Multidimensional poverty measurement with individual preferences*

Koen Decancq† Marc Fleurbaey‡ François Maniquet§

June 7, 2018

Abstract

We propose a new approach to multidimensional poverty measurement. To aggregate and weight the different dimensions of poverty, we rely on the preferences of the concerned individuals rather than on an arbitrary weighting scheme selected by the analyst. We provide an axiomatic characterization of an approach in which multidimensional poverty measures add up individual indices of poverty based on their multidimensional outcomes and their preferences. We discuss two families of these individual indices of poverty: quantity metrics and money metrics. Members of the first family evaluate individual poverty by the fraction of the poverty line vector to which the individual is indifferent. The second family considers the ratio between the income to which the individual is indifferent, for some fixed price vector, and the money value of the poverty line vector. We illustrate our approach with Russian survey data between 1995 and 2005. We find that, compared to standard poverty indices, our preference-sensitive indices lead to considerable differences in the identification of the poor.

JEL Classification: D63, D71.

Keywords: multidimensional poverty measurement, preferences.

*We thank Esfandiar Maasoumi, Christophe Muller, and Erik Schokkaert for helpful discussions as well as audiences in Barcelona, Bari, Bonn, Brighton, Glasgow, Leuven, Louvain-la-neuve, Marseille, Moscow, New Orleans, Namur, Odense, Paris, Princeton, Sevilla, Verona, and Warwick. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement n° 269831.

†CSB, University of Antwerp; CPNSS, London School of Economics; CORE, Université catholique de Louvain and CES, Katholieke Universiteit Leuven. Email: koen.decancq@uantwerp.be.

‡Princeton University. Email: mfleurba@princeton.edu.

§CORE, Université catholique de Louvain. Email: francois.maniquet@uclouvain.be.

1 Introduction

A growing consensus has emerged that well-being is multidimensional and that income —even suitably deflated for differences in prices— does not qualify as a good proxy for it.[1] As a consequence, well-being is increasingly measured as a multidimensional phenomenon. Anthony Atkinson, together with François Bourguignon, has been one of the pioneers of the multidimensional approach to compare distributions of well-being (see Atkinson and Bourguignon [6], amongst others). In a lucid article in the first issue of the *Journal of Economic Inequality*, Atkinson [5] discusses the measurement of multidimensional poverty by contrasting social welfare and counting approaches.

In the same spirit, the current paper aims to cross-fertilize the literature on the measurement of multidimensional poverty —which is currently largely dominated by counting approaches in the wake of Alkire and Foster [3, 4]— by new insights from social choice theory on fair social orderings. In particular, we discuss how multidimensional poverty can be measured in a way that is respectful for the preferences of the concerned individuals.

The common practice to measure multidimensional poverty consists of defining a threshold in each dimension of poverty, and claiming that an individual is deprived in a dimension if she experiences a lower level than the threshold. Measuring multidimensional poverty then requires a method to aggregate these deprivations across the different dimensions for all poor individuals.[2]

To do that, two central ethical choices need to be made. First, the so-called identification issue concerns the demarcation of the set of poor individuals. Some researchers adopt the union definition of poverty in which deprivation in at least one dimension is sufficient to qualify as poor, others follow the intersection definition and require individuals to be below the threshold in all dimensions. Intermediate positions can be taken in which deprivation in a limited number of dimensions is sufficient to qualify as poor, see Atkinson [5] for a discussion. Second, ethical choices about the relative importance of the dimensions (i.e. the weight assigned to the deprivation in each dimension) and whether the dimensions are seen as complements or substitutes have to be made before multidimensional poverty can be measured.[3]

[1] There are at least two reasons. First, private good markets, as well as labor markets, may fail to be competitive, so that individuals suffer from rationing. Second, some relevant goods may not be private and marketable (think of education, security or health, for instance).

[2] See, e.g., Tsui [51], Atkinson [5], Bourguignon and Chakravarty [13], Alkire and Foster [3, 4], Bossert, Chakravarty and D'Ambrosio [12], and Bosmans, Ooghe, and Lauwers [11].

[3] These choices are embedded in the choice of the shape of the so-called multidimensional poverty frontier (see, e.g., Duclos et al. [24], Maasoumi and Lugo [42], and Maasoumi and Racine [43]).

Typically, it is the researcher measuring poverty who makes both ethical choices. For obvious reasons, this practice can be criticized for being arbitrary. Ravallion [46] writes: "those with a stake in the outcomes will almost certainly be in a better position to determine what weights to apply than the analyst calibrating a measure of poverty." Turning to the opinions of the poor themselves, a large-scale participatory consultation by the World Bank at the end of the 1990s has indeed endorsed the view that poverty is a multidimensional phenomenon (see Narayan et al. [44]). At the same time, a wide diversity of views on the notion of multidimensional poverty by the poor themselves has been documented by this consultation. The question now arises whether poverty can effectively be measured as a multidimensional phenomenon, without becoming overly arbitrary on the embedded ethical choices.

We propose to use the individuals' own preferences to identify the poor and to aggregate across dimensions. That is, we enrich the model by considering that individuals have —possibly different— preferences over the different poverty dimensions. We discuss two broad families of measures capturing the idea that these preferences should be respected when measuring poverty. In these measures, the ethical choice of how to weight the goods or assessing their complementarity or substitutability is left to the individuals themselves. Respecting preferences also changes the outlook of the identification issue. The individuals identify themselves whether they are poor, by comparing their multi-attribute situation with the poverty line vector by means of their own preferences.

A difficulty with the preference-based approach, on the other hand, is the question how to make interpersonal well-being comparisons between individuals. For instance, we can no longer assume that two individuals with the same outcomes are equally poor. Likewise, we can no longer assume that an identical increase in one dimension has the same impact on two individuals if they have different preferences. We deal with both issues by resorting to two classical families of well-being indices that have recently been axiomatically characterized in social choice theory by Fleurbaey and Maniquet [30, 31, 32].[4] Following Deaton and Muellbauer [17, pp. 179-182], we call them *quantity metrics* and *money metrics*, respectively. Figure 1 provides an illustration of a quantity-metric multidimensional poverty index (in the left-hand panel) and a money-metric multidimensional poverty index (in the right-hand panel).

For the first family of quantity-metric multidimensional poverty indices, the specific index that evaluates an individual situation is one minus the fraction of the poverty line vector to which the individual is indifferent. This particular way of measuring individual

[4] See Fleurbaey and Maniquet [29] for a synthetic presentation of various results.

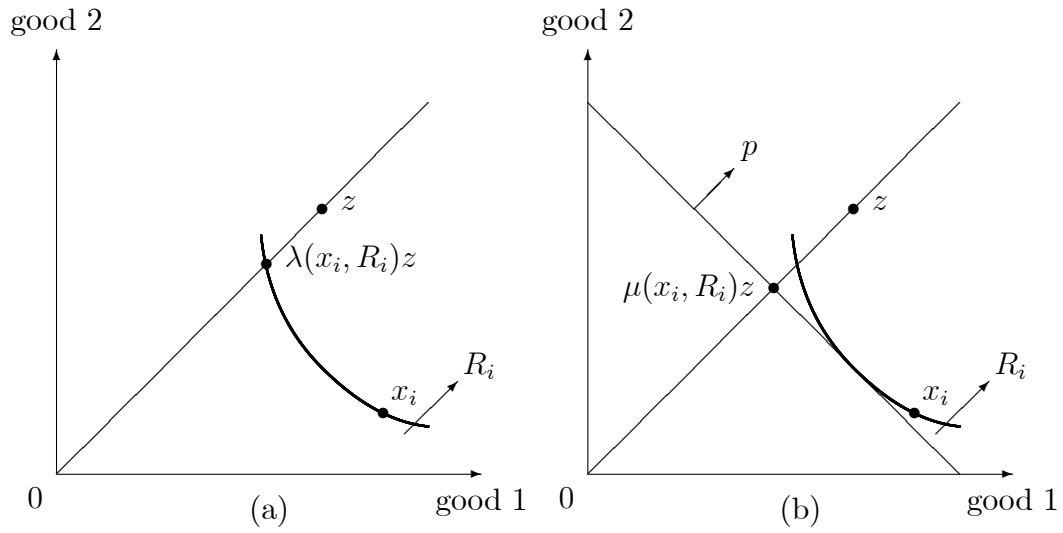


Figure 1: (a) A quantity-metric multidimensional poverty index: individual i is indifferent between x_i and $\lambda(x_i, R_i)z$; his contribution to global poverty is equal to $1 - \lambda(x_i, R_i)$. (b) A money-metric multidimensional poverty index: individual i is indifferent between x_i and income $p\mu(x_i, R_i)z$ facing prices p; his contribution to global poverty is equal to $1 - \mu(x_i, R_i)$.

well-being is similar to the ray index proposed by Samuelson [47], the distance function studied by Deaton [16], and the notion of egalitarian-equivalence due to Pazner and Schmeidler [45]. The second family of money-metric multidimensional poverty indices is inspired by money-metric utility, a concept proposed by Samuelson and Swamy [48] and Samuelson [47]. In this family, the specific index that evaluates an individual situation is one minus the fraction between the income to which the individual is indifferent, for some fixed price vector, and the money value of the poverty line bundle.[5]

Taking preferences into account in the measurement of poverty raises several new empirical problems, notably on the question how the information on the preferences of the concerned individuals can be obtained. We illustrate here one possible way in which one can tackle these problems. Using an existing Russian survey data set (RLMS-HSE) between 1995 and 2005, we estimate indifference maps in the space of equivalized household expenditures, health, housing quality, and unemployment based on a life satisfaction regression. This allows us to compute the proposed poverty measures, to assess the evolution of poverty and to compare our results with other unidimensional and multidimensional measures such as the counting approach. We find that taking preferences into account leads to important differences in the identification of the poor.

