The Economic Journal, **115** (*March*), C118–C132. © Royal Economic Society 2005. Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

HETEROGENEITY IN REPORTED WELL-BEING: EVIDENCE FROM TWELVE EUROPEAN COUNTRIES*

Andrew Clark, Fabrice Etilé, Fabien Postel-Vinay, Claudia Senik and Karine Van der Straeten

This paper models the relationship between income and reported well-being using latent class techniques applied to panel data from twelve European countries. Introducing both intercept and slope heterogeneity into this relationship, we strongly reject the hypothesis that individuals transform income into well-being in the same way. We show that both individual characteristics and country of residence are strong predictors of the four classes we identify. We expect that differences in the marginal effect of income on well-being across classes will be reflected in both behaviour and preferences for redistribution.

The relationship between income and utility or well-being is an important transversal question in social science. While the shape of the utility function is one of the keystones of microeconomics, most empirical estimation to date has been based on relatively simple specifications. One of the main drawbacks in much existing work is the lack of controls for unobserved individual heterogeneity.[1] This paper introduces two sources of heterogeneity into the relationship between income and utility.

Estimation requires that individual utility be measured. A common approach has been to appeal to well-being variables as measures of unobserved continuous utility. Such variables are increasingly found in representative household surveys. Some are global indices, such as happiness, life satisfaction or psychological stress,[2] others are domain specific, such as job or income satisfaction.

These measures have a number of particular characteristics. First, they are ordinal. A life satisfaction score of 6, on a scale of 1 to 7, does not correspond to twice as satisfied as a score of 3. In this ordinal world, 6 only means more than 5 and less than 7. Second, proxy utility measures are bounded. In our example above, someone with a satisfaction score of 7 last year has no way of indicating that she is even happier this year. As such, ordered probit or ordered logit estimation is required in cross-sections.

One worry regarding statistical analysis of subjective variables is that some people look at life pessimistically or optimistically, even though there is 'really' no difference in their level of well-being.[3] This 'anchoring effect' or intercept heterogeneity is a source of potential bias (Winkelmann and Winkelmann, 1998). Advances in econometric theory, and more pragmatically in the statistical packages

* We are grateful to participants at the Royal Economic Society conference (Swansea), Edward Morey, Catherine Sofer, Jon Temple and two anonymous referees for very useful comments.

[1] A recent survey is Senik (2005). See also Van Praag and Ferrer-i-Carbonell (2004), Frey and Stutzer (2002) and Easterlin (2001).

[2] The GHQ-12 score, used by Clark and Oswald (1994), is an example of an index of psychological stress or mental well-being.

[3] A more profound criticism suggests that replies to subjective questions are pure noise. However, one strand of the literature has appealed to panel data to show the predictive power of proxy utility measures, for example linking life satisfaction to future marriage (Lucas *et al.*, 2003; Stutzer and Frey, 2003) or job satisfaction to future quits (Clark, 2001; Freeman, 1978).

that the majority of economists use for their applied work, have largely overcome one part of this worry, at least in a technical sense. It is now simple to control for individual-specific effects in an ordinal regression. Recent examples of such estimation of well-being are Clark and Oswald (2002), Ferrer-i-Carbonell and Frijters (2004) and Senik (2004).

Using conventional fixed or random effects corrects for intercept heterogeneity. We go one step further and allow the parameters of the unobserved (latent) individual utility function to differ across individuals (Tinbergen, 1991; Sen, 1992), i.e. we model slope heterogeneity. In the context of this paper's subject matter, our approach therefore amounts to asking not only whether 'money buys happiness' but also 'for whom it buys the most happiness'.

We use latent class techniques to model intercept and slope heterogeneity simultaneously in the relationship between income and reported well-being across twelve European countries. This therefore represents an attempt at modelling heterogeneity in the marginal utility of income. The statistical model endogenously divides the observations (in a probabilistic sense) into separate classes or groups, which differ by the parameters (slope and intercept) of the relation between income and satisfaction.[4] These probabilities depend on time-invariant individual characteristics, including country dummies. A straightforward interpretation of our model is in terms of mixtures of distinct subgroups or classes of the population, who differ in their ability to transform income into well-being. This approach for modelling heterogeneity in a non-linear set-up is deeply rooted in the latent class analysis literature and its applications in various fields; see, *inter alia*, Uebersax (1999) in psychometry, Jedidi *et al.* (1997) and DeSarbo and Choi (1999) in marketing, or Eckstein and Wolpin (1999), Deb and Trivedi (2002) and Thacher and Morey (2003) in economics.

The data identify four classes of individuals; the hypothesis that the marginal effect of income on well-being is identical across classes is strongly rejected. As such, slope heterogeneity is important: models which control for intercept heterogeneity are not sufficient. The probabilities of class membership are correlated with individual characteristics, such as income, education and age. We also find some evidence of a North–South split and show that France, Germany and the UK are both close to the sample average and close to each other.

The rest of this paper is organised as follows. Section 1 describes the data that are used. Section 2 presents the methods implemented in order to reveal heterogeneity and Section 3 the results. Section 4 concludes.

1. Data

We use data from the European Community Household Panel (ECHP). The ECHP survey was conducted annually in EU Member States over the period

[4] We therefore do not consider exogenously-determined groups but rather let the data select them. Recent findings with exogenous groups include Lelkes (2005), who shows that the marginal utility of income is lower for the religious in Hungary, and Smith *et al.* (2004) who use HRS data to look at interactions between income and health (levels and changes) in a well-being function. An analogous application with willingness-to-pay data is Morey and Rossman (2003).

