TI 2020-069/III
Tinbergen Institute Discussion Paper

A multinomial and rank-ordered logit model with inter- and intra-individual heteroscedasticity

Anoek Castelein[1]
Dennis Fok[1]
Richard Paap[1]

[1] Econometric Institute, Erasmus University Rotterdam

A multinomial and rank-ordered logit model with inter- and intra-individual heteroscedasticity*

Anoek Castelein[a] Dennis Fok[a] Richard Paap[a]

[a] Econometric Institute, Erasmus School of Economics, Erasmus University Rotterdam

September, 2020

Abstract

The heteroscedastic logit model is useful to describe choices of individuals when the randomness in the choice-making varies over time. For example, during surveys individuals may become fatigued and start responding more randomly to questions as the survey proceeds. Or when completing a ranking amongst multiple alternatives, individuals may be unable to accurately assign middle and bottom ranks. The standard heteroscedastic logit model accommodates such behavior by allowing for changes in the signal-to-noise ratio via a time-varying scale parameter. In the current literature, this time-variation is assumed equal across individuals. Hence, each individual is assumed to become fatigued at the same time, or assumed to be able to accurately assign exactly the same ranks. In most cases, this assumption is too stringent. In this paper, we generalize the heteroscedastic logit model by allowing for differences across individuals. We develop a multinomial and a rank-ordered logit model in which the time-variation in an individual-specific scale parameter follows a Markov process. In case individual differences exist, our models alleviate biases and make more efficient use of data. We validate the models using a Monte Carlo study and illustrate them using data on discrete choice experiments and political preferences. These examples document that inter- and intra-individual heteroscedasticity both exist.

Key words: Scale, Heterogeneity, Markov, Logit scaling, Logit mixture, Dynamics, Conjoint, Fatigue, Markov switching.

*We thank Kathrin Gruber, Hans van Ophem, Michel Wedel, and participants at the 13th Netherlands Econometric Study Group (Amsterdam, 2019), the 1st International Econometrics PhD Conference (Rotterdam, 2019) and the 7th Econometric Institute PhD Conference (Rotterdam, 2019) for their helpful comments and suggestions. In this paper we make use of data of the LISS (Longitudinal Internet Studies for the Social sciences) panel administered by CentERdata (Tilburg University, The Netherlands).

1. Introduction

Understanding and predicting choices of individuals is important for numerous applications. These applications include predicting product demand, designing effective policies, and constructing meaningful product recommendations. In many cases, individuals are observed while making repeated choices over time. For example when responding to survey questions, choosing supermarket products across different visits, or completing a ranking by consecutively choosing best, second best, et cetera. To deduce an individual's preferences based on observed discrete choices, the (multinomial) logit model is often employed.

The logit model is based on a utility framework: an individual obtains utility from choosing a certain alternative/option and chooses the alternative which gives the highest utility (Manski, 1977). The utility is comprised of an explained part (the preferences/signal) and an unexplained part (the noise). The noise captures that the actual choice can differ from the choice that yields the highest signal. Hence, the noise can capture that (i) the signal fails to capture all preferences of an individual, and/or (ii) an individual can make 'mistakes' and choose an alternative that does not accord with her underlying preferences. Logit models assume an extreme value distribution for the noise.

The logit model has been extended in many ways to realistically capture certain aspects of individual behavior. One such aspect of behavior is that the randomness in the choice-making of individuals may vary over time. For example, during surveys, individuals may become fatigued and start responding more randomly to questions as the survey proceeds. Or when completing a ranking amongst multiple alternatives, individuals may be unable to accurately assign middle and bottom ranks, due to the required cognitive effort or lack of information. For supermarket purchases, an individual that is new to a certain product category (e.g. diapers) may at first pick alternatives quite randomly after which the preferences are learned and choices are more and more based on the underlying preferences.

The heteroscedastic logit model is able to estimate individual preferences while accounting for changes in the randomness in choices (Hausman and Ruud, 1987, Bradley and Daly, 1994). For this purpose, the model explicitly allows for changes in the relative importance of the explained and the unexplained part of utility (the signal-to-noise ratio). When choices become more random, this can be captured in the unexplained part becoming more dominant. Mathematically, the heteroscedastic logit model allows for changes in the signal-to-noise ratio via a time-varying scale parameter in the unexplained part of the utility specification.

The main drawback of the standard heteroscedastic logit model is that the scale parameter is specified at the population-level. Hence, the model assumes that the changes in the randomness in decision-making is equal across individuals, thereby only allowing for within-individual (intra-individual) heteroscedasticity. In the context of the earlier examples, this implies that each individual is assumed to become fatigued at the same time, or assumed to be able to accurately assign exactly the same ranks.

In this paper, we generalize the heteroscedastic logit model by allowing for differences across individuals in the changes in the scale parameter. That is, we allow for intra- and inter-individual heteroscedasticity (or *heterogeneous heteroscedasticity*): each individual has her own sequence of scale parameters over time, and the time-variation in the scale parameters can differ across individuals. For example, for some individuals the scale parameters may stay constant, for others the scale parameters may increase several times, and for again some others the scale parameters may first decrease and then increase.

In case such individual differences exist, using an individual-level instead of a population-level approach is beneficial for several reasons. First, existing population-level approaches generally lead to biased estimators for the preference parameters. That is, there will be a bias towards zero, because at each time period a number of individuals could be answering more randomly.[1] Second, population-level approaches make inefficient use of data. This is because it is assumed that at each time period, each individual provides the same amount of information in her choices. Finally, population-level approaches only give insight into the average time-variation in the scale parameter. In some cases, one might find a constant average scale parameter while in reality there is heterogeneous heteroscedasticity. An individual-level approach also gives more insight into the behavior of different individuals. To allow for individual differences, some structure is needed to model the heteroscedastic process.

We develop a multinomial logit model (MNL) and a rank-ordered logit model (ROL) that allow for heterogeneous heteroscedasticity. For this purpose, we include individual- and time/rank-specific scale parameters. We let the dynamics in the sequence of an individual's scale parameters be governed by a Markov process. We also allow for unobserved preference heterogeneity. For inference, we develop a maximum simulated likelihood estimation approach.

The Markov process assumes that, for subsequent choices to make or for consecutive ranks to assign, an individual can go through a number of phases. Each phase is marked

[1]Even when no individual differences exist, the existing estimators for the preference parameters in the heteroscedastic logit model (Bradley & Daly, 1994) are often biased away from zero, because the preference parameters are scaled such that the time period with lowest estimated signal-to-noise ratio has a scale of one. The estimator for our proposed model does not suffer from this shortcoming.

by a different scale parameter. When an individual moves to a phase with higher scale, the choices become more random. When an individual moves to a phase with lower scale, the choices become more predictable and more in line with the underlying preferences. For the example of fatigue during surveys, there may exist phases with a rather high scale. Respondents who become fatigued, enter these phases with high scale and answer more randomly as the survey proceeds. Respondents that do not become fatigued, remain in a phase with low scale.

In the literature, a related individual-level ROL has been proposed in Fok et al. (2012). They propose a latent class ROL where they allow for individuals to have different ranking abilities: different individuals may be able to assign a different number of top ranks accurately. When applied to the ROL, our approach can be seen as a generalization of Fok et al. (2012). First, we allow for unobserved preference heterogeneity and allow the model to be used for panel data. Especially in the context of non-constant scale, allowing for heterogeneity is important to avoid spurious findings. That is, when preference heterogeneity is unaccounted for, an individual who has preferences that deviate from the "average" individual is likely falsely classified as having a large scale parameter. Second, Fok et al. (2012) allow for individuals to assign a specific rank either accurately or completely randomly. Instead, our model allows for the decisions of individuals to be more in between, which is also possible in the standard heteroscedastic ROL. As a consequence, our model is a generalization of the heteroscedastic ROL, whereas the latent class ROL is not. Finally, our model straightforwardly and parsimoniously allows for individuals that might rank the middle ranks randomly, but both the top and bottom ranks accurately. This possibility was already provided as an extension in Fok et al. (2012) but requires work on top of the basic model specification.

