

ESRC Research Methods Programme
Working Paper No 8

*Comments on Tackling Social Exclusion:
Taking stock and looking to the future: a
response*

Nick Emmel

PLEASE DO NOT QUOTE THIS PAPER WITHOUT THE PERMISSION OF THE AUTHOR

This paper originally appeared as S3RI Applications and Policy Working Paper No. A05/10 and is included in this series by permission of the authors

Multilevel Modelling of Repeated Ordinal Measures: An application to attitudes to divorce

Ann Berrington¹, Yongjian Hu¹, Karla Ramírez-Ducoing², Peter Smith¹

¹ Southampton Statistical Sciences Research Institute, ² Secretaría de Desarrollo Social (Sedesol) de México (Mexican Social Development Ministry)

Abstract

This paper demonstrates the use of multilevel modelling techniques for an ordinal response. Using repeated measures of divorce attitude from the 1991-1996 waves of the British Household Panel Survey (BHPS) we investigate the factors predicting attitude to divorce and test whether changes in marital status are associated with changes in attitude to divorce. The paper discusses the methodological issues arising from the multilevel modelling of an ordinal outcome and compares the results obtained using marginal, quasi-likelihood and Bayesian methods. The paper demonstrates how the multilevel modelling approach deals with the complex pattern of attrition and intermittent non-response found in the BHPS.

Acknowledgements

Data from the British Household Panel Survey were originally collected by the UK Longitudinal Studies Centre and supplied by the UK Data Archive. Neither the original collectors of the data nor the Archive bear any responsibilities for the analyses or interpretations presented here. This work was supported in part by the ESRC grant “Measuring Attitude Stability and Change” (grant no. H333250026) as part of the ESRC Research Methods Programme. Initial analyses of the data were undertaken by Ramírez-Ducoing for her MSc Social Statistics dissertation at the University of Southampton under the supervision of Smith and Berrington.

1. Introduction

This paper has both substantive and methodological aims and examines the factors associated with divorce attitude among the British population using repeated measures of divorce attitude from a panel study. Little research has been carried out to establish the stability of individuals' divorce attitudes, how they reflect their own personal characteristics, and whether changes in marital status are associated with changes in attitude to divorce. This paper also investigates the relative importance of intra-household correlations in divorce attitude using hierarchical data from the British Household Panel Study. We hypothesize that there will be correlation within households in divorce attitudes as a result of assortative mating and the socialization of younger household members. Methodologically, the paper demonstrates the use of multilevel modelling techniques for an ordinal response, and demonstrates how the multilevel modelling approach to repeated measures copes with complex patterns of attrition and wave non-response as found in surveys such as the British Household Panel Study.

Section 2 introduces the data, discussing the ordinal nature of the response variable, the hierarchical structure and examining patterns of wave and item non-response in the British Household Panel Study. Section 3 introduces the methodology underlying the fitting and interpretation of ordinal logistic and multilevel ordinal logistic regression models. In section 4 we undertake separate cross-sectional analyses of individuals' attitude in 1992, 1994 and 1996. Then we use a multilevel analysis where the repeated attitude observation is nested within an individual respondent. We test for the presence of correlation at higher levels, i.e. we test for a random effect at the level of the household and primary sampling unit. In the longitudinal analysis, we are particularly interested in investigating whether changes in marital status over time are associated with a change in individuals' attitudes toward divorce.

2. Data

2.1 The British Household Panel Survey

The data used for our research come from the British Household Panel Survey (BHPS). The BHPS has been conducted since 1991 and was designed as an annual survey of each adult member (16 years old or above) of a nationally representative sample of more than 5,000 households and approximately 10,000 individual interviews. The main objective of the BHPS is

to gain the understanding of social and economic change at the individual and household level in Britain. A detailed description of the BHPS can be found in the User Guide (Taylor, 2005, <http://iserwww.essex.ac.uk/ulsc/bhps/doc/vola>).

We are primarily interested in respondents' answers to the question "It is better to divorce than to continue an unhappy marriage", which was recorded biennially in waves 2, 4, 6 etc. Although waves from 1991 to 2003 are available for analysis, our study only uses the first six waves. This provides us with three repeated measures of attitude towards divorce whilst reducing loss to the sample through attrition. Therefore, our sample contains those who responded at either waves 2, 4 or 6. Moreover, since we require socio-background information collected at the start of the survey the sample is also restricted to those respondents who had full interviews at wave 1. Since we are concerned with respondents who have left full time education, and those who are likely to have had some personal experience of partnership formation and dissolution we restrict our sample to those aged over 20 years.

2.2 Hierarchical structure of the data

The BHPS has a hierarchical data structure. In each wave individual respondents can be viewed as units clustered within households. In our sample there are 4864 households and 8005 individuals. The BHPS is a stratified sample of 250 Primary Sampling Units (PSU). Therefore, households are nested within PSUs. In addition, the panel design of the BHPS means that measurements are often repeated over time on the same subject and, therefore, observations are nested within individuals (or units) (Figure 1).

As revealed in Figure 1, a four-level structure needs to be considered with the repeated measures of attitude to divorce as the lowest level, and PSU as the top level. Since individuals within households are likely to be more similar than individuals in different households the standard modelling assumption of independence of individuals is unlikely to hold. Multilevel modelling techniques take this lack of independence into account and also allow us to assess whether each of these levels contributes towards the overall variation in divorce attitude.

