make use of a standard decomposition formula for the Gini coefficient that is convenient when the distribution can be partitioned by wealth level. Let G_{R} and G_{N} be the Gini coefficients for the wealth distribution among the rich and the wealth distribution amongst the non-rich. Also let G_{B} denote the between-group Gini (the Gini inequality that would arise if all wealth in the rich and non-rich groups were concentrated at the respective group means) which may be written as the wealth share of the rich minus population share of rich:

$$G_{\mathrm{B}} = s_{\mathrm{R}} - p_{\mathrm{R}}. \tag{4}$$

Then the Gini coefficient for the whole distribution G is the weighted sum of the rich Gini and the non-rich Gini plus between-group Gini:

$$G = p_{\mathrm{R}}s_{\mathrm{R}}G_{\mathrm{R}} + p_{\mathrm{N}}s_{\mathrm{N}}G_{\mathrm{N}} + G_{\mathrm{B}}. \tag{5}$$

Table 1 gives the breakdown for the four countries in three cases corresponding to different assumptions about which group constitutes the rich: p_{R} has been chosen as 10, 5 and 1 per cent respectively.[3] Perhaps surprisingly, the UK unambiguously exhibits the least inequality of the four countries; this applies to each of the components of wealth inequality and holds for all except the narrowest definition of "the rich" – in each of the of the five columns and in the first two parts of the table. Furthermore inequality between the rich and the non-rich groups is clearly larger than the inequality among the rich, again except for the narrowest definition of the rich.

[3] See also Jäntti et al. (2008) and OECD (2008), Chapter 10.

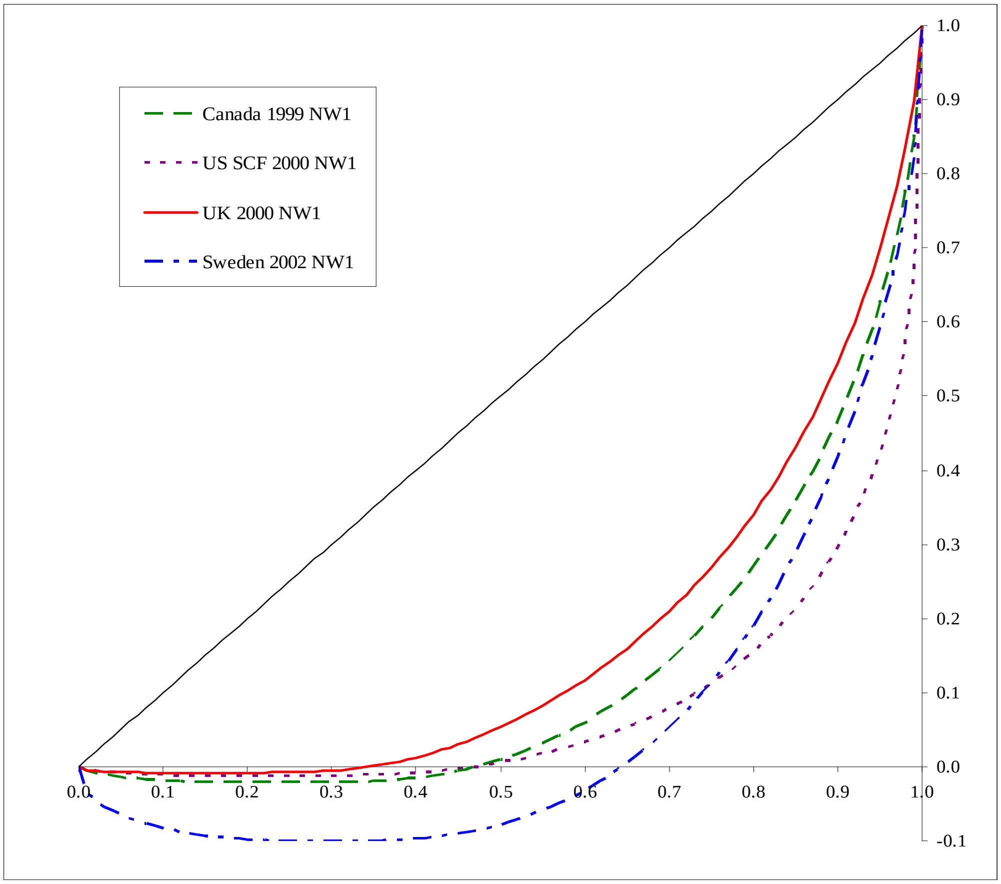

Figure 1: Lorenz Curves for Net Worth

3.2 Graphical representation

In the present case the most well-known graphical tool can be applied but it produces apparently strange results, as we can see if we examine the Lorenz curves in Figure 1. Net worth is negative where the slope of the curve is negative (within the zone where the curve passes below the horizontal axis.). How is one to interpret the substantial amount of negative net worth for some countries? Clearly, there could be some households who are in a desperate or precarious situation in terms of their long term wealth prospects; but there will be probably be many others for whom there is a less worrying interpretation. The wealth survey finds people at an arbitrary point in the life cycle and so it is to be expected that there will be some households in the sample that are currently in debt but whose long-term prospects are financially secure; they just happen to be observed at a point in their life where their mortgage debt is considerable or where they have not yet had sufficient years to have accumulated substantial resources. The extent to which such a household goes into debt will depend on the institutional arrangements for insurance and pension provision in old age – in the absence of state provision there will be greater need to save for one's own future. In view of this, in any given wealth survey, one could expect a significant proportion to report negative net worth, depending on the age structure and the institutions of the country in question. It is not surprising, for example, to see that the prevalence of negative values in Figure 1 is much higher for Sweden than for the United States in view of the substantial public pension provision in Sweden in contrast to the private arrangements in the US.

Of course, by focusing on inequality among the rich one can sidestep this problem of interpreting negative net worth. The picture of inequality among the top 10 percent (corresponding to the first case presented in Table 1) is provided in Figure 2. As we expect from the summary Gini statistics in Table 1), according to the surveys in the LWS database, the picture for the "rich" shows the UK to be unambiguously the least unequal, Sweden and Canada next (their Lorenz curves intersect) and the US most unequal.

However, we know that the tails of the wealth distribution may present difficulties of analysis and interpretation arising from data problems. So the question arises whether appropriate modelling of the upper tail of the wealth distribution would alter the simple conclusions about wealth inequality that we may be tempted to draw from Table 1 and Figures 1 and 2: is the overall wealth-inequality picture being distorted by errors in computing inequality among the wealthy?

4 Modelling the wealth distribution

What can be done to address the issue of missing or otherwise imperfect data in the upper tail? One possibility is to use a functional form to "patch in" the missing data.[4] Here, in common with other studies of the upper

[4] This is the "semiparametric" approach discussed in Cowell and Victoria-Feser (2008).

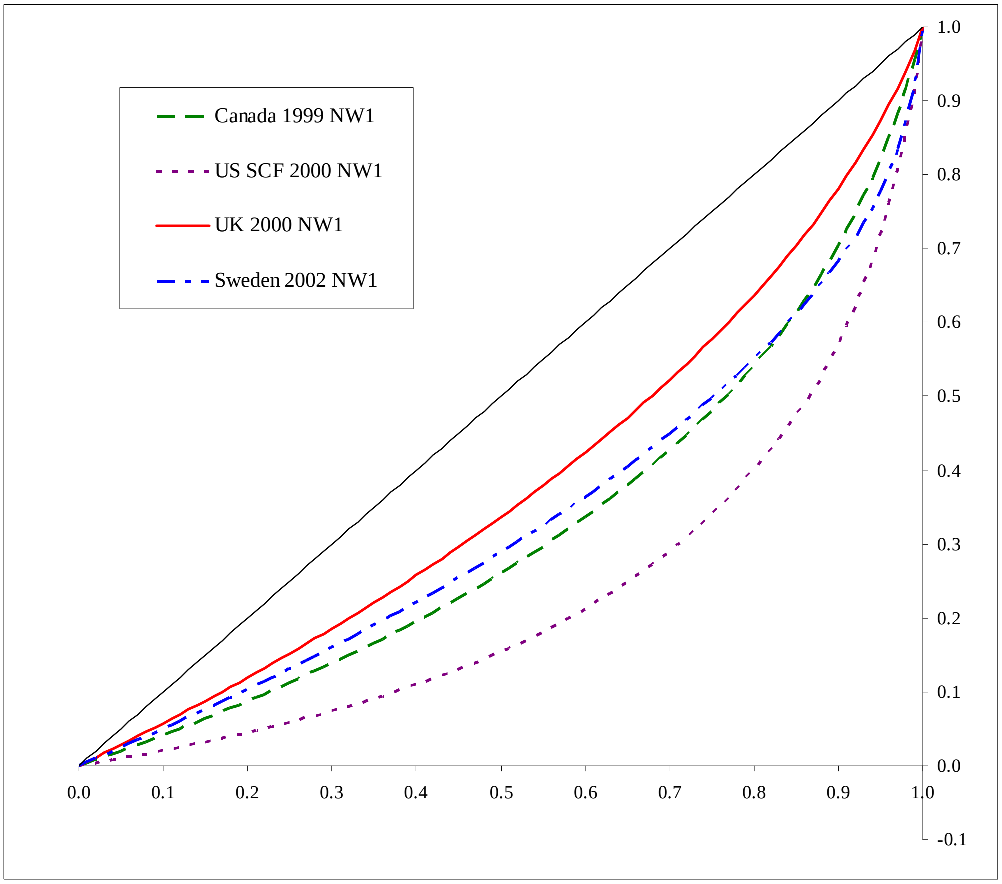

Figure 2: Lorenz Curves for Net Worth – Top 10%

tail of the wealth distribution we use the Pareto distribution.[5] The best way of introducing this widely accepted method of representing the wealth distribution among the rich is to follow through the graphical interpretation in Figure 3. This is a standard "Pareto diagram" – the horizontal axis in Figure 3 is net worth x plotted on a logarithmic scale and the vertical axis

[5]See, for example, Atkinson (1975), Atkinson and Harrison (1978), Clementi and Gallegati (2005), Cowell (2011), Johnson (1937), Klass et al. (2006), Soltow (1975) and Steindl (1965). Harrison (1981) discusses further how the distribution of income may be considered to be a combination of distributions, incorporating a Pareto tail. For the mathematics of the Pareto distribution see Kleiber and Kotz (2003).

is $P(x)$,[6] the proportion of people with net worth greater than or equal to x (also on a log scale).

