UK Wealth Inequality in International Context

Frank A. Cowell

STICERD, London School of Economics

December 2011

Abstract

Using the evidence from the Luxembourg Wealth Study it appears that the distribution of wealth in the UK is considerably less unequal than in Canada, the US or Sweden. But is this result an artefact of some peculiarities in the way components of net worth are reported? Or does it come from an underestimate of inequality among the wealthy and of the wealth differential between the rich and the rest? Using alternative definitions of wealth and 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.

■JEL: D31

■Keywords: wealth distribution

■Acknowledgements: I am grateful to the Nuffield Foundation for financial support and to John Hills for helpful comments on an earlier draft. Maximilian Eber and Paul Dolfen provided valuable research assistance.

■email: f.cowell@lse.ac.uk

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 chapter we will interpret these two questions as follows. First we address the question of how to make inequality comparisons of wealth 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 chapter is structured as follows. Section 2 discusses the international data source used to compare the UK situation with other European and North American countries. Section 3 discusses some important issues pertaining to wealth-inequality measurement and presents a first pass at the breakdown of inequality across countries. This initial look reveals some slightly surprising features and so Section 4 examines the inequality comparisons in more detail, focusing on alternative wealth concepts and a breakdown by population subgroups. Sections 5 and 6 show how a model of the upper tail of the wealth distribution may be used to clarify the international comparisons and to provide refined estimates of the wealth-inequality breakdown and Section 7 concludes.

2 The Data

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:

[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 the Treasury.

[2]For example, *w*hile the BHPS has 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 may affect the upper tail of the wealth distribution disproportionately.

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

Here and throughout the chapter "other investment property" is used synonymously with real estate other than the household's principal residence. Debt includes both home-secured debt (mortgages) and other forms.[3]

Let us see the importance of the different components of wealth in practice in the four countries, first taking the whole population of the country and then focusing on specific rich groups of the population. Table 1 presents the basic facts of the composition of total assets.

The entries in Table 1 can be read as follows. "Other investment property" formed 9 percent of household assets in the UK as a whole, 11 percent in Sweden, 13 percent in Canada and 17 percent in the USA; but if we focus just on the portfolios of the richest 10 percent, the proportion of assets held in this form was 15 percent for the UK, 17 percent for Sweden and so on. Overall there are three very striking points:

- The proportion of assets represented by the value of residence is consistently lower for the higher wealth groups. But this form of wealth holding is high for all groups in the UK: among the assets held by the top one percent slice of households the proportion of represented by the principal residence was 44 percent in the UK but only 4 percent in the US.

- Financial wealth represents a higher proportion of the total asset portfolio for the rich than for the general population (This is also true, with the exception of Canada, for other investment property).

[3] Again see Sierminska et al. (2006) for a detailed discussion of these wealth components. In the case of Sweden there is no separation between debt secured on one's home and other forms of debt.

	Whole Population	**Top 10%**	**Top 5%**	**Top 1%**
		Principal Residence		
UK	0.74	0.64	0.58	0.44
Sweden	0.61	0.48	0.42	0.23
Canada	0.64	0.41	0.32	0.20
US	0.45	0.13	0.07	0.04
		Other Investment Property		
UK	0.09	0.15	0.19	0.32
Sweden	0.11	0.17	0.21	0.32
Canada	0.13	0.20	0.22	0.11
US	0.17	0.23	0.22	0.26
		Financial Assets		
UK	0.17	0.21	0.23	0.24
Sweden	0.28	0.34	0.37	0.45
Canada	0.22	0.39	0.46	0.69
US	0.38	0.64	0.70	0.71
		Debt		
UK	0.21	0.08	0.07	0.06
Sweden	0.35	0.14	0.13	0.14
Canada	0.26	0.06	0.05	0.02
US	0.21	0.04	0.03	0.04

Source: LWS

Table 1: Proportions of total assets represented by main components of net worth: for whole population and for the rich

- Debt as a proportion of total assets is non-negligible for all countries and is very high in Sweden. This may have important consequences for the way in which we make wealth comparisons of the four countries.

We will come back to these points after a preliminary examination of wealth-inequality comparisons.

3 Wealth inequality: A first look

Should we measure the inequality of wealth the same way as we measure the inequality of other things? Clearly it would be helpful if we were to apply tools that are familiar and accepted in related 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. In our first look at inequality comparisons across the four countries we will briefly consider the use of inequality measures and of standard graphical presentations.

3.1 Tools: 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 the range of application of others.

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. But, as we have seen, debt represents at least one fifth of the value of total assets; so 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

	Gini Coefficient	Share in net worth of... top 10%	top 5%	top 1%
UK	0.665	0.456	0.301	0.101
Sweden	0.893	0.582	0.406	0.175
Canada	0.747	0.532	0.374	0.151
US	0.836	0.705	0.575	0.329

Source: LWS

Table 2: Inequality of net worth: overview

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 coefficient and supplement this with statistics on the shares in total net worth held by various key groups in the population – a summary of results is provided in Table 2.

At first sight there are two striking things about the results in Table 2. First the UK inequality of net worth appears relatively modest in the UK. It is true that the top ten percent own more than 45 percent of net worth, but in the US the top 10 percent own more than 70 percent; and while the top one percent own more than ten percent of net worth in the UK, the top one percent in the US own almost a third of net worth! Second the ordering of countries by inequality may not be what one might have expected before glancing at the figures: for each of the last three columns (representing the share of the richest in overall net worth) we have UK, Canada, Sweden, US in ascending; in the first column (representing overall inequality) there is a slight change of ordering so that Sweden emerges as the most unequal. What is going on?

3.2 Tools: Graphical representation

The most well-known graphical tool for representing inequality – the Lorenz curve – can be used to give a richer picture of the international context of UK wealth inequality. But again, as with Table 2, it appears to produce strange results, as we can see in Figure 1. Immediately one is struck by the fact that all of the Lorenz curves pass below the horizontal axis, substantially so in the case of Sweden. This follows from the large number of households that have negative net worth: net worth is negative where the slope of the curve is negative (within the zone where the curve passes below the horizontal axis).

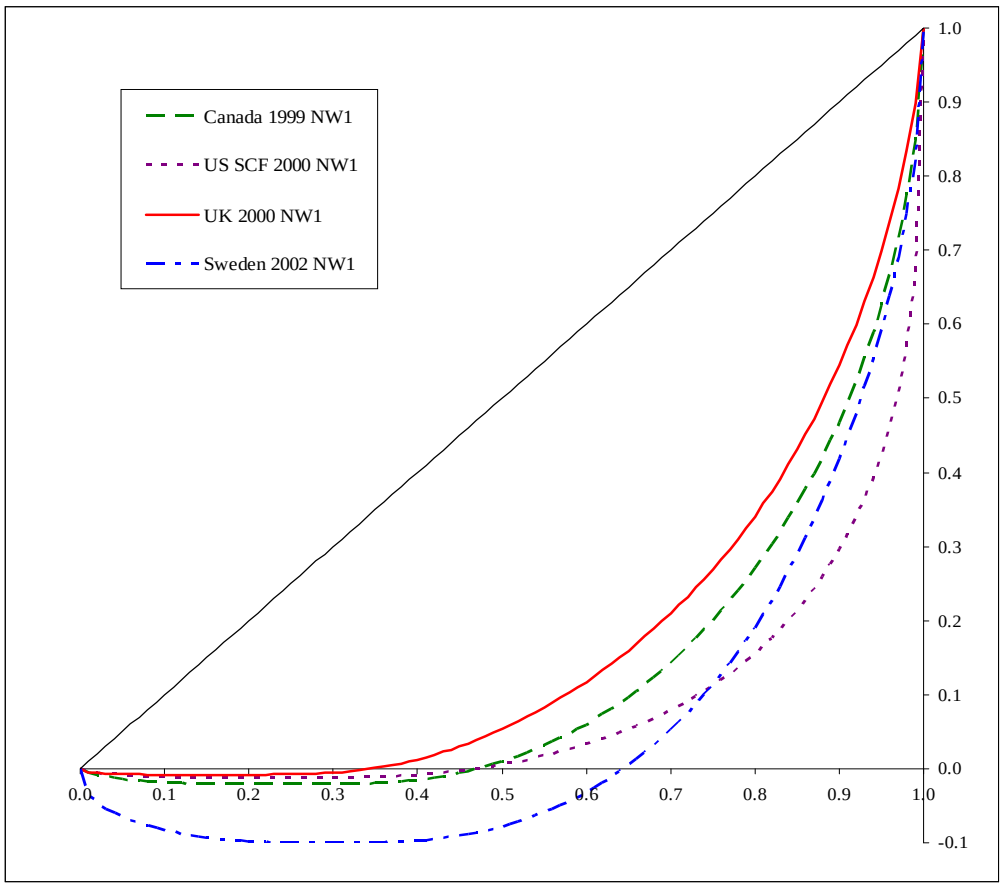

Figure 1: Lorenz Curves for Net Worth

How is one to interpret this? 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. The generous Swedish pension provision is not taken into account in the data, but it means that people do not need to save to provide for their old age; the net result may be that the rich possess a higher proportion of financial assets than in countries with less generous public pensions.