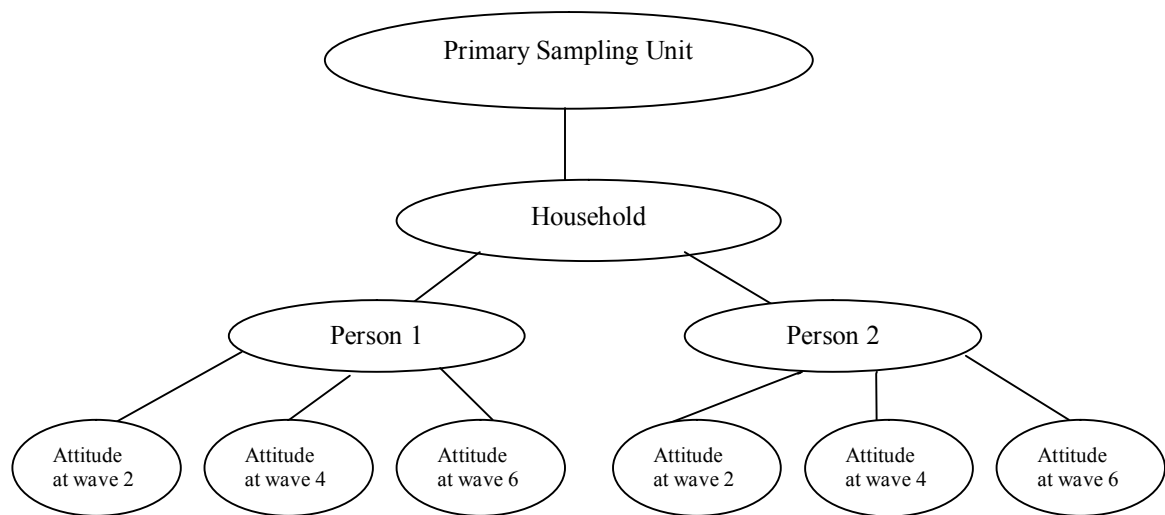


Figure 1: Data structure for modelling repeated attitudes in the BHPS

2.3 Ordinal responses

As mentioned in the first section, our research aims to identify the major predictors of attitude toward divorce. Subjects are asked to respond to the question “It is better to divorce than to continue an unhappy marriage” on a five point ordinal scale from ‘strongly agree’ to ‘strongly disagree’. It is generally recognised that ordinal models are under used in the social sciences. Typically, the response variable is either dichotomised, and modelled using binary logistic regression, or scored, and modelled using ordinary least squares regression (Liu and Agresti, 2005). The former may lose information, whilst the latter imposes a scale on the response and assumes that it is normally distributed.

2.4 Explanatory variables

Table 1 presents the names and coding of the explanatory variables considered in our model selection. These variables were chosen based their observed correlation with divorce attitude in preliminary analyses. Note that region is grouped since only Southern England and Scotland were the only standard regions to have a significantly different attitude to divorce than London.

Except for the marital status variable, the explanatory variables are time-constant variables; they are fixed at wave 1 or wave 2. The marital status variable combines information about *de-facto* and *de-jure* marital status available in the BHPS and distinguishes whether people have ever been married, whether they are currently married, and if currently not married whether

they are currently living with a cohabiting partner. The variable is time-varying in being updated at each wave to reflect current marital status.

Table 1: Names and coding of the variables

Variable name	Coding	Remark
Divorce attitude	Strongly disagree=1, disagree=2, neither disagree nor agree=3, agree=4, strongly agree=5	Response variable
Sex	Males =0, females=1	Fixed at wave 1
Age	21-29 years =1, 30-44=2,45-64=3, and 65+=4	Fixed at wave 2
Race	White=1, Black=2, Asian and others=3	Fixed at wave 1
Education	Degree=1, A level=2,O level=3, CSE=4,No qualifications=5	Fixed at wave 2
Region	London/Midlands / North England / Wales=1, Southern England =2, Scotland=3	Fixed at wave 1
Whether a parent	Yes=1, No=0	Fixed at wave 2
Marital status	Currently married=1, never married but currently cohabiting=2, previously married and currently cohabiting =3 previously married but currently no partner=4, never married and currently no partner=5	Time-varying; this variable reflects the respondents' current marital status
Time	1992=1, 1994=2, and 1996=3	Used only in longitudinal analysis

2.5 Missing values

Missing data is a common problem in survey data. Missing data may mean that no record is made for a whole unit being surveyed (unit non-response) or that only some of the items for a unit are available (item non-response). In our quantitative analysis, both unit non-response and item non-response are referred to as missing data. No imputation was undertaken and these individuals are deleted from the analyses.

Tables 2 to 4 present the pattern of missing values, and the number of observations available for use in each of the cross-sectional analyses at waves 2, 4 and 6. In these tables the missing value indicator is 0 if the variable was observed and 1 if the variable was missing. The actual number of observations available for analysis at each cross-section is shown at the end of the first row in bold. For example, in Table 2, we see that there are 7834 observations with no missing data for the modelling for wave 2. The difference between 8005 and 7834 is due to item non-response among those who did take part in wave 2. Most of the item non-response is

non-response to the attitude question (146 individuals). At wave 4, the amount of missing data increases, primarily because of wave non-response; 978 individuals do not have a valid marital status or divorce attitude. Some of these individuals will have been permanently lost to the BHPS, whilst others do respond at a later wave. A similar pattern can be seen at wave 6, where the effective sample size for analysis is 6345.

Table 2: Pattern of missing values at wave 2

Missing value indicator for:				Frequency
Divorce attitude	Race	Education	Marital status	
0	0	0	0	7834
0	0	1	0	17
0	1	0	0	7
1	0	0	0	146
1	0	1	0	1
Total				8005

Note: the variables already a parent, sex and age have no missing values.

Table 3: Pattern of missing values at wave 4

Missing value indicator for:				Frequency
Divorce attitude	Race	Education	Marital status	
0	0	0	0	6757
0	0	0	1	8
0	0	1	0	15
0	1	0	0	6
1	0	0	0	237
1	0	0	1	978
1	0	1	0	3
1	1	0	1	1
Total				8005

Note: the variables already a parent, sex and age have no missing values.

Table 4: Pattern of missing values at wave 6

Missing value indicator for:				Frequency
Divorce attitude	Race	Education	Marital status	
0	0	0	0	6345
0	0	0	1	10
0	0	1	0	15
0	1	0	0	4
1	0	0	0	163
1	0	0	1	1462
1	0	1	0	1
1	0	1	1	2
1	1	0	1	3
				8005

Note: the variables already, sex and age have no missing values.

Deletion of cases with missing values causes a reduction in sample size. However, as we will see, one advantage of using a multi-level modelling approach to analyse repeated data is that we do not need a balanced data structure. That is to say each individual in the dataset may contribute just one, two or three observations. Hence, we can keep in the analyses respondents who are present in earlier waves who are lost through attrition at some later point. Furthermore, the multi-level model allows individuals who were lost at wave 4 to re-enter the sample and provide information from wave 6. Therefore, the total number of observations (individuals × number of repeated measures) available for our longitudinal analysis is 20,936 (i.e. the sum of the cross-sectional total just described (Table 5)).