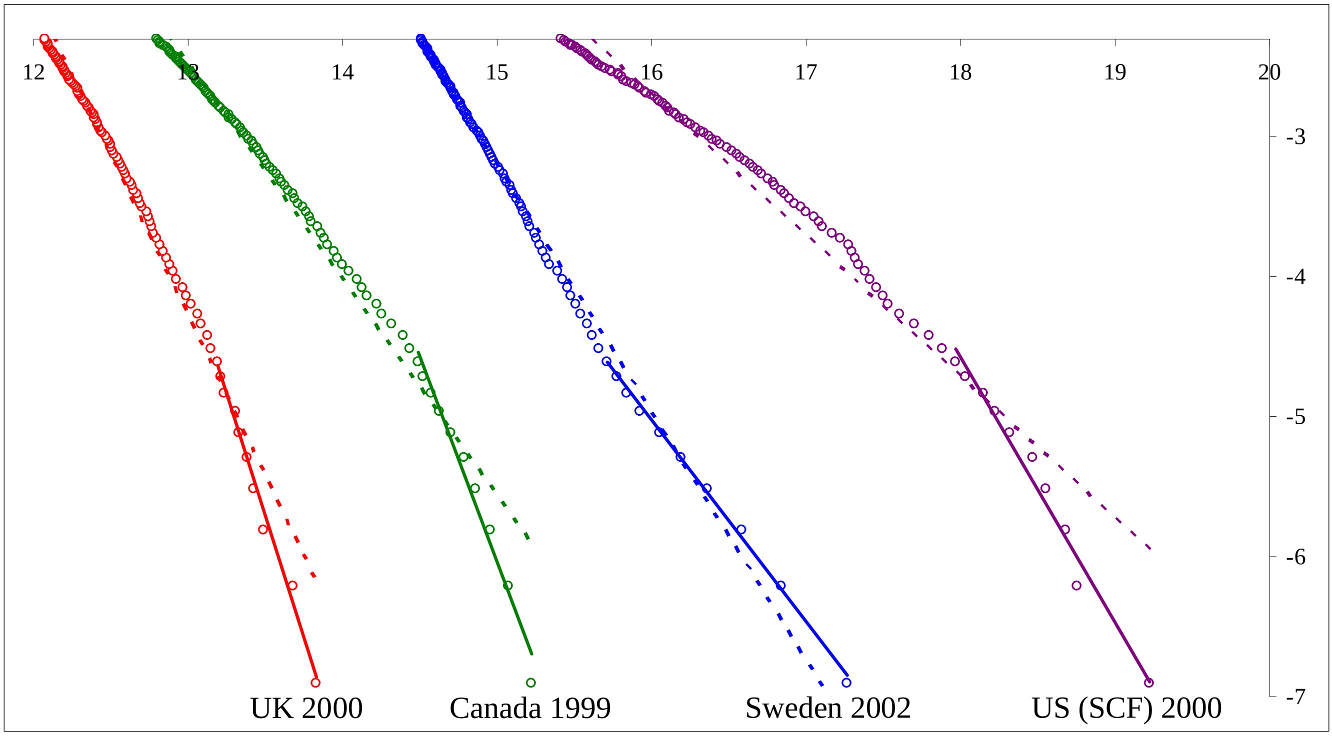

Figure 3: The Pareto model: a first look

The Pareto hypothesis is that there is a straight-line relationship:

$$\log P(x) = k - \alpha \log x \tag{6}$$

where the slope α is related to the inequality of the distribution and the intercept k is a location parameter. For an intuitive interpretation of α,

[6]Writing this in terms of the more familiar distribution function F we have $P(x) = 1 - F(x)$. Figure 3 is in fact four diagrams (for the separate countries) in one. Figures 4 to 13 for individual countries all have log of wealth on the horizontal axis and, since we do not need to compare wealth levels across countries, wealth has been left in the national currency. The same practice has been followed here: the common horizontal axis plots log wealth in each of the four national currencies.

consider an arbitrary reference level, or "base" level of wealth b. For the Pareto distribution it is true that the average wealth of all those with wealth at level b or more (the *conditional mean* for $x \geq b$) is given by

$$\frac{\alpha}{\alpha - 1} b. \tag{7}$$

So the "average/base" ratio is just

$$\beta := \frac{\alpha}{\alpha - 1}, \tag{8}$$

a constant for a true Pareto distribution. For any distribution this average/base idea gives a simple concept of inequality and we can see immediately from (7) that the higher is α, the lower is this ratio: high-inequality Pareto distributions have low values of α. We will discuss the inequality associated with the Pareto model further in section 5.

If we plot the LWS net-worth data so as to take a preliminary look at the Pareto hypothesis we can see that the straight-line model in this diagram is reasonable for these data, as indeed is usually the case for wealth data. But we can also see that for none of the four countries is the straight-line hypothesis completely satisfactory. The fitted Pareto model is also shown in Figure 3 where the regression line for the top 10% is depicted by a broken line and that for the top 1% by a solid line: it is clear that the estimate of α – and hence of wealth inequality – depends on the definition of the rich.[7]

[7] Details are given in Table 2 below.

4.1 How good is the model?

Because the estimates for α appear to depend on the size of the group that we consider as "the rich" it would be useful to have a feel for the suitability of the Pareto model. Simply taking the R^2 from the estimated regressions is not going to be helpful because, with highly skewed empirical distributions, the R^2 statistic is almost certain to be very high and of little use in evaluating the model.

For modelling problems that require heavy tailed distributions (as in the case of wealth) the following straightforward procedure is illuminating. Take the n sample wealth values and arrange them in descending order, so that $x_{(1)} \geq x_{(2)} \geq ... \geq x_{(n)}$. Then for a given value of k (between 1 and $n-1$ inclusive) the quantity

$$H_{k,n} := \frac{1}{k}\sum_{i=1}^{k}\frac{x_{(i)}}{x_{(k+1)}} \tag{9}$$

is a consistent estimator for $1/\alpha$ (Drees et al. 2000). Of course this immediately raises a further question: what value of k should be taken, in other words how many values to include in the estimate? The usual approach is to sidestep this issue and instead provide a graphical representation for a wide range of values of k: the appropriate value of α can be inferred if there is a "window" within which the estimates of (9) appear stable. An example of the standard representation is provided in Figures 4-7 which plot the inverse of (9) against k. This means that on the vertical axis in each graph we have an estimate of α and on the horizontal axis the number of observations used to provide the estimate (counting from the top) or, equivalently the cut-off wealth value used. In each figure the point estimate and a 95% confidence interval is provided.

For Canada and the US it is clear that there is a substantial window within which the relationship in (9) is stable, but that for k less than about 350 the relationship breaks down and the estimate of α increases sharply. In the case of Sweden the Hill plot again shows remarkable stability for k greater than 500, but then the estimate moves fairly sharply in the opposite direction from that of Canada or the US. For the UK data it is clear that the curve rises steadily as k falls which suggests that for the BHPS data we do not have a true Pareto distribution. Nevertheless, the Pareto model may be a sufficiently good approximation to the model generating the data to provide reasonable estimates of inequality.