1994–2001. In the first wave, in 1994, a sample of some 60,500 nationally representative households – approximately 130,000 adults aged 16 years and over – were interviewed in the then 12 Member States. Austria joined the survey in 1995 and Finland in 1996. The ECHP is an extensive, sample-based panel survey in which the same households and individuals are interviewed annually. Data from the surveys are available on the main aspects of welfare, including income and employment, housing, education, social relationships and health. The data come from a standardised questionnaire and are designed to be cross-nationally comparable. Details of the ECHP are available on the Eurostat web site.[5]

We take a 50% random sample of the data to ease the computational burden. We have three waves of the ECHP data, 1994–6. This yields a balanced sample of 109,425 observations (36,475 individuals over three waves and 12 countries).

Our key variables are satisfaction with financial situation, which is our proxy measure of utility, and income. The former is formulated as follows: '*Could you indicate, on a scale going from* 1 – *'not satisfied at all'* – *to* 6 – *'very satisfied'* – *your degree of satisfaction concerning your financial situation?*'.[6] The latter is given by net household income in Euros, converted between countries using PPPs. This income is further transformed into a household equivalent measure, using the modified OECD scale: weights of 1 for the first adult, 0.5 for subsequent adults (aged over 14), and 0.3 for children. The first and last percentile of the distribution of raw household income have been dropped, due to worries about the accuracy of the reported data. The distribution of all variables is presented in the first column of the Appendix Table. In the statistical analysis, we combine Belgium and Luxembourg, due to relatively small sample sizes in the latter.

Figure 1 below shows the results from non-parametric estimation of the probability of being satisfied (reported satisfaction of 5 or 6 on the six-point scale) on the log of per capita income. Results from four countries are shown: the Netherlands, the UK, France and Portugal. The estimated relationship is mostly positive, but it is obvious that neither the intercepts nor the slopes are the same across countries (although the French curve is remarkably close to the UK).

Our suspicion is that this graph shows heterogeneity in the relationship between income and reported well-being. However, bivariate correlations cannot prove anything, due to composition effects. The remainder of the paper presents and applies a latent class multivariate model which provides a robust test of this hypothesis.

2. Modelling Heterogeneity in Well-being

Our statistical model applies the latent class approach, in which sub-groups of the population are identified endogenously, to an ordered dependent variable, here satisfaction with financial situation.

[5] http://forum.europa.eu.int/irc/dsis/echpanel/info/data/information.html

[6] We would have preferred a general life satisfaction measure here; the ECHP does not contain one.

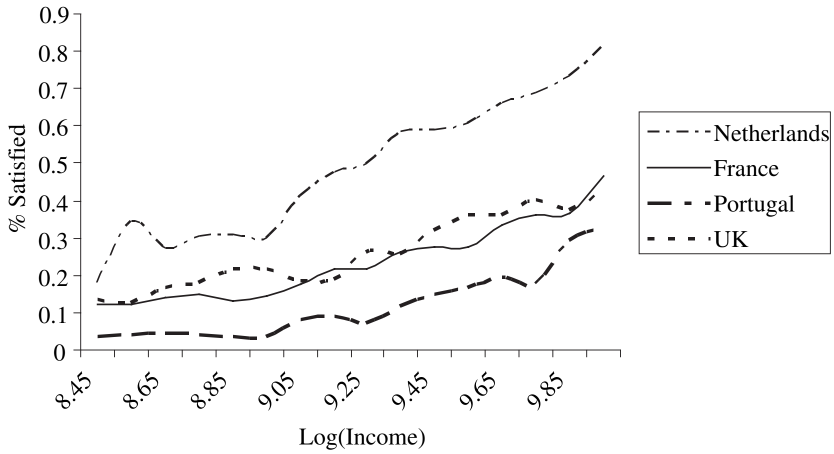

Fig. 1. *Per Capita Income and Well-Being in Europe*
Note: Non-parametric estimation of percentage satisfied (satisfaction = 5 or 6) as a function of the log of per capita income

Given that we do not observe utility directly, but rather infer it from a satisfaction variable, we are potentially faced with two types of heterogeneity. Interpreting subjective responses requires

(*i*) associating utility to observable characteristics, and
(*ii*) relating discrete verbal satisfaction judgements to latent continuous utility.

Figure 2 illustrates the process.[7]

Although we are not able to separate the two types of heterogeneity, i.e. the two sides of Figure 2, the specificity of this paper is to allow for heterogeneity (of both types) in the intercept of the regression line between income and reported satisfaction, but also in the slope of this curve.

An explanatory variable in the left-hand side of the Figure, such as income, is correlated with (unobservable) utility. Individual heterogeneity likely makes an appearance at this point, in the sense that the utility function is not the same across individuals: both intercept and slope heterogeneity can play a role. The right-hand side of the Figure shows the transformation of utility into reported satisfaction levels. Again the relationship between latent utility and verbal satisfaction labels is unlikely to be the same for everybody. The model described below, which is an extension of the standard ordered probit, identifies this second element with intercept heterogeneity. We are not able to distinguish empirically between heterogeneity in the utility function (translating income into utility) and heterogeneity in the expression function (turning utility into reported well-being). We can, however, propose a test of joint heterogeneity in the transformation of income into well-being.

[7] There is an obvious parallel between the two sides of Figure 2 and the phenomena of the hedonic and satisfaction treadmills underscored in Danny Kahneman's work (Kahneman, 2000).

