

Welfare, Inequality and Poverty Analysis with `rtip`: An Approach Based on Stochastic Dominance

by Angel Berihuete, Carmen D. Ramos and Miguel A. Sordo

Abstract Disparities in economic welfare, inequality and poverty across and within countries are of great interest to sociologists, economists, researchers, social organizations and political scientists. Information about these topics is commonly based on surveys. We present a package called `rtip` that implements techniques based on stochastic dominance to make unambiguous comparisons, in terms of welfare, poverty and inequality, among income distributions. Besides providing point estimates and confidence intervals for the most commonly used indicators of these characteristics, the package `rtip` estimates the usual Lorenz curve, the generalized Lorenz curve, the TIP (Three I's of Poverty) curve and allows to test statistically whether one curve is dominated by another.

Introduction

Surveys of income provide policy makers, researchers, social organizations and general public with a rich source of data to address and understand the topics of welfare, income inequality and poverty. In one approach, income differences among states, populations or groups are summarised by univariate indices. The most popular aggregate inequality and poverty indices (such as the Gini index and the at-risk-of-poverty rate) summarize inequality or poverty by a univariate index. Such an index is easy to compare over time, across regions or countries but may be insufficient for more detailed research and refined decision making. Increasing availability of data sets from a variety of statistical sources (including national and international statistical offices) allows scientists and policy-makers to carry out detailed comparisons by developing and applying more sophisticated tools and graphical representations from real data sets.

This paper presents an R package, called `rtip`, that implements methods and techniques based on stochastic dominance to make unambiguous comparisons, in terms of welfare, poverty and inequality, of income distributions. Stochastic dominance requires unanimity in rankings for large classes of indices (rather than simpler rankings based on a single index). We regard one income distribution as dominating another only if the same ranking is obtained for them with an entire family of indices. Stochastic dominance rules for income distributions can be easily implemented by seeking a dominance relation between three graphical representations: the Lorenz curve, the generalized Lorenz curve and the TIP (Three I's of Poverty) curve (also called the cumulative poverty gap (CPG) curve or the poverty profile curve, see Barrett et al. (2016) for other names). Two income distributions can be unambiguously ranked in terms of inequality if their Lorenz curves do not intersect (see Atkinson, 1970). Similarly, two income distributions can be unambiguously ranked with respect to a wide class of reasonable social welfare functions (see Shorrocks, 1983, for details) if their respective generalized Lorenz curves do not intersect. The literature also contains results (see Jenkins and Lambert, 1998a,b; Sordo and Ramos, 2011) connecting unambiguous poverty rankings by general classes of poverty measures with non-intersections of TIP curves. The package `rtip` estimates these curves directly from data sets and implements some tests to assess any statistical significance.

We can find some functions, in different software languages, addressing the analysis of income distributions via generalized Lorenz curves and their associated indices of inequality. For example, the Stata commands¹ `glcurve` (Jenkins and Van Kerm, 2004; Van Kerm and Jenkins, 2001), `svylorenz` (Jenkins, 2006), `clorenz` (Araar, 2005), `alorenz` (Azevedo and Franco, 2006) and `lorenz` (Jann, 2016), estimate, draw and provide variance estimates for generalized Lorenz curves. The command `povdeco` (Jenkins, 1999) estimates the poverty indices from the Foster et al. (1984) family. With `rtip`, we want to make these methods available to R users². We can find some packages in the R environment to analyse inequality and poverty of income distributions. For example, various indices of poverty and inequality are included in the packages `IC2` (Plat, 2012) and `ineq` (Zeileis, 2014), which also provide some graphical tools, including the Lorenz curve. A comprehensive collection of indicator methodology is included in the package `laeken` (Alfons and Templ, 2013), which implements functions and generalized Lorenz curves to estimate a wide set of social inclusion and poverty indicators in

¹user-written functions in Stata are called commands.

²This code was developed in the context of a project undertaken by the authors (Sordo et al., 2014) to compare, by using stochastic dominance techniques, inequality and poverty in Andalusia with other European regions. `rtip` has been also applied to the study of inequality by using tax return data (Sordo et al., 2017).

complex surveys. An interesting book combining EU-SILC surveys (the survey on income and living conditions conducted by Eurostat, the statistical office of the EU) and R is Longford (2015).

Some strengths of using the package **rtip** presented in this paper compared with others are the following:

- (a) **rtip** provides functions to load microdata having EU-SILC format (see Datasets section).
- (b) **rtip** evaluates income distributions on several equivalence scales and with adjustment for differences in household composition. The choice of the equivalence scale is important and can considerably affect the results (see, for example, Buhmann et al. (1988); Jenkins and Cowell (1994); De Vos and Zaidi (1997), who showed that the choice of the equivalence scale is not innocuous). Although the package **rtip** employs by default the OECD-modified scale used by Eurostat, it also allows to use the parametric scale in Buhmann et al. (1988), which covers almost all equivalence scales used in practice.
- (c) Besides providing point estimates and confidence intervals for the most commonly used indicators for inequality, welfare and poverty, the package **rtip** implements estimation of the usual Lorenz curve, the generalized Lorenz curve and the TIP curve and provides a way, based on the distribution-free test for generalized Lorenz and deprivation dominance suggested by Xu (1997) and Xu and Osberg (1998), respectively, to test statistically whether a curve is dominated by another.

The remainder of this paper is organized as follows. First, we explain the processes of loading data and setting up the surveys. Then, we briefly describe the indices and curves included in the package, explaining how to estimate them with **rtip**. Description of the statistical procedures implemented to test dominance of these curves is also presented and illustrated with examples. Finally we provide conclusions as well as a section to explain to users how to load their own datasets.

Datasets

The **rtip** package offers the possibility to load data extracted from EU-SILC using the `loadEUSILC` function. Users can modify `loadEUSILC` in order to load their own datasets. For instance, `loadLCS` function is a modification of `loadEUSILC` to load data from the Spanish Living Conditions Survey (see the Appendix). Note that **rtip** is not restricted to datasets with EU-SILC format and users can load their own datasets.

The **rtip** package contains two datasets: `eusilc2` and `LCS2014`.

```
library(rtip)
data(eusilc2)
data(LCS2014)
```

The `eusilc2` dataset is a modification of the synthetic `eusilc` dataset in the package **laeken** (see Alfons and Templ, 2013). The aim of this modification³ is to set up the variables according to expectations of the **rtip** package (see Table 1). The `eusilc` dataset is related to Austrian EU-SILC data from 2006, and was generated to comply with EU-SILC confidentiality rules. The `LCS2014` dataset⁴ is a data frame containing a selection of variables from the living conditions survey released in 2014 by the Spanish National Statistics Institute (INE in Spanish).