The remainder of the paper is organized as follows. In Section 2, we define the model and derive a representation theorem that extends the classical one-dimensional result of Foster and Shorrocks [35]. In Section 3, we combine this theorem with the recent work by Fleurbaey and Maniquet [30, 31, 32] on well-being measurement to derive two families of poverty measures. We discuss how the measure can be implemented with non-standard goods that are discrete or bounded. Then we turn to the empirical illustration of the theory. In Section 4, we use data from the Russian Longitudinal Monitoring Survey (RLMS-HSE) to show how preferences can be estimated. In Section 5, we illustrate that different individuals are identified as poor when preferences are taken into account. In Section 6, we show how the proposed families of poverty measures allow us to study the evolution of poverty in Russia in a robust way. In Section 7, we give some concluding comments. Econometric details of the empirical application and the proof of the theorem are provided in two appendices.

[5] These two approaches have been axiomatically characterized by Sprumont and Zhou [50] and Gevers [37] in terms of allocation rules (i.e., selection of the best allocation), by Fleurbaey [26] and Fleurbaey and Maniquet [28] in terms of social orderings (i.e., rankings of all allocations), and by Fleurbaey and Maniquet [30] and Fleurbaey and Tadenuma [33] in terms of indexes of individual well-being. Money metric utilities have received recent attention by Bosmans et al. [10].

2 The model and a basic result

Introducing preferences implies going from multi-dimensional attributes to a unidimensional notion of well-being. This section formalizes this simple observation. Let us first briefly introduce the framework.

An economic situation is a pair (x_N, R_N), where N is the population. Let $x_N = (x_i)_{i \in N}$, where each bundle x_i is a ℓ-dimensional vector in a consumption set $X = \mathbb{R}^\ell_+$, describing the situation of individual i by looking at several "goods" (equivalized household expenditures, health status, housing quality and employment status in the empirical illustration of this paper). Further, let $R_N = (R_i)_{i \in N}$, where each R_i is i's preference relation over the set X of possible individual situations. The corresponding strict preference and indifference relations are denoted P_i and I_i, respectively. Let us be precise about what we call preferences. Ideally, we would like preferences to describe how the different dimensions contribute to the well-being of an individual. As is typical in economics, we consider that well-informed and rational choices reveal what makes an individual better-off. However, because we are interested in including dimensions the levels of which cannot be chosen, such as health, we cannot resort to the standard revealed preference approach. In our empirical illustration we will therefore use survey information about what makes people better-off, by means of a life satisfaction regression.

We are looking for a *poverty index*, i.e., a function P such that for every (x_N, R_N) in the relevant domain $\mathcal{S}$, $P(x_N, R_N)$ is a real number measuring the extent of poverty in (x_N, R_N). It is assumed that $P(x_N, R_N)$ is continuous in x_N (which excludes simple headcount measures, but these can be considered a limit case of the measures we obtain). Observe that by means of $P(x_N, R_N)$ economic situations can be compared in which preferences R_i, not just individual vectors x_i, are different. Here we follow Fleurbaey and Tadenuma [33], who advocate that social choice should extend beyond the comparison of options for a given population. As in their paper, the possibility to compare populations with different preferences is a direct side-product of the construction of the index, because even for a given population the index needs to compare individuals with different preferences.

We make the following assumptions on the domain $\mathcal{S}$ of economic situations (x_N, R_N). First, we assume that R_i may be any member of the set $\mathcal{R}$ of preferences which are continuous, monotonic (that is, for two bundles $x_i, x_i' \in \mathbb{R}^\ell_+$, if $x_i \leq x_i'$, then $x_i' \, R_i \, x_i$, and if $x_i \ll x_i'$, then $x_i' \, P_i \, x_i$)[6] and convex. Second, the population N may have any positive

[6] The three vector inequalities are denoted $\leq$, $<$ and $\ll$.

finite size, and $\mathcal{N}$ denotes the set of possible populations. When the population has only one individual, we use the simpler notation $P(x_i, R_i)$ instead of $P(x_{\{i\}}, R_{\{i\}})$.

In the remainder of this section, we introduce basic requirements which constrain our search for a poverty index P. We start by addressing the question "who is poor?". In our framework, it amounts to answering the question: which bundles make individual i with preferences R_i poor? We allow the answer to depend on individual preferences in a very general sense. Individuals are poor if they consider themselves worse off compared to a poverty line vector z which is allowed to depend on their preferences, i.e., to be a function $z(R_i)$.

Our first axiom captures the idea that an individual is not poor, therefore irrelevant for P, if she prefers her situation x_i to $z(R_i)$. We call it *Focus*, by reference to Sen's Focus axiom (Sen [49]). The axiom requires that the poverty index, at the individual level, be independent of any change in the situation of a non-poor individual.

Axiom 1 FOCUS
There is a function $z : \mathcal{R} \to X$ such that for all $(x_N, R_N) \in \mathcal{S}$, $i \in N$, $x'_i \in X$, $R'_i \in \mathcal{R}$, if

$$x_i \, P_i \, z(R_i), \; x'_i \, P'_i \, z(R'_i)$$

then

$$P((x'_i, x_{-i}), (R'_i, R_{-i})) = P(x_N, R_N).$$

We now adapt the classical requirement that an improvement in the situation of one individual cannot increase poverty. The core of our contribution in this paper is to use individual preferences to evaluate improvements in terms of individual well-being. We propose to apply a Pareto axiom, restricted to the poor: If the preference satisfaction of all poor individuals weakly increases, then poverty weakly decreases. If, in addition, the preference satisfaction of at least one poor individual strictly increases, then poverty strictly decreases.

Axiom 2 PARETO AMONG THE POOR
For all $(x_N, R_N), (x'_N, R_N) \in \mathcal{S}$, if for all $i \in N$ such that $z(R_i) \, P_i \, x_i$, $x'_i \, R_i \, x_i$, then

$$P(x'_N, R_N) \leq P(x_N, R_N).$$

If, in addition, there is $j \in N$ such that $z(R_j) \; P_j \, x_j$ and $x'_j \, P_j \, x_j$, then

$$P(x'_N, R_N) < P(x_N, R_N).$$

Pareto efficiency is necessary if we don't want to claim that poverty has decreased in cases where individual well-being, defined in a way that is consistent with individual preferences, has increased. This calls for the following remark, echoing the discussion in the Introduction. Let us assume that we are interested only in marketable goods. If prices are constant and consumers maximize their preferences, then income is a proxy of well-being that is consistent with individual preferences. To put it differently, under these assumptions, income poverty measurement cannot conflict with Pareto efficiency. Here we follow a recent trend in the literature and we consider that income is not a good proxy for well-being, notably because not all goods are marketable. This view calls from describing individual situations in terms of individuals' outcomes in the relevant dimensions instead of incomes. By doing so, however, we create the risk of evaluating poverty in opposition to individual preferences. This axiom allows us to ward off this problem.

The next two axioms are standard axioms of the classical poverty measurement theory that do not need much adjustment to our framework. The next axiom is *Subgroup Consistency*. It requires that overall poverty decreases if it decreases in a subgroup of the population and the situation does not change for the other individuals.

Axiom 3 Subgroup Consistency
For all $(x_N, R_N), (y_M, R_M), (y'_M, R'_M) \in \mathcal{S}$, $P(y_M, R_M) \geq P(y'_M, R'_M)$ *if and only if*

$$P((x_N, y_M), (R_N, R_M)) \geq P((x_N, y'_M), (R_N, R'_M)).$$

This *Subgroup Consistency* axiom is a standard and powerful decomposability requirement. Observe that the decomposition it proposes does not allow the separation of the bundle from the preferences of an individual. That is the only difference with the standard axiom.

The above three axioms enable us to derive the representation theorem that we will use in the remaining of the paper. This result can even be simplified if the axiom of *Replication Invariance* is added. It requires that the poverty measure remains the same if the population is replicated and each replica of the current population exhibits the same characteristics as the current one. We need the following additional terminology. Let $r \in \mathbb{N}_{++}$ be a positive integer. The economic situation (x_N^r, R_N^r) is a replica of (x_N, R_N) if the set of individuals is r times larger than N and is partitioned in r subgroups, one of which is N, and each subgroup has the same distribution of goods and preferences as N.

Axiom 4 Replication Invariance

For all $(x_N, R_N) \in \mathcal{S}$, *all* $r \in \mathbb{N}_{++}$,

$$P(x_N, R_N) = P(x_N^r, R_N^r).$$

We are now equipped to state and prove the following representation result.[7] It is a straightforward generalization of a result obtained by Foster and Shorrocks [35] in the standard one-dimensional framework. If we gather the above axioms, the resulting poverty index needs to be additively separable in individual characteristics (x_i, R_i).

Theorem 1 *A poverty index P satisfies* Focus, Pareto among the Poor, *and* Subgroup Consistency *if and only if there exist*

- *a continuous function* $G : \mathbb{R} \times \mathcal{N} \to \mathbb{R}$, *strictly increasing in its first argument,*
- *for all* $N \in \mathcal{N}$, *for all* $i \in N$, *a function* $\phi_i^N : X \times \mathcal{R} \to \mathbb{R}$ *such that* ϕ_i^N *is continuous in its first argument,* $\phi_i^N(x_i, R_i) > \phi_i^N(x_i', R_i)$ *whenever* $z(R_i)\, R_i\, x_i'\, P_i\, x_i$, *and* $\phi_i^N(x_i, R_i) = 0$ *whenever* $x_i\, R_i\, z(R_i)$,

so that for all $(x_N, R_N) \in \mathcal{S}$,

$$P(x_N, R_N) = G\left[\sum_{i \in N} \phi_i^N(x_i, R_i), N\right].$$

Moreover, if Replication Invariance *is added to the axioms, the poverty index can be simplified into*

$$P(x_N, R_N) = G\left[\frac{1}{|N|} \sum_{i \in N} \phi(x_i, R_i)\right], \tag{1}$$

for a continuous and strictly increasing function $G : \mathbb{R} \to \mathbb{R}$.

The proof of this theorem is given in the second Appendix. Theorem 1 illustrates how taking preferences into account affects the definition of a multidimensional poverty index. The first consequence is that we return to a one-dimensional individual measure of poverty, i.e., the $\phi(x_i, R_i)$ measure, which is ordinally equivalent to $P(x_i, R_i)$. Preferences provide a powerful way of aggregating several dimensions into one complete order, and, therefore, into a one-dimensional individual measure. The second consequence is that the ϕ function is left unspecified. Constructing it requires to take preferences into account, so that, for instance, two individuals consuming the same bundle of goods can be assumed to experience different levels of poverty (i.e., they have different values of the ϕ function) if their preferences differ. The next section is devoted to the construction of the ϕ function.