We illustrate the usefulness of the newly proposed hidden Markov model specifications using a Monte Carlo study and two empirical applications. In the Monte Carlo study, we find that our proposed model works well and that the estimator seems unbiased in various settings. Furthermore, this study clearly illustrates the bias in the estimator for the preference parameters for the standard heteroscedastic logit model. Depending on the data generating process, the bias is either towards zero due to neglecting individual differences, or away from zero due to scaling the preference parameters based on the minimum of the estimated scale parameters. Furthermore, the estimator for the standard MNL is biased towards zero in case heteroscedasticity is present, because heteroscedasticity leads to more random-looking choice-making of respondents. Our proposed estimator and model alleviate these biases.

In the first empirical application, we consider binomial choices during a discrete choice experiment on healthy food choices. We allow for multiple phases to capture possible learning and fatigue effects. We find that accounting for individual differences in learning and fatigue leads to a much better fit of the data, while needing less free model parameters than the standard heteroscedastic logit model. In the second empirical application, we consider rank-ordered data from a survey on political preferences to capture possible differential capabilities in ranking. Again, allowing for individual differences in the dynamics of the heteroscedasticity leads to a much better fit of the data.

This paper is set up as follows. In Section 2, we discuss the background and related literature. In Section 3, we develop the hidden Markov MNL and ROL, and discuss identification and estimation. In Section 4, we report the results of a Monte Carlo study. In Sections 5 and 6, we report the results of the two empirical applications. Finally, we provide a discussion and conclude.

2. Background

In this section, we discuss the additive random utility framework (ARUM) we employ in our paper. This framework is central in deciding how to model individual-specific dynamics in the signal-to-noise ratio. We illustrate the identification problem that may arise, and discuss related papers that have proposed solutions to this. We also indicate how our approach differs from current specifications dealing with individual-specific dynamics in the signal-to-noise ratio.

The additive random utility framework of Manski (1977) is a useful and popular tool to model choices of individuals. It relies on the assumption that an individual obtains utility from a certain alternative and that an individual chooses the alternative that gives the highest utility. The utility is assumed to be an additive function of the signal (based on observed variables and unobserved parameters) and some noise

$$\text{Utility} \quad = \quad \text{Signal} \quad + \quad \text{Noise.}$$

Mathematically, we can write this utility specification in a general form as

$$U_{itj} = x'_{itj}\beta_{it} + \sigma_{it}\varepsilon_{itj}, \tag{1}$$

where U_{itj} is the (unobserved) utility that individual i obtains from choosing alternative j

at time t, x_{itj} is a vector of covariates representing the attributes of alternative j, β_{it} is a vector with preference parameters of individual i at time t, $\sigma_{it} > 0$ is a scale parameter for individual i at time t, and ε_{itj} is an i.i.d. error term with fixed variance.[2] The individual chooses the alternative that gives the highest utility. The multinomial logit and probit models are special cases of the ARUM.

Because only choices are observed and not utility, and the scale parameter does not vary over alternatives, we obtain an equivalent model for choices by rescaling the utility

$$U^*_{itj} = x'_{itj}\frac{\beta_{it}}{\sigma_{it}} + \varepsilon_{itj}, \tag{2}$$

where now the alternative is chosen with the highest scaled utility $U^*_{itj} = U_{itj}/\sigma_{it}$. The equivalence between the utility specifications in Equations (1) and (2) imply that only the signal-to-noise ratio (β_{it}/σ_{it}) is identified, and not the absolute values of the signal and the noise. Hence, if we would allow for both $\beta_{it} = \beta_i$ and $\sigma_{it} = \sigma_i$ to be individual-specific, separate identification of the two parameters can only come from distributional assumptions on these two parameters (Hess & Rose, 2012). The same holds when we allow both $\beta_{it} = \beta_t$ and $\sigma_{it} = \sigma_t$ to be time-dependent.

Therefore, for identification, the proposed models in the literature often allow for heterogeneity and time-variation in either β_{it} or σ_{it}. For example, the heteroscedastic multinomial and rank-ordered logit models allow for a (possibly) individual-specific β_i and a time-dependent scale σ_t (Hausman and Ruud, 1987, Bradley and Daly, 1994, DeSarbo et al., 2004).

We generalize the heteroscedastic multinomial logit model by allowing for the time-variation in the scale to be different across individuals (σ_{it}). We allow for individual-specific preference parameters in β_i, but exclude time-variation in this parameter. We ensure identification by letting the sequence of scale parameters of an individual $\{\sigma_{i1}, \sigma_{i2}, ...\}$ be governed by a Markov process with the scale of one state normalized at one, as will be shown later.[3]

Alternatively, one can allow for individual-specific heteroscedasticity by letting β_{it} be

[2] A more general specification can be obtained by allowing for correlation across the error terms ε_{itj} over individuals, time periods, and/or alternatives.

[3] Bhat and Castelar (2002) propose a multinomial logit model with individual-specific preference parameters β_i and individual- and time-specific scale parameters σ_{it}. However, their formulation is highly restrictive. The scale parameter σ_{it} can take on only one of two values and which of the two values it takes on is determined deterministically: $\sigma_{it} = 1$ in case observation t of individual i corresponds to a revealed preference observation, and $\sigma_{it} = \lambda$ in case it corresponds to a stated preference observation, with λ a parameter to be estimated. Hence, the variation in scale parameters only allows for the scale to be different across different types of data.

individual- and time-specific, and keeping $\sigma_{it} = \sigma$ constant over individuals and time. An advantage of this alternative approach is that it can model the change in the signal-to-noise ratio to be different across attributes, and can thus capture choice strategies where choices are made based on different subsets of attributes as time progresses or where preferences change over time. The main disadvantage is that, in small T settings, estimation uncertainty and overfitting become problematic.

There are three papers that propose discrete choice models with individual-specific time-variation in β_{it}: Hess and Rose (2009), Bhat and Sidharthan (2011), and Danaf et al. (2020). These three papers all propose a model with constant scale σ and preference parameters of the form $\beta_i + \beta_{it}$. Both β_i and β_{it} are allowed to follow arbitrary distributions, with the restriction that the unconditional mean of β_{it} is zero. The papers differ in the type of model (logit versus probit), estimation approach and the distributional form used for β_i and β_{it}. These approaches are quite general, but have the main disadvantage that they assume an additive specification $\beta_i + \beta_{it}$. For individual-specific heteroscedasticity in discrete choice models, a multiplicative specification via $\beta_i\beta_{it}$ (or β_i/σ_{it}) is more suitable, as choices becoming more random directly affect the signal-to-noise *ratio*. That is, increased randomness in choice-making leads to signal-to-noise ratios that become closer to zero. In a multiplicative specification, this can be modeled by a low β_{it} (or a high σ_{it}). Instead, with an additive specification, a given β_{it} could shrink the $\beta_i + \beta_{it}$ of one individual to zero, whereas for another individual it can make it more extreme or let it flip signs. Due to the additive nature of these approaches, they are less suited to model heterogeneous heteroscedasticity. Instead, we use a multiplicative specification.

A related strand of literature considers (time-invariant) scale heterogeneity: some individuals may choose more randomly throughout the observed period than others. Fiebig et al. (2010) propose a so-called generalized multinomial logit model that includes both individual-specific preferences β_i and an individual-specific scale parameter σ_i. Separate identification of the two parameter is achieved by imposing parametric population distributions on β_i and σ_i (Hess & Rose, 2012). Our approach differs crucially as we focus on the time-variation in the scale parameter, to allow for changes in individual behavior over time.

3. Methodology

In this section, we develop the hidden Markov multinomial logit model and the hidden Markov rank-ordered logit model to capture inter- and intra-individual heteroscedasticity. The methods are highly similar, the main difference is that for the MNL the heteroscedasticity refers to the change in the scale parameter as time progresses, and for the ROL the heteroscedasticity refers to the change in the scale parameter across consecutive ranks for the same ranking task.

Let us introduce some basic notation. We index the individuals by $i = 1, ..., N$, the observations for individual i by $t = 1, ..., T$, and the alternatives that individual i can choose between, or need to rank, at time t by $j = 1, ..., J$.[4] Furthermore, we denote by x_{itj} a $(K \times 1)$ vector of covariates representing the attributes of alternative j at time t for alternative j, and by β_i a $(K \times 1)$ vector with individual-specific preference parameters corresponding to x_{itj}.

