THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

SUBJECTIVE AND PREFERENCE-SENSITIVE MULTIDIMENSIONAL WELL-BEING AND INEQUALITY

LIN YANG

Thesis submitted to the Department of Geography and Environment
of the London School of Economics and Political Science
for the degree of Doctor of Philosophy, August 2016.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I declare that this thesis consists of 34,428 words (excluding abstract, acknowledgements, lists of contents, tables and figures, and bibliography).

Abstract

This thesis proposes a comprehensive framework that allows analysis of preference-sensitive well-being and inequality. It draws together complementary aspects of attempts to operationalise a more inclusive and multidimensional definition of well-being, through subjective well-being measurement, social welfare theory, and multidimensional indices of well-being and inequality. Theoretical proposals and empirical strategies are put forward, with illustrations using data from the British Household Panel Survey.

Chapter 1 examines the underlying structure of subjective well-being, and the relationship between these subjective components of well-being and commonly targeted objective well-being indicators. A key finding is that subjective well-being follows a time-consistent dual structure of underlying 'life satisfaction' and 'emotional well-being' components. Additionally, the 'life satisfaction' component appears more strongly associated than the 'emotional well-being' component to changes in objective indicators of well-being.

The 'preference index approach', the central proposal of the thesis, is introduced in Chapter 2. Preference comparisons are inspected at the individual and subgroup level, and a preference-sensitive index of multidimensional well-being is proposed. The chapter then uses the results of Chapter 1 to support the use of longitudinal life satisfaction regression to estimate the heterogeneous preferences between objective dimensions of life.

Chapter 3 illustrates the properties of the preference index approach in terms of multidimensional inequality analysis. The main contribution is the incorporation of preference inequality as well as distributional inequality, and the ability to quantify their interdependent contributions to overall inequality in multidimensional well-being.

Acknowledgements

First and foremost, I am hugely grateful to my supervisors Giles Atkinson and Frank Cowell, who have always been there to offer their guidance and support, a nudge in the right direction, or sometimes just a much-needed chat. Thank you both for your encouragement and mentorship over the years – it has given me the open-mindedness, self-confidence and stamina to complete this thesis.

I am also indebted to Marc Fleurbaey, who was very generous to host and supervise me in Princeton when I began work on what is now the second chapter. His influence has inspired the direction of much of this thesis – thank you for your patience and for welcoming me to new concepts and people.

Many other experiences have benefited my intellectual and personal development over the course of the PhD, especially through interactions with the people I have met. In particular, I thank Milorad Kovacevic and colleagues at the Human Development Report Office for my enjoyable time there as an intern; Koen Decancq and François Maniquet for insightful discussions during my visit to Belgium; Janet Gornick and colleagues at the Luxembourg Income Study Center in New York for welcoming me to a new field; and here at the LSE, Simon Dietz at the Grantham Research Institute for helpful comments and suggestions as my review supervisor, and my newest colleagues at CASE for being supportive through the final months of my PhD. I am also grateful to the Economic and Social Research Council for the funding that made much of this possible.

Finally, I am thankful for the friendships of those both close by and far away, and for the unconditional love and support of my mother Dong-Ling, father Jian-Bo, and now husband Steve. You have kept me going when the going got tough. Thank you for putting up with me, and for teaching me about life.

Contents

List of Tables

List of Figures

Introduction

The twenty-first century has seen large portions of the developed and developing world attain ever higher per capita incomes and living standards, according to traditional measures. At the same time, however, concern has grown around inequality, social exclusion, levels of personal satisfaction and emotional health. An illustrative, but by no means isolated, example of this incongruity is Egypt's experience leading to and after the Arab Spring. While per capita GDP grew from 2006-2012 (World Bank, 2016) uninhibited by the social and political turmoil, data on reported life satisfaction from the Gallup World Poll told an altogether different story – the proportion of Egypt's population reporting a positive outlook on life[1] fell year-on-year from almost 30% to under 10% in the same period.

The spotlight within research on well-being and living standards has noticeably shifted in response to these concerns, away from a traditionally unidimensional focus on mean incomes and towards more inclusive and multidimensional definitions of well-being. Although the idea of well-being as a multidimensional concept is not new, having been advocated several decades ago by Sen (1985, 1993), Stewart (1985) and Nussbaum (2000) among many others, it has recently grown in prominence and popularity. This has owed much to the high-profile Stiglitz-Sen-Fitoussi Commission report of 2008, the establishment of the OECD's Your Better Life Index (2013) and the independent initiatives of several governments to incorporate subjective and multidimensional well-being into their national accounts (the UK, Canada and Bhutan, for example).

Among economists, the most well-known conceptualisation of multidimensional well-being is Sen's capability approach (Sen, 1985), according to which an individuals's capabilities reflect the combinations of functionings that she can attain. Sen

[1]Defined as those giving responses of 7+ when asked to rate views of their present life situation and 8+ for the next five years on a 10-point scale.

defines functionings as "the things that he or she manages to do or be in leading a life" (Sen, 1999, p. 31), and multidimensional well-being is measured in terms of an individual's capability to attain these valuable functionings. At a practical level, the capability approach is often associated with the UNDP Human Development Index (HDI) – an example of a composite index, comprised of population-level indicators of health, education and income. Despite the many criticisms levelled at the HDI since its original conception in 1990 and revised form in 2010 (in particular see McGillivray and White (1993); Noorbakhsh (1998); Sagar and Najam (1998); Ravallion (2011, 2012)), it is still by far the most famous and widely cited composite index of multidimensional well-being and is used extensively in policy and research.[2] As an application of the capability approach, however, the HDI has been criticised as being "a pale reflection of the general and ambitious methodology proposed by the capability perspective" (Fleurbaey and Blanchet, 2013, p. xiv).

More broadly, objections have been raised against the composite index approach in general. One objection is that such measures reflect only average population performance, without revealing anything about more detailed inequalities among individuals. Indeed a vast majority of composite indices of well-being that have been proposed simply add up population-level average indicators (Yang, 2014), which fails to differentiate between groups of individuals with cumulative disadvantages concentrated in multiple dimensions of well-being, such as poor health, low income, lack of education etc., and groups for whom disadvantages are spread over individuals more sparsely. Data limitations can make such information about joint attainments in multiple dimensions difficult to obtain. However, the theory underpinning composite indices often fails to consider this issue at all, reflecting a more fundamental weakness. To overcome this weakness, a method must be used that first summarises individual situations before aggregating to the population level, whereas the overwhelming approach in the literature on composite indices is to first aggregate by taking population averages in each dimension and then sum these population-level averages. Following from this, another objection is that there is no clear theoretical framework for how to aggregate dimensions at this population level. As a consequence, the implied trade-offs resulting

[2]At the time of writing, a Google search of "Human Development Index" returns over 1,800,000 results. In comparison, "index of well-being" and "well-being index" jointly return 483,000 results, and "multidimensional well-being" returns just 6,040 results. The same searches on Google Scholar return similar proportions of results for each.

from such measures are accused of being arbitrary and may not in fact represent the priorities of anyone in the population.

Despite the theoretical weaknesses, the development of new composite indices of social performance has proliferated in recent years. Bandura (2006) provides a comprehensive inventory of over 400 country-level indices covering a range of socio-economic themes. A more recent inventory by the UNDP (Yang, 2014) details 101 international composite measures of well-being and social performance. This proliferation has in part been facilitated by improved data resources, and has led to new research priorities incorporating population heterogeneity and non-income dimensions of well-being that contribute to quality of life. With this has come doubt on the reliability and relevance of data on average economic outcomes.

At the same time as the explosion in multidimensional measures of objective performance, there has been increasing interest in using subjective well-being (SWB) to assess social performance (Diener, 1994; Helliwell, 2003; Layard, 2005; Conceição and Bandura, 2008; Diener et al., 2009; Graham, 2011; MacKerron, 2012). Proponents of SWB argue that direct measures of 'happiness' should be used as a barometer of social progress. The argument goes that happiness is the ultimate goal that individuals and societies strive for, and such measures are therefore a catch-all for everything that multidimensional measures of well-being attempt to aggregate, whilst avoiding the thorny issues surrounding conceptual and mathematical construction that the latter approach entails. On the other hand, however, raw SWB scores come with a different set of problems. The known effects of adaptation to different circumstances and different tastes are not reflected in such measures – recognised by Sen (1985) as "physical condition neglect". Recent developments in the study of SWB covariates has meant that there is now a method that attempts to correct for this omission (Ferrer-i Carbonell and Frijters, 2004), and indeed such a method is incorporated in Chapter 2 of this thesis. The decision to correct for such adaptation and tastes in fact exactly implies taking a more objective, resource-based perspective on measuring well-being. As such, the approach to well-being measurement developed in this thesis is compatible with the aim of SWB research to unpack the components of a good life. The key difference is that while SWB advocators take evidence of links between life satisfaction and these components as supporting SWB as a direct measure of performance, the perspective taken here is that a measure based on the objective components themselves can provide fairer interpersonal comparisons, which is essential for such a performance measure. The problem of physical condition neglect can be explicitly corrected for by taking this approach, while still providing scope for subjective information to be incorporated by way of the aggregation procedure and weighting structure of the final measure. This alternative approach offers a way for the SWB approach to inform the more resource-based approach associated with Sen (1999) which places emphasis on objective attainments, and lies at the heart of this thesis.

With the exception of broad surveys of well-being measurement,[3] research in SWB, composite indices and multidimensional inequality have tended to evolve as somewhat separate areas, with specific proposals within each area often made in isolation of the context of the others. This thesis draws together important aspects of each of these areas, with the view that designing a widely accepted measure of well-being requires issues across disciplines to be addressed in a complementary way. The combination of interdisciplinary techniques and perspectives in the following chapters is a product of this view.

Thesis structure

In Chapter 1, the concept and structure of SWB is first empirically investigated. Since SWB is central to the chosen implementation of the 'preference index approach' introduced and developed in Chapters 2 and 3, it is crucial to test the structure of the SWB concept and its relevance to designing the proposed index. In economic studies involving SWB, it is often taken as given that the structure of the underlying components of SWB are fixed over time, and that two components of SWB in particular – 'life satisfaction' and 'affect', or 'emotional well-being' – are consistently defined such that they may be used unambiguously in further analysis. A prominent example of such a perspective is the World Happiness Report (Helliwell et al., 2015). This may be due, at least in part, to the increasing availability of micro and macro level datasets that incorporate single-item indicators of SWB, along with the other socio-economic variables of interest to the analyst. On the other hand, the psychology literature tends

[3]Fleurbaey and Blanchet (2013) provide such a survey with an explicit and comprehensive discussion of the overarching theoretical connections between different approaches.

to delve further into the composition of the different components of SWB (Diener et al., 1999), requiring more specialised multi-item survey instruments. In Chapter 1, exactly such a multi-item instrument is exploited using the British Household Panel Survey (BHPS), which offers response data for several emotional and satisfaction variables in different domains at the individual level and over time. If SWB is to be used to incorporate subjective preferences into a multidimensional measure of well-being, then it is important that the chosen SWB indicator does in fact measure the same phenomenon over time. The first contribution of this chapter is therefore to investigate the composition of SWB in terms of underlying items and latent components, and whether this composition remains consistent over time. The findings of the chapter indicate that this is the case to a suitable degree, an encouraging result that allows progress on to further analysis using SWB in a longitudinal context, which is required in Chapter 2. The second contribution of Chapter 1 is an investigation into the links between latent components of SWB and objective dimensions of life (health, education and income). This investigation uses structural equation modelling to examine the responsiveness of the latent 'life satisfaction' and 'emotional well-being' components of SWB to the objective indicators. The 'life satisfaction' component of SWB is found to be more responsive than the 'emotional well-being' component to attainment in the objective indicators examined, supporting the use of life satisfaction response data to uncover the implicit preferences of individuals between objective dimensions of life. These objective indicators have not been collectively studied alongside the latent longitudinal structure of SWB in this way before.

The 'preference index approach', the central proposal of this thesis, is introduced in Chapter 2. Whereas the analyses in Chapter 1 are conducted at the population average level, Chapter 2 goes further to inspect preference patterns at the individual and subgroup level, and proposes a framework to use these preferences in order to design a preference-sensitive index of multidimensional well-being. The chapter uses the invariance results of Chapter 1 to support the use of longitudinal life satisfaction regression to estimate preferences between dimensions of life. It is argued, using theory proposed according to the 'equivalence approach' (Pazner and Schmeidler, 1978; Fleurbaey and Maniquet, 2011), that these preferences can be used to characterise a measure of multidimensional well-being that reflects interpersonal preference heterogeneities whilst being interpersonally comparable. This proposed 'preference index' is described axiomatically and empirically illustrated in further analysis using the BHPS data. This includes estimating the preferences of individuals separated according to age group and education level, and an initially unexpected finding of weaker precedence of the health dimension within older groups compared to younger groups. Another key result is the strong prioritisation of good health across all preference types, compared to relatively weak preferences for income. It is shown that the picture of well-being in the UK is quite different if preference heterogeneity is taken into account, compared to the picture painted by solely income or SWB, or by standard multidimensional measures. The 'preference index' proposal challenges the popular practice of assuming existence of a readily available cardinal measure of well-being, identically specified across individuals. It also challenges the popular practice in composite index approaches of using population averages to seek an aggregate assessment over the population. In these respects, the theoretical framework proposed and empirical operationalisation of a preference-sensitive multidimensional well-being index, as in Chapter 2, are new contributions.

Chapter 3 elaborates the properties of the preference index approach in terms of multidimensional inequality analysis that integrates preference heterogeneities. The incorporation of distributional inequality as well as preference inequality, and the ability to quantify these in an overall multidimensional analysis framework is highlighted as the main contribution of this chapter. The approach is presented in the context of existing multidimensional concepts, for example in relation to sensitivity to cumulative deprivations, inequality aversion, and conventionally proposed properties of multidimensional inequality indices. Given the importance of being able to break down inequality into its contributing influences, it is demonstrated that decomposition analysis can be performed using the preference index measure to break down the contributions to inequality by subgroup and by preference-sensitive dimension contributions. These two forms of decomposition are developed from proposals suggested by Lasso de la Vega and Urrutia (2005) and Shorrocks (2013) respectively. Empirically, it is shown that the new tools proposed can provide insights into the composition and evolution of multidimensional inequality, including dimension and preference interrelationships, that existing approaches cannot offer. Health inequality is found to be the major dimension contributor to preference-sensitive inequality, partly explaining the lower preference-sensitive inequality among the young and well-educated (who are generally quite healthy). Education inequality among the well-educated proves not to matter so much for preference-sensitive inequality in this group, whereas education inequality among the less educated matters notably more in this respect. Interestingly, diversity in preferences themselves outweighs inequality in incomes in terms of their contributions to preference-sensitive inequality in the population overall. Overall preference-sensitive inequality is shown to have slightly worsened over the 1996-2008 period analysed, with such inequality between individuals in the same preference type outweighing that of individuals between different preference types. The evolution of this within-preference inequality and between-preference inequality appears to have moved in opposing directions over the analysis period, with the evolution of overall preference-sensitive inequality appearing to be more a product of the changing distribution of dimension attainments rather than shifts in the preference structure due to changing subgroup population shares.

These analyses, along with the concrete theoretical proposals and surrounding issues of the proposed preference index approach, will be discussed in detail in the following chapters. First, if such an approach is to incorporate subjective preference information, then a proper understanding of the relevant aspects of SWB needed to estimate these subjective preferences is required; this is the subject of the first chapter.

Chapter 1

On the structure of subjective and objective well-being, using data from the British Household Panel Survey

1.1 Introduction

Since the establishment of the United Nations Human Development Report and its accompanying Human Development Index in the 1990s, policy-makers have come to take ever more seriously the need for alternative measures of human progress and well-being other than economic growth. There have since been numerous attempts to construct more encompassing measures of social and economic well-being by combining different dimensions thought to influence quality of life. Such initiatives include the UK's proposal for a General Well-being Index, "Beyond GDP" measures launched by a number of other countries, and the OECD's Better Life Index, to name but a few.

The economics of well-being is a burgeoning field[1] and, increasingly, interest is not only limited to measuring material and objective aspects of well-being but also subjective well-being (SWB). The concept of SWB is a broad one, and is concerned with the "on-line" (Diener et al., 1999) psychological and emotional experience of events

[1] For a survey of the literature, see Boarini et al. (2006), as well as recent well-publicised contributions from Helliwell et al. (2015), Graham (2011) and Layard (2005) among others.

that occur in life, as well as the satisfaction and meaningfulness gained from domains of life overall. These various elements of SWB are often classified into 'affect', 'life satisfaction' and 'eudaimonia' (OECD, 2013b). The precise boundaries, however, are sometimes blurred and in practice subject to data limitations. In particular, the idea of eudaimonia – the good psychological state of deeming one's life to be meaningful – has proved especially problematic in applied contexts, with relatively little evidence on the reliability of eudaimonic measures (OECD, 2013b). Developing a line of literature from Bradburn and Caplovitz (1965) and Andrews and Withey (1976), Lucas et al. (1996) proposed a structure for SWB involving only an 'affective' and 'cognitive' component – the affective component composed separately of both positive and negative emotional states, and the cognitive component concerning the evaluation of satisfaction in various domains of life. These correspond respectively to the 'affect' and 'life satisfaction' elements of SWB, dispensing with the more nebulous notion of 'eudaimonia'. For the reasons just stated, this more practical two-element notion of SWB is preferred here. Note, however, that although a distinction is often made between 'positive affect' and 'negative affect', the findings in this chapter do not find a demarcation between these two aspects of 'affect'. Therefore, references to the dual structure for SWB pertain to 'affect' as a single component and 'life satisfaction' as the other. As Diener et al. (1999) remark, "the degree of independence between momentary pleasant and unpleasant affect is still debated".

The contribution of this chapter is two-fold. First, an empirical examination of the existence and time-consistency of this dual SWB structure is presented, that is, its constitution of 'affect' (or 'emotional well-being') and 'life satisfaction' components. This examination uses factor analysis of data from the British Household Panel Survey (BHPS), with SWB as the latent construct of interest. To gain an understanding of whether its constitution remains constant over time, longitudinal factor invariance of the latent SWB construct is evaluated under varying degrees of strictness. The finding is that factor invariance does hold to a certain degree, to the extent that the strength of associations between the factors is consistent over time. In the second contribution, objective policy indicators (of health, education and income) are introduced to investigate how these are linked to variation in the dual SWB construct. This investigation uses a structural equation model (SEM) to examine the responsiveness of the 'life satisfaction' and 'emotional well-being' components of SWB to the objective well-being indicators. The finding here is that the 'life satisfaction' component of SWB generally appears to be more responsive than the 'emotional well-being' component to attainment in the objective indicators examined. This is consistent with broader findings in fragmented parts of the literature, however there do not appear to be any studies that examine these objective indicators collectively along with the latent longitudinal structure of SWB.

The rest of the chapter is structured as follows: Section 1.2 introduces the SWB data and variables; the longitudinal factor analysis of SWB is presented in Section 1.3; the objective well-being variables and SEM are presented and interpreted in Section 1.4; and Section 1.5 concludes.

1.2 Capturing subjective well-being

The use of self-reported surveys has proliferated in recent years in the field of well-being economics, the rationale being that there is no-one better placed to judge one's well-being than the individual herself (Diener, 1994). Although there are valid criticisms concerning the use of self-reported measures, such as contextual influences and biases, research across large samples of individuals across countries and over time has established that these have a limited impact on generally consistent patterns in the determinants of SWB (Schimmack and Oishi, 2005; Graham, 2011).

The idea of SWB itself is defined by Veenhoven (1984) as the degree to which an individual judges the overall quality of her life as a whole in a favourable way, with the subjective element reflecting that social and economic environments do not, by themselves, fully characterise quality of life (Diener and Suh, 1997). As discussed, this chapter works with a modified form of the structure proposed by Lucas et al. (1996) of a two-element SWB construct, involving on the one hand a cognitive component of evaluating satisfaction in various domains of life, and on the other hand an affective component of emotional well-being. In attempts to capture and measure SWB, early instruments for evaluation typically only used a single question (Albuquerque et al., 2012). Developments in research over recent decades in the psychology of well-being has led to improved multi-item instruments, allowing higher reliability and validity than is possible with a single item (Diener et al., 2009). The availability of multiple indicators of a latent construct for each person is crucial to be able to estimate the required structural equation models, and additionally the availability of these multiple indicators over multiple time periods allows the longitudinal invariance of the SWB construct to be tested.

1.2.1 UK data

The British Household Panel Survey (BHPS) is one such multi-item data resource, containing in it questions about various aspects of respondents' SWB. The BHPS is a rich panel of SWB and socio-economic data at both individual and household level, comprising 18 waves of observations from 1991 to 2008. Data from waves 6-18 (corresponding to the years 1996-2008) are used, since SWB questions were not incorporated in the waves prior to wave 6. Specifically comparisons are made for the six-year periods between waves 6, 12 and 18, in order to cover the whole range of the data whilst maintaining a manageable degree of complexity for the longitudinal analyses. The final dataset contains 151,721 observations over the whole period.

1.2.2 Subjective well-being questions

The analysis makes use of data extracted for the following items: Firstly, responses to questions asking "have you recently…

- Felt that you were playing a *useful* part in things?
- Felt constantly under *strain*?
- Been able to *enjoy* your normal day-to-day activities?
- Been losing *confidence* in yourself?
- Been feeling reasonably *happy*, all things considered?"

All item responses are coded from 1 to 4, with lower scores indicating a more negative response (i.e. feeling less useful, more under strain, enjoying activities less, losing confidence more, feeling less happy) and higher scores indicating a more positive response.

Secondly, responses to questions asking respondents to indicate "how dissatisfied or satisfied you are with the following aspects of your current situation:

- Your health
- The income of your household
- Your house/flat
- Your partner
- Your job
- Your social life"

Item responses are again coded from low to high, with 1 indicating "not satisfied at all", and 7 indicating "completely satisfied". To facilitate the comparison of model parameters of items reported on different scales, items originally reported on the 1 to 7 scale are linearly rescaled to range from 1 to 4. The BHPS does also contain an item on satisfaction with life overall, which is used in Chapter 2. However, it is excluded as an item in the following factor analysis, since the purpose of this analysis is precisely to test the multiple items for reliability of the overall underlying concept of satisfaction measured. This would not be possible using the single overall satisfaction item. The latent satisfaction factor that emerges from the subsequent analysis as one of the two underlying components of subjective well-being is, however, correlated with the overall satisfaction item to check for consistency between the two. This is presented at the end of Section 1.3.2.

1.3 Factor analysis of latent well-being components

1.3.1 One-factor model for the satisfaction items

To begin the investigation into the structure of SWB, as a starting point the six satisfaction items for *health*, *household income*, *housing*, *partner*, *job* and *social life* are first examined cross-sectionally for wave 6. Analyses are restricted to respondents with complete response data over the waves, since cross-sectional models are carried forward to the longitudinal analysis. Following the recommendations of Costello and Osborne (2005), maximum likelihood is used as the factor analysis extraction method, with multiple test runs to check for possible meaningful factors and oblique rotation to aid interpretation.

Since the manifest items under examination are ordered categorical variables measured on Likert-type scales, ideally an approach that deals explicitly with ordinal data in latent variable models should be used. For example, the item response function approach is a generalisation of the logit or probit factor model for binary items to the case of many ordered categories, and the underlying variable approach is factor analysis using the polychoric correlation matrix – the correlation matrix between the theorised normally distributed continuous latent variables underlying the observed ordinal variables. In applications where the number of categories is large for all items (six or seven as a rule of thumb), the ordinal items are often treated as interval variables, with standard linear factor models fitted using the Pearson's correlations between these "pseudo-continuous" items (Bartholomew et al., 2008). While this may not seriously affect the broad conclusions of the analysis, the estimates of factor loadings may be biased. Since several of the items have only four categories, the underlying factor approach using the polychoric correlation matrices is compared alongside the standard factor analysis model to check whether there is a substantive difference between the loadings.

	Health	*Income*	*House*	*Partner*	*Job*	*Social*
Health	1					
Income	0.338	1				
House	0.228	0.363	1			
Partner	0.086	0.091	0.091	1		
Job	0.236	0.210	0.005	0.211	1	
Social	0.360	0.362	0.334	0.110	0.121	1

Table 1.1: Wave 6 Pearson's correlation matrix for the satisfaction items

An inspection of the Pearson and polychoric correlation matrices given in Tables 1.1 and 1.2 reveal that all six items are indeed positively correlated with each other, with slightly higher polychoric than Pearson's correlations for all items. Recall that a higher score for each item indicates that an individual feels more positively about her own attainment in that item. *Partner*, however, has only weak correlations with the other items in both correlation matrices, and for this reason it may be expected that *partner* is unlikely to share any common factors. *Job* is also only weakly correlated

	Health	*Income*	*House*	*Partner*	*Job*	*Social*
Health	1					
Income	0.374	1				
House	0.263	0.416	1			
Partner	0.119	0.122	0.158	1		
Job	0.272	0.249	0.016	0.247	1	
Social	0.398	0.397	0.392	0.183	0.162	1

Table 1.2: Wave 6 polychoric correlation matrix for the satisfaction items

with *house* and *social*, and this may also have implications on the resulting model.

Preliminary testing of a two-factor and single-factor exploratory model was carried out using maximum likelihood estimation for both the standard and polychoric correlation matrices. The two-factor models resulted in Heywood cases, which occur when the maximum likelihood estimation method converges to unique variance values that are negative. This is an indication that these mdoels attempted to fit too many factors. Consequently, a two-factor model was rejected and only the single-factor model results are reported. The eigenvalue of the common factor is 1.438 and 1.661 for the standard and polychoric models respectively. The factor loadings and uniquenesses are presented in Tables 1.3 and 1.4.

	Factor 1	Uniqueness
Health	0.551	0.696
Income	0.639	0.591
House	0.499	0.751
Partner	0.186	0.966
Job	0.277	0.924
Social	0.605	0.634

Table 1.3: Factor loadings and unique variances for the single-factor exploratory satisfaction model

With the exception of *partner* and to some extent *job,* the positive and large factor loadings point to a conclusion that the common factor represents a general summary of the remaining items. The high uniqueness of *partner* and *job* further indicate that

	Factor 1	Uniqueness
Health	0.576	0.669
Income	0.661	0.563
House	0.552	0.695
Partner	0.260	0.932
Job	0.315	0.901
Social	0.649	0.578

Table 1.4: Factor loadings and unique variances for the polychoric exploratory single-factor satisfaction model

neither item is informative in this single-factor model, and therefore both are dropped from the satisfaction model in subsequent analyses. This is not a surprising result given the initial observations from the correlation matrices.

	Health	*Income*	*House*	*Partner*	*Job*	*Social*
Health	-0.000					
Income	-0.013	0.000				
House	-0.048	0.044	-0.000			
Partner	-0.017	-0.027	-0.001	-0.000		
Job	0.084	0.033	-0.132	0.159	-0.000	
Social	0.025	-0.024	0.032	-0.005	-0.047	0.000

Table 1.5: Correlation residuals for the polychoric single-factor satisfaction model

The item with the lowest uniqueness, or highest communality,[2] and therefore best represented by the common factor is *income*. *House* is least well-represented, though the differences in communality are not large. Since likelihood ratio tests are known to be over-sensitive in the case of large sample sizes, model fit is determined by examining residuals between observed and expected correlations, presented in Table 1.6. As a widely used rule of thumb, correlation residuals with absolute values greater than 0.1 suggest that the model does not sufficiently explain the corresponding sample correlation (Kline, 2010, p. 171). By this measure, with the exception of two residuals for *job,* which have been elected to be dropped in any case, the single-factor satisfaction

[2]The communality of a variable is the sum of the loadings of this variable on all factors in the model (Rietveld and van Hout, 1993). The uniqueness is equal to 1 - communality.

	Health	*Income*	*House*	*Partner*	*Job*	*Social*
Health	-0.000					
Income	-0.006	0.000				
House	-0.055	0.051	-0.000			
Partner	-0.031	-0.050	0.014	-0.000		
Job	0.091	0.041	-0.157	0.165	-0.000	
Social	0.024	-0.032	0.034	0.015	-0.042	0.000

Table 1.6: Correlation residuals for the single-factor satisfaction model

model provides a satisfactory fit for the data.

1.3.2 Two-factor model for subjective well-being

Continuing with the analysis of wave 6, the exploratory analysis can now be expanded to include items relating to the affective aspects of SWB. Figure 1.1 represents the two-factor exploratory model of interest with a path diagram. Standard path diagram notation is used: rectangles represent observed variables x; ellipses represent latent variables ξ; single-headed arrows represent hypothesised directional effects of one variable on another, such as factor loadings; and curved double-headed arrows represent undirected relationships between independent variables, such as covariances or correlations. Residual variances δ are a factor analytic term for the variances of observed variables not explained by the factors, and can be considered special types of latent variable. These are not placed in ellipses to aid representational clarity and are represented by single-headed arrows pointing to endogenous variables.

Although the same comparisons were carried out between standard and polychoric models for the two-factor SWB model as for the single-factor satisfaction model, for the sake of brevity only the polychoric model results are reported since these explicitly take into account the ordinal nature of the variables. The substantive conclusions produced by both models turned out to be the same, though communalities and loadings were generally larger under the polychoric model than the standard model, indicating that the underlying factor model provides a stronger representation of these ordinal items than the standard model.

