KU LEUVEN

DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES
DPS16.10

JUNE 2016

Well-being inequality and preference heterogeneity

Koen DECANCQ, Marc FLEURBAEY and Erik SCHOKKAERT

Public Economics

Faculty of Economics
and Business

Well-being Inequality and Preference Heterogeneity*

Koen Decancq[†] Marc Fleurbaey[‡] Erik Schokkaert[§]

June 8, 2016

Abstract

Standard measures of multidimensional inequality (implicitly) assume common preferences for all individuals and, hence, are not sensitive to preference heterogeneity among the members of society. In this paper, we measure the inequality of the distribution of equivalent incomes, which is a preference-sensitive multidimensional well-being measure. To quantify the contribution of preference heterogeneity to well-being inequality, we use a decomposition method that calculates well-being inequality in different counterfactual distributions. We focus on four sources of well-being inequality: the correlation between outcomes and preferences, the preference heterogeneity, the correlation between the outcome dimensions, and the inequality within each of the outcome dimensions. We find that preference heterogeneity accounts for a considerable part of overall well-being inequality in Russia for the period of 1995 to 2005.

Keywords: Well-being Inequality, Equivalent Incomes, Preference Heterogeneity.

JEL classification: D60, D71, I31.

*We thank the participants of a presentation in Leuven for helpful comments.

[†]Centre for Social Policy (University of Antwerp), CORE (Université catholique de Louvain-la-Neuve), and Center for Health and Wellbeing (Princeton University).

[‡]Center for Health and Wellbeing (Princeton University).

[§]Department of Economics (KULeuven) and CORE (Université catholique de Louvain-la-Neuve).

1 Introduction

It is now widely accepted that a concern for inequality in society should go beyond an exclusive focus on the income distribution and should also consider the distribution of other dimensions that make life go well (see Stiglitz et al. (2009)). This immediately raises the issue of how to aggregate the different life dimensions into a single measure of well-being inequality. Should one take individual preferences into account in this aggregation procedure? If one decides to use a preference-sensitive measure of multidimensional well-being inequality, how large is the contribution of preference heterogeneity to overall inequality?

The standard approach to multidimensional inequality measurement generalizes the Pigou-Dalton transfer principle to more dimensions and then directly imposes it in the multidimensional space of outcomes.[1] Although it is not always made explicit, most existing multidimensional inequality measures perform a two-step aggregation with one aggregation across dimensions and another across individuals. There are two possible sequences to perform these two aggregations and, in general, both sequences lead to different results (see Kolm (1977), Dutta et al. (2003) and Decancq and Lugo (2012)). In the first sequence, one first aggregates across the individuals in each dimension and then across the dimensions (see, e.g., Gajdos and Weymark (2005)). From a normative point of view, this procedure has a crucial drawback. It does not capture the cumulative deprivation that occurs if the positions of the individuals across the different dimensions are correlated. This problem can be solved if one follows the second sequence, in which one first aggregates the outcomes across the dimensions of well-being into a measure of well-being for each individual and then aggregates the well-being measures across individuals. However, the specification of the individual well-being measure that is used in the first step is typically determined by axioms that are formulated over the entire aggregation process and does not necessarily relate to the literature on the measurement of individual well-being (see Decancq et al. (2015b)).

A crucial question about these standard multidimensional inequality measures is whether they respect individual preferences over the life dimensions and their heterogeneity. In fact, it has been shown that they do not and even that they cannot. There is a deep conflict between respecting the multidimensional Pigou-Dalton transfer principle and respecting individual preferences (Fleurbaey and

[1] See Weymark (2006), Aaberge and Brandolini (2015), and Chakravarty and Lugo (2016) for overviews of the literature on the measurement of multidimensional inequality.

Trannoy (2003)). We briefly explain the issue in Section 2. This impossibility result brings the literature to a crossroad. One route is to keep the multidimensional Pigou-Dalton transfer principle and, consequently, to neglect individual preferences and their heterogeneity. This is the route taken by the standard approach to multidimensional inequality measurement. Alternatively, one takes preferences seriously and calculates the inequality in the distribution of a preference-sensitive well-being measure. This route leads to inequality measures that do not satisfy the multidimensional Pigou-Dalton transfer principle, but a unidimensional transfer principle in the space of well-being measures. This is the route that we explore in this paper. As an interpersonally comparable measure of well-being we use the so-called equivalent income, which we also introduce in the second section.[2]

In the empirical part of this paper we first measure the inequality in equivalent incomes in the Russian Federation between 1995 and 2005. To compute equivalent incomes and the inequality in their distribution, we estimate in Section 3 the - potentially heterogeneous - preferences of the respondents over their expenditures, health, housing quality, unemployment, and wage arrears on the basis of a life satisfaction equation.[3]

We are particularly interested in measuring the empirical relevance of preference heterogeneity on well-being inequality. In Section 4 - the core of the paper - we therefore construct various counterfactual distributions to decompose the inequality in equivalent incomes into four components: the correlation between outcomes and preferences, the preference heterogeneity, the correlation between the outcome dimensions, and the inequality within each of the outcome dimensions. We find that, along with inequality in the expenditure and health dimension, preference heterogeneity accounts for a considerable part of well-being inequality.

Section 5 confirms the importance of preference heterogeneity through a decomposition of well-being inequality within and between population subgroups with the same preferences. Section 6 discusses how multidimensional dominance approaches relate to our measure of well-being inequality and how they tackle (or do not tackle) preference heterogeneity. In Section 7, we conclude and briefly

[2]Fleurbaey and Blanchet (2013) and Decancq et al. (2015a) discuss the axiomatic underpinnings of the equivalent income measure.

[3]We use panel data from the Russia Longitudinal Monitoring Survey (RLMS-HSE) between 1995 and 2005. This data set has also been used to compute equivalent incomes by Decancq et al. (2015a). Compared to that paper, we include two additional periods in the analysis and use a more flexible specification of the life satisfaction equation.

discuss the normative implications of our findings.

2 Measuring well-being inequality and respecting preferences

Let there be a society of $n > 1$ individuals. The outcome vector $\ell_i = (\ell_i^1, \ell_i^2, \ldots, \ell_i^m)$ of each individual i contains her outcomes in the $m > 1$ dimensions of life. We assume that the first dimension ℓ_i^1 can be interpreted as "income" and the remaining $m-1$ dimensions as "non-income dimensions". Let L denote the $(n \times m)$ outcome matrix, of which each cell ℓ_i^j represents the outcome of individual i in dimension j. The column with the outcomes of all individuals for the $j-$th dimension is denoted ℓ^j.

We assume that each person i has a well-behaved preference ordering R_i over the set of her outcome vectors. We interpret these preferences as the well-considered judgements of the individual about what she considers "a good life".[4] The corresponding strict preference and indifference ordering are denoted P_i and I_i. We model the preference of each individual $R_i = R(a_i)$ as a function of a preference vector of k individual parameters $a_i = \left(a_i^1, a_i^2, \ldots, a_i^k\right)$. Let A denote the $(n \times k)$ preference matrix which contains all n preference vectors in the society.

We are interested in a measure of well-being inequality in the society $I(L, A)$ that uses the outcome matrix L and the preference matrix A as its arguments. We follow the so-called two-step procedure to measure well-being inequality. In the first step of this approach, a well-being measure $WB(\ell_i, a_i)$ is computed for each individual, and then in the second step a standard one-dimensional inequality index is applied to the well-being indices of the first step:

$$I(L, A) = I\left(WB\left(\ell_1, a_1\right), \ldots, WB\left(\ell_n, a_n\right)\right). \tag{1}$$

The standard approach to multidimensional inequality measurement, on the contrary, neglects the information contained in the preference matrix A and uses a common well-being measure for all individuals, which depends on a vector of parameters a, but not on their individual preferences. It can therefore be written

[4] As we know from the booming literature on behavioural economics, these well-considered judgements are not necessarily revealed in choice behaviour.

as

$$I(L) = I(WB(\ell_1, a), \ldots, WB(\ell_n, a)), \tag{2}$$

with different specific proposals corresponding to specific choices of $I(\cdot)$ and of the well-being measure $WB(\cdot, \cdot)$.[5] In the rest of this section, we first explain why the simplification embodied in (2) is unavoidable if one wants to respect the multidimensional Pigou-Dalton transfer principle, and then we introduce the equivalent income measure as one specific preference-sensitive proposal to measure well-being.

2.1 The impossibility of a Paretian egalitarian

A natural generalization of the Pigou-Dalton transfer principle into a multidimensional framework is the following (see, e.g., Fleurbaey and Maniquet (2011)):

Multidimensional Pigou-Dalton Transfer Principle. *(L', A') is strictly better than (L, A), if for all individuals $k \neq i, j$ we have that $\ell_k = \ell'_k$ and for individuals i and j we have that for $\delta \in \mathbb{R}^m_+ \setminus \{0\}$, $\ell'_i = \ell_i + \delta \leq \ell_j - \delta = \ell'_j$.*[6]

A situation is preferred to another situation if a positive bundle δ is transferred from a donor whose outcomes are at least as good in all dimensions of life as the receiver.[7]

The idea of respecting preferences and their heterogeneity can be expressed by the Weak Pareto Principle:

Weak Pareto Principle. *(L', A) is strictly better than (L, A), if for all individuals i we have that $\ell'_i \, P(a_i) \, \ell_i$.*

The Weak Pareto Principle and the Multidimensional Pigou-Dalton Transfer Principle conflict as soon as at least two individuals have different preferences.

[5] Maasoumi (1986) has proposed a two-step multidimensional generalized entropy inequality measure, for instance. Bosmans et al. (2015) interpret a normative two-step inequality measure as a measure of the social welfare loss due to the suboptimal distribution of outcomes after removing the social welfare loss due to its inefficiency.

[6] Let $<, \leq$, and $\ll$ denote the standard vector inequalities.