3.1. Hidden Markov multinomial logit model

For the multinomial logit model, we let the scalar $y_{it} \in \{1, 2, ..., J\}$ denote the alternative that individual i chooses at time t, and let Y_{it} denote the corresponding random variable. The latent utility that individual i obtains from choosing alternative j at time t is given by

$$U_{itj} = x'_{itj}\beta_i + \sigma_{it}\varepsilon_{itj}, \tag{3}$$

where $\sigma_{it} > 0$ is an individual- and time-specific scale parameter and the error terms ε_{itj} follow independent type I extreme value distributions with location 0 and scale 1. In case the scale parameter is equal across individuals ($\sigma_{it} = \sigma_t$), we obtain the heteroscedastic multinomial logit model. In case the scale parameter is also equal over time ($\sigma_{it} = \sigma = 1$), we obtain the standard (mixed) multinomial logit model.

At each time t, an individual chooses the alternative that yields the highest utility. Given the utility specification in Equation (3), it follows that the conditional probability that individual i chooses alternative j at time t is given by (McFadden, 1973)

$$\Pr[Y_{it} = j | \beta_i, \sigma_{it}] = \frac{\exp\left(\frac{1}{\sigma_{it}}(x'_{itj}\beta_i)\right)}{\sum_{l=1}^{J} \exp\left(\frac{1}{\sigma_{it}}(x'_{itl}\beta_i)\right)}. \tag{4}$$

The scale parameter σ_{it} captures heteroscedasticity. The higher σ_{it}, the lower the signal-to-noise ratio and the more random the choice of individual i at time t becomes. For example, when an individual becomes tired during a survey and starts to answer more randomly, this can be modeled by a sequence of scales $\{\sigma_{it}\}_{t=1}^{T}$ that increases over time. In the extreme

[4] In this notation, the number of observations T is equal across individuals and the number of alternatives J is equal across observations and individuals. These assumptions can be easily relaxed.

case that σ_{it} tends to infinity, the choice becomes completely random. In the other extreme case that σ_{it} is close to 0, the choice can be perfectly explained by the signal $x'_{itj}\beta_i$.

We let the time variation in the sequence of an individual's scale parameters $\{\sigma_{it}\}_{t=1}^{T}$ be governed a Markov process. Such a process assumes that while an individual is making choices, she can go through a number of phases. Each phase is marked by a different scale parameter. When an individual moves to a phase with higher scale, the choices become more random. When an individual moves to a phase with lower scale, the choices become more predictable and more in line with the underlying preferences. For the example of fatigue during surveys, there may exist phases with a rather high scale. Respondents who become fatigued, enter these phases with high scale as they answer more randomly as the survey proceeds. Respondents that do not become fatigued, remain in a phase with low scale.

Let M denote the number of possible phases an individual can go through, with M set by the researcher. The number of different scale parameters is equal to M: $\sigma_{it} \in \{\tilde{\sigma}_1, \tilde{\sigma}_2, ..., \tilde{\sigma}_M\}$. Let s_{it} denote the phase that an individual i is in at time t. Then we have

$$\sigma_{it} = \tilde{\sigma}_{s_{it}}. \tag{5}$$

For parameter identification, the scale parameter of one of the phases needs to be fixed. This fixed scale parameter can be set to 1, such that the preference parameters β_i can be interpreted with respect to the corresponding phase.

The phase indicators $\{s_{it}\}_{t=1}^{T}$ describe how individual i moves through the M phases. These indicators are unobserved. We let the time variation in $\{s_{it}\}_{t=1}^{T}$ follow a first-order Markov process (Goldfeld & Quandt, 1973). Such a process describes how individuals move from one phase to another using transition probabilities. We denote the transition probabilities by

$$q_{mnt} \equiv \Pr[S_{i,t+1} = n | S_{it} = m], \tag{6}$$

which is the probability that individual i is in phase n at time $t+1$ given that she was in phase m at time t, and where S_{it} denotes the random variable associated with outcome s_{it}, for $m, n = 1, .., M$ and $t = 1, .., T-1$. We have that $0 \leq q_{mnt} \leq 1$ and $\sum_{n=1}^{M} q_{mnt} = 1$. Finally, we denote the initial phase probabilities by

$$\pi_m \equiv \Pr[S_{i1} = m], \tag{7}$$

with $0 \leq \pi_m \leq 1$ and $\sum_{m=1}^{M} \pi_m = 1$.

Depending on the information in the data and the type of application, it may be desired to impose restrictions on the parameters of the Markov process. For example, one can restrict the transition probabilities such that an individual can either stay in the current phase or move one phase up, also known as a change-point model (Chib, 1998). Furthermore, one may wish to restrict (some of) the values of the M different scale parameters. In case the dataset contains relatively few observations per individual, it is important to have only few phases M, e.g. $M \leq 4$. This helps to avoid overfitting, in particular finding a perfect fit phase in which the choices can be seemingly perfectly explained by the signal.

Finally, the parameters in β_i capture the preferences of individual i for the attributes in x_{itj}. We let β_i follow some distribution with density $f(\beta_i|\theta)$ where θ denotes the set of population parameters to be estimated. The type of distribution for β_i should be set by the practitioner. Examples are the multivariate normal distribution (where θ represents the mean and covariance matrix), lognormal distribution, a mixture of discrete distributions, and a mixture of normal distributions. The parameters of the distribution could also be allowed to depend on individual-specific characteristics. Moreover, in case one has large T per individual, one can directly estimate β_i without imposing a population distribution.

For identification, with sufficient variation in the variables x_{itj}, a sufficient condition for θ to be identified is that the total number of observations in the phase with fixed variance exceeds the number of parameters in θ. Additional observations in each phase are needed to identify the parameters of the Markov process.

3.2. Hidden Markov rank-ordered logit model

Next, we generalize the hidden Markov multinomial logit model in Section 3.1 to allow for rank-ordered choices. That is, we now model a partial or complete ranking over the alternatives instead of only the most preferred alternative. As the ROL uses more information than the MNL, the ROL allows for more efficient use of data. We provide the model specification for panel data but the model can also be used for cross-sectional data with $T = 1$.

Let the vector y_{it} denote the complete ranking provided by individual i at time t out of the J alternatives, and Y_{it} the corresponding random variable.[5] That is, $y_{it} = (y_{it1}, y_{it2}, ..., y_{itJ})'$, and y_{itj} denotes the alternative that was ranked j^{th}. For example, in case alternative three was ranked first, we have that $y_{it1} = 3$. Note that the MNL only models the first-ranked alternative y_{it1}.

In the rank-ordered logit model, jointly modeling the complete ranking of alternatives

[5] The specification can be easily extended to problems in which only the top J^* alternatives out of J alternatives need to be ranked.

(y_{it}) is equivalent to modeling the sequential ranking from the highest rank (y_{it1}) to the lowest rank (y_{itJ}) (Beggs et al., 1981, Chapman and Staelin, 1982). For each rank h, the choice between the "remaining" alternatives $\{y_{itl}\}_{l=h}^{J}$ can be modeled with a multinomial logit model. That is, the probability of observing y_{it} has the form

$$\begin{aligned}\Pr[Y_{it} = y_{it}|\beta_i] &= \prod_{h=1}^{J-1} \Pr[Y_{ith} = y_{ith}|y_{it1}, ..., y_{it,h-1}, \beta_i] \\ &= \prod_{h=1}^{J-1} \frac{\exp\left(x'_{ity_{ith}}\beta_i\right)}{\sum_{l=h}^{J} \exp\left(x'_{ity_{itl}}\beta_i\right)}.\end{aligned}$$

To allow for intra-individual heteroscedasticity — individuals may be more or less capable to assign the top ranks as compared to the middle and bottom ranks — Hausman and Ruud (1987) propose a heteroscedastic ROL. For this purpose, they introduce a scale parameter σ_h that may differ over ranks h.[6] More specifically, in the heteroscedastic ROL the probability of observing y_{it} is given by

$$\Pr[Y_{it} = y_{it}|\beta_i, \sigma_1, \sigma_2, ..., \sigma_{J-1}] = \prod_{h=1}^{J-1} \frac{\exp\left(\frac{1}{\sigma_h}(x'_{ity_{ith}}\beta_i)\right)}{\sum_{l=h}^{J} \exp\left(\frac{1}{\sigma_h}(x'_{ity_{itl}}\beta_i)\right)}.$$

The higher σ_h, the more random the assignment to rank h. Hence, a high σ_h indicates that individuals find it relatively difficult to assign an alternative to rank h.