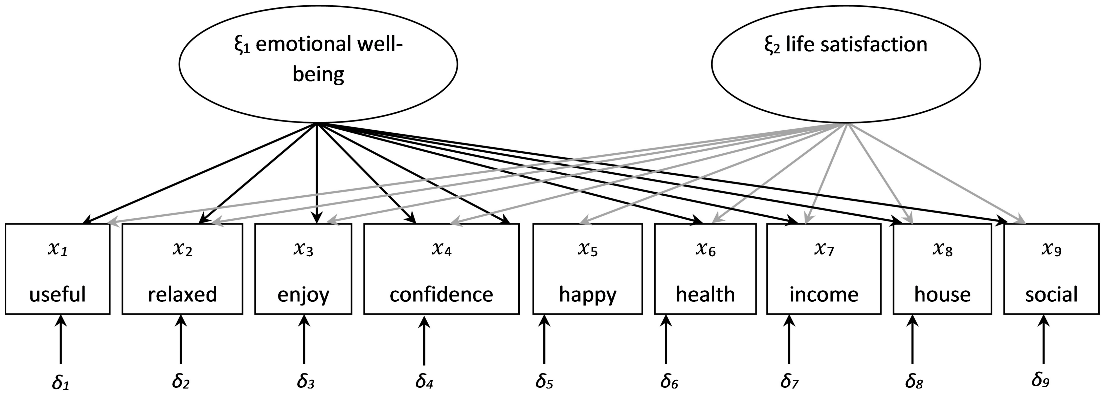

Figure 1.1: Path diagram of an exploratory model for SWB

	Useful	*Relaxed*	*Enjoy*	*Confidence*	*Happy*	*Health*	*Income*	*House*	*Social*
Useful	1.000								
Relaxed	0.198	1.000							
Enjoy	0.427	0.494	1.000						
Confidence	0.424	0.550	0.444	1.000					
Happy	0.454	0.492	0.640	0.533	1.000				
Health	0.246	0.306	0.325	0.356	0.286	1.000			
Income	0.145	0.276	0.186	0.252	0.220	0.370	1.000		
House	0.058	0.244	0.125	0.194	0.184	0.263	0.417	1.000	
Social	0.209	0.370	0.308	0.353	0.319	0.403	0.398	0.395	1.000

Table 1.7: Polychoric correlation matrix for all SWB items

An inspection of Table 1.7 confirms that there are indeed positive correlations between items, especially among the affect items. Three-factor, two-factor and single-factor models were tested, with the three-factor model being immediately rejected due to a Heywood case, a model-fitting error indicating an attempt to fit too many common factors. Between the two-factor and single-factor models, the two-factor model resulted in lower Bayesian and Akaike Information Criterion values and smaller residual correlations, indicating better model fit. The additional interpretive correspondence of the two-factor model with the proposed dual SWB structure leads the two-factor model to be retained over the single-factor model. The resulting eigenvalues for the first and second common factors are 3.158 and 0.864 respectively.

The factor loadings presented in Table 1.8 are interpreted as the correlations between

	Factor 1	Factor 2	Uniqueness
Useful	0.509	-0.203	0.700
Relaxed	0.660	0.017	0.564
Enjoy	0.725	-0.256	0.409
Confidence	0.688	-0.050	0.525
Happy	0.771	-0.252	0.343
Health	0.505	0.261	0.677
Income	0.431	0.486	0.578
House	0.358	0.492	0.630
Social	0.550	0.380	0.554

Table 1.8: Factor loadings and unique variances for the polychoric two-factor SWB model

the observed items and each respective latent common factor. It can be seen that all items have moderately large positive loadings on the first factor, indicating a general metric of SWB. The second factor seems to make the contrast between emotional well-being and life satisfaction.

Other than one exception between *relaxed* and *useful,* the correlation residuals, presented in Table 1.9, are smaller than 0.1 and indicate a satisfactory fit. The underlying structure can be further clarified by "rotating" the factors, in order to find an optimal expression of the solution with a simple structure so that each item has a large loading on one factor and very small loading on the other.

Table 1.10 details the resulting factor loadings from an oblique rotation, allowing for correlation between the factors. Correlated factors are highly plausible in the context of SWB, since it is intuitive that external conditions and cognitive mechanisms that determine emotional disposition would also influence an individual's outlook and satisfaction with life. The solutions quite clearly indicate that the items can be interpreted as relating to two underlying SWB concepts – the emotional well-being component modelled earlier, and a life satisfaction component. The simple structure observed from the oblique rotation shows substantial correlation of 0.512 between these two components, which confirms the intuition that day-to-day emotional well-being is positively related to how a person evaluates satisfaction with life overall. Using these factor loadings to explicitly construct variables for the two latent components, it is also

	Useful	*Relaxed*	*Enjoy*	*Confidence*	*Happy*	*Health*	*Income*	*House*	*Social*
Useful	0.000								
Relaxed	-0.135	0.000							
Enjoy	0.006	0.020	0.000						
Confidence	0.063	0.097	-0.067	0.000					
Happy	0.010	-0.013	0.017	-0.010	0.000				
Health	0.042	-0.032	0.026	0.022	-0.038	0.000			
Income	0.024	-0.017	-0.002	-0.020	0.011	0.025	0.000		
House	-0.024	-0.001	-0.008	-0.027	0.032	-0.046	0.023	0.000	
Social	0.006	0.001	0.008	-0.006	-0.009	0.026	-0.023	0.011	0.000

Table 1.9: Correlation residuals for the polychoric two-factor SWB model

	Oblique rotation		Initial solution	
	Factor 1	Factor 2	Factor 1	Factor 2
Useful	0.580	-0.070	0.509	-0.203
Relaxed	0.518	0.224	0.660	0.017
Enjoy	0.799	-0.063	0.725	-0.256
Confidence	0.596	0.157	0.688	-0.050
Happy	0.833	-0.044	0.771	-0.252
Health	0.188	0.449	0.505	0.261
Income	-0.061	0.679	0.431	0.486
House	-0.126	0.663	0.358	0.492
Social	0.124	0.596	0.550	0.380
Correlation	0.512		0	

Table 1.10: Comparison of polychoric factor model results after oblique rotation

useful to check their correlations with the overall life satisfaction variable available in the BHPS. The correlation between latent ‘emotional well-being’ and the BHPS ‘overall life satisfaction’ variable is 0.5798, and between latent ‘life satisfaction’ and the BHPS ‘overall life satisfaction’ variable is 0.7213. Being measures of the same proposed underlying concept, the stronger positive correlation between the latent and BHPS life satisfaction variables is as one would expect, and supports the consistency of the BHPS overall life satisfaction variable, which is used in Chapter 2.

1.3.3 Two-factor measurement model and longitudinal factor invariance

The conclusions obtained in Section 1.3.2 can now be used to fit a two-factor confirmatory model of the two distinct 'emotional well-being' and 'life satisfaction' components, with correlation allowed between the two factors. In contrast to the approach taken in Section 1.3, constraints are now imposed so that factor loadings of the 'emotional well-being' component on the life satisfaction items are zero, and similarly so that the factor loadings of the 'life satisfaction' component on the emotional well-being items are zero. The path diagram in Figure 1.2 provides a representation of the model to be estimated. This confirmatory model will form the measurement part of the structural equation model (SEM) to be introduced in Section 1.4, and the two latent components of SWB will then become endogenous variables in the SEM.

Before investigating the SEM however, it will be additionally verified whether SWB, as measured by the two-factor model, remains invariant in its structure over time. This is possible since the dataset contains multiple item indicators of the latent construct at each wave. In the case of only a indicator for each person at each wave, it would be necessary to rely on the assumption that the single item reflects the same construct over time. However since richer data are available, rather than assume that the model holds in a longitudinal context, a test can be provided for the longitudinal invariance of the latent two-factor SWB construct. If the construct can be shown to be invariant over time, then stronger conclusions can be warranted (Widaman et al., 2010).

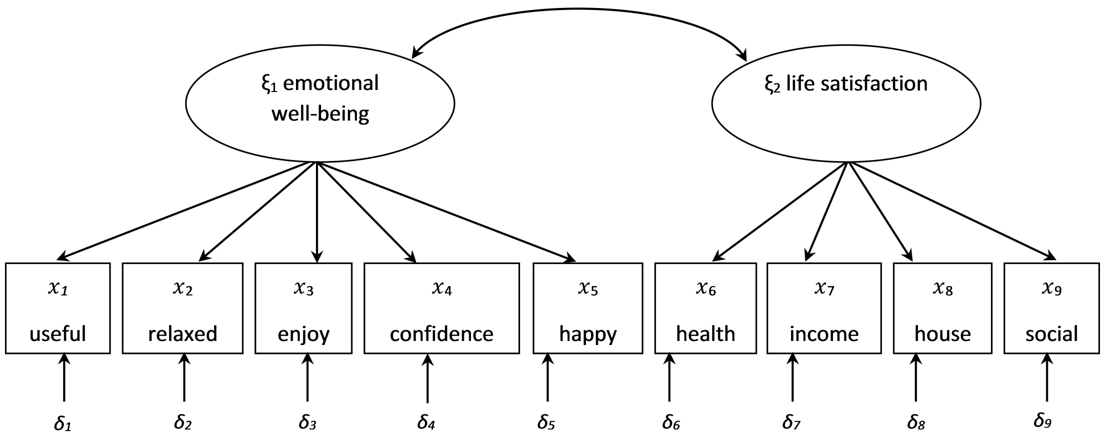

Figure 1.2: Path diagram of the SWB measurement model

Hierarchy of invariance and invariance testing

The two-factor confirmatory model is shown in Fig 1.2, with five observed variables for the emotional well-being component and four observed variables for the life satisfaction component. This model is used to measure the well-being construct on each of three occasions to span the period of data available, namely waves 6, 12 and 18, and then to evaluate how well the models respond to restrictions imposed by varying degrees of factor invariance across waves.

In more detail, factor invariance concerns the equivalence of the relationships between observed and latent variables, either across multiple groups or across time periods. The idea is to verify that the factors are indeed measuring the same underlying construct over the groups or periods. In this case, the aim is to check that the set of indicators measuring the latent SWB construct have the same factor structure over the waves of response data. Several parameters of the model can be tested for invariance, with the testing procedure following a hierarchy of nested models that each restricts more parameters to be equal across waves. The greater the number of restricted parameters the model can assimilate without sacrificing model fit, the higher the degree of invariance. This hierarchical procedure has been developed and refined by Meredith (1993); Byrne (1994); Cheung and Rensvold (1999) and Gregorich (2006) among others. On moving to each higher degree of invariance, constraints from the previous model are retained as additional constraints are added.

There are four degrees of invariance in the hierarchy of parameter constraints:

1. Configural invariance (no parameter restrictions)

2. Metric invariance (restrict loadings)

3. Strong invariance (restrict loadings and intercepts)

4. Strict invariance (restrict loadings, intercepts and residuals)

The degree of longitudinal factor invariance attained by the model is an important issue, since it has a bearing on the consistency of conclusions from cross-sectional analyses, in terms of their applicability to other waves and comparing conclusions across waves. Configural invariance would confirm that the same pattern of path restrictions imposed by the confirmatory model over time periods is indeed justified. Metric invariance would validate the comparison over time of the strength of associations between the latent and observed variables. Strong invariance would validate comparisons of the latent construct scores on the same scale over time. Strict invariance is necessary for analyses that require these scales to be equally reliable over time, such as statistical inference across time periods for an SEM with latent variables.

Configural invariance is tested by specifying the same measurement model for each wave, that is, including the same factors, observed items and relationships between factors and items. All parameters are estimated freely for each wave with no equivalence restrictions across waves – only that the same observed items load on the same factors in each wave (Meredith, 1993; Gregorich, 2006). Since this is the weakest type of invariance, this measurement model must provide a satisfactory fit in order for any degree of measurement invariance to hold. Configural invariance is required to conclude that the latent factors in a model are measured through the same qualitative, but not necessarily quantitative, structure over time.

If configural invariance is satisfied, then we can test for metric invariance (Meredith, 1993; Gregorich, 2006), which refers to the equivalence of factor loadings over time. This effectively requires that the scores on the scales of observed items retain the same meaning over time periods, or that unit changes in item scores are associated with the same change in factor scores at each wave Brown (2006). If this is the case, then it can be concluded that latent factors do manifest themselves quantitatively through the same item scores over time periods.

Assuming metric invariance is satisfied, the next step is then to test for strong invariance. Strong invariance refers to the restriction of both factor loadings and intercepts to be equal over time, and implies that not only the item scores, but the scaling of latent factors is equivalent over the time periods considered.

Finally, strict invariance requires satisfaction of all previous degrees of invariance, with the added requirement that item residuals are constrained to be equal over time. Invariance of the residuals implies that the scaling between items and factors are equally reliable in each time period examined.

Each of the invariance levels is tested by comparing pairs of models – one imposing the requisite cross-wave parameter constraints for the degree of invariance in question, and one estimated with the constraints of the next weakest invariance level in the hierarchy. The model with and without the stricter parameter constraints are then compared using a χ^2 difference test and changes in the Comparative Fit Index (CFI) value, as recommended by Cheung and Rensvold (2002) and Byrne and Stewart (2006). If the χ^2 difference for the two models is not statistically significant, or if there is no worsening of the CFI, then the fit of the model with stricter constraints is deemed comparable to the fit of the model with weaker constraints, and therefore the model satisfying stricter invariance is retained. Provided this is the case, then the same pair-wise model comparisons can be made for the incrementally stricter degrees of invariance. When the χ^2 difference or CFI *does* indicate a worsening of fit, it is still possible to test for partial invariance by constraining the parameters of only some items. Since the concern at present is with the SWB construct as measured by all of the items in the model, this line of analysis is not pursued. In addition to the χ^2 and CFI values, changes are examined for a number of other fit indices commonly used in the literature as a supplementary diagnostic tool.

Invariance test results

Estimation results from the fitted measurement models are presented in Table 1.11 under increasingly strict invariance models. To ensure model identification, variances of the latent factors are scaled to 1 with means of 0. The fit index and χ^2 values for each model are presented in the corresponding column of Table 1.12. Three fit indices are examined: the CFI as just discussed, and additionally the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC).

The CFI is an incremental fit index that measures the relative improvement in the fit of a model over that of a baseline model, typically the independence model (Kline, 2010). Index values fall inside the range 0-1 and good model fit is indicated by a CFI value of 0.9 or greater (Hu & Bentler, 1999). Models with a higher CFI value indicate better fit.

The AIC is a comparative fit measure, and is interpreted by examining the change in AIC value of one model relative to another. In selecting among models, the best fitting model is that with the lowest AIC value.

The BIC is interpreted similarly to the AIC, with lower values indicating better fit.

However, the BIC places a greater penalty on model complexity, and as a result it can be the case that in some circumstances the AIC favours a larger model than the BIC.

The model χ^2 values allow χ^2 difference tests to be carried out, which test the null hypothesis that the model with more parameter restrictions fits no worse than the model with less parameter restrictions. The relevant test statistic is the difference of the χ^2 values of the two hierarchical models in question. This χ^2 difference statistic is distributed with degrees of freedom equal to the difference of those of the two models. If the χ^2 difference statistic is significant, then the null hypothesis is rejected and the less restrictive model in the hierarchy is retained over the more restrictive model. On the other hand if the χ^2 difference statistic is insignificant, both models are deemed to fit equally well and the more restrictive model can be accepted without sacrificing model fit. In Table 1.12, χ^2 values are reported with degrees of freedom in parentheses, along with the χ^2 difference statistic and degrees of freedom.

Latent factor	Factor loadings	Free to vary (wave 6)	Metric invariance	Strong invariance	Strict invariance
Emotional	*Useful*	0.278 (0.007)	0.294 (0.004)	0.294 (0.004)	0.294 (0.004)
well-being	*Relaxed*	0.478 (0.009)	0.484 (0.005)	0.483 (0.005)	0.484 (0.005)
	Enjoy	0.397 (0.007)	0.409 (0.004)	0.409 (0.004)	0.409 (0.004)
	Confidence	0.491 (0.009)	0.504 (0.005)	0.504 (0.004)	0.505 (0.005)
	Happy	0.415 (0.006)	0.420 (0.004)	0.420 (0.004)	0.420 (0.004)
Life	*Health*	0.466 (0.010)	0.468 (0.005)	0.468 (0.005)	0.468 (0.005)
satisfaction	*Income*	0.491 (0.011)	0.488 (0.005)	0.487 (0.005)	0.487 (0.005)
	House	0.368 (0.010)	0.364 (0.005)	0.364 (0.005)	0.363 (0.005)
	Social	0.505 (0.009)	0.485 (0.005)	0.485 (0.005)	0.487 (0.005)

Table 1.11: Coefficients of the measurement model under different degrees of longitudinal invariance (standard errors in parentheses)

The first results column in Table 1.11 reports factor loadings under configural invariance, with model parameters, including factor loadings, free to vary across waves. Since factor loadings are estimated freely for each wave, results are presented for one wave only (wave 6 in this instance). Next, metric invariance is checked where factor loadings are constrained to be equal across waves. Observe that the metric invariance model is parameterised with slightly different factor loadings and standard errors, which is not

Fit statistic	Degree of invariance			
	Free	Metric	Strong	Strict
CFI	0.924	0.924	0.920	0.904
χ^2	4453.56 (78)	4481.10 (96)	4705.31 (114)	5676.35 (132)
$\Delta\chi^2$		27.54 (18)	224.21 (18)	971.04 (18)
AIC	436576.149	436567.691	436755.900	437690.941
BIC	437261.588	437106.250	437147.579	437935.741

Table 1.12: Fit statistics for the measurement model (degrees of freedom in parentheses)

surprising given that one model provides estimates for a single wave (wave 6) whereas the other provides estimates taking into account three waves. The model under metric invariance conditions fits just as well according to the CFI, and indeed both the AIC and BIC indicate that the metric invariance model even improves slightly on the fit of the configural invariance model. The χ^2 difference statistic with 18 degrees of freedom, reported in the $\Delta\chi^2$ row, is insignificant at the 5% significance level. The conclusion reached by the fit indices can therefore be reaffirmed, and the metric invariance model is retained over the weaker configural invariance model.

Proceeding to the next degree of invariance in the hierarchy, observe that the factor loadings and standard errors of the model under strong invariance conditions change only very slightly. However, the fit of the strong invariance model is worse than the fit of the metric invariance model according to all fit measures.[3] Strong invariance cannot therefore be assumed, nor any of the remaining invariance degrees in the hierarchy. Even though the requirements for satisfying further degrees of invariance are not met, the model results for strong and strict invariance are reported nonetheless, which further constrain intercepts and residuals respectively. Although the fit indices clearly show a worsening of fit for each further invariance model, the loadings and standard errors do not in fact change very much at all.

In conclusion, while there is evidence for only metric longitudinal invariance and

[3]The BIC of the strong invariance model is still lower than that of the weak configural invariance model, however a known particularity of BIC being a function of N may explain this inconsistency. As strong invariance implicitly groups the data into within-wave clusters, the choice of N between the pooled or within-group value becomes non-trivial, which in turn affects the calculation of BIC. (A detailed explanation of this problem is given in StataCorp LP (2015) along with examples.) This inconsistency is therefore noted but interpreted with caution.

none of the higher degrees in the invariance hierarchy, the parameters of those more restrictive models are in fact almost identical to those of the metric invariance model. In this case, the lack of stronger invariance is much less concerning than if the more restrictive model parameters differed substantially, and therefore comparisons over time of both the strength and scale of associations between latent and observed variables can be made relatively confidently. Carrying this confirmatory model forward as the measurement part of the SEM in Section 1.4, metric invariance is therefore retained in the specification of the respective measurement model.

1.4 SEM of objective and subjective well-being

Having identified how the two measurable components of SWB – emotional well-being and life satisfaction – are manifest in the response items of the BHPS, let us now investigate the associations of this SWB construct with indicators of several commonly targeted objective policy outcomes. The outcomes of interest are income, education and health. A multiple indicators multiple causes (MIMIC) model is fitted, with indicators of the objective outcome variables treated as exogenous to the model. That is, the model does not specify how income, education and health are generated, only that they exert some degree of directional influence on the components of SWB.

The reasons for interest in the relationship between SWB and these policy outcomes are two-fold:

1. As the prospect of SWB as a targeted measure gains prominence on the policy stage, it would be beneficial to understand how it relates to changes in outcomes in more established policy areas.

2. The relative strength and directions of these associations may shed light on respondents' underlying priorities among these objective outcomes. More specifically, outcomes with larger positive associations with the components of SWB may indicate higher priority outcomes than outcomes with lower or insignificant associations.

These three chosen objective policy outcomes appear to be widely accepted in both theory and in practice, and are frequently used in multidimensional applications to evaluating well-being. The UNDP Human Development Index perhaps set the precedent for incorporating education and health into comparisons of living standards (Stanton, 2007), recognising that these areas cannot be fully captured by measures of solely economic resources such as income. Inspired by this, a large body of work has grown out of this framework for measuring multidimensional well-being outcomes Yang (2014), and through these choices of data and corresponding goals, "the international community has identified a strong and shared view on the key dimensions of human well-being" (Dietz et al., 2007, p. 32). Following in this vein, focusing on the chosen outcomes helps to frame the analysis consistently with this literature. Section 1.4.1 details the outcome indicators used to expand the analysis of the previous BHPS items.

1.4.1 Exogenous objective well-being indicators

Income

Equivalised household income is used, calculated by taking the BHPS variable for total annual household income from all sources, and dividing by the square root of number of household members. This adjusts household income to account for economies of scale as resources are spread among additional household members, and makes household income comparable at the individual level. Although other more complex equivalisation scales are available, the derivation of all such scales depends on the assumptions made and judgements about needs, which are open to debate. The square root is the preferred equivalence scale of researchers for the Luxembourg Income Study, Eurostat, and more recently the OECD and many other individual countries (Chanfreau and Burchardt, 2008; OECD, 2013a), and is used here due to its popularity and wide comparability. Including equivalised incomes >£120,000 produced significant outliers, whilst including equivalised incomes <£100 produced bunching of values below this value, which is likely the result of reporting errors since this figure should include benefit payments, for which <£100 per person per annum is implausibly low. Therefore, individuals with equivalised incomes >£120,000 and <£100 are excluded from the analysis. The remaining equivalised household incomes are normalised to a $[0, 1]$ unit scale using the following commonly used min-max goalpost approach (Lugo, 2005; UNDP, 2013), where x_{it} is the original equivalised income value, $x_{\min}$ is the minimum equivalised income

£100, $x_{\max}$ is the maximum equivalised income £120,000, and $\hat{x}_{it}$ is the normalised value:

$$\hat{x}_{it} = \frac{x_{it} - x_{\min}}{x_{\max} - x_{\min}} \tag{1.1}$$

Education

Although education has appeared in many lists of basic well-being dimensions and on many policy agendas, the effect of educational attainment on SWB has been a subject of contention (Dolan et al., 2008; Michalos, 2008). MacKerron (2012, p. 721) concluded in his survey of the "happiness economics" literature that "the impact of education varies between studies: in some it has no significant effect, whereas in others highest [SWB] is variously associated with lower, higher, and intermediate levels of education." Although good education is often upheld as decisive in life, it seems that "empirical evidence remains quite divided and ambiguous when it comes to answers about... what people value in education." (Gibbons et al., 2009, p. 1). This is echoed in the findings of Decancq et al. (2015), who observe an insignificant relationship using Russian data between educational attainment and the life satisfaction component of SWB. Nevertheless, the inclusion of this association in the analysis will be of particular interest. For the education dimension, the categorical variable of highest education qualification is selected. Similarly to the income values, the education variable is normalised to a $[0, 1]$ unit scale with ordinal levels.

Health

For health, a composite indicator is derived using BHPS variables for the following health indicators: whether an individual has been a hospital inpatient in the last year, whether an individual has problems with limbs, with chest or breathing, with heart or blood pressure, with stomach or digestion, with diabetes, with migraines, and with anxiety or depression. The composite measure is derived based on predicted values of the linear index from an ordered logit model of subjective health satisfaction. The weights in the composite measure are the rescaled coefficients of the logit regression, with the rescaling used to normalise the composite measure between 0 and 1. Subjective health satisfaction is not used as a direct measure of health since this risks endogeneity with the SWB items, as discussed by Ferrer-i Carbonell and Frijters (2004).

1.4.2 SEM results

The measurement part of the model is pre-specified according to metric invariance constraints, following from the findings in Section 1.3.3. That is, the factor loadings are constrained to be equal across waves. The latent factor variances for 'emotional well-being' and 'life satisfaction' are scaled to 1. A similar invariance testing process is then carried out for the structural part of the model, for which the results and fit statistics are presented in Tables 1.13 and 1.14. Factor loadings for the measurement part of the model are also presented in Table 1.13, and these generally corroborate in magnitude with the earlier two-factor confirmatory model, although there are some adjustments due to the addition of the structural part. Note that *obj income* in the structural part of the table (upper half) refers to the objective equivalised income variable, whereas *subj income* in the measurement part (lower half) refers to the subjective income satisfaction item. Similarly, *obj health* refers to the objective health variable and *subj health* refers to the subjective health satisfaction item. Inspecting the fit statistics in the same fashion as Section 1.3.3, the structural model fails to satisfy metric invariance. Therefore, the model satisfies only the weakest form of invariance, for which the parameters of wave 6 are reported in Table 1.13.

Fit statistic	Degree of invariance	
	Free	Metric
CFI	0.836	0.835
χ^2	11260.60 (159)	11325.16 (171)
$\Delta\chi^2$		64.56 (12)
AIC	377521.806	377562.362
BIC	378205.153	378148.087

Table 1.14: Fit statistics for the complete structural model

This weaker degree of invariance is perhaps not surprising, since it would be expected that the structural model relationship between SWB and *health*, *education* and *income* is different in nature from the measurement model relationship between SWB and its manifest variables as analysed in Section 1.3. The hope is to find invariance

Latent factor	Coefficients	Free to vary (wave 6)	Metric invariance
Emotional well-being	*Obj income*	0.422* (0.174)	0.224* (0.076)
	Education	-0.053 (0.059)	-0.054 (0.033)
	Obj health	1.201* (0.048)	1.235* (0.027)
Life satisfaction	*Obj income*	2.557* (0.200)	1.569* (0.085)
	Education	-0.552* (0.066)	-0.530* (0.037)
	Obj health	1.953* (0.060)	1.808* (0.036)
Emotional	*Useful*	0.280 (0.004)	0.280 (0.004)
	Relaxed	0.452 (0.005)	0.452 (0.005)
	Enjoy	0.387 (0.003)	0.387 (0.003)
	Confidence	0.474 (0.005)	0.474 (0.005)
	Happy	0.394 (0.003)	0.394 (0.003)
Life	*Subj health*	0.467 (0.004)	0.468 (0.004)
	Subj income	0.409 (0.005)	0.408 (0.005)
	House	0.281 (0.005)	0.281 (0.005)
	Social	0.395 (0.005)	0.396 (0.005)

Table 1.13: Coefficients of the complete model under metric measurement invariance and different degrees of structural invariance (standard errors in parentheses)

in the latter since a consistent measurement of the latent SWB construct is desired over cross-sections of the data, and indeed this is the case with the finding of metric invariance. On the other hand, there are no claims that the cross-sectional relationship between the components of SWB and objective life outcomes should be the same in each wave. This structural relationship might be expected to be much more imperfect than the measurement relationship, and indeed the CFI value below 0.9 shows that overall model fit does weaken upon adding the structural part to the model, compared to the measurement model alone. A conclusion to draw from this is that invariance testing may not be the most appropriate tool to address the longitudinal aspect of the structural model. This is further pursued in Chapter 2, where a different perspective is taken to analysing this relationship across waves using non-linear fixed effects methods.

Figure 1.3 presents the information from the first results column of Table 1.13 in a path digram, for the complete wave 6 model under metric invariance for the meas-

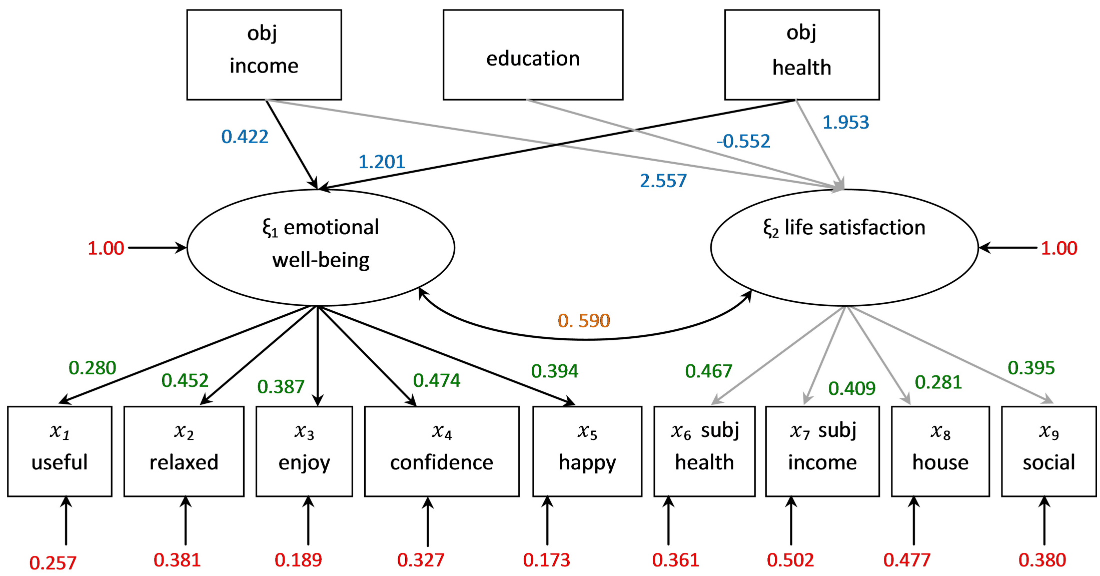

Figure 1.3: MIMIC path diagram of endogenous latent SWB components and objective well-being indicators

urement part and configural invariance for the structural part of the model. Figure 1.3 also displays additional model parameters, including correlation between the latent SWB factors, their variances (standardised to 1), and residual variances of the observed manifest variables in the measurement part of the model. Although the accepted parameters are for the configural invariance model, the directions and relative magnitudes of coefficients and loadings under the metric invariance model in the last column of Table 1.13 are comparable. The structural coefficients for *income* and *health* are positive and significant for both 'emotional well-being' and 'life satisfaction'. Recognising that individual heterogeneities are not controlled for and the model is therefore a descriptive one, this indicates that net relationships between greater financial resources and better health on the one hand, and higher emotional and evaluative SWB on the other, are positive, as might be expected. In terms of magnitude, the absolute value of coefficients of all objective indicators are larger for 'life satisfaction' than for 'emotional well-being'. The 'life satisfaction' component of SWB therefore appears more strongly related than the 'emotional well-being' component to attainments in the objective indicators. Intuitively, whilst endowments in finance and health facilitate the conditions and tools necessary to pursue a satisfying life, they cannot so easily guarantee hedonic utility and emotional well-being in the day-to-day sense.