Table 5: Missing pattern in the data used for multilevel analysis

Missing value indicator for:				Frequency
Divorce attitude	Race	Education	Marital status	
0	0	0	0	20936
0	0	1	0	47
0	1	0	0	17
0	0	0	1	18
1	0	0	0	546
1	0	1	0	5
1	0	0	1	2440
1	0	1	1	2
1	1	0	1	4
Total				24015

Note: the variables already, sex and age have no missing values

Attrition may lead to differences in the characteristics of those who are followed up from those present at the start of the study. In order to minimize the effect of attrition on our analyses we include covariates such as age, marital status and education which are associated with propensity to be lost to follow-up.

2.6 Frequency of the response variable: Attitude to divorce

Table 6 presents the frequency distribution of the divorce attitude. Over three-quarters of respondents strongly agreed or agreed with the statement “It is better to divorce than to continue an unhappy marriage”. Between 1992 and 1996 the distribution is fairly stable with only a slight reduction in the percentage who ‘strongly agree’. Relatively few people ‘strongly disagree’ with the statement, with less than one percent falling in this category. This skewness suggests that the normality assumption required for ordinary least squares regression of a scored response may not be tenable.

The relative small number of observations in the ‘strongly disagree’ category should be borne in mind when undertaking an ordinal regression analysis. Statistical packages, such as STATA and MLwiN, normally choose the last category as reference group. However, in practice it is better if the reference category is based on a larger sample size. One possibility would be to collapse the ‘strongly disagree’ and ‘disagree’ categories in order to have a sufficiently large sample in the reference group. We prefer, however, to preserve this information and to reverse the coding for the response variable. That is to say ‘strongly disagree’ is now coded 1 and ‘strongly agree’ is now coded 5. Hence, the ordinal model will now use ‘strongly agree’ as the reference group for the response variable.

Table 6: Frequency distribution of divorce attitude at waves 2, 4 and 6.

	Wave 2		Wave 4		Wave 6	
	Frequency	Percent	Frequency	Percent	Frequency	Percent
Strongly disagree	60	0.77	69	1.02	38	0.6
Disagree	351	4.48	329	4.87	268	4.22
Neither agree nor Disagree	1,244	15.88	1,114	16.49	1,082	17.05
Agree	4,270	54.51	3,726	55.14	3,581	56.44
Strongly agree	1,909	24.37	1,519	22.48	1,376	21.69
Total	7,834	100	6,757	100	6,345	100

3. Methodology

In section 3.1, we specify an ordered logit model without taking into account the hierarchical structure of data. This is used in our cross-sectional analyses. The model also provides a base for specifying a multilevel ordered logit model as described in section 3.2.

3.1 Specification of an ordered logit model

As discussed by Long (1997) the ordinal regression model was developed independently in the social sciences (in terms of an underlying latent variable with observed, ordered categories) and in biostatistics (where it is referred to as a proportional odds model). Below we follow Long (1997) and introduce the ordered logit model in the form of the latent variable model.

Suppose that a response variable has M categories, indexed by $m = 1, \dots, M$. The observed ordinal response y is thought of as providing incomplete information about an underlying y^* according to the measurement equation

$$y_i = m \text{ if } \tau_{m-1} \leq y_i^* < \tau_m \text{ for } m = 1, \dots, M. \quad (1)$$

The subscript i represents individual respondent. The τ are thresholds or cutpoints that divide the y^* into the five values 1 to 5 corresponding to strongly agree to strongly disagree. The observed y_i is related to y_i^* according to the measurement model

$$y_i = \begin{cases} 1 \Rightarrow \text{Strongly agree} & \text{if } y_i^* \leq \tau_1 \\ 2 \Rightarrow \text{Agree} & \text{if } \tau_1 < y_i^* \leq \tau_2 \\ 3 \Rightarrow \text{Neither agree or disagree} & \text{if } \tau_2 < y_i^* \leq \tau_3 \\ 4 \Rightarrow \text{Disagree} & \text{if } \tau_3 < y_i^* \leq \tau_4 \\ 5 \Rightarrow \text{Strongly disagree} & \text{if } y_i^* > \tau_4 \end{cases}$$

In our example, the observed responses to the divorce attitude question are related to an assumed underlying continuous scale which indicates an individual's degree of support for divorce.

The structural part of the model can be expressed as

$$y_i^* = \mathbf{x}_i' \boldsymbol{\beta} + \varepsilon_i, \quad (2)$$

where \mathbf{x}_i is the vector containing the values of the explanatory variables for individual i and $\boldsymbol{\beta}$ is the vector of regression coefficients. The ε is the disturbance term, which is assumed to be logistically distributed with mean 0, variance $\pi^2 / 3$ and cumulative distribution function

$$F(\varepsilon) = \frac{\exp(\varepsilon)}{1 + \exp(\varepsilon)} = \frac{1}{1 + \exp(-\varepsilon)}. \quad (3)$$

Therefore,

$$\begin{aligned} \Pr(y^* \leq \tau_m) &= \Pr(\mathbf{x}'\boldsymbol{\beta} + \varepsilon \leq \tau_{m-1}) \\ &= \Pr(\varepsilon \leq \tau_m - \mathbf{x}'\boldsymbol{\beta}) \\ &= F(\varepsilon \leq \tau_m - \mathbf{x}'\boldsymbol{\beta}) \\ &= \frac{1}{1 + \exp(-\tau_m + \mathbf{x}'\boldsymbol{\beta})} \end{aligned} \quad (4)$$

Hence, the probability of any observed outcome y equaling m , given \mathbf{x} , is

$$\begin{aligned} \Pr(y = m) &= \Pr(\tau_{m-1} < y^* \leq \tau_m) \\ &= \Pr(y^* \leq \tau_m) - \Pr(y^* \leq \tau_{m-1}) \\ &= \frac{1}{1 + \exp(-\tau_m + \mathbf{x}'\boldsymbol{\beta})} - \frac{1}{1 + \exp(-\tau_{m-1} + \mathbf{x}'\boldsymbol{\beta})}. \end{aligned} \quad (5)$$