[7] Let us observe that, in the case $r = 1$, *Replication Invariance* boils down to the classical anonymity requirement: names of the agents do not matter, that is, only the list of bundles of goods and preference relations can influence poverty.

3 Inter-preference poverty comparisons

From the previous section we know that for poor individuals we have that $\phi(x_i, R_i) \leq \phi(x_i', R_i)$ if and only if $x_i R_i x_i'$. In this section, we review two ways of extending this to poverty comparisons across preferences, i.e., between $\phi(x_i, R_i)$ and $\phi(x_i', R_i')$. Comparing $\phi(x_i, R_i)$ with $\phi(x_i', R_i')$ in a way that is respecting the preferences is equivalent to building numerical representations of preferences R_i and R_i' and comparing the resulting numbers. The literature has proposed several ways of constructing these numerical representations. Among these proposals, the quantity-metric utility of Deaton [16] and the money-metric utility of Samuelson [47] and Samuelson and Swamy [48] fit our framework particularly well as they only depend on outcomes and preferences. These two indices have recently been axiomatically characterized by Fleurbaey and Maniquet [30, 31, 32]. These characterization results can almost immediately be applied to the ϕ function. In this section, we define two families of poverty indices that can be derived from these contributions and we briefly comment on their main properties for the measurement of poverty.

A first lesson that can be drawn from these axiomatic characterizations is that the nature of the goods matters. Depending on whether quantities of goods are unbounded, bounded or discrete, and depending on whether convex combinations of baskets of goods are meaningful, the relevant axioms lead to different results. Let us begin with the standard case in which quantities of goods are unbounded and convex combinations are meaningful, which is the case studied by Deaton [16], Samuelson [47] and Samuelson and Swamy [48]. We return to non-standard cases below.

We start by considering the family of quantity-metric poverty indices. This family is illustrated in the left-hand panel of Figure 1. In this case, let us define the ϕ function as follows:

$$\phi(x_i, R_i) = f\left(1 - \min\{1, \lambda_z(x_i, R_i)\}\right), \tag{2}$$

where $f : [0,1] \to [0,1]$ is continuous, decreasing and convex and where $\lambda_z(x_i, R_i) = \lambda$ if and only if $x_i \, I_i \, \lambda z$, for a common poverty line vector z. Remember that Theorem 1 leaves open the possibility of having different poverty line vectors for different individuals.

We discuss three main properties of these quantity-metric poverty indices. First, if we fix some bundle x_i, the preferences R_i for which quantity-metric multidimensional poverty is the largest are the Leontief preferences with their kinks along the z-ray (the ray of all bundles proportional to z), see Fleurbaey and Maniquet [30]. That means that the level of poverty depends on the inability to trade-off among the different dimensions. For all bundles outside the z-ray, the only way to decrease the poverty of an individual

with these Leontief preferences is to assign this individual with some additional quantity of the good in which he is most deprived.

The second property is that, along the z-ray, all individuals with the same bundle have the same poverty, irrespective of the shape of their preferences. Contrary to objective multidimensional poverty indices, the axiom *Pareto among the Poor* prevents us from concluding more generally that two individuals with the same bundle always have the same level poverty, irrespective of their preferences. Yet, quantity-metric indices of poverty guarantee this independence of the preferences along the z-ray. As a result, this approach coincides with an objective approach when all individuals' outcomes are proportional to z.

The third property, which is due to the convexity of f, is that along this ray a transfer of resources from a less poor to a poorer individual decreases poverty. As is known from the work by Fleurbaey and Trannoy [34], it is not possible to guarantee that these kind of transfers always decreases poverty for measures that satisfy *Pareto among the Poor*, but it can be guaranteed along the z-ray.

By substituting equation 2 in the multidimensional poverty index (equation 1), we obtain the quantity-metric family of poverty indices. Formally, let there exist a continuous and strictly increasing function $G : [0,1] \to \mathbb{R}$ and a continuous, decreasing and convex function $f : [0,1] \to [0,1]$ so that for all $(x_N, R_N) \in \mathcal{S}$,

$$P(x_N, R_N) = G\left[\frac{1}{|N|}\sum_{i\in N} f\left(1 - \min\{1, \lambda_z(x_i, R_i)\}\right)\right].$$

Of course, there are as many multidimensional poverty indices as poverty line vectors z. This is why we speak of a family of multidimensional poverty indices.

We now turn to the second family, the money-metric poverty indices. They are dual to the first one[8] and are illustrated in the right-hand part of Figure 1. We now define the ϕ function as follows:

$$\phi(x_i, R_i) = g\left(1 - \min\{1, \mu_{z,p}(x_i, R_i)\}\right), \tag{3}$$

where $g : [0,1] \to [0,1]$ is continuous, decreasing and convex and where $\mu_{z,p}(x_i, R_i) = \mu$ if and only if $x_i\, I_i\, \max(R_i, \{x_i' \in X | px_i' \leq \mu pz\})$.[9]

[8] This duality is formally grounded on the fact that the set of lower contour sets and the set of upper contour sets are lattices. See Fleurbaey and Maniquet [30] for a formal treatment.

[9] For $B \subset X$ and $R_i \in \mathcal{R}$, we write $\max(R_i, B)$ to denote any bundle in B that maximizes R_i over B, that is, $\max(R_i, B) = x_i$ only if $x_i \in B$ and $x_i\, R\, x_i'$ for all $x_i' \in B$.

Again, we can sketch several properties of this way of measuring poverty. First, if we fix some bundle x_i, the preferences R_i for which money-metric multidimensional poverty is the lowest are always the linear preferences with a slope proportional to p (see Fleurbaey and Maniquet [30]). That means that, again, poverty is closely related to the inability to trade-off among the different dimensions, but this time it is expressed by identifying the least poor individual at some fixed bundle as that individual who is most able to trade-off among goods, that is, who has linear preferences.

Compared with the previous measure, we lose the property that well-being does not depend on preferences for some set of bundles. Moreover, as can be understood from Figure 1, some individuals with non-linear preferences may strictly prefer their outcome to the poverty line vector z and still qualify as poor. This observation will have the consequence that if we use the same z for the two families of multidimensional poverty indices, the number of poor will necessarily be larger with the money-metric index.

Finally, it is quite intuitive that this second way of measuring poverty is closer to the classical income poverty approach. In particular, if all individuals face the same price vector p and if this vector is used in the measure, then we are back to classical income poverty measurement and, in this case, poverty comparisons across individuals no longer depend on their preferences.

When the goods are of a different, non-standard, nature, either because they take values in a bounded domain or because they come in discrete quantities —as it is often the case in applications[10]— the characterization results that justify the indices above do not immediately generalize. As shown by Fleurbaey and Maniquet [32], in the presence of non-standard goods, the relevant question becomes: does there exist a value of the non-standard goods that an individual prefers? In many cases, the answer is yes: think of a perfect health, or one's preferred employment status, etc. It turns out that with such goods, it is natural to take the preferred value of the non-standard good as the benchmark, and measure poverty by reference to the bundle containing the preferred quantity of these goods to which the individual is indifferent. This idea can be applied to extend both the quantity-metric or money-metric poverty indices to the case of non-standard goods.

Formally, let the outcome set now be $X = \mathbb{R}^{\ell}_{+} \times A$, and let us describe the outcomes of an individual by a list (x_i, a_i) such that $x_i \in \mathbb{R}^{\ell}_{+}$ and $a_i \in A$. The first family of poverty indices requires redefining the λ_z function in equation (2) as

$$\lambda_z((x_i, a_i), R_i) = \lambda \Leftrightarrow (x_i, a_i)\, I_i\, (\lambda z, \tilde{a}_i),$$

[10] See, amongst others, Alkire and Foster [3] and Bossert et al. [12].

where $\tilde{a}_i \in A$ is individual i's preferred value of the variables in A at quantity λz for the consumption of ordinary goods.

Similarly, the second family of poverty indices requires redefining the $\mu_{z,p}$ function in equation (3) as

$$\mu_{z,p}((x_i, a_i), R_i) = \mu \Leftrightarrow (x_i, a_i)\, I_i\, \max(R_i, \{(x_i', \tilde{a}_i) \in X | p x_i' \leq \mu p z\}).$$

These are the two poverty indices that we implement in the next sections. To do that, we need to choose the precise values of the normative parameters z, p, f and g. Our axiomatic theory is silent about the precise choice of these parameters. In the empirical sections below, we offer additional considerations guiding the choices for z and p. We also propose a test to make robust poverty comparisons that hold for all continuous, decreasing and convex functions f or g.

We end this section by looking in more detail at the nature of the four goods that will play a central role in the empirical illustration of the next sections: equivalized household expenditures, housing, health, and employment status. These considerations echo the reasons why income may not be considered a good index of poverty.

The first good, equivalized household expenditures (on consumption goods), is considered as a standard good.

The second good, housing, is also standard, but we add the idea that the housing market may not be competitive. We think of demand rationing that may be based on observable characteristics, such as ethnicity, or unobserved ones, such as the probability to find a job. As a consequence, we will apply the money-metric with a shadow price that may differ from the observed market price.

The third good, health, is clearly not a marketed good. This is where our estimation of preferences needs to follow another approach than the classical revealed preference approach. The health variable that we construct may be considered as continuous, but it is bounded by the value of the variable at perfect health. Two options are available. Either we disregard the boundedness of the value, because very few individuals enjoy a perfect health, and we treat health as a standard good. Alternatively, we can take the boundedness of the variable into account and treat perfect health as the reference value with respect to which equivalence is measured. We will study and compare both options in our empirical illustration.

By treating the fourth good, employment status, as a non-standard good, we refer to two features. First, we acknowledge that the labor market is typically non-competitive. Second, we define it as a discrete good. As a result, we need to take the value that the

individuals themselves prefer as their reference value. As we explain below, it turns out that some individuals in our sample prefer to stay unemployed (everything else remaining equal, including the amount of expenditure on consumption goods) while others do not.

