

BONN ECON DISCUSSION PAPERS

Discussion Paper 16/2004

Structural Stability of the Joint Distribution of Income and Wealth

by

Michał Paluch

December 2008



Bonn Graduate School of Economics
Department of Economics
University of Bonn
Adenauerallee 24 - 42
D-53113 Bonn

The Bonn Graduate School of Economics is
sponsored by the

Deutsche Post  World Net
MAIL EXPRESS LOGISTICS FINANCE

Structural Stability of the Joint Distribution of Income and Wealth

Michał Paluch*

First draft: 30. September 2004

This version: December 8, 2007

Abstract

This paper analyzes the change over time in the distribution of households' income and financial wealth in Great Britain. Empirical analysis based on the British Family Resources Survey data from the period 1996-2001 examines whether the sequence of these distribution is structurally stable in the sense related to Malinvaud (1993). In order to do this, we look for the local time-invariance of a distribution derived after applying simple transformations like scaling or standardizing to the original distribution. In our study we make use of kernel density estimation to identify the changes in shapes of the aforementioned distributions and to perform a nonparametric density time-invariance test as proposed by Li (1996). Our main result is that accounting only for the changes in the vector of means of the original distribution is not sufficient to obtain the desired local time-invariance. In fact, this can be achieved by accounting for changes in the vector of means *and* dispersion parameters of the original distribution.

Keywords: evolution of a distribution, structural stability, kernel density estimation, test for equality of distributions, distribution of income and wealth

*Bonn Graduate School of Economics and Forschungsgruppe Hildenbrand, University of Bonn, Adenauerallee 24-42, D-53113 Bonn, Germany. URL: <http://www.uni-bonn.de/~mpaluch1>, E-mail: paluch@uni-bonn.de.

Financial support by Deutsche Forschungsgemeinschaft, grant Hi-232/2, is gratefully acknowledged.

1 Introduction

The notion of *structural stability* can be found in many fields of economic research. However, its definition turns out to be different for different fields of research. From the econometric point of view, for example, one could regard a postulated model to be structurally stable, if no structural breaks occur in the sense that parameter values are assumed to be constant over time, see e.g. Chow (1983) or Hansen (1992). A slightly different definition is used in game theory, where a game is considered to satisfy the property of structural stability, if small perturbations of the payoff matrix do not alter the qualitative nature of the outcome, see e.g. Palis and Smale (1970). In this paper we will confine ourselves to the notion of structural stability in the context of aggregation theory.

The concept of structural stability has been present in aggregation theory since the papers of Malinvaud (1993).² Unlike typical macroeconomic models that link aggregate response to aggregate explanatory variables, Malinvaud's idea was to model aggregates in terms of the entire joint distribution of all individual variables. This distribution was assumed to have a certain parametric form (structure), e.g., the log-normal distribution in case of the distribution of income or the firm size. In modeling changes over time in this distribution, he made use of the empirical fact that its structure does not change over time, i.e., the log-normality prevails, and its entire evolution can be well captured by changes in only few of its parameters like the mean or the variance. It is this phenomenon which Malinvaud refers to as structural stability.³

In fact, the concept of structural stability as stated by Malinvaud (1993) cannot be applied to distributions which are poorly approximated by a parametric form.⁴ If one does not want to impose any assumptions on the parametric form of the analyzed distributions, one is forced to find a more flexible (nonparametric) counterpart of Malinvaud's idea. Instead of keeping the parametric structure constant and allowing for changes in few parameters, one can fix the values of some parameters and allow the shape of the dis-

²Malinvaud (1993) was in the main the English translation of his paper in French from 1956.

³This empirical regularity has been mentioned not later than in the 19th century for the case of income distributions by Pareto (1896-1897).

⁴The assumption of the log-normality of the income distribution is violated for variety of countries because of its multimodality.

tribution to vary. This can be achieved by simple transformations of the distribution like centering, scaling or standardizing. This concept has been formulated by Hildenbrand and Kneip (1999). Their definition of structural stability of a sequence of distributions states that, by applying a simple transformation to the original distribution, the local time-invariance of the sequence of transformed distributions can be achieved. Hence, the local time-invariance holds if the period-to-period changes in the sequence of transformed distributions can be regarded as statistically insignificant. Therefore, if a transformed distribution turns out to be locally time-invariant, the complicated evolution of the original distribution can be captured completely by the changes in the parameters used for the transformation.⁵

The most important implication of structural stability is the possibility to predict the shape of the future distributions. Indeed, if structural stability holds, the original distribution in period $t + 1$ is completely determined by the original distribution in period t and the parameters, like the mean or the variance, which have been used for transformation, in period $t + 1$. As a consequence, the very complex modeling of the short-run evolution of this distribution can be reduced to the modeling of changes in the parameters. Interestingly, despite the arising new possibilities of modeling aggregate behavior on the basis of structural stability, one can hardly find applications of this concept in the literature.⁶ Indeed, to the author's knowledge, there is only one theory that models aggregation under structural stability. In order to model a relative change in an aggregate in an economy, Hildenbrand and Kneip (1999 and 2005) propose an approach based on the evolution over time of distributions of observed and unobserved explanatory variables.