Putting $\tau_0 = -\infty$ gives for the first category

$$\Pr(y = 1) = \frac{1}{1 + \exp(-\tau_1 + \mathbf{x}'\boldsymbol{\beta})}$$

and putting $\tau_M = \infty$ gives for the last category

$$\Pr(y = M) = 1 - \frac{1}{1 + \exp(-\tau_{M-1} + \mathbf{x}'\boldsymbol{\beta})}.$$

Under this model, the logit of the cumulative response probabilities is

$$\begin{aligned} \text{logit}\{\Pr(y \leq m)\} &= \log\left\{\frac{\Pr(y \leq m)}{1 - \Pr(y \leq m)}\right\} \\ &= \log\left\{\frac{\Pr(y^* \leq \tau_m)}{1 - \Pr(y^* \leq \tau_m)}\right\} \\ &= \log\left(\frac{1}{1 + \exp(-\tau_m + \mathbf{x}'\boldsymbol{\beta})} \bigg/ \left[1 - \left\{\frac{1}{1 + \exp(-\tau_m + \mathbf{x}'\boldsymbol{\beta})}\right\}\right]\right) \\ &= \log\{\exp(\tau_m - \mathbf{x}'\boldsymbol{\beta})\} \\ &= \tau_m - \mathbf{x}'\boldsymbol{\beta}. \end{aligned} \quad (6a)$$

Hence, the model is also known as the proportional odds model because the odds of the event is independent of category m , and the odds ratios are assumed to be constant for all categories. Note the negative sign, which is a consequence of the structural model (2), means that a positive β coefficient has the natural interpretation, where an increase in the explanatory variable corresponds to an increase in the response variable.

An alternative specification of this model is

$$\text{logit}\{\Pr(y \leq m)\} = \tau_m + \mathbf{x}'\boldsymbol{\beta}^* \quad (6b)$$

Note that some statistical packages, including STATA, estimate the coefficients in model (6a), whereas others, including MLwiN, estimate those in model (6b). However, since model (6a) is equivalent to model (6b) with $\boldsymbol{\beta} = -\boldsymbol{\beta}^*$, it is easy to calculate the $\boldsymbol{\beta}$ with the natural interpretations.

3.2 Specification of multilevel ordered logit model

As mentioned in section 2, the BHPS is a survey with a hierarchical structure: each individual is clustered within households and households within Primary Sampling Units. To take into account this hierarchical structure, we need to use multilevel modelling techniques. By using multilevel modelling (e.g., Goldstein, 2003), we can investigate the extent to which the variation of attitude toward divorce can be attributed to individuals, households and primary sampling unit levels, and identify whether their contributions are significant or not. It is worth noting that in our research we only consider random intercept effect and therefore assume that the within individual correlation between any two of the repeated measure is the same.

To start with, we specify a two-level ordinal logit model. The first level (i) is the repeated observation, the second level (j) is the individual respondent. In this model an individual-level random effect, u_j , is added to model (6a) to give

$$\text{logit}\{\Pr(y_{ij} \leq m)\} = \tau_m - \mathbf{x}'_{ij}\boldsymbol{\beta} + u_j, \quad (7)$$

where y_{ij} is the observed ordinal response of the i th measurement for the j th individual, with corresponding explanatory variables \mathbf{x}_{ij} , and u_j is normally distributed with mean zero and variance σ_u^2 .

As mentioned above, multilevel modeling allows the residual variation in the response variable to be partitioned into components that correspond to the different levels. For discrete response models there are a number of approaches to at least approximate this partition (Goldstein, 2003). If one considers the underlying latent variable, then its level one residual variance is $\pi^2 / 3$ and the total residual variance is $\pi^2 / 3 + \sigma_u^2$. Therefore, the proportion of this total variance which can be attributed to variation between individuals (level two unit) is

$$\frac{\sigma_u^2}{\pi^2 / 3 + \sigma_u^2}. \quad (8)$$

Equation (7) can be easily extended to a three or four-level model. For example, if we also take into account the clustering of individuals within households, equation (7) can be written as

$$\text{logit}\{\Pr(y_{ijk} \leq m)\} = \tau_m - \mathbf{x}'_{ijk} \boldsymbol{\beta} + v_k + u_{jk}, \quad (9a)$$

where v_k is the household-level random effect, normally distributed with mean zero and variance σ_v^2 .

Recall that MLwiN uses the alternative specification of this model, where $\boldsymbol{\beta} = -\boldsymbol{\beta}^*$:

$$\text{logit}\{\Pr(y_{ijk} \leq m)\} = \tau_m + \mathbf{x}'_{ijk} \boldsymbol{\beta} + v_k + u_{jk}. \quad (9b)$$

There are a number of approaches for estimating the parameters for multilevel ordinal logit models. MLwiN provides four quasi-likelihood approaches: first and second order marginal quasi-likelihood (MQL) and first and second order penalised quasi-likelihood (PQL) (Rasbash et al. 2004). If one again considers the underlying latent variable, then for a three-level model the total residual variance is $\pi^2 / 3 + \sigma_u^2 + \sigma_v^2$. Therefore, the proportion of this total variance which can be attributed to variation between individuals (level two unit) is

$$\frac{\sigma_u^2}{\pi^2 / 3 + \sigma_u^2 + \sigma_v^2} \quad (10a)$$

and to between households (level three unit) is

$$\frac{\sigma_v^2}{\pi^2 / 3 + \sigma_u^2 + \sigma_v^2}. \quad (10b)$$

4. Estimation results

4.1 Cross-sectional analysis

In this section we discuss the results from our ordered logit regression analysis. We used STATA to fit the models for waves 2, 4 and 6 separately. The results are given in Table 7. As mentioned in section 3.1, STATA expresses an ordered logit model by using Formula (5a). Therefore, a positive coefficient in the table indicates an increased chance that an individual will be observed in a higher category of the response variable, for a higher value of the corresponding explanatory variable. Given our recoding of the divorce attitude, the highest category is 'strongly agree'. Hence, more positive parameter coefficients are associated with more positive views on divorce. The odds ratio of being in the higher of any two adjacent categories off the response, for a given category of an explanatory variable, relative to the baseline category of the explanatory variable, can be derived by taking the exponential of the estimated coefficient.