4 Estimating preferences

To compute the proposed multidimensional poverty measures with real-world data, one needs to know the preferences of the concerned individuals. In this section, we illustrate how this information can be retrieved from an existing household panel data set, the Russian Longitudinal Monitoring Survey (RLMS-HSE). We use data between 1995 and 2005, which was a particularly turbulent period due to the fast Russian transition towards a market economy and the severe financial crisis of August 1999. Although the RLMS-HSE data set is designed to be representative, it is unlikely that the subsample that we use is representative as well. All our findings should therefore be interpreted as concerning the sample, rather than the underlying population of the Russian Federation. As mentioned before, we include four goods in the multi-attribute description of the respondents: a measure of equivalized household expenditures, a measure of health, housing quality, and unemployment. We provide more details on the data set and the construction of the variables in Appendix 1.

Various approaches can be followed to estimate preferences. First, preferences can be estimated from observed choice behavior. Yet, these revealed preference methods are only applicable to dimensions over which individuals make actual choices. Arguably this is not the case for the four goods considered here. Second, one can simply ask the respondent's opinions on the most appropriate trade-offs between the goods.[11] Such a stated-preference procedure may be cognitively demanding for the respondents, however. Moreover, it requires a specific battery of survey questions which are not included in the RLMS-HSE. Third, preferences of different sociodemographic groups can be estimated based on self-reported life satisfaction information.[12] Although imperfect, this life satisfaction approach

[11] Fleurbaey et al. [27] use a contingent valuation method to elicit another point on the indifference curve of the respondents in the space of monthly personal income and health outcomes. An alternative approach is to retrieve information about trade-offs by relying on a sequence of hypothetical binary choices, see Adler et al. [1], Benjamin et al. [7] and Decancq and Nys [21], for instance.

[12] Clark and Oswald [15] offer an introduction to the life satisfaction approach and Dolan and Fujiwara [23] a critical survey. Decancq et al. [18], [19], Decancq and Schokkaert [22] and Decancq and Neumann [20] use this approach to estimate ordinal preferences for well-being and inequality comparisons. The former two papers use a similar econometric model as the one presented here.

seems the most attractive for our problem given the nature of the considered goods and the data available in the RLMS-HSE. We return to the underlying assumptions of the life satisfaction approach at the end of this section.

To describe the econometric life satisfaction model, we denote the self-reported life satisfaction of individual i in period t as S_{it}. We start from a standard happiness regression with life satisfaction as the explained variable and a series of usual explanatory variables, including the vector of individual outcomes for the four dimensions of poverty (X_{it}) after a dimension-specific Box-Cox transformation, a time trend (γ_t) and some observable sociodemographic characteristics (Z_{it}) such as education, social status, marital status, average expenditures and employment level in a small geographical reference group, and the presence of wage arrears, which used to be a common phenomenon during the late nineties in Russia.[13] As unobservable personality traits are likely to influence self-reported life-satisfaction, we control for these time-invariant factors by including individual fixed effects (α_i) in the regression. We allow for preference heterogeneity by including interaction effects between the outcome vector and a vector of five dummies (D_{it}) capturing whether the respondents are young (below the age of 33), male, living in a rural area, obtained higher education and have a minority status. This leads to the following model:

$$S_{it}^* = \alpha_i + \gamma_t + (\beta + \Lambda D_{it})' X_{it} + \delta' Z_{it} + v_{it}, \tag{4}$$

where S_{it}^* is a latent satisfaction variable, β and δ are vectors of direct effects and Λ a matrix with interaction effects to be estimated. The idiosyncratic error term v_{it} is assumed to follow a logistic distribution function. We observe the reported life satisfaction $S_{it} = k$ for k in $\{1, 2, \ldots, 5\}$ if the latent life satisfaction (S_{it}^*) lies within an interval between η_{k-1} and η_k :

$$S_{it} = k \text{ if } \eta_{k-1} < S_{it}^* \leq \eta_k. \tag{5}$$

The thresholds η_k are allowed to depend on the individual fixed effects, the observable sociodemographic characteristics and the time trend (more estimation details are provided in Appendix 1).

The relevant estimation results are given in Table 1. Model 1 is estimated without including interaction effects and provides a benchmark where all individuals have the same preferences. More relevant for our purpose is Model 2 which includes interaction effects. In line with the literature, improvements in expenditures, health and housing

[13] A variable like education can both be seen as a dimension over which individuals have preferences and as a sociodemographic characteristic that affects the use of the response scale. With the satisfaction approach it is impossible to disentangle the two effects and identify preferences for education.

quality lead to higher life satisfaction, whereas being unemployed is found to decrease life satisfaction for most sociodemographic groups. To reach a parsimonious and tractable model, the least significant interaction effects in Λ have been consequently dropped until all the remaining interaction effects are significant at the 10% level. The coefficients of the remaining interaction terms are presented in the second part of the table. It can be seen that young people give relatively less importance to health and relatively less importance to their employment status, men care relatively more about health and about being employed than women, and so on.

Table 1: Happiness equation

	life satisfaction (Model 1)		life satisfaction (Model 2)	
expenditures (Box-Cox: $0,055$)	0,394***	(0,0328)	0,440***	(0,0351)
health (Box-Cox: $0,485$)	0,723***	(0,0928)	0,557***	(0,124)
house (Box-Cox: $-0,356$)	0,249**	(0,0842)	0,259*	(0,128)
unemployed	-0,424***	(0,0529)	-0,172	(0,123)
young × health			-0,315*	(0,159)
young × unemployed			0,190*	(0,0881)
male × health			0,401*	(0,179)
male × unemployed			-0,361***	(0.0882)
rural × house			0,397*	(0,169)
higher educated × expenditures			0,0441**	(0,0146)
higher educated × unemployed			-0,201+	(0,107)
higher educated × house			-0,240*	(0,109)
minority × health			0,655*	(0,263)
minority × expenditures			-0,382***	(0,0849)
N	53873		53873	
pseudo R^2	0,071		0,073	

Standard errors (in parentheses) are clustered at the household level.
Coefficients are obtained after controling for education level, social status, marital status, reference group expenditures, reference group employment level, the presence of wage arrears year dummies, and individual fixed effects.
$+\ p < 0,10$, $*\ p < 0,05$, $**\ p < 0,01$, $***\ p < 0,001$

The (McFadden) pseudo R-squared of the estimation is around 0,073, which is com-

parable to other studies using panel data (see, for instance Graham et al. [38]). This magnitude highlights that only a small part of the variation in life satisfaction can actually be explained. One may wonder at this point how problematic this finding is for the idea of respecting preferences when identifying the poor. People have different opinions on how to trade-off dimensions of poverty. If the estimated preferences were only meant to capture this heterogeneity in opinions or in behavior, then the relatively small share of the explained variation in life satisfaction should be a source of concern, as we should aim at approaching the actual preferences of people as close as possible. One may argue, however, that actual preferences are too idiosyncratic to be normatively compelling, for instance because people may make mistakes. Consequently, the actual preferences should be laundered before they are used in a normative judgement. This is precisely what the estimation carries out. We replace the actual individual preferences with the average preferences of the group to which the individual belongs, so that we end up only taking account of facts like the relative concern of elderly people for their health condition, the relatively low worry of higher educated people about housing conditions, and so on. However, if the sample size would have allowed us to increase the number of these groups and to take account of more relevant characteristics, that would certainly have been desirable.

Figure 2 illustrates the resulting indifference maps in the expenditures-health space for two population subgroups: the young, higher educated women (their outcomes are depicted with a diamond) and the old, lower educated men (the triangles in the figure). In line with expectations, we see that the elder individuals are on average in worse health. The illustrated indifference maps show that preferences of the latter group (depicted by the solid curves) are generally steeper, meaning that their willingness to pay for an increase in health is also higher.

Before turning to the identification of the poor individuals in the next section, we discuss shortly some modelling assumptions, which are implicit in the satisfaction approach (see also Dolan and Fujiwara [23] for a critical discussion). First, there is the central assumption that responses to a life satisfaction question are consistent with the preferences of the respondents. This so-called consistency assumption is discussed by Decancq et al. [18]. In general, the consistency assumption seems hard to test empirically.[14] Second, the econometric model summarized by equation (4) and (5) embeds several para-

[14]Testing the consistency assumption requires a comparison between life satisfaction responses and preferences. To the best of our knowledge, very few such comparisons have been performed, mostly in rater specific contexts. Benjamin et al. [7] compare satisfaction responses and stated preferences in hypothetical choice scenarios. They find considerable, but incomplete consistence between both approaches. Comparisons with revealed preferences are made by Benjamin et al. [8] and Akay et al. [2].

metric assumptions about the preferences and the way the response scale is used by the respondents. All preferences are assumed to be separable and within a sociodemographic subgroup preferences are assumed to be identical. An additional noteworthy assumption in equation (5) is that the use of the response scale is assumed not to be affected by outcomes in the dimensions of poverty (see Beegle et al. [9] for a discussion). Finally, the satisfaction approach hinges on a causal interpretation of the coefficients of the dimensions of poverty in the life satisfaction regression. Yet, even in the turbulent period in Russia that is covered by our data, the non-experimental variation in the dimensions of poverty is unlikely to be perfectly exogenous after the inclusion of observable controls and unobservable time-invariant personality traits. A causal interpretation of the coefficients may therefore not be warranted, yet we see no *a priori* reason to suspect that the coefficients would be systematically biased in one way or another.

5 Identifying the poor

The first step to measure poverty is to identify the poor, i.e., to make a comparison between the bundle of each individual and the poverty line vector z. As we have seen, both families of preference-sensitive poverty measures make this comparison in a different way. Consequently, they will identify different people as being poor. We illustrate this phenomenon using the RLMS-HSE data in the year 2000, the aftermath of the Russian financial crisis.