Comparing this to broader conclusions in the literature, this echoes the influential finding of Kahneman and Deaton (2010) that "high income improves evaluation of life but not emotional well-being", and similarly the findings of Brief et al. (1993, Table 1) of stronger correlations between objective health and life satisfaction than between objective health and positive and negative affect.[4] Smith et al. (2002) also examine the predictive role of objective and subjective health in accounting for the variance in components of SWB, although they study only those aged 70-100 and investigate different components of SWB than those investigated here. They find that objective and subjective health taken together accounted for 32% of the variance in ageing satisfaction, 20% in life satisfaction, 18% in depressivity, 14% in negative affect, and 13% in positive affect. Deaton (2008) on the other hand finds that "neither life satisfaction nor health satisfaction responds strongly to objective measures of health", using data from the Gallup World Poll. However, Deaton investigates only coarse measures of objective health such as life expectancy and the prevalence of HIV infection, which arguably mask some pertinent health considerations.

Interestingly the coefficient for education is negative for both factors, although statistically significant only in terms of effect on 'life satisfaction' and insignificant for 'emotional well-being'. That is, the net relationship between higher levels of education and general evaluation of life is negative, and there is no association with day-to-day happiness. This runs counter-intuitive to the expectation, or at least hope, that education has a positive impact on SWB and the common policy emphasis on expanding access to education. An explanation could be due to the effect of education on raising life expectations, and therefore introducing interpersonal variation in the interpretation of the life satisfaction scale. A negative net affect could be observed when measured without taking this into account. This issue is further pursued in Chapter 2, where individual fixed effects and an alternative strategy for identifying the effect of education is taken. As discussed in Section 1.4.1 however, the counter-intuitive education result is in fact a common finding in the SWB literature. Michalos (2008) notes that if education is defined, as it is here, as highest level of formal education attained, then the overwhelming evidence, including a review of 90 American studies (Witter et al., 1984) and more recently the findings of Layard (2005), is that "education has very little influence on happiness". However, he suggests that if the definition of education is

[4] The explicit relationships modelled by these authors further on in their paper are, however, not comparable to this study.

expanded to include also non-formal and broader aspects of education, then it does have an influence on the eudaimonic component of subjective well-being. However, given the difficulty in empirically implementing the concept of eudaimonia as discussed in Section 1.1, it is difficult to do better than refer back to Michalos' list of education's cited positive influences on variables *indirectly related* to SWB, such as health, employment and crime outcomes. It is recognised that the negative coefficient for education on life satisfaction is likely due to confounding of the effects of prior education, that have manifested themselves indirectly later in the life course through positive income effects or negative aspiration effects. However, with the data available and the SEM modelling approach it is difficult to do better.[5]

A caveat should be made regarding the application here, and indeed most applications of factor analysis in the social sciences, in that linear models are used exclusively which attempt to explain relationships only through covariances and correlations (Yalcin and Amemiya, 2001). The possibility of non-linear relationships is therefore neglected due to the analytical limitations of non-linear factor analysis methods. In this context of latent components of SWB, whilst factor analysis has confirmed the idea of two underlying latent factors, it assumes that the form of the relationships to the observed items takes that of weighted sums of the underlying factors. It is worth noting that this assumption of a linear structure may not be the "true" structure, and it may be that a non-linear relationship exists between underlying factors and observations. The motivation for using standard linear factor analysis therefore does not lie in uncovering a particular model form, but more in distinguishing the basic monotone relationships between observed items through the underlying factors.

The model also allows for correlation between latent life satisfaction and emotional well-being, and the finding is that attainment in the two component parts of SWB are interlinked. However, theory and intuition do not provide a definite answer as to whether there is causality between the two, or in which direction(s) such causality may run. That is, it is not clear if emotionally content people have the capacity to lead more satisfying lives as a result of their emotional disposition, or if emotional well-being is the direct result of high satisfaction in the various domains of life. In reality, it is likely that there is a complex two-way relationship which is difficult to precisely identify.

[5] Refer to Chapter 2 for a different approach to the issue of education and SWB, made possible by the ability of other regression techniques to cope with binary education variables.

Consequently, the correlation is left as such in the model, without attempting to insert a directional path between the two latent SWB components.

To summarise concisely, the total effects in the structural part of the MIMIC model are given by the equations:

$$\text{`Emotional well-being'} = 0.422(\text{equivalised income}) + 1.201(\text{health})$$

$$\begin{aligned}\text{`Life satisfaction'} &= 2.557(\text{equivalised income}) + 1.953(\text{health}) \\ &\quad -0.552(\text{education})\end{aligned}$$

1.5 Conclusions

In this chapter, the structure of SWB and its latent components were investigated with multiple self-reported response items on SWB from the BHPS. Using factor analysis, a simplified form of the structure proposed by Lucas et al. (1996) was found, consisting of a latent 'emotional well-being' component and 'life satisfaction' component. The longitudinal invariance of this latent SWB structure was then tested, to check that the same concept was indeed being measured over time and that it was consistently interpreted by individuals. Metric invariance for the SWB structure was found, indicating that the latent components manifest themselves quantitatively through the same unit changes in item scores over time. This is an important finding that supports the use of longitudinal SWB regression in Chapter 2 to uncover individual and group-level preferences over objective well-being dimensions.

The relationships between the two components of SWB (emotional well-being and life satisfaction) and objective dimensions of well-being (income, education and health) were also investigated in this chapter at the population level, in particular to see whether these relationships differed in nature from one component to the other. This was achieved by extending the factor analysis of SWB into a structural MIMIC model to include policy indicators of objective well-being. The findings were that while the 'life satisfaction' component was influenced by all three exogenous well-being variables, income, education, and health, the 'emotional well-being' component was not influenced by education, and to a lesser extent by the other two variables. Practically, the implication is that the 'life satisfaction' component of SWB seems to be more sensitive than the 'emotional well-being' component to changes or interventions in attainment in the objective indicators examined. Given this result and the result that emotional well-being was positively correlated with life satisfaction, it can be argued that the life satisfaction component of SWB is better for evaluating objective aspects of social performance such as the indicators examined, rather than the hedonic 'emotional well-being' component. Indeed, a recent amendment to *The Green Book* of the UK Treasury, the official guidance to other government agencies on evaluating policy proposals, stated that the 'life satisfaction approach' "will be important in ensuring that the full range of impacts of proposed policies are considered, and may provide added information about the relative value of non-market goods compared with each other" (HM Treasury, 2011, p. 58). Looking ahead to Chapter 2's approach of using SWB information to retrieve preferences for a multidimensional index, the greater policy relevance of life satisfaction as a yardstick for objective attainments justifies the use of 'life satisfaction' as the dependent variable in the SWB regressions.

More research is needed, for example, on whether different indicators of objective dimensions of well-being produce differing results. In particular education provides an interesting case for exploring alternative approaches. Scaling down to the intra-country level, introducing interaction effects would also make it possible to compare the well-being of different types of individuals, and whether the composition of multidimensional well-being and the implied trade-offs between dimensions differs between such individuals. Some of these extensions are explored in Chapter 2, although of course these are just a few suggestions from a plethora of potential research avenues.

Chapter 2

Towards an index of multidimensional well-being with heterogeneous preferences

2.1 Introduction

The question of how to define multidimensional well-being has been gaining prominence in research and policy agendas in recent years. GDP centric growth policies that served post-war economies relatively well seven decades ago no longer point in a direction that captures the concerns of modern society, and the broad consensus now is that a multidimensional measure of progress is needed. Indeed a plethora of just such measures has been proposed in recent years; a good overview of these developments can be found in Fleurbaey and Blanchet (2013) and Aaberge and Brandolini (2014). As Fleurbaey and Blanchet note however, this literature has been overwhelmingly concerned with aggregating average population-level indicators or distributions, and has been largely silent on how such measures can be made to capture the heterogeneous preferences, or differing 'recipes', of individuals for a good life. This is a troubling omission. After all, another pertinent question in the quest for redefining progress is: whose conception of progress should be measured? The view taken here is that in theory and as far as applied work allows, real and likely differing preferences of individuals in the population should be accounted for – not those of an arbitrary, 'representative' agent. From this perspective, there should not be a one-size-fits-all measure of progress, and in order for policy-makers to find the best way of stimulating progress and promoting well-being, we must first identify what counts towards the well-being of different individuals.

In current practice in the measurement of well-being and human progress, two types of simplification are often implicitly made. First, it is generally taken as given that there exists a readily available "measure of well-being which is capable of being expressed on a cardinal ratio scale" and that "the individual well-being functions are identical across individuals" (Dolan and Tsuchiya, 2009). This is a common assumption used in applied welfare and public economics. The second simplification is that not only is well-being defined identically across individuals, but in seeking an aggregate assessment over the entire population, a "representative agent" approach can be taken by taking the simple average of well-being levels across individuals. This issue is investigated in the theoretical literature on inequality measurement, but is then largely overlooked in the business of developing GDP-alternatives.

In this chapter a modified notion of multidimensional well-being indices is proposed, combining the intuitive idea of an index as a type of summary statistic with recent literature on axiomatic approaches to multidimensional well-being in welfare economic and social choice theory. Concrete proposals are made for unpacking and operationalising the processes obscured by the two simplifications highlighted above, putting individuals' own preferences centre stage. The finding is that such an index is on the one hand not so far removed from many of the GDP-alternatives that have failed to gain support from theorists; on the other hand it can in fact be seen as an application of a theoretical approach that has been rigorously defended. Namely, the proposed index can be seen as an application of the equivalence approach, developed by Pazner and Schmeidler (1978), and as a modified notion of the equivalent income approach (Fleurbaey, 2005, 2011; Decancq et al., 2015), itself an application of the equivalence approach. For the sake of clarity, the approach developed in this chapter will be referred to as the 'preference index approach', and is the first proposal and empirical application of a preference-sensitive well-being and inequality measure in the form of a multidimensional index.

The original contribution of equivalent income is in its use of individual-specific preferences to define well-being, an approach which had previously only been considered in the context of consumption-based "money metric utility" (Samuelson, 1974; Deaton, 1980), and in a medical context for the measurement of health status.[1] The preference index approach borrows from the contribution of equivalent income in that respect. However, contrary to the equivalence approach and equivalent income, it dispenses with a money-metric representation of well-being. Rather, it takes as its reference of variation the dimension-neutral unit space. At the same time, it bridges a gap in the theory between popular (but theoretically weak) synthetic indices of well-being and several strands of the economic literature on welfare measurement, *viz.* the capabilities approach, subjective well-being (SWB), equivalent income and multidimensional inequality measurement. In this sense, it is an attempt at addressing "the frequent gap between foundations and applied measures, between concepts and statistics" (Fleurbaey, 2008, p. 1), as well as gaps between literatures that have often evolved rather separately. It will be seen that the preference index approach results in a two-step multidimensional index akin to the specification proposed by Maasoumi (1986). The crucial difference is that an explicit derivation is given for the aggregation rule from an axiomatic point of view, with the help of the equivalent income framework, and that an empirical strategy is presented for how to measure heterogeneous parameter values for different individuals.

It should be noted at the outset that the departure from a monetary measure is not due to abhorrence for monetising aspects of life. That is to say, no objection is made to using money as a numeraire for making relative comparisons between dimensions of well-being. It has been noted that "[the] situation is not fundamentally different when none of the aggregated variables are monetary. Aggregation always implies assuming some more or less important substitution possibilities between the items that are aggregated" (Fleurbaey and Blanchet, 2013, p. 14). The rationale is rather that it results in a measure of well-being that is not contingent on a reference level in each dimension from which comparisons are made. This is a separate issue from the numeraire issue, and does have a bearing on making interpersonal comparisons if individuals have different preferences (*i.e.* their indifference curves cross). Additionally, in value systems where income has no value in well-being, none of the things that do have value can be measured in terms of trade-offs with income. For example, consider the notion of

[1]Specifically, the reference here is to the Health Utilities Index (HUI) developed by McMaster University. The Health Utilities Index elicits preferences about various health states from a representative sample of individuals within a community, and as such captures the views of society concerning health status. (Horsman et al., 2003)

pure "Buddhist" preferences under which income plays no role in the individual's definition of a good life – any attempt at implementing a monetary measure of well-being then breaks down since money is unable to capture any trade-offs between dimensions of life. Therefore, abandoning money as a well-being metric enables a wider array of value systems to be accommodated.

The essential purpose of the preference index approach is to use ordinal and non-comparable information about individual preferences to construct an interpersonally comparable index of well-being. In brief, the procedures in the proposed method can be broken down in four stages:

1. Mapping out individuals' indifference curves, each of which represents her preferences over dimensions of well-being.

2. Projecting any given individual's actual bundle of multidimensional well-being attainment onto an equivalent bundle along the individual-specific indifference curve, so that equivalent bundles of all individuals lie along a reference path in the multidimensional space.

3. Assigning individual index values to bundles on the reference path (analogous to the inverse of the distance function put forward by Deaton (1979)). Evaluation of bundles along the path provides interpersonal comparability between individuals with different preferences and different levels of attainment in dimensions of well-being.

4. Finding a suitable population index satisfying a certain set of axioms, to aggregate individual index values into an overall assessment of the multidimensional distribution of well-being.

The first step is an empirical question of eliciting preferences, whereas the other steps involve normative judgements in some form. Since the equivalence approach underpins the rationale for the preference index method as a whole, an introduction to the equivalence approach and how it relates to the proposed method is given in Section 2.2. The last axiomatic step draws heavily from the framework laid out in Fleurbaey and Maniquet (2011) and Decancq et al. (2014a). This is given in Section 2.3. The key theoretical argument is that a synthetic index of well-being can be composed in a way that is consistent with welfare economic theory, incorporating considerations of heterogeneous interpersonal preferences, fairness in evaluation, and inequality in distribution. Section 2.4 provides an empirical illustration of steps 1, 2, and 3, and the rich analysis possibilities it allows. One such analysis is a comparison of the preferences of different individuals and how the incorporation of these preferences paints a different picture of well-being from other welfare measures such as income and SWB. Section 3.5 concludes.

Regarding the extent of the contribution of this chapter, the axiomatics borrow from the existing literature on multidimensional inequality and poverty measurement, and in that respect is not original. However, as already discussed, the standard practice in this literature has been to make the very restrictive and often implicit assumption of a common individual well-being function, *i.e.*, to assume identical preferences across all individuals. Exceptionally, Decancq et al. (2014a) apply modified axioms to derive a class of preference-sensitive multidimensional poverty indices, and this chapter draws heavily from their work. Decancq and Neumann (2016) also include a similar approach to deriving individual-level well-being as an intermediate stage in the so-called "extended preference approach" (Adler, 2014), which stops short of considering inequality. In the field of money metric utility, consumer preferences have long been used in the calculation of consumption-based welfare and poverty measures (Samuelson, 1974; Deaton, 2008). However, the consideration of preference differences in a non-money metric composite index of multidimensional well-being and inequality, as in this thesis, is a new endeavour.

2.2 The equivalence approach

This section provides an overview of the equivalence approach and its conceptual role in the definition of the proposed preference-sensitive index. In its basic form in the context of a two-agent two-good exchange economy model, the equivalence approach defines fair allocations to be all the Pareto efficient allocations such that each individual is indifferent between her actual bundle and an egalitarian distribution of goods, *i.e.* each individual getting an identical bundle that is some fraction of the social endowment.[2]

[2] Note that such an egalitarian distribution need not be feasible, *i.e.* the fractions of the social endowment can sum to greater than one, since this distribution is only a hypothetical one to which

Pazner and Schmeidler (1978) coin the term Pareto-efficient and egalitarian-equivalent allocations (PEEEAs) to describe this resulting set of allocations. The egalitarian distribution can be seen as a fair hypothetical world to which actual distributions of bundles are compared. In this way PEEEAs are identified which are equally as good as this hypothetical egalitarian distribution, but in which individuals may not necessarily have identical bundles due to differing marginal rates of substitution between goods from individual to individual.

Following a modified line of reasoning in relation to the equivalence approach, the preference index approach proposed here amounts to comparing individuals in a hypothetical world in which they are just as satisfied as in their actual situation, but in which their bundle of attainments are all situated somewhere along a defined reference path. In this approach the path takes the form of a ray extending from the vector of minimum attainable values at the origin to the vector of maximum attainable values in each dimension. Figure 2.1 provides a two-dimensional representation, where $0B$ is the ray as defined. This ray represents fractions of the maximum attainment bundle, and is analogous to the fractions of the social endowment in the original equivalence approach. The key difference is that the original equivalence approach prescribes identical fractions across individuals in order to identify PEEEAs, whereas the purpose of the preference index approach here is to use differences in these fractions across individuals to compare their well-being. In other words, whereas the original approach looks for egalitarian distributions of bundles in the hypothetical world, the preference index approach uses the hypothetical world as a tool for comparing different bundles under different preferences.

This hypothetical world is constructed so that all individuals are indifferent between their actual situations and their equivalent situations on the ray in the hypothetical world, by moving along indifference curves. In this hypothetical world, bundles can be compared in terms of their distance from the origin along the ray, where there is no ambiguity in evaluation since attainments in all (normalised) dimensions are equalised across dimensions for a given individual. This "equally distributed equivalent" is then used as the individual well-being index measure. In Figure 2.1, $0B$ is the ray as defined, 0 is the origin defining the minimum attainment bundle and B is the maximum

bundles are being compared.

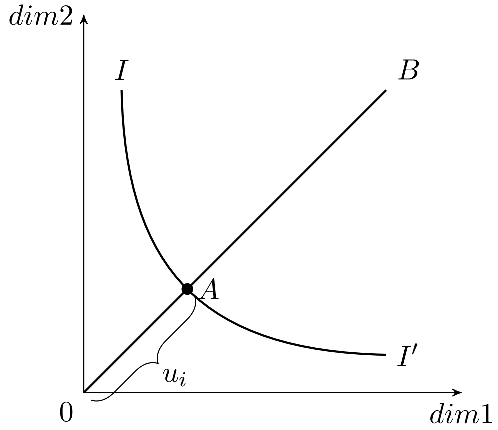

Figure 2.1: Simple illustration of the equivalence approach

attainment bundle. The distance $0A$, labelled u_i, is individual i's index value for any bundle on indifference curve II', which is equivalent to the hypothetical bundle A. Pazner and Schmeidler (1978) prove that the concept carries over to the case of multiple dimensions.

As noted, a main difference between the original equivalence approach and the preference index approach is that the objective of the former is a fair allocation rule, whereas that of the latter is a rule for making comparisons in potentially unfair allocations. Another difference is that since this application of the equivalence approach is to well-being attainment rather than goods allocation in a closed economy, the element of rivalry disappears. This is because the resources that enable well-being attainment are not explicitly modelled as being scarce, although it should be recognised that this is not entirely realistic. In the preference index model there is no longer a social endowment of goods to be divided among agents; instead there is a maximum level of attainment in each dimension of well-being, and it is assumed that attainment of one individual is not constrained by the number of other individuals or the attainment levels of those other individuals.

Comparing the preference index approach to the equivalent income approach developed by Fleurbaey in a series of papers (Fleurbaey, 2005, 2011; Fleurbaey and Maniquet, 2011), this one is based on a different extension to the original equivalence approach. The equivalent income approach also looks to evaluate and compare different bundles such that each individual is indifferent between the actual bundle and a bundle on a reference path. However, whereas the preference index approach introduced here compares situations to fractions of the maximum attainment bundle as the reference path, the equivalent income approach requires reference values of each non-income dimension to be defined, and optionally different values for each individual. In terms of Figure 2.1, whereas the diagonal ray $0B$ from the origin is the relevant multidimensional reference to which individual situations are compared and evaluated, in order to evaluate situations using the equivalent income approach a reference level must be fixed for each non-income dimension. This is visualised as a horizontal (or vertical) path extending perpendicularly from the chosen reference value on each non-income axis. An individual's well-being is then found by computing income minus the cumulative amount of income he or she would be willing to give up to attain the reference value of each other dimension (Decancq et al. (2015), for example, choose the reference value of perfect health). The resulting equivalent income measure is therefore an adjusted-income measure. The measure proposed here, on the other hand, is a composite index measure, which is not reliant on a monetary dimension such as income in order to be defined.

The argument behind using dimension-wise reference levels under the equivalent income approach is that, for each non-income dimension and each individual, there exists some optimal level of attainment for which it is ethically defensible to compare the situations of individuals, irrespective of their preferences. Again taking the example of perfect health as the reference level for the health dimension as in Decancq et al. (2014a), the argument would be that trade-offs between health and other dimensions are irrelevant when individuals are in perfect health. Equivalent income is then defined as each individual's actual income adjusted down for the loss in well-being associated with a less-than-perfect level of health. It can be interpreted as the individual's actual income minus her willingness-to-pay (or to give up income) to reach perfect health from her actual health status. The same applies for reaching the optimal reference levels in the other well-being dimensions. Intuitively this can be understood as having a hypothetical multidimensional baseline situation of optimal attainments that individuals reach by giving up income, whereby the equivalent income is defined as the remaining income upon reaching this hypothetical "optimal" situation.

The rationale for the preference index approach makes two counter-arguments against this equivalent income approach. First, if our objective is to obtain a well-being measure that is generaliseable to all societies, then preferences should be considered that place no weight on income in an individual's definition of well-being. Under such pure "Buddhist" preferences, any attempt at implementing a monetary measure of well-being then breaks down, since income is unable to capture the trade-offs between dimensions of life. Therefore, abandoning money as a well-being metric allows an approach to well-being measurement that is not reliant on a particular dimension (namely income) being present in the definition of well-being.

Second, it is not clear that taking perfect health, or a chosen "optimal" reference value in each non-income dimension for each individual, is the most convincing hypothetical baseline situation. This implies the extreme view that income as the slack variable should be reduced to the full extent of attaining this situation of optimality in the other dimensions. The preference index approach instead defines the hypothetical baseline situation to be where an individual's attainments are equally distributed across all the well-being dimensions. The corresponding willingness-to-pay interpretation is that the preference index is characterised by an individual's willingness to sacrifice attainment in higher-attaining dimensions in order to raise attainments in the lower-attaining dimensions to a hypothetical situation of equally distributed attainment across all dimensions. The ethical argument is that a balanced bundle of attainments across the dimensions is an unambiguous position from which to evaluate an individual's well-being, since her measure of well-being is simply the common degree of attainment across dimensions. This does not depend on income being present in the well-being definition. It also does not imply an extreme hypothetical situation for making comparisons in which one dimension, income, is used to compensate for all deviations below an optimal level in each of the other dimensions.

2.3 Axioms

Now that a conceptual basis has been laid out for evaluating individual indexes, the aggregate index can be specified in relation to the properties of these individual indexes. These axioms set out the properties the proposed preference-sensitive well-being index should satisfy, and are largely ethical in nature. The axioms that are put forward for the multidimensional index have unidimensional counterparts (see, for example, Chakravarty (1990)) and counterparts in the multidimensional context without preference sensitivity. Starting with modified versions of Pareto efficiency and Separation (subgroup consistency), the modifications, as well as some additional principles, allow for the consideration of interpersonal comparisons between heterogeneous preferences and for inequality aversion.

2.3.1 Theoretical framework

Consider a simple framework consisting of a population denoted by the non-empty and finite set $N \subset \mathbb{N}_{++}$ of individuals, for whom there are m relevant dimensions of well-being contained in the set M. Let $\mathcal{N}$ denote the set of non-empty finite sub-sets of $\mathbb{N}_{++}$. Each individual $i \in N$ has well-defined preferences over personal attainment bundles x belonging to the individual's potential attainment set $X_i \subseteq \mathbb{R}^m_+$. It is conceivable that X may vary from person to person, due for example to genetic reasons, but in this analysis individuals are treated identically in this respect with the same potential attainment set X.

Let R_i denote individual i's complete preference ordering over the set X. Preferences are assumed to be complete and transitive. When i prefers bundle x_i as least as much as bundle x_i', this is denoted by $x_i R_i x_i'$. Strict preference is denoted by $x_i P_i x_i'$ and indifference by $x_i I_i x_i'$. Let $\mathcal{R}$ denote the set of preferences over X that are continuous, monotonic and convex. By continuity, it is meant that for all $x_i \in X$, the sets $\{x_i' \in X | x_i R_i x_i'\}$ and $\{x_i' \in X | x_i' R_i x_i\}$ are closed. By monotonicity, it is meant that greater attainment in a given dimension of well-being entails strict preference (*i.e.* for two bundles $x_i, x_i' \in X$, if $x_i \geq x_i'$, then $x_i R_i x_i'$). By convexity, it is meant that for two bundles $x_i, x_i' \in X$, if $x_i R_i x_i'$ then $(\delta x_i + (1-\delta)x_i') R_i x_i'$ for all $\delta \in [0, 1]$. Note that here, x is assumed to be cardinal. In practice, some empirical literature has treated ordinal data as cardinal for tractability and model flexibility (see, for example, Allan (1976) and Labovitz (1970), and Harwell and Gatti (2001) for a discussion in the context of educational data), or because this can provide additional insights into useful relationships (Moses et al., 1984). In the later empirical application that follows, the education dimension will be treated in this manner.

Given a set N of individuals, a distribution refers to an N-list of attainment bundles $x_N = (x_i)_{i \in N} \in X^N$, where x_i refers to individual i's bundle for all $i \in N$. An index of well-being is a function $W : S = \bigcup_{N \subset \mathcal{N}} X^N \times \mathcal{R}^N \to \mathbb{R}_+$, such that $W(x_N, R_N)$ gives the level of well-being in distribution x_N when preferences in the population are R_N. (x_N, R_N) can thus be thought of as a social state. Given this framework, a definition of well-being can be specified that is consistent with the preferences of individuals in the population by satisfying the following axioms.

2.3.2 Axioms

Population well-being index comparisons

One of the contributions in this chapter is the use of individual preferences to evaluate changes in well-being. This adds to a recent original stream of literature on this subject (Fleurbaey, 2005, 2011; Fleurbaey and Maniquet, 2011; Fleurbaey and Schokkaert, 2012; Fleurbaey and Blanchet, 2013; Decancq et al., 2014a, 2015). In this context, the standard Pareto principle is modified to require that if the preference satisfaction of all individuals weakly increases then overall well-being must weakly increase, and if the preference satisfaction of at least one individual strictly increases then well-being must strictly increase. This is captured in the following axiom:

Axiom 1. Pareto

For all (x_N, R_N), $(x'_N, R_N) \in S$, if for all $i \in N$ we have that $x'_i R_i x_i$, then
$W(x'_N, R_N) \geq W(x_N, R_N)$.
Additionally, if there is at least one $j \in N$ such that $x_j P_j x'_j$, then
$W(x_N, R_N) > W(x'_N, R_N)$.

For a population partitioned into two sub-groups N and M of fixed sizes, it is also required that if sub-group N has the same bundle in two distributions, the ranking of these two distributions should remain the same if this sub-group were excluded from the evaluation. This ensures that a change in well-being within a subpopulation does not depend on the rest of the population that is not affected:

Axiom 2. Separation

For all (x_N, R_N), (y_M, R_M), $(y'_M, R'_M) \in S$, $W(y_M, R_M) \geq W(y'_M, R'_M)$ if and only if
$W((x_N, y_M), (R_N, R_M)) \geq W((x_N, y'_M), (R_N, R'_M))$.

Additionally, it is required that the index of well-being is continuous with respect to individuals' attainment bundles:

Axiom 3. CONTINUITY
For all $N \in \mathcal{N}$, $i \in N$, $W(x_N, R_N)$ is continuous in x_i.

A population distribution and its k-fold replication should be evaluated as having the same level of well-being, where k is any positive integer and each replica preserves the same individual characteristics (x_i, R_i) as the original. The implication of this is that population size does not matter for such an evaluation. This is a restrictive but widely used property in inequality and welfare measurement, though a significant body of literature has developed around reasons and ways for relaxing this axiom (see Blackorby et al., 2005). For example, a setting in which individuals attain the same level in each dimension of well-being but there are more individuals in total may be considered worse, for instance if the larger population makes the same attainments less valuable in terms of being able to live a good life. Recognising that there is more work to be done on extending this approach to model such contexts, the more simplistic replication invariance axiom is retained here:

Axiom 4. REPLICATION INVARIANCE
For all $(x_N, R_N) \in S, W(x_N, R_N) = W(x_{kN}, R_{kN})$ where (x_{kN}, R_{kN}) is the k-fold replication of (x_N, R_N).

Scale invariance is imposed, or equivalently homotheticity of $W(x_N, R_N)$, referring to invariance of the shape of the contour maps at different levels of well-being. It ensures that the well-being ranking of two population distributions is unaffected if the dimension attainments of each individual are rescaled by the same factor in both populations:

Axiom 5. SCALE INVARIANCE
For all $(x_N, R_N), (x'_N, R_N) \in S$ and $\lambda > 0$, if $W(x'_N, R_N) \geq W(x_N, R_N)$ then $W(\lambda x'_N, R_N) \geq W(\lambda x_N, R_N)$ where $(\lambda x_N, R_N)$ is the social state where individual attainments are λ-fold rescalings of attainments in (x_N, R_N).

Heterogeneous individual well-being index comparisons

Fleurbaey and Maniquet (2011, Theorem 2.1) prove that an incompatibility arises in the multidimensional context between the Pareto principle and Pigou-Dalton Transfer principle. To overcome this incompatibility a weakened form of Transfer axiom is necessary, allowing the set of bundles among which transfers increase overall well-being to be a subset of X. The relatively weak form of the Transfer axiom here is a result of that subset restriction.