[7] Lasso de la Vega et al. (2010) derive a class of multidimensional inequality measures consistent with this version of the Multidimensional Pigou-Dalton Transfer Principle. However, in the axiomatic literature on multidimensional inequality it is more common to consider transfers where the transferred bundle is a fraction of the difference between the outcome vectors of the donor and recipient of the transfer, and to drop the restriction that the outcomes of the donor should be at least as good as the outcomes of the recipient in all dimensions (see Weymark (2006), for instance). These modifications of the multidimensional Pigou-Dalton Transfer Principle do not change the impossibility result discussed in this section, however.

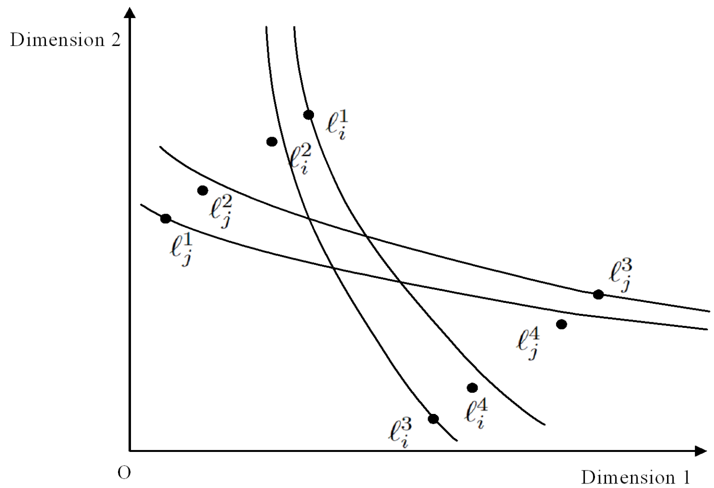

Figure 1: The impossibility of a Paretian egalitarian

Figure 1 illustrates this conflict (see Fleurbaey and Trannoy (2003) and Fleurbaey and Maniquet (2011)). According to the Weak Pareto Principle, distribution matrix L^1 is strictly better than L^4, because for all individuals the outcome vector in L^4 is below the indifference curve containing their outcome vector in distribution matrix L^1. For the same reason, L^3 is strictly better than L^2. According to the Multidimensional Pigou-Dalton Transfer Principle L^2 is strictly better than L^1, and L^4 is strictly better than L^3. Combining these judgements creates a cycle.

In fact, the conflict between these two principles is intuitive. The Weak Pareto Principle requires that individual preferences are respected, whereas the Multidimensional Pigou-Dalton Transfer Principle advocates equalizing transfers irrespective of individual preferences. It only uses information on the outcome matrix L and neglects all information concerning the preference matrix A. If one assumes that inequality can be meaningfully measured using only the outcome information in L, one implicitly imposes that the social evaluation is anonymous in the space of outcomes, i.e., indifferent between all permutations of the individual outcome vectors (Kolm, 1977). As a result, the well-being measures used to aggregate across dimensions must be identical for all individuals. This is the assumption that leads to the move from eq. (1) to eq. (2). Alternatively, if one wants to respect preferences and keep the more general framework of eq. (1), one must choose a particular interpersonally-comparable representation of the preference ordering R_i for each individual as a well-being measure. We now turn to that issue.

2.2 Inequality in well-being: equivalent incomes

By definition, a well-being measure that respects individual preferences is a utility function that represents the preference ordering, i.e., it satisfies

$$WB(\ell_i, a_i) \geq WB(\ell_i', a_i) \Leftrightarrow \ell_i \; R(a_i) \; \ell_i'. \tag{3}$$

In general, what we would like to capture with a preference-based well-being measure is the extent to which outcomes "match" individual preferences. Taking preferences into account implies that two individuals with different preferences may reach a different level of well-being, even if they are in the same objective situation. Consequently, we have that $I(L, A)$ and $I(L, A')$ may differ. Consider, for instance, two individuals with the same income, both living in a high-quality house, but with a low level of health. While their objective situation is the same, a preference-sensitive measure of well-being should be designed so that the individual who cares (relatively) less about health and (relatively) more about housing will reach a higher level of well-being.

It can be argued, however, that there are situations in which preference differences should not matter for interpersonal well-being comparisons. Let us define $(\widehat{\ell}_i^2, \ldots, \widehat{\ell}_i^m)$ as the vector that contains the optimal value in the non-income dimensions for individual i. With heterogeneous preferences, these optimal values may be different for different individuals. If preferences with respect to life dimension j are monotonic, the optimal value is the highest possible value for that dimension. If individual i prefers to be as healthy as possible, for instance, then her optimal value will be equal to perfect health. For other life dimensions, however, the assumption of preference monotonicity may be less realistic. Consider as an example the number of hours worked, where the optimal value may differ for a typical academic and a typical low-skilled blue-collar worker (Fleurbaey and Blanchet (2013)).

Now consider two individuals who both reach their own optimal outcome level in all non-income dimensions. Our basic assumption for making interpersonal well-being comparisons is that when comparing the well-being level of these two individuals we can restrict ourselves to comparing their incomes, independently of their actual preferences. Why should their preferences matter if they reach their optimal outcome level in all non-income dimensions? As shown in Decancq et al. (2015a), combining this assumption with respect for preferences characterizes the well-being ordering that compares individuals in terms of their equivalent incomes.[8] The equivalent income is formally defined as the solution ℓ_i^{1*} to the equation

$$(\ell_i^{1*}, \widehat{\ell}_i^2, \ldots, \widehat{\ell}_i^m)\, I(a_i)\, (\ell_i^1, \ell_i^2, \ldots, \ell_i^m). \tag{4}$$

In other words, it is the hypothetical level of income that, combined with the optimal outcome level in the non-income dimensions, keeps the individual on the indifference curve corresponding to her actual situation. If preferences are monotonic with respect to income, this equivalent income ℓ_i^{1*} cannot be larger than the actual income level ℓ_i^1 because $(\ell_i^1, \widehat{\ell}_i^2, \ldots, \widehat{\ell}_i^m)\, R(a_i)\, (\ell_i^1, \ell_i^2, \ldots, \ell_i^m)$.

The equivalent income can be interpreted as the income corrected for the loss in well-being associated with a suboptimal outcome level for the non-income dimensions. This is an intuitively attractive way of capturing the idea of multidimensional deprivation, we believe. Moreover, it is conveniently measured in monetary units, which provides a simple and familiar cardinal scale. In the following, we measure well-being by means of equivalent incomes, i.e., we have that $WB(\ell_i, a_i) = \ell_i^{1*}$.

Once we have calculated an equivalent income for all individuals in society, we can implement eq. (1) using any unidimensional inequality measure. In our empirical application we will work with the Generalized Entropy class of inequality measures (see Cowell (2011), and the references therein):

$$GE_\alpha(L, A) = \frac{1}{\alpha(\alpha-1)n} \left[\sum_{i=1}^{n} \left(\frac{WB\,(\ell_i, a_i)}{\mu} \right)^\alpha - 1 \right], \tag{5}$$

where μ is the average equivalent income $\frac{1}{n}\sum_{i=1}^{n} WB\,(\ell_i, a_i)$. The lower the value of the parameter α, the more we focus on the bottom part of the distribution of well-being measures. We will concentrate on the mean logarithmic deviation ($\alpha = 0$) and the Theil-index ($\alpha = 1$) in this paper, because these inequality measures have attractive decomposability properties.

This approach obviously satisfies the Pigou-Dalton Transfer Principle in the space of well-being measures. Given the impossibility result discussed in the

[8] We do not discuss the normative strengths and weaknesses of this proposal here, but we refer the interested reader to Decancq et al. (2015a,b) for a discussion.

previous section, we know that eq. (5) does not satisfy the Multidimensional Pigou-Dalton Transfer Principle. However, the measure does satisfy a Restricted Pigou-Dalton Transfer Principle that only applies to situations involving the optimal outcome levels in the non-income dimensions $(\widehat{\ell}_i^2, \ldots, \widehat{\ell}_i^m)$:

Restricted Pigou-Dalton Transfer Principle. *(L', A') is strictly better than (L, A), if for all individuals $k \neq i, j$ we have that $\ell_k = \ell'_k$ and for individuals i and j with outcomes $\ell_i = (\ell_i^1, \widehat{\ell}_i^2, \ldots, \widehat{\ell}_i^m)$ and $\ell_j = (\ell_j^1, \widehat{\ell}_j^2, \ldots, \widehat{\ell}_j^m)$ we have that for $\delta \in \mathbb{R}_{++}$, $\ell'_i = (\ell_i^1 + \delta, \widehat{\ell}_i^2, \ldots, \widehat{\ell}_i^m), \ell'_j = (\ell_j^1 - \delta, \widehat{\ell}_j^2, \ldots, \widehat{\ell}_j^m)$ with $\ell_i^1 + \delta < \ell_j^1 - \delta$.*

3 Equivalent incomes in Russia between 1995 and 2005

To compute equivalent incomes with real-world data, one needs information about the preferences of the concerned individuals. For that purpose we will exploit the ordinal information that can be derived from a life satisfaction equation (see Decancq et al. (2015a) for a similar procedure). We use data from the nine waves of the Russia Longitudinal Monitoring Survey (RLMS-HSE) between 1995 and 2005.[9] In this period, the Russian economy underwent sharp changes, including a deep financial crisis in August 1998. In the RLMS-HSE, life satisfaction is measured by the question: "*To what extent are you satisfied with your life in general at the present time?*", with answers on an ordinal five point-scale ranging from "not at all satisfied" to "fully satisfied". We first discuss the estimation of the life satisfaction equation and then briefly explain how one can compute equivalent incomes on the basis of these estimates.

3.1 Estimation of the life satisfaction equation

Let us denote the latent variable underlying the life satisfaction responses of individual i in period t by S_{it}^*. We can then specify the life satisfaction equation as follows

$$S_{it}^* = \alpha_i + \gamma_t + \sum_{j=1}^{5} (\beta^j + \mu^{j\prime} D_{it}) \times \Gamma^j(\ell_{it}^j) + \delta' Z_{it} + u_{it}. \quad (6)$$

[9] No data were collected in 1997 and 1999.