Both datasets, `eusilc2` and `LCS2014`, have seven variables which are briefly described in Table 1 (for further information see Eurostat, 2007). The records in these datasets are households, not individuals. Thus the 14,827 individual-level records in the `eusilc` dataset are condensed to 6,000 household-level records in the `eusilc2` dataset.

Setting up the surveys

The income variable most commonly studied for the inequality and poverty assessment is the disposable household income. Most official statistics (and **rtip**) use the equivalised disposable household income, which is the total disposable income of the household adjusted by taking into account its

³The code used to modify the `eusilc` dataset can be found at https://github.com/AngelBerihuete/rtip/blob/master/data-raw/eusilc_eusilc2.R.

⁴The dataset can be obtained at http://www.ine.es/dyngs/INEbase/en/operacion.htm?c=Estadistica_C&cid=1254736176807&menu=ultiDatos&idp=1254735976608.

Name of the variable	Meaning
DB010	Year of the survey
DB020	Country
DB040	Region
DB090	Household cross-sectional weight
HX040	Household size
HX050	Equivalised household size
HX090	Equivalised disposable income

Table 1: Variables needed before setting up the dataset. Variables HX050 and HX090, provided by Eurostat, are calculated using the OECD-modified scale.

composition (number of adults and children). The adjustment is made by dividing the disposable household income by the equivalised household size, which is defined as the number of household members converted into equivalent adults by using a specific equivalence scale. The equivalisation factor employed by Eurostat is the OECD-modified scale, which gives a weight of 1.0 to the first person aged 14 or more, a weight of 0.5 to other persons aged 14 or more and a weight of 0.3 to persons aged 0-13. `rtip` uses this equivalisation factor by default, but can also use the parametric scale of [Buhmann et al. \(1988\)](#). In this case, the equivalised household size is given by n^s , where n is the number of members of the household and s is a parameter known as elasticity of equivalence, with $0 \leq s \leq 1$. The value $s = 0.5$ is frequently used to make comparisons between countries (when $s = 0$ the composition is irrelevant).

The information is held in two files: basic household register (H-file) and household data (D-file). These files are loaded by functions `loadEUSILC` or `loadLCS` from EU-SILC or INE surveys, respectively. Next, the data is set up by using the function `setupDataset` which has one mandatory argument (`dataset`) and five optional arguments: the country to be analysed (`country`), the region of the country (`region`), the equivalence scale (`s`), a deflator (`deflator`) and the purchasing power parity rate (`pppr`). All optional arguments have NULL default value, except for `country = 'ES'`.

The country and the region are expressed using the nomenclature of territorial units for statistics⁵. All regions of the country will be selected by default (`region = NULL`) but it is possible to select one or more regions using a character string. Income is expressed in current monetary units. If `deflator` is not NULL the value assigned will be used as a deflator. Finally, by setting up the ratio of the purchasing power parity conversion factor to market exchange rate (`pppr`) we can compare the income across countries.

```
dataset <- setupDataset(eusilc2, country = "AT", region = NULL,
  s = NULL, deflator = NULL, pppr = NULL)
```

Next, we form a data frame with a new variable `ipuc`. The `ipuc` variable is the income per unit of consumption⁶ (or equivalised disposable income) which also takes into account the deflator and the purchasing power parity rate. If `deflator = NULL`, `pppr = NULL` and `s = NULL`, `ipuc` is set to `HX090` (see Table 1).

Indicators and curves

In this section we briefly describe some indicators and curves that are widely used in the study of poverty, inequality and welfare. The `eusilc2` dataset contains all the data necessary for estimating them.

Let X be a random variable⁷ with cumulative distribution function F , and let $\zeta_p = F^{-1}(p)$ be the quantile function with $0 \leq p \leq 1$. Let n be the number of observations in the sample, let $x := \{x_i\}_{i=1}^n$ be individual incomes sorted into ascending order so that $x_1 \leq x_2 \leq \dots \leq x_n$, and let $\omega := \{\omega_i\}_{i=1}^n$ denote the corresponding sample weights⁸. Weighted quantiles for the estimation of the population values are given by $\hat{\zeta}_p(x, \omega) = x_{(r)}$, where $x_{(r)}$ is the r -th order statistic such that $r = \lfloor p \sum_{i=1}^n \omega_i \rfloor$ is the closest integer not greater than $p \sum_{i=1}^n \omega_i$. $\hat{\zeta}_p(x, \omega)$ is the sample quantile level such that $100p$ percent

⁵Nomenclature of territorial units can be found at <http://ec.europa.eu/eurostat/web/nuts/overview>.

⁶The number of units of consumption is the number of household members converted into equivalent adults by using a specific equivalence scale.

⁷Unless otherwise stated it refers to equivalised disposable household income.

⁸Following indications in [Alfons and Templ \(2013\)](#), we take into account sample weights in the estimation of indices and curves.

of the weighted sample of observations is less than or equal to $\hat{\zeta}_p(x, \omega)$ and $100(1 - p)$ percent of the weighted observations is greater (Beach and Kaliski, 1986).

Indicators of poverty, inequality and welfare

Poverty curves and indicators are based on the poverty threshold that distinguishes between poor and non-poor households. The median income, $\zeta_{0.5}$, is frequently used for this purpose. On the recommendation of Eurostat this threshold, called the *at-risk-of-poverty threshold*, is set at 60% of the national median equivalised disposable income ($arpt = 0.6 \cdot \zeta_{0.5}$) and is estimated by

$$\widehat{arpt} = 0.6 \cdot \hat{\zeta}_{0.5}(x, \omega), \quad (1)$$

using the function `arpt()`. The default setting of 60% can be easily changed via the argument `pz`. Optional variable names can be set up for the income per unit of consumption (`ipuc`), the household cross-sectional weight (`hhcsw`) and the household size (`hhsiz`). Default values correspond to EU-SILC names. For the sake of clarity we use the default variable names in the following examples. A sample call to `arpt` would look like:

```
arpt(dataset, ipuc = "ipuc", hhcsw = "DB090", hhsiz = "HX040", pz = 0.6)
```

The indicator called *at-risk-of-poverty rate*, defined as the proportion of persons with an equivalised disposable income below the *at-risk-of-poverty threshold*, is estimated using the function `arpr()` (see Table 2). A sample call to `arpr` would look like:

```
arpr(dataset, arpt.value = arpt(dataset, pz = 0.6))
```