Axiom 6. TRANSFER

There exists a convex subset $T \subset X$ *such that for indifference curve* $I(x_i, R_i)$, $I(x_i, R_i) \cap T \neq \emptyset$ *for all* $x_i \in X$ *and all* $R_i \in R$. *For all* (x_N, R_N), $(x'_N, R_N) \in S$, *if for some* $j, k \in N$ *we have that*

$x_j, x_k, x'_j, x'_k \in T$,

$x_j \gg x_k$ *and* x'_j, x'_k *is obtained from* x_j, x_k *by a multidimensional Pigou-Dalton transfer, i.e.* $x'_j = (1-\delta)x_j + \delta x_k$ *and* $x'_k = (1-\delta)x_k + \delta x_j$ *for some* $\delta \in (0, \frac{1}{2})$,

$W(x'_j, R_j) \geq W(x'_k, R_k)$, *and*

for all $i \neq j, k : x_i = x'_i$

then it is an improvement:

$W(x'_N, R_N) \geq W(x_N, R_N)$.

Convexity of the subset T ensures that bundles resulting from a Pigou-Dalton transfer between bundles in T also belong to T, meaning further such transfers would still be desirable. Although the Transfer axiom is most often used in the context of income inequality, with a direct policy analogue in the tax-transfer system, one can generalise to other non-monetary dimensions of well-being by considering in-kind targeted provision and policy directed at access to good-quality health and education, for example.

T is further defined in terms of a normalised scale:

Axiom 7. NORMALISATION

T *contains the minimum and maximum attainment bundles in the potential attainment set* $X \subseteq \mathbb{R}^m_+$, *with the minimum normalised to* $W(x_{min}, R_i) = 0$ *and the maximum normalised to* $W(x_{\max}, R_i) = 1$ *for all* R_i.

2.3.3 A preference-sensitive multidimensional well-being index

The following multidimensional index specification is proposed, which satisfies the desired axioms in Section 2.3.2 and is to be interpreted as an application of the equivalence approach:

$$W(x_N, R_N) = \left(\frac{1}{|N|} \sum_{i=1}^{|N|} \left(\phi(x_i, R_i)\right)^{\alpha} \right)^{\frac{1}{\alpha}},$$

where ϕ is increasing, continuous and concave in x_i, and is the function which aggregates the multiple dimensions in bundle x_i according to preferences R_i into $\phi(x_i, R_i)$, the individual indexes of well-being.

The proof is closely analogous to Decancq et al. (2014a, Theorem 1) and is not reproduced here. $W(x_N, R_N)$ is additively separable in individual situations (x_i, R_i), and these can be ranked according to the equivalence approach using $\phi(x_i, R_i)$ to map onto the subset T, which must be a ray that connects the minimum and maximum attainment bundles. In other words, the subset T defines a reference path along which individuals situations can be compared, according to the equivalence approach. Figure 2.2 provides an illustration and is further discussed below. In addition, the results of Blackorby and Donaldson (1982, Theorem 2) mean that we arrive at a generalised mean specification for the aggregation of individuals and $W(x_N, R_N)$ is homothetic. Note that this means that this measure is also ordinally related to the unidimensional Generalised Entropy class of measures (Cowell, 1977; Cowell and Kuga, 1981).

At the level of the individual indexes $\phi(x_i, R_i)$, it is proposed to further restrict the Transfer axiom to the domain of homothetic preferences. The rationale is motivated by observing that the axiom does not take into account "leaky bucket" transfers among changing shapes of indifference curves, which may be relatively harmful to the transfer's donor whilst benefiting the receiver relatively little (Fleurbaey and Maniquet, 2011; Fleurbaey and Tadenuma, 2014). Homothetic preferences have scale invariant indifference contours and are therefore not susceptible to this problem.

Making this restriction to the domain of homothetic R_i (individual preferences), *i.e.* $x_i R_i x_i^{'} \Leftrightarrow \alpha x_i R_i \alpha x_i^{'}$, we have then that $\phi(\cdot)$ will be homogeneous and ordinally equivalent to

$$\begin{array}{ll} \left(\sum_{k=1}^{m} w_{ik} x_{ik}^{\rho} \right)^{\frac{1}{\rho}} & \rho \neq 0 \\ \prod_{k=1}^{m} x_{ik}^{w_{ik}} & \rho = 0 \end{array}$$

for dimensions $k = 1, ..., m$ and the generalised mean function is obtained for individual indexes. The generalised mean is a linearly homogeneous constant elasticity of substitution (CES) function. It is quasi-concave if and only if $\rho \leq 1$, with elasticity of substitution greater than one if and only if $\rho \geq 0$ and less than one if and only if $\rho \leq 0$, where elasticity $\sigma = \frac{1}{1-\rho}$. The constant elasticity assumption will later be relaxed to investigate what repercussions this has empirically. Coupling homothetic individual preferences with the axioms for the aggregate index, a double generalised mean function is obtained for the final overall index of well-being, given by the following specification:

$$W(\phi(x,R)) = \begin{cases} \left(\frac{1}{n}\sum_{i}^{n}(\phi(x_i,R_i))^{\alpha}\right)^{\frac{1}{\alpha}} & \alpha < 1, \alpha \neq 0 \\ \prod_{i}^{n}(\phi(x_i,R_i))^{\frac{1}{n}} & \alpha = 0. \end{cases} \tag{2.1}$$

where $\phi(x_i, R_i)$ are the individual indexes of well-being given by:

$$\phi(x_i,R_i) = \begin{cases} \left(\sum_{k=1}^{m} w_{ik}x_{ik}^{\rho}\right)^{\frac{1}{\rho}} & \rho < 1, \rho \neq 0 \\ \prod_{k=1}^{m} x_{ik}^{w_{ik}} & \rho = 0. \end{cases} \tag{2.2}$$

A graphical representation of the individual indexes is given in Figure 2.2. If the indifference curves of two individuals do not cross, then their situations can be compared unambiguously because an individual on a higher indifference curve will always be considered more satisfied than an individual on a lower indifference curve. The restricted application of the Transfer axiom allows for interpersonal comparisons between individuals with heterogeneous preferences, even if indifference curves of the different individuals cross. By the reference path T to which the Transfer axiom is restricted, it is deemed that $\phi(x_j, R_j) > \phi(x_i, R_i)$. By Normalisation, T goes through the minimum and maximum attainment bundles.

The rest of the chapter lays out the procedures for empirically deriving the individual indexes, with an illustration using data from the British Household Panel Survey. Empirical treatment of the aggregate index has been necessarily relegated to Chapter 3, where other substantial issues are duly addressed that cannot be covered here such as inequality aversion and measurement, multidimensional decomposition of well-being inequality in population subgroups and in well-being dimensions.

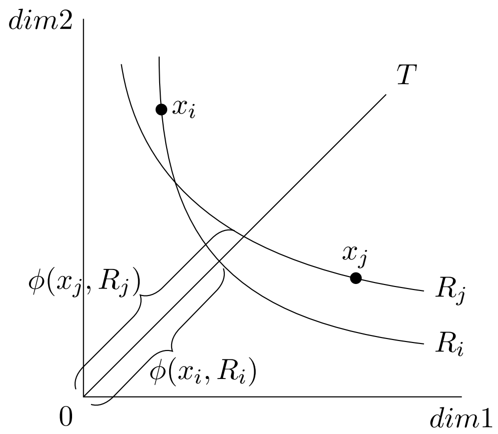

Figure 2.2: Restricted Transfer and inter-preference comparison

2.4 Multidimensional well-being in the UK

2.4.1 Methodological discussion

To empirically demonstrate the preference index approach outlined in the preceding sections, this section estimates preferences using a life satisfaction regression approach with micro level data from twelve waves of the British Household Panel Survey (BHPS). More generally however, several methods are theoretically possible for preference estimation:

1. Stated preference, such as choice experiments and utility models, as well as multi-criteria decision analysis which can be classified as an attribute-based stated preference method. Although these methods are widely used to elicit preferences on particular aspects of, for example, environmental projects or new consumer products, it is arguably more difficult for individuals to accurately and explicitly weigh up the aspects of a good life.

2. Revealed preference, such as observed market transactions, and decision utility inferred from observed choices, such as within behavioural experiments using monetary payments. These are difficult to implement in the case of non-market gains and losses, however, such as in health and education. There are models that use data on purchases of complementary goods, such as health insurance and school fees, however this possibility is not pursued here.

3. Subjective well-being (SWB) regression, the type of method chosen here. This method seems to resonate most in terms of its objective to expand our under-

> standing of social performance. Although the SWB approach advocates a different way of operationalising this objective, it shares some similar motivations with the proposed preference index approach. With the alignment of SWB research as a contender in the search for alternative performance measures, it seems natural to incorporate SWB regression as the preference estimation method here.

The BHPS is a representative sample of individuals aged over 16 in the UK. New entries and attrition means that the panel is unbalanced, with an average of 6 panels per individual. Wave 7 in 1996 marked the introduction of an additional self-completion questionnaire to the BHPS, asking individuals to indicate on a scale from 1 to 7 (very dissatisfied to very satisfied respectively) their satisfaction with various domains of life and life overall. Therefore, data from 1996 to the final wave in 2009 is used, excluding 2001 which omitted this life satisfaction question. This encompasses all waves of the BHPS containing the variables necessary for the analysis. Three dimensions of well-being are chosen for the analysis, similar to the main dimensions of the Human Development Index and in line with those investigated in Chapter 1: equivalised household income, health and education. This reflects an intent to frame the proposed concept of well-being around development objectives and public policy, therefore excluding private issues in family and social domains from the index dimensions for the present.

Equivalised household income is constructed by dividing annual household income by the square root of number of household members, and individuals with equivalised household incomes $>$£120,000 and $<$£100 are excluded from the analysis. The rationale for these choices are given in Chapter 1 Section 1.4.1. Again consistent with Chapter 1, the health dimension uses a composite measure derived from the following individual health indicators: whether an individual has been a hospital inpatient in the last year, whether an individual has problems with limbs, with chest or breathing, with heart or blood pressure, with stomach or digestion, with diabetes, with migraines, and with anxiety or depression. The composite measure is derived using rescaled predictions of the linear index from an ordered logit model of subjective health satisfaction, and is similar to the approach taken in Decancq et al. (2014a) and van Doorslaer and Jones (2003).

For the education dimension, a known problem with the indicator of highest education qualification, though commonly used, is that it confounds effects of prior education that have manifested themselves indirectly later in the life course, either through positive income effects or negative aspiration effects. For example, while some evidence points to a small positive association between education and life satisfaction (Veenhoven, 1996; Frey and Stutzer, 2002; Oswald and Powdthavee, 2007), contradictory findings in other studies have been suggested to be the result of raised aspirations that are unfulfilled or by the higher educated taking on more high-stress occupations (Stutzer, 2004; Ferrante, 2007; Sebates and Hammond, 2008). As mentioned in Section 1.4.1, the empirical evidence on education and life satisfaction is therefore quite mixed. In addition, since there is little variation in education level over the years, the education variable of highest education qualification provides limited information under the individual fixed effects model used in the following analysis. Therefore, for the education dimension an alternative variable is exploited for the SWB regression – the dummy variable indicating whether an individual has obtained a new qualification in the last year. Since this variable narrows the time frame under consideration to one year earlier as opposed to over the entire life course so far, the indirect life course effects through income and aspirations are removed. The estimated education dummy coefficients for each preference are then used to derive the education measure, by assigning the coefficient value to each additional qualification level an individual has attained and summing these values.

In relation to the definition of x in Section 2.3, we therefore treat all three indicators as continuous variables, though it should be acknowledged that the treatment of education is not as satisfactory as for the other dimensions due to the small number of education levels. To aid interpretation, analysis is conducted using normalised variables so that results that follow are interpretable with respect to a normalised $[0, 1]$ unit scale consistent with the Normalisation axiom.

2.4.2 A life satisfaction approach to estimating preferences

The proposed index of well-being requires that an ordinal representation of indifference curves is first estimated, in order to derive the individual indexes. This involves assuming that for each individual there is a stable mapping from dimension attainments to the latent variable that determines reported life satisfaction, and that this applies in all years of the survey.[3] This implies that an individual's rank according to her individual well-being index will correspond to her rank according to life satisfaction, and therefore the q-th quantile of the distribution of individual well-being will correspond to the q-th quantile of the distribution of life satisfaction (van Doorslaer and Jones, 2003).

More concretely let the following model be defined for life satisfaction, S_{it}, which is the outcome of attainment in dimensions of well-being and individual characteristics. This is

$$S_{it} = \alpha_i + \beta' X_{it} + \delta' Z_{it} + u_{it} \tag{2.3}$$

S_{it} is the self-reported life satisfaction of individual i in year t. X_{it} is the vector of attainment in the ℓ well-being dimensions of interest, in this case income, health and education. Z_{it} contains observed socio-demographic variables such as age, employment and marital status. These components of the model comprise a standard life satisfaction regression.

The empirical strategy further exploits the panel nature of the data to estimate an ordered logistic model of life satisfaction with individual fixed effects and individual-specific thresholds. This "fixed effect ordered logit" was developed in Ferrer-i Carbonell and Frijters (2004) and is further discussed in Frijters et al. (2006) and Jones and Schurer (2011). In practice, the model is estimated as a modification of the Chamberlain (1980) binary conditional fixed-effects logit model, with the modification allowing for an individual-specific rather than common life satisfaction threshold for each individual. This results in a much smaller loss of information compared to the original Chamberlain model since all individuals with any variation in satisfaction over time can be included, not just those with variation crossing over a fixed threshold. The resulting model results in a loss of only 8% of the observations.[4] A Hausman test confirms that fixed effects rather than random effects are appropriate, in line with a finding in the SWB literature that most panel studies examining determinants of life satisfaction have rejected the random effects assumption *i.e.*, the unobservable individual effects have been found in fact to be correlated with the explanatory variables (Frijters et al., 2006).

[3]Chapter 1 examines the invariance of this mapping in a simplified framework. Note that although it must remain stable within individuals, such a mapping will be allowed to vary from one individual to the next as explained in the following.

[4]Many thanks to Ada Ferrer-i-Carbonell for sharing Stata code for the simplified implementation of this model.

$$S_{it}^* = \alpha_i + \gamma_t + (\beta + \Lambda D_{it})'\Phi(X_{it}) + \delta' Z_{it} + u_{it} \tag{2.4}$$

The full model is given by (2.4), which further expands (2.3) to allow estimation of heterogeneous preferences. S_{it}^* is the latent life satisfaction variable such that reported life satisfaction $S_{it} = q$ for $q = 1, 2, ..., 7$ if S_{it}^* falls between thresholds η_{q-1} and η_q, where the η_q are individual-specific. α_i and γ_t capture unobserved individual and time fixed effects respectively, such as personality traits or aggregate shocks to the population. D_{it} is a vector of dummy variables to allow the estimation of heterogeneity in different partitions of the population. In theory one could conceive of using ever finer partitions to the extent of estimating heterogeneity at the individual level, however given the current data this is not possible. Furthermore it is arguably more useful from a policy perspective to learn about group-level rather than individual-level preferences, since policy-making can usually only be targeted at particular population sub-groups as identified by some socio-demographic characteristic. Φ is a function to be estimated, capturing the degree of elasticity of substitution between dimensions.

In order to pin down a suitable Φ, a generalised additive model (GAM) of (2.4) is first fitted using spline functions for the dimension variables. This allows for capturing flexible functional forms, using plots of the resulting relationships to give an initial idea of function curvatures. Note that splines cannot be fitted to the education indicator since it is a dummy variable, and must enter the model linearly. In a second step, the fit of this non-parametric model is compared to a fully parametric model for which an optimal power transformation is found, with the restricting assumption that there is a common transformation parameter for each dimension. Such a parametric specification allows the closest CES representation of the ordinal preferences to be determined, helping to pin down the more tractable index specification proposed in Section 2.3. The third step is then to search for the best-fitting parametric specification allowing the transformation parameter for each dimension to vary independently. This is done by searching over a fine m-dimensional grid of values. In this case $m = 2$ for the two continuous dimensions, income and health. By comparing the results of the latter two approaches, a picture can be obtained of how restrictive an assumption it is to impose CES preferences as opposed to allowing a more flexible and data-driven estimation of preference elasticities.

2.4.3 Results and comparison with other measurement approaches

The analysis for the first step in this three-step approach, the GAM estimation, is as follows. The GAM spline plots, given in Figures 2.3 and 2.4, suggest that non-linear transformations are appropriate for both the income and composite health indicators. These spline plots allow smooth piecewise polynomial curves, or splines, to be fitted to the chosen well-being variables as opposed to confining to only linear functions as in standard linear models. Cubic splines are used, ensuring smooth joining at the knots where the curves join. It can be seen from Figures 2.3 and 2.4 that the health spline exhibits a more obvious non-linear relationship with life satisfaction compared to the income spline, though it cannot be judged simply by looking at the plots whether this is significant enough to warrant different transformation parameters for each variable. This is further examined in the second and third steps.

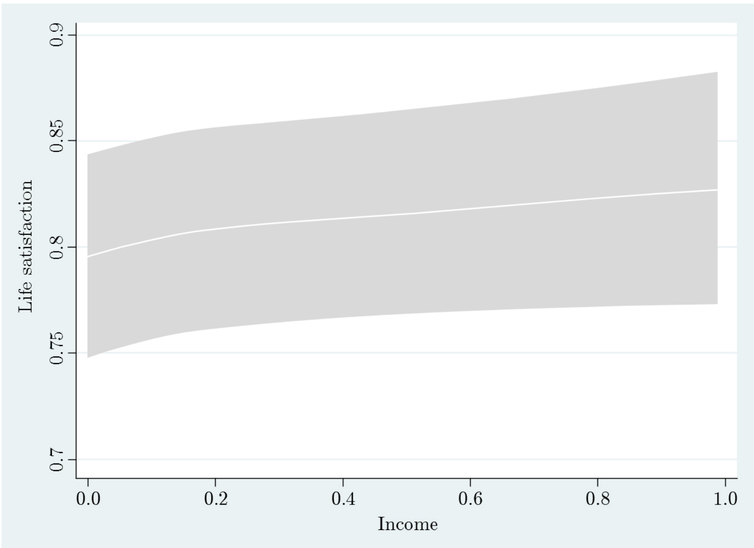

Figure 2.3: Spline function for the income variable

The second step is to compare the fit of this non-parametric model with fully parametric estimations imposing a common transformation parameter for each of the continuous dimensions. Model fit is compared using Akaike's Information Criterion (AIC), with lower AIC values indicating better fit. In the range $[-2, 2]$ of Box-Cox power transformation parameters tested, a common parameter of 0.2 gives the best-fitting model as shown in Table 2.1. The fit of this model is compared to those under integer-value parameters of 0 and 1 in turn, to identify the closest naturally interpretable approxim-

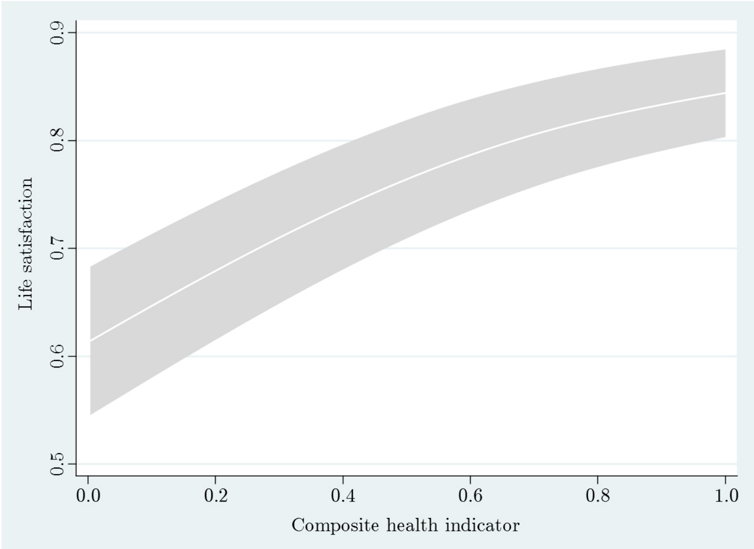

Figure 2.4: Spline function for the health variable

Transformation parameter	Observations	Degrees of freedom	AIC
Cubic spline	117358	31	104073.8
1	117358	29	104088.4
0.2	117358	29	104002.4
0	117358	29	104020.7

Table 2.1: Model fit for different curvatures of the well-being dimensions

ation to the 0.2 optimum value. A value of 0 reduces to a logarithmic transformation whereas a value of 1 equates to a linear relationship. The AIC values reveal that 0 provides a much closer approximation, which additionally gives a very palatable and tractable interpretation as the natural log transformation or equivalently Cobb-Douglas preferences or a weighted geometric mean of these two dimensions.

The third step is to see how restrictive this common-parameter assumption is in contrast to dropping the restriction and allowing for more flexible representation of the data. To do this a two-dimensional grid search is performed over 0.2 increments in the range $[-2, 2]$ for each dimension. The log-likelihood function is shown in Figure 2.5. Interestingly the likelihood-maximising values are again 0.2 simultaneously for both dimensions even without imposing the CES restriction, matching those obtained in the previous step. It may be concluded therefore that in this case the more tractable CES assumption has not imposed any restrictions on the model compared to the flexible GAM estimation. As discussed in the previous paragraph, given the closest integer approximation of 0 with its tractable logarithmic interpretation, in what follows this integer approximation will be used for the substitution elasticities of both the continuous dimensions, health and income.

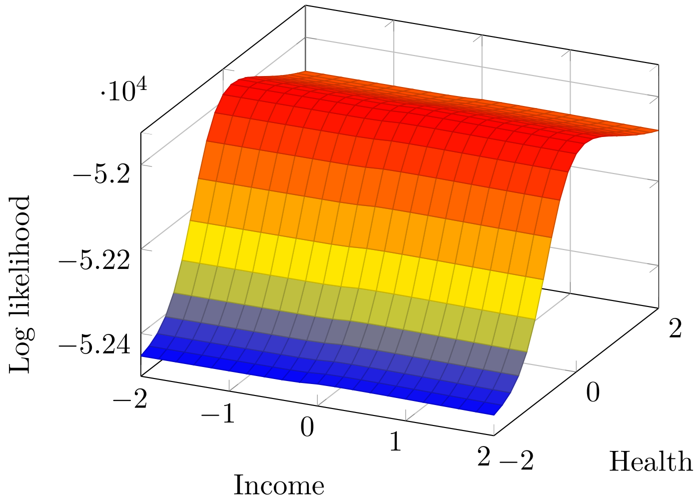

Figure 2.5: Log likelihood for two-dimensional grid search of transformation parameters

Following the preferred specification (2.4), the estimation results first without preference heterogeneity $\Lambda D'_{it}X_{it}$ are presented in the first column of Table 2.2 as a comparative "representative agent" approach. The second column reports the results when $\Lambda D'_{it}X_{it}$ interactions that gave significant effects were included in the regression. The pseudo R^2 values for both models are small, but in line with other fixed-effects studies of SWB (Graham et al., 2004, for example). With the exception of the age variables,[5] the socio-economic control variables were originally coded as categorical variables with each ranging from four to ten categories. For clarity of exposition, it was found that these variables could be reduced to two-category dummy variables without much loss of interpretation or change in magnitude of the dimension variable coefficients. It is the results using the simplified variables that are presented in Table 2.2.[6]

The effects of the socio-demographic variables are all as would be expected of a typical satisfaction regression. In the dimension variable coefficients, of note is the effect of education once interactions are added – the previously insignificant effect of education on satisfaction under the homogeneous preference model becomes very

[5]The age^2 variable is continuous and the age categories variable contains 5 categories.

[6]Details of the simplification procedure are available on request.

Satisfaction	Homogeneous model	Heterogeneous model
Equivalised income	0.038***	0.044***
	(0.015)	(0.016)
Health	0.544***	0.518***
	(0.018)	(0.019)
Education	0.044*	-0.009
	(0.026)	(0.036)
Young $\times$*income*		-0.044***
		(0.016)
Higher educated $\times$*income*		0.051**
		(0.024)
Young $\times$*health*		0.131***
		(0.040)
Young $\times$*education*		0.124**
		(0.052)
Unemployed	-0.600***	-0.604***
	(0.035)	(0.036)
Couple	0.314***	0.316***
	(0.041)	(0.042)
Separated	-0.242***	-0.236***
	(0.057)	(0.057)
Urban	-0.103**	-0.112***
	(0.041)	(0.041)
Age^2	0.000**	0.000
	(0.000)	(0.000)
Age categories	yes	yes
Social status class	yes	yes
Individual fixed effects	yes	yes
Year	yes	yes
N	117,353	116,267
Pseudo R^2	0.0195	0.0196

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.2: Satisfaction regression (standard errors in parentheses)

significant for young people under the heterogeneous model. This satisfaction effect for young people is lost when differences in preferences are not accounted for, and is in part due to the greater incidence of young people obtaining new qualifications compared to older people, but also highlights the underlying importance that young people place on educational qualifications. By way of a caveat, it may be that the increased life satisfaction of young people upon obtaining a qualification could be in part a result of having a positive collegiate experience rather than a fundamental result of the education itself. The imperfect nature of estimating such preferences must therefore be acknowledged in proceeding with the analysis, which is a general caveat and applicable especially to the well-being effects of education which are difficult to measure.

Inspecting the other interaction coefficients it can be seen that the young care less about income, and the higher educated care more about income. To test how these results depend on the precise definition of the dummy partition, a sensitivity analysis of the "young" dummy is carried out (unreported). The results of this sensitivity analysis indicate that the estimated age-related preferences for income start to turn from caring less to caring more at around 44 years, and for education these preferences turn from caring to not caring also at around 44 years. This is around the age when adults raising young children may also need to start caring for ageing parents – a combination which significantly increases the financial burden of supporting a family. It is therefore plausible that income becomes a priority at this stage, whilst the relative importance of accumulating additional qualifications falls. Interestingly, for the interaction of age and health, it appears that older people care less about health than younger people. This may be less counterintuitive than at first sight. In a Taiwan panel study by Collins et al. (2007) of 3,363 older persons, the authors find similar results suggesting that higher life satisfaction and optimism may indicate the presence of adaptive coping mechanisms, and that higher life satisfaction and optimism may in turn contribute to better health practices and to better physiological functioning in the longer term. In this analysis, the sensitivity analysis indicates that age-related preferences over health begin to turn towards caring less after the age of 68.

Comparing this health finding with comparable analysis in Decancq et al. (2015), they find the opposite using Russian data – there is a larger weight of health in the preferences of the old. Far from being a problematic inconsistency, these contrasting findings highlight the central argument for taking account of heterogeneous preferences. The implication is that older people living in the UK are less concerned about health

Preference type	Income	Health	Education
Young, lower educ	0.001	0.839	0.160
Young, higher educ	0.062	0.788	0.150
Older, lower educ	0.064	0.756	0.180
Older, higher educ	0.129	0.703	0.168
Representative agent	0.060	0.869	0.071
HDI approach	1/3	1/3	1/3
Income	1	0	0
Life satisfaction	-	-	-

Table 2.3: Preferences

relative to younger people, whereas in Russia older people are relatively more concerned about health. Examining the underlying fundamentals of health care and ageing in these two countries provides some insight to this result. Russia's social programmes and care for the elderly are plagued by meagre pensions and poor healthcare services; in the UK on the other hand, the existence of a high quality National Health Service and state and occupational pensions provide assurance for the health of the elderly. Russian sociologist Gennady Tikhonov captures the general sentiment, that "The difference between Russia and the West is that in our country old age is considered to be a time of loss and reminiscence, whereas in the West it's a time for new possibilities" (RT News, 2011). This seems to resonate with the observation of Deaton (2008) that whereas in the United States and Britain, health satisfaction actually improves with age after 50, in the the former Soviet Union health satisfaction falls very rapidly in the elderly. This is a difference that this preference-sensitive approach is able to capture.

Table 2.3 shows coefficients after a simple linear rescaling to give a clearer idea of the relative importance of each dimension. The coefficients in the table sum horizontally to one, however this rescaling can be chosen arbitrarily since the index treats preferences as ordinal and therefore only the *relative* weights are required. For all preference types income receives the lowest weight, though with some variation across groups. This result is consistent with the assertion of Deaton (2008, p. 54) that "many studies comparing people within countries have found only a small effect of income on life satisfaction relative to other life circumstances", citing as examples Helliwell (2003) and Blanchflower and Oswald (2004).

Health, on the other hand, receives a very high weight. Again this squares with similar findings in the literature on health and SWB, for example those of Campbell et al. (1976) that health was rated by subjects in the US as the most important factor on happiness. Interestingly however, other studies on the statistical association between *objective* health and SWB have tended to find that the relationship is a relatively weak one (Brief et al., 1993).

Calculating the marginal rate of substitution between income and health under representative agent preferences using the weights in Table 2.3, an individual with mean attainment in income and health would be willing to give up £2,130 in equivalised household income to improve her health attainment by 1 point on the health scale. This improvement roughly equates to eliminating an average of one problem from the BHPS list of health problems: limbs, chest or breathing, heart or blood pressure, stomach or digestion, diabetes, migraines, or anxiety or depression. For comparison, in their critical study of willingness-to-pay as a measure of health state preferences, O'Brien and Viramontes (1994, Table 3) find a willingness-to-pay of C$165 per month, or C$1,980 per year, for a therapy offering healthy lung functioning for individuals with household income between C$20,000 and C$39,999. The review of willingness-to-pay and health-status by Reed Johnson et al. (1997) finds estimates ranging from $1.18 per day, or $430.70 per year, for a mild cough, to $164.99 per day, or $60,221.35 per year, for severe angina. Clearly there is wide a range of estimates in the health literature for a spectrum of health conditions and severities, and the estimate implied here by the preference index application to health falls in a very reasonable position within that range.

Note that the education coefficient for the young group is more relevant as a measure of the estimated education effect, since people tend to obtain educational qualifications when young, and so the direct effect of education is not confounded with the effect of income later in life which will tend to run in tandem. Again however, the caveat applies as to whether this can be interpreted as the value of the education itself or simply a positive collegiate experience.