Table 7: Proportional odds regression models for divorce attitude at waves 2, 4 and 6

Variable	Wave 2			Wave 4			Wave 6		
	Coef.	Std. Err.	p-value	Coef.	Std. Err.	p-value	Coef.	Std. Err.	p-value
Sex (<i>male as ref.</i>)									
Female	0.0919	0.0449	0.0410	0.0869	0.0484	0.0730	-0.0437	0.0504	0.3860
Age (<i>21-29 as ref.</i>)									
30-44	0.0415	0.0714	0.5610	-0.0077	0.0749	0.9180	-0.0099	0.0768	0.8980
45-64	-0.1469	0.0759	0.0530	-0.2565	0.0798	0.0010	-0.2706	0.0819	0.0010
65+	-0.3123	0.0846	0.0000	-0.4535	0.0908	0.0000	-0.4719	0.0952	0.0000
Race (<i>white as ref.</i>)									
Black	0.2358	0.2213	0.2870	0.5016	0.2697	0.0630	0.5099	0.2829	0.0710
Asian and others	-0.3773	0.1627	0.0200	-0.1343	0.1736	0.4390	0.0890	0.1852	0.6310
Education (<i>degree as ref.</i>)									
A level	0.0601	0.0850	0.4800	-0.0081	0.0893	0.9280	0.0470	0.0921	0.6100
O level	0.2361	0.0650	0.0000	0.1961	0.0690	0.0040	0.2802	0.0719	0.0000
CSE level	0.4068	0.0790	0.0000	0.3301	0.0857	0.0000	0.5113	0.0888	0.0000
No qualifications	0.3783	0.0601	0.0000	0.4462	0.0650	0.0000	0.5393	0.0676	0.0000
Region (<i>London/Midlands /North England/Wales as ref.</i>)									
South	-0.1554	0.0482	0.0010	-0.1187	0.0518	0.0220	-0.1633	0.0537	0.0020
Scotland	0.2121	0.0780	0.0070	0.2747	0.0840	0.0010	0.2041	0.0892	0.0220
Already a parent (<i>No as ref.</i>)									
Yes	-0.0624	0.0663	0.3470	0.0374	0.0691	0.5880	-0.0983	0.0708	0.1650
Marital status (<i>currently married as ref.</i>)									
never married but currently cohabiting	0.4313	0.1218	0.0000	0.3151	0.1376	0.0220	0.3741	0.1400	0.0080
previously married but currently cohabiting	1.1895	0.1453	0.0000	1.5737	0.1597	0.0000	1.3296	0.1528	0.0000
previously married but currently no partner	0.5081	0.0655	0.0000	0.6442	0.0692	0.0000	0.7773	0.0714	0.0000
never married currently no partner	0.0530	0.0864	0.5400	0.1961	0.0927	0.0340	0.1035	0.0986	0.2940
Cut points									
τ_1	-4.6863	0.1568		-4.3522	0.1497		-5.0124	0.1851	
τ_2	-2.7125	0.1018		-2.5451	0.1022		-2.8769	0.1060	
τ_3	-1.1212	0.0929		-0.9991	0.0932		-1.1463	0.0937	
τ_4	1.3928	0.0935		1.5715	0.0947		1.5140	0.0948	
No. of observations			7834			6757			6345

It can be seen from Table 7 that the estimated coefficients of the variable sex have positive signs in the models for waves 2 and 4, indicating that women are more likely to agree with the statement 'it is better to divorce than continue in an unhappy marriage'. The result is significant at the 5 % level in wave 2, while it is significant at 10% level in wave 4. In wave 6, the estimated coefficient for the variable sex has a negative sign, but the result is not significant. The gender effect is however small in comparison with the age effects. Age is entered as a categorical variable with those aged 21-29 as the reference group. Compared to those in their twenties, middle aged and older people (45-64 and 65+) are less likely to agree with the statement, and hence have more conservative divorce attitudes in all three waves. For example, in wave 2, for those aged 45-64, the odds of 'strongly agree' as opposed to 'agreeing' are 86% ($100 \times \exp(-0.1469)$) of those aged in their twenties. For those aged 65+, the odds for them to 'strongly agree' rather than 'agreeing' are just 73 % ($100 \times \exp(-0.3123)$) of those in their twenties.

With regard to ethnicity, differences are small and inconsistent. Whilst there is some evidence from wave 2 that Asian and other ethnic groups have more conservative attitudes to divorce than the White population, no significant differences are found in subsequent waves.

There is a strong educational gradient in attitude towards to divorce even after controlling for individuals' own experiences of marital dissolution. We demonstrate the magnitude of these differences by calculating the predicted probability of "strongly agreeing" that divorce is better than an unhappy marriage at wave 2 for men with different levels of education (Figure 2). Other characteristics are held at the baseline and hence the probabilities refer to married men in their twenties who are childless and who are living in London/Midlands/North/Wales.

The probability for men with degree level qualifications is 0.199 ($1 - (1/1 + \exp(-1.3928))$). Whilst the probability for those with no educational qualifications is 0.266 ($1 - (1/1 + \exp(-1.3928 + 0.3783))$).