We extend the approach of Hausman and Ruud (1987) to additionally allow for inter-individual heteroscedasticity: the ranking capabilities may differ across individuals. More specifically, we let the probability of observing a complete ranking y_{it} be given by

$$\Pr[Y_{it} = y_{it}|\beta_i, \sigma_{i1}, \sigma_{i2}, ..., \sigma_{i,J-1}] = \prod_{h=1}^{J-1} \frac{\exp\left(\frac{1}{\sigma_{ih}}(x'_{ity_{ith}}\beta_i)\right)}{\sum_{l=h}^{J} \exp\left(\frac{1}{\sigma_{ih}}(x'_{ity_{itl}}\beta_i)\right)}, \tag{8}$$

where intra- and inter-individual heteroscedasticity is allowed for via the rank- and individual-specific scale parameter σ_{ih}.

As with the hidden Markov MNL in Section 3.1, we let the sequence of an individual's scale parameter $\{\sigma_{ih}\}_{h=1}^{J-1}$ be governed by a Markov process. This implies that we explicitly allow for "blocks" of consecutive ranks to be assigned based on the same amount of randomness. The sizes and locations of these blocks may differ across individuals. For example,

[6] In their paper, Hausman and Ruud (1987) use the notation σ_h to denote the inverse of the scale parameter.

some individuals may assign the top three and bottom three ranks accurately and the middle ranks more randomly. Others, may assign the top two ranks and the lowest rank accurately, and the remainder more randomly. Our model allows for all these individual differences.

We let β_i follow a distribution with density $f(\beta_i|\theta)$ and use a Markov process with M phases to govern the dynamics in $\{\sigma_{ih}\}_{h=1}^{J-1}$. That is, we have the latent phase indicator s_{ih} denoting the phase that individual i is in when assigning rank h, and we have that $\sigma_{ih} = \tilde{\sigma}_{s_{ih}}$. We let s_{ih} follow a first-order Markov process with transition and initial phase probabilities

$$q_{mnh} \equiv \Pr[S_{i,h+1} = n|S_{ih} = m], \tag{9}$$

$$\pi_m \equiv \Pr[S_{i1} = m], \tag{10}$$

with $0 \leq q_{mnh} \leq 1$, $\sum_{n=1}^{M} q_{mnh} = 1$ for $m = 1, .., M$ and $h = 1, .., J-2$, $0 \leq \pi_m \leq 1$ and $\sum_{m=1}^{M} \pi_m = 1$.

Our hidden Markov ROL generalizes the latent class ROL of Fok et al. (2012). To see the equivalence: the latent class ROL has a parameter p_j denoting the proportion of individuals that can rank exactly the first j alternatives correctly and the remaining $J-j$ alternatives randomly. Hence, the hidden Markov ROL is equivalent to the latent class ROL in case we take two phases with $\tilde{\sigma}_1 = 1$ and $\tilde{\sigma}_2 = \infty$, and do not allow individuals to move from phase two to phase one ($q_{22h} = 1$ for all h). Then $p_0 = \pi_2$, $p_1 = \pi_1 q_{121}$ and $p_j = \pi_1 q_{12j} \prod_{h=1}^{j-1} q_{11h}$ for $j = 2, ..., J-1$. Also, equivalently to testing for an empty class in the latent class ROL ($p_j = 0$) one can test $\pi_1 = 0$ (class 0) or $q_{11,j-1} = 1$ (classes 1 up to $J-1$). Moreover, with the hidden Markov ROL one can test for equal transition probabilities across ranks ($q_{11,j} = q_{11,j+1}$).

3.3. Parameter estimation

To estimate the parameters of the hidden Markov MNL (HM-MNL) in Equations (3)-(7) and of the hidden Markov ROL (HM-ROL) in Equations (8)-(10), we rely on maximum simulated likelihood estimation. The likelihood functions of the models are given by

$$\begin{aligned} p(y|\theta,q,\pi,\tilde{\sigma}) &= \prod_{i=1}^{N} p(y_i|\theta,q,\pi,\tilde{\sigma}) \\ &= \prod_{i=1}^{N} \left[\int p(y_i|\beta_i,q,\pi,\tilde{\sigma}) f(\beta_i|\theta) d\beta_i \right] \\ &= \prod_{i=1}^{N} \left[\int \left(\sum_{s_i^* \in \mathcal{S}} \Pr[S_i = s_i^*|q,\pi] p(y_i|\beta_i,\tilde{\sigma},s_i^*) \right) f(\beta_i|\theta) d\beta_i \right], \end{aligned} \tag{11}$$

where $y_i = \{y_{it}\}_{t=1}^{T}$, $y = \{y_i\}_{i=1}^{N}$, $s_i = \{s_{it}\}_{t=1}^{T}$ (HM-MNL), $s_i = \{s_{ih}\}_{h=1}^{J-1}$ (HM-ROL), $\mathcal{S}$ is a set of all possible sequences of phases $s_i \in \mathcal{S}$, and

$$p(y_i|\beta_i,\tilde{\sigma},s_i) = \begin{cases} \prod_{t=1}^{T} \frac{\exp\left(\frac{1}{\tilde{\sigma}_{s_{it}}}(x'_{itj}\beta_i)\right)}{\sum_{l=1}^{J}\exp\left(\frac{1}{\tilde{\sigma}_{s_{it}}}(x'_{itl}\beta_i)\right)}, & \text{(HM-MNL)}, \\ \prod_{t=1}^{T}\prod_{h=1}^{J-1} \frac{\exp\left(\frac{1}{\tilde{\sigma}_{s_{ih}}}(x'_{ity_{ith}}\beta_i)\right)}{\sum_{l=h}^{J}\exp\left(\frac{1}{\tilde{\sigma}_{s_{ih}}}(x'_{ity_{itl}}\beta_i)\right)}, & \text{(HM-ROL)}. \end{cases}$$

The expression to sum over all possible sequences $s_i^* \in \mathcal{S}$ in Equation (11) seems computationally intensive. However, it can be rewritten as a sequential filter which is computationally efficient (Hamilton, 1989), see Equations (12) and (13) in Appendix A. Moreover, the probability of observing a sequence s_i^* is a straightforward function of q and π.

For a general density $f(\beta_i|\theta)$, the integral in the likelihood function in Equation (11) cannot be solved analytically. To approximate the integral, we use Monte Carlo integration. That is, we obtain R draws $\beta_i^{(r)}$ from a distribution with density $f(\beta_i|\theta)$ and approximate the integral by the average of $p(y_i|\beta_i^{(r)},q,\pi,\tilde{\sigma})$ over these R draws, see appendix A for more details. We use scrambled Halton draws to ensure good coverage of $f(\beta_i|\theta)$ (Bhat, 2003, Bhat, 2001, Braaten and Weller, 1979). In case the distribution over β_i is taken to be discrete, the integral over β_i can be written as a sum and the log-likelihood function can be directly maximized without needing Monte Carlo integration.

The use of the sequential filter allows us to directly maximize the (simulated) log-likelihood function without needing to augment the likelihood function with s_{it} (or s_{ih}) to enable an Expectation Maximization (EM) type of algorithm (Dempster et al., 1977, Hamilton, 1990). This direct maximization requires less computations in a single iteration of the optimization than an EM algorithm, and also does not depend on a given draw of s_{it} (or s_{ih}) which may possibly slow down convergence due to the extra iterations needed.

Specialized code is written in C++ and R (R Core Team, 2013, Eddelbuettel and François, 2011) to obtain the scrambled Halton draws and to evaluate the (simulated) log-likelihood function and compute its analytic gradients. The details are given in Appendix A. We take the standard errors equal to the square root of the diagonal elements of the inverse of the negative Hessian of the log-likelihood function. We approximate the Hessian using the outer-product-of-gradients approximation.

The probability that an individual i is in a phase m at time t conditional on observed choices y_i, $\Pr[S_{it} = m|y_i, \theta, q, \pi, \tilde{\sigma}]$, can be computed after the maximum likelihood estimates have been obtained. Details are in Appendix B.