Figure 2.6 illustrates two groups of indifference curves – the older higher educated and the younger lower educated. To illustrate the point empirically that taking ac-

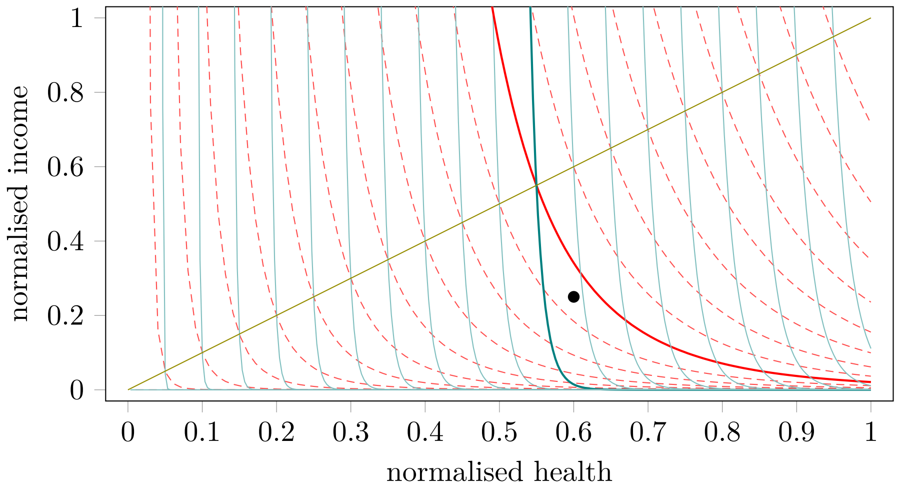

Figure 2.6: Indifference curves

count of heterogeneous preferences is important when measuring well-being, consider an individual situated at the attainment bundle marked by the black circle. If this individual were an older higher educated person (with the dashed indifference curves), this would be a position of lower preference satisfaction than if the individual were a younger lower educated person (with the steeper, solid indifference curves) situated at the same bundle. In stark contrast to conventional measures of well-being, with the preference index it is possible that two individuals with identical attainment can have differing ideas about their degree of well-being.

Quintiles of income	Quintiles of preference-sensitive measure 1	2	3	4	5
1	0.33	0.23	0.21	0.12	0.11
2	0.28	0.22	0.21	0.17	0.12
3	0.18	0.21	0.20	0.24	0.17
4	0.12	0.19	0.19	0.25	0.25
5	0.08	0.15	0.19	0.23	0.35

Table 2.4: Crosstabulation of different measures

To get a better idea of how the picture of well-being using the preference index measure corresponds with a number of other popular measures of welfare, some comparisons are presented in the following tables. Table 2.4 contains a cross-tabulation of

	Preference sensitive (1)	Represent-ative agent (2)	HDI approach (3)	Income (4)	Life sat-isfaction* (5)
< 60% of median attainment (%)	19.3	22.0	12.1	12.1	3.1
Equivalised income (£)	16,536	17,269	10,903	6,507	17,480
Health (0-1 scale)	0.24	0.26	0.32	0.65	0.42
Life satisfaction (1-7)	4.72	4.76	4.71	5.09	1.62
Male (%)	36.0	36.6	35.0	37.4	40.0
Young (%)	10.0	11.0	11.2	30.7	24.4
Higher educated (%)	12.1	13.2	2.5	10.1	15.9
Urban (%)	71.4	70.7	72.2	69.3	73.6
Unemployed (%)	18.1	17.1	22.9	16.5	34.8

* Median life satisfaction was 5, so figures are for those who responded 3 or lower.

Table 2.5: Average characteristics of the least well-off in 2008/9

quintiles of the preference index with quintiles of income as the sole well-being measure. It is immediately obvious on inspection of the diagonal that there is limited agreement between the two measures on the rankings of individual well-being positions. Table 2.5 expands the number of measures compared to include the homogeneous-preference "representative agent" index (2), the United Nations' Human Development Index (HDI) measure (3), and the raw life satisfaction measure (5), focusing on the policy-relevant problem of identifying the least well-off. This is defined as those individuals with < 60% of median attainment according to each measure. Let those identified as least well-off according to the preference index measure be referred to as the 'preference poor'.

In terms of dimension attainments, those identified as preference poor suffer from much worse health and somewhat lower life satisfaction than the solely income poor, whereas income does not seem to have such a bearing on preference-sensitive well-being since the preference poor have relatively high incomes. The preference index measure is more in line with the other multidimensional measures (2) and (3) and also with the raw satisfaction measure (5) though to a slightly lesser extent. Interestingly, all measures considered paint a similar socio-demographic picture of the average least well-off member of society – these tend to be older, lower educated female workers living in urban places. This is not to say that the measures necessarily identify the *same* individuals, only that the majority of individuals identified possess these similar characteristics.

2.5 Conclusion

The main objective of this chapter was to formulate a multidimensional measure of well-being that is generalisable to other dimensions, useful from a policy perspective, and that reduces sacrifices in the representation of interpersonal preference heterogeneities. The end goal was not to prescribe a definitive well-being measure, nor to make definitive conclusions about quality of life. However, the empirical illustration did demonstrate some of the interesting analysis possibilities that the proposed approach provides.

First, the theoretical 'equivalence approach' (Pazner and Schmeidler, 1978; Fleurbaey and Maniquet, 2011) underpinning the proposed preference index was introduced, as well as how the preference index approach differs from other implementations of the equivalence approach. The preference index was then axiomatically presented, drawing from existing work in welfare economic theory and incorporating considerations of interpersonal heterogeneity, fairness in evaluating different situations, and inequality in the distribution of well-being. An empirical illustration was then presented of how the preference index approach could be implemented and used for types of analysis that traditional measures cannot offer. In the chosen method of operationalising the approach, generalised additive modelling was first used to flexibly model the relationships between life satisfaction and dimensions of well-being. This non-parametric model was then compared to a fully parametric CES model, and a parametric model allowing differing elasticities between dimensions. It was found that CES in fact gave the closest parametric representation of preferences between dimensions, meaning that the CES form of the preference index specification was not restrictive in this case. The parameters of the resulting model pointed to a weighted geometric mean functional form. Interaction and individual fixed effects were further used to uncover the differing weights and therefore differing preferences between heterogeneous types of individuals, according to the equivalence approach.

A comparison of preferences was illustrated by separating individuals and their preferences according to age group and education level. Among other results, an interesting finding was that the older group had weaker preferences for health compared to the younger group. It was also shown how consideration of this kind of preference heterogeneity potentially changes our understanding of well-being in society compared with unidimensional measures such as income and SWB, and other composite measures. The preference index measure was able to reflect strong subjective preferences for good health across all individuals compared to relatively weak preferences for income. This was reflected in the wide disparity with income, compared to lesser disparities with the raw life satisfaction measure of SWB and other composite measures in terms of identification of the least well-off in society.

It is argued that this preference index approach, based on the work of Fleurbaey and others, is superior to its composite index predecessors because it is grounded in economic theory rather than using aggregation procedures for which there is no convincing theoretical basis. The main feature of the preference index measure illustrated in this chapter is its ability to take account of preferences within the population over the various dimensions of life, while at the same time not losing the more beneficial features of a composite index such as a normalised scale and mean-family functional form. A second feature is that it is able to take account of inequality in the distribution of well-being, and overlapping deprivations in multiple dimensions among the least well-off. The advantages of the preference index approach in these aspects are the topics of Chapter 3.

Chapter 3

Decomposing multidimensional inequality with heterogeneous preferences

3.1 Introduction

The study of inequality has long been concerned with the analysis of one dimension of well-being in particular – income. Indeed income has often served, either implicitly or explicitly, as a proxy for well-being as the objective of economic modelling and policy-making. More recently, however, academic research has moved towards greater recognition of the need for a more comprehensive notion of well-being beyond income. In tandem, greater policy emphasis has recognised the multidimensional nature of well-being as vital to advancing our understanding of how to study and shape "better policies for better lives".[1] In order to make practical use of this multidimensional concept of inequality, it must therefore first be measured. To this end, large strides have been made in recent decades, especially in the formalisation of procedures in constructing synthetic indices of multidimensional inequality. Whilst different measures have evolved from a focus on differing aspects of inequality measurement – for example by requiring that a measure satisfies certain mathematical and ethical properties – all have the similarity that they require every individual in the analysis to be identical in terms

[1] Quoted here is the motto of the OECD, which recently launched its own extensive research initiative with the aim of establishing more inclusive measures of well-being and progress.

of preferences between dimensions of life. This is not the case in practice however. Individuals have different priorities depending on their particular circumstances and characteristics. The contribution of this chapter in combination with the previous chapters lies in demonstrating (i) how multidimensional measures can be modified to relax this restriction on the treatment of preferences, (ii) the implications of doing so, and (iii) how such a proposed measure (the 'preference index' measure) can be used to develop a much richer picture of well-being and inequality.

When the object of inequality lies along a single dimension, there are no preferences as such to take account of, other than the fact that more is usually preferred to less. Moving to a multidimensional concept, however, requires the degree of commensurability of different dimensions of life to be considered, and the fact that this may vary from person to person. Chapter 2 makes a contribution to this end by explicitly characterising the aggregation rules of dimensions from an axiomatic point of view with the help of an equivalence approach framework (Pazner and Schmeidler, 1978; Decancq et al., 2015), and presenting an empirical strategy for how to measure heterogeneous preference parameters in a composite index of well-being. Most importantly, it is reasoned that this framework allows consistent comparisons to be made across individuals when well-being preferences differ, taking into account personal cognitive and circumstantial differences. Whereas Chapter 2 focused on multidimensional well-being at the level of the individual, the purpose of this chapter is to demonstrate that the proposed individual-level measures can be used in the analysis of multidimensional inequality and well-being at the *social* level to overcome conventional limitations of multidimensional indices with respect to subjective differences in preferences. It is shown that by borrowing tools from the literatures of income inequality and coalition game theory, it is possible to disentangle the contributions of differing preferences, dimension attainments, and interactions between dimensions towards overall inequality. Changes in these contributions can be monitored over time, providing a valuable insight into underlying trends behind the evolution of preference-sensitive well-being.

The tools in this chapter are not new, however they have never before been refashioned in the way proposed here for the decomposition of multidimensional inequality. In particular, the application of the Shapley (1953) value to decomposing multidimensional inequality with heterogeneous preferences is a new contribution. The shift from one dimension to multidimensional inequality brings complications that usually necessitate restrictions on the choice of analytical tools available for empirical analysis. This chapter ties together existing literatures, in order to both theoretically and empirically address the problem of accounting for preferences in multidimensional well-being analysis.

The rest of the chapter is laid out as follows: Section 3.2 presents the standard multidimensional framework, how the proposed preference index framework relates to this, and the specifics of this approach in theory; Section 3.3 explains how the preference index measure can overcome the conventional preference restrictions required by classic inequality decompositions by subgroup and factor source; Section 3.4 illustrates the measurement and decomposition approach from Sections 3.2 and 3.3 empirically using individual-level data for the UK; and Section 3.5 concludes.

3.2 Approaches to multidimensional inequality

A standard multidimensional framework is first introduced, around which the rest of the analysis can be based.

Consider a population denoted by the non-empty and finite set $N \subset \mathbb{N}_{++}$ of individuals, and let $\mathcal{N}$ denote the set of non-empty finite sub-sets of $\mathbb{N}_{++}$. Each individual $i \in N$ considers m dimensions of life that matter for her well-being. In the empirical illustration, the examples of health, education and income are used, but in general one may arrive at such a list of dimensions by various deliberation methods.[2] Attainment in dimension k by individual i is a positive real number x_{ik}, and the personal attainment bundle of individual i is an m-dimensional vector $x_{i\cdot} = (x_{i1}, x_{i2}, ..., x_{im})$. Let n denote the size of the set N. The distribution of attainments can then be represented as an $n \times m$ distribution matrix L with attainment bundle $x_{i\cdot}$ at the i-th row. The k-th column is then the n-dimensional vector $x_{\cdot k}$ of all individual attainments in dimension k. In future, the dot subscripts "$\cdot$" are dropped for notational convenience. The set of all distribution matrices is $\mathcal{D}$.

A simple way to examine the multidimensionality of well-being and inequality is to consider the evolution of each dimension one at a time. Examples of this approach

[2]Note that in contrast to any type of money metric well-being measure, there is no a priori requirement that income must be part of the well-being concept (although most people would probably agree that it is), and as a result this framework is a truly generalisable one.

can be found in Atkinson et al. (2002) and World Bank (2005), among others. With reference to the framework above, this would equate to examining the column vectors of matrix L and analysing the distribution of the elements x_k column by column. This approach does indeed provide a more multi-faceted picture of well-being than is possible with a sole focus on income distribution. However, it is clear that observing the dimensions as separate columns offers no regard for the interrelationship between columns and indeed across rows. In substantive terms, these translate into a neglect of interrelationships between dimensions and interpersonal comparability of well-being across individuals respectively. For example, compare two societies; one in which deprivations tend to cumulate across dimensions for certain members of society whilst advantage tends to cumulate for others, and another in which the same degree of deprivation and advantage in society as a whole is spread across different people over the multiple dimensions. The ability to differentiate these contrasting interrelationships would seem very important for a multidimensional approach to inequality, and something that the dimension-by-dimension approach fails to deliver.

An approach that does have the capacity to account for cumulative deprivation and advantage is the use of inequality indices, which having first been introduced in the univariate context of income inequality, now have multidimensional extensions (Maasoumi, 1986; Bourguignon, 1999; Tsui, 1999). The discussion here will focus on multidimensional indices based on the normative approach to inequality measurement, that is, indices derived from an explicit social evaluation (or social evaluation function as a representation thereof) of the possible distribution matrices in $\mathcal{D}$. [3] It has become standard that such multidimensional indices should satisfy a number of basic properties, following from generalisations of such considerations in the unidimensional case derived on ethically, intuitively and mathematically attractive grounds. One set of such properties is concerned with exactly the distributional interrelationships discussed in the previous paragraph. There will be a discussion later of some of these properties and where the proposed preference index approach diverges from those conventionally considered.

The social evaluation functions used to construct inequality indices in the normative approach often have an underlying two-step aggregation procedure, with each step

[3] A social evaluation is defined as a binary preference relation that allows one distribution to be ranked as "socially preferred" to another.

consisting of a unidimensional aggregation – either across the n individuals in society or across the m dimensions of well-being. Of the two possible ways of sequencing these aggregations, only one gives the desired sensitivity to dimensional interrelations and cumulative attainments (Kolm, 1977). These two aggregation steps are now examined in more detail to better place the proposed preference index approach in relation to existing literature.

In the aggregation sequence that offers the desired properties, the initial step is to aggregate across the m dimensions for each individual i resulting in a distribution of individual-specific measures of well-being. The second step is then to aggregate these individual well-being measures, either to obtain an overall social evaluation (function), or to directly construct the multidimensional inequality index by applying a univariate inequality index to the distribution of well-being measures obtained from the initial aggregation. The sequencing of this procedure results in an overall measure that reflects the cumulative effect of dimension attainments on individual well-being. Maasoumi (1986) was the first to propose a direct inequality index based on this particular two-step aggregation, and it will be seen that the preference index measure proposed here bears a resemblance to Maasoumi's final index. While this procedure for obtaining a multidimensional inequality index does coincide with the normative approach to inequality, Maasoumi's index is derived on an information-theoretic basis rather than a normative welfare theoretic one.[4]

In the alternative sequencing procedure, the order of the aggregations is reversed so that in the initial step, aggregation is done over the n individuals for each dimension k to arrive at m dimension-wise summary statistics. The second step then aggregates these summary statistics across dimensions into an overall measure. It is clear that this procedure cannot be sensitive to the interrelationships between dimensions since the distributional information is collapsed dimension by dimension in the initial aggregation step, independently of the distributional content of other dimensions. Despite this there exist prominent examples of this approach, such as the UNDP Human Development Index (UNDP, 2013), Gajdos and Weymark (2005) and Jones and Klenow (2010) to name a few. The approach essentially consists of modelling the population as a single "representative agent". And herein lies another problem with this interpretation: there

[4] Many thanks to Koen Decancq for pointing out that Bosmans et al. (2015) provide a retrospective justification for Maasoumi's approach from within a normative framework.

is no clear theoretical framework for how to aggregate dimensions at this population level, which is necessary in the second step. As a consequence, the implied trade-offs resulting from the second aggregation step of dimension summary statistics will either be arbitrary, or based upon a hypothetical representative agent which, in reality, may not in fact represent any individual in the population.

3.2.1 Respecting differing preferences

On the previous point of theoretical bases for aggregation, the first way of sequencing is therefore the preferred procedure here: aggregating first across the m dimensions and then over the i individuals. By considering dimensions at the individual level in the initial step there is a clear basis for this aggregation, since individuals go through life having to consider trade-offs between dimensions of well-being, and therefore must realistically evaluate their lives in these terms by having preferences between the dimensions. First aggregating dimensions according to individual-specific preferences and then carrying out the aggregation over members of the population therefore gives an "individualistic" rather than "representative agent" approach. The individualistic view is arguably the one that provides the stronger theoretical and ethical basis for making social evaluations using a multidimensional index . Standard inequality indices do not follow this view of preferences, however, as they implicitly assume the same preferences for all individuals through use of an *Anonymity* property. *Anonymity* requires symmetric treatment of all individuals so that exchanging the preferences of individuals does not affect overall social well-being. Therefore, this type of index presents a clash with respect of differing preferences as a central principle in the measurement of well-being.[5] The purpose of this chapter is not to convince readers that preferences should take precedence in the measurement of well-being. This is essentially an ethical position, and normative arguments for the validity and attractiveness of incorporating heterogeneous preferences in well-being comparisons are covered in Chapter 2 and elsewhere.[6] There purpose here is to propose tools that will hopefully be useful for the analysis and decomposition of inequality in the multidimensional well-being space, *given* that variation in preferences as well as variation in distribution are to be taken into consideration.

[5]As an exception, Cowell (1980) does suggest a class of "partially symmetric" additive inequality measures, which would allow for different preferences between subgroups of individuals.

[6]Readers are referred to Fleurbaey and Maniquet (2011), Decancq et al. (2014b) and Decancq et al. (2015).

Recalling the standard multidimensional framework introduced earlier, the following representation of individual preferences is added to that purely distributional framework, corresponding to Chapter 2 and originally based on Fleurbaey and Maniquet (2011). Each individual can be thought of as having personal values that guide an individual-specific view of what makes a good life, and these views can be represented as preferences over attainment bundles x belonging to the potential attainment set $X \subseteq \mathbb{R}_+^m$ (which is common to all individuals and represents the realm of possible well-being attainments). Let R_i denote individual i's complete preference ordering over the set X. When i prefers bundle x_i at least as much as bundle x_i', this is denoted by $x_i R_i x_i'$. Strict preference is denoted by $x_i P_i x_i'$ and indifference by $x_i I_i x_i'$. Let $\mathcal{R}$ denote the set of preferences over X that are continuous, monotonic and convex.

Now recalling the preferred individualistic aggregation procedure, individual i's preference ordering R_i is embedded into the first aggregation step by using ordinal information about these preferences to inform the functional form of the aggregation. By using indifference curve analysis from the equivalence approach (Pazner and Schmeidler, 1978; Fleurbaey and Maniquet, 2011) as a theoretical basis for representing R_i, it is possible to construct interpersonally comparable preference-sensitive measures of well-being $\phi(x_i, R_i)$ for all individuals $i \in N$ which are then carried forward into the second aggregation step across the individuals.[7] In the standard approach without consideration for preferences, it is assumed that two individuals with identical attainment bundles x experience the same degree of well-being. However, once the aggregation is over dimensions with possibly heterogeneous preferences, this is no longer the case and the $\phi(\cdot)$ functions capture ordinal information about these preferences R_i.[8]

An index of well-being is defined as a function $W : \mathcal{S} = \bigcup_{N \subset \mathcal{N}} X^N \times \mathcal{R}^N \to \mathbb{R}_+$, such that $W(x_N, R_N)$ gives the level of well-being in distribution x_N when preferences in the population are R_N. Given this, the definition of well-being obtained is consistent

[7] To avoid repetition, the reader is referred to Chapter 2 for a theoretical and empirical explanation of the construction of these individual-level measures.

[8] The function $\phi(\cdot)$ is different from a generic "utility function", since a utility function refers to any representation of ordinal preferences. The function given in Equation 3.2 in Section 3.2.2, which is a weighted constant elasticity of substitution function, is a particular cardinalisation of preferences that satisfies the particular axiomatic characterisation given in Chapter 2 Section 2.3. Although the proposed individual-level measure was derived on the basis of the equivalence approach, other authors have arrived at similar aggregation functions from alternative starting points, which seems to coincide with and support the proposed approach. For example, Maasoumi's index of multidimensional inequality based on information theory begins by minimising a multivariate Generalised Entropy measure of divergence to arrive at a CES function.

with the preferences of individuals in the population. The corresponding index of inequality is then $I : \mathcal{S} = \bigcup_{N \subset \mathcal{N}} X^N \times \mathcal{R}^N \to \mathbb{R}_+$, such that $I(x_N, R_N)$ gives the level of well-being inequality in distribution x_N when preferences in the population are R_N. Thus, given a set N of individuals and denoting its size by n, the distribution used in standard inequality indices refers to the n-list of attainment bundles $x_N = (x_i)_{i \in N} \in X^N$, whereas the distribution relevant to the proposed preference index approach is the n-list of preference-sensitive well-being measures defined over those attainment bundles $\phi_N(x_N, R_N) = \phi(x_i, R_i)_{i \in N} \in \mathcal{S}$. Given this representation of preferences via individualistic well-being measures using the $\phi(\cdot)$ functions, this can be made explicit by using $W(\phi(x, R))$ and $I(\phi(x, R))$ to denote the index of well-being and inequality respectively.

3.2.2 The Atkinson-Kolm-Sen approach

The key conceptual difference has been distinguished between standard multidimensional inequality indices, which evaluate purely distributional information, and the proposed preference-sensitive inequality index, which takes into account the preferences of the individuals to whom the distributions refer. The incorporation of these preferences was possible only by taking the individualistic aggregation procedure discussed, and by embedding preferences into the initial aggregation step using an equivalence approach framework. Presented again now is the complete index of multidimensional well-being introduced in Chapter 2, for which the second aggregation step is now examined in more detail, along with its normative interpretation for the analysis of inequality and its consideration of interrelationships between dimensions. The well-being index, which serves as a social evaluation function, is given by the following specification:

$$W(\phi(x, R)) = \begin{cases} \left(\frac{1}{n} \sum_{i}^{n} (\phi(x_i, R_i))^{\alpha} \right)^{\frac{1}{\alpha}} & \alpha < 1, \alpha \neq 0 \\ \prod_{i}^{n} (\phi(x_i, R_i))^{\frac{1}{n}} & \alpha = 0. \end{cases} \tag{3.1}$$

where

$$\phi(x_i, R_i) = \begin{cases} \left(\sum_{k=1}^{m} w_{ik} x_{ik}^{\rho} \right)^{\frac{1}{\rho}} & \rho < 1, \rho \neq 0 \\ \prod_{k=1}^{m} x_{ik}^{w_{ik}} & \rho = 0. \end{cases} \tag{3.2}$$

Equation 3.2 gives the general specification of the individual multidimensional well-being measures following the initial aggregation step. Now since the interest is directly in inequality rather than in making social evaluations, on arriving at the distribution of these measures instead of defining a social evaluation function in the $\mathcal{S}$ space as in Equation 3.1, the second step directly applies a unidimensional inequality index to the distribution of individual well-being measures following the normative approach. This normative approach to inequality attempts to find indices that are based on a set of reasonable, ethically attractive social value judgements, in which an explicit relationship is established between indices of inequality and social evaluation orderings over distributions. Following the pioneering works of Atkinson (1970), Kolm (1969) and Sen (1973) (AKS), normative inequality measurement in the unidimensional context focuses on finding the fraction of total goods or income that could be discarded without changing social well-being if the remainder were redistributed equally among all individuals. The greater is this fraction, the greater is the degree of prevailing inequality. Equivalently, social well-being can be thought of as comprising two parts: the average level of well-being taken over all individuals representing the optimum potential well-being of the society, and the loss in social well-being from this optimum due to unequal distribution of well-being across individuals. This normatively motivated and intuitively parameterised interpretation is highly attractive for practical applications, and is the reason why applying the univariate AKS framework to individual well-being measures is favoured here over, for example, Tsui's direct multidimensional AKS generalisation with less intuitive parameterisation (Tsui, 1995)[9] or generalised entropy approach to multidimensional inequality (Tsui, 1999).

In the multidimensional context, Kolm (1977) has proposed an extension to this unidimensional approach. A measure of multidimensional inequality is defined as the fraction of the aggregate amount of each dimension that could be discarded if every dimension of the distribution matrix were equalised across individuals, without changing the social well-being of the resulting distribution from the original distribution matrix. Tsui (1995) and others have since further developed this idea,[10] and there now exists a standard set of axioms which are usually considered in the construction of

[9]Brandolini (2008, pp. 14-15) provides an interpretation of a reformulated version of Tsui's (1995) index.

[10]See Weymark (2006) for a survey.

multidimensional inequality indices. Two of these are discussed in the following subsection. However, due to the sole focus of the standard approach on distribution matrices, again this type of approach cannot provide an answer to the pursuit of a measure that respects preferences. Maasoumi (1986) does also arrive at Equation 3.2 as a representation of individual well-being using the concept of relative entropy, arguing that this functional form provides a distribution as close as possible to the distribution of constituent attainments x_{ik}. However, Maasoumi suggests the use of principal components to derive the weights, while acknowledging this imposes a "rather ad hoc" restriction to linear aggregation Maasoumi (1986, p. 996). No other rationale is given for the elasticities of substitution in such a functional form, nor for the alternative suggestion of "any desired functional form". This is where preferences provide the answer, and the unidimensional AKS approach provides a natural framework into which the distribution of preference-derived individual well-being measures can be nested. Dropping the subscript i for notational parsimony, this preference-sensitive definition of inequality can be represented by:

$$W(\phi(x,R)) = \mu_\phi(1 - I(\phi(x,R)), \tag{3.3}$$

or equivalently

$$I(\phi(x,R)) = 1 - \frac{W(\phi(x,R))}{\mu_\phi}, \tag{3.4}$$

where μ_ϕ is the arithmetic mean of the individual well-being measures $\phi(x_i, R_i)$ taken over all individuals $i \in n$. From the equation above, it is clear that specifying the social evaluation function $W(\cdot)$[11] automatically pins down the inequality index $I(\cdot)$ and vice versa. This relationship is illustrated graphically in Figure 3.1 for the case of a two-person society.

Points along the dashed 45° line from the origin represent mean well-being distributions where each member of society attains equal multidimensional well-being according to the proposed preference- index measure. For any given distribution such as point

[11] $W(\cdot)$ is differentiated from a social welfare function (SWF) in that the 'welfarist' approach underpinning a SWF would only be concerned with the vector of individual well-being measures, but not with the underlying (multivariate) distribution that generates individual well-being. Since great lengths are taken here to account for this by explicitly modelling well-being preferences over multidimensional attainments in the first aggregation step, this is acknowledged in referring to $W(\cdot)$ as a social evaluation function.

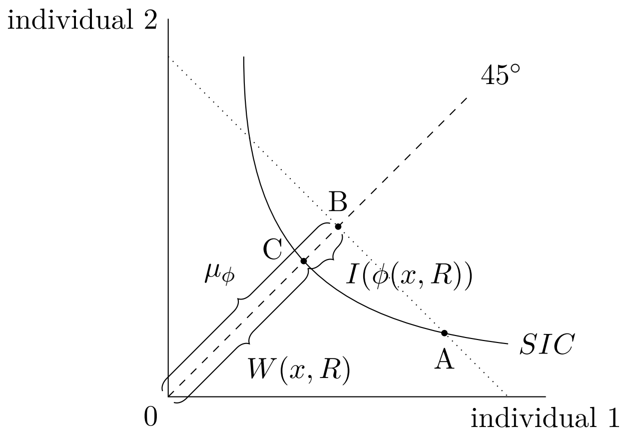

Figure 3.1: Illustration of the Atkinson-Sen-Kolm Approach

A, the dotted line gives all alternative hypothetical redistributions starting from distribution A, and hence its intersection B with the dashed 45° line gives the hypothetical optimum μ_ϕ where preference-sensitive well-being is equalised over individuals. The social indifference contour SIC gives well-being distributions that are equally as good as distribution A given a particular degree of aversion to inequality, with monotonically increasing social well-being as the contours move further away from the origin. Point C where the SIC coincides with the 45° line gives the distribution which is no better or worse than A, but in which $\phi(x_i, R_i)$ is attained equally by all individuals. This is the "equally distributed equivalent" (EDE) notion in the AKS approach, and can be interpreted as the smallest proportion of total well-being $\sum_{i=1}^{n} \phi(x_i, R_i)$ in the prevailing distribution that, if distributed optimally, would leave society on an equal ranking according to the social evaluation of $W(\phi(x, R))$. The EDE provides a particular cardinalisation for making social evaluations, and the index $W(\phi(x, R))$ hinges on this EDE interpretation. $I(\phi(x, R))$ can therefore be interpreted as the proportion of overall well-being *lost* to the society at distribution A due to inequality.

Note that the final well-being index inherits the cardinal properties of a relative index from the AKS approach used in the second aggregation step, whereas the cardinalisation of individualistic well-being measures from the first aggregation step is only a representation of ordinal information about preferences. Other ordinally equivalent representations of these preferences are therefore possible.

3.2.3 Distributional axioms

In the standard normative approach, distributional concerns are formalised by requiring that the inequality measure satisfy a number of well-known axioms. Comparing indices characterised by differing sets of axioms allows comparison of the differing value judgements implicit in these indices. Attention is focused here on three such multidimensional axioms, to the extent that they cannot be satisfied if a preference-sensitive index of inequality is to be pursued. Recall, however, that the preference index *does* satisfy the transfer principle, Axiom 6, proposed in Chapter 2 Section 2.3.2. This is defined in the unidimensional space of individuals' preference-sensitive well-being, $\phi(x_i, R_i)$, having obtained these measures in an initial multidimensional aggregation step at the individual level. In contrast, the axioms below define the transfer principle directly in the multidimensional space of population attainments in a single step. Since these axioms bypass preferences and deal only with attainments, they are not commensurable with a preference-sensitive approach, as will now be discussed.

- The first is the *Anonymity* axiom, sometimes referred to as *Symmetry,* which was mentioned previously.
- The second is a pair of variants of the multidimensional Pigou-Dalton Transfer Principle – *Uniform Pigou-Dalton Majorization* and *Uniform Majorization.*
- The third is *Correlation-Increasing Majorization*, which addresses correlation-sensitivity between the dimensions.

