

Thinking about Inequality

Personal Judgment and
Income Distributions

Y. AMIEL AND F. A. COWELL



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge CB2 1RP, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge CB2 2RU, UK <http://www.cup.cam.ac.uk>
40 West 20th Street, New York, NY 10011-4211, USA <http://www.cup.org>
10 Stamford Road, Oakleigh, Melbourne 3166, Australia

© Y. Amiel and F. A. Cowell 1999

This book is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 1999

Printed in the United Kingdom at the University Press, Cambridge

Typeset in Monotype Times 10/13pt [SE]

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Amiel, Y. (Yoram)

Thinking about inequality / Y. Amiel and F. A. Cowell.

p. cm.

Includes bibliographical references (p.).

ISBN 0 521 46131 6 (hb) – ISBN 0 521 46696 2 (pb)

1. Income distribution. 2. Equality. 3. Poverty. I. Amiel, Yoram. II. Title.

HB523.C694 1999

339.2–dc21 99-13081 CIP

ISBN 0 521 46131 6 hardback

ISBN 0 521 46696 2 paperback

Contents

	<i>List of figures</i>	<i>page</i> viii
	<i>List of tables</i>	x
	<i>Preface</i>	xiii
1	Introduction	1
	1.1 A look at inequality analysis	1
	1.2 A second look	3
	1.3 A guide to the book	6
2	What is inequality? The economists' view	8
	2.1 The axiomatic approach	8
	2.2 Inequality rankings and orderings	9
	2.3 The transfer principle	11
	2.4 Income and population	12
	2.5 Decomposability	15
	2.6 Summary	17
3	An investigative strategy	18
	3.1 What are we investigating?	18
	3.2 Experiments	21
	3.3 Questions	23
	3.4 A new approach	24
	3.5 Implementing the approach	27
	3.6 Summary	30
4	What is inequality? The students' view	31
	4.1 Drawing an inequality map	31
	4.2 An introduction to the questionnaires	32

vi	Contents	
	4.3 Inequality and changes in income and population	35
	4.4 Transfers and the structure of inequality comparisons	38
	4.5 Do the answers make sense?	42
	4.6 More on the transfer principle	45
5	Income and welfare	49
	5.1 What is welfare?	49
	5.2 Social welfare	50
	5.3 Empirical results	57
	5.4 Summary: welfare judgments and inequality comparisons	66
6	Income change	69
	6.1 Introduction: comparing cakes	69
	6.2 Uniform enrichment	71
	6.3 The dependence hypothesis	75
	6.4 Unbalanced enrichment	78
	6.5 Policy appraisal	86
7	Poverty	89
	7.1 Introduction	89
	7.2 What does 'poverty' mean?	89
	7.3 The poverty questionnaires	94
	7.4 Income distributions and poverty	96
	7.5 Conclusions: the approach to poverty comparisons	111
8	A cross-cultural perspective	114
	8.1 Introduction	114
	8.2 A statistical approach	116
	8.3 Principles of distributional judgments	117
	8.4 Direct and indirect approaches to inequality	122
	8.5 Does economics matter?	123
	8.6 An appraisal	126
9	Thinking again about inequality	127
	9.1 Second thoughts about second thoughts	127
	9.2 Applying inequality judgments	128
	9.3 Where next?	130
	9.4 A final word	133
	Appendix A Inequality analysis: a summary of concepts and results	136
	A.1 The axiomatic approach	136

	Contents	vii
A.2 Inequality and welfare rankings		139
A.3 Poverty comparisons		140
Appendix B The questionnaires		143
<i>References</i>		173
<i>Index</i>		178

Figures

1.1	A simple distributional experiment	<i>page 4</i>
1.2	A simple distributional experiment: second view	5
1.3	A simple distributional experiment: third view	5
2.1	An inequality ranking	9
2.2	Inequality in a two-person world	10
2.3	Inequality comparisons in a three-person world	11
2.4	Scale independence	13
2.5	Translation independence	14
2.6	The population principle	14
2.7	Population replication – has inequality fallen?	15
2.8	Decomposability	16
3.1	Two extreme approaches to identity in the distribution problem	25
3.2	The identity problem with some information	26
4.1	The framework for a three-person inequality map	32
4.2	The framework for the two-person projection of the inequality map	33
4.3	Verbal questions on scale and translation independence	34
4.4	Numerical problems on changes in income and population	35
4.5	The question on the population principle	35
4.6	Scale-independent and translation-independent iso-inequality lines	36
4.7	Deviations from scale transformation	37
4.8	Numerical problems on the transfer principle and decomposability	39
4.9	Verbal questions on the transfer principle and decomposability	39
4.10	The transfer principle in the two-person projection	40
4.11	Two-dimensional projection of the inequality map	41
4.12	The Lorenz curve	43
4.13	Lorenz ranking	43
5.1	The anonymity principle	51