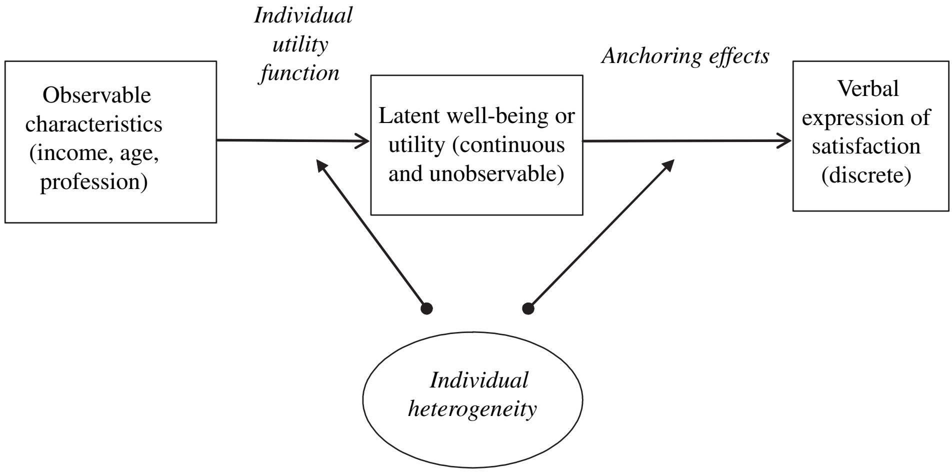

Fig. 2. *Heterogeneity Problems with Subjective Variables*

2.1. *Econometric Modelling*

Consider an agent i who reports her well-being at time t using P different 'naturally' ordered labels such as excellent, very good, good etc. Denote a_{it} her answer, which belongs to the ordered set of labels $L = \{L_1, L_2, \ldots, L_P\}$. The most common way to model this choice assumes that there exists an underlying continuous utility function U^*_{it} and $P+1$ ordered individual threshold parameters $s^0_i = -\infty, s^1_i, \ldots, s^p_i, \ldots, s^P_i = +\infty$ such that:

$$a_{it} = L_p \Leftrightarrow s^{p-1}_i \leq U^*_{it} < s^p_i \tag{1}$$

Here we model utility as:

$$U^*_{it} = \alpha_i Y_{it} + \boldsymbol{\beta}\mathbf{X}_{it} + \varepsilon_{it} \tag{2}$$

where Y_{it} is log income, $\mathbf{X}_{it}$ is a vector of labour market status variables and wave dummies, and ε_{it} is a shock which is independent across individuals and time, and is distributed standard normal.[8] In this specification, the way in which income affects utility is individual-specific. Heterogeneity is thus twofold, first because the marginal utility of income (α_i) is idiosyncratic and, second, because individuals may use different labels to translate utility (U^*_{it}) into reported well-being (the thresholds $s^1_i, \ldots, s^p_i$ are individual-specific).[9]

[8] There is no intercept in (2) as this is included in the threshold parameters. In this model, it is not possible to disentangle the effects of a set of variables on the utility on the one hand and the transformation of utility into reported well-being on the other hand. For such a task, we would need additional arbitrary identifying restrictions (such as a common anchoring point: see Groot, 2000). An alternative route is to identify the conditional distribution of unobserved utility by combining a model for an observable behaviour, presumably driven by utility, and a measurement model of the type considered in this paper.

[9] Clearly this is not the most general model one can think of. In particular, because we want to focus on the effect of income on well-being, we constrain the effects of the labour force status regressors $\mathbf{X}_{it}$, to be equal across individuals. Another possible extension is to account for temporal correlation in the error term.

We use a finite mixture approach to model heterogeneity. That is, we assume that the parameter vector $\mathbf{v}_i = (\alpha_i, s_i^1, \ldots, s_i^p, \ldots, s_i^{P-1})$ is distributed over a finite number of points C: any given individual i in the sample belongs to one of C classes, where each class c is defined by a common value $\mathbf{v}_c$ of the vector $\mathbf{v}_i$.[10]

As we do not observe class membership, we posit that individual i has probability ω_{ic} of belonging to class c (we shall specify ω_{ic} below). The data provide us with empirical probabilities $\mathrm{PR}(a_{i1}, \ldots, a_{iT} | Y_{i1}, \mathbf{X}_{i1}, \ldots, Y_{iT}, \mathbf{X}_{iT})$. Summing over the support of the distribution of the random vector $\mathbf{v}_i$ yields the following decomposition of the individual contribution to the sample likelihood:

$$\Pr(a_{i1}, \ldots, a_{iT} | Y_{i1}, \mathbf{X}_{i1}, \ldots, Y_{iT}, \mathbf{X}_{iT}) = \sum_{c=1}^{C} \omega_{ic} \left[\prod_{t=1}^{T} \Pr(a_{it} | Y_{it}, \mathbf{X}_{it}, \mathbf{v}_i = \mathbf{v}_c) \right]. \tag{3}$$

Note that this multiplicative decomposition of the individual likelihood results from the assumption of no autocorrelation in the error term: replies a_{ij} and a_{ik} at different periods j and k are independent given class membership and contemporaneous values of the covariates.

We now turn to the specification of the individual class membership probabilities ω_{ic}. We assume the following multinomial logit distribution:

$$\omega_{ic} = \Pr(\mathbf{v}_i = \mathbf{v}_c | Z_i) = \frac{\exp(\gamma_c Z_i)}{\sum_{c'=1}^{C} \exp(\gamma_{c'} Z_i)} \tag{4}$$

where $\gamma_1 = 0$ and Z_i is a set of time-invariant individual characteristics including gender, birth cohort, marital status, number of children, education and country of residence.