Surprisingly, even in the empirical literature the explicit verification of structural stability is very seldom. For example, the evolution of individual or cross-country relative income distribution has been studied extensively in the economic literature. Empirical

⁵Consequently, one can distinguish several versions of structural stability depending on the strictness of this assumption, e.g. the local time-invariance of a standardized distribution is a weaker assumption than the corresponding assumption for the centered or relative distribution.

⁶Schumpeter (1951), as cited by Malinvaud (1993), regrets that researchers do not exploit the potentialities of structural stability:

"Few if any economists seem to have realized the possibilities that such invariants hold out for the future of our science... nobody seems to have realized that the hunt for, and the interpretation of, invariants of this type might lay the foundations of an entirely novel type of theory"

work on this topic, e.g. Cowell, Jenkins and Litchfield (1996), Quah (1997) or Sala-i-Martin (2002), however, was targeted mainly at the aspect of changing inequality, poverty, and convergence of these distributions.⁷ Indeed, we are aware of only two papers that studied empirical validity of structural stability of the distribution of households' income. In Hildenbrand, Kneip and Utikal (1999), graphical analysis of the evolution of relative and standardized income distribution for Great Britain is presented. It turns out that simple transformations of this distribution like scaling or standardizing can remove a huge part of its variation over the years. Pittau and Zelli (2001) analyse trends in income distribution in Italy both graphically and by means of a statistical test and show that the distribution of relative incomes is locally time-invariant for many periods.

The aforementioned empirical studies concerned only univariate distributions. However, in the formulation of structural stability, Malinvaud mentions the *joint distribution* of all individual exogeneous variables. This motivates our paper, which extends the empirical study of Hildenbrand et al (1999) on income distribution in two aspects. First, we incorporate an additional variable, namely wealth of a household. Consequently, in this paper we will study the short-run dynamics of the joint distribution of households' income and wealth. In particular, we try to find local time-invariance in this distribution after exposing it to scaling or standardizing transformations. Second, to endorse graphical arguments and to check whether the observed changes over time in this distribution are statistically significant, a nonparametric time-invariance test as suggested by Li (1996) is performed.

The remainder of this paper is organized as follows. We give a motivation for the study of the joint distribution of income and wealth and its evolution in Section 2.2. A brief description of one particular application of the aggregation model formulated by Hildenbrand and Kneip (2005) with emphasis on the hypothesis of structural stability is given. In Section 2.3 we present the data from the Family Resources Survey used in our empirical analysis and report some descriptive statistics of the underlying population

⁷The mentioned papers apply kernel density estimation and are therefore not the typical ones in the empirical literature on convergence and changing inequality of the income distribution. Usually, the analysis of these issues is based solely on the study of the changes in the characteristic parameters of this distribution, like the Gini-coefficient, variance of log-income, Atkinson (1970) indices or the mean-median ratio. One example of papers following this approach is Gottschalk and Smeeding (2001) that contains an international comparison of the income inequality and its changes over time.

of British households. Furthermore, we describe the econometric methods which are employed in this paper to analyze the short-run dynamics of distributions. Finally, we look for a transformation of the original distribution that is sufficient to yield the local time-invariance of the resulting distribution in Section 2.4. A short summary and conclusions are provided in Section 2.5.

2 A motivating example: aggregation of households' consumption expenditure

The aim of the aggregation model in Hildenbrand and Kneip (2005) is to explain the relative change in an aggregate over time. The starting point of this model is the behavioral relation of the microunit, which links explanatory variables to the individual response variable. The modeling occurs, amongst others, in terms of changes in the distribution of *observable* and *unobservable* individual exogeneous variables across the whole population. In particular, the joint distribution of all *observable* micro-specific variables across the whole population is assumed to be structurally stable.