Using the function `arpr()` and changing the argument `pz` to 40%, 50% and 70% in `arpt()`, we obtain the *dispersion around the at-risk-of-poverty threshold* (percentage of persons with an equivalised disposable income below 40%, 50% and 70% of the national median equivalised disposable income, respectively). For example:

```
arpr(dataset, arpt.value = arpt(dataset, pz = 0.4))
```

Table 2 contains a brief description of some other well-known poverty indicators such as the Foster, Greer and Thorbecke (FGT_1) poverty index (Foster et al., 1984). This index is calculated by the function `s1()`. Its normalized version, also called the poverty gap ratio, is obtained by setting `norm = TRUE`. In this case, the index provides the average of the ratio of the poverty gaps⁹ to the at-risk-of-poverty threshold.

```
s1(dataset, arpt.value = arpt(dataset), norm = TRUE)
```

The `rtip` package also provides the function `s2()` to calculate the Sen-Shorrocks-Thon (SST) poverty index (Shorrocks, 1995; Zheng, 1997). It is estimated as twice the area below the TIP curve (see TIP curve in next section).

For income inequality, we calculate the most commonly used indices of inequality, the *Gini index*, and the *quintile share ratio* (see Table 2). The *mean income per person, per household and per unit of consumption* are estimated by the respective functions `mip()`, `mih()` and `miuc()`.

Confidence intervals can be obtained for all the indicators in `rtip` package by a bootstrap method. We use the `boot` package (Canty and Ripley, 2016) to generate bootstrap replicates. The user can change the number of replicates with the parameter `rep` which, by default, is set to 500. If `verbose = TRUE` we obtain a plot showing the histogram of the indicator estimations (replicates) and the Standard Normal quantile-quantile plot of bootstrap estimates. For instance, in case of the `arpt` function and 98% confidence interval, a typical call would look like:

```
arpt(dataset, pz = 0.6, ci = 0.98, rep = 500, verbose = TRUE)
```

Curves for inequality, welfare and poverty

Lorenz and generalized Lorenz curve

The Lorenz curve (Gastwirth, 1971) is defined as

⁹A poverty gap is the difference between the *at-risk-of-poverty threshold* and the equivalised disposable income, with the non-poor being given a difference of zero.

Index name, function	Description	Formula
<i>At-risk-of-poverty rate</i> , arpr()	Proportion of persons with an equivalised disposable income below the <i>at-risk-of-poverty threshold</i>	$\frac{\sum_{i=1}^n \omega_i \mathbf{I}(x_i \leq \widehat{arpt})}{\sum_{i=1}^n \omega_i} \cdot 100 \quad (2)$
<i>Relative median at-risk of-poverty gap</i> , rmpg()	Difference between the <i>at-risk-of-poverty threshold</i> and the median equivalised disposable income of people below the <i>at-risk-of-poverty threshold</i> , expressed as a percentage of the threshold	$\frac{\widehat{arpt} - \hat{\zeta}_{0,5}^p(x, \omega)}{\widehat{arpt}} \cdot 100 \quad (3)$
Foster, Greer and Thorbecke, s1()	Average of the absolute poverty gaps (difference between the <i>at-risk-of-poverty threshold</i> and the equivalised disposable income, with the non-poor being given a difference of zero)	$\frac{\sum_{i=1}^n \omega_i (\widehat{arpt} - x_i)^+}{\sum_{i=1}^n \omega_i} \quad (4)$
Gini, gini()	Relationship of cumulative proportions of the population arranged according to the level of equivalised disposable income, to the cumulative proportions of the equivalised total disposable income they receive	$\left[\frac{\sum_{i=1}^n \left(\omega_i x_i \sum_{j=1}^i \omega_j \right) - \sum_{i=1}^n \omega_i^2 x_i}{\sum_{i=1}^n \omega_i \sum_{i=1}^n \omega_i x_i} - 1 \right] \cdot 100$
Quintile share ratio, qsr()	The ratio of total income received by the 20 percent of the population with the highest income to that received by the 20 percent of the population with the lowest income	$\frac{\sum_{i=1}^n \omega_i x_i \mathbf{I}(x_i > \hat{\zeta}_{0.8}(x, \omega))}{\sum_{i=1}^n \omega_i x_i \mathbf{I}(x_i \leq \hat{\zeta}_{0.2}(x, \omega))} \cdot 100 \quad (5)$

Table 2: Common indicators of poverty and inequality coded in **rtip** package. In formula 4 we define $(x)^+ = x$ if $x \geq 0$ and $(x)^+ = 0$ if $x < 0$. In formulae 2 and 5, $\mathbf{I}(\cdot)$ is the usual indicator function that equals 1 if the bracketed expression is true, and 0 otherwise. In formula 3, $\hat{\zeta}_{0,5}^p(x, \omega)$ is the median income of poor persons.

$$L(p) = \frac{1}{\mu} \int_0^p \zeta_q dq,$$

and the generalized Lorenz (GL) curve is given by $GL(p) = \mu L(p)$ (Shorrocks, 1983) where μ denotes the mean income. The values of the GL curve are estimated on a regular grid of K points selected such that $p_i = i/K$, and their population quantiles denoted by $\zeta_{p_i} = F^{-1}(p_i)$ with $i = 1, 2, \dots, K$. The conditional mean of income less than or equal to ζ_{p_i} is denoted as $\gamma_i = \mathbb{E}[X|X \leq \zeta_{p_i}]$, for $i = 1, 2, \dots, K$ ($\gamma_K = \mu$). The $K \times 1$ vector of GL ordinates at p_1, p_2, \dots, p_K is given by $\theta = [p_1 \gamma_1, p_2 \gamma_2, \dots, p_K \gamma_K]'$ and can be estimated consistently by

$$\hat{\theta} = [p_1 \hat{\gamma}_1, p_2 \hat{\gamma}_2, \dots, p_K \hat{\gamma}_K]', \tag{6}$$

where the sample counterpart of γ_i is

$$\hat{\gamma}_i = \frac{\sum_{j=1}^{r_i} \omega_j x_j}{\sum_{j=1}^{r_i} \omega_j},$$