The fact that the Lorenz curve for Sweden intersects that for the US implies that, for some interpretations of inequality, Sweden has higher wealth inequality than the US (as we see from the Gini estimates in Table 2); for other interpretations the US will appear more unequal than Sweden. Should we take this at face value or is it an artefact of the unusual picture of debt in Sweden? Clearly this may merit closer examination and, along with this, we ought to see if the remarkable picture of the UK as the least unequal of the four countries also bears closer scrutiny.

4 Wealth inequality: A second look

The "debt puzzle" that emerged in the previous section suggests two possible ways forward. On the one hand it may be sensible to focus on parts of the wealth distribution that are likely to be less affected by debt. On the other we might wonder whether the net worth concept, although theoretically appealing, presents practical problems and whether a clearer picture could emerge if we looked at other wealth concepts. So, for a second look at

international comparisons we will examine more closely the composition of wealth inequality by groups in the population and by the main constituent parts of net worth.

4.1 Wealth inequality and wealth groups

Of course, by focusing on inequality only among the rich, one can sidestep this problem of interpreting negative net worth. Focusing on the rich may also yield additional insights on the wealth-inequality comparisons between countries. We will again adopt the pragmatic definition of the rich as a given percentage "top slice" of the distribution of net worth, but again we will try out more than one value for this given percentage.

Does the inequality ranking change as we focus on progressively more narrowly defined groups? The picture of inequality among the top 10 percent, 5 percent, 1 percent (corresponding to the three cases presented in Table 3) is provided in Figures 2-4. If we focus on the top 10 percent 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; if we narrow the focus down to the top 5 percent or the top 1 percent the relative position of the UK stays the same, but the rankings of the other countries change. More details on this are available from Table 3 below.

But just zeroing in on inequality within a narrowly defined group conveys only part of the story of wealth inequality broken down by wealth-defined groups. So let us broaden the focus to the inequality of the rich and its relation to overall wealth inequality. To do this we can use a standard method to give an exact decomposition of the Gini coefficient that is convenient when the distribution can be partitioned by wealth level.

The procedure is as follows. Given the definition of the rich, mean net worth in the population can then be expressed as the weighted average of the mean net worth of the rich and the non-rich groups (where the weights are the population proportions of the two groups). The Gini coefficient for the whole population is then the weighted sum of the rich Gini and the non-rich Gini plus between-group Gini; for each group, rich and non-rich, the correct weight is the group's population share times its wealth share; 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) is simply

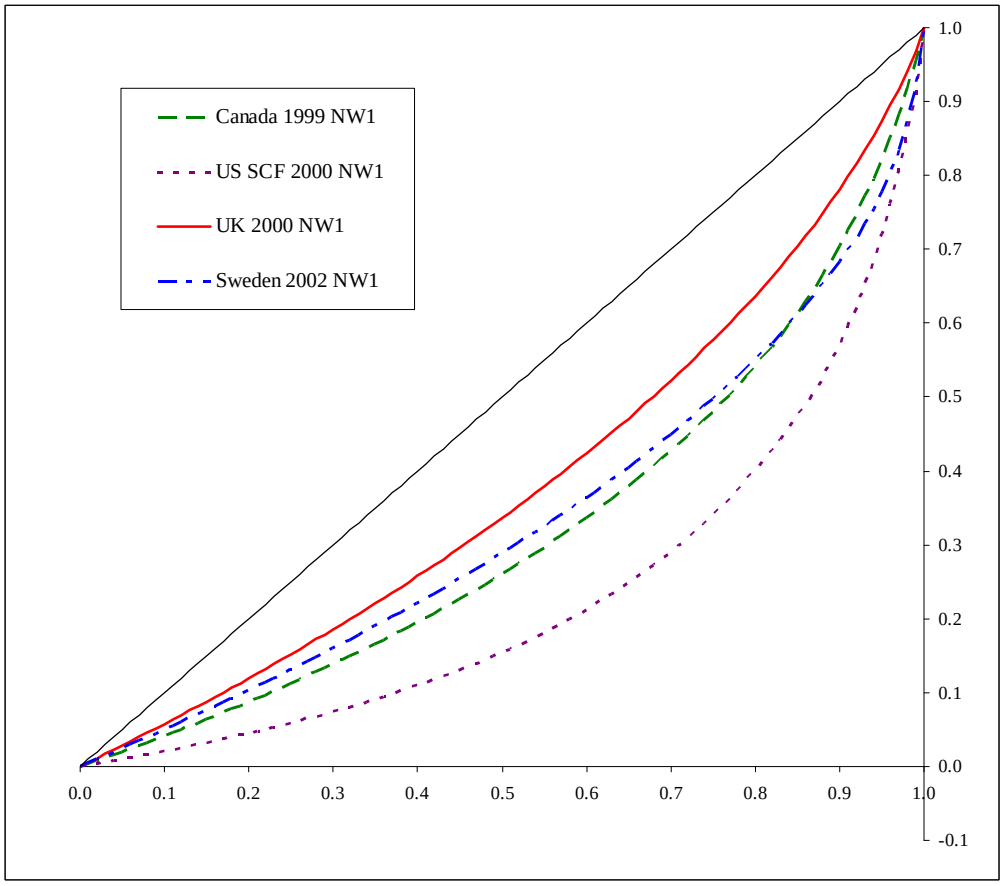

Figure 2: Lorenz Curves for Net Worth – Top 10%

the wealth share of the rich minus population share of rich.[4]

Table 3 uses this method to the data summarised in Table 2:[5] it gives the breakdown for the four countries in three cases corresponding to the

[4]The formal version of this breakdown of the totals into constituent parts can be expressed as follows. Let p_{R} and p_{N} be the proportions of the population considered as rich and non-rich respectively, where $p_{\mathrm{N}} = 1 - p_{\mathrm{R}}$. Correspondingly let μ, μ_{R}, μ_{N} be the mean net worth overall and in the two groups and G, G_{R}, G_{N} be the Gini coefficient for net worth overall and in each of the two groups. Then we have $\mu = p_{\mathrm{R}}\mu_{\mathrm{R}} + p_{\mathrm{N}}\mu_{\mathrm{N}}$ and $G = p_{\mathrm{R}}s_{\mathrm{R}}G_{\mathrm{R}} + p_{\mathrm{N}}s_{\mathrm{N}}G_{\mathrm{N}} + G_{\mathrm{B}}$, where $s_{\mathrm{R}} = p_{\mathrm{R}}\mu_{\mathrm{R}}/\mu$, $s_{\mathrm{N}} = p_{\mathrm{N}}\mu_{\mathrm{N}}/\mu$ (the shares of the rich and of the non-rich in overall net worth) and $G_{\mathrm{B}} = s_{\mathrm{R}} - p_{\mathrm{R}}$.

[5]Here is the working for the first row of Table 3: $0.665 = 0.1 \times 0.456 \times 0.260 + 0.9 \times (1 - 0.456) \times 0.607 + 0.356$.