As already mentioned in the Introduction, one application of the model stated in Hildenbrand and Kneip (2005) is the aggregation of households' consumption expenditures. For this particular case, the whole population in period t - denoted by H_t - consists of households h , who have to decide about the level of their consumption expenditure. Therefore, their behavioral relation links following explanatory variables: income, wealth, prices, interest rates, preference parameters of the utility function, expectational variables like expected future income, life expectancy etc. to the response variable, i.e., the consumption expenditure of a household. The consumption theoretical application presented in Hildenbrand and Kneip (2005) treats only two of the variables mentioned above as observable⁸ and micro-specific. These two variables are the household's income and wealth denoted by y_1^h and y_2^h , respectively, and are captured in the vector of observable micro-specific variables of household h , which is denoted by y^h . Consequently, for this particular application of the model, the joint distribution of income and wealth across the whole population, denoted by $\text{distr}(y | H_t)$, is assumed to evolve in the struc-

⁸The main criterion to consider a variable to be observable is the availability of the data on this variable. It is often the case that even if the variable is observable in reality, e.g. some aspects of wealth, households are either not asked for or they just do not know its exact value.

turally stable way. Hildenbrand and Kneip (2005) state this assumption in terms of the the standardized distribution, i.e.,

Hypothesis: Structural stability of $\text{distr}(y | H_t)$

The standardized joint distribution of log-income and log-wealth across the whole population⁹ changes sufficiently slowly over time in the sense that this distribution can be considered as approximately equal for two periods that are close to each other.

In the empirical part of this paper we will study the evolution of the relative and standardized joint distribution of log-income and log-wealth. Therefore, the empirical results can be used to verify the hypothesis of structural stability of the joint distribution of log-income and log-wealth as formulated above by Hildenbrand and Kneip (2005).

3 Data treatment and methodology

Our empirical analysis is based on cross-sectional data from the British Family Resources Survey (henceforth referred to as FRS). This survey was started in 1992 by order of the Department of Social Security. For each individual in the household it collects information on income, savings and financial assets and on a variety of socio-economic and demographic variables like age or employment status of each household's member. Each year about 25,000 households are interviewed. The information gained by this survey is mainly used by non-governmental organizations to simulate and analyze the response of the population to new policy measures. Furthermore, basically due to the large sample sizes, the FRS data is gaining popularity in empirical research being a reliable basis for studies on dynamics of income and wealth, see e.g. Piachaud and Sutherland (2002) or Ginn and Arber (2000).

The variables used for the search of structural stability are income and financial wealth. Unfortunately, due to inconsistency problems in the definitions of these two variables, the time horizon for the analysis had to be reduced to six years, i.e. 1996-2001. As we look for local and not global time-invariance of the distribution, the span of only six years data is adequate for analysis.

⁹For the precise definition of the standardized joint distribution of log-income and log-wealth, see Section 2.4.2.

The income variable used in this paper is household's weekly disposable non-property income, which is defined as the intrahousehold sum of total net earnings from all sources (excluding property income), net pensions and various state transfers like benefit income, income in kind, etc. As far as financial wealth is concerned, balances from following accounts are included: current accounts, savings accounts, gilts, trusts, stocks, shares, national saving certificates, save-as-you-earn contributions, yearly plans, premium bonds, pensioner guaranteed income bonds, etc., whereas life insurance is not included. The value of household's financial wealth is obtained in the following way. At the beginning of the interview about household's wealth, the head of family is asked whether its total amount of capital is between £1500 and £20000. Should it lie within this interval, further questions regarding the composition and amount of financial wealth are asked. Otherwise, the amount of capital is approximated by dividing the yearly investment income from aforementioned accounts by the corresponding account specific interest rates.

It is a well known empirical fact that the distributions of income and wealth are right-skewed. The analysis of the time-invariance of a distribution is much simpler if it is symmetric, because such a distribution can be easier characterized by its moments like mean, variance, etc. Furthermore, at the outset of our empirical study, the large changes in the distributions of income and wealth can be noticeably reduced by using logarithmic transformation. Therefore, for the analysis in this paper we use the log-values of income and financial wealth. The desired effect achieved by the logarithmic transformation can be seen in Figure 1, where the kernel density estimates of the distributions of income and log-income for years 1996-2001 are plotted.

However, the verification of the hypothesis of structural stability of the joint distribution of log-income and log-wealth creates the following problem. Typically, not all households hold financial assets. Because of the use of log-values of income and wealth, the joint distribution $\text{distr}(y | H_t)$ is defined only for strict positive values of y . This forces us to conduct a separate analysis for subpopulation H_t^1 containing all households in the population H_t with positive wealth¹⁰ and subpopulation H_t^0 , which contains the remain-

¹⁰We treat all household with the capital amount of less than £100 (in prices of 1988) as if they had no wealth. This is motivated by the fact that for each household that claims its financial wealth to be less than £1500, the value of financial wealth is approximated by the division of household's




Figure 1: Kernel density estimates of income and log-income distributions across H_t for 1996-2001.