4. Monte Carlo study

In this section, we illustrate the performance of our hidden Markov multinomial logit model with a Monte Carlo study. The study consists of two parts. We first evaluate the small-sample performance of the model and estimator under correct model specification. Next, we evaluate the performance of the model under model misspecification.

For the first part of the study, we consider three data generating processes (DGPs). We use 1,000 Monte Carlo replications per DGP. For each DGP, we consider 1,000 individuals, 15 observations per individual, and 2 alternatives per observation. We consider three explanatory variables: x_{1itj}, x_{2itj} from a standard normal distribution and x_{3itj} from a Bernoulli distribution with probability 0.5 of outcome one. Furthermore, in the DGPs we draw the individual-specific preference parameters from a multivariate normal distribution

$$\beta_i \sim MVN(b, \Sigma_\beta),$$

where Σ_β is a positive definite covariance matrix.

In the first DGP, the HM-MNL is the true model. In this DGP, we aim to mimic the possible learning and fatigue behavior that individuals may experience when completing a survey. We use three phases $\tilde{\sigma} = (\infty, 1, \infty)$. Individuals in the first phase still need to learn (e.g. about their preferences) and answer randomly, individuals in the second phase answer most accurately according to their true preferences (the *minimum variance phase*), and individuals in the third phase answer randomly due to fatigue. We consider initial phase probabilities $\pi = (0.2, 0.7, 0.1)$, and transition probabilities $q_{11t} = 0.50$ and $q_{22t} = 0.99$ for all t. Based on π and q, the percentage of observations in phases one to three are 2.7%, 81.6%, and 15.8%, respectively. Furthermore, 21.5% of individuals reach phase three. At $t = 5$, the percentage of individuals in the minimum variance phase (phase 2) is largest: 85.6%. Finally, we take Σ_β diagonal.

The second and third DGPs are altered versions of the first DGP. In DGP two, the true model is the MNL: we set $\pi_2 = 1$ and $q_{22t} = 1$ for all t. In DGP three, instead of a diagonal covariance matrix as in DGP one, we add correlation across the preference parameters by letting Σ_β be a full positive definite covariance matrix with implied correlations ρ_β. These correlations can capture time-invariant scale heterogeneity: some individuals may choose more randomly throughout the survey than others. Time-invariant scale-heterogeneity shows itself in preference parameters of the same individual to either all tend to more extreme values than b or to all tend to 0. To incorporate time-invariant scale-heterogeneity in the DGP, we let the implied correlations have an absolute level of 0.7. That is, when two b parameters are both positive or both negative we set the correlation to 0.7, when one of them is positive and the other negative we set the correlation to -0.7.

For each replication, we estimate the parameters of three models: (i) a MNL, (ii) a heteroscedastic MNL (H-MNL)[7], and (iii) our HM-MNL. For all three models, we let $\beta_i \sim MVN(b, \Sigma_\beta)$ and use 250 scrambled Halton draws per individual. For the H-MNL, we fix $\sigma_1 = 1$ during estimation and, for each replication, after estimation we scale b and Σ_β such that the lowest scale parameter is equal to one. For the HM-MNL, in estimation we use three phases with known scales $\tilde{\sigma} = (\infty, 1, \infty)$. We restrict the transition probabilities such that individuals can only stay in the current phase or move one phase up and let the transition probabilities be equal over time.

In the second part of the study, we check the robustness of our model to misspecification of the Markov process. We consider three extra DGPs (DGPs four to six) in which there are more phases in the DGP's Markov process to capture more complex forms of heterogeneous heteroscedasticity. The details of the DGPs and the results are in Appendix C.

4.1. Results

The results of the first part of the Monte Carlo study are given in Table 1. We report the mean across replications of the parameter estimates and the corresponding root mean squared error (RMSE) in parentheses. We do this for the three models estimated: the standard (mixed) MNL, the heteroscedastic MNL and the hidden Markov MNL.

In the first DGP, the HM-MNL is the true model with Σ_β diagonal. For this DGP,

[7]The heteroscedastic MNL (Bradley & Daly, 1994) incorporates time-dependent scale parameters $\{\sigma_t\}_{t=1}^T$ that are freely estimated except for one. We set the first scale parameter equal to one: $\sigma_1 = 1$.

Table 1: Mean and RMSE (in parentheses) of the parameter estimates for the Monte Carlo study. Based on 1,000 Monte Carlo replications per DGP.

	DGP 1: HM-MNL ($\rho = 0$)				DGP 2: MNL				DGP 3: HM-MNL ($\rho \neq 0$)			
Parameter	True	MNL	H-MNL	HM-MNL	True	MNL	H-MNL	HM-MNL	True	MNL	H-MNL	HM-MNL
b_1	1.00	0.76	0.94	1.00	1.00	1.00	1.17	1.02	1.00	0.76	0.94	1.00
		(0.24)	(0.09)	(0.06)		(0.03)	(0.18)	(0.03)		(0.24)	(0.08)	(0.06)
b_2	0.30	0.23	0.28	0.30	0.30	0.30	0.35	0.30	0.30	0.22	0.27	0.30
		(0.07)	(0.03)	(0.03)		(0.02)	(0.06)	(0.02)		(0.08)	(0.04)	(0.03)
b_3	-0.50	-0.38	-0.48	-0.50	-0.50	-0.50	-0.58	-0.51	-0.50	-0.37	-0.46	-0.50
		(0.12)	(0.06)	(0.05)		(0.04)	(0.10)	(0.04)		(0.13)	(0.06)	(0.05)
$\sigma_{\beta,1}$	0.50	0.52	0.65	0.49	0.50	0.50	0.59	0.49	0.50	0.52	0.65	0.49
		(0.03)	(0.15)	(0.06)		(0.03)	(0.10)	(0.04)		(0.03)	(0.15)	(0.06)
$\sigma_{\beta,2}$	0.40	0.33	0.41	0.40	0.40	0.40	0.47	0.40	0.40	0.33	0.41	0.40
		(0.07)	(0.04)	(0.03)		(0.03)	(0.08)	(0.03)		(0.07)	(0.04)	(0.03)
$\sigma_{\beta,3}$	0.70	0.59	0.72	0.70	0.70	0.70	0.81	0.70	0.70	0.58	0.71	0.70
		(0.12)	(0.07)	(0.07)		(0.05)	(0.13)	(0.05)		(0.13)	(0.07)	(0.07)
$\rho_{\beta,12}$	0.00	0.21	0.21	-0.02	0.00	0.00	0.00	-0.03	0.70	0.69	0.69	0.71
		(0.22)	(0.22)	(0.15)		(0.08)	(0.08)	(0.10)		(0.07)	(0.07)	(0.10)
$\rho_{\beta,13}$	0.00	-0.20	-0.20	0.02	0.00	0.00	0.00	0.03	-0.70	-0.68	-0.68	-0.71
		(0.21)	(0.22)	(0.14)		(0.09)	(0.09)	(0.10)		(0.08)	(0.08)	(0.10)
$\rho_{\beta,23}$	0.00	-0.10	-0.10	0.00	0.00	-0.01	-0.01	0.00	-0.70	-0.72	-0.72	-0.70
		(0.14)	(0.15)	(0.12)		(0.09)	(0.09)	(0.09)		(0.10)	(0.10)	(0.10)
π_1	0.200			0.204	0.000			0.014	0.200			0.207
				(0.047)				(0.027)				(0.046)
π_2	0.700			0.699	1.000			0.978	0.700			0.696
				(0.055)				(0.035)				(0.054)
π_3	0.100			0.097	0.000			0.008	0.100			0.097
				(0.042)				(0.015)				(0.042)
q_{11}	0.500			0.477	-			0.003	0.500			0.483
				(0.170)								(0.163)
q_{22}	0.990			0.990	1.000			0.999	0.990			0.990
				(0.003)				(0.001)				(0.004)

the standard MNL underestimates the mean of the preference parameters b. This is as expected, as heteroscedasticity leads to more random-looking choice making of individuals. The H-MNL also slightly underestimates the preference parameters. Hence, the bias towards zero, due to 14.4% of individuals choosing randomly at the minimum variance task 5, seems stronger than the bias away from zero, due to scaling back to the minimum variance task. The parameter estimator for the HM-MNL seems to have negligible small sample bias. Hence, the model seems well able to capture and distinguish between the individual-specific preferences and heteroscedasticity.