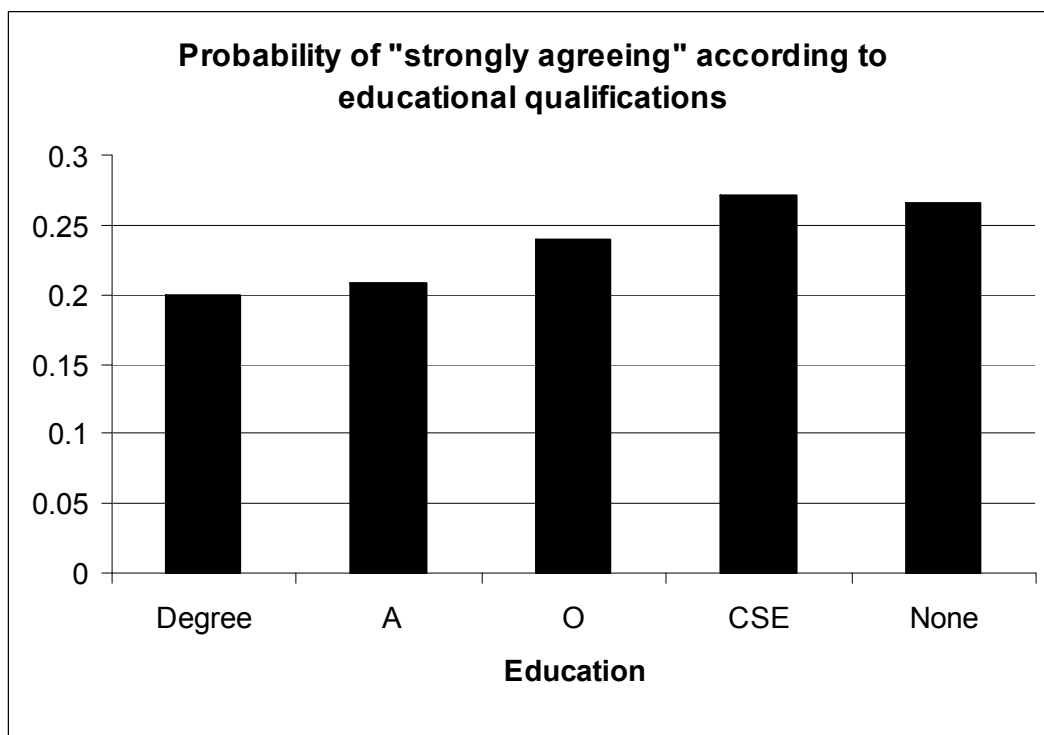


Figure 2: Predicted probability of strongly agreeing with the statement “Divorce is better than an unhappy marriage” in wave 2, according to highest educational qualification. (Note that remaining characteristics held at baseline level.)

There is a significant variation across geographical regions in terms of attitude toward divorce. The effect size is comparable to the effect of age group. The results suggest that people in Southern England are more likely to disagree, whereas people in Scotland are more likely to agree with divorce. No difference in divorce attitude is seen according to whether the respondent has had children.

Differences in attitude to divorce according to marital status are, not surprisingly, large. Compared to currently married persons, all other marital status groups, apart from those who are never married and not currently with a partner, are significantly more likely to approve of divorce. For example, in wave 2, almost half of those who have previously divorced or separated and who were then cohabiting strongly agreed with the statement, compared to 28% of those who had never been married but who were cohabiting, and 21% of those who had never been married and who were currently without a partner (Figure 3). These marital status effects are relatively stable across the three waves of the BHPS.

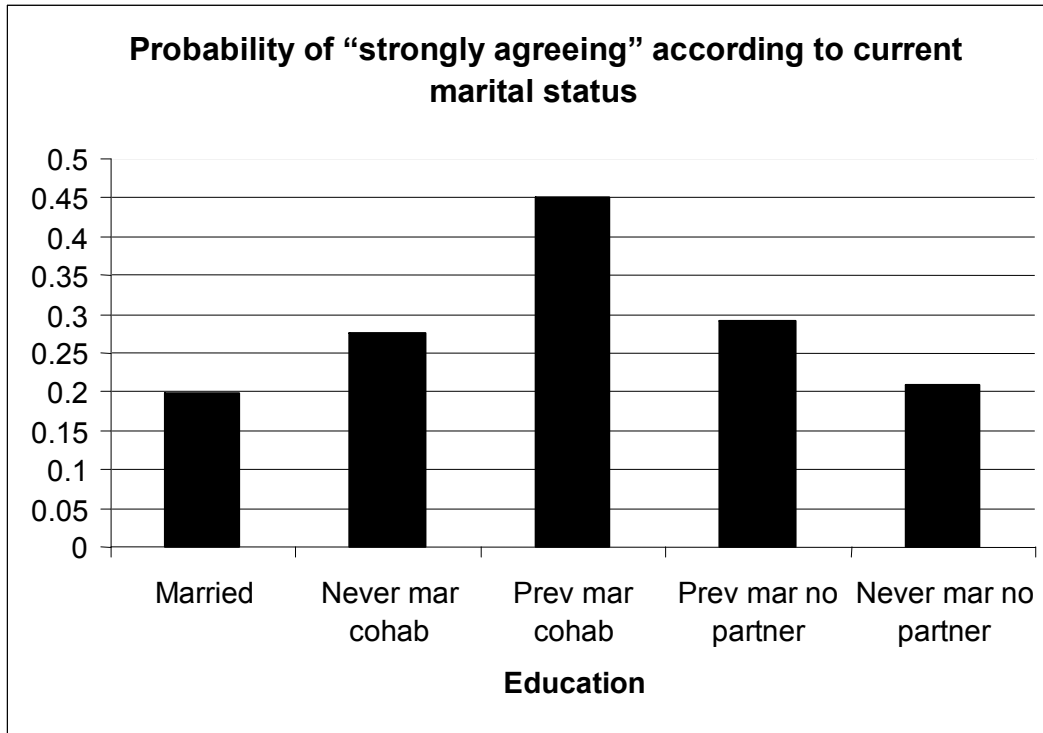


Figure 3: Predicted probability of strongly agreeing with the statement “Divorce is better than an unhappy marriage” in wave 2, according to current marital status. (Note that remaining characteristics held at baseline level.)

We can formally test whether the variables ‘race’ and ‘already a parent’ significantly improve model fit using the likelihood ratio (LR) test. Comparison of the full model with the model without race suggests that in wave 2 the variable race does (just) improve model fit at the five percent level (LR=6.57, p-value=0.0374), whereas in waves 4 and 6 it does not improve model fit. The LR test suggests that the inclusion of the variable indicating whether the respondent was a parent or not is not significant in any wave. Given these findings we decided not to include these two variables in our multilevel analyses.

4.2 Longitudinal analysis: multilevel ordered logit model

In this section we exploit the longitudinal nature of the data and investigate whether there is significant change in attitude over time, whether the predictors of divorce attitude change over time, and whether changes in marital status are associated with change in divorce attitude. Finally we test whether, once the individual characteristics are controlled, there are similarities in attitude within a household not captured in our model. We first re-shape the data from wide

to long-form. That is to say each line of data now corresponds to an observation. For an individual who was observed in waves 2, 4 and 6 there are three lines of data. For an individual observed in just waves 2 and 4, two lines of data and so on.