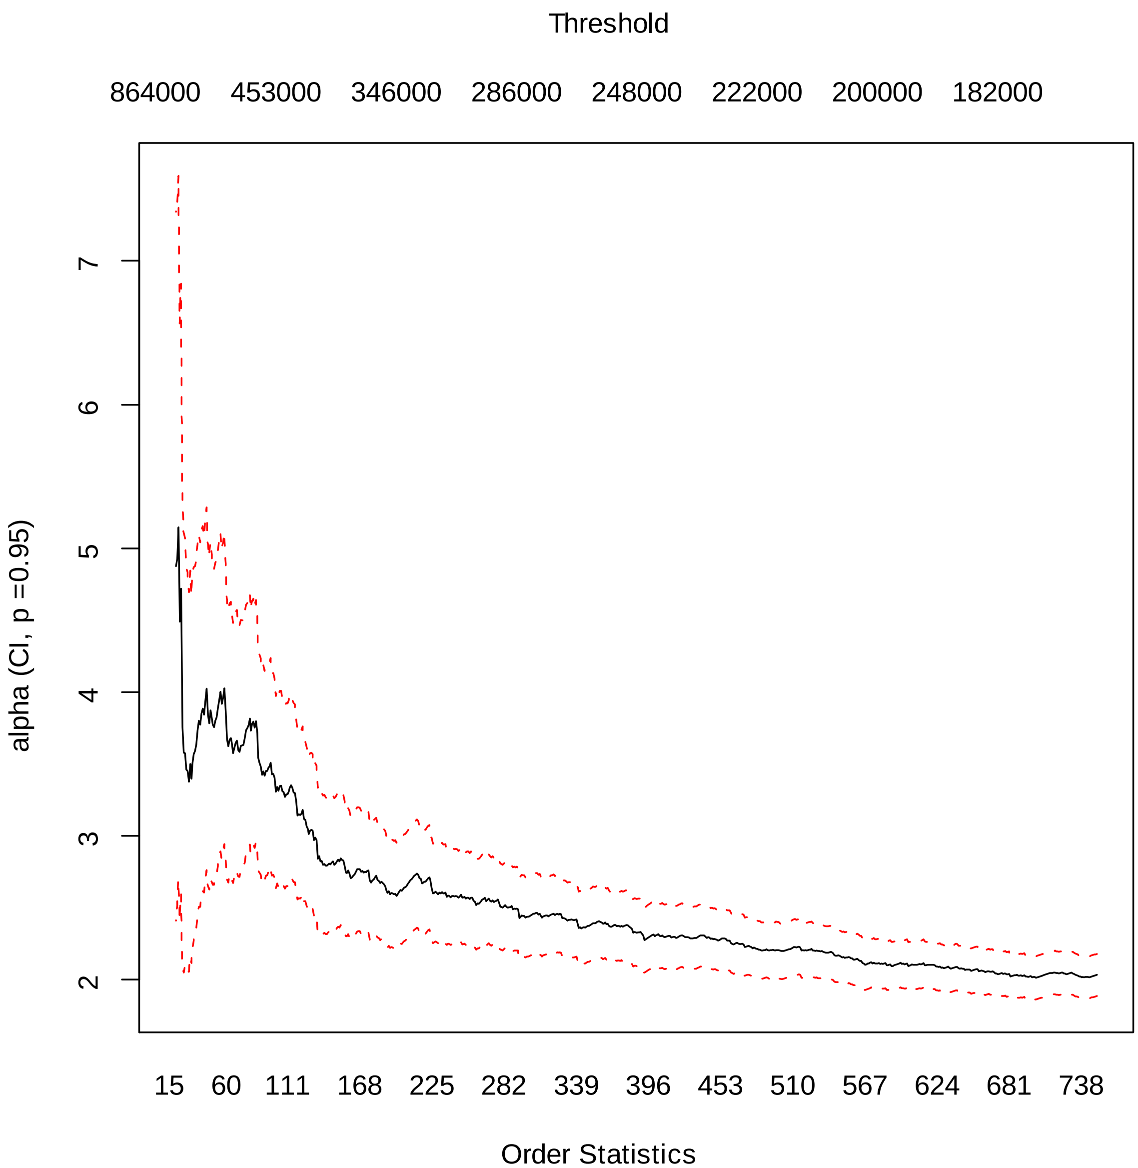

Figure 4: Hill Plot, UK 2000

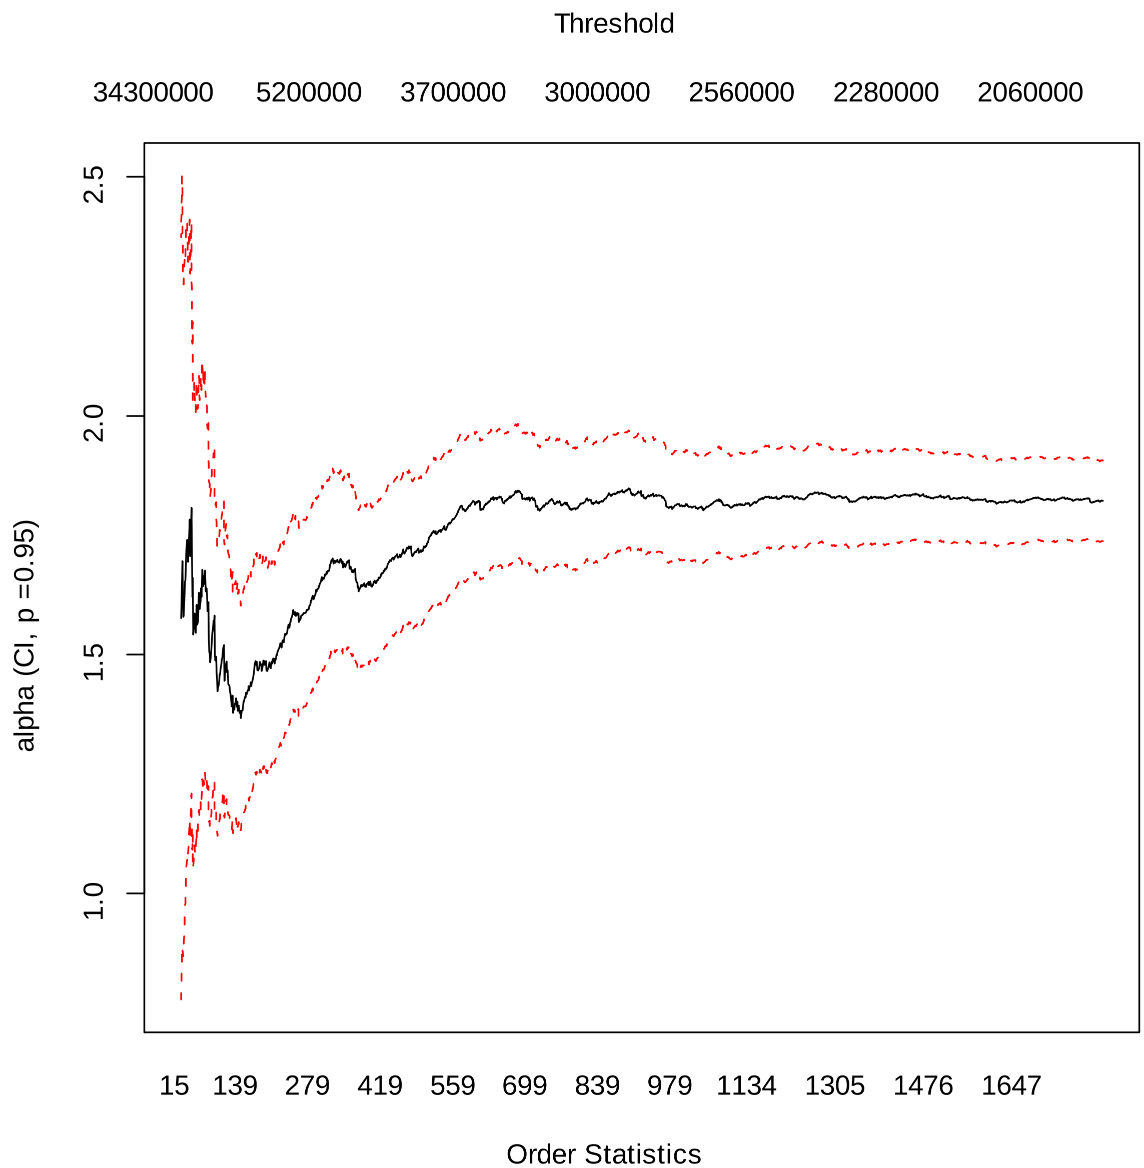

Figure 5: Hill Plot, Sweden 2002

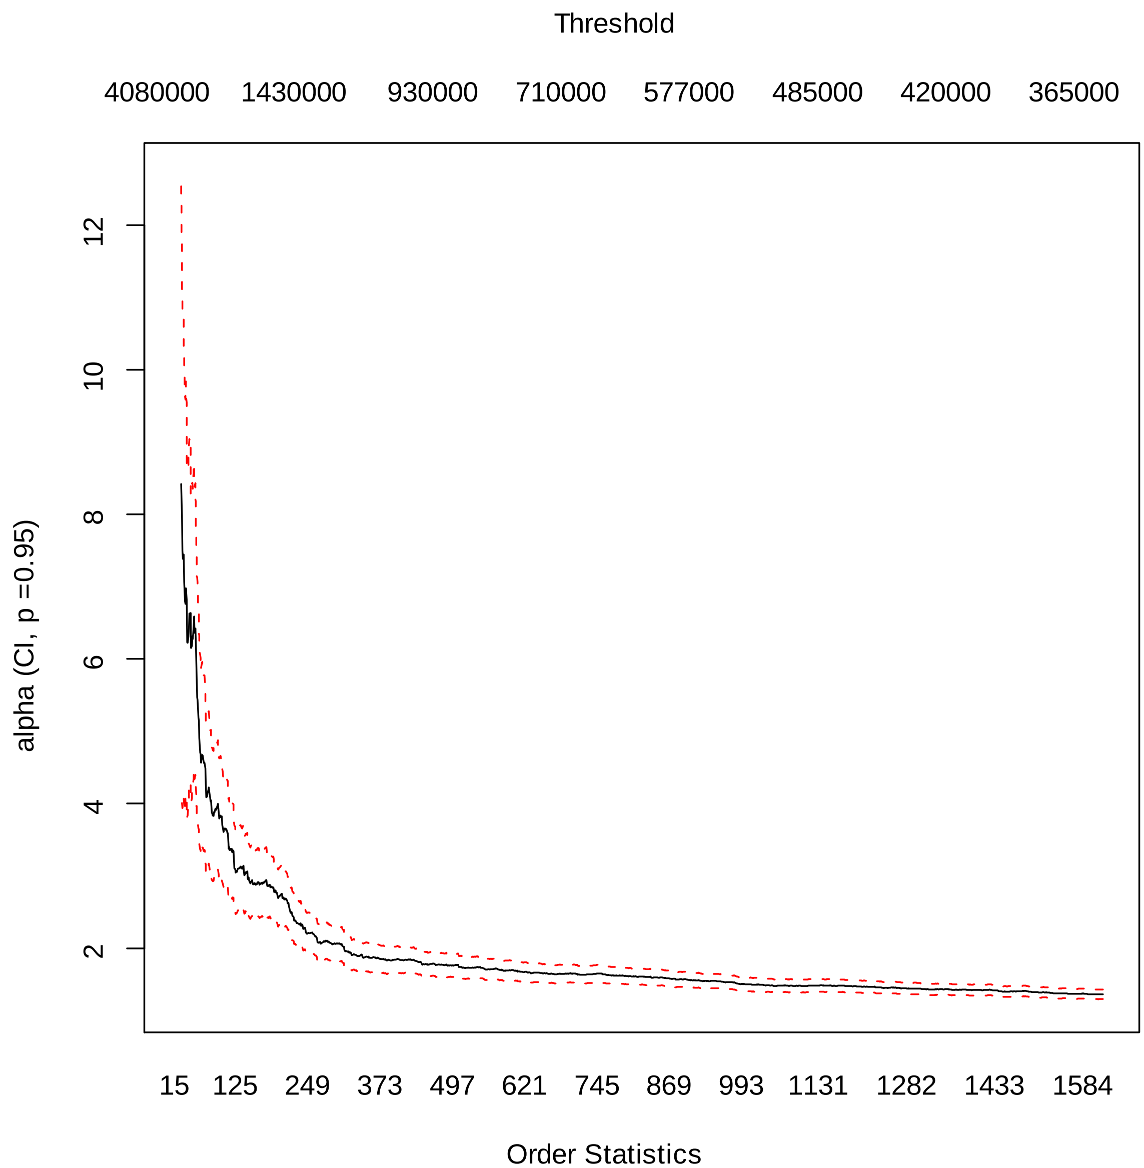

Figure 6: Hill Plot, Canada 1999

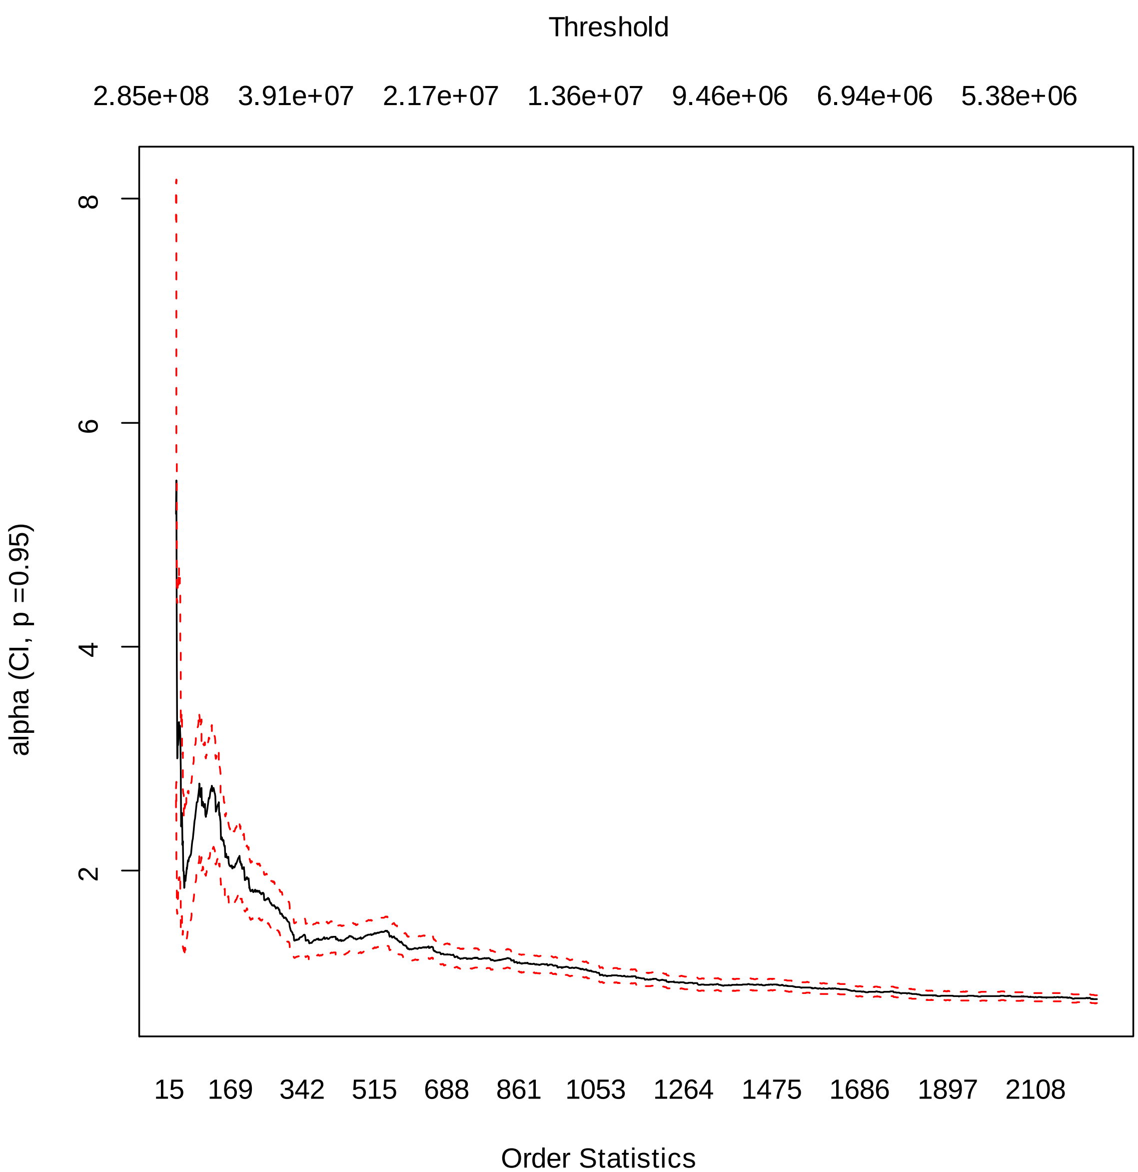

Figure 7: Hill Plot, US (SCF) 2000

	Ordinary Least Squares			Robust estimation		
	$p_{\mathrm{R}} = 0.10$	$p_{\mathrm{R}} = 0.05$	$p_{\mathrm{R}} = 0.01$	$p_{\mathrm{R}} = 0.10$	$p_{\mathrm{R}} = 0.05$	$p_{\mathrm{R}} = 0.01$
UK	2.092	2.232	0.422	1.942	2.092	0.559
Sweden	1.634	1.410	0.278	1.852	1.855	0.372
Canada	1.536	1.590	0.520	1.333	1.379	0.568
US	1.026	1.142	0.538	0.768	1.038	0.705