5.2	Anonymity and history	52
5.3	Monotonicity	53
5.4	Dominance and monotonicity	54
5.5	The setting for the social welfare numerical problems	57
5.6	Contours of a standard social welfare function	58
5.7	A social welfare function which satisfies the transfer principle, but not decomposability	63
5.8	Contours of a non-monotonic welfare function	65
6.1	Two pie distributions, before and after income growth	70
6.2	Additions to Irene's and Janet's incomes	72
6.3	Enlargement of figure 6.2	72
6.4	Scale independence	73
6.5	Translation independence	74
6.6	Intermediate-type independence	75
6.7	Transformation direction dependent on income (1)	75
6.8	Transformation direction dependent on income (2)	76
6.9	Introduction to questionnaire A3	77
6.10	Unbalanced enrichment	79
6.11	Inequality and growth: first view	80
6.12	Inequality and growth: second view	81
6.13	Extract from questionnaire A1	81
6.14	Extract from questionnaire A2	83
6.15	Mean income and inequality as incomes grow from \$5 to \$10	85
6.16	Mean income and inequality as incomes grow from \$1 to \$10	85
7.1	A fundamental partition of the population	91
7.2	The poverty line, incomes and poverty gaps	92
7.3	Counting the poor (1): all the poor are equal	93
7.4	Counting the poor (2): poverty is proportional to poverty gap	93
7.5	Counting the poor (3): sensitivity to inequality amongst the poor	94
7.6	A poverty interval	95
7.7	Numerical problems in the first poverty questionnaire	96
7.8	Weak monotonicity	97
7.9	The monotonicity question	98
7.10	Income transfers and the poverty count	100
7.11	Transfer principle (weak version)	100
7.12	Positions for an additional population member	104
7.13	Numerical problems in the second poverty questionnaire	106
7.14	Extract from the verbal questions in the second poverty questionnaire	107

Tables

4.1	Inequality and proportionate and absolute income differences	<i>page</i> 38
4.2	The effect on inequality of cloning the distributions	38
4.3	The transfer principle	41
4.4	Decomposability	42
4.5	What happens to inequality if you add or subtract a fixed sum?	44
4.6	Agreement with the transfer principles for different types of transfer: numerical responses	46
4.7	Verbal agreement with the transfer principle	46
4.8	Agreement with basic axioms: summary	48
5.1	Agreement with transitivity of inequality and social welfare orderings	59
5.2	The anonymity principle	59
5.3	The transfer principle again: numerical responses	60
5.4	Transfer principle: verbal responses on social welfare questionnaire	61
5.5	The effect on social welfare of cloning the distribution	62
5.6	Decomposability of social welfare?	63
5.7	Agreement with monotonicity: numerical questions	64
5.8	Agreement with monotonicity: verbal questions	65
5.9	Agreement with basic axioms on social welfare: summary	66
5.10	Support for basic principles of inequality comparisons	67
6.1	What income change will leave inequality unchanged?	77
6.2	What happens to inequality when you increase people's incomes?	78
6.3	Perceived inequality change in the growth process of figure 6.13	82
6.4	Comparing extremes in the growth process	84
7.1	What happens to poverty if a poor person gets \$1 more income?	99
7.2	What happens to poverty if \$1 is taken from fairly poor Irene and given to very poor Janet?	101