We then estimate two, three and four-level models with PSU as the highest level. A number of estimation techniques were tried; first and second order marginal and penalized quasi-likelihood (MQL1, MQL2, PQL1 and PQL2) and Bayesian Monte Carlo markov chain (MCMC) algorithms using the MLwiN software (Browne, 2003). The random effects were found to be larger using the penalized quasi-likelihood and MCMC techniques. The estimation bias that we found using MQL is consistent with that shown for binary responses by Rodríguez and Goldman (1995; 2001). We were not able to get 2nd order PQL estimation to converge and hence present two sets of results: estimates from 1st order PQL estimation and estimates from MCMC simulation. No significant random effect was found at the PSU level and hence we present the results for the three-level model (Table 8).

It is worth noting that MLwiN specifies an ordered logit model by using Formula (9b). Therefore, a negative coefficient indicates an increased chance that a subject will be observed in a higher category of the response, for a higher value of the explanatory variable. A positive coefficient indicates that the chance that a subject with higher score on the explanatory variable will be observed in a lower category. As before, odds ratios can be derived by taking the exponent of the estimated coefficients. MLwiN provides the standard errors of the coefficients. Z statistics can be calculated by dividing each coefficient by its standard error. For a two-tailed test at the 5% level, the critical values of the Z statistic are -1.96 and 1.96.

For the fixed effects in Table 8, the substantive conclusions from the 1st-order PQL and MCMC estimation are the same. The coefficients from the MCMC estimation are larger than those from the PQL estimation (as found previously by Rodríguez and Goldman, 2001) but so are the standard errors meaning that the significance is similar. Comparison of Table 7 and Table 8 suggest that the size and direction of the coefficients for the fixed covariates are similar to the repeated cross-sections. However, for time-varying covariates, such as marital status, our substantive interpretation changes: We now interpret the coefficient for previously married and currently cohabiting (1.77 from MCMC) as the log odds ratio associated with **moving** from being married to previously married and currently cohabiting. In other words the odds of strongly agreeing with the statement that “divorce is better than an unhappy marriage” are six

times higher for respondents who divorced and began living with a new partner in the two years between biennial waves as compared with respondents who remained married.

Table 8: Three-level random intercept estimation of attitude toward divorce

Variable	1 st – order PQL			MCMC – 25,000 iterations		
	Coefficient	Std. Err.	Z-value	Coefficient	Std. Err.	Z-value
Sex (<i>male as ref.</i>)						
Female	-0.096	0.045	-2.133	-0.110	0.051	-2.157
Age (<i>21-29 as ref.</i>)						
30-44	0.024	0.075	0.320	0.035	0.087	0.402
45-64	0.308	0.079	3.899	0.370	0.091	4.066
65 +	0.534	0.092	5.804	0.633	0.107	5.916
Education (<i>high level as ref.</i>)						
A level	-0.012	0.090	-0.133	-0.012	0.103	-0.117
O level	-0.251	0.070	-3.586	-0.288	0.082	-3.512
CSE level	-0.475	0.085	-5.588	-0.562	0.097	-5.794
No qualifications	-0.484	0.065	-7.446	-0.558	0.076	-7.342
Region (<i>London/Midlands/North England/Wales as ref.</i>)						
South	0.182	0.056	3.250	0.215	0.064	3.359
Scotland	-0.288	0.091	-3.165	-0.345	0.105	-3.286
Marriage (<i>currently married as ref.</i>)						
Never married currently cohabiting	-0.349	0.109	-3.202	-0.384	0.122	-3.148
Previously married and currently cohabiting	-1.503	0.134	-11.216	-1.771	0.151	-11.728
Previously married but currently no partner	-0.716	0.064	-11.188	-0.840	0.074	-11.351
Never married currently no partner	-0.170	0.076	-2.237	-0.194	0.087	-2.230
Time (<i>1992 as ref.</i>)						
1994	0.125	0.035	3.571	0.149	0.037	4.027
1996	0.127	0.035	3.629	0.153	0.038	4.026
Cut points						
_cut1	-5.886	0.124		-6.638	0.140	
_cut2	-3.669	0.093		-4.249	0.110	
_cut3	-1.564	0.087		-1.913	0.104	
_cut4	1.890	0.088		2.202	0.106	
Individual level random effect variance	1.883	0.084		2.855	0.138	
Household level random effect variance	0.852	0.077		1.186	0.112	
No. of observations			20,936			20,936

In the longitudinal model we also introduce time as a covariate. Here two dummies are included for 1994 and 1996. In MLwiN the positive coefficients associated with the subsequent time period suggest that over time, BHPS respondents become slightly less positive in their divorce attitude. By including interactions between time and the other fixed covariates, e.g. gender, it is possible to test whether certain groups e.g. men, are more likely to change their attitude.

However, in this analysis no interactions with time were found to be significant. Our conclusion is that the predictors of divorce attitude are stable in their effect over time.

The random effect variances at the individual and household level are both significant suggesting that there is unmeasured heterogeneity at the level of the individual and household which is not captured by the covariates in the model. We can demonstrate the magnitude of a random effect by looking at the odds ratio associated with a value of the random effect one standard deviation above and below its mean. Using the estimates from the MCMC analysis, the standard deviation of the individual random effect is $\sqrt{2.855} = 1.690$. Hence, a value of the random effect one standard deviation above its mean corresponds to an odds ratio of 5.42 ($\exp(1.690)$), whereas a value of the random effect one standard deviation below the mean corresponds to an odds ratio of 0.18 ($\exp(-1.690)$). The effect of the random intercept is therefore larger than the fixed time-constant covariates and of a similar magnitude to the time-varying marital status covariate. This suggests that there are important individual level factors not measured in the model which result in some individuals being much more prone to approve of divorce than others.

It is sometimes of interest to compare the proportions of the residual variance of the underlying latent variable explained by the various levels of the hierarchy, and to compare these with results from other studies. From equation 10, the proportion attributed to the between individual variation is $2.855 / (\pi^2 / 3 + 4.041) = 0.39$; and the proportion attributed to the between household variation is $1.186 / (\pi^2 / 3 + 4.041) = 0.16$. (The remaining, between time point, variation thus accounts for $1 - 0.39 - 0.16 = 0.45$ of the residual variation.) This tells us that there is unexplained variability between households but that within households there are significant unmeasured differences between individuals.