Table 2: Estimates of Pareto's alpha for different definitions of the Rich

4.2 The influence of outliers

Because of the apparent sensitivity of results to precise assumptions made about the upper tail of the wealth distribution it seems advisable to investigate alternative ways of fitting the statistical model to the data. Here we use two methods: (1) Ordinary Least Squares (OLS), the method used to generate the simple illustration in Figure 3, (2) robust estimation, a procedure that downweights outliers. The estimates are summarised in Table 2.

For easy comparison Figures 8-11 also show the outcome of least squares (solid line) and robust estimation (broken line) for the four countries where the Pareto model has been fitted to the top 10% group. Notice that in each case the robust method "pulls" the regression line away from the data points in the extreme right hand corner: this is exactly what is meant by the downweighting. So the robust estimate of α is higher than the OLS estimate in the case of Sweden and the robustly estimated α is lower than the OLS estimate in each of the other countries. This has an important implication if we use a semi-parametric approach to "patch" the top end of the wealth distribution (Cowell and Victoria-Feser 2008). For the UK, Canada and the US, if we estimate the model robustly this will imply higher within-group inequality for the wealthy if we use standard OLS; for Sweden the opposite holds. More on this in Section 5.

For the UK the following two figures also show how the relationship between OLS and robust regression methods alters when we change the size of the "rich" group. For the top 5% the effect of switching from OLS to robust regression is similar to the 10% case. However, if we fit the Pareto model to the top 1% the relationship is reversed and one can see clearly the role of the topmost data points.