The joint specification of (2) and (4) requires some comment. Note first that only time-varying regressors are included in the right-hand side of (2), while only time-invariant regressors[11] appear as determinants of class membership in (4). Equation (4) thus imposes a simple parametric structure on the correlation between these latter variables and class membership, which allows all moments of the distribution of unobserved individual heterogeneity v_i to be affected by observed fixed individual characteristics.

Given independence between $\mathbf{v}_i$ and ε_{it}, the distribution of a_{it} conditional on Y_{it}, $\mathbf{X}_{it}$ and $\mathbf{v}_i$ is the standard ordered probit. Denoting the standard normal c.d.f. by Φ, the parameters $\mathbf{v}_c$, $\boldsymbol{\beta}$ and γ_c are obtained via the maximisation of the following log-likelihood:

[10] Alternatively, both intercept and slope heterogeneity can be modelled as continuous joint-normal variables. The method we use has the advantage, in terms of our paper's subject, of providing a clear typology of individuals.

[11] Over our three-year observation period, it is reasonable to consider marital status and number of children as time invariant.

$$\sum_i \log\left[\sum_{c=1}^{C}\left(\frac{\exp(\gamma_c Z_i)}{\sum\limits_{c'=1}^{C}\exp(\gamma_{c'} Z_i)}\prod_{t=1}^{T}\left\{\prod_{p=1}^{P}\left[\Phi(s_c^p-\alpha_c Y_{it}-\boldsymbol{\beta}\mathbf{X}_{it})-\Phi(s_c^{p-1}-\alpha_c Y_{it}-\boldsymbol{\beta}\mathbf{X}_{it})\right]^{1\{a_{it}=L_p\}}\right\}\right)\right] \tag{5}$$

2.2. *Identification and Model Selection*

Individual class membership of individuals being unobserved, we have a standard problem of missing data. This is solved by using a variant of the standard iterative EM (Expected Maximisation) algorithm for missing data (Dempster *et al.*, 1977), the Simulated Annealing EM algorithm (Celeux *et al.*, 1995), which allows a better detection of a global maximum of the sample likelihood and avoidance of saddle points.[12] An interesting by-product of any EM algorithm is a fuzzy classification of observations into classes. For each individual, we compute for each of the C classes the following posterior conditional probability:

$$\begin{aligned} w_{ic} &= \Pr(\mathbf{v}_i=\mathbf{v}_c|a_{i1},\ldots,a_{iT},Y_{i1},\mathbf{X}_{i1}\ldots,Y_{iT},\mathbf{X}_{iT}) \\ &= \frac{\omega_{i,c}\Pr(a_{i1},\ldots,a_{iT}|\mathbf{v}_i=\mathbf{v}_c,Y_{i1},\mathbf{X}_{i1}\ldots,Y_{iT},\mathbf{X}_{iT})}{\Pr(a_{i1},\ldots,a_{iT}|Y_{i1},\mathbf{X}_{i1}\ldots,Y_{iT},\mathbf{X}_{iT})}. \end{aligned} \tag{6}$$

The problem of theoretical identification is important: is it possible to find several sets of parameters that would fit the data equally well (i.e. produce the same likelihood)? Uebersax (1999) proposes an order condition: the number of parameters in the model should be less than the number R of empirical patterns of response. Here we have C slope parameters, $C-1$ probabilities ω_{ic} (the C weights sum up to 1) and $(P-1)C$ threshold parameters, making a total of $C(P+1)-1$. There are P response modalities and T waves. Moreover, conditional on class membership, the probabilities of response are time-independent. Hence, $R=(P+T-1)!/T!(P-1)!$. With $P=6$ and $T=3$, we have $R=56$, so that the order condition inequality becomes $7C-1<56$, and the maximum number of classes we can identify is 8.[13]

To select the empirical optimal number of classes, we compare information criteria such as the BIC and the AIC, and the normalised entropy criterion for 1, 2 or more points. The BIC and AIC statistics are commonly used in order to balance the gain in log-likelihood through an increase in C and the loss of degrees of freedom from the greater number of parameters. The normalised entropy criterion assesses the accuracy of the classification, lower values indicating better class identification (McLachlan and Peel, 2000).

[12] We implement the Simulated Annealing EM algorithm in Stata.

[13] While this condition is obviously necessary, Uebersax (1999) does not prove its sufficiency. Intuitively, identification also requires the presence of a continuous right-hand side regressor (for instance income). Otherwise, it is always possible to classify individuals perfectly according to their response patterns and any set of discrete characteristics.

2.3. *Testing for Homogeneity*

One of our ambitions is to challenge the existing literature somewhat, by showing that there is slope heterogeneity in the income well-being relationship. Homogeneity assumes that

(*i*) the slope parameters α_i (i.e. the marginal effect of income on utility) and
(*ii*) the threshold parameters $(s_i^1, \ldots, s_i^{P-1})$

are the same across individuals. Since the ordered probit specification imposes normalisation of the error-term variance to 1, and of the intercept in the utility function to 0, U_{it}^* is identified up to a linear transformation. As such, slope parameters are only identified up to a multiplicative constant (the unobserved variance). The interpretation of differences in estimated parameters as slope heterogeneity therefore relies on the assumption of homoscedasticity.

We can however test for intercept heterogeneity, even under heteroscedasticity, by showing that the thresholds for class c are not a linear transformation of those for class c'. This holds if and only if the column vectors of thresholds for each class s_c, $\mathrm{s}_{c'}$ and the $(P-1) \times 1$ vector of ones, $\mathbf{1}$, are not collinear. We use the test of rank proposed by Robin and Smith (2000) to test this condition for each couple $(\mathbf{c},\mathbf{c}')$. This is a test of heterogeneity in the well-being effect of income, which is robust to heteroscedasticity.