This life satisfaction equation includes five life dimensions.[10] The first three life dimensions are measured by continuous variables: ℓ_{it}^{1} denotes real equivalized household expenditures (with the square root of household size as the equivalence scale); ℓ_{it}^{2} denotes health, measured as a composite index of objective disease indicators, using the weights obtained from an ordered logit regression with self-assessed health as the dependent variable; ℓ_{it}^{3} captures housing quality, measured as the predicted value of a hedonic regression of self-reported housing values on a number of housing characteristics (after controlling for regional price differences, a time trend and household size). The final two life dimensions are binary indicators of unemployment (ℓ_{it}^{4}) and "wage arrears" (ℓ_{it}^{5}). The latter indicator captures the phenomenon that wages were often not paid on time in Russia during the late nineties.

We allow for the non-linearity of the life satisfaction equation (and hence for less than perfect substitutability between the dimensions) through a so-called Box-Cox transformation of the continuous dimensions $j = 1, \ldots, 3$ (see Box and Cox (1964)):

$$\Gamma^{j}(\ell_{it}^{j}) = \begin{cases} \left(\left(\ell_{it}^{j} \right)^{\theta^{j}} - 1 \right) / \theta^{j} & \text{when } \theta^{j} \neq 0. \\ \ln\left(\ell_{it}^{j} \right) & \text{when } \theta^{j} = 0. \end{cases}$$

For the other two binary indicators, Γ^{j} is the identify function so that these dimensions are not transformed. The scaling of life satisfaction in eq. (6) is allowed to be influenced by a number of socio-demographic characteristics Z_{it} (education, social status, and marital status) that are introduced together with time dummies γ_t as control variables. Moreover, we include individual fixed effects α_i to control for unobserved individual heterogeneity in time-invariant characteristics including personality traits. To model preference heterogeneity we include interaction effects between the outcomes and four dummy variables contained in D_{it}. These dummies capture whether the respondent is living in a rural area, is young (below the age of 33), is male and has obtained higher education. Finally, u_{it} is a disturbance term.

The scalars β^{j} and θ^{j}, as well as the vectors μ^{j} and δ, are coefficients to be estimated. Since the observed life satisfaction responses are measured on an ordinal scale, we estimate an ordered logit model. We incorporate individual fixed effects into the estimation using the approximation proposed by Jones

[10] More detailed information on the construction of the data can be found in Decancq et al. (2015a) or obtained from the authors on request.

and Schurer (2011) of the method discussed by Ferrer-i-Carbonell and Frijters (2004), Frijters et al. (2004), and Frijters et al. (2006). The three Box-Cox parameters of the continuous variables are chosen on the basis of a grid search to maximize the overall fit of the model. Standard errors are corrected for clustering at the household level.

The estimation results are summarized in Table 1. The results for the full model are shown in the rightmost column (Model 4). The results for the life dimensions and the socio-demographic control variables are in line with what is usually found in the literature. Interesting for our purposes are the interaction coefficients μ^j, as these coefficients capture the heterogeneity in preferences. These coefficients will be treated as the preference parameters a_i in the well-being measure. Even with our restricted set of five life dimensions, many interactions would need to be estimated. We therefore simplify the model by including only the interaction effects that are significant at the 10% level. In fact, not dropping the insignificant coefficients would lead to imprecisely computed equivalent income well-being measures and well-being distributions. The remaining interaction terms in Model 4 can be interpreted easily. Unemployment has a stronger negative effect on the life satisfaction of older, higher educated males. Expenditures are relatively more important for the higher educated in urban areas. Young females give a relatively smaller weight to health. Housing matters more in rural areas, and less for the highly educated respondents. Finally, the Box-Cox parameters of the continuous variables are shown at the bottom of the column and are equal to -0.08, 0.46, and -0.36 (for expenditures, health, and housing quality, respectively).

The other columns in Table 1 contain the results for restricted versions of the full model. Model 3 imposes that $\theta^j = \theta$ for the continuous dimensions. The resulting estimate of θ equals 0.05. This restriction is close to being rejected by a standard likelihood ratio test ($\chi^2(2; 4.32) = 0.115$), and is significant from an economic point of view. Model 2 keeps the differentiated Box-Cox parameters but removes all preference heterogeneity. This restriction is clearly rejected on statistical grounds ($\chi^2(9; 59.24) = 0.000$). Model 1 is the most restricted model without interactions and with the same Box-Cox parameter for the three first life dimensions. It is therefore very close to the common homothetic specification of well-being without preference heterogeneity that is often (implicitly) used in the standard multidimensional inequality measures (as in eq. (2)). This restricted model is strongly rejected with respect to the full Model 4.

	Model 1		Model 2		Model 3		Model 4	
Expenditures	0.255***	(0.0209)	0.255***	(0.0209)	0.259***	(0.0268)	0.259***	(0.0270)
Health	0.477***	(0.0644)	0.711***	(0.0912)	0.394***	(0.0801)	0.636***	(0.117)
Housing	0.182**	(0.0653)	0.247**	(0.0843)	0.186+	(0.0981)	0.242+	(0.129)
Unemployed	-0.446***	(0.0529)	-0.448***	(0.0529)	-0.209+	(0.123)	-0.206+	(0.123)
Wage arrears	-0.297***	(0.0393)	-0.298***	(0.0393)	-0.309***	(0.0394)	-0.311***	(0.0394)
Education	-0.0817***	(0.0163)	-0.0809***	(0.0163)	-0.0996***	(0.0172)	-0.0999***	(0.0173)
High status	0.401***	(0.0883)	0.402***	(0.0883)	0.401***	(0.0883)	0.404***	(0.0882)
Middle status	0.276***	(0.0411)	0.274***	(0.0411)	0.278***	(0.0412)	0.278***	(0.0412)
Married	0.192*	(0.0871)	0.193*	(0.0872)	0.187*	(0.0871)	0.186*	(0.0871)
Living as married	0.0569	(0.0875)	0.0590	(0.0876)	0.0532	(0.0873)	0.0537	(0.0874)
Divorced	-0.266**	(0.0931)	-0.262**	(0.0931)	-0.264**	(0.0930)	-0.261**	(0.0930)
Widowed	-0.255*	(0.103)	-0.260*	(0.103)	-0.276**	(0.103)	-0.282**	(0.103)
1996	-0.0884+	(0.0477)	-0.0877+	(0.0477)	-0.0803+	(0.0477)	-0.0785+	(0.0477)
1998	-0.319***	(0.0540)	-0.318***	(0.0540)	-0.290***	(0.0543)	-0.286***	(0.0544)
2000	0.265***	(0.0506)	0.268***	(0.0506)	0.288***	(0.0508)	0.295***	(0.0509)
2001	0.578***	(0.0507)	0.582***	(0.0506)	0.603***	(0.0510)	0.612***	(0.0510)
2002	1.110***	(0.0524)	1.114***	(0.0524)	1.139***	(0.0529)	1.148***	(0.0530)
2003	0.906***	(0.0528)	0.909***	(0.0528)	0.935***	(0.0535)	0.945***	(0.0535)
2004	1.115***	(0.0550)	1.119***	(0.0549)	1.147***	(0.0557)	1.158***	(0.0558)
2005	1.213***	(0.0564)	1.216***	(0.0563)	1.247***	(0.0572)	1.258***	(0.0572)
Rural × Expenditures					-0.0791+	(0.0442)	-0.0787+	(0.0443)
Higher Educated × Expenditures					0.0314**	(0.0109)	0.0333**	(0.0116)
Young× Health					-0.191	(0.136)	-0.320*	(0.158)
Male× Health					0.349**	(0.127)	0.415*	(0.176)
Rural× House					0.320*	(0.135)	0.387*	(0.170)
Higher Educated × House					-0.172*	(0.0839)	-0.223*	(0.108)
Young× Unemployed					0.202*	(0.0880)	0.195*	(0.0881)
Male× Unemployed					-0.360***	(0.0881)	-0.360***	(0.0881)
Higher Educated × Unemployed					-0.193+	(0.107)	-0.193+	(0.107)
Box-Cox parameter Expenditures	-0.05		-0.08		-0.05		-0.08	
Box-Cox parameter Health	-0.05		0.46		-0.05		0.46	
Box-Cox parameter Housing	-0.05		-0.36		-0.05		-0.36	
N	53873		53873		53873		53873	
pseudo R^2	0.071		0.071		0.072		0.072	
log. likelihood	-21844.5		-21842.0		-21814.5		-21812.4	

Clustered standard errors in parentheses.

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1: Four life satisfaction models: Model 1 and 2 without interactions, Model 1 and 3 with common Box-Cox parameter

We will use the estimates of Model 4 as presented in the right-most column of Table 1 to compute the inequality in equivalent incomes in the next section, except for one simplification. Since the Box-Cox parameter of expenditures is close to zero, we equalize it to zero and work with the logarithm of equivalized expenditures as the relevant monetary life dimension. This brings us closer to the specification that is used in the bulk of the literature on life satisfaction (See Layard et al. (2008), for instance). As will become clear in the following subsection, this restriction also prevents the equivalent incomes from becoming negative, which causes computational problems when calculating the Generalized Entropy inequality measure.

3.2 Computation of the equivalent incomes

Let the satisfaction function S_{it} be the function that maps the outcome vectors and preference parameters to the response to the life satisfaction question by individual i in period t. These answers can be used to estimate preferences and compute equivalent incomes if they are consistent with the preferences of the respondents, i.e., if the following consistency assumption holds:

$$S_{it}(\ell_{it}, a_i) \geq S_{it}(\ell'_{it}, a_i) \Leftrightarrow \ell_{it} \; R(a_i) \; \ell'_{it}. \tag{7}$$

Under this consistency assumption, the life satisfaction function is one possible utility function that provides a representation of the preference ordering of individual i, just as for the equivalent income well-being measure. Equivalent incomes and the life satisfaction function have different ways of attaching a label to the indifference curves, however. In Decancq et al. (2015a), we have argued that equivalent incomes are interpersonally comparable in a normatively attractive way, whereas the life satisfaction functions are not.