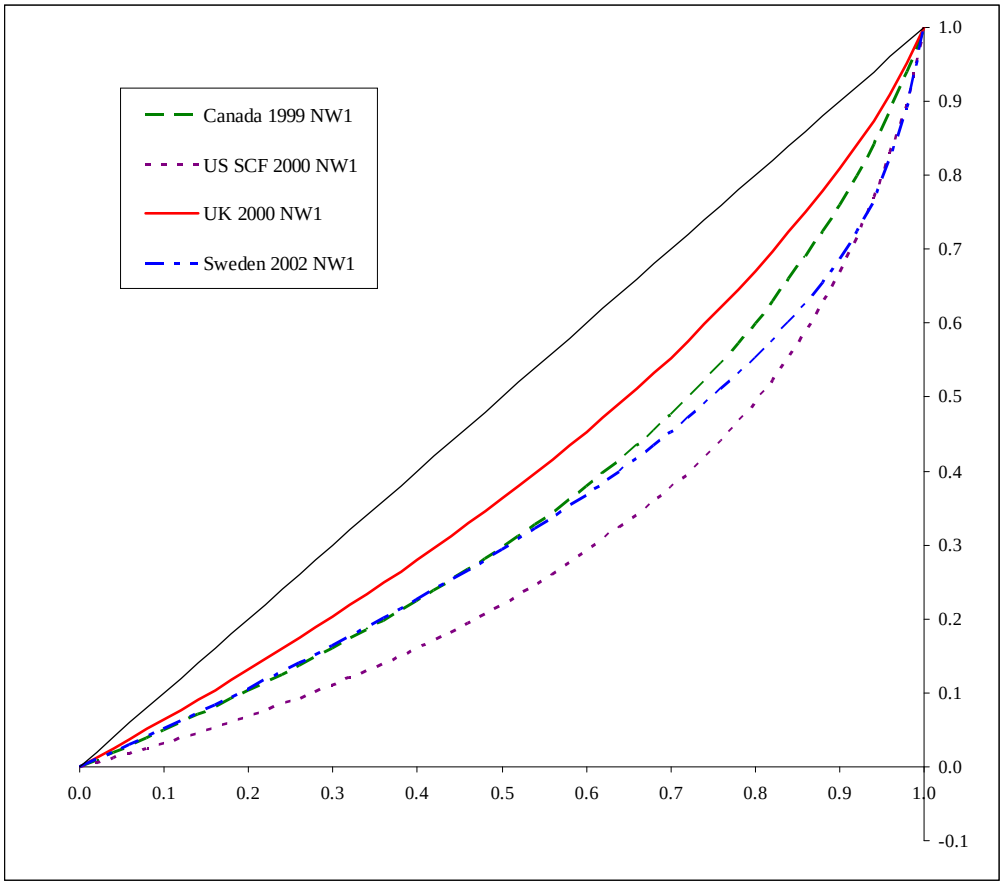

Figure 3: Lorenz Curves for Net Worth – Top 5%

three different assumptions used earlier about the definition of the rich: the top 10 percent, 5 percent and 1 percent respectively.[6] Four conclusions are immediately apparent:

- Inequality between the rich and the non-rich groups is clearly larger than the inequality among the rich, except for the narrowest definition of the rich.
- The magnitude of wealth inequality in Sweden within the non-rich group comes as little surprise: this is to be expected from the con-

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

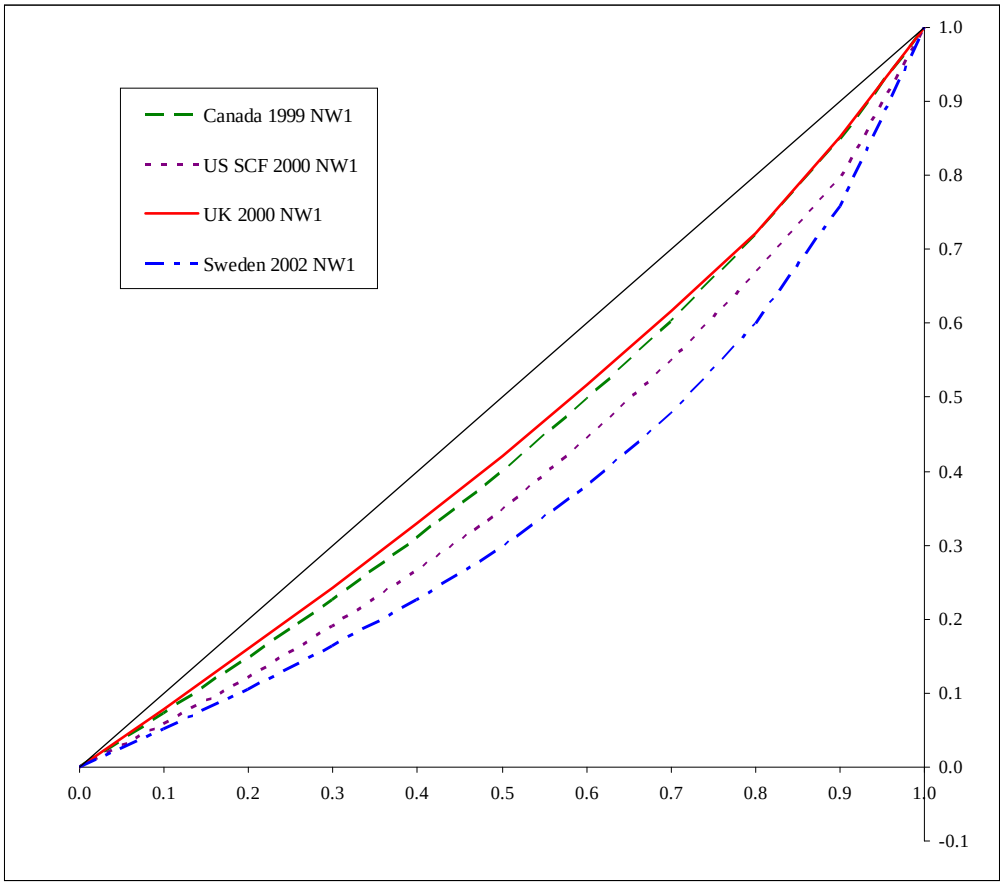

Figure 4: Lorenz Curves for Net Worth – Top 1%

siderable amount of negative net worth that is evident in Figure 1.[7]

- By contrast the very high inequality in Sweden within each of the rich groups is rather remarkable. If we focus just on the situation within the top 1 percent group then Sweden is unambiguously more unequal then the US (although the top 1 percent in the US has a large share of total net worth than does the top 1 percent in Sweden).

[7]It is also unremarkable to find that, when the rich is taken to mean the top 10 percent, the Gini coefficient for the non-rich (the bottom 90 percent) is 1.045. If a large proportion of the population have substantial negative net worth it can fairly easily happen that the Gini exceeds 1.

	Gini overall	Share rich	Gini of rich	Gini of non-rich	Gini between groups
		Top 10 percent			
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
		Top 5 percent			
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
		Top 1 percent			
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 3: Net worth: Gini decomposition for the top 10 percent, top 5 percent, top 1 percent

- Once again, perhaps surprisingly, the UK unambiguously exhibits the least inequality of the four countries. This conclusion 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.

However, some may be sceptical about the last two conclusions. 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 3 and Figures 1 to 4: is the overall wealth-inequality picture

being distorted by errors in computing inequality among the wealthy? This will be addressed below in Section 5.

4.2 Alternative wealth concepts

As we have noted the high proportion of households with negative net worth in Sweden, and the substantial amount of debt in the composition of net worth (as noted in the discussion of Table 1 above) may have had serious consequences for the overall picture of wealth-inequality comparisons. Perhaps the way debt has been computed and imputed makes Sweden appear to be much more unequal relative to other countries including the UK.[8] So the question naturally arises, does the inequality ranking change if we switch to other definitions of wealth?

Tables 4 and 5 provide a summary picture of the inequality of the wealth embodied in the households' principal residence, the inequality of other investment property, the inequality of financial assets and the inequality of total assets: Table 4 (shares) shows the proportions of various assets owned by the rich, defined as those holding the top 10 percent, 5 percent or 1 percent of net worth.. The Lorenz curves for each of the separate asset types are given in Figures 5 - 8. Using both the tables and Figures 5 - 8 we can immediately draw the following conclusions on the inequality of each wealth component:

- The *Principal Residence* results are of special interest in studying the UK because this asset type forms such a large proportion of households' portfolios. Of the four countries the UK has the highest between-group inequality of principal residence, but the lowest inequality within the rich group and the lowest inequality overall.

- Although the inequality of *Other Investment Property* is highest in the UK out of the four countries, this is not particularly significant in terms of the overall picture of wealth-inequality comparisons, for two reasons. First, inequality of this wealth component is very high in every country

[8]Note that the Swedish data are derived from information collected in connection with wealth-tax assessment. The value of the principal residence is a value reported for tax purposes and then inflated by a factor calculated by Statistics Sweden; the result is claimed to be close to market value. Debt on the tax statement was recorded as single amount as an offset against one's taxable wealth.

– the UK does not stand out from the rest. Second, in the UK at that time other investment property formed a very small part of the total asset portfolio

- The overall inequality of *Financial Assets* is a little higher in the UK than it is in Sweden but much less than in North America. But the financial-asset inequality between the rich and the non-rich is consistently lower in the UK than in other countries.

Finally let us examine the inequality of total assets, in other words wealth inequality ignoring debt. In almost every case – within-group, between group and overall – the ordering of countries by ascending order of total-asset inequality is UK, Canada, Sweden, US (the exceptions relate to within-group inequality for the top 5 percent and top 1 percent where Canada's inequality is low and Sweden's is very high). So, even if we ignore debt altogether, the UK still appears to be the least unequal of the four and Sweden's inequality is still high, although not perhaps the highest of all.[9]

[9]If we had combined the principal residence data with the other investment property data so as to get total non-financial assets, the Gini coefficents would have been: UK 0.584, Canada 0.616, US 0.702, Sweden 0.703; exactly the same ordering of countries by inequality as that found in Table 2 for net worth.