ing households in the population. Interestingly, the relative size of H_t^1 , i.e. H_t^1/H_t , does not change substantially over time. The descriptive statistics for the whole population H_t and the coefficient of correlation between log-income and log-wealth across H_t^1 are given in Table 1.

year	group size			mean log-income		mean	corr.
	H_t^0	H_t^1	H_t^1/H_t	H_t^0	H_t^1	log-wealth	
1996	9401	16019	63.01%	4.832 (0.587)	5.230 (0.716)	7.979 (1.671)	0.105
1997	8911	14387	61.75%	4.870 (0.596)	5.255 (0.725)	7.848 (1.658)	0.075
1998	8816	13951	60.65%	4.884 (0.591)	5.270 (0.733)	7.848 (1.649)	0.097
1999	9895	14929	60.13%	4.929 (0.589)	5.288 (0.737)	7.899 (1.689)	0.079
2000	9763	13813	58.58%	5.061 (0.674)	5.243 (0.720)	7.914 (1.677)	0.065
2001	10196	14931	59.42%	5.014 (0.630)	5.367 (0.716)	7.805 (1.606)	0.067

Terms in parentheses are standard deviations of log-values.

Table 1: Descriptive statistics and the coefficient of correlation between log-income and log-wealth across H_t^1 .

As far as econometric methods applied in this paper are concerned, all distributions have been estimated nonparametrically using the adaptive bandwidth kernel density

yearly investment income by the interest rate. The breaking point of £100 corresponds to the negligible household's weekly investment income of £0.10 if one assumes that the interest rate is at 5%.

estimator with the second order Gaussian kernel function. The pilot bandwidth was chosen according to Sheather and Jones (1991) plug-in method.

Once densities are estimated, an important question arises, whether the observed changes over time in the estimates are statistically significant. In order to answer this question, we apply a nonparametric test of closeness between two distribution functions as proposed by Li (1996). Given the observations¹¹ $X = (X_1, \dots, X_n)$ and $Y = (Y_1, \dots, Y_n)$ drawn from the corresponding unknown density functions f_X and f_Y the test is based on the integrated squared difference between f_X and f_Y denoted by I and defined by

$$\begin{aligned} I &= \int [f_X(t) - f_Y(t)]^2 dt = \int [f_X^2(t) + f_Y^2(t) - 2f_X(t)f_Y(t)] dt \\ &= \int f_X(t) dF_X(t) + \int f_Y(t) dF_Y(t) - 2 \int f_Y(t) dF_X(t). \end{aligned}$$

In our paper the densities f_X and f_Y correspond to the distributions from different time periods, e.g. f_X and f_Y are the relative log-income distributions in period t and $t + 1$ respectively. The feasible estimator of I , denoted by I_n , can be obtained, if one substitutes the density functions f_X and f_Y by their kernel estimates \hat{f}_X and \hat{f}_Y , i.e.,

$$\hat{f}_X(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x - X_i}{h}\right) \quad \text{and} \quad \hat{f}_Y(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x - Y_i}{h}\right).$$

Using these estimates and replacing F_X and F_Y by their empirical distribution functions, one can write $I_n = I_{1n} + I_{2n}$, where

$$I_{1n} = \frac{2K(0)}{nh} - \frac{2}{n^2h} \sum_{i=1}^n K\left(\frac{X_i - Y_i}{h}\right) = c(n) + \mathcal{O}(n^{-1})$$

and

$$I_{2n} = \frac{1}{n^2h} \sum_{i=1}^n \sum_{\substack{j=1 \\ i \neq j}}^n \left[K\left(\frac{X_i - X_j}{h}\right) + K\left(\frac{Y_i - Y_j}{h}\right) - K\left(\frac{Y_i - X_j}{h}\right) - K\left(\frac{X_i - Y_j}{h}\right) \right].$$

¹¹For the sake of simplicity of the presentation, we assume the samples of observations on X and Y to be of equal sizes and to be drawn from univariate densities f_X and f_Y . However, the extension of the test for the case of different sample sizes and multivariate distributions is easy. Furthermore, the random variables X and Y need not to be independent in the sense that the possible dependence does not change the asymptotic distribution of the test statistic.

The test structure is as follows:

$$H_0: f_X(x) = f_Y(x) \text{ almost everywhere}$$

$$H_1: f_X(x) \neq f_Y(x) \text{ for some } x.$$

Under the null hypothesis of time-invariance and assuming that for $h \rightarrow 0$ and $nh \rightarrow \infty$ Li (1996) has shown that $T_n := nh^{1/2} \frac{I_n - c(n)}{\hat{\sigma}_0} \rightarrow^d N(0, 1)$, where

$$\hat{\sigma}_0 = \frac{2}{n^2 h} \sum_{i=1}^n \sum_{j=1}^n \left[K \left(\frac{X_i - X_j}{h} \right) + K \left(\frac{Y_i - Y_j}{h} \right) + 2K \left(\frac{X_i - Y_j}{h} \right) \right] \left[\int K^2(u) du \right]$$

and $c(n) = 2K(0)/nh$.