7.3	If we permute the incomes, does poverty stay the same?	101
7.4	What happens to poverty if we clone the economy?	102
7.5	Population decomposability	103
7.6	What happens to poverty if the 'rich' get richer?	103
7.7	The effect of introducing one new person	105
7.8	What happens to poverty if there is one more non-poor person?	105
7.9	What happens to poverty if there is one more poor person?	105
7.10	What happens to poverty when the basic needs income level increases?	108
7.11	What happens to poverty when you double incomes and basic needs?	109
7.12	Shifting incomes and the poverty line by a fixed sum	110
7.13	What is poverty?	111
7.14	Support for standard axioms in inequality, social welfare and poverty analysis	112
8.1	Breakdown of views on the transfer principle: direct approach	118
8.2	Do X and Y have the same pattern of responses on the transfer principle? Direct approach	119
8.3	Agreement with monotonicity: does B exhibit higher social welfare than A?	120
8.4	Breakdown of verbal responses on the monotonicity principle	121
8.5	Breakdown of views on the poverty line: verbal question	122
8.6	Do X and Y have the same pattern of verbal responses on the poverty line question?	122
8.7	Breakdown of views on the transfer principle: indirect approach	124
8.8	Do X and Y have the same pattern of responses on the transfer principle? Indirect approach	124
8.9	Do X and Y have the same pattern of responses on monotonicity?	125
9.1	Standard axioms in three related fields	129
9.2	Standard axioms in the analysis of income and probability distributions	131
9.3	What happens to risk when you increase people's incomes?	132
9.4	What income change will leave risk unchanged?	132
B.1	Breakdown of the combined sample	144

1 Introduction

1.1 A look at inequality analysis

Thinking about inequality is not always a fashionable topic amongst economists. But thinking about inequality actually goes on all the time. Perceptions of inequality affect economic choices and political decisions. A sensitivity to inequality coupled with compassion for the poor motivates charitable giving by individuals and states. Notions about inequality appear to inform popular views about the appropriateness or otherwise of pay awards. And any parent with two or more children needs no formal analysis to be persuaded of the importance of distributive justice. Fashionable or not, thinking about inequality plays a part in the judgments and actions of politicians, planners and ordinary people.

Of course the study of economic inequality has not just been a matter of fashion. It has been an integral part of the general historical development of political economy and economics, and the approach to the topic has changed with the passage of time. While this is not the place for an extensive treatise on the history of economic thought about inequality, a brief sketch to introduce conventional wisdom on the subject may help to put into context what we want to tackle in this book.

This century has witnessed a shift in emphasis in thinking about inequality. It used to be commonplace to set the analysis of economic injustice within a particular social or institutional framework – such as Ricardo’s or Marx’s class-based theories of political economy. Alternatively, issues of inequality used to be cast in terms of specific models of income distribution – such as Vilfredo Pareto’s famous laws of distribution. However, in more recent times, there has been a move away from these narrowly focused perceptions of the problem to an approach founded upon general principles. What principles?

A cursory review of recent literature suggests that the principles encompass a wide range of theoretical and applied economics. But all the same it is possible

to simplify them down to a relatively few essentials. In practice we may usefully distinguish four major building blocks that are required in the analysis of income distributions:

- *The definition of income.* We need to specify carefully, or to be told clearly, what the thing called ‘income’ is.
- *The income recipient.* We also need to be clear about the nature of the entities – persons, families, households or whatever – that receive those incomes.
- *The reference group.* We should explicitly define the ‘universe’: the collection of persons or groups within which inequality comparisons are to be made.
- *The calibration system.* The ‘inequality thermometer’ – the inequality measurement tool – has to be precisely specified.

These four main components of inequality analysis get unequal treatment in the literature. As a sweeping generalisation we may state that items one, two and four in this list get a lot of attention; item three does not.

For example, theoretical economists focus principally on the fourth item in the list: the specification of a system of calibration. This forms a natural extension to a substantial literature on social choice and welfare economics. The way the analytical problems are formulated has close connections with other related issues such as the assessment of risk, the meaning of individual utility and the construction of index numbers of prices and income. In fact, inequality presents a classic theoretical measurement problem, and is typically treated in a classic fashion by setting out a system of axioms that appear to be reasonable and by formulating key propositions that follow from the axiomatic base.

On the other hand, applied economists and statisticians usually pay close attention to the first two issues: it is widely recognised that practical matters in defining income (or wealth, consumption expenditure, or whatever) or the family unit that is the income receiver are essential to understanding levels and trends of inequality within most economies.