The *Anonymity* axiom requires that any permutation of attainment vectors from one individual to another should make no difference to the social evaluation of the society, and hence the evaluation of inequality. In the standard multidimensional approach this axiom allows the problem of making social evaluations to be reduced to the comparison of distribution matrices alone and disregards any differences in individual needs or preferences over the dimensions. By definition, if individual heterogeneity is acknowledged then individuals cannot be treated as anonymous, since permuting one individual's attainment vector with another individual will in general change their individual well-being, unless they have either identical preferences or identical attainments.

An inequality index $I(\cdot)$ satisfying *Uniform Pigou-Dalton Majorization (UPD)* evaluates inequality in a distribution matrix L as greater than inequality in distribution matrix LT, where T is a non-permutation matrix that performs a finite series of multidimensional Pigou-Dalton transfers. Formally, $T = \lambda U + (1-\lambda)Q$, $0 < \lambda < 1$, where U is an identity matrix and Q is a permutation matrix. Then we have that $(L, LT) \in UPD \Rightarrow I(L) > I(LT)$. *Uniform Majorization (UM)* is a slightly stronger condition than *UPD* and replaces the matrix T in the definition of *UPD* with a bistochastic matrix B in the definition of *UM*. As a result *UM* is stronger than *UPD* since any T matrix is also a non-permutation bistochastic matrix B, but there exists some B that cannot be expressed as the product of T matrices, that is, as the product of a series of Pigou-Dalton transfers.

The following example, adapted from Decancq and Lugo (2009), helps to illustrate *Uniform Majorization.* Consider the following matrices, which summarise the attainments of three individuals (rows) in two dimensions of well-being (columns):

$$B = \begin{bmatrix} 0.75 & 0.25 & 0 \\ 0.25 & 0.75 & 0 \\ 0 & 0 & 1 \end{bmatrix}; L = \begin{bmatrix} 50 & 80 \\ 90 & 20 \\ 10 & 50 \end{bmatrix} \text{ and } BL = \begin{bmatrix} 60 & 65 \\ 80 & 35 \\ 10 & 50 \end{bmatrix}$$

B is a bistochastic matrix, and BL is obtained from L by a bistochastic transformation. An inequality index satisfying *Uniform Majorization* evaluates distribution matrix L as more unequal than BL.

Whereas the two *Uniform Majorization* axioms impose criteria for evaluations based on the spread of the distribution of dimension attainments, *Correlation Increasing Majorization (CIM)* is a criterion based on the correlation structure between dimension distributions. *CIM* stipulates that for two distribution matrices K and L, if L can be derived from K by rearranging rows and making a finite number of correlation-increasing transfers between individuals, then distribution L is more unequal than distribution K. That is, $(K, L) \in CIM \Rightarrow I(L) > I(K)$. This axiom captures the concern for cumulative deprivation and advantage in individuals, so that given marginal distributions of dimension attainments, the greater the correlation between dimensions the more unequal the distribution is considered.

Once again, an example helps to illustrate. Consider the following matrices, again taken from Decancq and Lugo (2009):

$$K = \begin{bmatrix} 50 & 80 \\ 90 & 20 \\ 10 & 50 \end{bmatrix} \text{ and } L = \begin{bmatrix} 50 & 20 \\ 90 & 80 \\ 10 & 50 \end{bmatrix}.$$

Distribution matrix L is obtained from K by a correlation-increasing transfer, rearranging rows between the first two individuals. Of the first two individuals in L, the first individual has the lowest attainments in all dimensions whereas the second individual has the highest in all dimensions. *Correlation Increasing Majorization* stipulates that L is more unequal than K.

Divergence from the two *Uniform Majorization* axioms and *Correlation-Increasing Majorization* axiom necessarily follow from divergence from *Anonymity.* As stated, the asymmetry between individuals introduced by preferences means that the distribution matrix alone no longer defines how a society is ranked. However, by definition this is exactly what the *Majorization* axioms require since they are defined solely in the space of attainments. In the theory of fair allocation, Fleurbaey and Maniquet (2011) have proved that any approach prioritising the *Pareto Principle* by evaluating well-being in terms of individual indifference curves may conflict with a multidimensional *Pigou-Dalton Transfer Principle*, on which the *Majorization* axioms hinge. This is because the Pareto Principle prioritises preferences, whereas the multidimensional Pigou-Dalton Transfer Principle judges the desirability of transfers irrespective of preferences, and as discussed in Section 2.3.3 the two are therefore incompatible in the multidimensional context. A full illustration and proof is provided in Theorem 2.1 of Fleurbaey and Maniquet (2011).

The divergence from these standard axioms highlights that respect for preferences is a fundamental rather than simply ideological difference between the approach followed here and the standard distribution-only approach to inequality measurement. Practically speaking, this divergence points to deficiencies of the albeit neat standard approach by recognising that it cannot capture distributional concerns stemming from differing needs and personal values.

3.3 Decomposing preference-sensitive inequality

In the literature on inequality decomposition, two main avenues of analysis have emerged in seeking to decompose the sources of inequality. One has been the study of subgroup decomposition – the population's composition of subgroups partitioned by identifiable characteristics. Along this line, various desirable decomposition properties for indices have been put forward aiming to capture the contribution of each subgroup to the overall degree of inequality. The class of Generalised Entropy (GE) indices has been characterised as the only class satisfying perfect additive subgroup decomposability, allowing separation of the overall inequality measure into two component inequalities: the contribution of inequalities within each subgroup and the contribution between subgroups. The sum of these two components exactly equals overall inequality (Bourguignon, 1979; Cowell, 1980; Shorrocks, 1980). This intuitive decomposition property has made the GE class a popular choice in practical applications.[12] Returning to the focus on normative indices following the AKS approach however, a recent multiplicative subgroup decomposition for the Atkinson index has been proposed by Lasso de la Vega and Urrutia (2005), allowing an analogous decomposition into the exact *product* of within- and between-group components using the complementary *equality* index to the index of inequality. This multiplicative decomposition will be revisited in the following subsection in the context of the proposed preference index approach.

The second avenue of analysis has been concerned with the decomposition of factor sources. This has mostly been confined to the study of income inequality where incomes can be thought of as linearly composed of m separate income sources (or factors), so that the contribution of each factor to the overall degree of inequality can be identified. Factor source decomposition has been investigated for the GE class by Shorrocks (1982) and for the Gini index by Lerman and Yitzhaki (1985). Other authors include Cowell and Fiorio (2011), Shorrocks (2013) and Morduch and Sicular (2002). Maasoumi suggested an analogous decomposition by dimension for his multidimensional index of inequality (1986) based on the GE factor source decomposition. However, these methods are incompatible with the unconventional preference-sensitive approach

[12] Additive subgroup decompositions have also been proposed for the Gini and the Atkinson index. However, none of these are exact and leave behind an unattractive "residual" component. Numerous decompositions have been proposed for the Gini starting with Soltow (1960) and, prominently, Lambert and Aronson (1993). For additive subgroup decomposition of the Atkinson index see Blackorby, Donaldson, and Auersperg (1981).

proposed here, and therefore a different method of decomposition is proposed. This decomposition method uses the Shapley (1953) value, a well-known tool in cooperative game theory, which was first proposed to compute income source contributions independently by Chantreuil and Trannoy (2013) and Shorrocks (2013). It has not been adapted and applied in a multidimensional setting, however. This decomposition is dealt with in Section 3.3.2.

3.3.1 Multiplicative decomposition by population subgroups

Let us first turn to the subgroup decomposition of preference-sensitive inequality. Notice from Equations 3.1 and 3.4 (Section 3.2.2) that the second aggregation step can be interpreted as applying a unidimensional Atkinson index of inequality to the derived individualistic well-being measures. As a result, subgroup decomposition of the preference index measure can proceed as in the unidimensional case, with the unidimensional distribution consisting of the n $\phi_i(x_i, R_i)$ measures. The Atkinson index is characterised by an alternative decomposition propoerty proposed by Lasso de la Vega and Urrutia (2005) as "multiplicative decomposability", versus the popular additive decomposition property of GE indices.

Under multiplicative decomposability, the complementary *equality* index of the index of inequality can be separated into:

- A between-group component defined as the equality level of a hypothetical distribution, in which individual well-being is replaced by the arithmetic mean well-being in each subgroup.
- A within-group component defined as the weighted generalised mean of subgroup equality levels, with weights summing to one.

The product of these between-group and within-group components is exactly the overall degree of equality, with equality defined as the complement to inequality, $E(\phi(x, R)) = 1 - I(\phi(x, R))$, where $I(\cdot)$ is the index of inequality.

The authors show that the Atkinson class[13] is the only class of inequality indices satisfying this multiplicative decomposition property. Compared to the additive decomposition for the GE class, this multiplicative decomposition has an analogous in-

[13] The authors actually define an "extended" Atkinson class. However, indices belonging to the new tail of this extended class are not widely used.

terpretation for the between-group and within-group components, in equality terms, for the Atkinson class. It should be mentioned that Blackorby et al. (1981) have proposed a different multiplicative decomposition of the Atkinson class, also in terms of equality indices. However, they use subgroup-level "equally distributed equivalent" (EDE) attainment levels to determine the between-group component of overall inequality, whereas the decomposition adopted here retains the traditional subgroup arithmetic mean definition of between-group inequality.

As Lasso de la Vega and Urrutia (2005) explain, their multiplicative approach has several advantages. First, the sum of the decomposition coefficients is equal to 1 so that if the level of equality coincides in all groups, the within-group component is equal to that level. This is consistent with the traditional definition of the within-group component. Second, while the traditional within-group component in additive decomposition is defined as the arithmetic mean of subgroup levels, the multiplicative decomposition broadens this definition to the generalised mean of order α, with $\alpha < 1$. If the level of equality coincides in all groups, the arithmetic mean and other generalised means are equivalent; however, the greater the difference in subgroup equality levels, the more the generalised means penalise differences in subgroup equality levels. Third, the multiplicative decomposition allows the impact of marginal changes of a given group to be evaluated in terms of its effect on overall equality. This analysis is carried out in Section 3.4.4. Additive decomposition, on the other hand, must rely on approximations for such analysis of marginal changes (Theil and Sorooshian, 1979). Finally, the decomposition does not result in an extra, difficult-to-interpret interaction term, as has been pointed out with previous additive decompositions proposed for both the Gini and Atkinson class (Pyatt, 1976; Das and Parikh, 1981; Mookherjee and Shorrocks, 1982; Cowell, 1988).

Formally, the Lasso de la Vega and Urrutia (2005) multiplicative decomposition property is given by

$$E(\phi(x,R)) = \begin{cases} \left(\sum_{j=1}^{J} \omega_j \left[E\left(\phi(x,R)^j\right)\right]^{\alpha}\right)^{\frac{1}{\alpha}} \times E\left(\overline{\phi(\cdot)}^1, \overline{\phi(\cdot)}^2, ..., \overline{\phi(\cdot)}^J\right), & \alpha < 1, \alpha \neq 0 \\ \prod_{j=1}^{J} \left[E\left(\phi(x,R)^j\right)\right]^{\omega_j} \times E\left(\overline{\phi(\cdot)}^1, \overline{\phi(\cdot)}^2, ..., \overline{\phi(\cdot)}^J\right), & \alpha = 0 \end{cases} \tag{3.5}$$

where $\overline{\phi(\cdot)}^j$ denotes the mean well-being in subgroup j, the weights ω_j are a function of the subgroup population shares p_j and subgroup income shares s_j (or in this context, well-being shares) of subgroup j given by $\omega_j = \frac{p_j^{1-\alpha} s_j^{\alpha}}{\sum_{j=1}^{J} p_j^{1-\alpha} s_j^{\alpha}}$, and $\sum_{j=1}^{J} \omega_j = 1$.

In Equation 3.5, the first right-hand-side term is the within-group equality component and the second term is the between-group component. It is apparent from Equation 3.5 that for the case of $\alpha < 1, \alpha \neq 0$, the weights ω_j depend on both p_j and s_j, whereas for the case of $\alpha = 0$ this reduces to $\omega_j = p_j$. Since the within-subgroup component depends in turn on these weights (ω_j), if for example between-group inequality changes (which by definition means at least one s_j has changed), this will cause a change in the within-group component through ω_j for the case of $\alpha < 1, \alpha \neq 0$, even if actual within-group inequality has not changed. This is analogous to the "path independence" condition pointed out by Foster and Shneyerov (2000), Shorrocks (1980) and Anand (1983) that within- and between-subgroup components of inequality are specified independently of each other in only one case of the GE class and for the variance of logarithms. Of the Generalized Entropy class, only the case with the GE parameter $\alpha = 0$ satisfies path independent additive subgroup decomposition. Lasso de la Vega and Urrutia (2005) show that similarly for the Atkinson class, only the case of $\alpha = 0$ provides independent multiplicative decomposition components. As a result, this is the specification used in the empirical analysis. Note that the GE class is ordinally equivalent to the Atkinson class. However, only the Atkinson class is derived from explicit normative foundations with a resulting intuitive interpretation of inequality in terms of welfare loss.

3.3.2 Shapley value decomposition by dimension contributions

In contrast to subgroup decomposition, the multidimensional inequality index $I(\phi(x, R))$ is not readily decomposable by dimensions since heterogeneity is allowed between individual well-being specifications. The index does not, therefore, satisfy the standard *Anonymity* axiom, although it does satisfy the *partial* anonymity property proposed in Cowell (1980). A decomposition solution exists, however, by looking to a different literature.

In the field of cooperative game theory, the Shapley value is a well-known concept in the analysis of superadditive games in which players can form coalitions to improve their payoff. A key question in cooperative game theory is then how to distribute the surplus achieved through cooperation. Shapley (1953) proposed dividing the coalition payoff according to players' individual expected marginal contributions to this payoff. For each player, such a contribution is measured as the average marginal increase in the payoff of any coalition, resulting from the addition of this player to the coalition. Shapley showed that this concept – now known as the Shapley value – is the only payoff distribution mechanism satisfying certain desirable normative conditions. Translating this concept for the decomposition of multidimensional inequality, instead of calculating payoffs for players, it is possible to use the Shapley value to calculate the contribution of each dimension of well-being towards overall inequality. More precisely, the Shapley value of each dimension can be interpreted as the average marginal contribution made by that dimension to overall inequality taking all dimensions together. This concept has notably been proposed in the context of income inequality and unidimensional poverty analysis (Shorrocks, 2013), but this is the first application to multidimensional inequality, and in particular to the treatment of heterogeneous preferences.

Let $I(\phi(x, R))$ denote the multidimensional inequality measure, recalling that for each individual $i \in N$ in the population, attainment bundle x_i consists of $x_i = \{x_{i1}, \ldots, x_{im}\}$ for the m dimensions of well-being contained in the set M over which the index is computed. Let $\mu_1, \ldots, \mu_m$ denote the mean dimension attainments – these are situations where inequality in dimension $k \in M$ is eliminated, or "switched off". The actual distributions of x_k for $k \in M$ are the situations where inequality is "switched on" – it is usually the case that there is some degree of inequality within a population. There are 2^m ways of switching inequality in different combinations of dimensions on and off. For each dimension k, from these 2^m combinations pairs of combinations must be compared which are identical except that inequality in dimension k is "on" in the first combination and "off" in the second. Let set S denote a given combination of dimensions with inequality switched "on", with the other dimensions "off". Then the set $S \cup \{k\}$ denotes the same combination, now with additional inequality from dimension k switched on. Therefore, comparing (hypothetical) inequality in $S \cup \{k\}$ and inequality in S gives one possible marginal contribution of dimension k to overall inequality. To calculate the inequality in S, a function $v(S)$ is defined which recomputes individual well-being measures $\phi(x_i, R_i)$ given by Equation 3.2 and the inequality index given by Equations 3.1 and 3.4, with the attainments in dimensions $j \notin S$ set to μ_j.

Given the definitions above, the Shapley value for computing the contribution of dimension k to multidimensional inequality in all m dimensions is:

$$\Phi_k(m, v) = \frac{1}{m!} \sum_{S \subseteq m \backslash \{k\}} |S|!(m - |S| - 1)! \left[v(S \cup \{k\}) - v(S)\right] \tag{3.6}$$

As an example, let dimension k be the health dimension, and m is the full number of dimensions in the set M (in this case $m = 3$ for health, education and income). First consider S as the set of remaining variables when dimension k is dropped from set M. $[v(S \cup \{k\}) - v(S)]$ measures the difference between inequality over all m dimensions, and inequality having eliminated health inequality from the computation. This gives one possible marginal contribution of health inequality. However, it must also be considered that the set S (containing the two other dimensions education and income in this example) could have been formed in 2! different sequences prior to the introduction of health. In general, the number of possible sequences is $|S|!$. Then, repeating the elimination of health inequality over other possible sets of dimensions S gives the number of all possible marginal contributions: $\sum_{S \subseteq m \backslash \{k\}} |S|!(m - |S| - 1)!$. $m!$ marginal contributions are then obtained for dimension k. Averaging these marginal contributions by multiplying them by $\frac{1}{m!}$ allows the expected marginal contribution of health inequality to be determined. This computation can be used to find the find the contribution of all k dimensions, providing an additive decomposition of the inequality measure by dimension.

An important characteristic of the application of the Shapley decomposition procedure to the proposed preference-sensitive measure is that, whereas for a standard multidimensional index $v(S) = 0$ for $S = \{\emptyset\}$, this is no longer the case with the introduction of preferences. Intuitively, if inequality is "switched off" in all k dimensions by replacing attainments in each dimension by the dimension mean μ_k (i.e. so that the set S of dimensions with inequality switched "on" is empty), then a standard multidimensional index would consider there to be no inequality remaining in the distribution. For the preference index measure however, even if the distribution of attainments is equalised across individuals in all dimensions, there is still inequality arising from the heterogeneity of preferences, so that in general $v(S) < 0$ for $S = \{\emptyset\}$. The importance of this lies in the observation that, as well as using the Shapley value to find the contribution of dimensions to overall inequality, the $v(\{\emptyset\})$ component can also be used to determine the contribution of preference heterogeneity to overall inequality.

3.4 Application to UK data

The proposed multidimensional inequality measure and decomposition tools are now applied to data from the British Household Panel Survey (BHPS). Following the empirical application in Chapter 2, data is taken from the 1996-2008 waves of the BHPS to examine three widely-investigated dimensions from the literature: health, education and income. Details of the variables and preliminary data processing procedures used are covered in Chapter 2 Section 2.4.1.

Recall from Equations 3.3 and 3.4 that inequality in the normative approach is identified as the difference in well-being given by the actual (univariate) distribution and that given by the hypothetical optimal distribution, where all individuals receive the mean. With that in mind, the mean trends are first considered in the separate dimensions and in the multidimensional measure of individual well-being W, shown in Figure 3.2. All k dimensions are normalised using the same min-max goalpost approach described in Chapter (1) Section (1.4.1), to range between 0 and 1 as follows:

$$\hat{x}_{itk} = \frac{x_{itk} - x_{k\min}}{x_{k\max} - x_{k\min}} \tag{3.7}$$

The results from Chapter 2 Section 2.4 revealed that all empirically estimated preference types from the BHPS place significant weight on health in the multidimensional measure; it is therefore unsurprising that the trend for W closely follows the trend for health in Figure 3.2. Mean attainment in income and education rose from 1996-2008 whereas mean health and mean multidimensional well-being remained largely unchanged aside from a slight dip leading up to 2000 and a subsequent gradual recovery afterwards.

Next the inequality trends in the separate dimensions and in W are examined. In all following instances of empirical calculations involving inequality index $I(\phi(x, R))$, unless otherwise noted the specification $\alpha = 0$ for Equation 3.1 is used. This is due to the "path independence" property discussed in Section 3.3.1, for which only the specification $\alpha = 0$ satisfies independent multiplicative within- and between-group de-

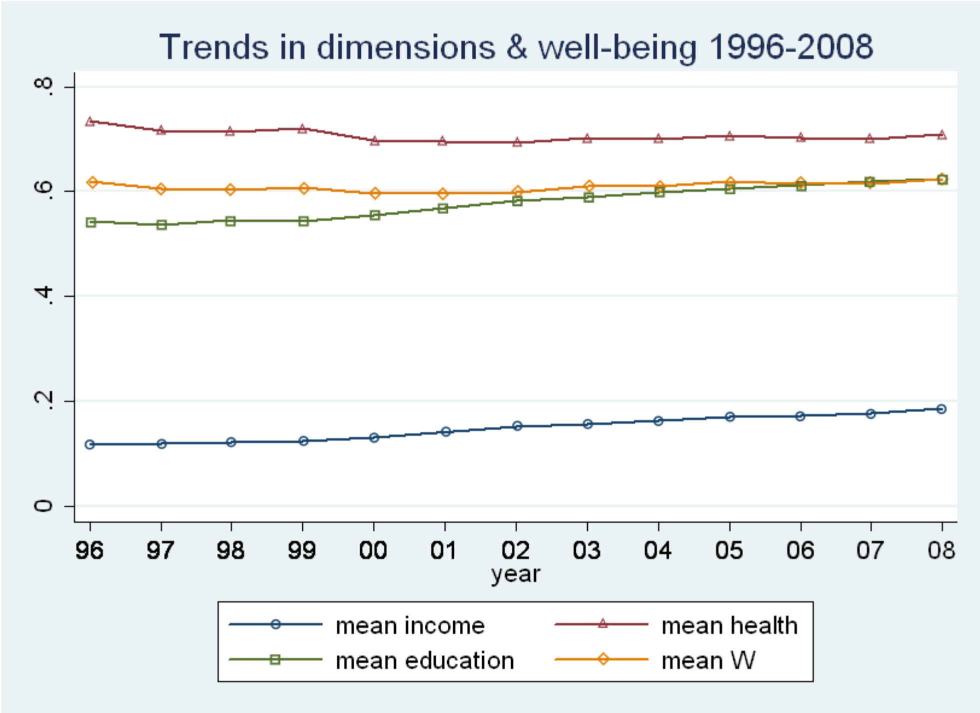

Figure 3.2: Trends in mean dimensions and mean multidimensional well-being

composition. Independence of the two components allows the contributions of different groups to changes in overall inequality to be calculated and interpreted unambiguously, the analysis and results for which are presented later. Figure 3.3 shows an equalising trend in income and education attainment over the period 1996-2008, whilst attainment in health and multidimensional well-being become increasingly unequal.

Although the BHPS data shows a trend of falling income inequality, it has been noted in the unidimensional UK income literature that differences in trends arise between different sources due to varying coverage of incomes at the top and bottom of the distribution. Jenkins (2010), for example, highlights differences in BHPS income inequality trends compared to the larger and more specialised Households Below Average Income (HBAI) survey data, and in turn Burkhauser et al. (2016) compare differences between the HBAI data and tax return data from the Survey of Personal Incomes. The trends presented here are not, therefore, a definitive indication of the wider UK context beyond the BHPS.

3.4.1 The inequality aversion parameter α

Although in theoretical terms, the distributional axioms discussed in Section 3.2.2 are not satisfied by the proposed preference index, the relevant properties can still often

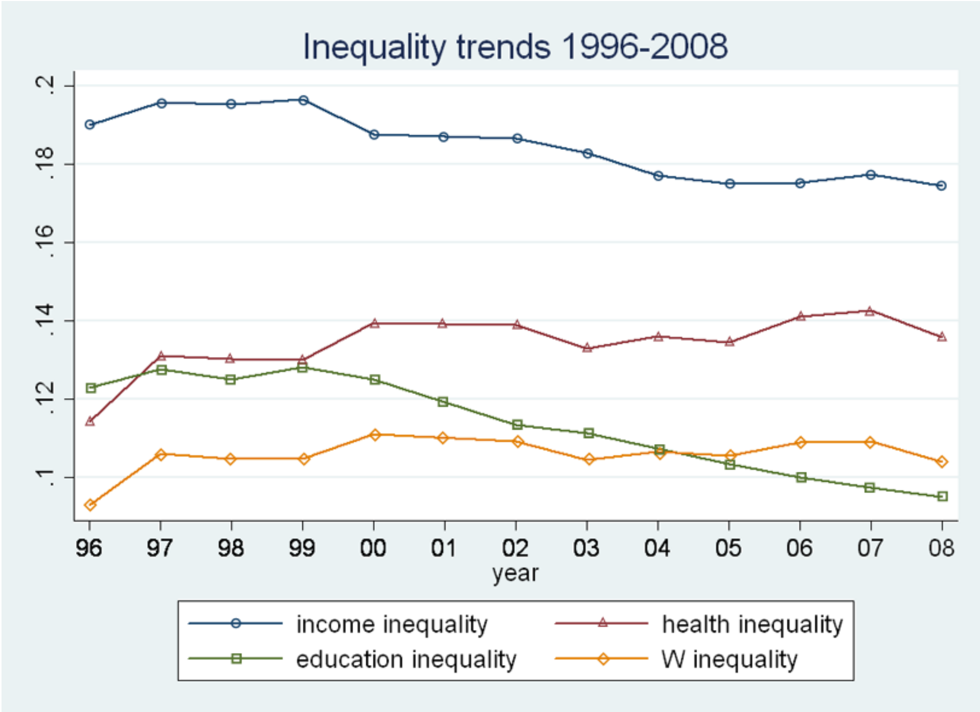

Figure 3.3: Trends in dimension inequality and inequality in multidimensional well-being

be observed in practice. It can be shown that this multidimensional measure is indeed sensitive to correlation-increasing and correlation-decreasing rearrangements, using a dominance criterion proposed by Lasso de la Vega and Urrutia (2008) for ordering distributions in terms of inequality. This dominance criterion will be explained and illustrated using the BHPS data, however a brief discussion of inequality aversion is first warranted.

The degree of inequality aversion built into the preference index is captured by the α parameter in the second aggregation step given by Equation 3.1, when the distribution of individualistic well-being measures is subsumed into the overall social evaluation function or inequality index. Although the inequality aversion parameter $\alpha = 0$ was chosen in this instance for its desirable decomposition properties, in general since different indices of inequality vary in the way they treat inequality in different parts of the distribution, it is possible that they may contradict one another in their social evaluations of pairs of distributions. The canonical method of establishing agreement between all indices in the class of relative, transfer-sensitive inequality indices is to check that the Lorenz curves of the distributions do not intersect. If this is the case, then one distribution can be ranked as unambiguously more unequal, and is said to be

Lorenz-dominated by the other. Lasso de la Vega and Urrutia (2008) have also shown that, for the subset of aggregative inequality indices, in some cases where Lorenz curves do intersect and therefore cannot be used to rank the distributions, agreement on their ranking may still be reached by comparing an alternative curve. This is the curve of the inequality level drawn as a function of the inequality aversion parameter α, and an analogous dominance criterion using this curve can be used to help ascertain whether the social ranking of distributions changes with α. The following analysis gives the results of this approach.

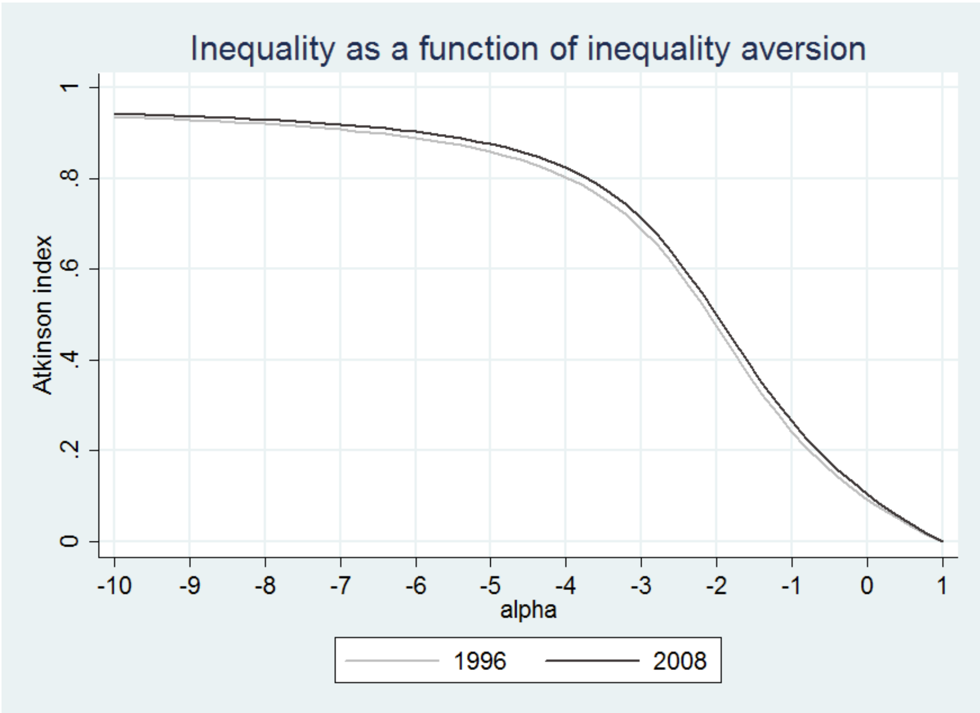

Figure 3.4: Inequality as a function of the inequality aversion parameter α

	Health	Income	Education
Health	1.0000		
Income	0.1733	1.0000	
Education	0.2188	0.3400	1.0000

Table 3.1: Correlations between well-being dimensions

Figure 3.4 shows inequality measured by $I(\phi(x, R))$ as a function of the inequality aversion parameter α in the first and the last year of the data. Inequality in 2008 is everywhere greater than inequality in 1996 irrespective of the level of inequality aversion. Now, for the year 2006, a hypothetical correlation-*increasing* redistribution is

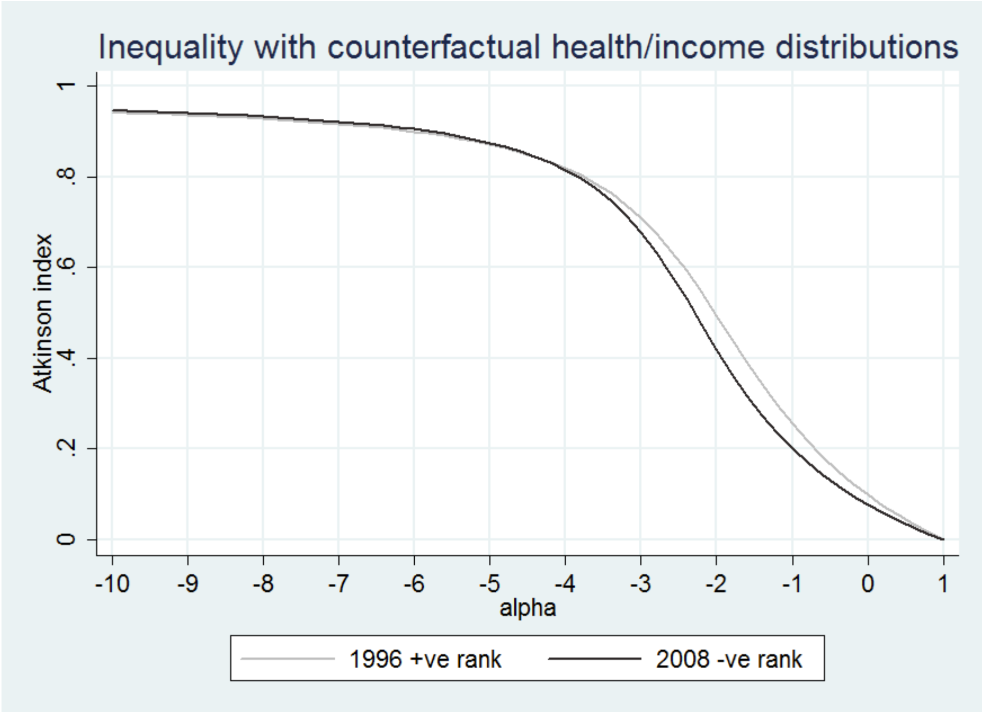

Figure 3.5: Inequality aversion dominance curves following alternate correlation rearrangements

carried out so that high attainment is cumulative across dimensions at one end of the distribution, and similarly across the spectrum of attainments so that low attainment is also cumulative across dimensions. In contrast, for the year 2008, a hypothetical correlation-*decreasing* redistribution is carried out so that the higher an individual's attainment in one dimension the lower the attainment in another dimension. The original correlation structure from the data is given in Table 3.1. The correlation rearrangements are implemented with respect to the health and income dimensions, so that health-income correlation is 1 after the correlation-increasing rearrangement and -1 after the correlation-decreasing rearrangement.[14] The resultant re-ranking of inequality in 2006 and 2008 in Figure 3.5 reflects the sensitivity of the $I(\phi(x, R))$ measure to greater multidimensional inequality as a result of the correlation-increasing rearrangement in 2006, in contrast to the greater multidimensional equality in 2008. Recalling the discussion of distributional axioms in Section 3.2.3, this is exactly the essence of the *Correlation Increasing Majorization* axiom.