To compute equivalent incomes, we first determine for each individual her optimal outcome level in the non-income dimensions $\left(\widehat{\ell}_i^2, \ldots, \widehat{\ell}_i^m\right)$. The estimates presented in Table 1 lead to preferences which are monotonic with respect to all life dimensions, so that the optimal values will be the same maximal value for all respondents. To be precise, they are set at being in perfect health, having a high housing quality[11], not being unemployed, and not suffering from wage arrears.

[11] To avoid the results from being overly sensitive to outliers, we select the 90th percentile value of the estimated housing values.

Under the consistency assumption of eq. (7), we can use the definition of equivalent incomes given by eq. (4) and the econometric specification of the life satisfaction equation given by eq. (6) to write:

$$\begin{aligned} S^*_{it} &= \alpha_i + \gamma_t + (\beta^1 + \mu^{1\prime} D_{it}) \times \ln\left(\ell^1_{it}\right) + \sum_{j=2}^{5} (\beta^j + \mu^{j\prime} D_{it}) \times \Gamma^j(\ell^j_{it}) + \delta' Z_{it} + u_{it} \\ &= \alpha_i + \gamma_t + (\beta^1 + \mu^{1\prime} D_{it}) \times \ln\left(\ell^{1*}_{it}\right) + \sum_{j=2}^{5} (\beta^j + \mu^{j\prime} D_{it}) \times \Gamma^j(\widehat{\ell}^j_{it}) + \delta' Z_{it} + u_{it}, \end{aligned}$$

which yields

$$\ell^{1*}_{it} = \ell^1_{it} \times \exp\left[\sum_{j=2}^{5} \frac{\beta^j + \mu^{j\prime} D_{it}}{\beta^1 + \mu^{1\prime} D_{it}} \times \left(\Gamma^j(\ell^j_{it}) - \Gamma^j(\widehat{\ell}^j_{it})\right)\right]. \tag{8}$$

The shape of the indifference curves as measured by the marginal rate of substitution between the non-income dimensions and the income dimension $(\beta^j + \mu^{j\prime} D_{it})/(\beta^1 + \mu^{1\prime} D_{it})$ is of crucial importance in eq. (8). On the contrary, the conditioning variables Z_{it}, the fixed effects α_i, the time trends γ_t, and the idiosyncratic disturbance term do not appear in eq. (8). These variables only shift the level of reported life satisfaction upwards or downwards, without affecting the marginal rates of substitution between the life dimensions. These shifts can be interpreted as changes in aspirations and expectations, and are considered irrelevant in making well-being comparisons by means of equivalent incomes.[12]

4 Decomposing well-being inequality

Once we have computed the equivalent incomes for all individuals in the sample, we can immediately calculate the inequality $I(L, A)$ or, more specifically, $GE_\alpha(L, A)$ as shown in eq. (5). We now want to analyse how sensitive this measure is to the various components of the measure that have been explained in Section 2. Does preference heterogeneity matter? How important is the issue of cumulative deprivation, i.e., correlation between the outcomes? Does the answer to the latter question depend on whether preference heterogeneity is taken into account or not? We explain in section 4.1 how we simulate different counter-

[12] See Decancq et al. (2015a) for a more extensive discussion of the computation of equivalent incomes.

factual distributions to give an empirical answer to these questions. The results are discussed in section 4.2.

4.1 Construction of counterfactual well-being distributions

The central idea of our approach is to compare the well-being inequality $I(L, A)$ in the sample with the inequality in different counterfactual well-being distributions that are constructed by neutralizing one or more sources of well-being inequality. As noted before, the matrix of preference parameters A contains the estimates of the interaction coefficients from the life satisfaction equation. Given that vector D_{it} contains four dummy variables, we only have 16 different preference groups, and the preference matrix A contains a large number of identical rows. We construct the following four counterfactual matrices for each considered period:

Reshuffled preference matrix $\widetilde{A}$. The matrix $\widetilde{A}$ is a permuted version of the preference matrix A, i.e., $\widetilde{A}=P \cdot A$, where P is an $(n \times n)$ permutation matrix. This operation reshuffles entire vectors of preference parameters a_i across individuals. Each individual is randomly assigned a new preference vector from the sample. Clearly, the resulting preference matrix $\widetilde{A}$ is not unique. In our empirical application, we will therefore generate 200 of these permutation matrices and then provide information about the resulting distribution of the inequality measures.

Equalized preference matrix $\overline{A}$. The matrix $\overline{A}$ is an averaged version of A, i.e., $\overline{A}=Q \cdot A$ for Q the $(n \times n)$ bistochastic matrix with $1/n$ in each cell. Note that the resulting preference ordering $R(\overline{a})$ is in some sense artificial, since it is obtained by averaging the preference parameters and it does not necessarily occur in the sample.

Reshuffled outcome matrix $\widetilde{L}$. The matrix $\widetilde{L}$ is a (dimension-wise) permutation of the outcome matrix L. Each dimension is obtained by $\widetilde{\ell^j}=P^j \cdot \ell^j$ for P^j an $(n \times n)$ permutation matrix for dimension j. We randomly assign to each individual an outcome from the sample. Since the resulting outcome matrix $\widetilde{L}$ is not unique, we will again generate 200 different reshuffled outcome matrices.

Equalized outcome matrix $\overline{L}$. The matrix $\overline{L}$ is a (dimension-wise) averaged version of the outcome matrix L. Each dimension is obtained by $\overline{\ell}^j=Q \cdot \ell^j$ for Q the $(n \times n)$ bistochastic matrix with $1/n$ in each cell. We perform this averaging dimension by dimension. Let $\overline{L}_1$ denote the outcome matrix where only the incomes are equalized. The matrix $\overline{L}_2$ denotes the outcome matrix where income and health are equalized. Similarly, $\overline{L}_3$ denotes the outcome matrix where income, health, and housing are equalized and $\overline{L}_4$ denotes the outcome matrix where income, health, housing, and unemployment are equalized. Finally, let $\overline{L}_5 = \overline{L}$ denote the outcome matrix where all five dimensions are equalized.

For each counterfactual matrix, reshuffling neutralizes the correlation and averaging neutralizes the heterogeneity. With these counterfactual matrices as building blocks, we can construct the first decomposition of well-being inequality, which we will call the "preferences-first" decomposition:

$$\begin{aligned} I(L,A) \quad = \quad & \underbrace{(I(L,A) - I(L,\widetilde{A}))}_{\substack{correlation \\ outcome-pref.}} + \underbrace{(I(L,\widetilde{A}) - I(L,\overline{A}))}_{\substack{preference \\ heterogeneity}} + \\ & \underbrace{(I(L,\overline{A}) - I(\widetilde{L},\overline{A}))}_{\substack{outcome \\ correlation}} + \underbrace{(I(\widetilde{L},\overline{A}) - I(\overline{L},\overline{A}))}_{\substack{outcome \\ inequality}}. \end{aligned} \tag{9}$$

Note that we provide a full decomposition with $I(\overline{L},\overline{A}) = 0$. While it is natural to first neutralize the correlation by reshuffling and then to neutralize inequality by taking averages, there is no *a priori* reason to start by considering the preference matrix A first rather than the outcome matrix L. We will therefore also consider an alternative decomposition of well-being inequality, which we will call the "outcomes-first" decomposition:

$$\begin{aligned} I(L,A) \quad = \quad & \underbrace{(I(L,A) - I(\widetilde{L},A))}_{\substack{outcome \\ correlation}} + \underbrace{(I(\widetilde{L},A) - I(\overline{L},A))}_{\substack{outcome \\ inequality}} + \\ & \underbrace{(I(\overline{L},A) - I(\overline{L},\widetilde{A}))}_{\substack{correlation \\ outcome-pref.}} + \underbrace{(I(\overline{L},\widetilde{A}) - I(\overline{L},\overline{A}))}_{\substack{preference \\ heterogeneity}}. \end{aligned} \tag{10}$$

Note that in the outcomes-first decomposition we have that $I(\overline{L},A) = I(\overline{L},\widetilde{A})$

by construction. When all individuals have the same outcome vector, permuting the preference vectors does not affect inequality. So, in the outcomes-first decomposition there is no effect of the correlation between outcomes and preferences on well-being inequality.

As is common with this type of decomposition, the results are path-dependent. Therefore, the results of the preferences-first and outcomes-first decompositions will be different. Yet, we do not consider this path-dependence problematic since our aim is not to obtain a unique decomposition, but rather to understand the contribution of the different components and the interactions between them. Combining the results from both decompositions will therefore yield useful additional insights.

4.2 Empirical results

Based on the preference estimates in Table 1, we can construct $\widetilde{A}, \overline{A}$, $\widetilde{L}$, and $\overline{L}$ for each considered period. With these building blocks we then construct various counterfactual well-being distributions and compute the corresponding Generalized Entropy inequality measures. The empirical results can be summarized in four figures. Figures 2 and 3 show the results for the preferences-first decompositions for GE_0 and GE_1 respectively. Figures 4 and 5 display the results for the outcomes-first decompositions for the same inequality measures. We now discuss the results for each of the four steps in the sequence of the decomposition.

4.2.1 Correlation between outcomes and preferences

We look first at the effect of the correlation between outcomes and preferences on well-being inequality. This effect can be measured by looking at the first term in the preferences-first decomposition given by eq. (9). This term quantifies the difference between the well-being inequality, computed using the actual outcome and preference matrix $I(L, A)$, and the counterfactual well-being inequality, obtained from the actual outcome matrix and the reshuffled preference matrix $I(L, \widetilde{A})$.