	Share of ...		
	Top 10%	**Top 5%**	**Top 1%**
	Principal Residence		
UK	0.339	0.201	0.051
Sweden	0.346	0.209	0.050
Canada	0.269	0.146	0.036
US	0.168	0.073	0.024
	Real Estate		
UK	0.653	0.540	0.302
Sweden	0.680	0.579	0.386
Canada	0.645	0.493	0.096
US	0.784	0.606	0.415
	Financial Assets		
UK	0.484	0.346	0.120
Sweden	0.534	0.401	0.213
Canada	0.743	0.609	0.357
US	0.977	0.863	0.506
	Total Assets		
UK	0.392	0.256	0.085
Sweden	0.440	0.303	0.133
Canada	0.419	0.291	0.114
US	0.580	0.468	0.271

Table 4: Shares of the rich in alternative wealth concepts

	Gini (all)	Gini within... Top 10%	Top 5%	Top 1%
	Principal Residence			
UK	0.559	0.274	0.284	0.308
Sweden	0.708	0.369	0.355	0.429
Canada	0.603	0.372	0.435	0.416
US	0.645	0.484	0.500	0.449
	Other Investment Property			
UK	0.966	0.860	0.793	0.582
Sweden	0.949	0.847	0.839	0.827
Canada	0.930	0.727	0.672	0.674
US	0.959	0.700	0.747	0.750
	Financial Assets			
UK	0.799	0.584	0.548	0.543
Sweden	0.778	0.587	0.593	0.509
Canada	0.860	0.616	0.553	0.195
US	0.899	0.510	0.417	0.294
	Total Assets			
UK	0.578	0.263	0.230	0.171
Sweden	0.666	0.332	0.336	0.358
Canada	0.626	0.293	0.262	0.132
US	0.748	0.345	0.324	0.195