But as far as the third point on the list is concerned – the appropriate reference group – one is immediately struck by the lack of references in the mainstream economics literature. Why this apparent neglect of one of the main components of income distribution analysis? Perhaps the answer is that to many researchers the issue seems obvious or self-defining. For example, in an empirical study, the sample is what it is. The population which the sample represents – so it might be argued – does not really need more than the most cursory discussion. Yet in principle the ‘universe set’ on which income distributions are to be defined and inequality to be assessed is a matter of theoretical as well as practical debate. For example, in the world of Plato or Aristotle the issue of distributional justice was applied only to free men since, in a social system that tolerated slavery, economic injustice for slaves was not a particularly relevant concept (and, of course, women did not get a look in). Similar difficulties have been raised in connection with modern theories of justice: who is to be counted within the

ambit of such theories, or who is to be party to the social contract? The voting public? All adults? The whole population? If the jurisdiction of nation-states can abruptly change, even this last broad definition may be imprecise. Matters become yet more complicated if we try to take account of all the citizens of the world or persons yet unborn. The question has also been raised as to whether the principles that are applied to people should also be applied to cats, dogs and other animals.¹

Even on the empirical level the issue of the reference group can have a dramatic impact on the picture that emerges about the pattern of world inequality. As a simple instance of this consider the study of international income comparisons by Summers and Heston (1988, 1991). Their hundred-plus countries are divided into six broad groups (Africa, North and Central America, South America, Asia, Europe and Oceania) so that it is possible to obtain a broad-brush picture of world income inequality in 1985 and 1988. But at first glance this broad-brush picture looks rather extraordinary: we find that in 1985 per capita income in Oceania was remarkably low – below Asia and South America so that it ranked fifth out of the six world regions (were the New Zealanders and Australians really so hard-pressed?); but in 1988 per capita income in Oceania had seemingly jumped so that it ranked third out of six (after Europe, but above South America). The answer to this conundrum is not hard to find: in the 1985 data compilation the relatively poor Indonesia – with its 160 million inhabitants – was classified as being part of Oceania; in 1988 it was lumped in with Asia. So, by re-specifying the groups only very slightly – in effect just relabelling one country – a substantially different story emerges of income inequality among different regions of the world. Clearly too, whether one counts Indonesia as an Asian country or part of Oceania is going to have a dramatic impact on the perceived inequality within Oceania.

This brief mention of theoretical and practical difficulties is not intended to imply that clear comparisons of inequality are usually impossible or meaningless. But it serves to highlight the importance of what might appear to be mere background features of the problem in making sensible inequality comparisons.

1.2 A second look

In our view there are deeper problems associated with the issue of the reference group. In fact it is arguable that the issue lies at the root of some of the more intractable problems in the assessment of income distribution. One of these problems – which we shall be taking up later in the book – is the relationship between the analysis of economic inequality and the analysis of poverty. Over recent years each of these two related topics has been extensively developed in terms of a mathematical approach founded upon a set of formal assumptions or axioms. But they have been developed separately, each using a distinct set of axioms as an intellectual basis. The intellectual divorce between the two branches of the

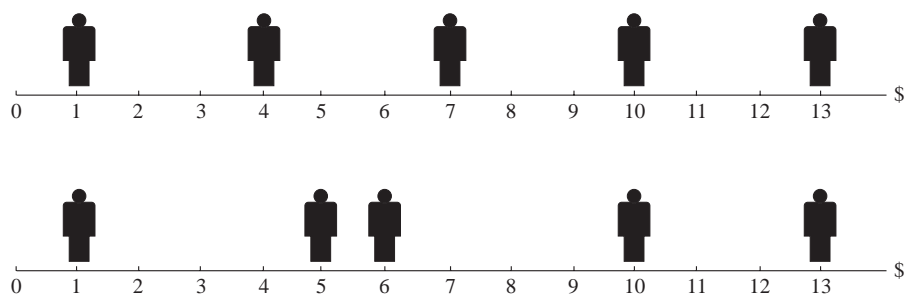


Figure 1.1. A simple distributional experiment.

subject can to some extent be explained in terms of different approaches to the idea of a reference group as we shall see further in chapter 7.