We start our analysis by setting the poverty line vector at 60% of the bundle that consists of the pooled median value in expenditures, health and housing. As we have discussed in Section 3, there are two alternative ways to treat the health variable. We start by considering health as a standard good. As reference value for the non-standard discrete (binary) unemployment variable the best possible value is selected according to each respondent's own preferences (as reported by Model 2 in Table 1). For most sociodemographic groups, the best possible value is "being employed". However, for respondents who are young, female, and lower educated the best possible value is "being unemployed".[15]

Let us return to Figure 2 to illustrate the identification of the poor. In this figure the poverty line vector in the income-health space is represented by the black dot. We consider first the family of quantity-metric multidimensional poverty measures in the top panel of

[15] Our sensitivity analysis has shown that assinging the reference value "being employed" to all respondents has only a very small effect on the results.

Figure 2. In this family, individuals are identified as being poor whenever they consider themselves worse off than the poverty line vector. Once the preferences are known for each respondent (using the life satisfaction method discussed in the previous section or any other method), it is straightforward to check whether they consider themselves worse off compared to the poverty line vector, or not. In the figure, the highlighted indifference curves distinguish the poor and non-poor outcome bundles for both preferences. Clearly, taking account of preference heterogeneity matters in the identification of the poor. Consider the young, higher educated women (the diamonds) situated in the south-east of the poverty line vector between both highlighted indifference curves. These individuals consider themselves to be worse off than the poverty line vector. However, if these individuals had the steeper (solid) indifference map, then they would not have considered themselves to be poor.

Second, we consider the family of money-metric multidimensional poverty measures. The measures in this family identify the poor by comparing the income to which an individual is indifferent, for some fixed reference price vector, with the money value of the poverty line bundle. Compared to the first family, an additional normative parameter needs to be set: the reference price vector p. For our illustration we select the price vector which is precisely tangent to the indifference curve of the average preference in the poverty line vector z (i.e., the indifference curve corresponding to Model 1 in Table 1, not shown in the figure). The bottom panel of Figure 2 illustrates the identification of the poor according to the money-metric measures, with the highlighted indifference curves distinguishing the poor and non-poor outcome bundles for both preferences. Also here preference heterogeneity clearly matters for the identification of the poor.

We also measure poverty under the assumption that health is a non-standard good and that perfect health is the reference value. That case would be represented in Figure 2 by a point at the extreme right of the figure, when the objective health indicator is equal to 1. Let us note that in this case quantity-metric and money-metric measures treat health in the same way. The difference between the two families is now restricted to the way they treat the trade-offs between expenditures and housing.

Next, we compare the characteristics of the poorest 16,1% individuals according to both preference-based measures and according to both ways in which health can be treated, with standard measures from the literature.[16] Each column of Table 2 gives the

[16]The counting approach leads to a discrete measure of intensity of mulitidimensional poverty, i.e., the number of dimensions for which the individual falls below the threshold. To avoid arbirariness in the identification of the poor for the counting approach, we take 16,1% as cut-off for the other methods as

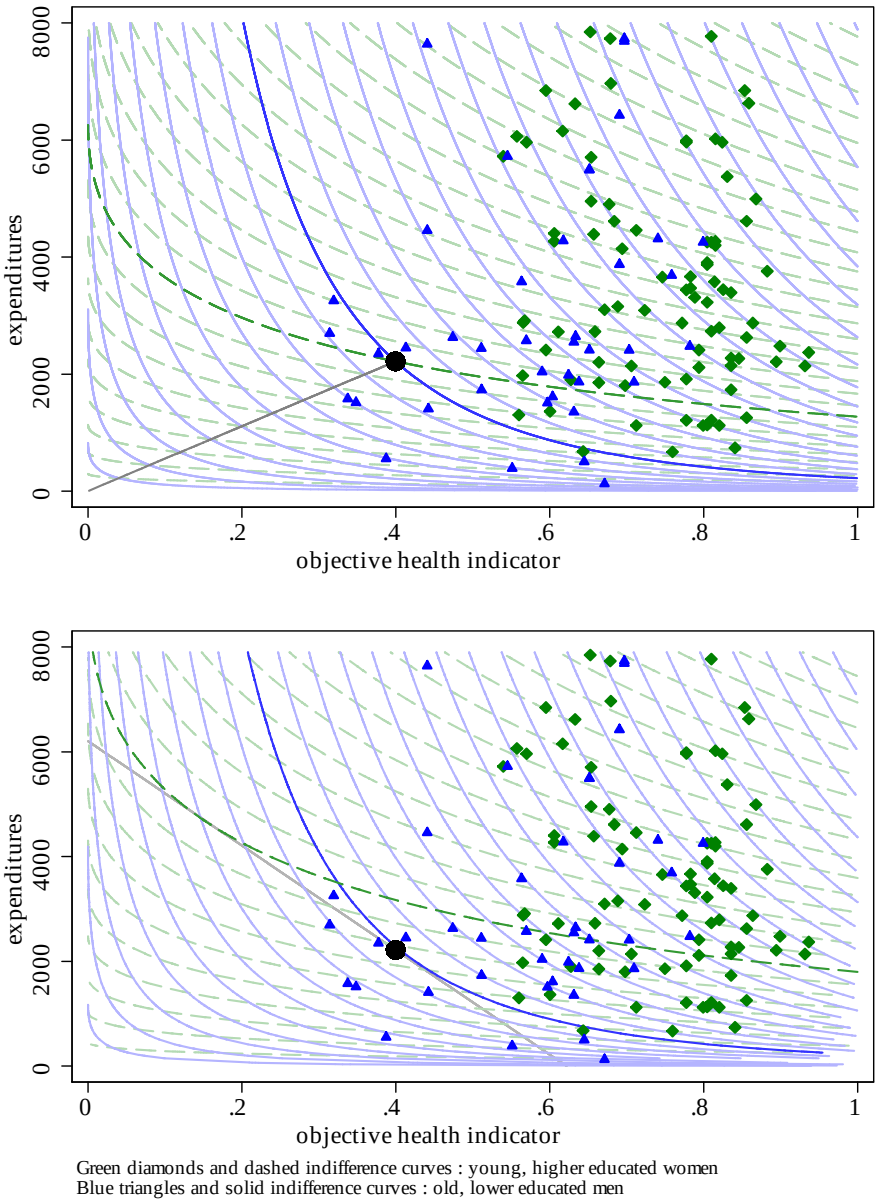


Figure 2: Indifference map of two subgroups in 2000. The top panel illustrates identification of the poor according to the quantity-metric measure, and the bottom panel illustrates identification of the poor according to the money-metric measure.

average characteristics of the poorest 16,1% in 2000 according to a specific identification method. The first column shows the characteristics of the full sample as benchmark. The second column shows the characteristics of the poor, identified according to the quantity-metric multidimensional poverty measure treating health as standard good. The poor are relatively old, predominantly female and in bad health. The third column zooms in on the characteristics of the poorest individuals according to the money-metric multidimensional poverty measure. The fourth and the fifth column presents the characteristics of the poorest, when health is treated as a non-standard bounded variable and the reference value is perfect health for all individual. As could be intuitively expected, the groups that are more sensitive to their health situation (relatively to the other goods) are more likely to end up in the poor population. The poor are now more old, higher educated, urban, and belonging to a minority compared to the previous columns. As this analysis heavily depends on the shape of the indifference curves in a part of the consumption set with few observations, we tend to prefer the case in which health is treated as a standard good.

Column six of Table 2 shows the characteristics of the poorest individuals identified according to a classical expenditure poverty measure. They have low expenditures, and are more often male and are younger compared to the preference-based methods treating health as a standard good. Finally, the seventh column provides the results according to the popular counting approach to multidimensional poverty. In the counting approach, the identification of the poor is based on the number of dimensions for which the individual falls below the threshold (see Alkire and Foster [3, 4]). In this illustration, we consider individuals as poor when they are below the threshold for at least two out of the four dimensions. The counting method identifies more people as poor who are unemployed, living in a relatively low quality house and who are lower educated. As different people are indeed identified as poor according to the three methods, it is clear that the question whether and how to take their preferences into account may have strong implications for the design of targeted poverty alleviation programs.

We conclude this section by considering in Figure 3 the degree of overlap between the bottom 16,1% individuals according to the four preference-based poverty measures and the standard expenditure and counting approaches. In the panels at the top, health is considered as a standard good. Slightly more than one third of the individuals are identified as belonging to the poor according to all measures. All other individuals are considered worst off by at least one method, but not by all methods. In fact, one finds

well. This value corresponds to the number of individuals who fall below the threshold in two dimensions.

Table 2: Portrait of the poor in 2000

	total sample	quantity metric	money metric	quantity metric	money metric	expenditure poverty	counting approach
treatment of health		standard	standard	non-stan.	non-stan.		
expenditure (in rubbles)	4302	1597	1774	3324	3418	985	1541
health (on 0-1 scale)	0,65	0,52	0,55	0,57	0,57	0,62	0,56
house (in 100.000 rubbles)	2,92	2,36	2,40	2,91	2,92	2,34	2,04
unemployed (in %)	8,4	26,4	28,1	24,1	23,7	14,4	26,8
life satisfaction	2,41	1,96	1,98	2,21	2,21	2,01	1,96
male (in %)	43,9	31,9	28,4	52,4	51,5	39,8	38,6
young (in %)	36,8	26,4	34,4	15,9	15,2	32,8	29,2
higher educated (in %)	66,2	58,7	65,8	82,6	82,4	55,4	50,8
rural (in %)	27,8	35,8	34,7	8,9	10,8	38,6	46,1
minority (in %)	14,4	8,7	15,3	55,3	58,4	14,7	17,1