The results are shown in Figures 2 and 3. The dark grey area around the line $I(L, \widetilde{A})$ shows the 95% confidence interval caused by "reshuffling variance", i.e., variance which originates from the non-uniqueness of the reshuffling procedure. This confidence interval is derived from the empirical distribution of the 200

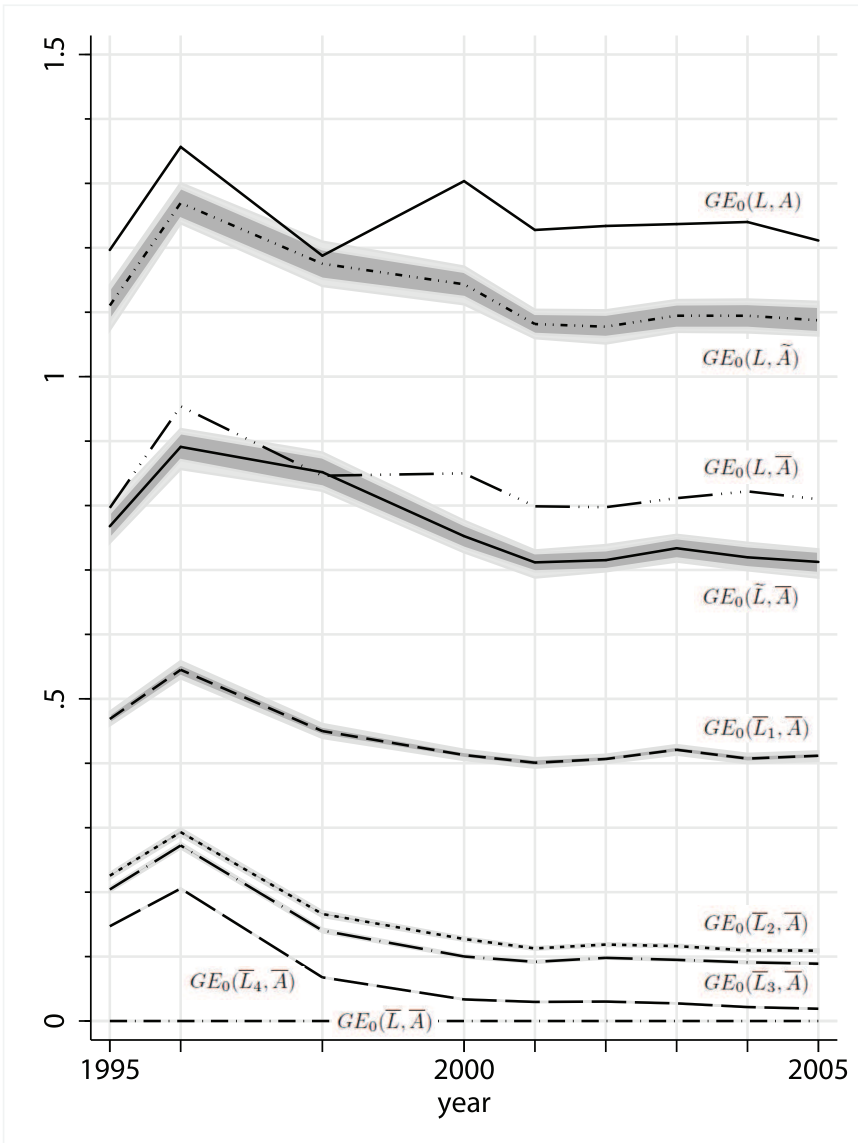

Figure 2: Preferences first: $GE_0(L, A), GE_0(L, \widetilde{A}), GE_0(L, \overline{A})$, $GE_0(\widetilde{L}, \overline{A})$, and $GE_0(\overline{L}, \overline{A})$

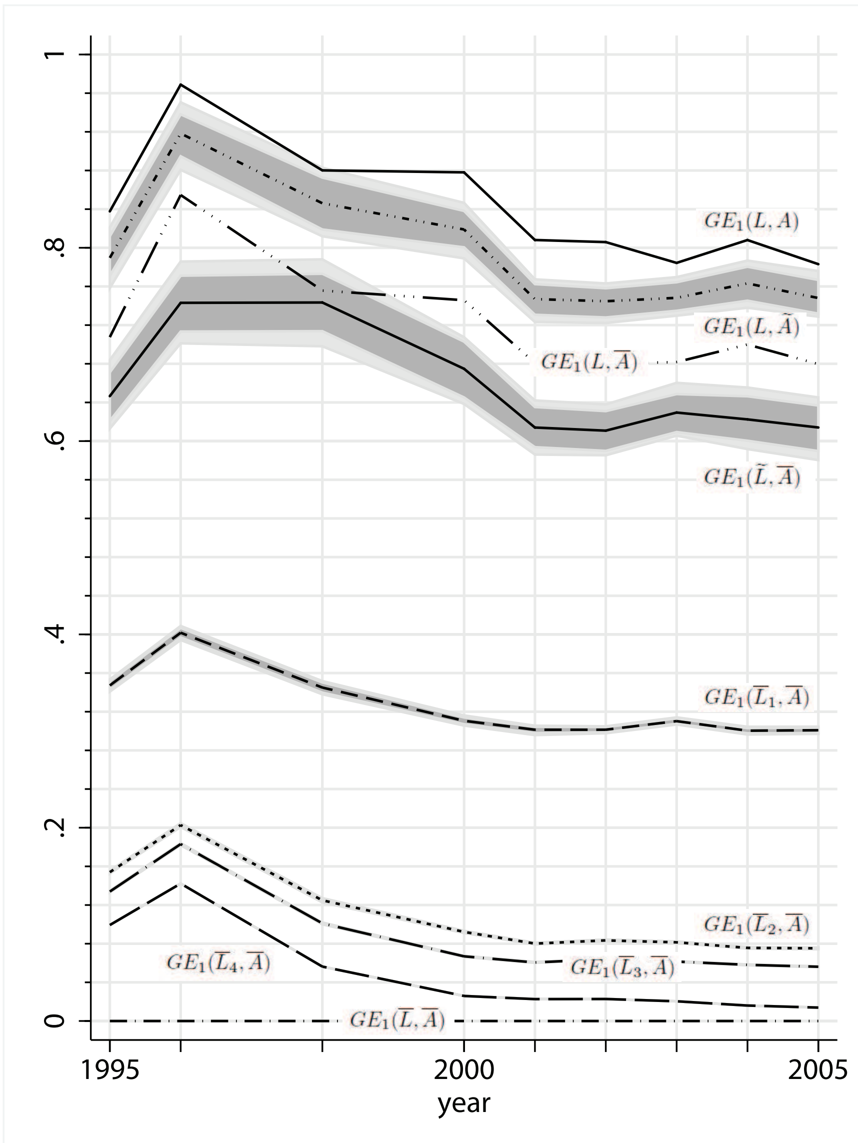

Figure 3: Preferences first: $GE_1(L,A), GE_1(L,\widetilde{A}), GE_1(L,\overline{A})$, $GE_1(\widetilde{L},\overline{A})$, and $GE_1(\overline{L},\overline{A})$

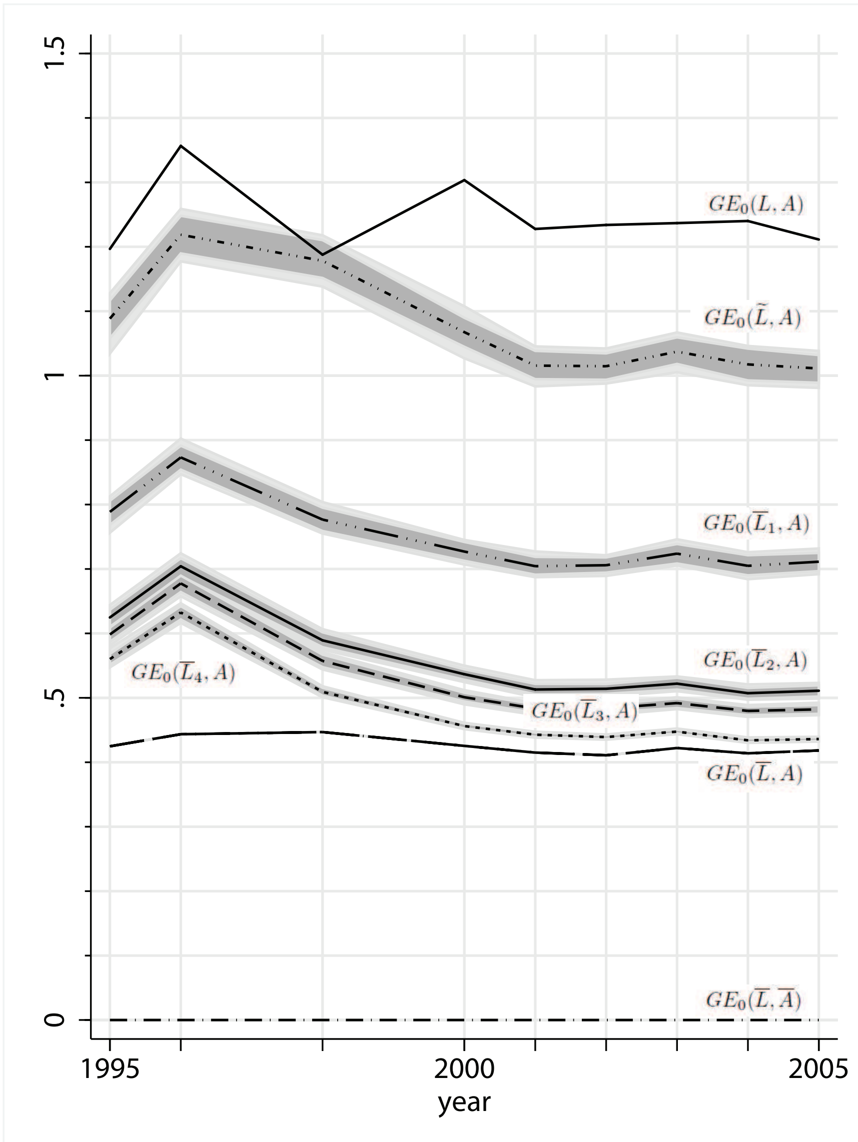

Figure 4: Outcomes first: $GE_0(L,A)$, $GE_0(\widetilde{L},A), GE_0(\overline{L},A), GE_0(\overline{L},\widetilde{A})$, and $GE_0(\overline{L},\overline{A})$