Interestingly, the MNL and H-MNL spuriously find a positive correlation between b_1 and b_2 ($\rho_{\beta,12} > 0$) and negative correlations $\rho_{\beta,13}$ and $\rho_{\beta,23}$. This implies that individuals with an extremer value for b_1 also tend to have extremer values for b_2 and b_3, and vice versa. These correlations thus try to capture part of the individual-specific time-variation in the scale parameter via (spurious) individual-specific time-invariant correlations.

In the second DGP, the standard MNL is the true model. The estimator for the standard MNL seems to be unbiased. In contrast, the H-MNL clearly overestimates the preference parameters b. This illustrates that the parameter estimator for the H-MNL is biased away from zero due to estimation uncertainty in $\{\sigma_t\}_{t=1}^T$ of which the lowest is used to scale the preference parameters. The estimator for the HM-MNL seems almost unbiased: there is a slight bias away from zero. This is because the model assigns, on average, a small 2.2% fraction of individuals to start in phases one and three. The mean of the estimated probability of staying in the minimum variance phase (q_{22}) is close to the true 1. The RMSEs indicate that the loss in efficiency in estimating the HM-MNL instead of the correctly specified MNL is almost negligible.

Finally, in DGP 3, the HM-MNL is the true model and there is correlation across the preference parameters to allow for time-invariant scale heterogeneity. The estimator for the HM-MNL seems to be unbiased for both the mean of the preference parameters (b) and the covariance matrix (Σ_β), indicating that this model can distinguish between time-invariant scale heterogeneity and time-varying scale heterogeneity.

To summarize, the parameter estimator for the preference parameters in the H-MNL specification seems biased for all three DGPs. Depending on the DGP and whether individual differences exist, the bias can be towards zero or away from zero. Also, the estimator for the standard MNL specification is biased towards zero in case heteroscedasticity is present. The HM-MNL alleviates these biases.

The results of the second part of the study, the performance of the HM-MNL under misspecification of the Markov process, are in Appendix C. In the three DGPs considered, the true Markov process contains more phases than the three used in estimation, and in many of the phases the scale parameter is between one and infinity. Such a Markov process could be more realistic than the process assumed in estimation. We find that the estimators for the three models all underestimate the preference parameters. The bias towards zero is largest for the standard MNL, followed by the H-MNL. The estimator for the HM-MNL is most accurate in estimating the preference parameters. This indicates that our proposed HM-MNL works comparatively well when the true underlying Markov process is more complex than assumed in the model.

5. Case study I: learning and fatigue during discrete choice experiments

In this section, we illustrate our hidden Markov multinomial logit model with data obtained from a discrete choice experiment (DCE). During DCEs, respondents are repeatedly asked to make a hypothetical choice among a set of alternatives, where each alternative is described by a number of attributes (Green, 1974, Louviere and Woodworth, 1983). These experiments are used to elicit the preferences of respondents. The results can be used in product design and in predicting product demand (Rao, 2014). During DCEs, respondents might still need to learn about their preferences or the choice task at hand (Plott, 1993, Braga and Starmer, 2005), or may become tired, bored, or irritated while completing the choice tasks (Lavrakas, 2008). This latter process is known as *fatigue*. Due to learning and fatigue, a respondent may respond more randomly at some tasks. This randomness will lead to unpopular products to be more often selected and, if unaccounted for in the model, overestimation of their potential demand.

The papers that have examined the presence of learning and fatigue during discrete choice experiments have so far only used population-level approaches for the learning and fatigue process.[8] Using different datasets, they find mixed results: some find evidence of learning

[8]The only exception is the individual-level model of Campbell et al. (2015), which is a rather restrictive model. Campbell et al. (2015) *a priori* divide the choice tasks into early (E), middle (M) and late (L) tasks. To model the choices for the three different types of tasks, they specify a latent class model with seven classes. There are three different vectors of preference parameters β_E, β_M, β_L and three scale parameters σ_E, σ_M and σ_L. The first class of the latent class model has constant preferences β_M and constant scale σ_M for early, middle and late tasks (hence, no learning and fatigue). Classes 2 to 4 have a constant σ_M but different combination of β's: class 2 has β_E for early tasks and β_M for the remaining tasks (only learning), class 3 has β_L for late tasks and β_M for the remaining tasks (only fatigue), and class 4 has β_E for early tasks, β_M for middle tasks and β_L for late tasks (learning and fatigue). Classes 5 to 7 have a constant β_M and a

(DeSarbo et al., 2004, Holmes and Boyle, 2005, Czajkowski et al., 2014), some of fatigue (Bradley and Daly, 1994, Koppelman and Sethi, 2005, Savage and Waldman, 2008) and some of neither (Savage and Waldman, 2008, Hess et al., 2012). Because of the population-level approaches, these papers only provide insight into the aggregate scale per choice task, and thus cannot distinguish between different respondents at the same choice task: those that answer accurately, those that need to learn, and those that are fatigued. Hence, findings based on an individual-level model may totally differ.

To examine learning and fatigue during DCEs, we use data obtained from a discrete choice experiment on food choices conducted in the Netherlands (Koç & van Kippersluis, 2017).[9] During the experiment, the respondents had to complete 18 choice tasks. At each choice task, a respondent was asked to choose between two meals: "Which of the two meals would you eat regularly (at least twice a week)?".

The meals were described by the attributes price, taste, cooking time, and health consequences. Each attribute could take on three levels, with a clear ordering between the levels. For example, the price of the meal was either 2 Euro, 6 Euro, or 10 Euro. The respondents were divided into three groups. The groups differed in the attributes and information they obtained during the DCE about the health consequences of the meal. For the first respondent group, the health consequences of the meal were described by one explicit attribute: a meal could either be healthy, health neutral, or unhealthy. For the second and third group, the health consequences were described by three implicit health attributes: number of calories, grams of saturated fat, and grams of sodium. Furthermore, group 2 obtained health information describing what levels of these attributes constitute a healthy meal. Group 3 did not obtain this health information. For an overview of the attributes and corresponding levels per respondent group, see Table 2. The ordering of the tasks within each respondent group were randomized over the respondents and there was no overlap of respondents across groups.

We retain all respondents who filled in at least two choice tasks, also when a respondent dropped out. The three respondent groups contain the responses of 1,206, 1,154 and 1,185 respondents, respectively. In the model, we include the attribute levels as different dummy variables. For each attribute, we take the baseline level to be the first attribute level.

We consider three models: (1) a MNL, (2) a heteroscedastic MNL (H-MNL) (Bradley

similar combination of σ as classes 2 to 4 have for β. This model does not allow for unobserved preference heterogeneity and the timing of learning and fatigue is fixed across respondents by *a priori* dividing the tasks into three sets.

[9] The dataset was obtained from the LISS (Longitudinal Internet Studies for the Social sciences) panel administered by CentERdata (Tilburg University, The Netherlands).

Table 2: Attributes and attribute levels for the three respondent groups. The final column indicates which respondents groups (1,2 or 3) obtained which attributes in the choice experiment.

Attribute	Attribute levels			Groups
Price	2 Euro	6 Euro	10 Euro	1, 2, 3
Cooking time	10 min	30 min	50 min	1, 2, 3
Taste	OK	Good	Very good	1, 2, 3
Health consequence	Unhealthy	Health neutral	Healthy	1
Number of kilocalories	800	1,100	1,400	2, 3
Grams of saturated fat	10	20	30	2, 3
Milligrams of sodium	900	1,200	1,500	2, 3

& Daly, 1994), and (3) our HM-MNL. For all three models, we take a multivariate normal distribution for β_i as given by

$$\beta_i \sim MVN(b, \Sigma_\beta),$$

where Σ_β is a full positive definite covariance matrix. In estimation, we use 250 scrambled Halton draws per respondent and 30 starting values per model.[10]

For the HM-MNL, we consider two specifications. Both specifications have three phases $\tilde{\sigma} = (\infty, 1, \infty)$: respondents in the first phase still need to learn and answer randomly, respondents in in the second phase (the *minimum variance phase*) answer most accurately, and respondents in the third phase answer randomly due to fatigue. We restrict the transition probabilities such that a respondent can either stay in the current phase or move one phase up. For the first specification, we let the transition probabilities be equal over tasks: $q_{11t} = q_{11}$ and $q_{22t} = q_{22}$ for all t. For the second specification, we allow for the transition probabilities from the minimum variance phase to the fatigue phase to be different over tasks.