Table 5: Gini coefficient for components of net worth and for total assets

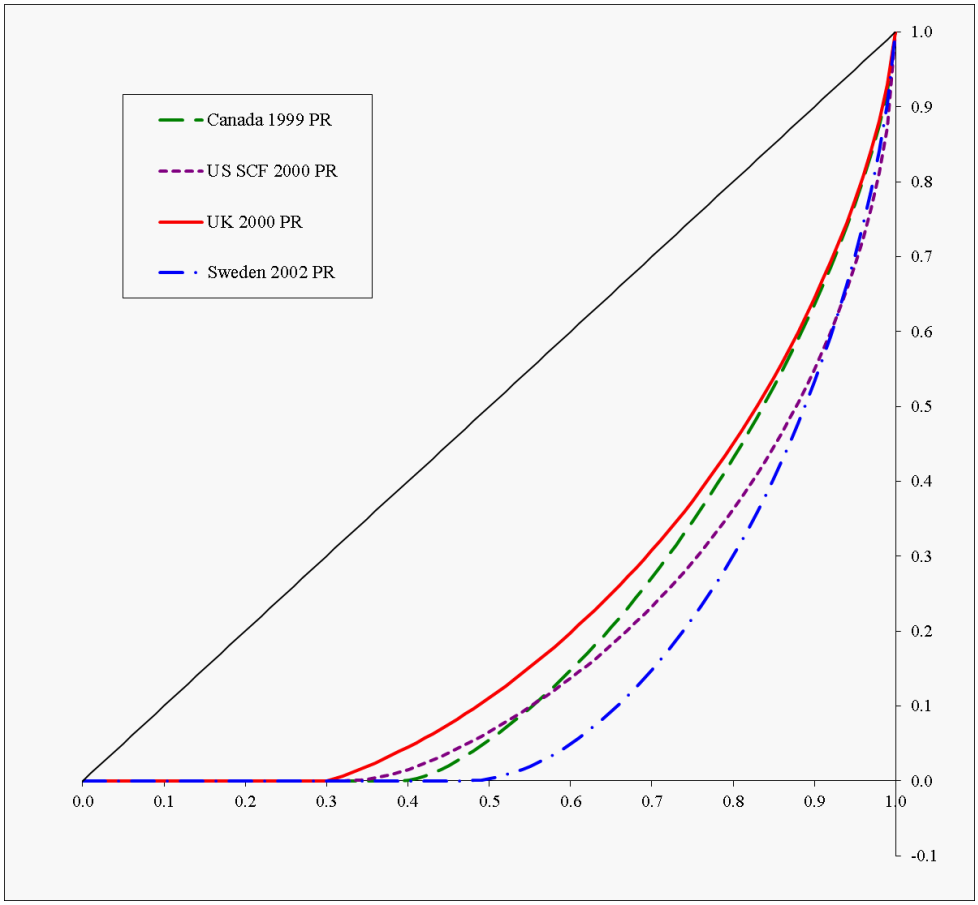

Figure 5: Lorenz Curves for Value of Principal Residence

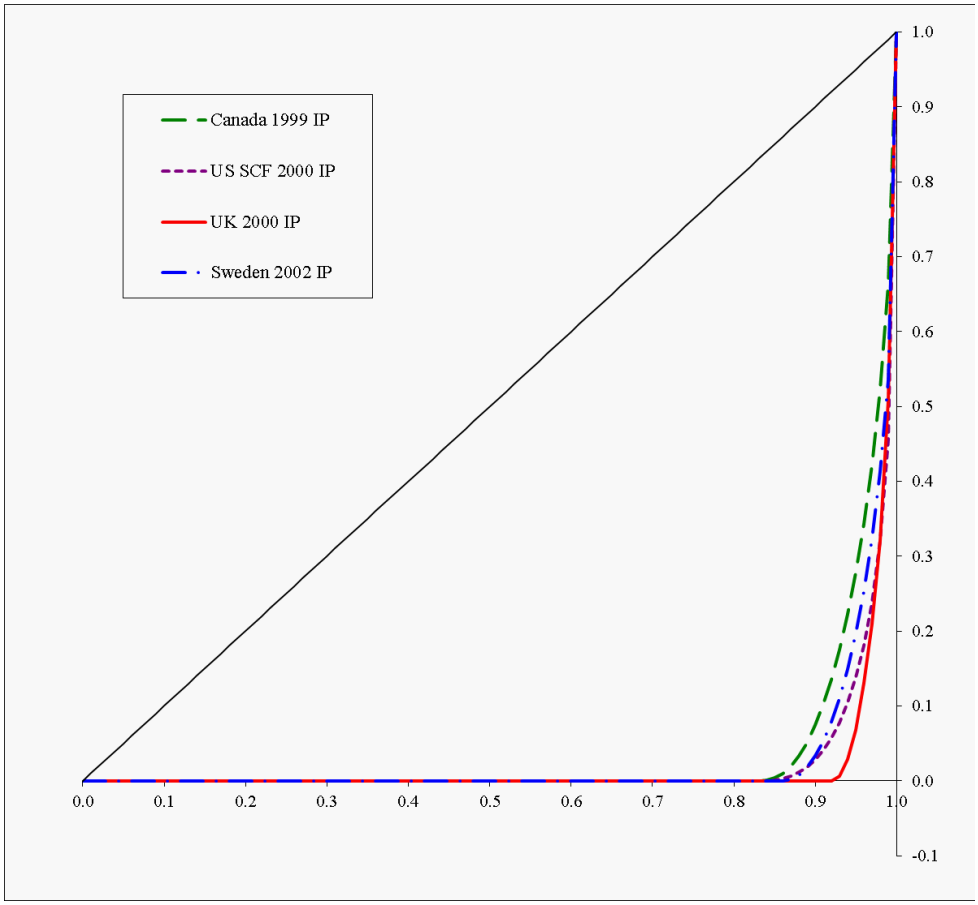

Figure 6: Lorenz Curves for Other Investment Property

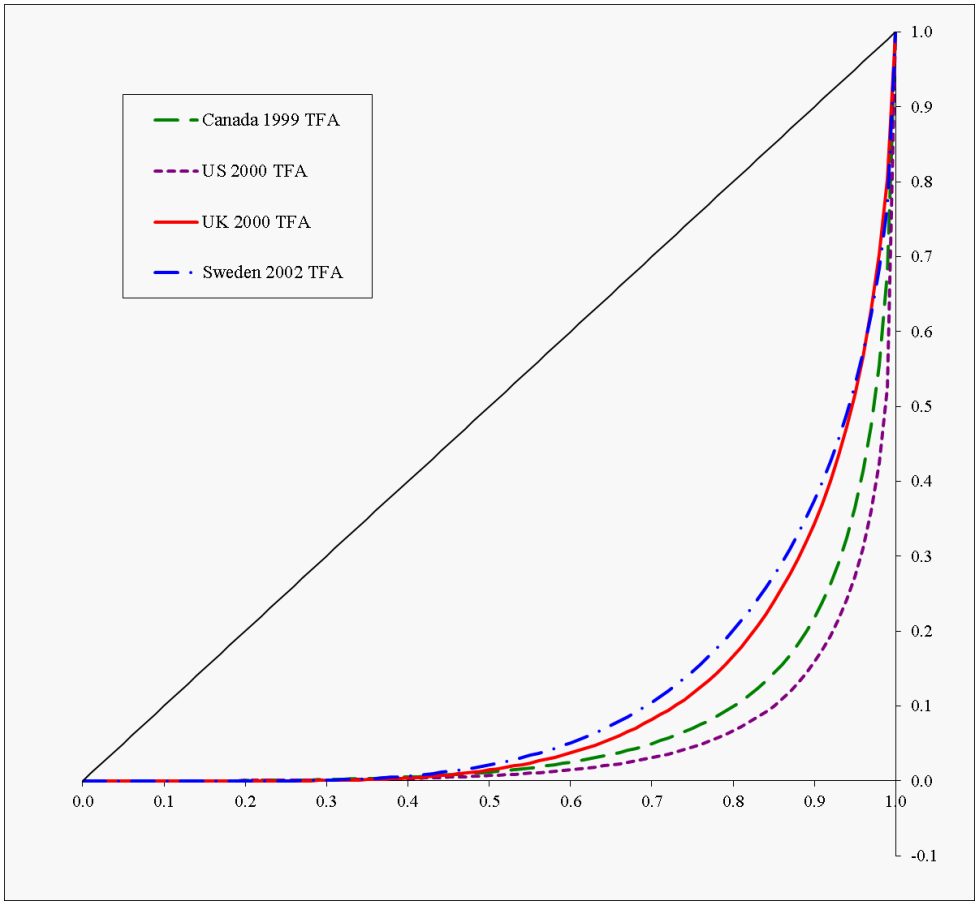

Figure 7: Lorenz Curves for Financial Assets

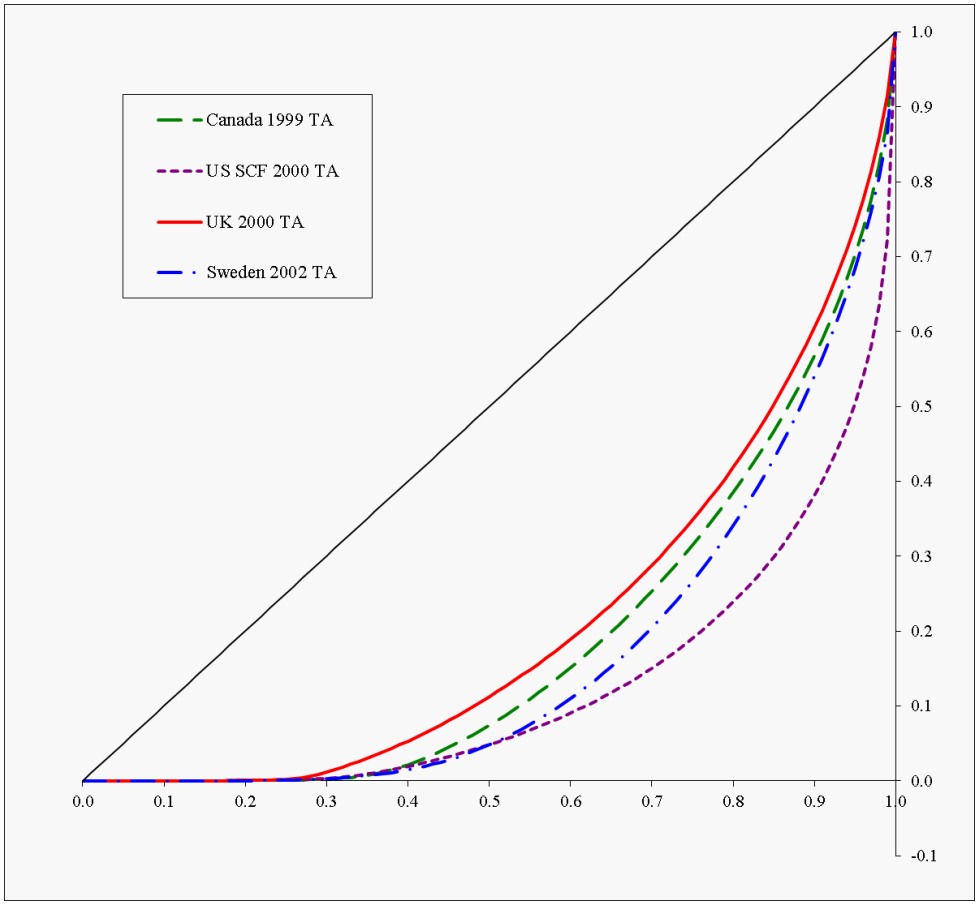

Figure 8: Lorenz Curves for Total Assets

5 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.[10] Here, in common with other studies of the upper tail of the wealth distribution, we use the Pareto distribution.[11] To introduce this widely accepted method of representing the wealth distribution among the rich, follow through the graphical interpretation in Figure 9. This is a standard "Pareto diagram" – the horizontal axis in Figure 9 is net worth x plotted on a logarithmic scale and the vertical axis is $P(x)$,[12] the proportion of people with net worth greater than or equal to x (also on a log scale).

The Pareto hypothesis is that the underlying relationship should be exactly a straight line, the slope of which is known as the Pareto coefficient, α.[13] For an intuitive interpretation of α, 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* relative to the base b) is given by $\frac{\alpha}{\alpha-1}b$. So the "average/base" ratio is just $\frac{\alpha}{\alpha-1}$ which must be a constant, independent of the base level b, for a true Pareto distribution. For any distribution this average/base idea gives a simple concept of inequality and we can see immediately that the higher is α, the lower is the average/base ratio: high-inequality Pareto distributions have low values of α. We will discuss the inequality associated with the Pareto model further in section 6.

If we plot the LWS net-worth data so as to take a preliminary look at

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

[11] 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).

[12] Writing this in terms of the more familiar distribution function F we have $P(x) = 1 - F(x)$. Figure 9 is in fact four diagrams (for the separate countries) in one. Because the diagram plots the log of wealth on the horizontal axis and we do not need to compare wealth levels across countries, wealth has been left in the national currency. So the common horizontal axis plots log wealth in each of the four national currencies: changing the rate of exchange for a country's currency would just displace the country's plots and regression lines horizontally by a constant amount and would not affect any of our conclusions.

[13] Formally we have $\log P(x) = k - \alpha \log x$ where the slope α is related to the inequality of the distribution and the intercept k is a location parameter.