The way that reference groups are perceived also has a bearing upon some basic propositions in inequality analysis. In effect, what people mean by inequality can be crucially dependent on their perception of the relevant reference groups and in the ways that these groups are interlinked. The problem of the reference group and the way in which it relates to people's thinking about inequality is actually a convenient introduction to the case for a second look at the basics of inequality and income distribution analysis.

As an example of what is involved here, try a simple experiment. Figure 1.1 shows two possible income distributions in a very elementary economy. Each distribution contains five persons who have been arranged on an income scale in positions corresponding to their incomes, and the two distributions have the same total income (\$35). The units of income are irrelevant in the experiment (the '\$' sign has an unspecified value) but let us suppose that income tells us all that we might need to know about the 'well-offness', economic status or whatever of the persons; and we might as well assume that the five anonymous persons are as identical as the caricature suggests them to be. The experiment is simply this: write down which of the two distributions appears to you at first sight to be the more unequal and, if possible, give reasons for your answer.

Now, noting that the difference between the two distributions directly affects only two of the persons in the experiment, consider the slight modification of the diagram that is presented in figure 1.2. Here we have explicitly divided the population of five into two component groups, left and right, as indicated by the shading, but the distributions are in reality just the same as in figure 1.1. Notice that in each of the two subgroups taken separately it is arguable that the situation at the bottom of the diagram represents greater inequality than that at the top. The richest person in the left-hand group has a higher income (\$5 rather than \$4) and the poorest person in the right-hand group has a lower income (\$6 rather than \$7); so in both cases the income gaps within each reference group widen as we go from the top of the diagram to the bottom. However, that is not the end of the story.

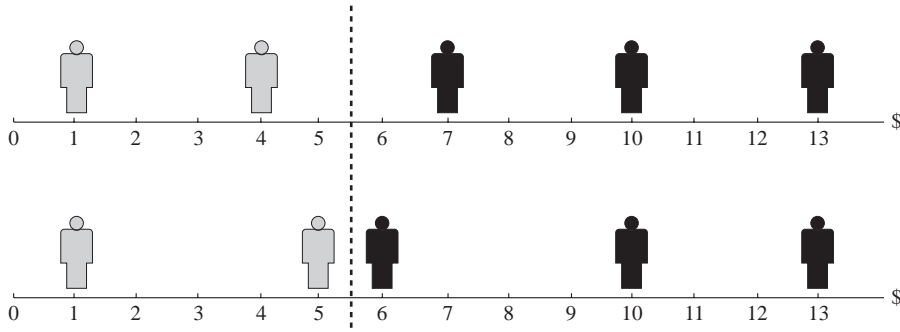


Figure 1.2. A simple distributional experiment: second view.

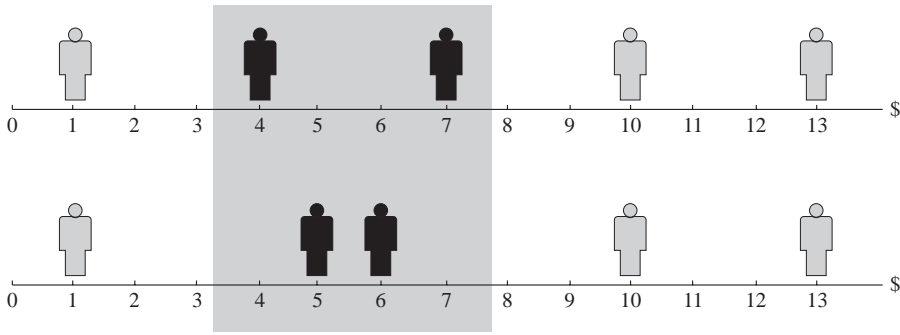


Figure 1.3. A simple distributional experiment: third view.

As a final step in the experiment have a look at figure 1.3, which again merely retouches the picture that was originally displayed in figure 1.1. In this case we have highlighted just the two persons whose incomes are directly affected in going from one distribution to the other. Put in this fashion there appears to be no argument whatsoever as to which distribution is the more unequal. Disregarding for the moment the persons whose income positions remain unchanged in going from one distribution to the other (the people with incomes \$1, \$10, \$13) it is clear that there has been an unambiguous reduction in the gap between the two remaining persons: the gap closes from (\$4, \$7) to (\$5, \$6). Put another way, if we consider the top income distribution as the ‘before’ picture of inequality, and the bottom distribution as the ‘after’ picture, then there has been a redistribution of \$1 from a richer to a poorer person: according to this view inequality *must* have fallen.