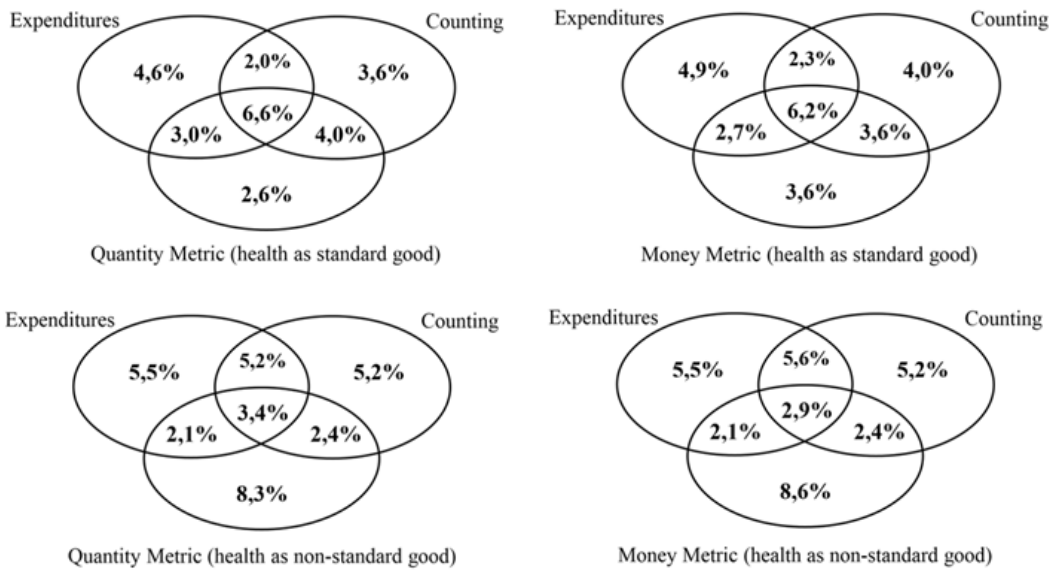


Figure 3: Overlap between the poor according to different approaches

remarkably little overlap between both multidimensional methods given that the same poverty line vector z is used. There is even less overlap in the panels at the bottom, when health is considered as a bounded variable and when preference heterogeneity plays a large role in the identification of the poor. This finding stresses once more the empirical implications of taking the preferences into account in the identification of the poor.

6 Robust poverty comparisons

In this section, we illustrate how poverty comparisons can be made that are robust for the specific choice of the f and g functions in the proposed families of multidimensional poverty measures. To do that, we make use of the so-called Three I's of Poverty" (TIP) dominance from unidimensional poverty analysis. For reasons of brevity, we focus on the quantity-metric poverty measure that treats health as a standard (unbounded) good.

Jenkins and Lambert [39] have shown that whenever there is TIP dominance of one distribution over another distribution, then there is unanimous agreement in the family of all additive poverty measures based on a decreasing and convex transformation function of the poverty gap that poverty is higher in the first distribution (see also Zheng [53]). Testing for TIP dominance can be done easily by the comparing the corresponding TIP curves. The TIP curve of the quantity-metric poverty measure plots the cumulative poverty gaps for the population ranked from poor to rich according to λ_z. In other words,

it consists of the pairs

$$\left(K/|N|, \sum_{k=1}^{K} 1 - \min(1, \lambda_z(x_{[k]}, R_{[k]})) \right) \text{ for all } K \leq |N|, \tag{6}$$

where for all respondents i and j we have that $i < j \Rightarrow \lambda_z(x_{[i]}, R_{[i]}) \leq \lambda_z(x_{[j]}, R_{[j]})$. Similarly we can define the TIP curve of the money-metric poverty measures by substituting λ_z for $\mu_{z,p}$ in equation (6).

In Figure 4 we have depicted the TIP curve of the quantity-metric poverty measure for each wave of the RLMS-HSE data between 1995 and 2005.[17] As these curves become flat above the poverty line, we only show their leftmost part. The most striking observation is that only few of the nine TIP curves cross, and, hence, that we get an almost complete ordering of the different waves according to multidimensional poverty. That means that the precise choice of the f or g function is not very decisive for the evaluation of the evolution of poverty in Russia over the considered period. Clearly, the TIP curve of 1998 is everywhere above the other curves. In other words, 1998 is unambiguously the year with most multidimensional poverty according to all poverty measures that belong to the family of quantity-metric multidimensional poverty measures. On the other hand, we see that the curve of 2005 is everywhere below the curve of 1995, indicating that poverty decreased over the considered period. The analysis can be repeated for the money-metric poverty measure. These TIP curves are not shown as they are very similar to Figure 4.

Finally, let us emphasize that our empirical application aims at illustrating the effect of adopting a preference-sensitive approach to the measurement of multidimensional poverty. It is not meant to be a definitive study on multidimensional poverty in Russia, which arguably requires a richer, larger and more tailored data set. Moreover, such a definitive study would require (bootstrapped) estimates of the sampling distribution of the poverty measures, the identification of the poor and the TIP curves. This goes beyond the scope of this paper. However, the findings of this section have shown that taking preferences into account makes a difference when identifying the poor in this sample and comparing poverty, even with a rather crude data set and approach to estimate preferences.

7 Conclusion

Measuring multidimensional poverty requires aggregating across dimensions and across individuals. In this paper, we have studied the consequences of aggregating across di-

[17] No data have been collected in 1997 and 1999.

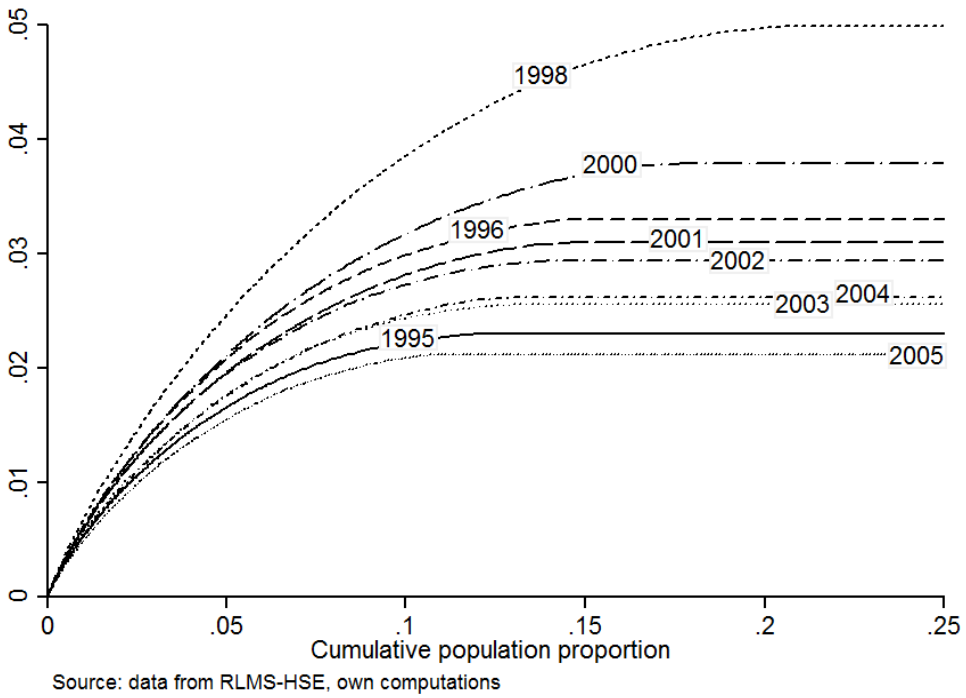


Figure 4: TIP curves for the quantity-metric multidimensional poverty measure for the different waves

mensions at the level of each individual by taking the individual's preferences as the aggregation device. This approach forced us to find new ways of aggregating across individuals, as individual levels of preference satisfaction cannot readily be compared. By introducing ways to build inter-preference poverty comparisons, we have been able to provide and characterize two families of poverty indices.

We have illustrated how the approach proposed in this paper can be implemented using existing Russian survey data from RLMS-HSE and we found some remarkable differences with standard (multidimensional) poverty measures. By taking preferences into account, different people are indeed identified as poor. The data that are needed to apply our approach are clearly more demanding than what is required to apply the other indices proposed in the literature. For instance, the counting approach is remarkably parsimonious in terms of the required data, whereas our approach requires a tailored data set that allows identification of the preferences in a wide set of dimensions. We believe that the data requirement of our approach is the price to pay to develop an attractive way to measure multidimensional poverty without relying on arbitrary weights or arbitrary assumptions on the nature of the goods.

Finally, taking preferences into account when measuring poverty has clear policy

implications, which we sketch here briefly. First, the different subgroups that are identified as poor with the preference-sensitive method call for a redirection of targeted poverty alleviation programs, e.g., those that are targeted to specific gender, age or education groups. In our empirical illustration with Russian data, for instance, a greater concern for the unemployed, the elderly, the urban, and women would be warranted. Second, the evaluation of poverty policies which have conflicting effects on different dimensions (e.g., improve health at some cost on income, or conversely) can be assessed in a more appealing way when population preferences are incorporated.

References

[1] Adler, M. D., P. Dolan and G. Kavetsos 2017, "Would you choose to be happy? Tradeoffs between happiness and the other dimensions of life in a large population survey" *Journal of Economic Behavior and Organization*, 139: 60–73.

[2] Akay, A., O. Bargain, X. H. Jara 2017, "'Fair' Welfare Comparisons with Heterogeneous Tastes: Subjective Versus Revealed Preferences" IZA Discussion Paper No. 10908.

[3] Alkire S., J. Foster 2011a, "Counting and multidimensional poverty measurement," *Journal of Public Economics* 95: 476–487.

[4] Alkire S., J. Foster 2011b, "Understandings and misunderstandings of multidimensional poverty measurement," *Journal of Economic Inequality* 9: 289–314.

[5] Atkinson A.B. 2003, "Multidimensional deprivation: contrasting social welfare and counting approaches," *Journal of Economic Inequality* 1: 51–65.

[6] Atkinson A.B., F. Bourguignon 1982, "The comparison of multidimensioned distributions of economic status," *Review of Economic Studies* 49: 183–201.

[7] Benjamin, D. J., O. Heffetz, M. Kimball, A. Rees-Jones 2012. "What do you think would make you happier? What do you think you would choose?" *American Economic Review*, 102: 2083–2110.

[8] Benjamin, D. J., O. Heffetz, M. Kimball, A. Rees-Jones 2014. "Can marginal rates of substitution be inferred from happiness data? Evidence from residency choices" *American Economic Review*, 104: 3498–3528.

[9] Beegle, K., K. Himelein, M. Ravallion 2012. "Frame-of-reference bias in subjective welfare" *Journal of Economic Behavior and Organization*, 81, 556–570.