and $r_i = [p_i \sum_{j=1}^n \omega_j]$ (the closest integer not greater than $p_i \sum_{j=1}^n \omega_j$) is an integer such that $\hat{\zeta}_{p_i}(x, \omega) = x_{(r_i)}$ is the r_i -th sample quantile level such that 100 p_i percent of the weighted sample of observations is less than or equal to $\hat{\zeta}_{p_i}(x, \omega)$, $i = 1, 2, \dots, K$ (Beach and Davidson, 1983; Beach and Kaliski, 1986). In the package **rtip**, the function `lc()` is implemented to estimate both Lorenz and GL curve ordinates. This function calculates the Lorenz curve for the number of abscissae p_i given by the argument `samplesize`. Following the examples in Beach and Davidson (1983), Beach and Kaliski (1986) and Xu (1997), we set `samplesize = 10` by default. If `samplesize = complete`, ordinates are computed in each value along the whole distribution. By setting `generalized = TRUE`, the GL curve is calculated. The left-hand panel of Figure 1 shows the Lorenz curve for income distribution of the Burgenland region. It is produced with:

```
Burgenland <- setupDataset(eusilc2, country = "AT", region = "Burgenland")
lorenz_curve <- lc(Burgenland, samplesize = 10, generalized = FALSE, plot = FALSE)

p1 <- ggplot(lorenz_curve, aes(x.lg, y.lg)) + geom_line() +
  geom_segment(aes(x = 0, y = 0, xend = 1, yend = 1),
    linetype = "dotted", color = "grey") +
  scale_x_continuous(expression(p)) +
  scale_y_continuous(expression(L(p))) + theme_bw()
print(p1)
```

The variable `lorenz_curve` contains the abscissae p_i and Lorenz curve ordinates. A less elaborate plot is obtained by setting `plot=TRUE` in `lc` function.

TIP curve

For an individual (or household-level) measure of deprivation Y , with distribution function F_Y , the deprivation profile for F_Y is (Shorrocks, 1995, 1998)

$$D(F_Y, p) = \int_{F_Y^{-1}(1-p)}^{\infty} y dF_Y(y) = \int_{1-p}^1 F_Y^{-1}(q) dq, \quad p \in [0, 1].$$

Let $z > 0$ be a poverty threshold and let X be an income random variable. If we consider the poverty gap $Y = (z - X)^+$ as the measure of deprivation, where $x^+ = \max\{x, 0\}$, the TIP (Three I's of Poverty) curve is obtained and denoted $TIP(p, z)$. Alternatively, we can write (Jenkins and Lambert, 1998a,b)

$$TIP(p, z) = \int_0^{r_z^X} (z - F^{-1}(t)) dt, \quad p \in [0, 1],$$

where $r_z^X = \sup\{F(x) : x < z\}$ is the proportion of people with income below z . Scaling by z the poverty gap, that is, using $Y = (z - X)^+ / z$ as the measure of deprivation, we obtain the normalized TIP curve, which is simply $TIP(p, z) / z$.

For its estimation, let the observations of poverty gaps $y := \{y_i\}_{i=1}^n$ be ordered in increasing order so that $y_1 \leq y_2 \leq \dots \leq y_n$ with $\hat{z} = \widehat{arpt}$ given by equation (1) and let $\omega := \{\omega_i\}_{i=1}^n$ denote the




Figure 1: Lorenz curve for Burgenland region (left). Dotted diagonal represents the benchmark for a perfectly equal income distribution. Normalized TIP curve for Burgenland region (right).

corresponding sampling weights. Using the relation between deprivation profile and the GL curve, $D(F_Y, p) = \mu(F_Y) - GL(F_Y, 1 - p)$ for $p \in [0, 1]$, the $K \times 1$ vector of TIP curve ordinates corresponding to $[p_1, p_2, \dots, p_K]'$ is estimated consistently by

$$\hat{\phi} = [(\hat{\gamma}_K - p_{K-1}\hat{\gamma}_{K-1}), (\hat{\gamma}_K - p_{K-2}\hat{\gamma}_{K-2}), \dots, (\hat{\gamma}_K - p_1\hat{\gamma}_1), \hat{\gamma}_K]', \tag{7}$$

where the sample counterpart of γ_i is

$$\hat{\gamma}_i = \frac{\sum_{j=1}^{r_i} \omega_j y_j}{\sum_{j=1}^{r_i} \omega_j},$$

and $r_i = [p_i \sum_{j=1}^n \omega_j]$ (the closest integer not greater than $p_i \sum_{j=1}^n \omega_j$) is an integer for which $\hat{\zeta}_{p_i}(y, \omega) = y_{(r_i)}$ is the r_i -th sample quantile level such that 100 p_i percent of the weighted sample of poverty gaps is less than or equal to $\hat{\zeta}_{p_i}(y, \omega)$, $i = 1, 2, \dots, K$ (Beach and Kaliski, 1986; Xu and Osberg, 1998). The function `tip()` estimates both the unnormalized and normalized TIP curve ordinates. Normalization is established by setting `norm = TRUE` and the estimated poverty threshold, \hat{z} , is computed by the function `arpt`. The number of ordinates computed is given by the the argument `samplesize`, and following the example in Xu and Osberg (1998), we set `samplesize = 50` by default. If `samplesize = complete`, `tip` ordinates are computed in each value along the whole distribution. The right-hand panel of Figure 1 shows the normalized TIP curve for income distribution of Burgenland region. It is produced with:

```
tip_curve <- tip(Burgenland, arpt(Burgenland), samplesize = 50, norm = TRUE)
p2 <- ggplot(tip_curve, aes(x.tip, y.tip)) + geom_line() +
  scale_x_continuous(expression(p)) +
  scale_y_continuous(expression(TIP(p, z))) +
  theme_bw()
print(p2)
```

As in the previous example, a less elaborate plot is obtained by setting `plot=TRUE` in the call to the `tip` function.

Dominance tests

Given two income distributions X_1 and X_2 , X_1 is said to Lorenz dominate X_2 if the Lorenz curve of X_1 lies everywhere above that of X_2 , which is interpreted as less inequality in X_1 than in X_2 . The normative aspects of Lorenz dominance have been studied by Atkinson (1970) and its relationship to other dominance criteria may be found in Arnold (1987). Similarly, X_1 is said to dominate X_2 in the generalized Lorenz sense if the generalized Lorenz curve of X_1 lies everywhere above that of X_2 , which is interpreted in terms of welfare (see Shorrocks, 1983). For the relationship of generalized Lorenz dominance to other dominance criteria, see Ramos et al. (2000). Given a poverty threshold $z > 0$, X_1

is said to TIP dominate X_2 if $TIP_{X_1}(p, z) \geq TIP_{X_2}(p, z)$ for all $p \in (0, 1)$, which means that there is less poverty in X_2 than in X_1 according to various wide classes of poverty indices, see [Jenkins and Lambert \(1998a,b\)](#) and [Sordo et al. \(2007\)](#). If we use different poverty thresholds for different income distributions, that is, z_1 in X_1 and z_2 in X_2 , non-intersection of normalized TIP curves is equivalent to unanimous poverty orderings by different classes of poverty indices based on normalized poverty gaps.