For the H-MNL, we fix $\sigma_1 = 1$ during estimation. After estimation, we scale b, Σ_β and $\{\sigma_t\}_{t=1}^T$ such that the minimum variance task has variance one, that is, $\min\{\sigma_t\}_{t=1}^T = 1$.

5.1. Results

The results for the first respondent group are shown in Table 3.[11] For this group, the meals were described by four attributes, explicit health information was given in the final attribute 'health consequences'. With the MNL, we find that individuals, on average, positively value

[10] The starting values for b and Σ_β for the H-MNL and HM-MNL are set equal to the maximum likelihood estimates of the MNL.

[11] The detailed results for the covariance matrix Σ_β are available upon request.

a low price and cooking time, a good taste, and a healthy meal.

The H-MNL and HM-MNL find similar patterns as the MNL in the preference parameters, although they are estimated further away from zero. Hence, learning and fatigue seem to both be present. For the population-level H-MNL, this is clearly indicated in the time-variation in the scale parameters. The variance increases at task 2, then decreases, and in the final couple of tasks increases again. The minimum variance task is estimated to be 14. Hence, in the first 14 tasks learning seems more prevalent than fatigue, whereafter fatigue seems more prevalent. The standard errors do imply that there is quite some estimation uncertainty and the minimum variance task could be anywhere from task 7 to 16.

The first HM-MNL, with transition probabilities restricted over tasks, also finds evidence of learning and fatigue. An estimated 17.4% of respondents start in the learning phase in which they reside on average five tasks.[12] Fatigue also occurs: 1.2% of respondents are estimated to answer randomly throughout the survey, and at each task an estimated 0.3% of respondents in the minimum variance phase gets fatigued. Based on the estimated initial and transition probabilities, 5.2% of respondents is fatigued at the final choice task.

The less restrictive second HM-MNL finds that an estimated 17.0% of respondents start in the learning phase and none of the respondents start in the fatigue phase. Furthermore, mainly after the first and second task, respondents seem to become tired. For tasks 3 up to 15, most respondents in the minimum variance phase seem to stay there, and at the final three tasks 16 to 18, more respondents seem to become fatigued. At the final choice task, an estimated 7.2% of respondents is fatigued.

According to the information criteria, both HM-MNL models are preferred over the standard MNL and the H-MNL. The first HM-MNL model seems the most preferred. Even though this HM-MNL has 13 parameters less to estimate than the H-MNL, the likelihood value indicates that it better fits the data. Thus, there seems to be quite some heterogeneity in learning and fatigue across respondents.

The results for the HM-MNL models indicate that at each task, a number of respondents answer randomly. This causes a bias towards zero in the preference parameters b in the H-MNL next to the bias away from zero due to scaling back to the minimum variance task. For this food choice dataset for the first respondent group, the biases seem to almost cancel each other with the bias towards zero seeming just a bit more dominant: the preference parameters b estimated by the H-MNL are slightly closer to zero than those estimated by the HM-MNL.

[12] The average number of tasks that someone who starts in the learning phase will remain in the learning phase is equal to $1/(1 - q_{11})$.

Table 3: Parameter estimates and standard errors (in parentheses) for group 1 from the food choice dataset. Baseline levels are price 2 euro, time 10 minutes, taste OK, and health unhealthy.

	MNL		H-MNL		HM-MNL[a]		HM-MNL[b]	
Price 6 euro	-0.74	(0.04)	-0.96	(0.07)	-0.98	(0.06)	-0.99	(0.06)
Price 10 euro	-2.15	(0.08)	-2.75	(0.18)	-2.76	(0.12)	-2.77	(0.13)
Time 30 min	-0.35	(0.04)	-0.44	(0.04)	-0.46	(0.05)	-0.46	(0.05)
Time 50 min	-1.23	(0.06)	-1.57	(0.11)	-1.59	(0.08)	-1.60	(0.08)
Taste good	0.66	(0.04)	0.83	(0.06)	0.85	(0.05)	0.85	(0.05)
Taste very good	1.18	(0.05)	1.51	(0.10)	1.50	(0.07)	1.51	(0.07)
Health neutral	3.50	(0.09)	4.49	(0.28)	4.55	(0.16)	4.58	(0.17)
Healthy	4.96	(0.13)	6.36	(0.39)	6.43	(0.23)	6.48	(0.24)
π_1					0.174	(0.024)	0.170	(0.024)
π_2					0.814	(0.033)	0.830	(0.039)
π_3					0.012	(0.015)	0.000	(0.025)
q_{11}					0.816	(0.043)	0.796	(0.046)
q_{22}					0.997	(0.001)		
σ_1			1.59	(0.13)				
σ_2 or $q_{22,1}$			1.76	(0.14)			0.990	(0.032)
σ_3 or $q_{22,2}$			1.59	(0.13)			0.968	(0.023)
σ_4 or $q_{22,3}$			1.64	(0.15)			1.000	(0.027)
σ_5 or $q_{22,4}$			1.33	(0.11)			1.000	(0.026)
σ_6 or $q_{22,5}$			1.51	(0.14)			1.000	(0.019)
σ_7 or $q_{22,6}$			1.16	(0.11)			1.000	(0.018)
σ_8 or $q_{22,7}$			1.21	(0.12)			1.000	(0.024)
σ_9 or $q_{22,8}$			1.14	(0.11)			1.000	(0.027)
σ_{10} or $q_{22,9}$			1.01	(0.10)			1.000	(0.018)
σ_{22} or $q_{22,10}$			1.21	(0.11)			0.990	(0.012)
σ_{12} or $q_{22,11}$			1.13	(0.11)			1.000	(0.014)
σ_{13} or $q_{22,12}$			1.22	(0.13)			1.000	(0.014)
σ_{14} or $q_{22,13}$			1.00	-			1.000	(0.019)
σ_{15} or $q_{22,14}$			1.10	(0.11)			1.000	(0.022)
σ_{16} or $q_{22,15}$			1.20	(0.11)			0.985	(0.015)
σ_{17} or $q_{22,16}$			1.25	(0.12)			1.000	(0.021)
σ_{18} or $q_{22,17}$			1.24	(0.11)			0.987	(0.021)
average σ_β	1.5		1.9		1.9		1.9	
# free parameters	44		61		48		64	
log-likelihood	-10,560		-10,524		-10,483		-10,477	
BIC	21,559		21,658		21,446		21,593	
AIC3	21,252		21,232		21,111		21,146	
AIC2	21,208		21,171		21,063		21,082	

[a] HM-MNL with 3 phases $\tilde{\sigma} = (\infty, 1, \infty)$: equal transition probabilities over tasks.
[b] HM-MNL with 3 phases $\tilde{\sigma} = (\infty, 1, \infty)$: transition probability to fatigue different per task.

Table 4: Parameter estimates and standard errors (in parentheses) for groups 2 and 3 from the food choice dataset. Baseline levels are price 2 euro, time 10 minutes, taste OK, calories 800, saturated fat 10 gram, and sodium 900 mg.