[10] Bosmans, K. K. Decancq and E. Ooghe 2018, "Who is afraid of aggregating money metrics?" *Theoretical Economics*, 13(2), 467-484.

[11] Bosmans K., E. Ooghe, L. Lauwers 2018, "Prioritarian poverty comparisons with cardinal and ordinal attributes," *Scandinavian Journal of Economics*, forthcoming.

[12] Bossert W., S. R. Chakravarty, C. D'Ambrosio 2013, "Multidimensional Poverty and Material Deprivation with Discrete Data," *Review of Income and Wealth* 59: 29–43.

[13] Bourguignon F., S. R. Chakravarty 2003, "The measurement of multi-dimensional poverty," *Journal of Economic Inequality* 1: 25–49.

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

[15] Clark, A. E. and A.J. Oswald 2002, "A simple statistical method for measuring how life events affect happiness" *International Journal of Epidemiology*, 31: 1139–1144.

[16] Deaton A. 1979, "The distance function in consumer behaviour with applications to index numbers and optimal taxation," *Review of Economic Studies* 46: 391–405.

[17] Deaton, A., J. Muellbauer 1980, *Economics and Consumer Behavior,* Cambridge: Cambridge University Press.

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

[19] Decancq K., M. Fleurbaey M., E. Schokkaert 2017. "Well-being inequality and preference heterogeneity" *Economica*, 84: 210–238.

[20] Decancq, K. and D. Neumann, "Does the choice of well-being measure matter empirically?" *in* M. Adler and M. Fleurbaey (Eds.) *The Oxford Handbook of Well-Being and Public Policy*, Oxford University Press.

[21] Decancq, K. and A. Nys 2018, "Nonparametric Well-being Comparisons" *Mimeo.*

[22] Decancq, K. and E. Schokkaert 2016, "Beyond GDP: Using equivalent incomes to measure well-being in Europe", *Social Indicators Research* 126: 21–55.

[23] Dolan, P., D. Fujiwara 2016, "Happiness-based policy analysis" *in* M. Adler and M. Fleurbaey (Eds.) *The Oxford Handbook of Well-Being and Public Policy*, Oxford University Press.

[24] Duclos J.-Y., D.E. Sahn, S.D. Younger 2006, "Robust multidimensional poverty comparisons," *Economic Journal* 116: 943–968.

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

[26] Fleurbaey M. 2007, "Two criteria for social decisions", *Journal of Economic Theory*, 134, 421–447.

[27] Fleurbaey, M., Luchini, S., Muller, C. and E. Schokkaert 2013, "Equivalent income and fair evaluation of health care" *Health Economics*, 22: 711–729.

[28] Fleurbaey M., F. Maniquet 2008, "Fair Social Orderings" *Economic Theory* 34: 25–45.

[29] Fleurbaey M., F. Maniquet 2011, *A Theory of Fairness and Social Welfare*, Cambridge: Cambridge University Press.

[30] Fleurbaey M., F. Maniquet 2017a, Fairness and well-being measurement, *Mathematical Social Sciences*, forthcoming.

[31] Fleurbaey M., F. Maniquet 2017b, Inequality averse well-being measurement, *International Journal of Economic Theory*, forthcoming.

[32] Fleurbaey M., F. Maniquet 2017c, Fairness and well-being measurement with non-classical goods, *Mimeo.*

[33] Fleurbaey M., K. Tadenuma 2014, "Universal social orderings," *Review of Economic Studies* 81: 1071–1101.

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

[35] Foster J., A. F. Shorrocks 1991, "Subgroup consistent poverty indices," *Econometrica* 59: 687–709.

[36] Frijters P., I. Geishecker, J.P. Haisken-DeNew, M. 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.

[37] Gevers, L. 1986, "Walrasian Social Choice: some simple axiomatic approaches". In: Heller, W., et al. (eds.) Social Choice and Public Decision Making. Essays in Honor of K.J. Arrow. Cambridge: Cambridge University Press.

[38] Graham, C., A. Eggers, S. Sukhtankar 2004, "Does happiness pay? An exploration based on panel data from Russia," *Journal of Economic Behavior and Organization* 55: 319–342.

[39] Jenkins, S.P., P.J. Lambert 1997, "Three I's of poverty curves, with an analysis of UK poverty trends," *Oxford Economic Papers* 49: 317–327.

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

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

[42] Maasoumi E., M.A. Lugo 2008, "The information basis of multivariate poverty assessments", in N. Kakwani and J. Silber (eds.) "Quantitative Approaches to Multidimensional Poverty Measurement", Palgrave-MacMillan.

[43] Maasoumi E., J.S. Racine 2016, "A solution to aggregation and an application to multidimensional 'well-being' frontiers" *Journal of Econometrics* 191(2): 374–383.

[44] Narayan, D., R. Patel, K. Schafft, A. Rademacher, S. Koch-Schulte 2000, *Voices of the Poor: Can Anyone Hear Us?* New York, N.Y.: Published for the World Bank, Oxford University Press.

[45] Pazner E., D. Schmeidler 1978, "Egalitarian equivalent allocations: A new concept of economic equity," *Quarterly Journal of Economics* 92: 671–687.

[46] Ravallion M. 2011, "On multidimensional indices of poverty," *Journal of Economic Inequality* 9: 235–248.

[47] Samuelson P.A. 1977, "Reaffirming the existence of "reasonable" Bergson-Samuelson social welfare functions," *Economica* 44: 81–88.

[48] Samuelson P.A. and S. Swamy 1974, "Invariant economic index numbers and canonical duality: survey and synthesis", *American Economic Review* 64: 566–593.

[49] Sen A. K. 1976, "Poverty: an ordinal approach to measurement," *Econometrica* 44: 219–231.

[50] Sprumont, Y., Zhou, L. 1999, "Pazner–Schmeidler Rules in large societies". *Journal of Mathematical Economics* 31, 321–339.

[51] Tsui K.Y. 2002, "Multidimensional poverty indices," *Social Choice and Welfare* 19: 69–93.

[52] van Doorslaer E., A.M. Jones 2003, "Inequalities in self-reported health: validation of a new approach to measurement," *Journal of Health Economics* 22: 61–87.

[53] Zheng B. 2000, "Poverty orderings," *Journal of Economic Surveys* 14: 427–466.

Appendix 1. Data and estimation

Data

In the empirical sections, we use the Russian Longitudinal Monitoring Survey (RLMS-HSE) data set between 1995 and 2005. Data are collected on a yearly basis, with exception of 1997 and 1999, when no data have been collected. The raw (unbalanced) data set contains 68.044 respondents in 9 waves, with yearly sample sizes between 5.395 (in 1995) and 9.172 (in 2002). The estimation sample contains 53.783 respondents. Although the RLMS-HSE data set is designed to be nationally representative for the Russian Federation, it is unlikely that our final estimation sample is nationally representative for several reasons. First, there may be non-random attrition from the panel. Second, we can only retain observations in our analysis for whom we have non-missing data for all variables of interest. Finally, as we will see below, the estimation method forces us to drop respondents with time-invariant life satisfaction scores. Given the non-random selection of our final estimation sample, our findings should be interpreted cautiously as results about the sample. For that reason, no sample weights have been used in the analysis.

The four goods we consider are constructed as follows.[18] First, we use the square root of household size as equivalence scale for the real household expenditures (the reference year is 1992). Second, the measure of individual health is a composite index that consists of various objective health indicators such as indicators of diabetes, heart attack, aneamia, hospitalization and recent operations. The weights of the indicators in the health index are derived from the coefficients of an ordered logit regression using the self-assessed health as explained variable (for a similar approach, see van Doorslaer and Jones [52]). Third, the measure of housing quality is the predicted value of a logarithmic hedonic housing regression based on self-reported housing values and a series of housing characteristics after controlling for regional price differences, time trends and the size of the household (by using again the square root equivalence scale).[19] Finally, the measure of unemployment is a binary indicator that takes 1 if the respondent is unemployed. The unemployment rate in the considered period fluctuates around 8%. Table 3 provides summary statistics. The correlation between the standard variables is overall relatively low: in 2000 it is highest between expenditures and housing quality (0.181) and lowest between health and housing quality (-0.064).

[18] Estimation results are available upon request.

[19] In principle, one can argue that, in the preference approach, equivalence scales, the measure of individual composite health and housing quality should all be determined using individual preferences.

Life satisfaction is measured in the RLMS-HSE by the following question: *"To what extent are you satisfied with your life in general at the present time?"*. Respondents use a five point-scale, ranging from "not at all satisfied" to "fully satisfied" to answer this question.

Estimation

To estimate the prefereces we estimate the following econometric model:

$$S_{it}^{*} = \alpha_i + \gamma_t + (\beta + \Lambda D_{it})' X_{it} + \delta' Z_{it} + v_{it}.$$

We observe the reported life satisfaction $S_{it} = k$ for k in $\{1, 2, \ldots, 5\}$ if the latent life satisfaction (S_{it}^{*}) lies within an interval between η_{k-1} and η_k :

$$S_{it} = k \text{ if } \eta_{k-1} < S_{it}^{*} \leq \eta_k.$$

The thresholds η_k are allowed to depend on the individual fixed effects, the observable socio-demographic characteristics and the time trend (for more details, see Jones and Schurer [40]). Finally, to allow for non-perfect substitutability between the three continuous dimensions of poverty, the outcomes of the individuals in these dimensions are transformed by a so-called Box-Cox transformation. For each dimension j in {expenditures, health, housing}, we have that:

$$X_{it}^{j} = \begin{cases} \left[\left(Y_{it}^{j} \right)^{\varepsilon_j} - 1 \right] / \varepsilon_j & \text{when } \varepsilon_j \neq 0 \\ \log \left(Y_{it}^{j} \right) & \text{when } \varepsilon_j = 0, \end{cases}$$

where Y_{it}^{j} is the observed outcome of individual i in period t in dimension j. For the binary unemployment variable, $Y_{it} = X_{it}$. The three dimension-specific transformation parameters ε_j are chosen over a fine grid to maximize the overall fit of the model (Box and Cox [14]).