Since the initial papers by [Beach and Davidson \(1983\)](#) and [Bishop et al. \(1989\)](#) focussing on statistical tests for Lorenz dominance and generalized Lorenz dominance, respectively, many studies have been conducted to implement these ranking criteria empirically (see Chapter 17 in [Duclos and Araar \(2006\)](#) for a review). Some tests in the literature are based on a two-stage testing strategy including multiple pairwise sub-tests. [Xu \(1997\)](#) offers an alternative to this approach by providing a joint and simpler procedure to test for generalized Lorenz dominance directly, which is adapted to testing for deprivation dominance in [Xu and Osberg \(1998\)](#). We have implemented [Xu \(1997\)](#) and [Xu and Osberg \(1998\)](#) procedures in `rtip`.

Generalized Lorenz dominance

To make statistical inference about GL dominance from sample GL curve estimates we have implemented the asymptotically distribution-free statistical inference procedure in [Xu \(1997\)](#). This article provides one test based on Theorem 1 in [Beach and Davidson \(1983\)](#) which derives the asymptotic joint variance-covariance matrix of GL curve ordinates. EU-SILC surveys involve sampling weights. We implemented an extension of the methodology in [Beach and Davidson \(1983\)](#) to samples which involve weighted observations (see [Beach and Kaliski, 1986](#)).

Given two income distributions, X_1 and X_2 , let θ_1 and θ_2 be the $K \times 1$ vectors of GL curve ordinates for X_1 and X_2 , respectively. The dominance relation tested by the null hypothesis is $H_0 : \theta_1 - \theta_2 \geq 0$ against the alternative hypothesis $H_1 : \theta_1 - \theta_2 \not\geq 0$. The test statistic, T , for the GL dominance is (see [Xu, 1997](#))

$$T = \tilde{\Delta}' \left[\frac{\hat{\Sigma}_1}{n_1} + \frac{\hat{\Sigma}_2}{n_2} \right]^{-1} \tilde{\Delta}, \tag{8}$$

where $\tilde{\Delta} = [(\hat{\theta}_1 - \hat{\theta}_2) - (\tilde{\theta}_1 - \tilde{\theta}_2)]$; n_i is the size of a random sample from X_i ; $\hat{\Sigma}_i$ is the estimated $K \times K$ covariance matrix for the unrestricted vector of GL ordinates $\hat{\theta}_i$ given by (6) while $\tilde{\theta}_i$ is the restricted estimate minimizing ($i = 1, 2$)

$$\Delta' \left[\frac{\hat{\Sigma}_1}{n_1} + \frac{\hat{\Sigma}_2}{n_2} \right]^{-1} \Delta \tag{9}$$

$$\text{s.t. } (\theta_1 - \theta_2) \geq 0$$

with $\Delta = [(\hat{\theta}_1 - \hat{\theta}_2) - (\theta_1 - \theta_2)]$. The auxiliary function `OmegaGL()` computes the empirical unrestricted vector of GL curve ordinates, $\hat{\theta}_i$, and its corresponding covariance matrix, $\hat{\Sigma}_i$, $i = 1, 2$. The function for testing generalized Lorenz dominance between two income distributions is `testGL()`. For both functions the number of ordinates, K , estimated by `rtip` and employed for testing dominance is controlled by the argument `samplesize`. Following the example in [Xu \(1997\)](#) the default value is `samplesize=10`. The upper-and lower-bounds of critical values (at the α significance level) for testing inequality restrictions are provided by [Kodde and Palm \(1986\)](#). If the value of the T statistic (called `Tvalue`) falls into an inconclusive region (between the lower- and upper-bounds) the simulated p -value is estimated following [Wolak \(1989\)](#), otherwise the p -value is set to NA. For instance, to test the null hypothesis that the income distribution of Burgenland dominates the income distribution of Carinthia in the generalized Lorenz sense we use the following procedure:

```
Burgenland <- setupDataset(eusilc2, country = "AT", region = "Burgenland")
Carinthia <- setupDataset(eusilc2, country = "AT", region = "Carinthia")
testGL(Burgenland, Carinthia, generalized = TRUE, samplesize = 10, alpha = 0.05)
```

The output produced is:

```
$Tvalue
      [,1]
[1,] 7723.701

$p.value
[1] NA
```

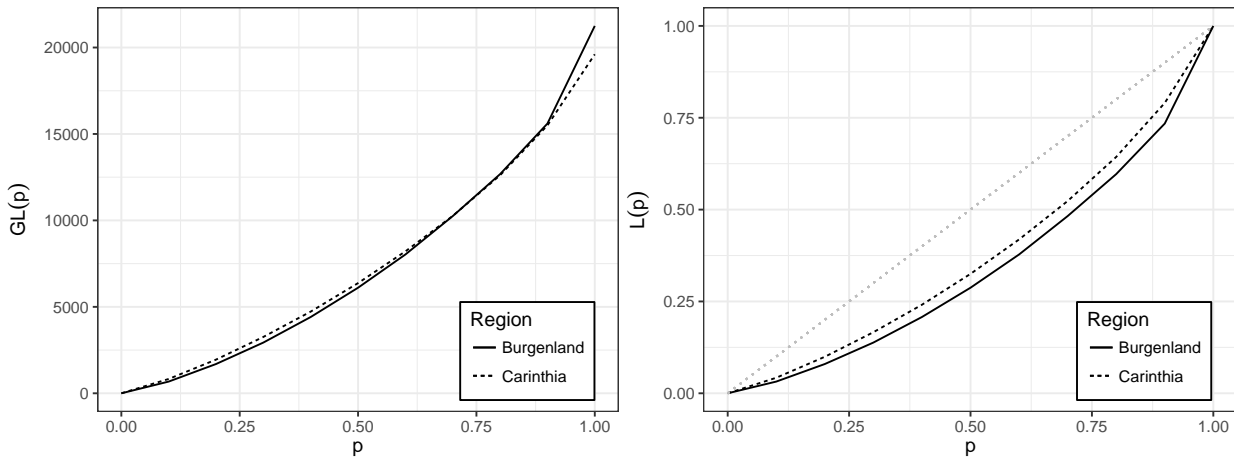



Figure 2: Generalized Lorenz curves for Burgenland and Carinthia (left). Lorenz curves for Burgenland and Carinthia Dotted diagonal represents the benchmark for a perfect equality (right).

```
$decision
[1] "Reject null hypothesis"
```

In this example, the null hypothesis of dominance is rejected with a significance level of 5%. The left-hand panel of Figure 2 displays the two estimated GL curves.