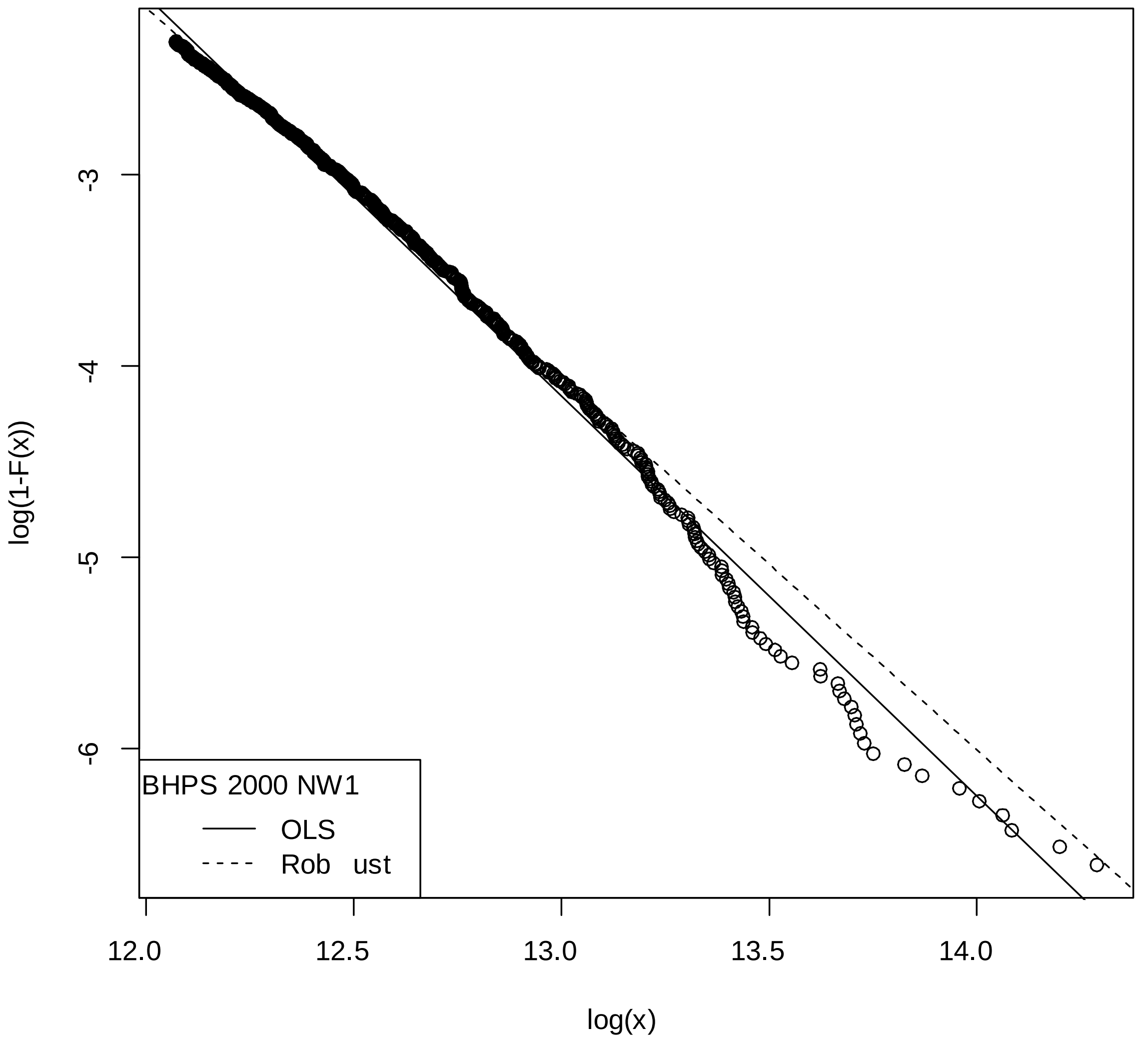

Figure 8: Pareto Model for Top 10%: UK 2000

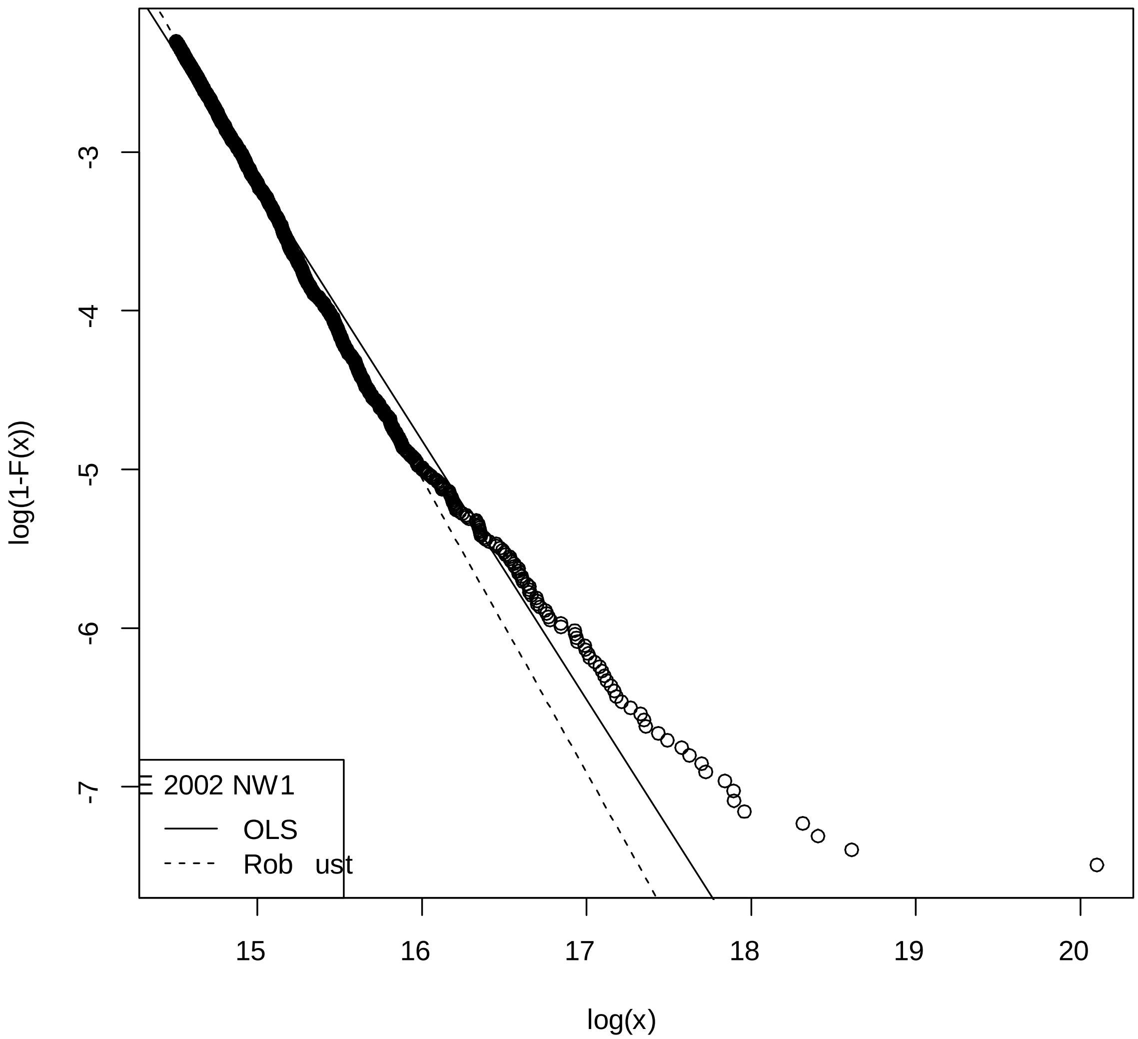

Figure 9: Pareto Model for Top 10%: Sweden 2002

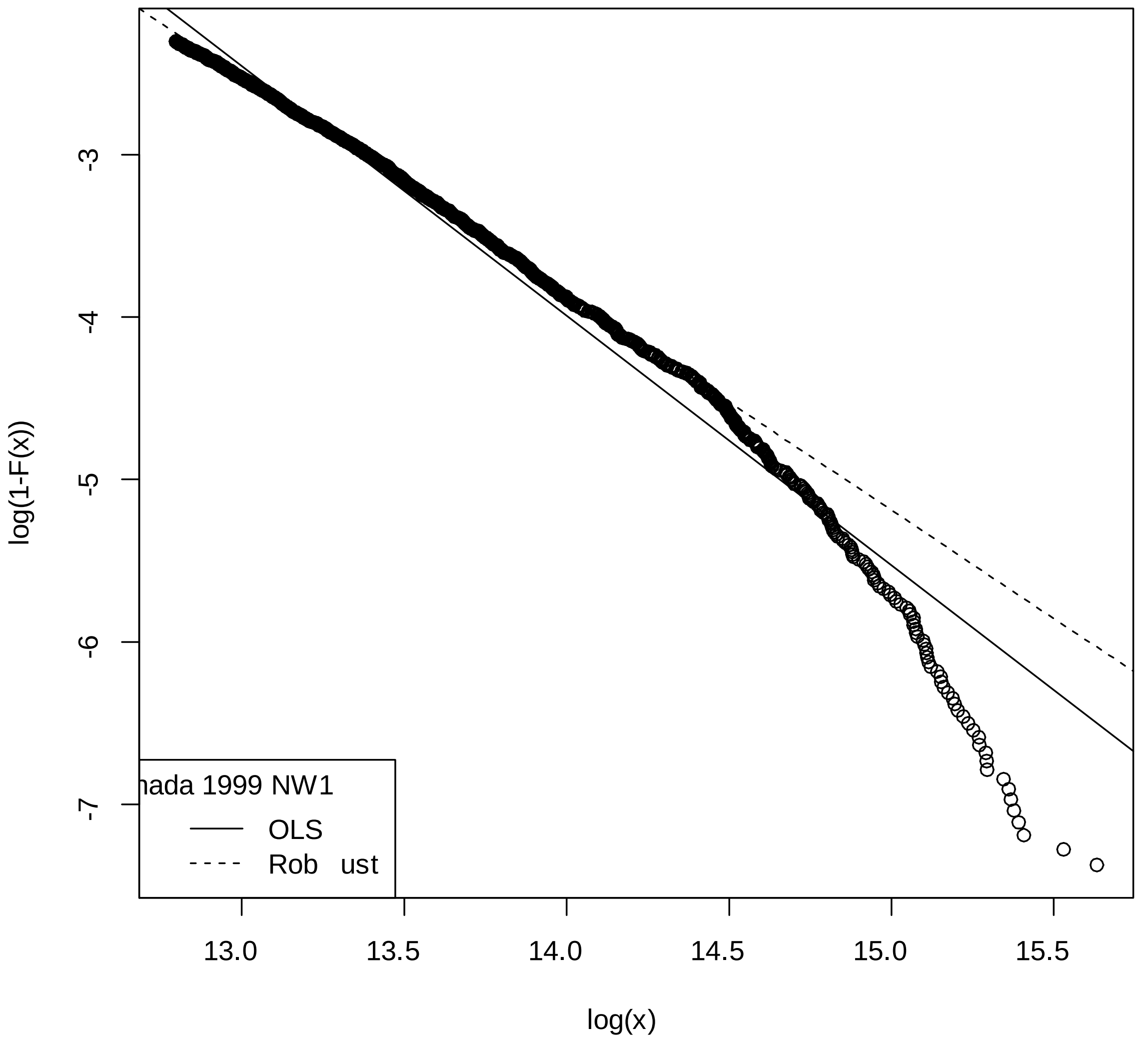

Figure 10: Pareto Model for Top 10%: Canada 1999

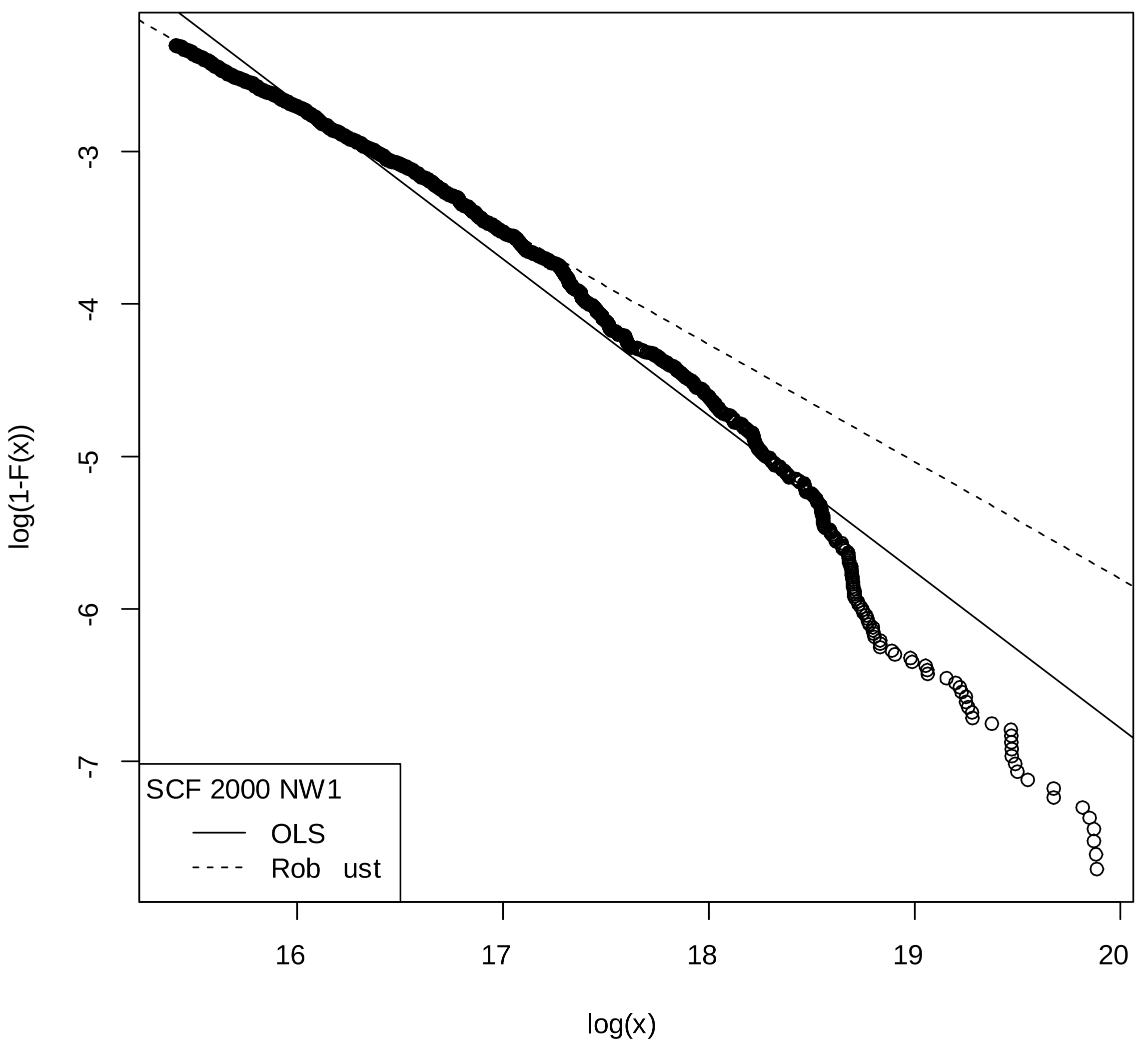

Figure 11: Pareto Model for Top 10%: US (SCF) 2000

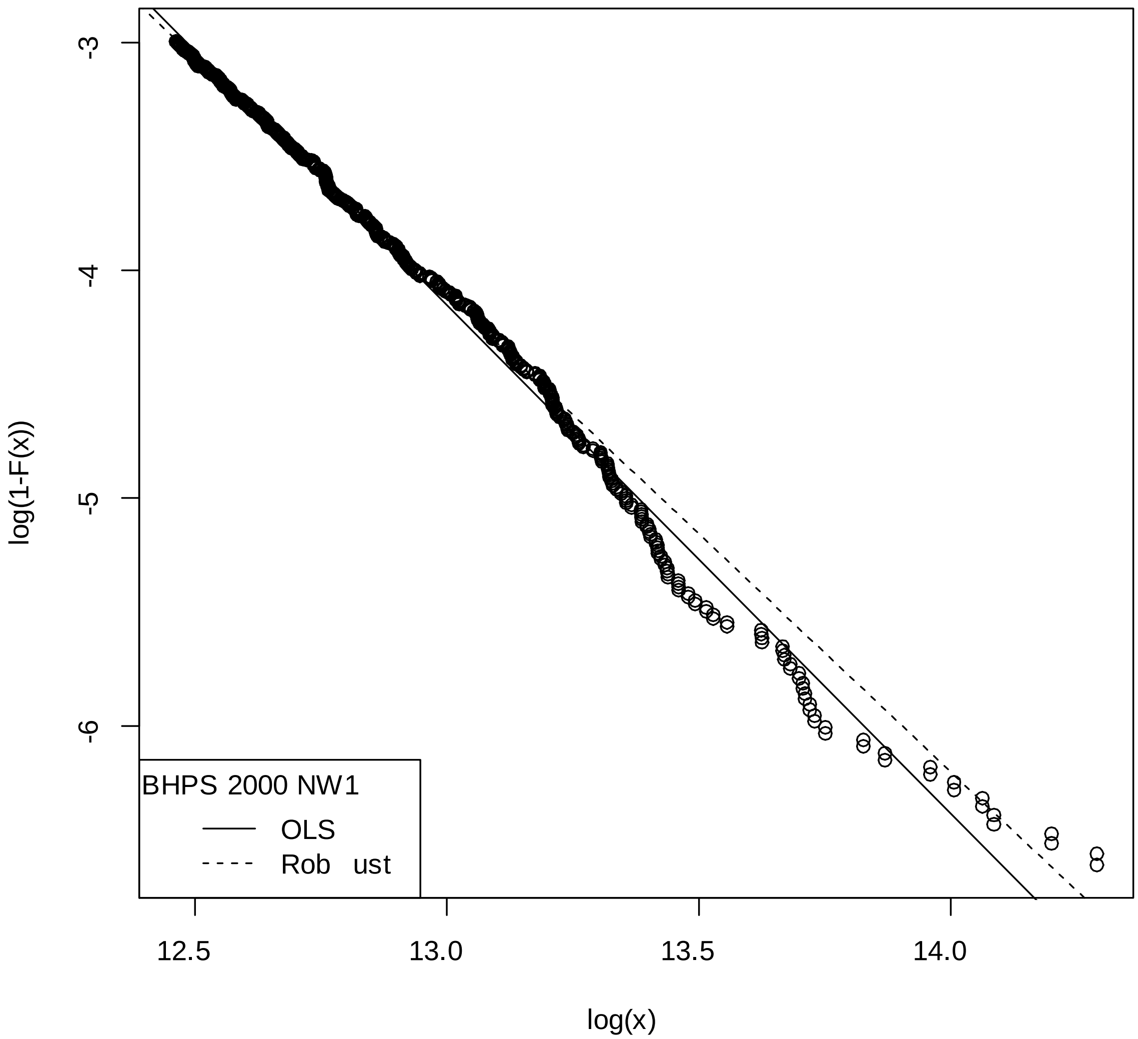

Figure 12: Pareto: UK, top 5%

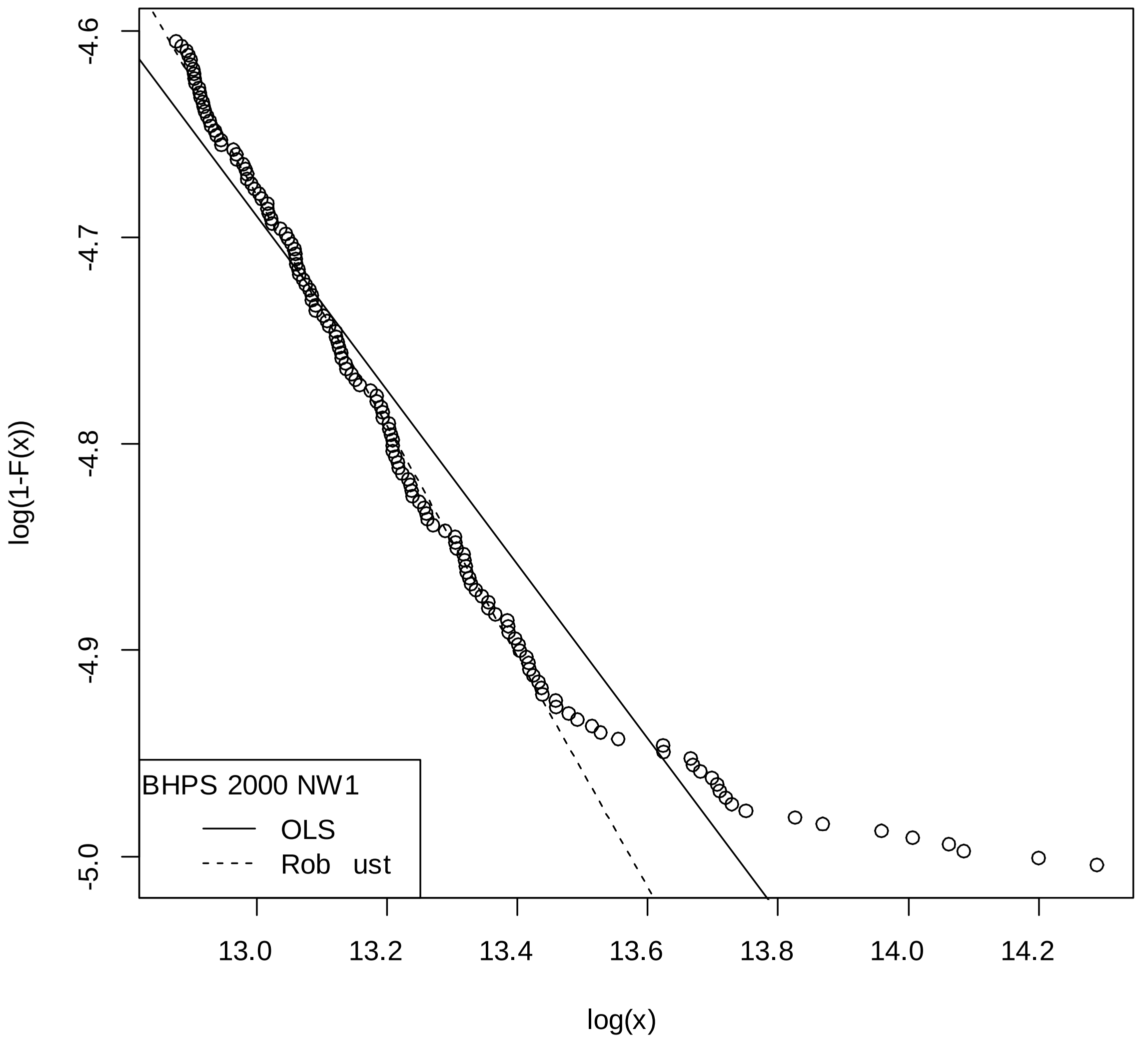

Figure 13: Pareto: UK, top 1%

5 Wealth inequality comparisons – a second look

We can now put the formal modelling to work. We can use the estimated Pareto distributions to model inequality among the rich so as to re-appraise the breakdown of wealth inequality in Table 1. Intuitively what we are doing is "splicing" the estimated parametric model of the rich into the empirical wealth distribution and then recomputing inequality as measured by the Gini coefficient. Obviously the overall effect both on how we interpret the "rich" in this exercise and on how we fit the model: the econometric method will affect inequality estimates (Cowell and Victoria-Feser 2007).

5.1 The Lorenz curve

First let us redraw the Lorenz curves using the model to patch the upper tail. Again write $P(x)$ for the proportion of the population with wealth greater than or equal to x. Analogously define $S(x)$ is the proportion of wealth held by those owning x or more. For a Pareto distribution the relationship between $P(x)$ and x is given by equation (6) and the relationship between $S(x)$ and x is just

$$\log S(x) = \frac{k}{\beta} - \frac{\alpha}{\beta} \log x \tag{10}$$

where, as before, $\beta = \alpha / [\alpha - 1]$ is the "average/base" ratio. The Lorenz curve is therefore given by :

$$\log S(x) = \frac{1}{\beta} \log P(x) .$$

Figures 14-17 show the "top right-hand corner" of the Lorenz curves for the four countries for the raw data (as in Figure 1) and for the Pareto-modelled data using the two different ways of estimating α. We are focusing on the top 10% snapshot of the whole distribution rather than looking at the top 10% as a sub-population as we did in Figure 2.