3. Results

Entropy measures suggest that there are four classes, whilst information criteria opt for five. We retain four classes, for reasons of parsimony.[14]

Table 1 presents the results relating satisfaction to income and labour market status, both for the whole sample and for each of the four groups.

Rank tests reject the hypothesis that the thresholds for one group are a linear transformation of the thresholds for another group.[15] Hence, the way individuals transform utility into well-being varies greatly in the sample and/or utility functions are heterogeneous. Table 1 shows that there are very sharp differences in the effect of income on declared satisfaction.

We are interested in the income well-being relationship within each class, conditional on the other control variables. Figure 3 shows the predicted probabilities of reporting a given satisfaction level, for the 'average' agent (having the average sample demographic characteristics), conditional on her membership of class $c\,(c = 1,\ldots,4$; LCOP means Latent Class Ordered Probit).

[14] The BICs for 4 and 5 classes are respectively −157,558 and −157,114, while the respective normalised entropy measures are 4.44 and 5.58.

[15] The critical values for the statistics proposed by Robin and Smith (2000) were obtained after simulation of their distribution functions. Collinearity was overwhelmingly rejected for all pairs of classes. The precision of the estimates produces critical values for the test statistics which are all under 1. The lowest test statistic we obtain is 850.9 for classes 3 and 4. Other statistics are available from the authors.

Table 1
Satisfaction with Financial Situation

Model	Ordered probit	Latent class ordered probit model – 4 classes: Class 1	Class 2	Class 3	Class 4
Latent index parameters: α_c, β					
ln(income): Y_{it}	0.624**	0.653**	0.558**	0.501**	0.803**
	(0.010)	(0.0003)	(0.0003)	(0.002)	(0.0003)
Labour force status (ref: inactive): $\mathbf{X}_{it}$					
Works over 15 hrs per week.	0.088**	0.151**			
	(0.011)	(0.001)			
Works under 15 hrs per week.	−0.045*	−0.019**			
	(0.024)	(0.007)			
Unemployed	−0.760**	−0.817**			
	(0.020)	(0.002)			
Other controls					
Wave dummies: $\mathbf{X}_{it}$	Yes	Yes			
Time-invariant characteristics: $\mathbf{Z}_i$	Yes	No			
Threshold parameters: $\mathbf{s}_c^1, \ldots, \mathbf{s}_c^{P-1}$					
Cut 1	4.399**	3.743**	2.678**	4.016**	4.943**
	(0.099)	(0.003)	(0.004)	(0.028)	(0.004)
Cut 2	5.066**	4.223**	3.000**	4.825**	6.195**
	(0.099)	(0.006)	(0.010)	(0.032)	(0.007)
Cut 3	5.840**	4.899**	3.661**	5.604**	7.528**
	(0.100)	(0.012)	(0.014)	(0.030)	(0.009)
Cut 4	6.660**	6.055**	4.165**	6.225**	8.990**
	(0.100)	(0.017)	(0.025)	(0.021)	(0.003)
Cut 5	7.622**	7.875**	5.067**	6.821**	10.074**
	(0.101)	(0.005)	(0.040)	(0.017)	(0.001)
% of the sample	100%	26.4%	8.9%	28.6%	36.2%

Note: * = Significant at the 5% level; ** = at the 1% level. Robust standard errors in parentheses clustered on households.

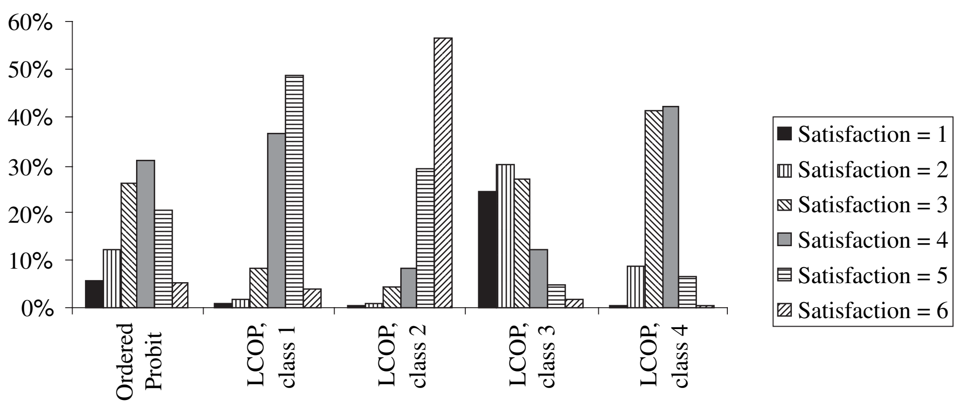

Fig. 3. *Predicted Satisfaction Probabilities.*
Note: Probabilities calculated at sample mean

Compared to predictions from the simple ordered probit model, shown on the left, those in classes 2 and (to an extent) 1 are more likely satisfied, while those in class 3 are more likely dissatisfied. Individuals in class 4 are average.

Table 2
Income Elasticities of Well-Being

Latent class ordered probit	$\bar{\Delta}$ (1)	$\bar{\Delta}$ (2)	$\bar{\Delta}$ (3)	$\bar{\Delta}$ (4)	$\bar{\Delta}$ (5)
LCOP – Class 1	−1.75	−1.47	−1.10	−0.54	−0.06
LCOP – Class 2	−1.59	−1.43	−1.11	−0.88	−0.50
LCOP – Class 3	−0.64	−0.36	−0.16	−0.07	−0.02
LCOP – Class 4	−2.28	−1.42	−0.63	−0.12	−0.01

Note: Calculated at sample mean.