So we seem to have at least two stories about what is going on in this trivial problem of inequality comparison. How does one go about resolving the apparently contradictory pictures of inequality that emerge from even a very simple experiment such as the one we have been considering? Indeed, is there any point in trying to resolve such contradictions? Evidently the way that one tries to

answer this sort of question will strongly influence one's entire conception of the meaning of inequality comparisons.

The conventional approach to the subject has been twofold. On the one hand there is the horny-handed practical approach to evaluating empirical income distributions: having argued about the right way to measure income, and whether we should view income receivers as individuals, families, households or other groups, you pick a measure of dispersion off the shelf and you estimate this along with mean income and other statistics (we shall be looking at some of these off-the-shelf measures later in the book). Alternatively there is a theoretical approach to the problem that appears to be full of the intellectual promise that accompanies analytical rigour: this method is to introduce a particular set of axioms which collectively define what is meant by inequality comparisons and hence what is meant by economic inequality itself. It is essentially the picture of inequality characterised by figure 1.3 that is taken to be the standard paradigm for the majority of theoretical and empirical work in the economics literature.

The ambiguity of answers from the simple experiment raises issues that are considerably wider than the particular principle or principles which may be brought to bear on the particular distributional problem highlighted in figure 1.1. It prompts the question as to whether the way in which inequality is conventionally presented in the literature on economics and in other related disciplines is in some sense 'appropriate'.

1.3 A guide to the book

Those who know the economics literature on inequality will say that we have presented the pictures in our little experiment the wrong way round. That is actually quite true. We deliberately put the unorthodox view of the inequality comparison experiment first, and followed up with the standard story. The reason for this has little to do with the grand sweep of the history of thought on the subject, and much to do with a small domestic incident.

A few years ago one of the authors, Yoram Amiel, was asked by his wife Hayka (who is a school-teacher and not an economist) to explain the topic of his research. To put the main ideas over concisely he gave her a little numerical example as an illustration – something close to the experiment that we have just been considering, in fact. Faced with the choice between the two distributions, Hayka gave the 'wrong' answer. Yoram concisely pointed this out: the answer should have been clear, unambiguous and the exact opposite to hers – according to the standard theory of inequality measurement. Hayka's reply was similarly concise: 'So change the theory.'

We make no claim to be changing the theory of inequality measurement in this book. But this issue did prompt an extensive research project which, amongst other things, resulted in the book. Along the way it also raised a number of provocative questions which we make no claim to have resolved. Is the standard

theory ‘right’ and, furthermore, what does it mean for a theory to be ‘right’ in this context? These questions have in turn prompted the theoretical and empirical analysis which is reported in the following chapters.

Chapters 2 to 4 are principally about the problem of inequality in its purest form, the problem of dividing a cake of fixed size amongst a fixed number of people. Chapter 2 gives a summary guide to the standard approach in the literature on inequality measurement, chapter 3 explains the method we used to investigate the assumptions underlying the approach – a series of specially structured questionnaires – and chapter 4 reports the results of these investigations. These chapters also deal with elementary issues of how one can compare situations that have different sizes of cake or different numbers of people sharing the cake.

Chapters 5 to 7 extend the approach to three areas closely related to the pure inequality problem: social welfare (chapter 5), the relationship between income growth and perceptions of inequality (chapter 6), and poverty (chapter 7). Each of these additional topics requires additional assumptions on top of the structure used for the pure inequality problem, and we subject these to the same sort of investigative strategy; they also provide us with an opportunity to check our results on the pure inequality issue.

Chapter 8 makes a comparison of responses to our various questionnaire studies across countries and across academic disciplines, while chapter 9 sums up and suggests directions in which thinking about inequality may yet go. Finally, those readers who like to have assumptions and propositions tidied up in a concise mathematical format may want to use appendix A which sets out the main results in the conventional approach to inequality measurement; all this material has been parked in this unglamorous location because, although it has its uses, it is no substitute for thinking about inequality.

Notes

- 1 See, for example, Sen’s discussion of Rawls’s concept of the ‘original position’ (Sen 1970, p. 124).