The asymptotic distribution of the test statistic T under the null hypothesis has a slow rate of convergence to the the standard normal distribution. In order to account for this finite sample bias, we perform the bootstrap procedure to approximate the distribution of T . We repeat a following procedure 500 times: Out of the pooled sample $\{X_1, \dots, X_{n_1}; Y_1, \dots, Y_{n_2}\} =: \{Z_1, \dots, Z_{n_1+n_2}\}$ two samples, $\{X_1^*, \dots, X_{n_1}^*\}$ and $\{Y_1^*, \dots, Y_{n_2}^*\}$, are randomly drawn with replacement. Then, based on the new samples the test statistic $T_{n,i}^*$ is computed. The empirical distribution of T under the null hypothesis is then estimated from the sample $\{T_{n,1}^*, \dots, T_{n,500}^*\}$. The bandwidth for testing purposes was obtained as an optimal bandwidth for density estimation for the pooled sample $\{Z_1, \dots, Z_{n_1+n_2}\}$ according to the Sheather and Jones (1991) plug-in method. A proof of consistency of this bootstrap in the context of testing our hypotheses can be found in Li, Maasoumi and Racine (2007).

4 Empirical results

4.1 The evolution of the relative joint distribution of log-income and log-wealth

The relative joint distribution of log-income and log-wealth across the population H^1 in period t is defined as the distribution of $\hat{y}_t^h = (\hat{y}_{t,1}^h, \hat{y}_{t,2}^h) := (y_{t,1}^h/m_{t,1}, y_{t,2}^h/m_{t,2})$, where $m_{t,1}$ and $m_{t,2}$ denote the mean log-income and mean log-wealth across H_t^1 , respectively. For the population H^0 the relative joint distribution of log-income and log-wealth is just the univariate distribution of relative log-income. Mean-scaling of the distribution implies the first moment of the resulting relative distribution to be constant over time

and equal to 1. Therefore, one can regard the relative distribution as a detrended one in which only higher moments like variance, skewness or kurtosis may change over time.¹² Consequently, if the shape of the relative distribution does not change significantly over time, the evolution of the original distribution is captured entirely by the changes over time in its mean.

4.1.1 Population H^1

Figures 2 and 3 show the kernel density estimates of $\text{distr}(\hat{y} | H_{1996}^1)$ and the associated density contours for years 1996 and 1997, respectively. As one can see in Figure 3, the density contours for these two years do not differ noticeably from each other. We have observed this feature also for other years of the sample. This fact can be seen more clearly on two dimensional graphs of marginal distributions of $\text{distr}(\hat{y} | H_t^1)$, i.e., the relative log-income distribution and relative log-wealth distribution across H_t^1 , which are presented in Figure 4.




Figure 2: Kernel density estimate of $\text{distr}(\hat{y} | H_{1996}^1)$.

¹²Pittau and Zelli (2001) use a different definition of the relative distribution, which is derived by dividing all observations by the sample median and not the mean. Note that in the case of median-scaling, the mean of this kind of relative distribution will not usually be not time-invariant.