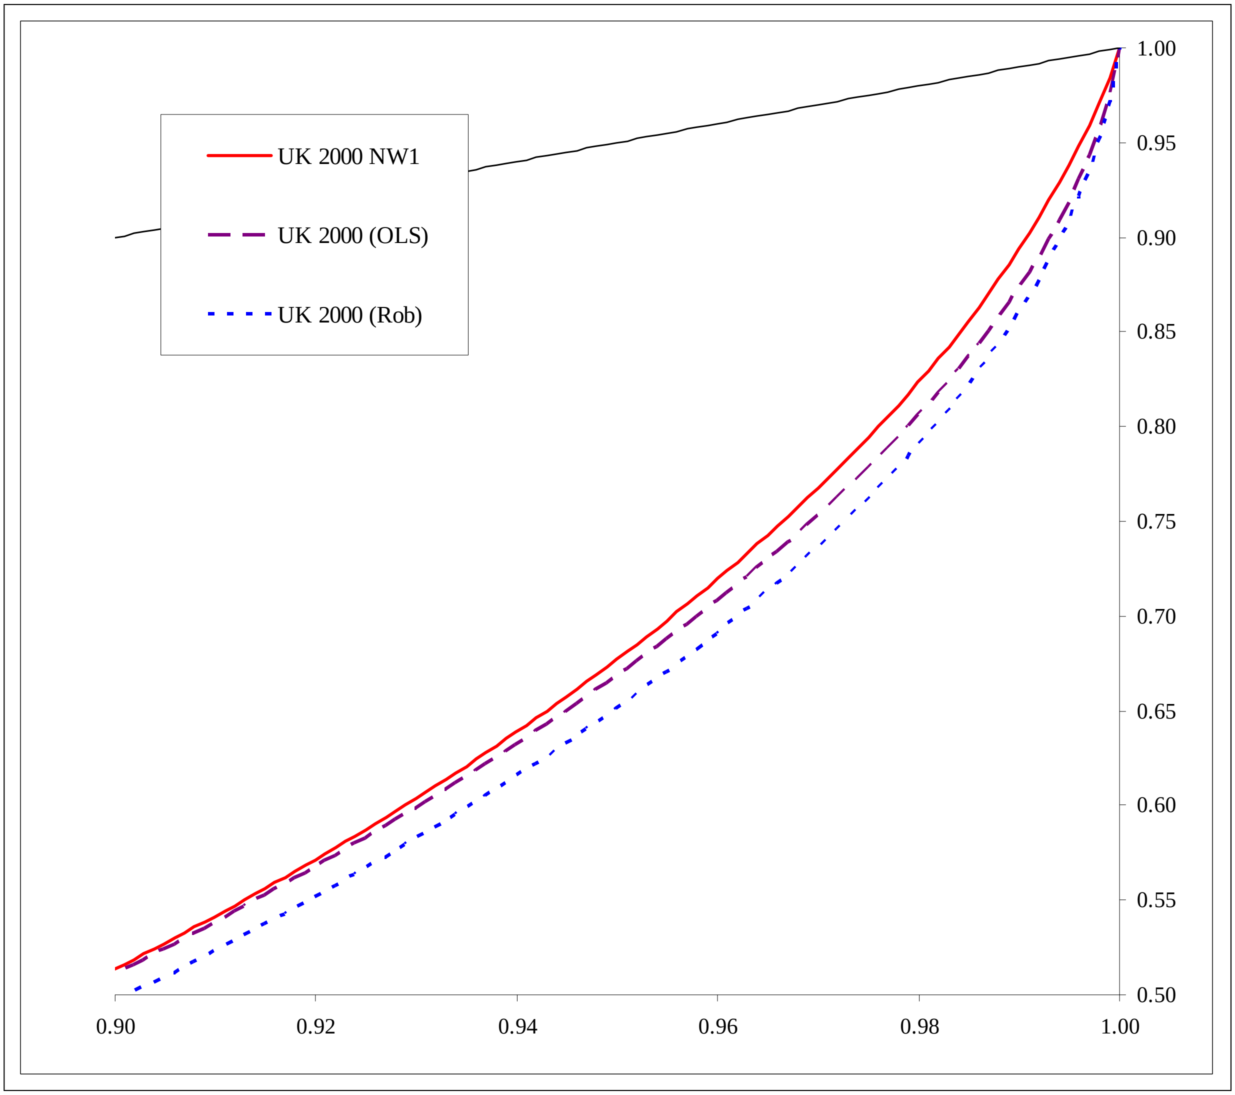

Figure 14: Raw Data and Pareto Model: UK 2000

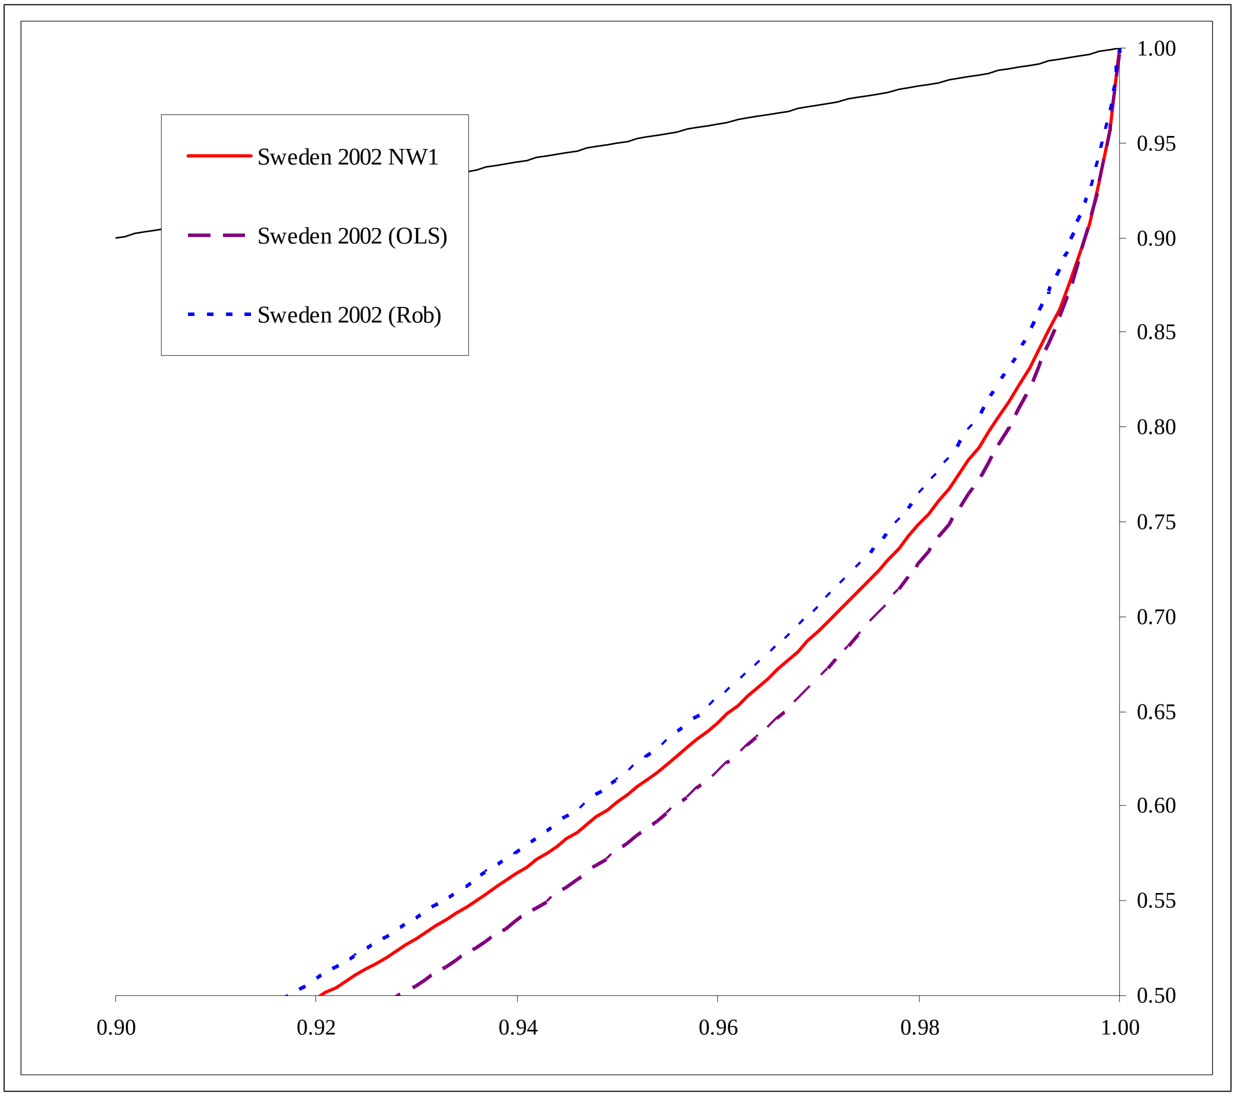

Figure 15: Raw Data and Pareto Model: Sweden 2002

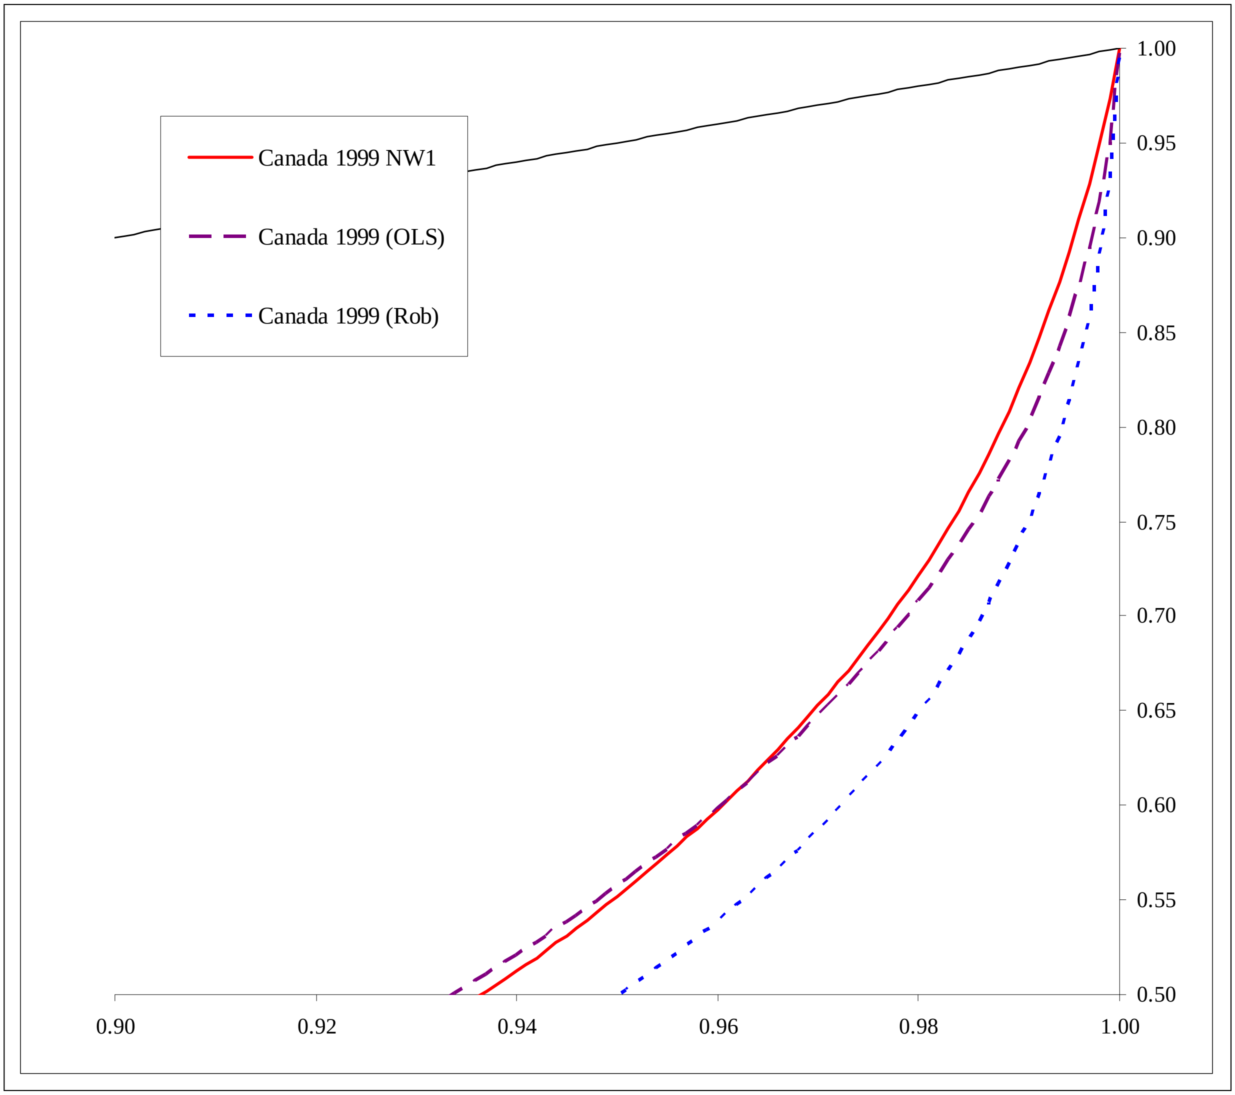

Figure 16: Raw Data and Pareto Model: Canada 1999

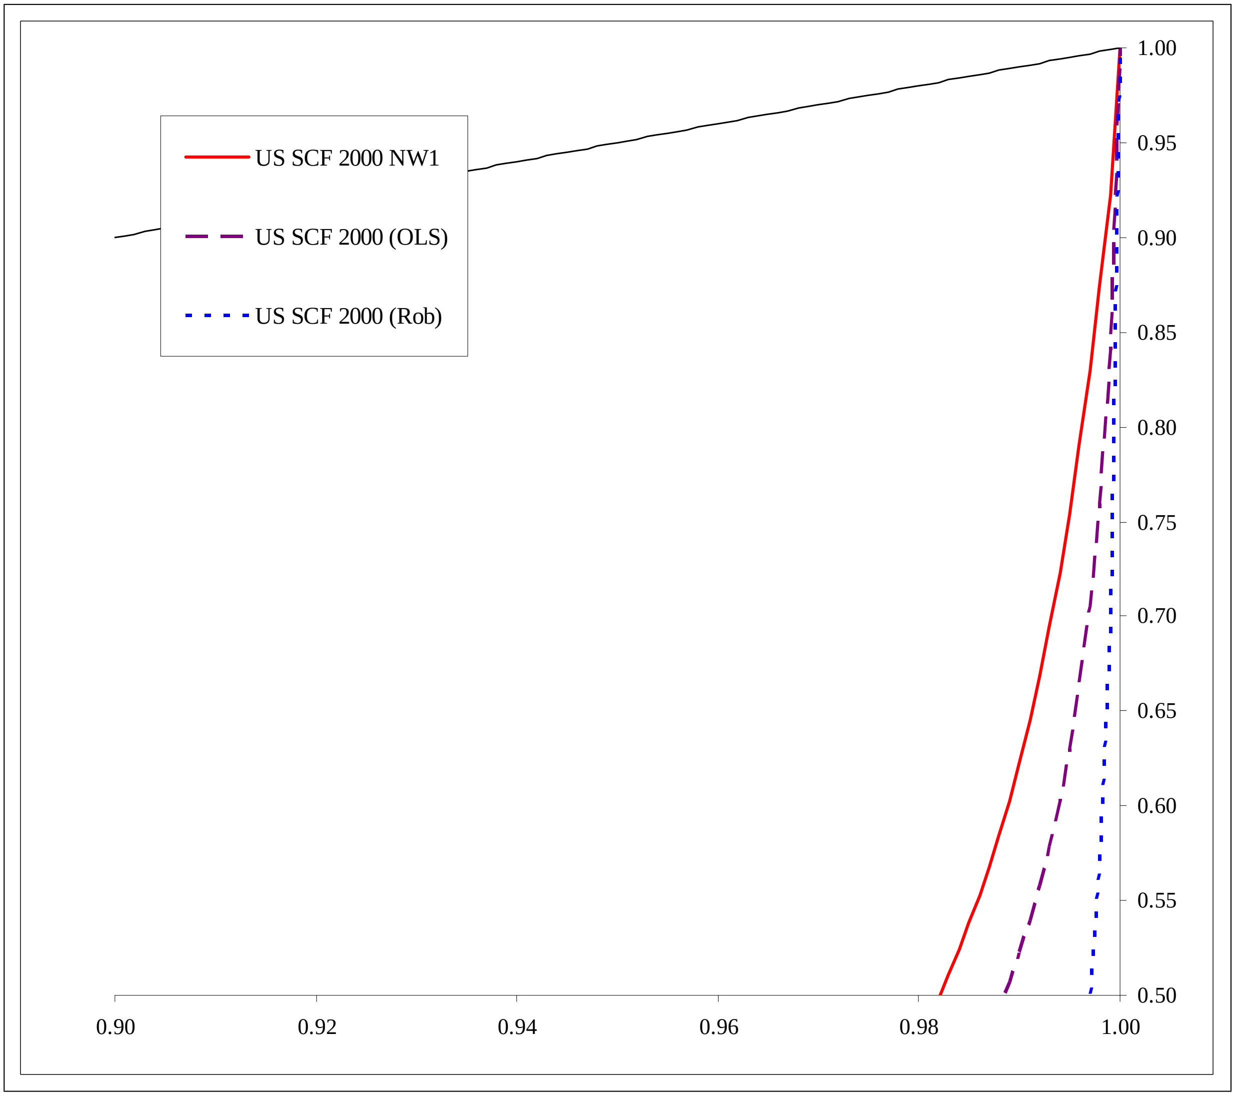

Figure 17: Raw Data and Pareto Model: US (SCF) 2000

In each case two of the Lorenz curves are easy to compare: if, for any two true Pareto distributions with coefficients α_1 and α_2, it is true that $\alpha_1 > \alpha_2$ then distribution 1 must Lorenz-dominate distribution 2 (Cowell 2011). So, using the fact that for Sweden the robust estimate of α is higher than the OLS estimate, the robust estimate of the Lorenz curve in Sweden must lie inside the OLS estimate; the reverse is true for the UK, Canada and the US. We can also see that in the case of the USA, which has very high inequality among the rich, the precise assumption made about the rich is crucial: at the top of the diagram the two modelled Lorenz curves lie far apart on either side of the Lorenz curve for the raw data. The other cases are much less extreme with the modelled Lorenz curves intersecting the raw-data Lorenz curves, usually near the top.

5.2 Recomputing inequality

It is important to see how the estimates of inequality and its components are affected by the modelling procedure, especially in view of the fact that the raw-data and modelled Lorenz curves intersect. There are two main effects to consider. The first effect has already been alluded to in Section 4.2 – the modelled distribution will affect the estimate of the within-group Gini for the top group. The second effect is on the wealth share of the top group, which in turn affects the between group inequality estimate – see (4).

5.2.1 Within group

This can be disposed of very briefly. We have just seen that the Lorenz curves for the Pareto distribution are ranked in the same order as the α values. In fact the Gini coefficient for a Pareto distribution with just $\frac{1}{2\alpha-1}$. So the within-group inequality if we model parameter the top $100p_{\mathrm{R}}\%$ of the distribution as a Pareto with estimated parameter α is

$$G_{\mathrm{R}}^{*} = \frac{1}{2\alpha - 1} \tag{11}$$

5.2.2 Between group

Let x_{R} be the p_{R}-quantile – the income boundary between the rich and non-rich groups. Using (6) we can see that the lower boundary x_{R} of the top

$100p_{\mathrm{R}}\%$ is found from the following:

$$\log x_{\mathrm{R}} = \frac{k - \log p_{\mathrm{R}}}{\alpha} \tag{12}$$

Then. using (7), mean wealth for the rich is

$$\mu_{\mathrm{R}}^* = \frac{\alpha}{\alpha - 1} x_{\mathrm{R}}$$

Mean wealth overall therefore becomes

$$\mu^* = p_{\mathrm{R}}\mu_{\mathrm{R}}^* + p_{\mathrm{N}}\mu_{\mathrm{N}} \tag{13}$$

and wealth shares in the rich and non-rich groups are now:

$$s_{\mathrm{R}}^* = \frac{p_{\mathrm{R}}\mu_{\mathrm{R}}^*}{\mu^*} \tag{14}$$

$$s_{\mathrm{N}}^* = \frac{p_{\mathrm{N}}\mu_{\mathrm{N}}}{\mu^*} \tag{15}$$

– compare this with equations (1)-(3). Therefore, between-group inequality with the modelled upper tail is now

$$G_{\mathrm{B}}^* = s_{\mathrm{R}}^* - p_{\mathrm{R}}. \tag{16}$$

5.2.3 Overall

Now let us combine both effects of the modelling to obtain total Gini inequality and its breakdown. Using the results for the new components of inequality (11), (14) and (11) we can compute adjusted overall inequality as follows

$$G^* = p_{\mathrm{R}} s_{\mathrm{R}}^* G_{\mathrm{R}}^* + p_{\mathrm{N}} s_{\mathrm{N}}^* G_{\mathrm{N}} + G_{\mathrm{B}}^*. \tag{17}$$

– compare this with equation (5).

The results of the Pareto modelling on the structure of wealth inequality are shown in Tables 3 and 4. First, a technical note. When the Pareto model is fitted there is no restriction on the estimated value of α which may give rise to problems in interpreting the decomposition formula (17). If the the value is less than one then the mean of the Pareto distribution is undefined and, of course, the average/base ratio (8), the adjusted share of the rich (14) have no meaning either; if the value is less than 0.5, then the adjusted rich-group

	Gini overall(G)	Share rich(s^*_{R})	rich(G^*_{R})	Gini non-rich(G_{N})	between(G^*_{B})
			($p_{\mathrm{R}} = 0.10$)		
UK	0.691	0.505	0.314	0.607	0.405
Sweden	0.904	0.614	0.441	1.045	0.514
Canada	0.799	0.635	0.483	0.710	0.535
US	0.988	0.982	0.951	0.779	0.882
			($p_{\mathrm{R}} = 0.05$)		
UK	0.687	0.352	0.289	0.618	0.302
Sweden	0.915	0.533	0.549	0.941	0.483
Canada	0.799	0.511	0.459	0.703	0.461
US	0.959	0.874	0.779	0.842	0.824

Table 3: Gini decomposition – adjusted with OLS Pareto model rich group

	Gini overall(G)	Share rich(s^*_{R})	rich(G^*_{R})	Gini non-rich(G_{N})	between(G^*_{B})
			($p_{\mathrm{R}} = 0.10$)		
UK	0.682	0.482	0.347	0.607	0.382
Sweden	0.892	0.537	0.370	1.045	0.437
Canada	0.817	0.660	0.600	0.710	0.560
			($p_{\mathrm{R}} = 0.05$)		
UK	0.669	0.308	0.314	0.618	0.258
Sweden	0.889	0.360	0.369	0.941	0.310