To interpret the effect of income on reported well-being, we calculate probability elasticities. These reflect both the estimated income coefficient and the thresholds, holding all other characteristics constant (at the overall sample mean). Table 2 presents the estimates of the following elasticity:

$$\bar{\Delta}(L_p) = \frac{\left\{\Pr\left[\widehat{WB}_{\bar{i}} \leq L_p | \bar{i} \in c, \bar{\mathbf{X}}, \ln(1.01\bar{I})\right] - \Pr\left[\widehat{WB}_{\bar{i}} \leq L_p | \bar{i} \in c, \bar{\mathbf{X}}, \ln(\bar{I})\right]\right\}}{\Pr\left[\widehat{WB}_{\bar{i}} \leq L_p | \bar{i} \in c, \bar{\mathbf{X}}, \ln(\bar{I})\right]}$$

where I is our income variable (so that $Y = \ln(I)$).

This Table should be interpreted as follows. The figures show the percentage change in the probability of reporting satisfaction lower than or equal to the number in parentheses in the column head; these can be thought of as exit rates from low satisfaction. A 1% rise in income decreases the probability that someone in class 2 (the 'happy' class, from Figure 3) has satisfaction of three or lower (on the one to six scale) by 1.11 percentage points. On the contrary, it has little effect on the same probability for someone in class 3 (the 'unhappy').

The results here are unambiguous. The marginal effect of income on subjective well-being depends on unobserved heterogeneities relating either to the underlying utility function or to the way people label their utility. Further, one group (class 2) is both highly satisfied and has large marginal effects of income on well-being, while another (class 3) is the least satisfied and has the lowest marginal effects of income on well-being. Classes 1 and 4 occupy intermediate positions.

The Appendix shows the distribution of observable characteristics across the four well-being classes in Table 1. Taking the two classes of most interest, we see that those in class 2 (satisfied, high marginal effect of income on well-being) are older, more affluent but actually less likely to be active in the labour market. Those in class 3 (dissatisfied, low marginal effect of income on well-being) are less likely to be married, but more likely to be unemployed. They are also less well-educated. With respect to countries, we broadly find that Northern countries are over-represented in class 2 and under-represented in class 3; the opposite holds for Southern countries.

Table 3 shows the results from the multinomial logit estimation of class membership, as in (4). The estimated coefficients γ_c show the log-odds of class c membership relative to class 1. All of the control variables are time-invariant.

We concentrate on the log-odds of class 2 membership (the satisfied with greater sensitivity to income) relative to class 3 membership (the dissatisfied who are less sensitive to income), i.e. $\gamma_2 - \gamma_3$. The results for the demographic variables show

Table 3

Estimated Class Membership Probabilities

Variables $\mathbf{Z}_i$	Class 2 (γ_2)	Class 3 (γ_3)	Class 4 (γ_4)
Male	−0.248**	0.081	0.125**
	(0.053)	(0.036)	(0.040)
(Age at wave 1)/10	−0.009	0.167*	0.105
	(0.120)	(0.078)	(0.087)
(Age at wave 1-squared)/100	0.037**	−0.042**	−0.027**
	(0.011)	(0.008)	(0.009)
Number of children at home under age 16 (at wave 1).	−0.005	0.037	−0.052
	(0.038)	(0.024)	(0.027)
Marital status at wave 1 (ref: single & has never been in couple)			
Married/Living in couple.	−0.089	−0.507**	0.143
	(0.104)	(0.068)	(0.077)
Widowed, Separated, Divorced.	−0.187	0.704**	0.634**
	(0.132)	(0.091)	(0.105)
Education (ref: no education/primary education)			
Higher Education	0.159	−1.026**	−0.350**
	(0.083)	(0.065)	(0.066)
Secondary Education	0.073	−0.575**	−0.178**
	(0.070)	(0.048)	(0.053)
Country effects (ref: France)			
Germany	1.838**	−0.128	−0.056
	(0.227)	(0.085)	(0.087)
Belgium-Luxembourg	2.794**	−0.324**	−0.379**
	(0.229)	(0.116)	(0.130)
Netherlands	2.323**	−1.679**	−1.495**
	(0.212)	(0.098)	(0.100)
Denmark	3.466**	−1.188**	−0.841**
	(0.217)	(0.129)	(0.122)
United Kingdom	2.564**	0.442**	0.337**
	(0.224)	(0.093)	(0.099)
Ireland	2.948**	0.580**	−0.123
	(0.224)	(0.093)	(0.116)
Italy	0.867**	0.936**	0.861**
	(0.264)	(0.082)	(0.085)
Portugal	2.001**	1.106**	2.387**
	(0.294)	(0.140)	(0.135)
Spain	1.707**	0.921**	0.199*
	(0.233)	(0.081)	(0.095)
Greece	0.956**	1.641**	1.645**
	(0.319)	(0.103)	(0.106)
Constant	−4.218**	0.532**	0.107
	(0.333)	(0.170)	(0.192)

Note: * = significant at the 5% level; ** = at the 1% level. Robust standard errors in parentheses clustered on households.

that couples, the educated, and older people are more likely to belong to class 2 than class 3; men, and those who were in couples but now are no longer, are more likely to belong to class 3 (the dissatisfied) than class 2. These are standard results in the literature (Kahneman *et al.*, 1999).