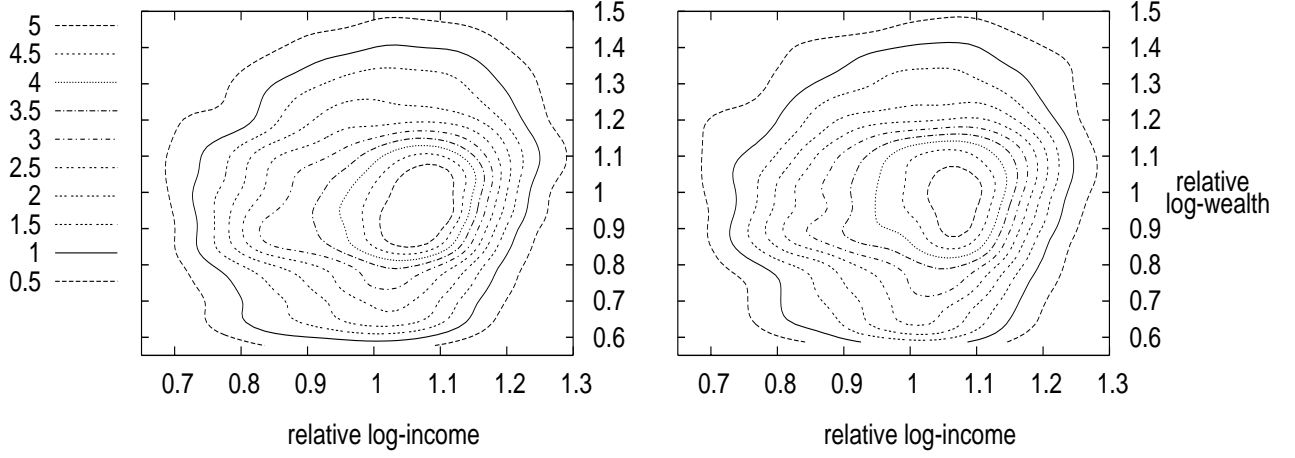


Figure 3: Density contours of $\text{distr}(\hat{y} | H^1_{1996})$ (left) and $\text{distr}(\hat{y} | H^1_{1997})$ (right).

4.1.2 Population H^0

The relative log-income distribution across H^0 , which is plotted in Figure 5, can be also regarded as stable over time. However, a huge increase in the dispersion of the original distribution in the year 2000 that can be seen in Table 1 is reflected in the estimate, which is quite different from that for other years. As the mean-scaling transformation does not account for changes in the dispersion, we can expect the changes during the transitions 1999-2000 and 2000-2001 to be highly significant.

4.1.3 Li (1996) test results for the relative distributions

The question, whether the observed year-to-year changes are significant or not, cannot be answered without applying proper statistical test. Therefore, in order to study the significance of changes in the relative joint distribution of log-income and log-wealth over time, we apply the Li (1996) test. The test results are given in Table .

As one can see in Table 2, the null hypothesis of equality of $\text{distr}(\hat{y} | H^1_t)$ and $\text{distr}(\hat{y} | H^1_{t+1})$ cannot be rejected for only one transition period, 1997-1998, which implies that the evolution of $\text{distr}(y | H^1)$ is too complex to be captured by only its first moment. As far as the distribution $\text{distr}(\hat{y}_1 | H^0)$ is concerned, one cannot reject the equality hypothesis for only two transition periods, 1997-1998 and 1998-1999. This motivates the attempt to incorporate further parameters that would account for changes in the dispersion of the original distribution. The most intuitive candidates for this are the




Figure 4: Kernel density estimator of the relative log-income distribution and the relative log-wealth distribution across H^1 for 1996-2001.

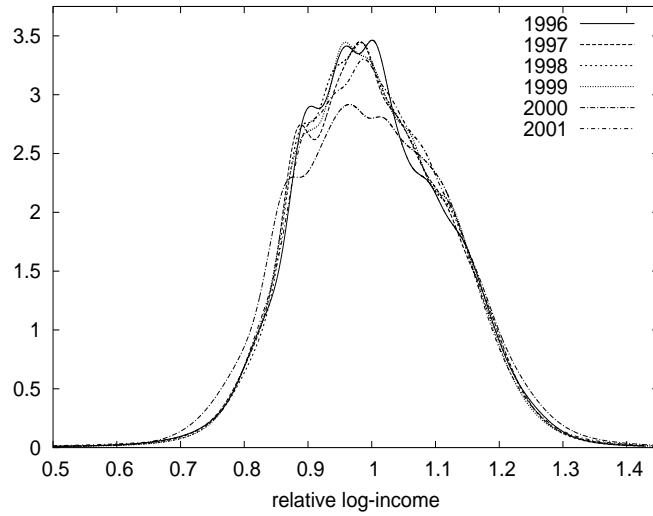


Figure 5: *Kernel density estimate of the relative log-income distribution across H^0 for 1996-2001.*

elements of the covariance matrix of the original distribution. In the next subsection, we will study the case of standardizing transformation as an example of such an extension.

transition period	Subpopulation H^1		Subpopulation H^0	
	T-stat	empirical p-value	T-stat	empirical p-value
1996 vs. 1997	3.934	0*	2.868	0.004*
1997 vs. 1998	1.061	0.107	-1.054	0.807
1998 vs. 1999	3.173	0*	-1.299	0.902
1999 vs. 2000	12.880	0*	17.354	0*
2000 vs. 2001	6.069	0*	14.816	0*

Asterisk indicate that equality is rejected at the 5% level.

Table 2: Li (1996) test results for the distributions $\text{distr}(\hat{y} | H^1)$ and $\text{distr}(\hat{y} | H^0)$ for years 1996-2001.