We follow the estimation method suggested by Jones and Schurer [40] that approximates the approach proposed by Ferrer-i-Carbonell and Frijters [25] and applied by Frijters et al. [36] to the RLMS-HSE. This method approximates the above model by a logit regression with as explained variable a binary indicator whether the respondent is above or below her personalized average life satisfaction score across all waves. About 6% of respondents have time-invariant life satisfaction scores and bring no variation to the model. They have been dropped from the sample.

Table 3: Summary Statistics

		mean	stand. dev.	minimum	maximum
1995	life satisfaction	2,216	1,058	1	5
	expenditure (in rubbles)	4.986	4.530	6,570	33.202
	health (on 0-1 scale)	0,648	0,169	0,031	1
	house (in 100.000 rubbles)	2,788	1,038	0,61	9,004
	unemployed (in %)	0,07	0,256	0	1
1996	life satisfaction	2,177	1,052	1	5
	expenditure (in rubbles)	4.781	4.737	78,75	35.941
	health (on 0-1 scale)	0,647	0,173	0,111	0,99
	house (in 100.000 rubbles)	2,783	1,057	0,622	8,394
	unemployed (in %)	0,083	0,276	0	1
1998	life satisfaction	2,043	1,030	1	5
	expenditure (in rubbles)	4.122	4.453	35,45	33.611
	health (on 0-1 scale)	0,638	0,18	0,087	0,995
	house (in 100.000 rubbles)	2,926	1,116	0,456	8,181
	unemployed (in %)	0,091	0,287	0	1
2000	life satisfaction	2,405	1,101	1	5
	expenditure (in rubbles)	4.302	4.692	30,37	40.623
	health (on 0-1 scale)	0,647	0,183	0,031	0,995
	house (in 100.000 rubbles)	2,922	1,195	0,456	17,79
	unemployed (in %)	0,084	0,278	0	1
2001	life satisfaction	2,588	1,114	1	5
	expenditure (in rubbles)	4.818	4.515	42,24	36.386
	health (on 0-1 scale)	0,635	0,184	0,004	0,99
	house (in 100.000 rubbles)	3,028	1,164	0,612	8,110
	unemployed (in %)	0,077	0,267	0	1
2002	life satisfaction	2,858	1,119	1	5
	expenditure (in rubbles)	5.028	4.787	10,39	40.330
	health (on 0-1 scale)	0,636	0,184	0,001	0,99
	house (in 100.000 rubbles)	3,049	1,196	0,551	17,9
	unemployed (in %)	0,082	0,275	0	1
2003	life satisfaction	2,762	1,137	1	5
	expenditure (in rubbles)	5.477	5.294	17,25	41.990
	health (on 0-1 scale)	0,63	0,187	0,018	0,99
	house (in 100.000 rubbles)	3,052	1,199	0,528	12,63
	unemployed (in %)	0,082	0,275	0	1
2004	life satisfaction	2,893	1,097	1	5
	expenditure (in rubbles)	5.645	5.597	50,49	50.103
	health (on 0-1 scale)	0,628	0,185	0,061	0,995
	house (in 100.000 rubbles)	3,077	1,203	0,514	12,63
	unemployed (in %)	0,084	0,277	0	1
2005	life satisfaction	2,916	1,094	1	5
	expenditure (in rubbles)	5.935	5.726	91,51	46.797
	health (on 0-1 scale)	0,625	0,185	0,061	0,99
	house (in 100.000 rubbles)	3,069	1,169	0,608	16,25
	unemployed (in %)	0,083	0,277	0	1

To reach a parsimonious and tractable model, the least significant interaction effects in Λ have been dropped consequently until all the remaining interaction effects are significant at the 10% level. All reported standard errors are corrected for clustering at the household level.

As can be seen from Table 1, the likelihood maximizing Box-Cox transformation parameters are respectively $-0,055$; $0,485$ and $-0,356$ for expenditures, health, and housing quality. The approximately logarithmic transformation of expenditures is a common finding in the happiness literature, suggesting important decreasing marginal returns of expenditures (see, amongst others, Layard et al. [41]). The marginal returns for health and housing are found to be respectively less and more decreasing.

Appendix 2. Proof of Theorem 1

Proof. We focus on the "only if" part. Let P be a poverty index satisfying the axioms. Let $N \subset \mathbb{N}_{++}$. By *Subgroup Consistency* applied by decomposing N into $\{i\}$ and $N \backslash \{i\}$, there exist functions $g : X^{N\backslash\{i\}} \times \mathcal{R}^{N\backslash\{i\}} \to \mathbb{R}$ and $f : \mathbb{R}^2 \to \mathbb{R}$ such that for all $(x_N, R_N) \in \mathcal{S}$,

$$P(x_N, R_N) = f(P(x_i, R_i), g(x_{N\backslash\{i\}}, R_{N\backslash\{i\}})).$$

By applying the same argument to decompose g and the following functions, we reach the conclusion that there exists a function $f : \mathbb{R}^{|N|} \to \mathbb{R}$ such that for all $(x_N, R_N) \in \mathcal{S}$,

$$P(x_N, R_N) = f((P(x_i, R_i))_{i\in N}).$$

Fix R_N for the moment. The line segment from 0 to $z(R_i)$ is a compact set. By *Focus*, the range of $P(x_i, R_i)$ for fixed R_i is the image of this line segment. By continuity, the range of $P(x_i, R_i)$ is a compact interval, of strictly positive length by *Pareto among the Poor*.

We claim that f is continuous on this domain. Indeed, take a sequence $p_N^k \to p_N^*$ in this domain and suppose $f\left(p_N^k\right)$ does not tend to $f\left(p_N^*\right)$. It then has a subsequence p_N^t such that $f\left(p_N^t\right)$ stays away from $f\left(p_N^*\right)$. There is a corresponding sequence x_N^t on the line segments from 0 to $(z(R_i))_{i\in N}$ such that $p_N^t = (P(x_i^t, R_i))_{i\in N}$ for all t. As the sequence is in a compact set, it has a converging subsequence $x_N^s \to x_N^*$. By continuity, $P(x_N^s, R_N) \to P(x_N^*, R_N)$. Also by continuity, $p_N^s \to p_N^*$ implies that $p_N^* = (P(x_i^*, R_i))_{i\in N}$. Therefore $f\left(p_N^*\right) = P(x_N^*, R_N)$ is the limit of $f\left(p_N^s\right)$. This contradicts the assumption that $f\left(p_N^s\right)$, a subsequence of $f\left(p_N^t\right)$, stays away from $f\left(p_N^*\right)$.

Now let us drop the restriction on R_N. By *Focus*, one can normalize $P(x_i, R_i) = 0$ whenever $x_i R_i z\left(R_i\right)$. Therefore the range of $P(x_i, R_i)$ is the union of the intervals for every possible R_i, all these intervals having 0 as their lower bound, so that their union is an interval. The f function being continuous on any of these (Cartesian products of) intervals, it is continuous on (the Cartesian product of) their union.

So, f is a real-valued function defined over the Cartesian product of connected and separable spaces (intervals). By *Pareto among the Poor*, each dimension of the domain of f is essential (f cannot be constant in any $P(x_i, R_i)$). Therefore, by *Continuity* and *Subgroup Consistency*, we can apply Debreu's theorem on additive representation of separable preferences: there exists a continuous and strictly increasing function $G_N : \mathbb{R} \to \mathbb{R}$ and for all $i \in N$ a continuous function $f_i : \mathbb{R} \to \mathbb{R}$ such that

$$P(x_N, R_N) = G_N \left[\sum_{i=1}^{|N|} f_i(P(x_i, R_i)) \right].$$

By defining functions $\phi_i^N : X \times \mathcal{R} \to \mathbb{R}$ as follows: $\phi_i^N = f_i \circ P$, we get

$$P(x_N, R_N) = G_N \left[\sum_{i=1}^{|N|} \phi_i^N(x_i, R_i) \right].$$

By continuity, ϕ_i^N are continuous (in their first argument). By *Pareto among the Poor*, $\phi_i^N(x_i, R_i) > \phi_i^N(x_i', R_i)$ whenever $z(R_i)\, R_i\, x_i'\, P_i\, x_i$. By *Focus*, we can normalize $\phi_i^N(x_i, R_i) = 0$ whenever $x_i\, R_i\, z(R_i)$. Note that defining functions G_N, one for each N, or defining one function G with N as an argument, as in the statement of the Theorem, are equivalent.

In order to prove the second part of the Theorem, let us begin by observing that any permutation of individuals in (x_N, R_N) yields a 1-replica of (x_N, R_N). That shows that *Replication Invariance* implies that P has to be anonymous: the name of the individuals does not matter. As a result, we can remove the subscript from the ϕ_i^N functions, and the second argument of the G function reduces to the number of individuals in N.

Let us assume that the claim of the theorem is wrong. For all G and ϕ^N such that $P(x_N, R_N) = G\left[\sum_{i=1}^{|N|} \phi(x_i, R_i), N\right]$, there exist $(x_N, R_N), (x'_{N'}, R'_{N'}) \in \mathcal{S}$ such that

$$\frac{1}{|N|} \sum_{i=1}^{|N|} \phi(x_i, R_i) = \frac{1}{|N'|} \sum_{i=1}^{|N'|} \phi(x_i', R_i') \tag{7}$$

and

$$P(x_N, R_N) \neq P(x'_{N'}, R'_{N'}). \tag{8}$$

Let $n = |N|$ and $n' = |N'|$. By *Replication Invariance*,

$$P(x_N, R_N) = P(x_N^{n'}, R_N^{n'}) = G\left[n' \sum_{i=1}^{|N|} \phi(x_i, R_i), nn' \right],$$

and

$$P(x'_{N'}, R'_{N'}) = P({x'_{N'}}^n, {R'_{N'}}^n) = G\left[n \sum_{i=1}^{|N'|} \phi(x_i', R_i'), nn' \right].$$

Eq. 7 implies that

$$n' \sum_{i=1}^{|N|} \phi(x_i, R_i) = n \sum_{i=1}^{|N'|} \phi(x_i', R_i'),$$

in contradiction to Eq. 8. ■