There is some initial *prima facie* evidence for country groupings, as shown by the estimated log-odds. It is particularly noticeable that Belgium-Luxembourg, Netherlands and Denmark have the greatest probability of being in class 2 relative to class 3. There then follows a medium group, given by Germany, the UK and

Ireland. Finally, the group with the lowest probability of being in the satisfied class relative to the unsatisfied class consists of the Southern European countries. The same country grouping pattern can be observed in the estimated probabilities of being in class 2 (the satisfied) compared to class 4 (the average), and being in class 1 (the fairly satisfied) compared to class 3 (the dissatisfied).

Although one should be cautious about country classifications, it is interesting to look at average ex-post probabilities of class membership (the w_{ic} in (6)) by country. We find, notably, that the three countries the closest to the sample average are France, Germany and the UK, sometimes considered as the motors of the EU economy.[16]

We also compare the average for individuals in the six founding member countries of the European Union to averages in those countries that joined later. Interestingly, predicted class membership in the UK is closest to that of the six founding members (and the UK indeed joined in the first accession wave). Generally, the closer a country's class probabilities to those of the founding members, the earlier the date at which the country joined the EU. The only exception to this pattern is Spain, which is more similar to the 1973 accession countries than to those which joined in the 1980s.

4. Conclusion

This paper models the relationship between income and self-reported well-being using latent class techniques applied to panel data from twelve European countries. We introduce both intercept and slope heterogeneity into this relationship.

The model we propose allows us to identify slope heterogeneity under the assumption of homoscedasticity, and intercept heterogeneity even if errors are heteroscedastic. However, we are not able to model separately the utility function (translating income into utility) and the expression function (turning utility into reported well-being).

Our results strongly reject the hypothesis that individuals carry out the income-well being transformation in the same way. Specifically, we find intercept heterogeneity and, under the assumption of homoscedasticity, slope heterogeneity.

We identify four classes of individuals and show that the marginal effect of income on well-being is very different across these classes. There are distinct demographic and country patterns between these classes. This has at least two important implications.

First, in a political economy sense, as the transformation of income into reported well-being differs sharply across classes (and classes are not independently distributed between countries), we would expect average opinion regarding economic policies to differ across countries. To the extent that we have identified country groups in Table 3, we *a priori* expect countries within these groups to vote similarly with respect to European-level reforms.

Second, there is a great deal of heterogeneity within countries as well, and predicted class membership at the individual level is likely to correlate with various

[16] The similarity of France and the UK is also apparent in the raw data: see Figure 1.

behaviours. In particular one class is satisfied, somewhat richer and has a high marginal income effect, while another is dissatisfied, somewhat poorer and has a low marginal income effect. We imagine that, within a country, class membership may be correlated with both preferences for redistribution and voting behaviour. This is a subject for ongoing research.

Individuals, who seem to fall naturally into a number of different classes, differ in ways that are far more complicated than those picked up by a simple fixed effect. We believe that future applied work in microeconomics will increasingly take slope heterogeneity into account in order to model individual behaviour better.

Appendix: Descriptive Statistics

The Table is to be read as follows. The percentage figures show the probability of having the demographic characteristic in question conditional on belonging to the different classes. For example, 2.2% of respondents in class 2 are unemployed, compared to 9.5% of respondents in class 3. The non-percentage figures are means.

Variables	Whole Sample	Class 1	Class 2	Class 3	Class 4
Log income	9.09	9.31	9.35	8.92	9.00
Over 15 hrs/week	47.8%	49.7%	34.3%	44.7%	52.2%
Under 15 hrs/week	2.9%	3.8%	3.7%	2.6%	2.2%
Unemployed	5.5%	3.6%	2.2%	9.5%	4.4%
Inactive (omitted category)	43.8%	42.9%	59.8%	43.2%	41.2%
Age	47.7	48.0	56.9	44.9	47.6
Number of children	0.64	0.65	0.44	0.70	0.64
Male	44.7%	45.5%	41.5%	43.2%	46.2%
Married, Living in couple	70.5%	75.7%	72.5%	60.7%	74.0%
Widowed, Separated, Divorced	12.0%	9.2%	16.3%	15.0%	10.7%
Single and never married (omitted category)	17.5%	15.1%	11.2%	24.3%	15.3%
Higher Education	13.7%	18.7%	19.8%	9.0%	12.4%
Secondary Education	30.7%	38.8%	34.7%	25.8%	27.7%
Primary Education (omitted category)	55.6%	42.5%	45.5%	65.2%	59.9%
Germany	8.7%	12.0%	7.7%	6.7%	8.1%
Belgium-Luxembourg	5.0%	4.8%	9.3%	2.7%	2.4%
Netherlands	10.0%	22.2%	21.6%	3.1%	3.7%
Denmark	5.4%	7.6%	23.3%	1.7%	2.5%
UK	7.0%	6.6%	10.3%	7.0%	6.6%
Ireland	7.0%	6.5%	12.7%	9.4%	4.1%
France (omitted category)	10.7%	14.3%	1.6%	10.5%	10.4%
Italy	14.0%	9.5%	2.7%	18.5%	16.4%
Portugal	10.3%	2.6%	2.5%	6.7%	20.7%
Spain	10.8%	9.1%	6.5%	17.1%	7.9%
Greece	12.4%	4.8%	1.6%	16.7%	17.2%
Satisfaction = 1	10.6%	1.5%	0.8%	32.8%	2.2%
Satisfaction = 2	14.3%	2.5%	0.9%	28.7%	14.7%
Satisfaction = 3	23.8%	8.5%	4.7%	22.9%	40.2%
Satisfaction = 4	25.3%	33.2%	8.0%	10.1%	35.6%
Satisfaction = 5	18.7%	48.7%	27.0%	4.0%	6.5%
Satisfaction = 6	7.3%	5.6%	58.5%	1.5%	0.7%

Note: There are 109,425 observations, representing 36,475 individuals over three waves.