4.2 The evolution of the standardized joint distribution of log-income and log-wealth

The standardized joint distribution of log-income and log-wealth across H^1 in period t is defined as the distribution of $\tilde{y}_t^h := \Sigma_t^{-1/2}(y_t^h - m_t)$, where m_t denotes the vector of means of log-income and log-wealth and Σ_t is the covariance matrix of log-income and log-wealth across H_t^1 . The correlation between log-income and log-wealth across the population H^1 presented in Table 1 is very small. Therefore, one can approximate this distribution by applying to the original distribution – $\text{distr}(y | H_t^1)$ – the simpler version of the standardization, so called coordinate-wise standardization. The coordinate-wise standardized distribution of y_t^h is then defined as the distribution of $(\tilde{\tilde{y}}_{t,1}, \tilde{\tilde{y}}_{t,2}) := \left(\frac{y_{t,1} - m_{t,1}}{\sigma_{t,1}}, \frac{y_{t,2} - m_{t,2}}{\sigma_{t,2}} \right)$, where $\sigma_{t,1}$ and $\sigma_{t,2}$ denote the standard deviations of log-income and log-wealth, respectively and m_t is the vector of corresponding means across the population H_t^1 .

We expect changes over time in the shape of the standardized distribution to be less significant as the corresponding changes in the relative distribution. This is due to the fact that the standardizing transformation (even the coordinate-wise one) implies not only the time-invariance of the vector of means (equal to 0) of the transformed distribution, but also the time-invariance of the variances (equal to 1) of its marginal distributions.

4.2.1 Population H^1

Kernel density estimates of $\text{distr}(\tilde{y} | H_{1996}^1)$ and the associated density contours for years 1996 and 1997 are presented in Figures 6 and 7, respectively. As in the case of the relative distribution, the density contours for these years do not change much over time, which also holds for other years. Marginal distributions of $\text{distr}(\tilde{y} | H_t^1)$, i.e. the standardized log-income distribution and the standardized log-wealth distribution across H_t^1 are presented in Figure 8 and reveal small variations in these distributions.




Figure 6: Kernel density estimate of $\text{distr}(\tilde{y} | H_{1996}^1)$.

4.2.2 Population H^0

Figure 9 comprises the evidence for the strength of structural stability in showing how even considerably different original distributions can be transformed to very similar ones by controlling for changes in only few parameters. The original distribution of log-income for the year 2000 differs much from that for other years, however, if one applies standardization, the resulting distributions are very similar for all years. Note that this is in contrast to the case of the corresponding relative distributions as shown in Figure 5.




Figure 7: Density contours of $\text{distr}(\tilde{y} | H^1_{1996})$ (left) and $\text{distr}(\tilde{y} | H^1_{1997})$ (right).

4.2.3 Li (1996) test results for the standardized distribution

The null hypothesis of equality of $\text{distr}(\tilde{y} | H^1_t)$ and $\text{distr}(\tilde{y} | H^1_{t+1})$ cannot be rejected for all years within the time period 1996-2001. These results, given in Table 3, indicate the possibility of capturing the evolution of the entire distribution $\text{distr}(\tilde{y} | H^1_t)$ by only few parameters, namely the means and the standard deviations. As for the population H^0 , the hypothesis of equality cannot be rejected at the 5% significance level for the transitions 1997-1998 and 1998-1999. Further, one cannot reject the equality at the 1% level for the transitions 1996-1997 and 1999-2000. The changes in the standardized distribution of log-income between 2000 and 2001 turn out to be statistically significant at the 1% level.

5 Conclusions

The main aim of this paper was to examine the short-run dynamics of the joint distribution of income and wealth of British households on the basis of the Family Resources Survey 1996-2001. The focal point of our analysis is the property of structural stability of this distribution – a notion that was formulated firstly by Malinvaud (1993) for distributions of a certain parametric form and was reformulated for the nonparametric case by Hildenbrand and Kneip (1999). In this paper, we want to avoid any assumptions on the shape of this distribution and we follow the latter approach. According to this concept,




Figure 8: Kernel density estimate of the standardized log-income distribution and the standardized log-wealth distribution across H^1 for 1996-2001.




Figure 9: Kernel density estimate of the standardized log-income distribution across H^0 for 1996-2001.

if a sequence of distributions can be exposed to a simple transformation in that manner that the sequence of the transformed distributions is locally time-invariant, then the sequence of original distributions is said to be structurally stable. In our search for a sim-

transition period	Subpopulation H^1		Subpopulation H^0	
	T-stat	empirical p-value	T-stat	empirical p-value
1996 vs. 1997	0.182	0.392	2.354	0.011*
1997 vs. 1998	0.073	0.468	-1.372	0.912
1998 vs. 1999	0.062	0.457	-0.945	0.715
1999 vs. 2000	0.160	0.391	2.004	0.017*
2000 vs. 2001	0.199	0.344	4.107	0**

Asterisks * (**) indicate the rejection of equality at the 5% (1%) level.