To test the null hypothesis that the income distribution of Carinthia dominates the income distribution of Burgenland in the Lorenz sense we use the function `testGL()` as follows:

```
testGL(Carinthia, Burgenland, generalized = FALSE)
```

This results in:

```
$Tvalue
      [,1]
[1,] 4.049037e-12
```

```
$p.value
[1] NA
```

```
$decision
[1] "Do not reject null hypothesis"
```

In this case we do not have evidence to reject the null hypothesis with a significance level of 5%. The right-hand panel of Figure 2, which displays the two estimated Lorenz curves, suggests (without statistical significance) that Carinthia may dominate Burgenland in the Lorenz sense.

TIP dominance

To make statistical inference about TIP dominance we have implemented the asymptotically distribution-free statistical procedure in [Xu and Osberg \(1998\)](#). As in the case of GL dominance, we have followed the methodology suggested by [Beach and Davidson \(1983\)](#) and [Beach and Kaliski \(1986\)](#).

Let's consider two poverty thresholds $z_1, z_2 > 0$; two income distributions, X_1 and X_2 , and the corresponding poverty gaps $Y_i = (z_i - X_i)^+, i = 1, 2$. Let ϕ_1 and ϕ_2 be the $K \times 1$ vectors of TIP curve ordinates for X_1 and X_2 , respectively. The dominance relation tested on the null hypothesis is $H_0 : \phi_1 - \phi_2 \geq 0$ against the alternative hypothesis $H_1 : \phi_1 - \phi_2 \not\geq 0$. Following the methodology in [Xu and Osberg \(1998\)](#) to test for TIP dominance between two TIP curves, the test-statistic implemented is such that

$$T = \tilde{\Delta}' \left[\frac{\hat{\Omega}_1}{n_1} + \frac{\hat{\Omega}_2}{n_2} \right]^{-1} \tilde{\Delta}, \tag{10}$$

where $\tilde{\Delta} = [(\hat{\phi}_1 - \hat{\phi}_2) - (\tilde{\phi}_1 - \tilde{\phi}_2)]$; n_i is, for $i=1, 2$, the size of a random sample from X_i ; $\hat{\Omega}_i = R\hat{\Sigma}_iR'$, $\hat{\Sigma}_i$ is the estimated $K \times K$ covariance matrix for the unrestricted vector of GL ordinates and R is the

$K \times K$ differencing matrix¹⁰ according to

$$\begin{bmatrix} 0 & \dots & 0 & -1 & 1 \\ 0 & \dots & -1 & 0 & 1 \\ 0 & \dots & 0 & 0 & 1 \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ -1 & \dots & 0 & 0 & 1 \\ 0 & \dots & 0 & 0 & 1 \end{bmatrix}$$

The matrix $\hat{\Omega}_i$ is, for $i = 1, 2$, the estimated $K \times K$ covariance matrix for the unrestricted vector of TIP ordinates $\hat{\phi}_i$ given by (7) while $\tilde{\phi}_i$ is the restricted estimate minimizing

$$\Delta' \left[\frac{\hat{\Omega}_1}{n_1} + \frac{\hat{\Omega}_2}{n_2} \right]^{-1} \Delta \tag{11}$$

$$\text{s.t. } (\phi_1 - \phi_2) \geq 0$$

with $\Delta = [(\hat{\phi}_1 - \hat{\phi}_2) - (\phi_1 - \phi_2)]$. Since the TIP curve becomes horizontal at p equals to the at-risk-of-poverty rate (*arpr*), the test has only been implemented over the interval $[0, \max\{arpr_1, arpr_2\}]$ where $arpr_i$ is the at-risk-of-poverty rate for $X_i, i = 1, 2$. Therefore, $\hat{\phi}_i$ are truncated with the same dimension $k = \max\{arpr_1, arpr_2\} < K$ and the dimension of $\hat{\Omega}_i$ is $k \times k$, for $i = 1, 2$. The auxiliary function `OmegaTIP()` computes the empirical unrestricted vector of TIP curve ordinates, $\hat{\phi}_i$, and its corresponding covariance matrix, $\hat{\Omega}_i, i = 1, 2$. The function for testing TIP dominance between two income distributions is `testTIP()`. Following the practical example in [Xu and Osberg \(1998\)](#) the default value is `samplesize=50`. The rules of rejection and non-rejection based on the value of the statistic T (called `Tvalue`) are in [Xu and Osberg \(1998\)](#). By setting the argument `norm` equal to `TRUE`, the function `testTIP()` uses normalized TIP curves. For example, to test the null hypothesis that the normalized TIP curve of Carinthia dominates the normalized TIP curve of Burgenland we use the following procedure:

```
testTIP(Carinthia, Burgenland, norm = TRUE, samplesize = 50, alpha = 0.05)
```

This yields:

```
$Tvalue
      [,1]
[1,] 11939.99

$P.value
[1] NA

$decision
[1] "Reject null hypothesis"
```

The null hypothesis of dominance is rejected with a significance level of 5%. Figure 3 displays the two estimated normalized TIP curves.

Conclusions

The package `rtip` presented in this paper compares income distributions in terms of welfare, poverty and inequality. Besides providing point estimates and confidence intervals for some commonly used indicators, `rtip` implements the methodology and techniques based on the stochastic dominance to compare income distributions. In particular, we can estimate with `rtip` the usual Lorenz curve, the generalized Lorenz curve, the TIP curve of income distributions and test statistically whether one curve is dominated by another. Although potential users of `rtip` may have their own data, the package allows to load microdata from EU-SILC survey and Spanish Living Conditions survey and offers different equivalence scales to adjust measures of income for differences in household composition. A development version of the `rtip` package can be found at [GitHub](#)¹¹.

¹⁰Recall that the deprivation profile $D(F_Y, \cdot)$ for F_Y is related to the GL curve $GL(F_Y, \cdot)$ of the deprivation measure for F_Y as follows: $D(F_Y, p) = \mu(F_Y) - GL(F_Y, 1 - p)$, for $p \in [0, 1]$.

¹¹<https://github.com/AngelBerihuete/rtip>.




Figure 3: Normalized TIP curves for Burgenland and Carinthia regions (the normalization factors are the respective poverty thresholds).