CNRS, DELTA and IZA
INRA-CORELA
INRA-LEA, CREST, CEPR and IZA
DELTA and University Paris IV
CNRS and CECO, Ecole Polytechnique

References

Celeux, G., Chauveau, D. and Diebolt, J. (1995). 'On stochastic versions of the EM algorithm', Working Paper no. 2514, Paris: Institut National de Recherche en Informatique et en Automatique.

Clark, A. E. (2001). 'What really matters in a job? Hedonic measurement using quit data', *Labour Economics*, vol. 8, pp. 223–42.

Clark, A. E. and Oswald, A. J. (1994). 'Unhappiness and unemployment', Economic Journal, vol. 104, pp. 648–59.

Clark, A. E. and Oswald, A. J. (2002). 'Well-being in panels', DELTA, mimeo.

Deb, P. and Trivedi, P. K. (2002). 'The structure of demand for health care: latent class versus two-part models', *Journal of Health Economics*, vol. 21, pp. 601–25.

Dempster, A. P., Laird, N. M. and Rubin, P. B. (1977). 'Maximum-likelihood from incomplete data via the E-M algorithm', *Journal of the Royal Statistical Society*, Series B, vol. 39, pp. 1–38.

DeSarbo, W. S. and Choi, J. (1999). 'A latent structure double hurdle regression model for exploring heterogeneity in consumer search patterns', *Journal of Econometrics*, vol. 89, pp. 423–55.

Easterlin, R. (2001). 'Income and happiness: towards a unified theory', Economic Journal, vol. 111, pp. 465–84.

Eckstein, Z. and Wolpin, K. I. (1999), 'Why youths drop out of high school: the impact of preferences, opportunities and abilities', *Econometrica*, vol. 67, pp. 1295–339.

Ferrer-i-Carbonell, A. and Frijters, P. (2004). 'How important is methodology for the estimates of the determinants of happiness?', Economic Journal, vol. 114, pp. 641–59.

Freeman, R. B. (1978). 'Job satisfaction as an economic variable', *American Economic Review*, vol. 68, pp. 135–41.

Frey, B. and Stutzer, A. (2002). 'What can economists learn from happiness research?', *Journal of Economic Literature*, vol. 40, pp. 402–35.

Groot, W. (2000). 'Adaptation and scale of reference bias in self-assessment of quality of life', *Journal of Health Economics*, vol. 19, pp. 403–20.

Jedidi, K., Jagpal, H. S. and DeSarbo, W. S. (1997). 'Finite-mixture structural equation models for response-based segmentation and unobserved heterogeneity', *Marketing Science*, vol. 16, pp. 39–59.

Kahneman, D., Diener, E. and Schwarz, N.(eds.) (1999). *Well-Being: The Foundations of Hedonic Psychology*, New York: Russell Sage Foundation.

Kahneman, D. (2000). 'Experienced utility and objective happiness: a moment-based approach', in (D. Kahneman and A. Tversky, eds.), *Choices, Values and Frames*, pp. 673–92, New York: Cambridge University Press and the Russell Sage Foundation.

Lelkes, O. (2005). 'Tasting freedom: happiness, religion and economic transition', *Journal of Economic Behavior and Organization*, forthcoming.

Lucas, R., Clark, A. E., Georgellis, Y. and Diener, E. (2003). 'Re-examining adaptation and the setpoint model of happiness: reaction to changes in marital status', *Journal of Personality and Social Psychology*, vol. 84, pp. 527–39.

McLachlan, G. and Peel, D. (2000). *Finite Mixture Models*, New-York: John Wiley & Sons.

Morey, E. and Rossmann, K. (2003). 'Using stated-preference questions to investigate variations in willingness to pay for preserving marble monuments: classic heterogeneity, random parameters, and mixture models', *Journal of Cultural Economics*, vol. 27, pp. 215–29.

Robin, J-M. and Smith, R. J. (2000). 'Tests of rank', *Econometric Theory*, vol. 16, pp. 151–75.

Sen, A. (1992). *Inequality Reexamined*, Oxford: Oxford University Press.

Senik, C. (2004). 'Relativizing relative income', DELTA Working Paper No. 2004–17.

Senik, C. (2005). 'Income and well-being: what can we learn from subjective data?', *Journal of Economic Surveys*, forthcoming.

Smith, D., Langa, K., Kabeto, M. and Ubel, P. (2004). 'Health, wealth, and happiness: financial resources buffer subjective well-being after the onset of a disability', University of Michigan, mimeo.

Stutzer, A. and Frey, B. S. (2003). 'Does marriage make people happy, or do happy people get married?', University of Zurich, mimeo.

Thacher, J. and Morey, E. (2003). 'Using individual characteristics and attitudinal data to identify and characterize groups that vary significantly in their preferences for monument preservation: a latent-class model', University of Colorado at Boulder, mimeo.

Tinbergen, J. (1991). 'On the measurement of welfare', *Journal of Econometrics*, vol. 50, pp. 7–15.

Uebersax, J. S. (1999). 'Probit latent class analysis with dichotomous or ordered category measures: conditional independence/dependence models', *Applied Psychological Measurement*, vol. 23, pp. 283–97.

Van Praag, B. and Ferrer-i-Carbonell, A. (2004). *Happiness Quantified*, Oxford: Oxford University Press.

Winkelmann, L. and Winkelmann, R. (1998). 'Why are the unemployed so unhappy? Evidence from panel data', *Economica*, vol. 65, pp. 1–15.