5. Discussion

This paper had both substantive and methodological aims. In terms of new substantive findings, this work has found that, in general, the British population are supportive towards divorce. Very few people are willing to report that they 'strongly disagree' that "divorce is better than an unhappy marriage". Between 1992 and 1996 there was little aggregate change in the

distribution of responses to this question. Furthermore, the significant predictors of divorce attitude remained the same over this six year period. New insights for Britain provided by this work include a persistent gender difference in divorce attitude whereby women are more favourable than men. This finding is consistent with work in the US where the gender difference has been explained by the fact that custody arrangements for children tend to impact more negatively on the relationship between men and their children than for women (e.g. Thornton, 1985).

We find evidence that attitudes to divorce are strongly linked to past and current marital status. Those who have experienced divorce themselves are more positive about divorce. This is especially the case for those who have been divorced and are now cohabiting with another partner. It is difficult to tease out from such data whether the positive divorce attitude facilitated divorce (a selection effect), or whether the individuals' attitudes are adapted to reflect their divorce experience.

We have also demonstrated strong educational differentials in attitude to divorce which remain even after the respondent's own marital status is taken account of. It is not clear to us what these educational differences reflect. Do they reflect differential exposure in their family, or social network to marital dissolution e.g. through the breakdown of their parent's marriage, or the marriage of other friends/relatives? Given that age at marriage is one of the strongest predictors of divorce risk in Britain (Berrington and Diamond, 1999) then there will be a greater experience and perhaps more acceptability of divorce among those with less education.

In addition, the BHPS reveals some small but persistent regional differences in divorce attitude which, according to our knowledge, have not previously been commented on. Future research is needed to examine reasons why Scottish respondents were more positive about divorce, for example to assess whether this extends from historically different divorce legislation.

Previous analyses have found that family members tend to have similar political views and behaviour (Brynin, 2000; Johnston et al, 2005). In this paper we have found substantial within-household agreement on attitudes to divorce. This within-household similarity is not accounted for by the observed characteristics of the individual members in terms of their age, marital history, educational level, or region of residence. Similarities in the attitudes of partners are likely to have existed before they lived together in the same household (as are result of marital

homogamy), but are also likely to have developed as a result of subsequent shared experiences. Young adults who remain in the parental home may be more likely to share their parents attitudes, for example through shared experience (although this might not necessarily follow - for example a past marital breakup may be viewed by the parent in a positive light but viewed in a more negative way by their children).

Our methodological aim was to demonstrate the application of multi-level models to repeated ordinal measures. Previous research in this area has tended to focus on binary outcomes. Our experience with fitting variance component models using MLwiN to ordinal outcomes suggests that random effects are likely to be underestimated using marginal and penalized quasi-likelihood estimation methods. Unlike for the binary case (Rodríguez and Goldman, 2001), MLwiN does not contain a facility for a parametric bootstrap to reduce the bias of the MQL and PQL estimates. Hence we believe it advisable to estimate such models using MCMC. We have also demonstrated how this approach to modelling repeated measures is useful in situations where you have complex patterns of attrition and wave non-response. By including respondents who were later lost to the survey either through attrition or wave non-response we substantially increased the sample size available for analysis.

The fact remains however, that our results will reflect differential response within the BHPS sample. The BHPS does in fact provide an individual level respondent weight which is available for those who took part in every wave up until the wave of interest. However, as discussed by Skinner and Holmes (2004) the methodology to incorporate weights in to the analysis has not yet been developed for ordinal outcomes and estimation is not available in MLwiN. We would speculate however, that our results would remain relatively unchanged given that many of the key variables which predict loss-to-follow up within the BHPS e.g. age, marital status, education, are already included as covariates in the model.

Finally, planned future extensions of this work include the inclusion of random slopes (which would allow subject-specific random variation for each of the covariates) and the testing of alternative model specifications (which would relax the assumption of proportionality in the odds model).

References

- Berrington, A. and Diamond, I. (1999) Marital dissolution among the 1958 British birth cohort: The role of cohabitation. *Population Studies*, 53, 19-38.
- Browne, W. (2003) *MCMC estimation in MLwiN*. Centre for Multilevel Modelling, Institute of Education, University of London.
- Brynin, M. (2000) Political values: a family matter? Pp 193-214 in R. Berthoud, J. Gershuny (Eds.) *Seven Years in the Lives of British Families: Evidence on the Dynamics of Social Change from the British Households Panel Study*. Bristol: The Policy Press.
- Goldstein, H. (2003) *Multilevel Statistical Models*. Third Edition. London:Arnold.
- Johnston, E., Jones, K., Propper, C., Sarker, R., Burgess, S. and Bolster, A. (2005) A missing level in the analyses of British voting behaviour: the household as context as shown by analyses of a 1992-1997 longitudinal survey. *Electoral Studies*, 24, 201-225.
- Liu, I. and Agresti, A. (2005) The analysis of ordered categorical data: an overview and a survey of recent developments. *Sociedad de Estadística e Investigación Operativa Test*, 14, 1-73.
- Rasbash, J., Steele, F., Browne, W., Prosser, B. (2004) *A user's guide to MLwiN version 2.0*. London, Institute of Education.
- Rodríguez, G. and Goldman, N. (1995) An assessment of estimation procedures for multilevel models with binary responses. *Journal of the Royal Statistical Society Series A*, 158, 73-89.
- Rodríguez, G. and Goldman, N. (2001) Improved estimation procedures for multilevel models with binary response: a case study. *Journal of the Royal Statistical Society Series A*, 164, 339-356.
- Skinner, C. J. and Holmes, D. J. (2003) Random Effects Models for Longitudinal Survey Data. Pp 205-219 in R. L. Chambers and C. J. Skinner (eds.) *Analysis of Survey Data*. Chichester: John Wiley & Sons.
- Taylor, Marcia Freed (ed.) with John Brice, Nick Buck and Elaine Prentice-Lane (2005) *British Household Panel Survey User Manual Volume A: Introduction, Technical Report and Appendices*. Colchester: University of Essex.



ESRC Research Methods Programme
Cathie Marsh Centre for Census and
Survey Research
Faculty of Social Sciences and Law
The University of Manchester
Manchester
M13 9PL
United Kingdom

Tel: 0161 275 4891
Email: methods@man.ac.uk
Web: <http://www.ccsr.ac.uk/methods/>

Director: Professor Angela Dale
Administrator: Ruth Durrell