These correlation rearrangements can also be used to illustrate the previous theoretical argument made for the insufficiency of a dimension-by-dimension approach.

[14] In fact the correlations will be slightly different from 1 and -1 since the rearrangements are implemented by year, whereas for brevity Table 3.1 presents the correlations over the pooled dataset.

Figure 3.6 repeats the multidimensional inequality trends presented in Figure 3.3 after correlation rearrangements of perfectly positively correlated attainments in the health and income dimensions (line with asterisk markers), and perfectly negatively correlated dimension attainments (line with no markers), as well as the separate dimension-wise inequality trends. Observe that the separate dimension trends fail to reflect either of the new correlation structures between the dimensions of well-being – something that the multidimensional measure W does pick up. Since the separate dimension inequalities remain impervious to these differing interrelationships, they provide no insight to such multidimensional inequalities. Considering that redistributive or spending policy changes and socio-economic shocks have the potential to change these correlation structures, and thus alter the trajectory of experienced multidimensional inequality within the wide bounds of these alternative distribution scenarios, sensitivity to such changes is an important feature of multidimensional inequality indices such as the one proposed.

3.4.2 Dimension contributions to multidimensional inequality

The correlation rearrangements show that dimension-by-dimension analysis is completely insensitive to differences in cumulative attainments across dimensions, whereas sensitivity to these differences is an important aspect of the multidimensional approach. It is evident, therefore, that these two approaches to analysing inequality in multiple dimensions are very different. Bearing this in mind the Shapley decomposition of the separate dimension contributions is now presented, which *does* take into account the interrelationships between dimensions, and also the role of preferences, when separating out their contributions to multidimensional inequality. Using Equation 3.6 to compute the dimension-wise inequality contributions, Figure 3.7 can be plotted, which shows the proportional contributions of health, education, income and preferences to multidimensional inequality.

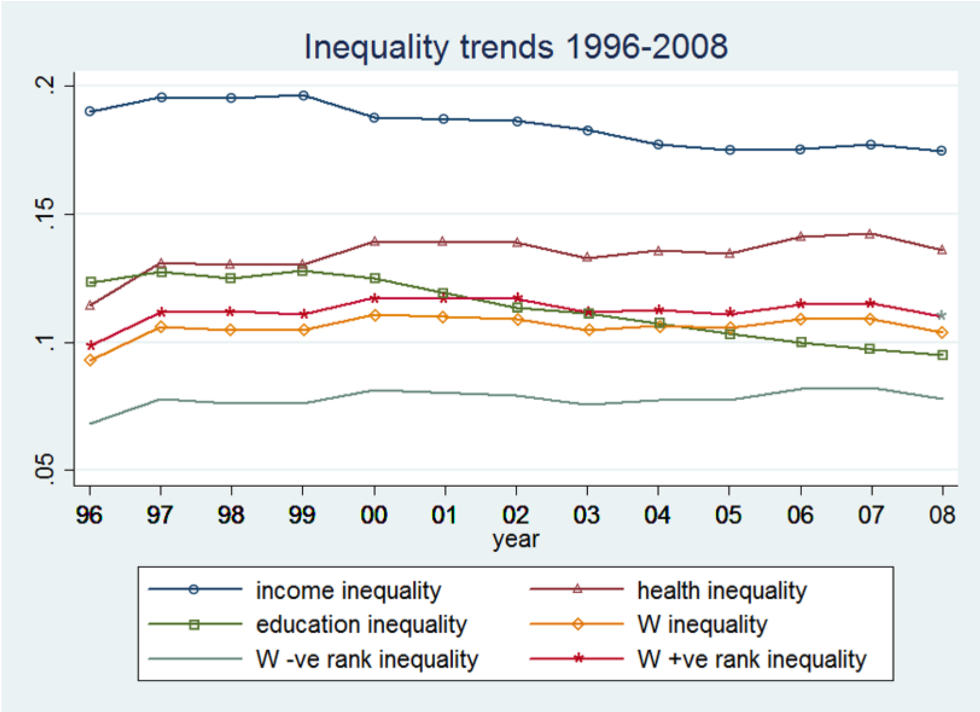

Figure 3.6: Inequality trends following correlation-decreasing and -increasing rearrangements

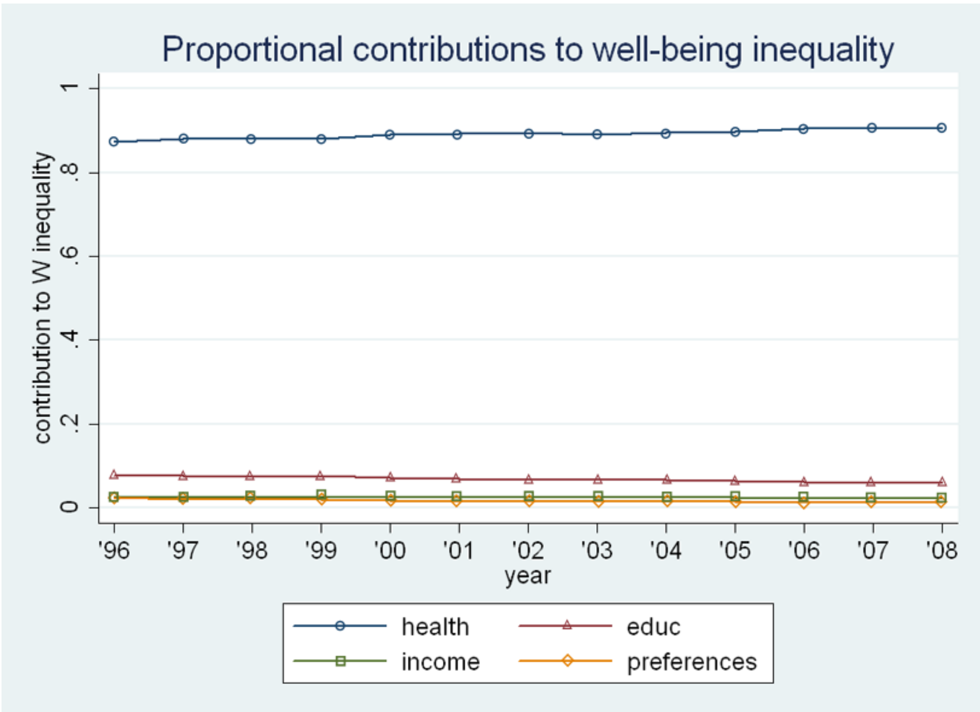

Figure 3.7: Dimension contributions to multidimensional well-being inequality when preference heterogeneity is considered

The proportional contributions are the rescaled contributions summing to one for ease of interpretation. Health is by far the largest contributor to well-being inequality, with its proportional contribution increasing throughout the period. Prominently, when heterogeneous preferences are taken into account as they are in Figure 3.7, the contribution of differing preferences to multidimensional inequality is comparable in magnitude to the contribution of income. This is a considerable finding, and somewhat undermines the emphasis in applied literature on income inequality whilst highlighting the significance of respecting heterogeneous preferences, often assumed away as an analysis inconvenience in favour of the simplifying assumption of identical preferences.

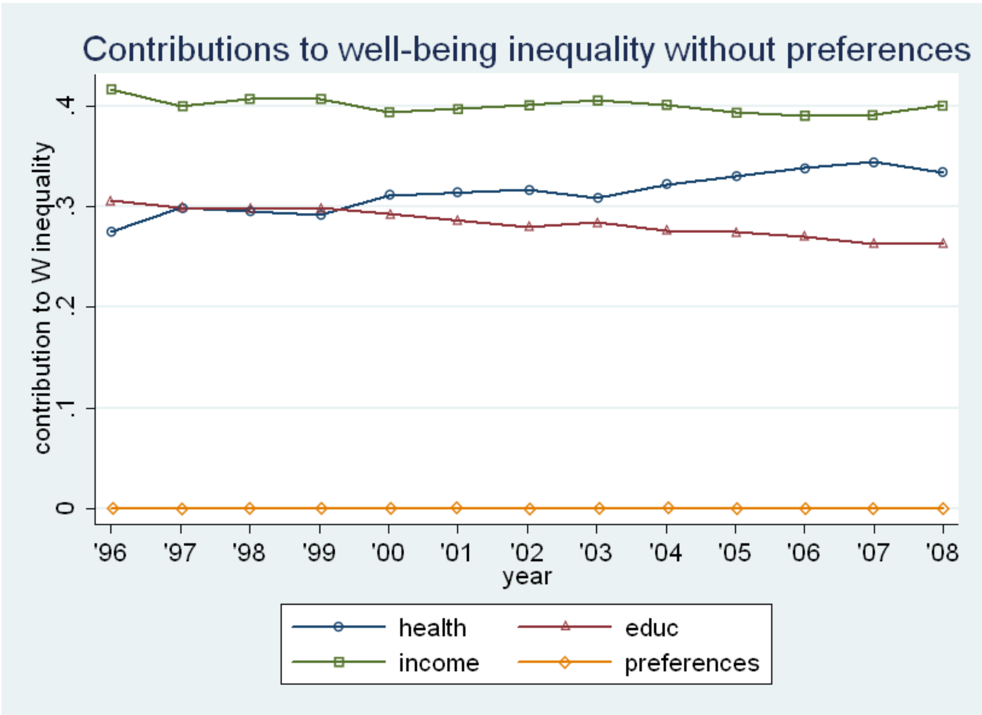

Figure 3.8: Dimension contributions to multidimensional well-being inequality when preferences are homogeneous

In contrast, Figure 3.8 repeats the Shapley decomposition when the empirically estimated heterogeneous preferences are removed altogether, and instead recomputes the contributions giving equal weight to all three dimensions. From Figure 3.8, preferences have been suppressed and of course are given zero contribution to multidimensional inequality. The same increase over time in the contribution of health can be seen as in Figure 3.7. However, income now appears as the largest contributor with a broadly unchanged contribution over time, and education by the end of the period is the smallest contributor. The stark difference in the contributions of income and health to inequality in the two approaches – with and without considering the derived preference weights – is an important finding. It strongly illustrates how the traditional emphasis on income inequality seems justified when the importance of other dimensions on well-being are not duly recognised, and yet when preferences are considered it becomes obvious that income plays a comparatively much less significant role in terms of multidimensional inequality.

So far, the inequality analysis has dealt with heterogeneous preferences by incorporating them into the multidimensional measures of individual well-being themselves. This has provided a population-level view of the role of preference heterogeneity and its determination of dimension contributions to multidimensional inequality. Now, the differing experiences of individuals belonging to the different preference types are explicitly examined by conducting a decomposition of dimension contributions by subgroup, where the subgroups are partitioned according to preference type. Note that this choice of preference type as the partitioning criteria is due to particular interest here in further examining preference heterogeneity, however such subgroup analysis need not be partitioned in this way and indeed any other criteria of interest may be used.

3.4.3 Dimension contributions by subgroup

In the previous analysis, the Shapley value was used to decompose preference-sensitive multidimensional inequality into its dimension contributions for the population *as a whole*. Now the Shapley decomposition analysis is applied subgroup-by-subgroup to examine what additional insights can be gained about different groups of individuals within the population, focusing on heterogeneous preference types. The evolution of preference-sensitive inequality by subgroup is first inspected without decomposing the dimension contributions, as shown in Figure 3.9.

Figure 3.9 reveals that while three of the four preference groups did experience the increasing well-being inequality observed in Figure 3.3 for the population as a whole, the younger more educated group saw preference-sensitive well-being inequality decline in general over that period. This can largely be attributed to the influence of collectively high and increasingly uniform health attainment of this group compared to other groups.

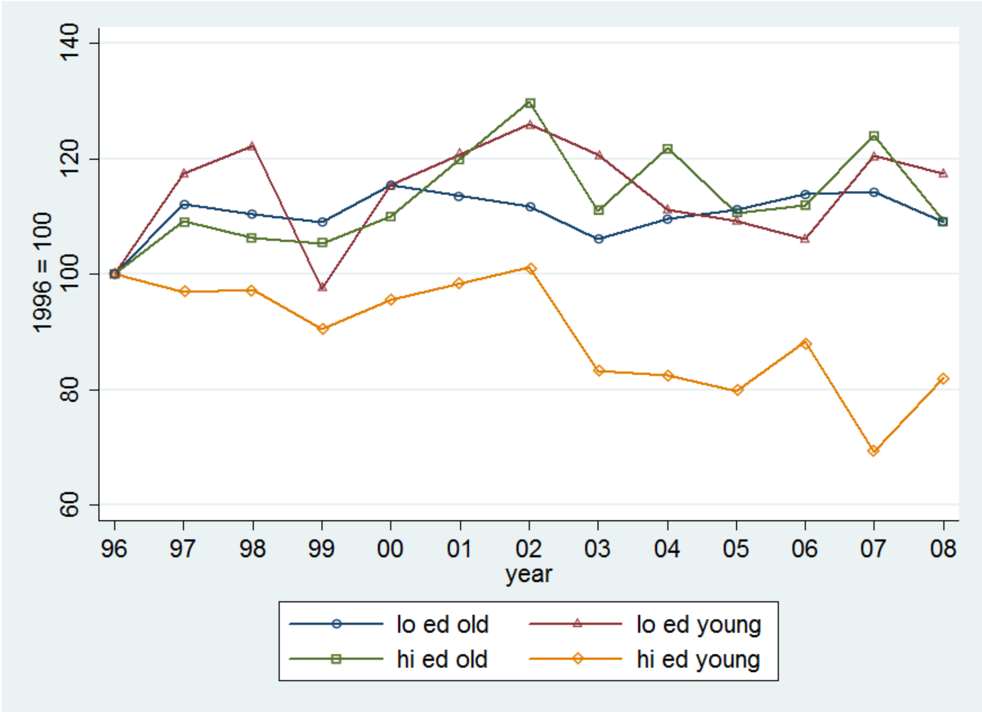

Figure 3.9: Evolution of preference-sensitive multidimensional inequality, by subgroup

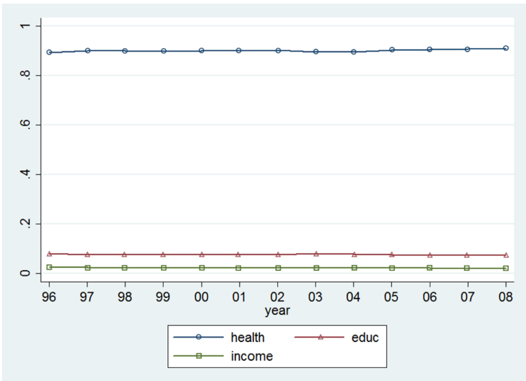

Figure 3.10: Dimension contributions for the older lower educated

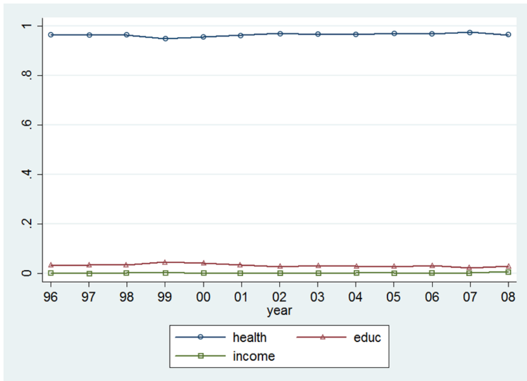

Figure 3.11: Dimension contributions for the younger lower educated

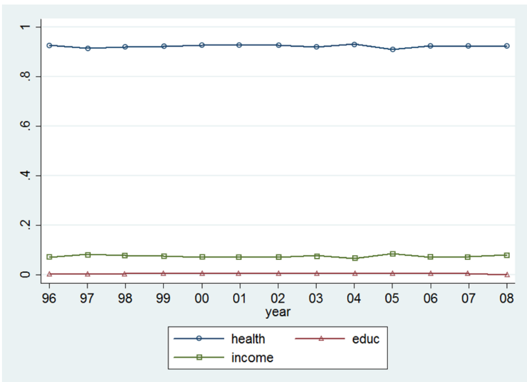

Figure 3.12: Dimension contributions for the older higher educated

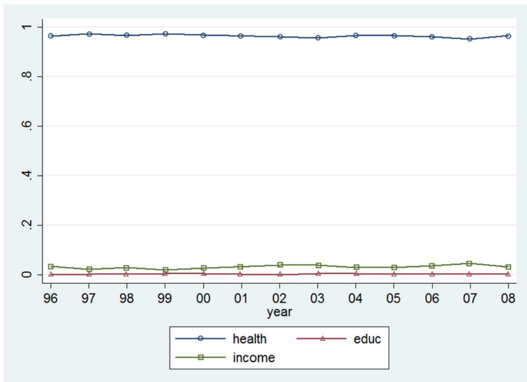

Figure 3.13: Dimension contributions for the younger higher educated

Within each subgroup, preference-sensitive inequality can be further decomposed into its dimension contributions in a similar fashion to Section 3.4.2. This information is given in Figures 3.10-3.13, and two main findings are highlighted here. The first is to do with differences in the contribution from inequality in education attainment. For both of the higher educated groups – that is, those with at least a first university degree – education attainment plays almost no role in terms of contribution to preference-sensitive inequality, conditional on some form of university education having been obtained. Education inequality among the well-educated therefore has comparatively little bearing on preference-sensitive well-being inequality among this group, relative to the other dimensions. In contrast, the lower educated groups have notably higher contributions from education such that education inequality surpasses income inequality in terms of its contribution to preference-sensitive inequality, particularly for the older less educated group. The indication is that variation in education attainment among those with zero to secondary school attainment is also associated with varying fortunes in terms of well-being, which is in contrast to the experience of those with some form of university education. The second finding is to do with differences in the contribution of health inequality between the younger groups and the older groups. The two younger groups experience larger contributions to preference-sensitive well-being from inequalities in health than do the older groups, and this is a reflection of the higher preference for health in the younger preference groups. Both these findings serve to highlight the insights that can be gained by the observation of heterogeneous preferences within such analysis.

The subgroup-by-subgroup decomposition of preference-sensitive inequality by dimension contributions provided some idea of the differing experiences of different types of individuals. Note however that this analysis is only a partial illustration of the full subgroup decomposition given by Equation 3.5 since it does not address the population shares of each group. This information is required to identify the within-group inequality, between-group inequality and to what degree each group is responsible for the within- and between-components relative to the population as a whole. Table 3.2 contains the full multi-decomposition of within- and between-group Shapley contributions, with the within-group contribution further broken down by population-weighted subgroup contribution according to Equation 3.5. Remember that this subgroup decomposition requires equality components to be used, (defined as $1-I(\phi(x,R))$), rather than the opposite inequality interpretation used so far in the empirical analysis). The Shapley *equality* contributions are therefore mostly negative, since the dimensions examined mostly contribute positively to multidimensional *inequality.*[15] This full multi-decomposition is carried out by subgroup and by dimension contributions for the 2008 wave of the BHPS, although the procedure could be repeated for other years.

Subgroup j	**Pop. share**	Dimension k contribution ϕ_k			**Equality**
	p_j	**Health**	**Education**	**Income**	$1+\sum\phi_k$
Older, lower educated	0.5487	-0.1045	-0.0084	-0.0023	**0.8848**
Younger, lower educated	0.2198	-0.0519	-0.0016	-0.0001	**0.9462**
Older, higher educated	0.1653	-0.0565	0.0001	-0.0048	**0.9388**
Younger, higher educated	0.0663	-0.0323	-0.0001	-0.0011	**0.9666**
Within-group		-0.0806	-0.0049	-0.0022	**0.9122**
Between-group		-0.0078	-0.0007	-0.0002	**0.9901**
Overall (*within* $\times$ *between*)					**0.9030**

Table 3.2: Shapley dimension contributions to within- and between-group equality in 2008

The three columns of dimension contributions in Table 3.2 correspond to the rescaled proportional contributions for 2008 in Figures 3.10-3.13 for the different preference types. Inspecting the rest of Table 3.2, it becomes apparent that the preference groups differ widely in population share, with the older less educated group representing over half of the overall population. It is therefore the case that the experience of this subgroup has a large influence on the resulting measures of within-group, between-group and overall preference-sensitive multidimensional equality (and thus inequality). It is also observed that within-group equality is lower than between-group equality and that this is due to greater reductions to the within-component than the between-component from all dimension contributions. Recalling that the subgroups in this analysis are partitioned by preference type, the interpretation here is that well-being inequality between individuals belonging to different preference types is outweighed by well-being inequal-

[15]However, note that Shapley inequality contributions can be negative if the dimension in question contributes to a decrease in multidimensional inequality. See Chantreuil and Trannoy (2013).

ity resulting from the unequal distribution of dimension attainments within preference types. Comparing this to other studies, greater within-group than between-group inequality is not an uncommon finding in the empirical literature on income inequality – Cowell and Fiorio (2011) have found, using Generalised Entropy measures, that most of the income inequality in the United States and Finland in the 1980s and 2004 was due to the within component of inequality when partitioned by education level; and income inequality within cohorts has been found to generally be greater than inequality between cohorts (Easterlin et al., 1993; O'Rand and Henretta, 1999). Similar empirical examples in a multidimensional context are, however, difficult to find.

3.4.4 Accounting for inequality changes over time by demographic vs. distributional factors

Finally, the multiplicative decomposability property for $\alpha = 0$ is applied to examine how changes in preference-sensitive inequality from the beginning to the end of the analysis period have been shaped by changes inside the within- and between-group components. As mentioned in Section 3.3.1, such an analysis is possible since the multiplicative decomposition allows the impact of marginal changes of a given group to be evaluated in terms of its effect on overall *equality*. Following a line of research popularised by Shorrocks and Mookherjee (1982), the analysis comprises of (a) changes in population shares (demographic factors), and (b) changes in inequality components themselves (the distributional factor). In the context of preference types as the population subgroup partitions, the demographic factor becomes particularly interesting in that it reflects changes in the preference structure within the population over time. In the static context of the multi-decomposition considered in Section 3.4.3, the role of population shares in determining overall preference-sensitive inequality was clear. However, in a dynamic context it is less clear what the role of changing population shares, and therefore a changing preference structure, are in altering preference-sensitive well-being inequality. Following Shorrocks and Mookherjee (1982), this question is addressed by means of a shift-share analysis using the specification given in Equation 3.8, adapted from Goerlich-Gisbert et al. (2009).

Referring back to Equation 3.5 for the case of $\alpha = 0$, let us first simplify the notation by letting $E_0(\phi)$ denote the preference-sensitive equality index $E(\phi(x, R))$, $E_{W_0}(\phi)$

denote the within-group equality component $\prod_{j=1}^{J}\left[E\left(\phi(x,R)^{j}\right)\right]^{\omega_j}$, and $E_{B_0}(\phi)$ denote the between-group equality component $E\left(\overline{\phi(\cdot)}^{1},\overline{\phi(\cdot)}^{2},...,\overline{\phi(\cdot)}^{J}\right)$ for $\alpha=0$. Then, transforming the multiplicative decomposition to be additive in logs, and recalling that for $\alpha=0$ the weights ω_j simply reduce to the subgroup population shares so that $\omega_j=p_j$, the changes $\triangle$ in the overall equality index can be decomposed by shift-share analysis in the following way:

$$\begin{aligned}
\Delta\log E_0(\phi) &= \Delta\log E_{W_0}(\phi)+\Delta\log E_{B_0}(\phi)\\
&= \sum_{j=1}^{J}\left(\Delta\left(p_j\log E_0\left(\phi^j\right)\right)\right)+\Delta\log E_{B_0}(\phi)\\
&= \underbrace{\sum_{j=1}^{J}\Delta p_j\left[\frac{\log E_0\left(\phi^j\right)_t+\log E_0\left(\phi^j\right)_{t-1}}{2}\right]}_{\text{within-subgroup population shares}}+\underbrace{\sum_{j=1}^{J}\left[\frac{p_{j,t}+p_{j,t-1}}{2}\right]\Delta\log E_0\left(\phi^j\right)}_{\text{within-subgroup equality}}\\
&\quad+\underbrace{\Delta\log E_{B_0}(\phi)}_{\text{between-subgroup equality}}
\end{aligned}\tag{3.8}$$

In the last line of Equation 3.8, the subscript t denotes a point in time, the first right-hand-side term captures the contribution to within-subgroup equality due to changing population shares (preference structure), the second term captures the contribution to within-subgroup equality due to changing well-being equality levels (distributional factor), and the last term captures changes in between-subgroup well-being equality. For convenience, the columns of Table 3.3 present the empirical values necessary to calculate the changes in the components of equality according to the decomposition in Equation 3.8. Subgroup j gives the preference types. Period t is the last year of the analysis period, 2008, and period $t-1$ is the initial year, 1996. $E_0(\phi)_t$ is the overall equality level in period t. $p_{j,t}$ is the population share of subgroup j in period t. $E_0\left(\phi^j\right)_t$ is the equality level in subgroup j in period t. $\overline{\phi(\cdot)}_t^{j}$ is the mean well-being in subgroup j in period t, and is necessary to calculate the change in the between-subgroup component $E_{B_0}(\phi)$. It can also be verified that for each period t and $t-1$, the within-subgroup and between-subgroup components calculated using the values in Table 3.3 according to Equation 3.5 do indeed multiply to give the overall level of equality $E_0(\phi)_t$.

Table 3.4 presents the three successive decompositions of the change in equality from 1996-2008 given by the three lines of Equation 3.8: 1) the contributions of changes

Subgroup j	**Period**	$E_0\left(\phi\right)_t$	$p_{j,t}$	$E_0\left(\phi^j\right)_t$	$\overline{\phi(\cdot)}_t^j$
Older, lower educated	t	0.0970	0.5487	0.1152	0.5599
	$t-1$	0.0911	0.5422	0.1057	0.5373
Younger, lower educated	t	0.0970	0.2198	0.0538	0.7716
	$t-1$	0.0911	0.2923	0.0459	0.7586
Older, higher educated	t	0.0970	0.1653	0.0612	0.6474
	$t-1$	0.0911	0.1022	0.0561	0.6129
Younger, higher educated	t	0.0970	0.0663	0.0334	0.7709
	$t-1$	0.0911	0.0633	0.0407	0.7393

Table 3.3: Values for the calculation of shift-share components

in the within- and between-group components, 2) the contributions of within-group changes by subgroup, and 3) the contributions of changes in the preference structure and distributional factor. Using the fact that $\%\triangle x \approx 100.\triangle \log(x)$, the changes in log values are presented in the table in terms of percentage changes.

The first row of Table 3.4 gives the change in overall equality from 1996-2008, and the contribution from between-subgroup and within-subgroup equality changes. A deterioration of -0.65% in overall preference-sensitive well-being equality is estimated, with an equality-increasing contribution of 0.27% from between subgroups and an equality-reducing contribution of -0.92% from within subgroups. Adding up the absolute value of the changes from between and within subgroups in the first row, as a proportion of this absolute value around 77% of the change in equality over the period can be attributed to the contribution of within-group changes. The contribution of between-group equality on the other hand, almost 23%, reflects an equalising trend in preference-sensitive well-being between preference groups. These proportional contributions are given in the second row of Table 3.4.