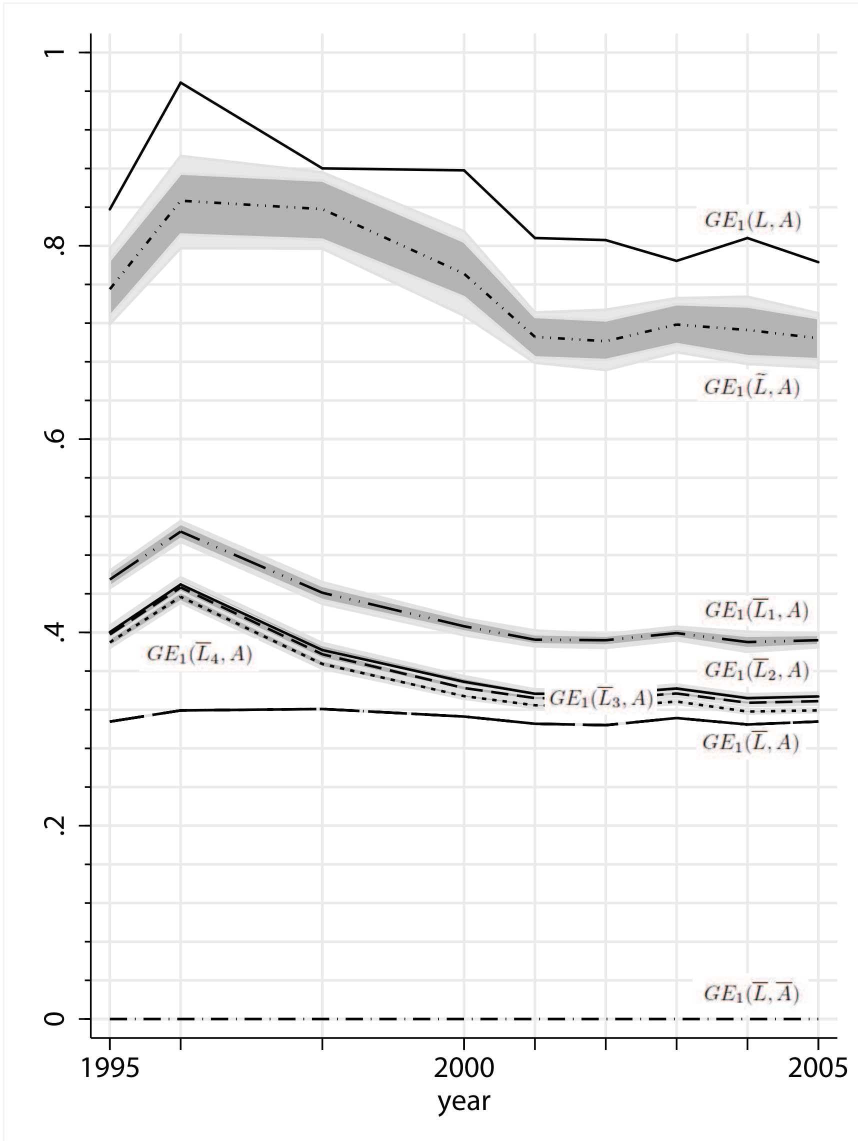

Figure 5: Outcomes first: $GE_1(L, A)$, $GE_1(\widetilde{L}, A), GE_1(\overline{L}, A), GE_1(\overline{L}, \widetilde{A})$, and $GE_1(\overline{L}, \overline{A})$

inequality indices that are each computed with a different reshuffled preference matrix $\widetilde{A}$. Similarly, the light grey area shows the 99% confidence interval.[13]

The figures show that neutralizing the correlation between the outcomes and preferences lowers well-being inequality. This inequality-reducing effect of the correlation between outcomes and preferences can be understood by consulting Table 2, which, for each non-income dimension of life ($j = 2, \ldots, 5$), shows the Spearman rank correlation coefficient between the outcomes ℓ_{it}^j and the individual-specific marginal rates of substitution between that dimension and the income dimension $(\beta^j + \mu^{j\prime} D_{it})/(\beta^1 + \mu^{1\prime} D_{it})$. This marginal rate of substitution measures the "willingness-to-pay" for a small improvement in dimension j. Individuals who care more about their outcomes in the non-income dimension (or less about their income) have larger marginal rates of substitution and steeper indifference curves. The rank correlation coefficients for the housing dimension in Table 2 are negative, for instance. Individuals who live in a relatively low quality house and are further away from their own optimal outcome level suffer relatively more from this hardship. The difference between the equivalent income ℓ_i^{1*} and the actual income ℓ_i^1 is the loss in well-being of not reaching the optimal outcome level; hence the negative correlations for the housing dimension further increase the difference between the well-being of those individuals who do better and those who do worse on that dimension. This finding presents a first indication that it may be worthwhile to take preference heterogeneity into account when measuring well-being inequality.

4.2.2 Preference heterogeneity

We now turn to the second term of the decomposition, which neutralizes the preference heterogeneity by constructing a counterfactual distribution with the averaged preference matrix $\overline{A}$.[14] The resulting counterfactual inequality measures in the preferences-first decomposition $I(L, \overline{A})$ use a common well-being

[13] This "reshuffling variance" stemming from the non-uniqueness of the reshuffling procedure should not be confused with sampling variance. In this paper we do not estimate the sampling variance of our results for several reasons. First, the additional confidence bounds would clutter the graphs and complicate the interpretation of results. Second, it is an open question how to deal with sampling variance in the computation of equivalent incomes. Finally, and most importantly, the RLMS-HSE data set does not provide sufficiently detailed information on the sampling procedure, so that the estimates of the sampling variance would at best provide rough approximations.

[14] An alternative approach would have been to use the coefficients of Model 2 in Table 1, which is estimated without interaction terms and in which preference differences also have been removed. The results for that approach are very similar to the ones presented here and can be obtained from the authors on request.

	1995	1996	1998	2000	2001	2002	2003	2004	2005
Health	0.0189	0.0030	-0.0232	-0.0103	-0.0132	-0.0163	-0.0056	0.0166	0.0253
Housing	-0.3233	-0.3323	-0.3831	-0.3521	-0.3620	-0.3875	-0.3458	-0.3505	-0.3532
Unemployment	-0.0351	-0.0644	-0.0655	-0.0842	-0.0674	-0.0754	-0.0781	-0.0797	-0.0812
Wage arrears	0.0556	0.0870	0.1326	-0.0080	-0.0213	-0.0277	-0.0324	-0.0624	-0.0377

Source: Own computations with RLMS-HSE

Table 2: Spearman rank correlation coefficient between outcome level and marginal rate of substitution with income for each non-income dimension

measure for all individuals and, in that sense, resemble the standard multidimensional inequality measures of eq. (2). The difference between this standard approach based on a common well-being measure and our heterogeneous approach shows up in the differences between $I(L, A)$ and $I(L, \overline{A})$ in Figures 2 and 3. After neutralizing the correlation between preferences and outcomes, the "net" effect of preference heterogeneity shows up in the differences between $I(L, \widetilde{A})$ and $I(L, \overline{A})$.

The results are striking. Removing preference heterogeneity leads to a substantial decrease in well-being inequality. The effect of substituting the averaged preference matrix $\overline{A}$ for the reshuffled matrix $\widetilde{A}$ is larger than that of substituting the reshuffled matrix $\widetilde{A}$ for the actual matrix A.

We make two further observations. First, the contribution of preference heterogeneity to well-being inequality remains quite stable over time. This is not surprising since we have assumed that preferences are constant over time for each person. Second, the contribution of preferences is relatively larger for GE_0 than for GE_1.[15] One possible interpretation is that individuals at the top of the well-being distribution score well on their non-income dimensions, so there is only limited room for preference heterogeneity to affect their well-being. For individuals at the bottom of the well-being distribution, on the contrary, the relative weighting of their different (larger) deprivations is more important.

The empirical relevance of preference heterogeneity on well-being inequality as well as the differences with the standard multidimensional inequality measures are further illustrated by the results of the outcomes-first decomposition given by eq. (10). The two bottom curves in Figures 4 and 5 show the evolution over time of $I(\overline{L}, A)$ (which is equal to $I(L, \widetilde{A})$) and $I(\overline{L}, \overline{A})$ (which equals 0). The former counterfactual captures inequality in the situation where all individuals have the same averaged outcomes, but their own preferences. According to the standard approach, which only uses information about L, this inequality is necessarily equal to zero. As we discussed in section 2.2, however, there may be inequality in well-being, even with identical outcomes, as soon as we introduce a concern for preference heterogeneity. In fact, the figures show that this inequality is substantial in our data. We return to the normative implications of this finding in the conclusion.

[15] Additional calculations, which are not shown here, confirm the pattern that preference heterogeneity has a larger effect as the inequality measure becomes more sensitive to the bottom of the well-being distribution.

		Expenditures	Health	Housing	Unemployment
1995	Health	0.0444			
	Housing	0.2296	-0.1062		
	Unemployment	-0.0633	0.1363	-0.0649	
	Wage arrears	0.0104	0.1027	-0.0248	-0.1353
2000	Health	0.1226			
	Housing	0.2904	-0.0734		
	Unemployment	-0.0744	0.1428	-0.0961	
	Wage arrears	0.0028	0.0538	-0.0658	-0.1094
2005	Health	0.1666			
	Housing	0.2023	-0.0821		
	Unemployment	-0.1284	0.1736	-0.0936	
	Wage arrears	-0.0046	0.0522	-0.0498	-0.0713

Source: Own computations with RLMS-HSE

Table 3: Spearman rank correlation coefficient between outcome dimensions

4.2.3 Correlation between outcomes

As described in the introduction, the phenomenon of cumulative deprivation, i.e., the correlation between the outcomes, has played a prominent role in the discussion on multidimensional inequality measurement (see Atkinson and Bourguignon (1982), Dardanoni (1996), and Tsui (1999), for instance).