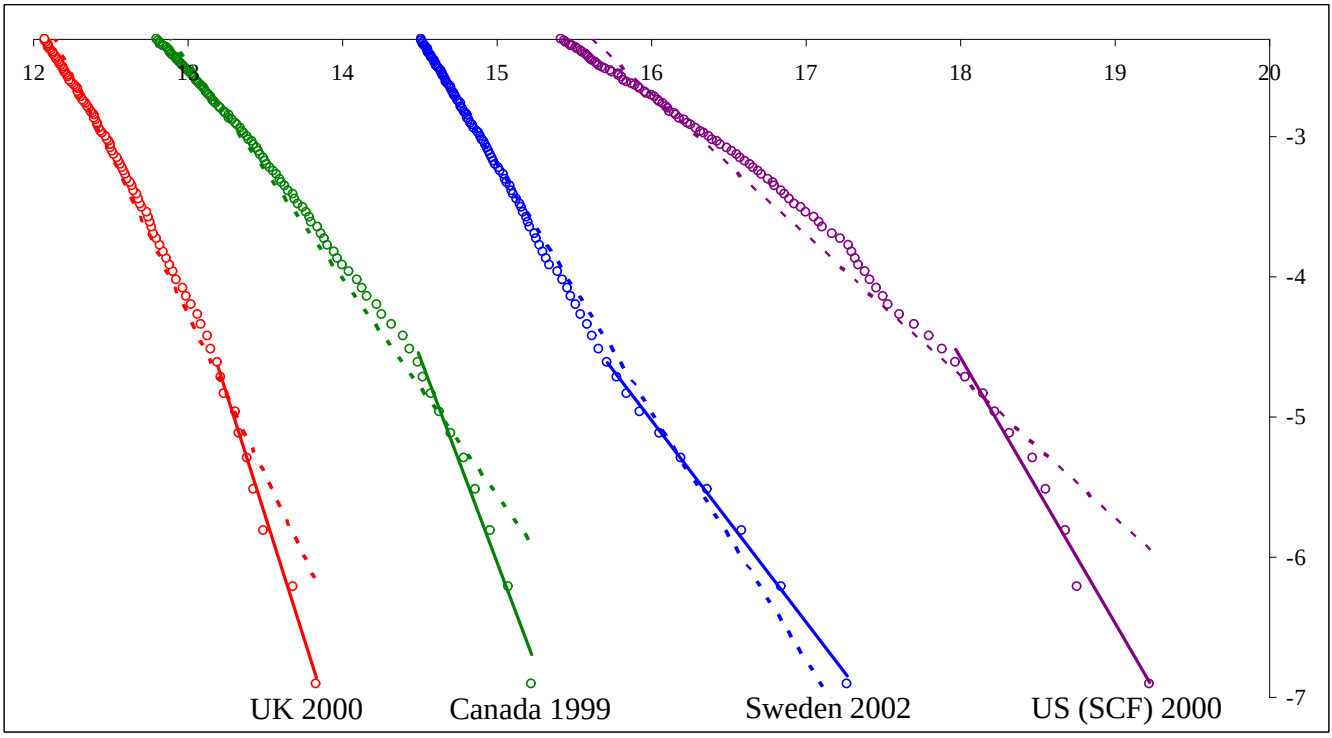

Figure 9: The Pareto model: a first look

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 9 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.[14] Nevertheless, the Pareto model may be a sufficiently good approximation to the model generating the data to provide reasonable estimates of inequality.

5.1 The influence of outliers

Because of the apparent sensitivity of results to precise assumptions made about the upper tail of the wealth distribution it makes sense to investigate

[14] Details are given in Table 6 below.

	OLS estimation			Robust estimation		
	Top 10%	**Top 5%**	**Top 1%**	**Top 10%**	**Top 5%**	**Top 1%**
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 6: Estimates of Pareto's alpha for different definitions of the Rich

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 9, and (2) robust estimation, a procedure that downweights outliers. The estimates are summarised in Table 6.

In each case the robust method "pulls" the regression line away from the data points at the bottom right-hand end of each country's scatter of observations in Figure 9: this is what is meant by "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 than if we use standard OLS; for Sweden the opposite holds. More on this in Section 6. Notice too that the estimate of α is sometimes below 1. In this case the mean of the Pareto distribution is undefined and so the Pareto model is not going to be suitable for the semi-parametric patch job. For the cases where we could use the Pareto model, Table 7 gives the intuitive "average/base" inequality estimates for both estimation methods. The UK again clearly emerges as the least unequal of the four countries.

6 Wealth inequality comparisons – a third 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

	OLS estimation		Robust estimation	
	Top 10%	**Top 5%**	**Top 10%**	**Top 5%**
UK	1.916	1.812	2.062	1.916
Sweden	2.577	3.439	2.174	2.169
Canada	2.866	2.695	4.000	3.636
US	39.462	8.042		27.027

Table 7: Estimates of average/base inequality for different definitions of the Rich

the breakdown of wealth inequality in Table 3. 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).

6.1 The Lorenz curve

First let us redraw the Lorenz curves using the model to patch the upper tail. Fortunately the Pareto model has associated with it a Lorenz curve that has a very simple shape. To explain this let us again use $P(x)$ to mean the proportion of the population with wealth greater than or equal to x; similarly let us use $S(x)$ to mean the proportion of total net worth held by those owning x or more. We know that for a Pareto distribution the relationship between $P(x)$ and x is double log-linear with slope α (the equation is in footnote 13); the relationship between $S(x)$ and x is also double log-linear, but with slope $\alpha - 1$. From this we may deduce that the wealth share $S(x)$ is found by raising the population share $P(x)$ to the power $1 - 1/\alpha$. The modelled part of the Lorenz curve is found by graphing $S(x)$ against $P(x)$ over the range of x for which we are assuming the model to be valid.[15] Figures 10-13 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

[15] Formally we have $\log S(x) = \frac{k}{\alpha}[\alpha - 1] - [\alpha - 1]\log x$ and the Lorenz curve is therefore given by $\log S(x) = \frac{\alpha - 1}{\alpha}\log P(x)$.

focusing on the top 10 percent snapshot of the whole distribution rather than looking at the top 10 percent as a sub-population as we did in Figure 2.

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.

6.2 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 5.1 – 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.

Within-group inequality. 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 parameter α is $\frac{1}{2\alpha-1}$. So if we use a Pareto distribution to model the wealth distribution among the rich – whether they are defined as the top 10 percent, 5 percent or 1 percent of wealth holders – within-group inequality is just $\frac{1}{2\alpha-1}$, where the value of α is that estimated for the particular definition of the rich. Of course inequality within the non-rich remains just as before.