Acknowledgements

The authors thank two anonymous referees and the Editor of the journal for their detailed comments which have led to significant improvements in both the content and presentation of the paper. Miguel A. Sordo and Carmen D. Ramos acknowledge the support received from Ministerio de Economía y Competitividad (Spain) under grant MTM2014-57559-P.

Bibliography

- A. Alfons and M. Templ. Estimation of Social Exclusion Indicators from Complex Surveys: The R Package **laeken**. *Journal of Statistical Software*, 54(15):1–25, 2013. URL <http://doi.org/10.2139/ssrn.2244876>. [p328, 329, 330]
- A. Araar. *Clorenz: Stata Module to Estimate Lorenz and Concentration Curves*, 2005. URL <https://ideas.repec.org/c/boc/bocode/s456515.html>. [p328]
- B. Arnold. *Majorization and the Lorenz Order: A Brief Introduction*. Springer-Verlag, 1987. [p334]
- A. Atkinson. On the Measurement of Inequality. *Journal of Economic Theory*, 2(3):244–263, 1970. URL [https://doi.org/10.1016/0022-0531\(70\)90039-6](https://doi.org/10.1016/0022-0531(70)90039-6). [p328, 334]
- J. P. Azevedo and S. Franco. **alorenz**: Stata Module to Produce Pen's Parade, Lorenz and Generalised Lorenz Curve, 2006. URL <https://ideas.repec.org/c/boc/bocode/s456749.html>. [p328]
- G. F. Barrett, S. G. Donald, and Y.-C. Hsu. Consistent Tests for Poverty Dominance Relations. *Journal of Econometrics*, 191(2):360–373, 2016. [p328]
- C. M. Beach and R. Davidson. Distribution-Free Statistical Inference with Lorenz Curves and Income Shares. *The Review of Economic Studies*, 50(4):723–735, 1983. URL <https://doi.org/10.2307/2297772>. [p333, 335, 336]
- C. M. Beach and S. F. Kaliski. Lorenz Curve Inference with Sample Weights: An Application to the Distribution of Unemployment Experience. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 35(1):38–45, 1986. [p331, 333, 334, 335, 336]
- J. A. Bishop, S. Chakraborty, and P. D. Thistle. Asymptotically Distribution-Free Statistical Inference for Generalized Lorenz Curves. *The Review of Economics and Statistics*, 71(4):725–727, 1989. URL <https://www.jstor.org/stable/1928121>. [p335]
- B. Buhmann, L. Rainwater, G. Schmaus, and T. M. Smeeding. Equivalence Scales, Well-Being, Inequality and Poverty: Sensitivity Estimates across Ten Countries Using the Luxembourg Income Study (LIS) Database. *Review of Income and Wealth*, 34(2):115–142, 1988. URL <https://doi.org/10.1111/j.1475-4991.1988.tb00564.x>. [p329, 330]
- A. Canty and B. D. Ripley. **boot**: *Bootstrap R (S-Plus) Functions*, 2016. R package version 1.3-18. [p331]

- K. De Vos and M. A. Zaidi. Equivalence Scale Sensitivity of Poverty Statistics for the Member States of the European Community. *Review of Income and Wealth*, 43(3):319–33, 1997. URL <https://doi.org/10.1111/j.1475-4991.1997.tb00222.x>. [p329]
- J. Y. Duclos and A. Araar. *Poverty and Equity: Measurement, Policy and Estimation with DAD*. Economic Inequality and Poverty Series. Springer-Verlag, 2006. ISBN 9780387258935. URL <https://books.google.es/books?id=KOwnYw4qvW4C>. [p335]
- Eurostat. *Description of SILC User Database Variables: Cross-Sectional and Longitudinal*. Unit F-3: Living conditions and social protection statistics, Directorate F: Social Statistics and Information Society, Eurostat, Luxembourg, 2007. [p329]
- J. Foster, J. Greer, and E. Thorbecke. A Class of Decomposable Poverty Measures. *Econometrica*, 52(3):761–766, 1984. [p328, 331]
- J. L. Gastwirth. A General Definition of the Lorenz Curve. *Econometrica*, 39(6):1037–1039, 1971. URL <https://EconPapers.repec.org/RePEc:ecm:emetrp:v:39:y:1971:i:6:p:1037-39>. [p331]
- B. Jann. **lorenz**: Stata Module to Estimate and Display Lorenz Curves and Concentration Curves, 2016. URL <https://ideas.repec.org/c/boc/bocode/s458133.html>. [p328]
- S. P. Jenkins. **povdeco**: Stata Module to Calculate Poverty Indices with Decomposition by Subgroup, 1999. URL <https://ideas.repec.org/c/boc/bocode/s366004.html>. [p328]
- S. P. Jenkins. **svylorenz**: Stata Module to Derive Distribution-Free Variance Estimates from Complex Survey Data, 2006. URL <https://ideas.repec.org/c/boc/bocode/s456602.html>. [p328]
- S. P. Jenkins and F. A. Cowell. Parametric Equivalence Scales and Scale Relativities. *The Economic Journal*, pages 891–900, 1994. URL <https://doi.org/10.2307/2234983>. [p329]
- S. P. Jenkins and P. J. Lambert. Ranking Poverty Gap Distributions: Further TIPs for Poverty Analysis. *Research on Economic Inequality: A Research Annual*, 8:31–38, 1998a. [p328, 333, 335]
- S. P. Jenkins and P. J. Lambert. Three 'Y's of Poverty Curves and Poverty Dominance: Tips for Poverty Analysis. *Research on Economic Inequality: A Research Annual*, 8:39–56, 1998b. URL <https://doi.org/10.1093/oxfordjournals.oep.a028611>. [p328, 333, 335]
- S. P. Jenkins and P. Van Kerm. **glcurve**: Stata Module to Derive Generalised Lorenz Curve Ordinates, 2004. URL <https://ideas.repec.org/c/boc/bocode/s366302.html>. [p328]
- D. A. Kodde and F. Palm. Wald Criteria for Jointly Testing Equality and Inequality Restrictions. *Econometrica*, 54(5):1243–48, 1986. [p335]
- N. T. Longford. *Statistical Studies of Income, Poverty and Inequality in Europe Computing and Graphics in R Using EU-SILC*. Statistics in the Social and Behavioral Sciences. Chapman & Hall/CRC Press, 2015. ISBN 978-1-4665-6832-7. [p329]
- D. Plat. **IC2**: *Inequality and Concentration Indices and Curves*, 2012. URL <https://CRAN.R-project.org/package=IC2>. R package version 1.0-1. [p328]
- H. M. Ramos, J. Ollero, and M. A. Sordo. A Sufficient Condition for Generalized Lorenz Order. *Journal of Economic Theory*, 90(2):286–292, 2000. URL <https://doi.org/10.1006/jeth.1999.2606>. [p334]
- A. F. Shorrocks. Ranking Income Distributions. *Economica*, 50:3–17, 1983. URL <https://doi.org/10.2307/2554117>. [p328, 333, 334]
- A. F. Shorrocks. Revisiting the Sen Poverty Index. *Econometrica*, 63(5):1225–1230, 1995. URL <https://doi.org/10.2307/2171728>. [p331, 333]
- A. F. Shorrocks. Deprivation profiles and deprivation indices. In S. P. Jenkins, A. Kapteyn, and B. M. S. van Praag, editors, *The Distribution of Welfare and Household Production: International Perspectives*, chapter 11, pages 250–267. Cambridge University Press, Cambridge, 1998. [p333]
- M. A. Sordo and C. D. Ramos. Poverty Comparisons When TIP Curves Intersect. *SORT-Statistics and Operations Research Transactions*, 35:65–80, 2011. [p328]
- M. A. Sordo, H. M. Ramos, and C. D. Ramos. Poverty Measures and Poverty Orderings. *SORT-Statistics and Operations Research Transactions*, 31(2):169–180, 2007. [p335]