Table 3: Li (1996) test results for the distributions $\text{distr}(\tilde{y} | H_t^1)$ and $\text{distr}(\tilde{y} | H_t^0)$ for years 1996-2001.

ple transformation of a original distribution, i.e. the joint distribution of log-income and log-wealth, that yields local time-invariance of the transformed distribution, we analyze two transformations. The first one, mean-scaling, which could control for the changes over time in mean log-income and mean log-wealth and resulted in the relative joint distribution of log-income and log-wealth, was not sufficient to support the hypothesis of structural stability. However, after applying the standardizing transformation, which accounted for changes in means and dispersion of the original distribution we obtained a sequence of distributions that was local time-invariant, i.e. the period-to-period changes in this sequence were statistically insignificant for almost all years in our sample. This fact empirically supports the hypothesis of structural stability of the joint distribution of income and wealth providing a justification for using this hypothesis in theoretical aggregation models such as the model in Hildenbrand and Kneip (2005).

References

- Atkinson, A. B.** (1970): “On the Measurement of Inequality”, *Journal of Economic Theory*, 2, pp. 244-263.
- Chow, G.** (1983): ”Econometrics”, Chapter 9, McGraw-Hill Book Company, New York.
- Cowell, F.A, S.P.Jenkins and Litchfield J.A.** (1996) “The Changing Shape of the UK Income Distribution: Kernel Density Estimates”, in J.Hills (ed) “New Inequalities: The Changing Distribution of Income and Wealth in the United Kingdom”, Cambridge University Press, Cambridge.
- Ginn, J. and Arber, S.** (2000): “Ethnic inequality in later life: variation in financial circumstances by gender and ethnic group”, *Education and Ageing*, 15, pp.65-83.
- Gottschalk, P and Smeeding T. M.** (2001): ”Empirical Evidence on Income Inequality in Industrial Countries”, in ”Handbook of Income Distribution”’, Atkinson, A. B., Bourguignon, F. (eds.), Amsterdam, Elsevier, pp. 261-304.
- Hansen, E. B.** (1992): ”Testing for Parameter Instability in Linear Regression Models”, *Journal of Policy Modeling*, 14, pp. 517-533.
- Hildenbrand, W., Kneip, A. und Utikal, K** (1999): “Une analyse non paramétrique des distributions du revenu et des caractéristiques des ménages”, *Revue de Statistique Appliquée*, 47, pp. 39-56.
- Hildenbrand, W. and Kneip, A.** (1999): “Demand Aggregation under Structural Stability”, *Journal of Mathematical Economics*, 31, pp. 81-110.
- Hildenbrand, W. and A. Kneip** (2005): “Aggregate Behavior and Microdata”, *Games and Economic Behavior*, 50, pp. 3-27.
- Li, Q.** (1996): “Nonparametric Testing of Closeness between Two Unknown Distribution Functions”, *Econometric Reviews*, 15, pp. 261-274.

Li, Q., Maasoumi, E. and J. Racine (2007): “A Nonparametric Test for Equality of Distributions with Mixed Categorical and Continuous Data”, *mimeo*.

Malinvaud, E. (1993): “A Framework for Aggregation Theories”, *Ricerche Economiche*, 47, pp. 107-135.

Palis J. and Smale S. (1970): “Structural Stability Theorems”, *Proceedings of the Institute on Global Analysis*, American Mathematical Society, 14, pp. 223-232.

Pareto, V. (1896-1897): “Cours d’économie politique”, Vol. 1-2, Rouge, Lausanne.

Piachaud, D. and Sutherland, H. (2002): “Changing poverty post-1997”, Centre for Analysis of Social Exclusion, London School of Economics, CASEPaper 63.

Pittau, M. G. and Zelli R. (2001): “Income distribution in Italy: A nonparametric analysis”, *Statistical Methods and Applications*, 10, pp. 175-189.

Quah, D. (1997): “Empirics for growth and distribution: Stratification, polarization, and convergence clubs”, *Journal of Economic Growth*, 2, pp. 27-59.

Sala-i-Martin, X. (2002) “The World Distribution of Income (estimated from Individual Country Distributions)”, NBER Working Paper No. 8933

Schumpeter, J. (1949). “Vilfredo Pareto, 1848-1923”, Reprinted in J. Schumpeter (1951), “Ten Great Economists”, Oxford University Press, Oxford.

Sheather, S.J. and M.C. Jones (1991). “A Reliable Data-Based Bandwidth Selection method for Kernel Density Estimation”, *Journal of the Royal Statistics Association, Series B* 53, 683-690.