Next, further decomposing the within-group component by looking at the disaggregated contributions of the different preference groups, the finding is that whilst within-group inequalities in the older preference types have tended to have a equality-reducing contribution to changes in well-being, the younger preference types have offset this with slight within-group equalisations in well-being. This is seen from the negative

	Overall	Between subgroups	Within subgroups
Changes 1996-2008	-0.65%	0.27%	-0.92%
Absolute changes		**22.86%**	**77.14%**
Contributions to changes in within-subgroup equality, by subgroup			
Older, lower educated			-0.66%
Younger, lower educated			0.16%
Older, higher educated			-0.45%
Younger, higher educated			0.04%
Total			**-0.92%**
Within-subgroup equality by factor			
Preference structure			-0.10%
Distributional factor			-0.82%
Total			**-0.92%**

Table 3.4: Shift-share subgroup decomposition of equality 1996-2008

percentage changes in contributions from the older subgroups and positive percentage changes from the younger subgroups.

Finally, inspecting the contribution of the different factors to these within-group equality changes, it can be seen that the deterioration in preference-sensitive within-group well-being equality has mostly been driven by changes in the distribution of well-being dimension attainments (the distributional factor), comprising -0.82% of the overall -0.92% change. Having said this, changes to the preference structure through shifts in population shares have worsened this deterioration by -0.10%, and this is a result of a greater proportion of individuals in the older preference groups by the end of the period compared to the beginning – a reflection of the more general long-term trend towards population ageing in advanced economies.

3.5 Conclusion

In this chapter it was shown how the preference index approach, the central proposal introduced in Chapter 2, can be used to analyse and decompose multidimensional inequality in a way that also considers the inequality in subjective preferences between individuals over dimensions of life. The underlying two-step nature of the proposed preference index was explained, along with how it relates to and builds on existing multidimensional frameworks, such as in terms of sensitivity to cumulative deprivations and inequality aversion. Also discussed was why it necessarily departs from conventionally held *Anonymity* and *Majorization* properties for multidimensional inequality indices. It was illustrated that inequality decomposition analysis, which has often used the Generalised Entropy class of measures due to its additive decomposition properties, can also be performed using the normatively-motivated preference index. Decomposition of inequality by subgroup contributions is possible through the multiplicative decomposition proposed by Lasso de la Vega and Urrutia (2005), and decomposition by preference-sensitive dimension-wise contributions is possible by using the Shapley algorithm modified from Shapley (1953) and Shorrocks (2013).

The components of the proposed theoretical approach were combined in an empirical application, and it was shown how the new tools proposed are able to provide insights into the evolution of multidimensional inequality, taking into account dimension and preference interrelationships that existing approaches do not offer. The contribution of health inequality to preference-sensitive inequality was highlighted, a reflection of the strong preference for good health across all preference types. Decomposing results by subgroup, this priority on health also contributed towards the well-educated younger, and therefore healthier, group enjoying higher average preference-sensitive well-being and lower preference-sensitive inequality. Education inequality among the well-educated proved to have comparatively little impact on preference-sensitive inequality, in contrast to higher contributions of education inequality among the lower educated groups, particularly the older less educated group. Another interesting finding was that across the population as a whole, inequality in preferences themselves outweighed inequality in incomes in terms of their contributions to preference-sensitive inequality. Preference-sensitive inequality between individuals in the same preference type was also found to outweigh that of individuals between preference types, with the shift-share analysis showing an increase in within-preference inequality but decrease in between-preference inequality over the period. Finally, the main driver of change in preference-sensitive inequality has been the changing distribution of dimension attainments rather than a changing preference structure in the population.

Conclusion

This concluding chapter provides a summary of the key contributions and findings of the preceding chapters, and the scope for future research possibilities.

Chapter 1 made two key contributions in its empirical investigation of the concept and structure of subjective well-being (SWB). The first contribution was an examination of the existence and time-consistency of a dual SWB structure, comprised of an 'emotional well-being' component and a 'life satisfaction' component. The second contribution was an investigation into the links between these two latent components of SWB and the objective well-being dimensions of health, education and income.

If SWB is to be used to incorporate subjective preferences into a multidimensional measure of well-being, then it is important that the chosen SWB indicator does in fact measure the same phenomenon over time. Using factor analysis of data from the British Household Panel Survey (BHPS), metric longitudinal invariance of the dual SWB structure was found, indicating that the concept was indeed being consistently interpreted by individuals and measured through the same quantitative changes in scores over time. This was an important finding, supporting the later use of longitudinal SWB regression in Chapter 2 to uncover individual and group-level preferences over objective well-being dimensions. Structural equation modelling was then used to examine the responsiveness of the latent 'life satisfaction' and 'emotional well-being' components of SWB to the objective indicators. The 'life satisfaction' component of SWB was found to be more responsive than the 'emotional well-being' component to attainments in the objective indicators examined, again supporting Chapter 2, this time in its speceific use of life satisfaction response data to uncover preferences. Longitudinal analysis of the latent structure of SWB and its relationship to objective well-being indicators, as in Chapter 1, has not been collectively studied in this way before. Since SWB is central to the implementation of the 'preference index approach' proposed and illustrated in

Chapters 2 and 3, this analysis was crucial to designing the index.

In Chapter 2, life satisfaction regressions were used to empirically implement the multidimensional index proposal central to the thesis. This ‘preference index approach’ combined the intuitive idea of an index as a type of summary statistic with recent literature on axiomatic approaches to multidimensional well-being. It is the first proposal and empirical application of a preference-sensitive well-being measure in the form of a multidimensional index.

The proposed index was specified to respect the fact that individuals may have different subjective preferences over dimensions of well-being, while still retaining comparability among the well-being of such individuals. Empirically, key findings using the BHPS data included an interesting result that older individuals in the sample had weaker preferences for health compared to younger individuals. It was also shown how consideration of this kind of preference heterogeneity potentially changes our understanding of well-being in society compared with unidimensional measures such as income and raw SWB measures, as well as other composite measures. The preference index measure was able to reflect strong subjective preferences for good health across all types of individuals, compared to relatively weak preferences for income. It was argued that this somewhat undermines the emphasis in applied literature on income as a yardstick for well-being in society.

Chapter 3 demonstrated the properties of the preference index approach in terms of multidimensional inequality analysis, with the integration of preference heterogeneities. The incorporation of distributional inequality as well as preference inequality, and the ability to quantify these in an overall multidimensional analysis framework was highlighted as the main contribution of this chapter. It was shown, through new ways of using existing analysis tools, that the preference index approach is able to take into account dimension and preference interrelationships that other approaches do not offer.

The contribution of health inequality to preference-sensitive inequality was highlighted, a reflection of the strong preference for good health across all types of preferences analysed. This priority on health contributed towards the well-educated younger, and therefore healthier, group enjoying higher average preference-sensitive well-being and lower preference-sensitive inequality. Education inequality among the well-educated proved to have comparatively little impact on preference-sensitive inequality, in contrast to higher contributions of education inequality among the lower educated. Another interesting finding was that across the population as a whole, inequality in preferences themselves outweighed inequality in incomes in terms of their contributions to preference-sensitive inequality. Shift-share analysis showed that the patterns in preference-sensitive inequality over time seem to have been mostly driven by the changing distribution of dimension attainments themselves, rather than a changing preference structure in the population.

In summary, the overarching theme of these chapters has been to coherently unify a range of well-being theories and methods of analysis from the areas of SWB, composite indicators, welfare theory and multidimensional inequality. This has resulted in the 'preference index approach', a proposal for multidimensional preference-sensitive well-being and inequality measurement, which has provided exciting new insights and analytical possibilities as illustrated in this thesis. As such, the purpose here has been to contribute to the quest for a richer, more meaningful measure of well-being, and ultimately better policy goals. Of course, this is only an incremental contribution in a field of interdisciplinary research with huge depth and breadth. There is no doubt scope for future improvements and expansion of the methodological particulars presented here, especially with regards to measuring the well-being value of education, and further exploration into the other methods of estimating preferences and whether these produce systematic differences in results. It is not unlikely, for example, that preference elicitation through stated preference surveys may produce slightly different findings from the life satisfaction regression approach. These differences would produce additional and valuable insights into how we think about and weigh up important dimensions of life.

Although this work has primarily drawn from areas related to economics broadly defined, there exist rich possibilities of overlapping and complementary research further afield in social policy, psychology, philosophy, political economy and environmental economics. For example, how can the idea of the sustainability of well-being be explicitly investigated using this approach? Such a line of investigation may require thinking about current determinants, or "stocks", affecting future well-being. Even further afield, what we can learn about aspects of our own well-being will surely be of great interest not only to academics and policy-makers, but also to anyone interested in the larger question of what constitutes progress in the context of a meaningful life.

Bibliography

Aaberge, R., Brandolini, A., September 2014. Multidimensional poverty and inequality. Economic Working Paper 976, Bank of Italy, Economic Research and International Relations Area.

Adler, M., 2014. Extended preferences and interpersonal comparisons: A new account. Economics and Philosophy 30 (2), 123–162.

Albuquerque, I., Lima, M. P. d., Matos, M., Figueiredo, C., 2012. Subjective well-being structure: Confirmatory factor analysis in a teachers' Portuguese sample. Social Indicators Research 105 (3), 569–580.

Allan, G. J. B., 1976. Ordinal-scaled variables and multivariate analysis: Comment on Hawkes. American Journal of Sociology 81 (6), 1498–1500.

Anand, S., 1983. Inequality and poverty in Malaysia. Oxford University Press.

Andrews, F., Withey, S., 1976. Social indicators of well-being. Plenum Press, New York.

Atkinson, A., 1970. On the measurement of inequality. Journal of Economic Theory 2 (3), 244–263.

Atkinson, A., Cantillon, B., Marlier, E., Nolan, B., 2002. The EU and social exclusion. Oxford University Press, Oxford.

Bandura, R., 2006. A survey of composite indices measuring country performance: 2006 update. Tech. rep., UNDP Office of Development Studies.

Bartholomew, D. J., Steele, F., Galbraith, J., Moustaki, I., 2008. Analysis of multivariate social science data, 2nd Edition. Chapman and Hall/CRC.

Blackorby, C., Bossert, W., Donaldson, D., 2005. Population issues in social choice theory, welfare economics and ethics. Cambridge University Press, New York.

Blackorby, C., Donaldson, D., 1982. Ratio-scale and translation-scale full interpersonal comparability without domain restrictions: Admissible social-evaluation functions. International Economic Review 23 (2), 249–68.

Blackorby, C., Donaldson, D., Auersperg, M., 1981. A new procedure for the measurement of inequality within and among population subgroups. Canadian Journal of Economics 14 (4), 665–85.

Blanchflower, D., Oswald, A., 2004. Well-being over time in Britain and the USA. Journal of Public Economics 88 (7-8), 1359–1386.

Boarini, R., Johansson, A., Mira d'Ercole, M., 2006. Alternative measures of well-being. OECD Economics Department Working Paper 476, OECD, Paris.

Bosmans, K., Decancq, K., Ooghe, E., 2015. What do normative indices of multidimensional inequality really measure? Journal of Public Economics 130 (C), 94–104.

Bourguignon, F., 1979. Decomposable income inequality measures. Econometrica 47 (4), 901–920.

Bourguignon, F., 1999. Comment to 'Multidimensional approaches to welfare analysis' by Maasoumi, E. In: Silber, J. (Ed.), Handbook of income inequality measurement. Kluwer Academic, Dordrecht, pp. 477–484.

Bradburn, N., Caplovitz, D., 1965. Reports of happiness. Aldine, Chicago.

Brandolini, A., 2008. On applying synthetic indices of multidimensional well-being: Health and income inequalities in selected EU countries. Working Paper 668, Bank of Italy, Economic Research and International Relations Area.

Brief, A., Butcher, A., George, J., Link, K., 1993. Integrating bottom-up and top-down theories of subjective well-being: The case of health. Journal of Personality and Social Psychology 64 (4), 646–653.

Brown, T. A., 2006. Confirmatory factor analysis for applied research. The Guilford Press, New York.

Burkhauser, R., Hérault, N., Jenkins, S., Wilkins, R., 2016. What has been happening to UK income inequality since the mid-1990s? Answers from reconciled and combined household survey and tax return data. Working Paper 2016-03, Institute for Social and Economic Research.

Byrne, B. M., 1994. Testing for the factorial validity, replication, and invariance of a measurement instrument: A paradigmatic application based on the Maslach Burnout Inventory. Multivariate Behavioral Research 29 (3), 289–311.

Byrne, B. M., Stewart, S. M., 2006. The MACS approach to testing for multigroup invariance of a second-order structure: A walk through the process. Structural Equation Modeling 13 (2), 287–321.

Campbell, A., Converse, P., Rodgers, W., 1976. The quality of American life. Russell Sage Foundation, New York.

Chakravarty, S. R., 1990. Ethical social index numbers. Springer-Verlag, Berlin Heidelberg.

Chamberlain, G., 1980. Analysis of covariance with qualitative data. Review of Economic Studies 47 (1), 225–238.

Chanfreau, J., Burchardt, T., September 2008. Equivalence scales: Rationales, uses and assumptions. Scottish Government.
URL `www.gov.scot/resource/doc/933/0079961.pdf`

Chantreuil, F., Trannoy, A., 2013. Inequality decomposition values: The trade-off between marginality and consistency. Journal of Economic Inequality 11 (1), 83–98.

Cheung, G. W., Rensvold, R. B., 1999. Testing factorial invariance across groups: A reconceptualization and proposed new method. Journal of Management 25 (2), 1–27.

Cheung, G. W., Rensvold, R. B., 2002. Evaluating goodness-of-fit indexes for testing measurement invariance. Structural Equation Modeling 9 (2), 233–255.

Collins, A. L., Goldman, N., Rodriquez, G., 2007. Are life satisfaction and optimism protective of health among older adults? Working Paper 2007/02, Princeton University, Office of Population Research.

Conceição, P., Bandura, R., 2008. Measuring subjective wellbeing: A summary review of the literature. UNDP Working Paper, United Nations Development Programme.

Costello, A., Osborne, J., July 2005. Best practices in exploratory factor analysis: Four recommendations for getting the most from your analysis.
URL `http://pareonline.net/getvn.asp?v=10&n=7`

Cowell, F., 1977. Measuring inequality. Philip Allan, Oxford.

Cowell, F., 1980. On the structure of additive inequality measures. Review of Economic Studies 47 (3), 521–531.

Cowell, F., October 1988. Inequality decomposition: Three bad measures. Bulletin of Economic Research 40 (4), 309–312.

Cowell, F., Fiorio, C., 2011. Inequality decompositions – A reconciliation. Journal of Economic Inequality 9 (4), 509–528.

Cowell, F., Kuga, K., 1981. Additivity and the entropy concept: An axiomatic approach to inequality measurement. Journal of Economic Theory 25 (1), 131–143.

Das, T., Parikh, A., June 1981. Decomposition of Atkinson's measure of inequality. Australian Economic Papers 20 (36), 171–178.

Deaton, A., 1979. The distance function in consumer behaviour with applications to index numbers and optimal taxation. Review of Economic Studies 46 (3), 391–405.

Deaton, A., 1980. The measurement of welfare: Theory and practical guidelines. Working Paper 7, Living Standards Measurement Study, World Bank.

Deaton, A., 2008. Income, health, and well-being around the world: Evidence from the Gallup World Poll. Journal of Economic Perspectives 22 (2), 53–72.

Decancq, K., Fleurbaey, M., Maniquet, F., January 2014a. Multidimensional poverty measurement with individual preferences. Research Paper 058-2014, Princeton University, William S. Dietrich II Economic Theory Center.

Decancq, K., Fleurbaey, M., Schokkaert, E., 2014b. Inequality, income, and well-being. CORE Discussion Paper 2014/18, Université Catholique de Louvain, Center for Operations Research and Econometrics (CORE).

Decancq, K., Fleurbaey, M., Schokkaert, E., 2015. Happiness, equivalent incomes and respect for individual preferences. Economica 82 (s1), 1082–1106.

Decancq, K., Lugo, A. M., August 2009. Measuring inequality of well-being with a correlation-sensitive multidimensional Gini index. Working Paper 124, ECINEQ, Society for the Study of Economic Inequality.

Decancq, K., Neumann, D., 2016. Does the choice of well-being measure matter empirically? In: Adler, M., Fleurbaey, M. (Eds.), Oxford Handbook on Well-being and Public Policy. Oxford University Press.

Diener, E., 1994. Assessing subjective well-being: Progress and opportunities. Social Indicators Research 31 (2), 103–157.

Diener, E., Oishi, S., Lucas, R. E., 2009. Subjective well-being: The science of happiness and life satisfaction. In: Snyder, C. R., Lopez, S. J. (Eds.), Oxford handbook of positive psychology. Oxford University Press, New York.

Diener, E., Suh, E., 1997. Measuring quality of life: Economic, social, and subjective indicators. Social Indicators Research 40 (1), 189–216.

Diener, E., Suh, E., Lucas, R., Smith, H., 1999. Subjective well-being: Three decades of progress. Psychological Bulletin 125 (2), 276–302.

Dietz, S., Hepburn, C., Stern, N., 2007. Arguments for a better world: Essays in honour of Amartya Sen. In: Basu, K., Kanbur, R. (Eds.), Economics, ethics and climate change. Vol. 2: Society, institutions and development. Oxford University Press, Oxford.

Dolan, P., Peasgood, T., White, M., 2008. Do we really know what makes us happy? A review of the economic literature on the factors associated with subjective well-being. Journal of Economic Psychology 29 (1), 94–122.

Dolan, P., Tsuchiya, A., 2009. The social welfare function and individual responsibility: Some theoretical issues and empirical evidence from health. Journal of Health Economics 28 (1), 210–220.

Easterlin, R. A., Macunovich, D. J., Crimmins, E. M., 1993. Economic status of the young and the old in the working-age population, 1964 and 1987. In: Bengston,

V. L., Achenbaum, W. A. (Eds.), The changing contract across generations. Sage Publications, Newbury Park, California, pp. 67–86.

Ferrante, F., 2007. Does human capital bring life satisfaction? Human capital, expectations and life satisfaction. Mimeo, faculty of Economics, University of Cassino.

Ferrer-i Carbonell, A., Frijters, P., 2004. How important is methodology for the estimates of the determinants of happiness? Economic Journal 114 (497), 641–659.

Fleurbaey, M., 2005. Health, wealth, and fairness. Journal of Public Economic Theory 7 (2), 253–284.

Fleurbaey, M., August 2008. Beyond GDP: Is there progress in the measurement of individual well-being and social welfare? Working paper, Commission on the Measurement of Economic Performance and Social Progress.

Fleurbaey, M., 2011. Willingness-to-pay and the equivalence approach. Revue D'Économie Politique 121 (1), 35–58.

Fleurbaey, M., Blanchet, D., 2013. Beyond GDP: Measuring welfare and assessing sustainability. Oxford University Press, New York.

Fleurbaey, M., Maniquet, F., 2011. A theory of fairness and social welfare. Cambridge University Press, New York.

Fleurbaey, M., Schokkaert, E., May 2012. Behavioral fair social choice. Working Paper 2012-012, University of Chicago, Human Capital and Economic Opportunity Working Group.

Fleurbaey, M., Tadenuma, K., 2014. Universal social orderings: An integrated theory of policy evaluation, inter-society comparisons, and interpersonal comparisons. Review of Economic Studies 81 (3), 1071–1101.

Foster, J., Shneyerov, A., 2000. Path independent inequality measures. Journal of Economic Theory 91 (2), 199–222.

Frey, B. S., Stutzer, A., 2002. Happiness and Economics. Princeton University Press, Princeton.

Frijters, P., Geishecker, I., Haisken-DeNew, J., Shields, M. A., 2006. Can the large swings in Russian life satisfaction be explained by ups and downs in real incomes? Scandinavian Journal of Economics 108 (3), 433–458.

Gajdos, T., Weymark, J., 2005. Multidimensional generalized Gini indices. Economic Theory 26 (3), 471–496.

Gibbons, S., Machin, S., Silva, O., 2009. Valuing school quality using boundary discontinuities. SERC Discussion Paper 18, SERC.

Goerlich-Gisbert, F., Lasso de la Vega, C., Urrutia, A., 2009. The 'extended' Atkinson family and changes in the expenditure distribution. Spain 1973/74-2003. Journal of Income Distribution 18 (1), 20–41.

Graham, C., 2011. The pursuit of happiness: An economy of well-being. Brookings Institution Press.

Graham, C., Eggers, A., Sukhtankar, S., 2004. Does happiness pay? An exploration based on panel data from Russia. Journal of Economic Behavior & Organization 55 (3), 319–342.

Gregorich, S. E., 2006. Do self-report instruments allow meaningful comparisons across diverse population groups? Testing measurement invariance using confirmatory factor analysis frameworks. Medical Care 44 (Suppl 3), S78–S94.

Harwell, M. R., Gatti, G. G., 2001. Rescaling ordinal data to interval data in educational research. Review of Educational Research 71 (1), 105–131.

Helliwell, J., 2003. How's life? Combining individual and national variables to explain subjective well-being. Economic Modelling 20 (2), 331–360.

Helliwell, J., Layard, R., Sachs, J. (Eds.), 2015. World Happiness Report 2015. New York: Sustainable Development Solutions Network.

HM Treasury, 2011. The Green Book: Appraisal and evaluation in central government.

Horsman, J., Furlong, W., Feeny, D., Torrance, G., 2003. The Health Utilities Index (HUI©): Concepts, measurement properties and applications. Health and Quality

of Life Outcomes 1 (54), date accessed: 3 May 2014.
URL `http://www.hqlo.com/content/1/1/54`

Jenkins, S., 2010. The British Household Panel Survey and its income data. Working Paper 2010-33, Institute for Social and Economic Research.

Jones, A. M., Schurer, S., 2011. How does heterogeneity shape the socioeconomic gradient in health satisfaction? Journal of Applied Econometrics 26 (4), 547–579.

Jones, C., Klenow, P., September 2010. Beyond GDP? Welfare across countries and time. NBER Working Paper 16352, National Bureau of Economic Research.

Kahneman, D., Deaton, A., 2010. High income improves evaluation of life but not emotional well-being. Proceedings of The National Academy of Sciences of the United States of America 107 (38), 16489–16493.

Kline, R. B., 2010. Principles and practice of structural equation modelling, 3rd Edition. Guildford Press, New York.

Kolm, S.-C., 1969. The optimal production of social justice. In: Margolis, J., Guitton, H. (Eds.), Public economics: An analysis of public production and consumption and their relations to the private sector. Plagrave Macmillan.

Kolm, S.-C., 1977. Multidimensional egalitarianisms. Quarterly Journal of Economics 91 (1), 1–13.

Labovitz, S., 1970. The assignment of numbers to rank order categories. American Sociological Review 35 (3), 515–24.

Lasso de la Vega, C., Urrutia, A., 2005. Path independent multiplicatively decomposable inequality measures. Investigaciones Económicas 29 (2), 379–387.

Lasso de la Vega, C., Urrutia, A., 2008. The ‘extended’ Atkinson family: The class of multiplicatively decomposable inequality measures, and some new graphical procedures for analysts. Journal of Economic Inequality 6 (2), 211–225.

Layard, R., 2005. Happiness: Lessons from a new science. Penguin Books.

Lerman, R., Yitzhaki, S., 1985. Income inequality effects by income source: A new approach and applications to the United States. Review of Economics and Statistics 67 (1), 151–156.

Lucas, R., Diener, E., Suh, E., 1996. Discriminant validity of well-being measures. Journal of Personality and Social Psychology 71 (3), 616–628.

Lugo, M. A., 2005. Comparing multidimensional indices of inequality: Methods and application. Working Paper 14, ECINEQ, Society for the Study of Economic Inequality.

Maasoumi, E., 1986. The measurement and decomposition of multidimensional well-being. Econometrica 54 (4), 991–997.

MacKerron, G., 2012. Happiness economics from 35,000 feet. Journal of Economic Surveys 26 (4), 705–735.

McGillivray, M., White, H., 1993. Measuring development? The UNDP's Human Development Index. Journal of International Development 5 (2), 183–192.

Meredith, W., 1993. Measurement invariance, factor analysis and factorial invariance. Psychometrika 58 (4), 525–543.

Michalos, A., 2008. Education, happiness and well-being. Social Indicators Research 87 (3), 347–366.

Mookherjee, D., Shorrocks, A., 1982. A decomposition analysis of the trend in UK income inequality. Economic Journal 92 (368), 886–902.

Morduch, J., Sicular, T., 2002. Rethinking inequality decomposition, with evidence from rural China. Economic Journal 112 (476), 93–106.

Moses, L. E., Emerson, J. D., Hosseini, H., 1984. Analyzing data from ordered categories. New England Journal of Medicine 311 (7), 442–8.

Noorbakhsh, F., 1998. The Human Development Index: Some technical issues and alternative indices. Journal of International Development 10 (5), 589–605.

O'Brien, B., Viramontes, J. L., 1994. Willingness to pay: A valid and reliable measure of health state preference? Medical Decision Making 14 (3), 289–97.

OECD, 2013a. Framework for integrated analysis. In: OECD framework for statistics on the distribution of household income, consumption and wealth. OECD, Paris.

OECD, 2013b. OECD guidelines on measuring subjective well-being. OECD, Paris.

O'Rand, A. M., Henretta, J. C., 1999. Age and inequality: Diverse pathways through later life. Westview Press, Boulder, Colorado.

Oswald, A., Powdthavee, N., 2007. Death, happiness, and the calculation of compensatory damages. IZA Discussion Paper 3159, Institute for the Study of Labour, Germany.

Pazner, E. A., Schmeidler, D., 1978. Egalitarian equivalent allocations: A new concept of economic equity. Quarterly Journal of Economics 92 (4), 671–687.

Pyatt, G., June 1976. Disaggregation of Gini coefficients. Economic Journal 86 (342), 243–255.

Ravallion, M., 2011. The Human Development Index: A response to Klugman, Rodriguez and Choi. Journal of Economic Inequality 9 (3), 475–478.

Ravallion, M., 2012. Troubling tradeoffs in the Human Development Index. Journal of Development Economics 99 (2), 201–209.

Reed Johnson, F., Fries, E. E., Spencer Banzhaf, H., 1997. Valuing morbidity: An integration of the willingness-to-pay and health-status index literatures. Journal of Health Economics 16 (6), 641–665.

Rietveld, T., van Hout, R., 1993. Statistical techniques for the study of language and language behaviour. Mouton de Gruyter, Berlin.

RT News, February 2011. What it means to get old in Russia. Date accessed: 4 February 2016.
URL `https://www.rt.com/news/old-russia-elderly-people/`

Sagar, A., Najam, A., 1998. The Human Development Index: A critical review. Ecological Economics 25 (3), 249–264.

Samuelson, P. A., 1974. Complementarity – An essay on the 40th anniversary of the Hicks-Allen revolution in demand theory. Journal of Economic Literature 12 (4), 1255–1289.

Schimmack, U., Oishi, S., 2005. The influence of chronically and temporarily accessible information on life satisfaction judgments. Journal of Personality and Social Psychology 89 (3), 395–406.

Sebates, R., Hammond, C., 2008. The impact of lifelong learning on happiness and well-being.

Sen, A., 1973. On economic inequality. Oxford University Press.

Sen, A., 1985. Commodities and capabilities. North-Holland, Amsterdam.

Sen, A., 1999. Development as freedom. Oxford University Press.

Shapley, L., 1953. A value for n-person games. In: Kuhn, H., Tucker, A. (Eds.), Contributions to the theory of games. Vol. II. Princeton University Press, pp. 307–317.

Shorrocks, A., 1980. The class of additively decomposable inequality measures. Econometrica 48 (3), 613–625.

Shorrocks, A., 1982. Inequality decomposition by factor components. Econometrica 50 (1), 193–211.

Shorrocks, A., 2013. Decomposition procedures for distributional analysis: A unified framework based on the Shapley value. Journal of Economic Inequality 11 (1), 99–126.

Shorrocks, A., Mookherjee, D., 1982. A decomposition analysis of the trend in UK income inequality. Economic Journal 92 (368), 886–902.

Smith, J., Borchelt, M., Maier, H., Jopp, D., 2002. Health and well-being in the young old and oldest old. Journal of Social Issues 58 (4), 715–732.

Stanton, E., February 2007. The Human Development Index: A history. Working Paper 127, University of Massachusetts Amherst, Political Economy Research Institute.

StataCorp LP, 2015. BIC note – Calculating and interpreting BIC.

Stutzer, A., 2004. The role of income aspirations in individual happiness. Journal of Economic Behavior & Organization 54 (1), 89–109.

Theil, H., Sorooshian, C., 1979. Components of the change in regional inequality. Economics Letters 4 (2), 191–193.

Tsui, K. Y., 1995. Multidimensional generalizations of the relative and absolute inequality indices: The Atkinson-Kolm-Sen approach. Journal of Economic Theory 67 (1), 251–265.

Tsui, K. Y., 1999. Multidimensional inequality and multidimensional Generalized Entropy measures: An axiomatic derivation. Social Choice and Welfare 16 (1), 145–157.

UNDP, 2013. Human Development Report 2013. The rise of the South: Human Progress in a Diverse World. United Nations Development Programme.

van Doorslaer, E., Jones, A. M., 2003. Inequalities in self-reported health: Validation of a new approach to measurement. Journal of Health Economics 22 (1), 61–87.

Veenhoven, R., 1984. Conditions of happiness. The Reidel Publishing Company, Dordrecht.

Veenhoven, R., 1996. Developments in satisfaction research. Social Indicators Research 37 (1), 1–46.

Widaman, K., Ferrer, E., Conger, R., 2010. Factorial invariance within longitudinal structural equation models: Measuring the same construct across time. Child Development Perspectives 4 (1), 10–18.

Witter, D., Okun, M., Stock, W., Haring, M., 1984. Education and subjective wellbeing: A meta-analysis. Educational Evaluation and Policy Analysis 6 (2), 165–173.

World Bank, 2005. World Development Report 2006: Equity and development.

World Bank, 2016. World Development Indicators. Date accessed: 12 February 2016. URL `http://data.worldbank.org/indicator/NY.GDP.PCAP.CD/countries/EG?display=graph`

Yalcin, I., Amemiya, Y., 2001. Nonlinear factor analysis as a statistical method. Statistical Science 16 (3), 275–294.

Yang, L., 2014. An inventory of composite measures of human progress. Occasional paper on methodology, UNDP Human Development Report Office.
URL http://hdr.undp.org/sites/default/files/inventory_report_working_paper.pdf