In Table 3, we present the Spearman rank correlation coefficients between each pair of dimensions for 1995, 2000, and 2005. In line with the findings of Decancq (2014), we see an increased rank correlation between the expenditure and health dimension. Individuals who are top-ranked in the expenditures distribution become more likely to also be top-ranked in the health distribution over the considered period in Russia. Overall, however, the pattern of the correlation coefficients is mixed.[16]

The contribution of the correlation between the outcome dimensions to well-being inequality can be seen in both decompositions. In the preferences-first decomposition of eq. (9) it is reflected by the term $I(L, \overline{A}) - I(\widetilde{L}, \overline{A})$ and in the outcomes-first decomposition of eq. (10) by $I(L, A) - I(\widetilde{L}, A)$. The former shows the effect of the correlation between outcomes after preference heterogeneity has

[16]The impact that the increasing correlation between outcomes has on well-being inequality is an empirical matter because it depends on the interplay between the degree of substitutability and inequality aversion in both aggregation steps (see, e.g., Dardanoni (1996), Bourguignon (1999), and Bosmans et al. (2015) for discussions).

been removed, whereas the latter takes preference heterogeneity into account. The results are shown in Figures 2 through 5, where the dark grey area around the curves for $I(\widetilde{L}, A)$ and $I(\widetilde{L}, \overline{A})$ shows the 95% confidence interval of the reshuffling variance originating from the 200 reshuffled outcomes matrices $\widetilde{L}$. The light grey area shows the 99% confidence interval.

We see that the correlation between the outcome dimensions increases well-being inequality and that the contribution increases over time. Moreover, the contribution of the correlation between the outcomes has a stronger effect in the outcomes-first decomposition when the preference heterogeneity has not yet been neutralized. The increase in well-being inequality in the tumultuous period between 1998 and 2000 seems to be largely driven by the contribution of the correlation between the outcomes. As can be seen from Figure 4, a counterfactual situation with a stable contribution of correlation over time would have led to a decrease rather than an increase in well-being inequality.

4.2.4 Inequality in outcomes

Let us finally look at the fourth term of the decomposition, which captures the contribution of the inequality in each of the outcome dimensions to overall well-being inequality. Again, the results for the preferences-first decomposition (Figures 2 and 3) are related to the results of the standard approach to multidimensional inequality (since the preference parameters are fixed for each individual at $\overline{a}$).

We neutralize the inequality in the different dimensions in a specific order: first we average expenditures, followed by, consecutively health, housing quality, unemployment, and wage arrears. In principle, this specific sequence may affect the results and other sequences may lead to different results. Yet, since the equivalent income well-being measure as defined by (8) is close to being additively separable, this effect is quite small, and reversing the sequence hardly changes the results.[17]

Our findings are similar in the four figures.[18] Overall, the most important contributors to overall well-being inequality are the inequality in the expenditure and health dimensions. Moreover, there is a remarkable increase in well-being

[17] Results are available from the authors on request.

[18] Since the non-averaged dimensions have been reshuffled in the previous step of the decomposition, the non-uniqueness of the reshuffling remains to cause some variance in the counterfactual inequality measures. The more dimensions that are averaged, however, the smaller this variance becomes.

inequality due to the presence of wage arrears around 1996. The effect of wage arrears tapers off over time, however.

5 Subgroup decomposition by preference groups

An alternative approach to investigate the importance of preference heterogeneity on well-being inequality is based on a classic between-within subgroup decomposition (see Cowell (2011)). We partition the sample into 16 preference subgroups that are based on the socio-demographic characteristics captured by the four dummies in D_{it}, i.e., the gender of the respondents, whether they have obtained some higher education, whether they live in a rural area, and whether they are young or not.

Following Cowell and Jenkins (1995), we look at the subgroup decomposition to understand the importance of this particular partitioning in preference groups for well-being inequality. We do that separately for $I(L, A)$, $I(L, \overline{A})$, and $I(\overline{L}, A)$. The mean logarithmic deviation (GE_0) and the Theil-index (GE_1), on which we have focused so far, have attractive decomposition properties. It is indeed well-known that GE_α can be additively decomposed in a within component GE_α^W and a between component GE_α^B:

$$GE_\alpha(L, A) = GE_\alpha^W(L, A) + GE_\alpha^B(L, A),$$

where the between component is computed by setting all equivalent incomes in each preference group equal to their group average, and the within component is given by

$$GE_\alpha^W(L, A) = \sum_{k=1}^{K} \left[(v_k)^{\,(1-\alpha)} \right] \times [(s_k)^\alpha] \times GE_\alpha(L_k, A_k),$$

with L_k and A_k being the outcome matrix and preference matrix for preference subgroup $k = 1, \ldots, K$, $v_k = n_k/n$ being the population share and s_k the equivalent income share. When $\alpha = 0$ the inequality within the preference groups is weighted by the population shares, whereas for $\alpha = 1$, the equivalent income shares are used.[19]

[19]In our data set, the equivalent income shares of the preference groups are more unequal compared to the population shares. Young, urban respondents have a larger equivalent income share and lower educated, rural respondents a lower equivalent income share compared to their population share.

	1995	1996	1998	2000	2001	2002	2003	2004	2005
$GE_0(L, A)$	1.20	1.36	1.19	1.30	1.23	1.23	1.24	1.24	1.21
$GE_0^B(L, A)$	0.41	0.47	0.38	0.49	0.44	0.49	0.46	0.47	0.45
$GE_0^W(L, A)$	0.79	0.88	0.81	0.81	0.79	0.74	0.78	0.77	0.76
$GE_0(L, \overline{A})$	0.80	0.95	0.85	0.85	0.80	0.80	0.81	0.82	0.81
$GE_0^B(L, \overline{A})$	0.12	0.17	0.13	0.19	0.17	0.18	0.18	0.18	0.17
$GE_0^W(L, \overline{A})$	0.68	0.79	0.71	0.66	0.63	0.61	0.63	0.64	0.64
$GE_0(\overline{L}, A)$	0.42	0.44	0.45	0.43	0.42	0.41	0.42	0.41	0.42
$GE_0^B(\overline{L}, A)$	0.42	0.44	0.45	0.43	0.42	0.41	0.42	0.41	0.42
$GE_0^W(\overline{L}, A)$	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Source: Own computations with RLMS-HSE

Table 4: Subgroup decomposition analysis for GE_0

Tables 4 and 5 present the results.[20] For $I(L, A)$ the within component dominates: since preferences are the same within each preference group, this component captures the differences in outcomes between the individuals within each preference group. The between component captures both outcome and preference differences between the preference groups.

Additional insights can be obtained by looking at the decomposition of $I(L, \overline{A})$. Since all preferences are equalized in this counterfactual distribution, the between component only reflects differences in outcomes between the preference groups. The between component is now much smaller, suggesting that preference heterogeneity is an important contributor to the differences between the preference groups. As can be seen from comparing both tables, the within-group inequality is larger for the mean logarithmic deviation when with, rather than without, preference heterogeneity; that is, $GE_0^W(L, A) > GE_0^W(L, \overline{A})$, while the opposite is true with the Theil-index, i.e., $GE_1^W(L, A) < GE_1^W(L, \overline{A})$. This is in line with our earlier finding that taking preference heterogeneity into account has a larger effect for inequality measures that focus more on the bottom of the well-being distribution.

Finally, the subgroup decomposition of $I(\overline{L}, A)$ confirms the findings of section 4.2. The within component now becomes zero, while the between component captures the effect of preference heterogeneity in the counterfactual situation when all outcomes are averaged.

[20]The values in Tables 4 and 5 are the same as the corresponding ones in Figures 4 and 5.

	1995	1996	1998	2000	2001	2002	2003	2004	2005
$GE_1(L,A)$	0.84	0.97	0.88	0.88	0.81	0.81	0.78	0.81	0.78
$GE_1^B(L,A)$	0.30	0.34	0.29	0.34	0.32	0.34	0.33	0.33	0.32
$GE_1^W(L,A)$	0.53	0.63	0.59	0.54	0.49	0.46	0.46	0.48	0.46
$GE_1(L,\overline{A})$	0.71	0.85	0.76	0.75	0.68	0.68	0.68	0.70	0.68
$GE_1^B(L,\overline{A})$	0.11	0.16	0.13	0.16	0.15	0.16	0.16	0.16	0.15
$GE_1^W(L,\overline{A})$	0.60	0.69	0.63	0.58	0.53	0.52	0.52	0.54	0.53
$GE_1(\overline{L},A)$	0.31	0.32	0.32	0.31	0.31	0.30	0.31	0.30	0.31
$GE_1^B(\overline{L},A)$	0.31	0.32	0.32	0.31	0.31	0.30	0.31	0.30	0.31
$GE_1^W(\overline{L},A)$	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Source: Own computations with RLMS-HSE

Table 5: Subgroup decomposition analysis for GE_1

6 The dominance approach and preference heterogeneity

When computing multidimensional inequality with a common well-being measure $WB(\ell_i, a)$, the results will depend on the specific choice of the common preference vector a. The dominance approach addresses this dependence by taking an agnostic position on the precise shape of the common preferences and by computing the results for classes of well-being measures that are characterized by restrictions on their cross-derivatives with respect to the outcomes (see Atkinson and Bourguignon (1982) and Trannoy (2006)). To some extent, this agnosticism moves us away from the perfectionism that is implicitly underlying the choice of a single well-being measure for all individuals. It comes at a price, however. The resulting well-being inequality ranking may turn out to be incomplete, meaning that some comparisons will be indecisive.

To illustrate, we implement the following well-being inequality ranking[21]:

$$L \text{ is more unequal than } L' \Leftrightarrow I(L,\widehat{A}) \geq I(L',\widehat{A}) \quad \forall \widehat{A} \in \widehat{\mathcal{A}}, \tag{11}$$

where $\widehat{\mathcal{A}}$ is the set of all preference matrices in which all individuals share a common preference vector $\widehat{a}$ that is observed in the sample. As we have seen in the previous section, there are 16 different preference groups in our

[21] A similar approach underlies sensitivity analyses, such as, for instance, the one by Maasoumi and Jeong (1985).