- M. A. Sordo, C. D. Ramos, and A. Berihuete. Bienestar, Desigualdad y Pobreza En Andalucía: Un Estudio Comparativo Con El Resto De España a Partir De Las Encuestas De Condiciones De Vida 2006 y 2012. *Statistics: Colección Actualidad*, 71, 2014. ISSN 1699-8294. [p328]
- M. A. Sordo, A. Berihuete, C. D. Ramos, and H. Ramos. On a Property of Lorenz Curves with Monotone Elasticity and Its Application to the Study of Inequality by Using Tax Data. *SORT-Statistics and Operations Research Transactions*, 41, 2017. URL <https://doi.org/10.2436/20.8080.02.50>. [p328]
- P. Van Kerm and S. P. Jenkins. Generalized Lorenz Curves and Related Graphs: An Update for Stata 7. *Stata Journal*, 1(1):107–112, 2001. URL <http://EconPapers.repec.org/RePEc:tsj:stataj:v:1:y:2001:i:1:p:107-112>. [p328]
- F. A. Wolak. Testing Inequality Constraints in Linear Econometric Models. *Journal of Econometrics*, 41(2):205–235, 1989. URL [https://doi.org/10.1016/0304-4076\(89\)90094-8](https://doi.org/10.1016/0304-4076(89)90094-8). [p335]
- K. Xu. Asymptotically Distribution-Free Statistical Test for Generalized Lorenz Curves: An Alternative Approach. *Journal of Income Distribution*, 7(1):45–62, 1997. URL [https://doi.org/10.1016/S0926-6437\(97\)80004-2](https://doi.org/10.1016/S0926-6437(97)80004-2). [p329, 333, 335]
- K. Xu and L. Osberg. A Distribution-Free Test for Deprivation Dominance. *Econometric Reviews*, 17(4):415–429, 1998. URL <http://doi.org/10.1080/07474939808800425>. [p329, 334, 335, 336, 337]
- A. Zeileis. **ineq**: *Measuring Inequality Concentration and Poverty*, 2014. URL <https://CRAN.R-project.org/package=ineq>. R package version 0.2-13. [p328]
- B. Zheng. Aggregate Poverty Measures. *Journal of Economic Surveys*, 11(2):123–162, 1997. URL <https://doi.org/10.1111/1467-6419.00028>. [p331]

Angel Berihuete
 Departamento de Estadística e Investigación Operativa
 Universidad de Cádiz
 Spain
 ORCID: 0000-0002-8589-4423
angel.berihuete@uca.es

Carmen Dolores Ramos
 Departamento de Estadística e Investigación Operativa
 Universidad de Cádiz
 Spain
 ORCID: 0000-0002-0134-605X
carmen.ramos@uca.es

Miguel Angel Sordo
 Departamento de Estadística e Investigación Operativa
 Universidad de Cádiz
 Spain
 ORCID: 0000-0003-1383-9051
mangel.sordo@uca.es

Appendix

Users can create their own load function using the `loadLCS` function as a template (see the example below). The load function must produce a data frame with the following variable names: DB010, DB020, DB040, DB090, HX040, HX050, HX090 (see Table 1 to set properly the variables). In the case of the Spanish Living Conditions Survey, for instance, the Spanish Statistical Office delivers four files with many variables. We only require some of them to work with **rtip**:

- ‘lcs_d_file.csv’:
 - DB010: a numeric vector containing the year of the survey.
 - DB020: a factor with one level which is the country considered.
 - DB030: a numeric vector containing the household ID.
 - DB040: a factor with as many levels as there are regions in the country.

- DB090: a numeric vector containing information about household cross-sectional weight.
- 'lcs_hfile.csv':
 - HB010: a numeric vector containing the year of the survey.
 - HB030: a numeric vector containing the household ID.
 - HX040: an integer vector containing information about the household size.
 - HX240: a numeric vector containing information about the equivalised household size. The scale employed is the modified OECD scale.
 - vhRentaa: a numeric vector containing the total disposable household income.

A function to load these files and variables is the following:

```
loadLCS <- function(lcs_d_file, lcs_h_file){

  dataset1 <- read.table(lcs_d_file, header=TRUE, sep= ',')
  dataset2 <- read.table(lcs_h_file, header=TRUE, sep= ',')

  check1 <- identical(dataset1$DB010, dataset2$HB010)
  check2 <- identical(dataset1$DB030, dataset2$HB030)

  if (!check1) {
    stop('Different years!')
  } else if (!check2) {
    stop('You do not have the same identification for homes')
  } else {
    subdataset1 <- subset(dataset1, select = c("DB010", "DB020", "DB030",
                                              "DB040", "DB090"))
    subdataset2 <- subset(dataset2, select = c("HB010", "HB030", "HX040",
                                              "HX240", "vhRentaa"))

    subdataset2$HX050 <- subdataset2$HX240
    subdataset2$HX090 <- subdataset2$vhRentaa/subdataset2$HX240

    dataset <- cbind(subdataset1, subdataset2)
    dataset <- subset(dataset, select = c("DB010", "DB020", "DB040",
                                          "DB090", "HX040", "HX050",
                                          "HX090"))

    return(dataset)
  }
}
```

The code required to reproduce the examples in this paper can be downloaded from: <https://gist.github.com/AngelBerihuete/7e88d55845044ce04a9e61edcd5954f2>.