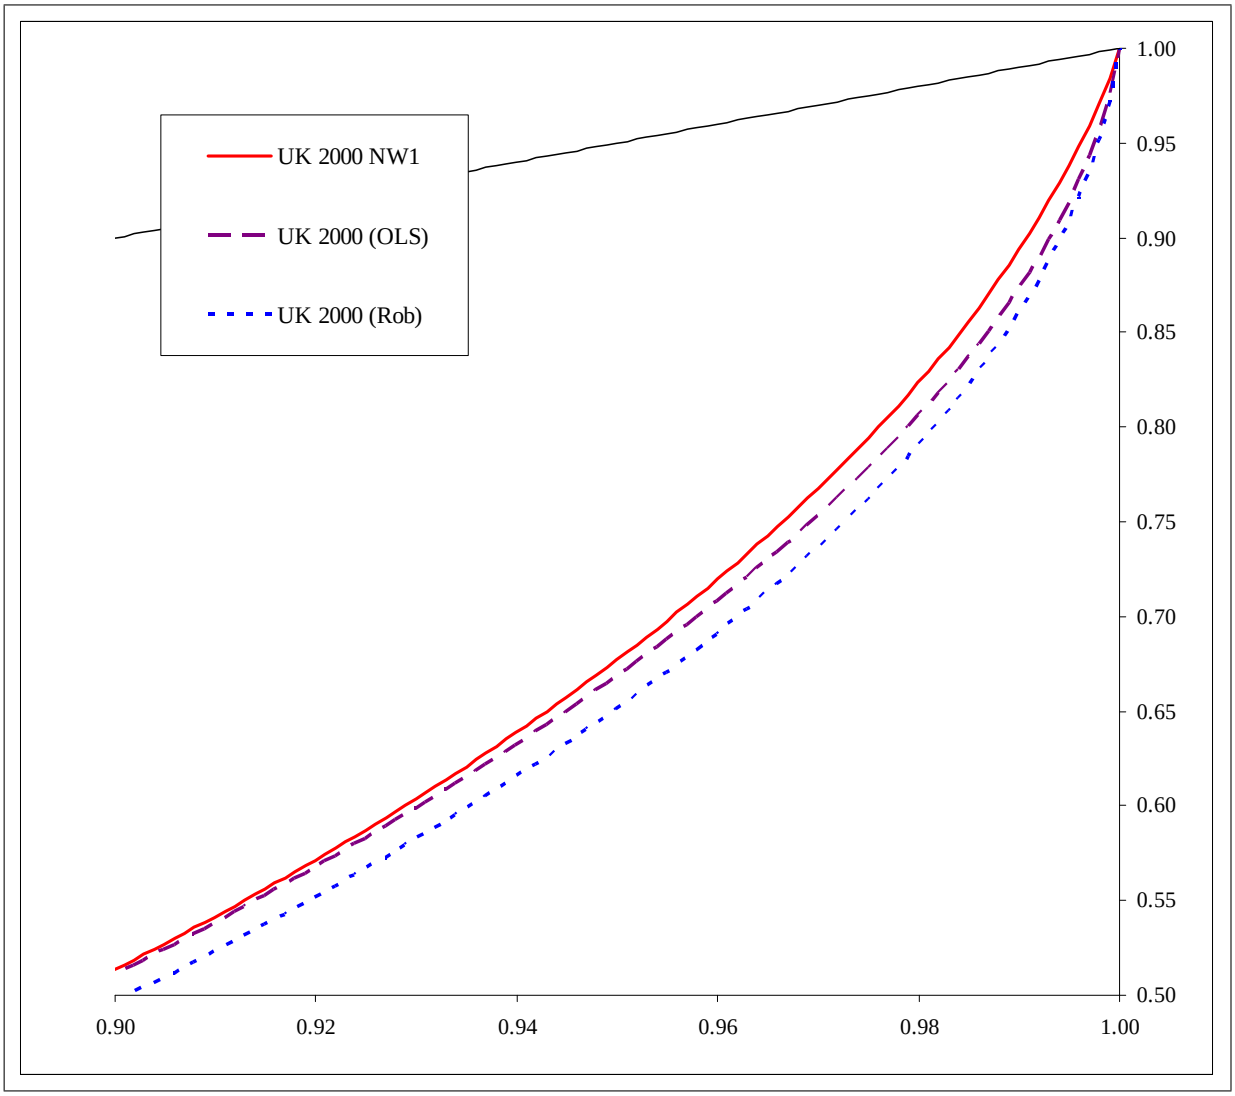

Figure 10: Raw Data and Pareto Model: UK 2000

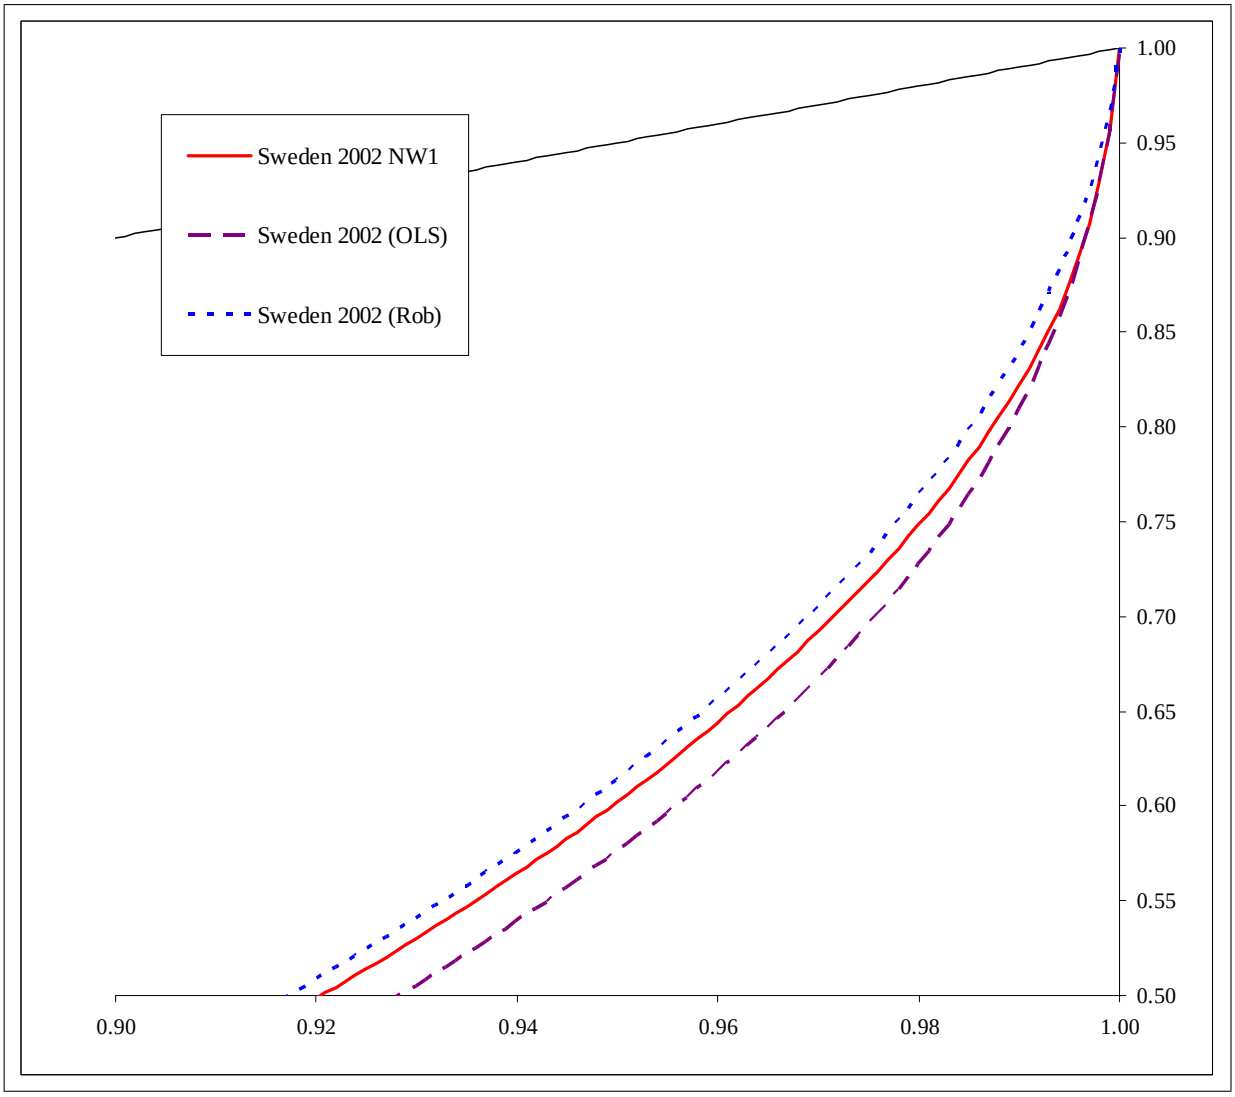

Figure 11: Raw Data and Pareto Model: Sweden 2002

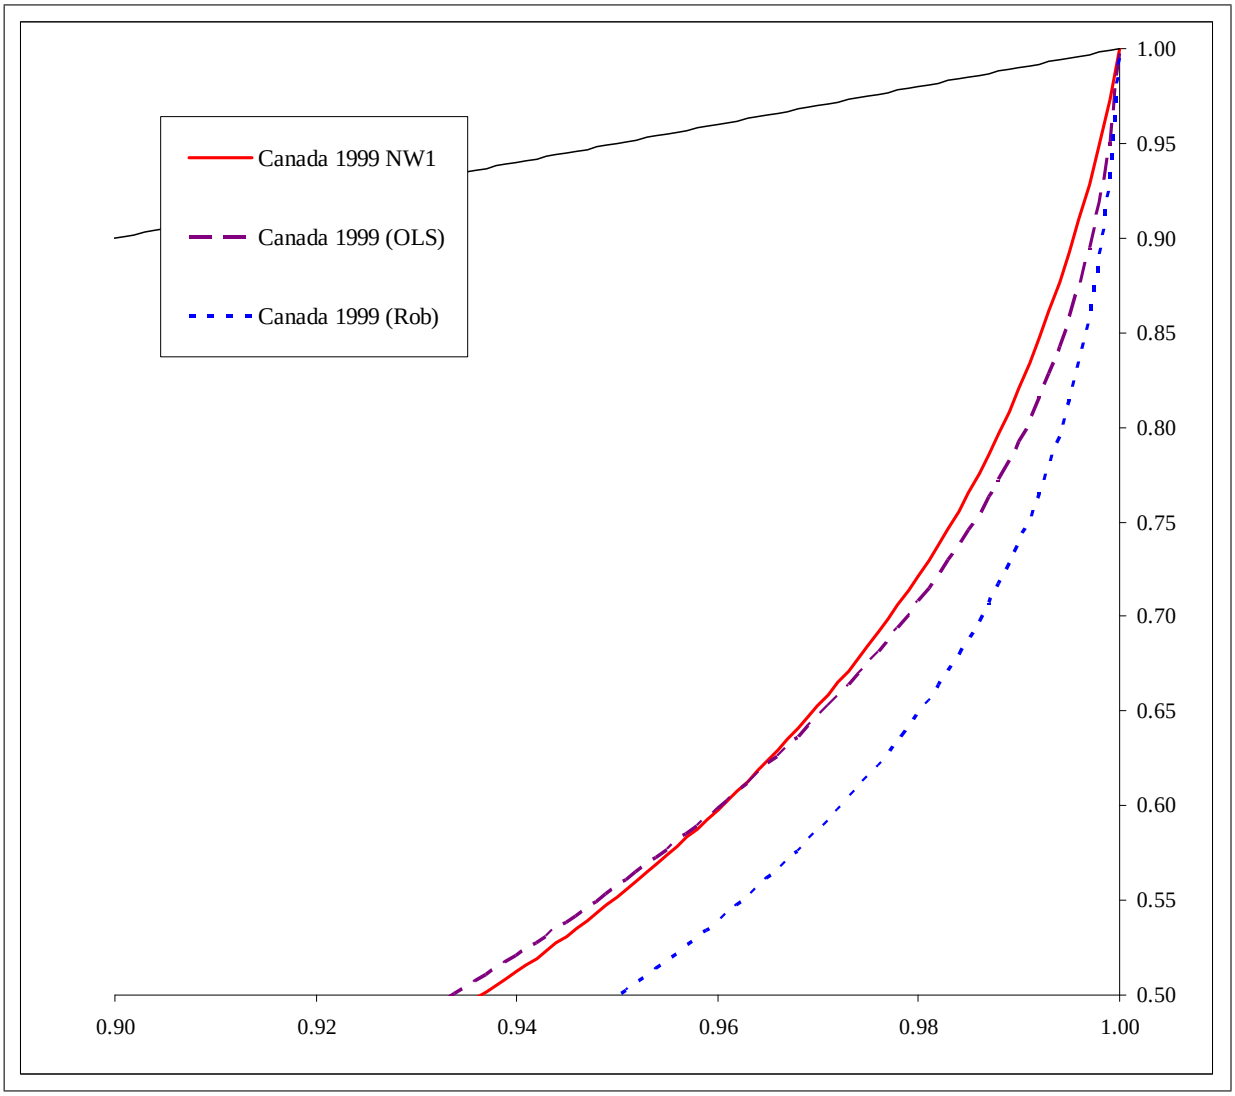

Figure 12: Raw Data and Pareto Model: Canada 1999

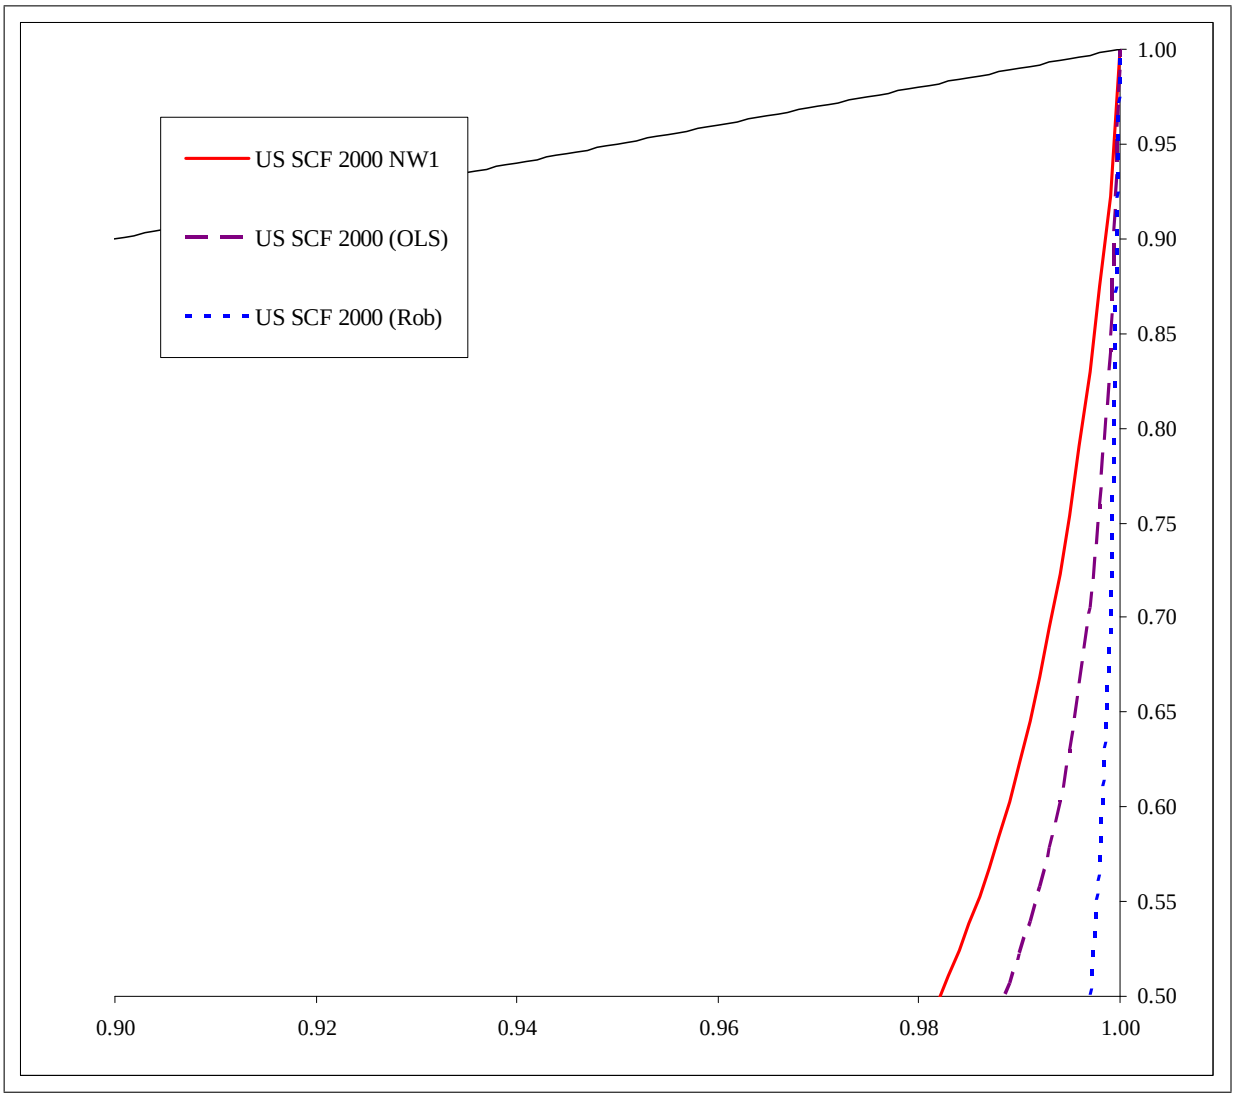

Figure 13: Raw Data and Pareto Model: US (SCF) 2000

Between-group inequality. In order to recalculate the between-group Gini we need to recompute the following:

- the revised mean of the rich group,
- the revised overall mean,
- the revised wealth share of the rich group.

The first of these is immediate once we know the quantile that forms the boundary between the rich and non-rich groups. For example if we define the rich to be the top 10 percent then we just use the formula for the Pareto distribution to calculate the top decile. From the value of the top decile we obtain the first bullet-point item; the second and third bullet-point items then follow.[16]

The between-group Gini is then, as before, just the difference between the wealth share and the population share of the rich group.

Overall inequality. Now let us combine the within-group and between-group effects of the modelling to obtain revised versions of total Gini inequality and its breakdown. We need do no more than re-use the decomposition formula for the Gini coefficient that was employed in Section 4.1 (see footnote 4 for the formal details), only this time using the recomputed wealth shares, between-group Gini and within-rich Gini, along with the unchanged Gini for inequality within the non-rich.

The results of the Pareto modelling on the structure of wealth inequality are shown in Tables 8 and 9. 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. If the value is less than one then the mean of the Pareto distribution is undefined and, of course, the average/base ratio, the adjusted share of the rich have no meaning either; if the value is less than 0.5, then the adjusted rich-group Gini is undefined. In such cases the decomposition of inequality has been omitted from the tables. It is clear from Table 6 that for both OLS and robust estimation such