	1995	1996	1998	2000	2001	2002	2003	2004	2005
1995	-	1	0	0	0	0	0	0	0
1996	0	-	0	0	0	0	0	0	0
1998	0	1	-	0	0	0	0	0	0
2000	0	1	0	-	0	0	0	0	0
2001	0	1	0	1	-	0	0	0	0
2002	0	1	0	1	0	-	0	0	0
2003	0	1	0	1	0	0	-	0	0
2004	0	1	0	0	0	0	0	-	0
2005	0	1	1	1	0	0	0	1	-

Own computations with RLMS-HSE

Table 6: Dominance test. A "1" means that the row year has a lower GE_0 than the column year for all observed preference parameters.

empirical analysis and, hence, 16 different preference matrices $\widehat{A}$. Clearly, this dominance idea could be further generalized by checking the inequality in eq. (11) for various members of some class of inequality measures - by testing Lorenz dominance, for instance. However, testing dominance is not the purpose of this paper and we only illustrate the approach for a single measure, which is the mean logarithmic deviation GE_0.

Table 6 presents the results for all pairwise year-by-year tests of the dominance test given by eq. (11). A cell with a "1" denotes that the row year has a lower well-being inequality according to GE_0 for each of the 16 common preference matrices $\widehat{A}$. We see that all years are less unequal than 1996, and that most of the years after 2000 are less unequal than 2000.

It is important to stress that the dominance approach does not take into account the diversity in preferences in a given society at a given point in time. There is an important difference between, on the one hand, looking for a unanimous inequality ranking for different well-being measures, each of them common to all individuals in society, and on the other hand measuring well-being inequality while respecting preference heterogeneity.

Although one has to interpret our findings cautiously, this difference can be illustrated by comparing the results in Table 6 with those in Figures 2 and 4. According to both approaches we find that 1996 is more unequal than 1995 (and more unequal than 1998). Yet, while 2000 was more unequal than 1998 based on the figures, this is not found in the dominance results. This difference may have to do with the underlying causes of the increases in well-being inequality. The increase in well-being inequality in 1996 is due to the sharp well-being loss as a result of the presence of wage arrears, which is an "objective" phenomenon that affects the inequality for all preference matrices $\widehat{A}$. The increase in 2000, however, is mainly due to an increase in the effect of the correlation between outcomes, and the correlation between outcomes and preferences. Preference differences are important for the evaluation of the former correlation and essential for the latter correlations. This may explain why the sharp inequality increase in 2000 is not reflected in the dominance results.

7 Conclusion

We have shown that preference heterogeneity constituted an important part of well-being inequality in Russia between 1995 and 2005. All-in-all, we have found that the main drivers of well-being inequality in the considered period were preference heterogeneity, expenditure inequality, health inequality, and wage arrears inequality (during the late 90s).

Some caveats apply, however. First, our empirical findings are based on one data set only. In the period between 1995 and 2005, Russia was a specific setting, characterized by large social and economic changes in a heterogeneous society. It is not clear whether preference heterogeneity would be equally important in other settings. Second, our method to estimate preferences on the basis of a satisfaction equation is arguably rather primitive. In particular, the consistency condition in eq. (7) is debatable and is hard to test empirically. Yet, one could argue that the fact that we cannot identify individual preferences with this method, but have to limit ourselves to only 16 different preference groups, strengthens our conclusion on the empirical relevance of preference heterogeneity.

More important than our specific findings for Russia, however, are the normative and methodological questions that are raised by these findings. Preference heterogeneity is completely neglected by the standard approach to multidimensional inequality measurement. Leaving pragmatic considerations of the availability of preference information aside, this position has been justified on normative grounds. There seems to be a certain distrust of individual preferences in the capability approach, for instance (Sen, 1985). The capability approach has been very influential in shaping the multidimensional approach towards the measurement of well-being, inequality, and poverty. Already before Sen (1985), Kolm (1977) suggested in his seminal article that a common well-being measure could be seen as "the observer's evaluation of the individual welfare", and Scanlon (1975) wrote that the common objective opinion on what a good life is and what constitutes well-being is rooted in some "reasoned social agreement on basic components of well-being and on the relative 'urgency' of claims to different goods". As we have seen, the dominance approach does not depart from the basic idea that there is one underlying common well-being measure, but introduces the additional twist that there may be uncertainty about this reasoned social agreement or a lack of consensus between different ethical observers.

As emphasized by preferentialists, the argumentation in favour of neglecting individual preferences has a strong perfectionist flavour. They claim that in a pluralist society with widely divergent opinions about what constitutes a good life, public policy in general, and inequality measurement in particular, cannot neglect these divergences and should therefore take up preference heterogeneity.

The difference between the two approaches is perhaps illustrated most strikingly by their evaluation of the hypothetical situation in which all individuals in society have the same objective outcomes. According to the standard multidimensional measurement literature, no ethically relevant inequality remains in that situation. If one takes preferences into account, however, the match between the outcomes and the preferences is brought into the picture, and it is seen as ethically relevant that different individuals can attach different weights to the different dimensions and may therefore have a different well-being, even when their objective outcomes are the same.

Multidimensional inequality measures and dominance approaches are arguably the best way to proceed if one believes that individuals do not have well-defined conceptions of the good life, or that, even when they exist, it is impossible to know them, or that, even when they exist and one can approximate them, one should not do so, but rather implement an objective conception of the good life (Decancq et al. (2015b)). If, on the other hand, one does believe that individuals can form a well-considered opinion about what is important in their own life, that these preferences can be reasonably (although imperfectly) approximated, and that they should be respected in a pluralist society, then one should introduce preference heterogeneity into the measurement of well-being inequality. This is essentially a normative debate, to which we did not contribute in this paper. What we have shown, however, are the stakes of the debate. The normative choices determining the role of preference heterogeneity have a crucial effect on the resulting well-being inequality. They do matter.

References

Aaberge, R. and A. Brandolini (2015): "Chapter 3 - Multidimensional Poverty and Inequality," in *Handbook of Income Distribution*, ed. by A. B. Atkinson and F. Bourguignon, Elsevier, vol. 2 of *Handbook of Income Distribution*, 141 – 216.

Atkinson, A. B. and F. Bourguignon (1982): "The comparison of multi-dimensioned distributions of economic status," *The Review of Economic Studies*, 49, 183–201.

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

Bourguignon, F. (1999): "Comment on 'Multidimensioned approaches to welfare analysis' by E. Maasoumi," in *Handbook of income inequality measurement*, ed. by J. Silber, London: Kluwer Academic, 477–84.

Box, G. E. P. and D. R. Cox (1964): "An analysis of transformations," *Journal of the Royal Statistical Society. Series B (Methodological)*, 26, 211–252.

Chakravarty, S. and M. Lugo (2016): "Multidimensional Indicators of Inequality and Poverty," in *Oxford Handbook on Well-being and Public Policy*, ed. by M. Adler and M. Fleurbaey, Oxford University Press, 246–285.

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

Cowell, F. A. and S. P. Jenkins (1995): "How Much Inequality Can we Explain? A Methodology and an Application to the United States," *The Economic Journal*, 105, 421–430.

Dardanoni, V. (1996): "On multidimensional inequality measurement," in *Research on Economic Inequality*, ed. by C. Dagum and A. Lemmi, London: JAI Press, vol. 6, 201–205.

Decancq, K. (2014): "Copula-based measurement of dependence between dimensions of well-being," *Oxford Economic Papers*, 66, 681–701.

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

——— (2015b): "Inequality, Income, and Well-Being," in *Handbook of Income Distribution*, ed. by A. B. Atkinson and F. Bourguignon, Elsevier, vol. 2 of *Handbook of Income Distribution*, 67 – 140.

Decancq, K. and M. A. Lugo (2012): "Inequality of Wellbeing: A Multidimensional Approach," *Economica*, 79, 721–746.

Dutta, I., P. K. Pattanaik, and Y. Xu (2003): "On measuring deprivation and the standard of living in a multidimensional framework on the basis of aggregate data," *Economica*, 70, 197–221.

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

Fleurbaey, M. and D. Blanchet (2013): *Beyond GDP: Measuring Welfare and Assessing Sustainability*, Oxford University Press.

Fleurbaey, M. and F. Maniquet (2011): *A Theory of Fairness and Social Welfare*, Cambridge University Press.

Fleurbaey, M. and A. Trannoy (2003): "The impossibility of a Paretian egalitarian," *Social Choice and Welfare*, 21, 243–263.

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

Frijters, P., J. P. Haisken-DeNew, and M. A. Shields (2004): "Investigating the patterns and determinants of life satisfaction in Germany following reunification," *Journal of Human Resources*, XXXIX, 649–674.

Gajdos, T. and J. Weymark (2005): "Multidimensional generalized Gini indices," *Economic Theory*, 26, 471–496.

Jones, A. M. and S. Schurer (2011): "How does heterogeneity shape the socioeconomic gradient in health satisfaction?" *Journal of Applied Econometrics*, 26, 549–579.

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

Lasso de la Vega, C., A. Urrutia, and A. de Sarachu (2010): "Characterizing multidimensional inequality measures which fulfil the Pigou–Dalton bundle principle," *Social Choice and Welfare*, 35, 319–329.

Layard, R., G. Mayraz, and S. Nickell (2008): "The marginal utility of income," *Journal of Public Economics*, 92, 1846 – 1857.

Maasoumi, E. (1986): "The measurement and decomposition of multi-dimensional inequality," *Econometrica*, 54, 991–997.

Maasoumi, E. and J. H. Jeong (1985): "The trend and the measurement of world inequality over extended periods of accounting," *Economics Letters*, 19, 295–301.

Scanlon, T. (1975): "Preference and urgency," *Journal of Philosophy*, 72, 655–669.

Sen, A. K. (1985): *Commodities and Capabilities*, Amsterdam and Oxford: North-Holland.

Stiglitz, J. E., A. Sen, and J.-P. Fitoussi (2009): "Report by the Commission on the measurement of economic performance and social progress," Mimeo.

Trannoy, A. (2006): "Multidimensional egalitarianism and the dominance approach: a lost paradise," in *Inequality and Economic Integration*, ed. by F. Farina and E. Savaglio, London: Routledge, 284–302.

Tsui, K. Y. (1999): "Multidimensional inequality and multidimensional generalized entropy measures: an axiomatic derivation," *Social Choice and Welfare*, 16, 145–157.

Weymark, J. A. (2006): "The normative approach to the measurement of multidimensional inequality," in *Inequality and Economic Integration*, ed. by F. Farina and E. Savaglio, London: Routledge.