Canada	0.797	0.496	0.569	0.703	0.446
US	0.972	0.928	0.929	0.748	0.878

Table 4: Gini decomposition – adjusted with robust Pareto model

Gini (11) is undefined. In such cases the decomposition of inequality has been omitted from the tables. It is clear from Table 2 that for both OLS and robust estimation such problems arise everywhere when applied to the top 1% group and that the problem arises for robust estimates in the case of the US.

Comparing the remainder of the results in Tables 3 and 4 with Table 1 it is clear that modelling the rich (either the top 10% or the top 5%) increases within-group ("Rich") inequality, in the case of the US, dramatically so: US "rich" inequality becomes much larger than between-group inequality. For Sweden and the UK robust modelling reduces the between-group Gini coefficient but, for Canada and the US, modelling the rich always increases between-group inequality.

Does modelling the upper tail of the distribution change the provisional conclusion that we drew from Table 1, that the UK wealth is unambiguously less unequally distributed than the other countries? It is clear that, for each version of the Pareto model (top 10% modelled or top 5%, OLS or robust regression) the countries are always ranked in ascending order of inequality as follows: UK, Canada, Sweden, US. Moreover in terms of each separate component of inequality the UK is always the least unequal of the four countries and the US is always the most unequal (Sweden and Canada swap second and third places for inequality among the rich and between-group inequality).

6 Conclusion

The Luxembourg Wealth Study enables us to get a clear picture of the comparative structure of wealth inequality across countries. Using LWS it is clear see that the inequality of net worth in the UK at the time of the new millennium was unambiguously lower than in Sweden, Canada or the US. This result might have been because of under-recording at the top of the distribution, but the evidence suggests that this is not the explanation. The conclusion is robust under modelling of the upper tail in order to allow for incomplete data and for outliers that may not represent the "true" wealth distribution. remarkably it is also true for subgroup decompositions of wealth inequality that can be undertaken using the Gini coefficient.

References

Amiel, Y., F. A. Cowell, and A. Polovin (1996). Inequality amongst the kibbutzim. *Economica 63*, S63–S85.

Atkinson, A. B. (1975). The distribution of wealth in Britain in the 1960s - the estate duty method re-examined. In J. D. Smith (Ed.), *The Personal Distribution of Income and Wealth*. New York: National Bureau of Economic Research.

Atkinson, A. B. and A. J. Harrison (1978). *Distribution of Personal Wealth in Britain*. Cambridge University Press.

Clementi, F. and M. Gallegati (2005). Pareto's law of income distribution: evidence for Germany, the United Kingdom, and the United States. In A. Chatterjee, S. Yarlagadda, and B. K. Chakrabarti (Eds.), *Econophysics of Wealth Distributions*. Berlin: Springer.

Cowell, F. A. (2011). *Measuring Inequality* (Third ed.). Oxford: Oxford University Press.

Cowell, F. A. and M.-P. Victoria-Feser (2007). Robust stochastic dominance: A semi-parametric approach. *Journal of Economic Inequality 5*, 21–37.

Cowell, F. A. and M.-P. Victoria-Feser (2008). Modelling Lorenz curves: robust and semi-parametric issues. In D. Chotikapanich (Ed.), *Modeling Income Distributions and Lorenz Curves*, Chapter 13, pp. 241–255. Springer.

Drees, H., L. de Haan, and S. Resnick (2000). How to make a hill plot. *The Annals of Statistics 28*, 254–274.

Harrison, A. J. (1981). Earnings by size: A tale of two distributions. *Review of Economic Studies 48*, 621–631.

Jäntti, M., E. Sierminska, and T. Smeeding (2008). The joint distribution of household income and wealth: Evidence from the Luxembourg Wealth Study. OECD Social, Employment and Migration Working Paper 65, OECD, Paris.

Johnson, N. O. (1937). The Pareto law. *Review of Economic Statistics 19*, 20–26.

Klass, O. S., O. Biham, M. Levy, O. Malcai, and S. Solomon (2006). The Forbes 400 and the Pareto wealth distribution. *Economics Letters 90*, 290–295.

Kleiber, C. and S. Kotz (2003). *Statistical Size Distributions in Economics and Actuarial Sciences*. Hoboken. N.J.: John Wiley.

OECD (2008). *Growing Unequal? Income Distribution And Poverty In OECD Countries*. Organisation For Economic Co-Operation And Development.

Sierminska, E., A. Brandolini, and T. M. Smeeding (2006). The Luxembourg Wealth Study - a cross-country comparable database for household wealth research. *Journal of Economic Inequality 4*, 375–383.

Soltow, L. (1975). The wealth, income and social class of men in large northern cities of the United States in 1860. In J. D. Smith (Ed.), *The Personal Distribution of Income and Wealth*. New York: National Bureau of Economic Research.

Steindl, J. (1965). *Random Processes and the Growth of Firms: A Study of the Pareto Law*. New York: Hafner Press.