[16] In general, let x_{R} be the p_{R}-quantile where p_{R} is the proportion of the population that is rich. Using the log-linear relationship given in footnote 13 we have $\log x_{\mathrm{R}} = [k - \log p_{\mathrm{R}}]/\alpha$. Then the revised mean of the rich group is $\frac{\alpha}{\alpha-1}x_{\mathrm{R}}$ and the rest follows from the relationships in footnote 4.

	Gini overall	**Share rich**	**Gini of rich**	**Gini of non-rich**	**Gini between groups**
			Top 10 percent		
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
			Top 5 percent		
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 8: Gini decomposition – adjusted with OLS Pareto model rich group

	Gini overall	**Share rich**	**Gini of rich**	**Gini of non-rich**	**Gini between groups**
			Top 10 percent		
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
			Top 5 percent		
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 9: Gini decomposition – adjusted with robust Pareto model

problems arise everywhere when applied to the top 1 percent group and that the problem arises for robust estimates in the case of the US.

Comparing the remainder of the results in Tables 8 and 9 with Table 3 it is clear that modelling the rich (either the top 10 percent or the top 5 percent) 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 3, that wealth in the UK is unambiguously less unequally distributed than in 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).

7 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.

It might have been surmised that this result could have arisen from some anomaly in the way that the components of net worth are recorded (particularly debt) or 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 alternative definitions of wealth. It is also robust when one models the upper tail in order to allow for incomplete data and for outliers that may not represent the "true" wealth distribution. Furthermore it is also true for the rich/non-rich 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.

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.