	Food 2								Food 3							
	MNL		H-MNL		HM-MNL[a]		HM-MNL[b]		MNL		H-MNL		HM-MNL[a]		HM-MNL[b]	
Price 6 euro	-0.61	(0.04)	-0.86	(0.06)	-0.91	(0.06)	-0.92	(0.06)	-0.71	(0.04)	-1.03	(0.07)	-1.03	(0.07)	-1.06	(0.07)
Price 10 euro	-1.54	(0.08)	-2.16	(0.13)	-2.19	(0.13)	-2.21	(0.13)	-1.82	(0.09)	-2.64	(0.16)	-2.68	(0.14)	-2.61	(0.15)
Time 30 min	-0.08	(0.04)	-0.13	(0.03)	-0.15	(0.05)	-0.16	(0.05)	-0.09	(0.04)	-0.14	(0.03)	-0.12	(0.05)	-0.15	(0.06)
Time 50 min	-0.65	(0.06)	-0.93	(0.07)	-0.92	(0.09)	-0.93	(0.09)	-0.73	(0.07)	-1.08	(0.08)	-1.09	(0.09)	-1.04	(0.10)
Taste good	0.69	(0.04)	0.99	(0.06)	0.98	(0.06)	0.98	(0.06)	0.89	(0.05)	1.30	(0.08)	1.39	(0.08)	1.40	(0.08)
Taste very good	1.19	(0.06)	1.69	(0.10)	1.64	(0.09)	1.65	(0.09)	1.40	(0.07)	2.08	(0.12)	2.19	(0.12)	2.26	(0.12)
Calories 1100	-0.81	(0.04)	-1.10	(0.07)	-1.11	(0.07)	-1.11	(0.07)	-0.58	(0.04)	-0.85	(0.06)	-0.81	(0.06)	-0.92	(0.07)
Calories 1400	-1.59	(0.06)	-2.20	(0.13)	-2.19	(0.11)	-2.20	(0.11)	-1.29	(0.06)	-1.89	(0.11)	-1.86	(0.10)	-2.04	(0.11)
Sat fat 20 gram	-0.66	(0.04)	-0.94	(0.06)	-0.94	(0.07)	-0.95	(0.07)	-0.43	(0.04)	-0.61	(0.04)	-0.60	(0.06)	-0.68	(0.06)
Sat fat 30 gram	-1.28	(0.06)	-1.78	(0.11)	-1.79	(0.09)	-1.81	(0.10)	-0.91	(0.05)	-1.31	(0.08)	-1.30	(0.08)	-1.44	(0.09)
Sodium 1200mg	-0.37	(0.04)	-0.52	(0.05)	-0.55	(0.06)	-0.55	(0.06)	-0.33	(0.04)	-0.48	(0.04)	-0.52	(0.06)	-0.55	(0.07)
Sodium 1500mg	-0.87	(0.05)	-1.21	(0.08)	-1.22	(0.08)	-1.23	(0.08)	-0.72	(0.05)	-1.04	(0.07)	-1.07	(0.07)	-1.11	(0.08)
π_1					0.218	(0.030)	0.217	(0.030)					0.273	(0.033)	0.269	(0.032)
π_2					0.765	(0.036)	0.782	(0.042)					0.713	(0.035)	0.706	(0.047)
π_3					0.017	(0.021)	0.002	(0.030)					0.014	(0.027)	0.025	(0.045)
q_{11}					0.816	(0.042)	0.802	(0.044)					0.838	(0.036)	0.831	(0.037)
q_{22}					0.996	(0.001)							0.995	(0.002)		
σ_1			1.80	(0.14)							1.90	(0.14)				
σ_2 or $q_{22,1}$			1.96	(0.16)			0.973	(0.040)			2.13	(0.17)			0.993	(0.064)
σ_3 or $q_{22,2}$			1.85	(0.15)			0.983	(0.035)			1.85	(0.14)			0.990	(0.052)
σ_4 or $q_{22,3}$			1.60	(0.15)			0.999	(0.037)			1.88	(0.16)			0.998	(0.046)
σ_5 or $q_{22,4}$			1.21	(0.10)			0.999	(0.035)			1.40	(0.12)			0.992	(0.030)
σ_6 or $q_{22,5}$			1.57	(0.13)			0.998	(0.034)			1.17	(0.10)			0.995	(0.033)
σ_7 or $q_{22,6}$			1.48	(0.12)			0.982	(0.024)			1.70	(0.14)			0.996	(0.029)
σ_8 or $q_{22,7}$			1.30	(0.11)			0.999	(0.028)			1.63	(0.13)			0.997	(0.029)
σ_9 or $q_{22,8}$			1.35	(0.11)			0.999	(0.026)			1.44	(0.13)			0.999	(0.028)
σ_{10} or $q_{22,9}$			1.16	(0.10)			0.997	(0.025)			1.29	(0.12)			0.987	(0.026)
σ_{22} or $q_{22,10}$			1.51	(0.13)			0.996	(0.026)			1.30	(0.12)			0.990	(0.028)
σ_{12} or $q_{22,11}$			1.32	(0.12)			0.999	(0.021)			1.17	(0.10)			0.990	(0.022)
σ_{13} or $q_{22,12}$			1.19	(0.11)			0.999	(0.026)			1.25	(0.11)			1.000	(0.022)
σ_{14} or $q_{22,13}$			1.00	-			0.995	(0.023)			1.26	(0.11)			0.999	(0.031)
σ_{15} or $q_{22,14}$			1.40	(0.12)			0.983	(0.022)			1.13	(0.09)			0.982	(0.022)
σ_{16} or $q_{22,15}$			1.10	(0.11)			0.998	(0.025)			1.50	(0.12)			0.993	(0.020)
σ_{17} or $q_{22,16}$			1.22	(0.11)			0.999	(0.025)			1.33	(0.12)			0.999	(0.018)
σ_{18} or $q_{22,17}$			1.18	(0.12)			0.999	(0.039)			1.00	-			1.000	(0.033)
average σ_β	1.1		1.5		1.4		1.4		1.1		1.6		1.5		1.6	
# free parameters	90		107		94		110		90		107		94		110	
log-likelihood	-10,989		-10,946		-10,900		-10,897		-11,474		-11,433		-11,360		-11,354	
BIC	22,872		22,955		22,734		22,887		23,844		23,933		23,656		23,804	
AIC3	22,248		22,212		22,082		22,124		23,217		23,188		23,001		23,037	
AIC2	22,158		22,105		21,988		22,014		23,127		23,081		22,907		22,927	

[a] HM-MNL with 3 phases $\tilde{\sigma} = (\infty, 1, \infty)$: equal transition probabilities over tasks.
[b] HM-MNL with 3 phases $\tilde{\sigma} = (\infty, 1, \infty)$: transition probability to fatigue different per task.

The results for the second and third respondent groups are in Table 4. For these two groups, health consequences were described by three attributes: (i) number of calories, (ii) amount of saturated fat, and (iii) amount of sodium. For group two, health information on the attributes was provided in the text, for group three no health information was provided.

The HM-MNL models find that as the amount of information decreases from group one to three, the percentage of respondents that start in the learning phase increases from 17% to 22% to 27%. Moreover, according to the first HM-MNL specification, the probability to become fatigued once in the minimum variance phase increases from group one to three, from 0.3% to 0.4% to 0.5% per choice task, although there is some uncertainty in these estimates. Hence, summarizing health information in one attribute or providing health information in the text seems to reduce the need for learning and the risk of fatigue.

With the H-MNL we also find evidence of learning for respondent groups two and three. Remarkably, for all three groups, we find an initial increase in the variance from task 1 to 2, after which the variance decreases. This is not due to the order of the choice tasks as they are randomized over respondents. Hence, it seems that there is a relatively large group of respondents who try to answer accurately at task 1, but from task 2 onwards do not so anymore. This behavior can be seen more explicitly by the second HM-MNL specification, where we find that there is relatively large group of respondents that move to the fatigue phase after tasks one and two. Combined with the reduction of respondents that need to learn in these two tasks, the variation in the scale parameters in the H-MNL can be explained.

In summary, the hidden Markov MNLs provide the best fit of the data for all three respondent groups. Moreover, these models provide interesting and plausible insights into the presence of learning and fatigue during the discrete choice experiment.

6. Case study II: differential capabilities in ranking

In this section, we illustrate our hidden Markov rank-ordered logit model with rankings obtained from a survey on cultural opinions conducted in the Netherlands (Sociaal en Cultureel Planbureau, 2004, Fok et al., 2012). One of the questions in the survey asked the respondents to rank 16 political goals from most to least desired. In total, 2,261 individuals aged sixteen years and older completed the ranking. The initial presented ordering of the political goals in the survey was randomized over respondents.

We estimate and compare three models: (1) a ROL, (2) a heteroscedastic ROL (H-ROL) (Hausman & Ruud, 1987) and (3) our HM-ROL. For all models, we take a multivariate normal distribution for β_i as given by

$$\beta_i \sim MVN(b, \Sigma_\beta),$$

with full positive definite covariance matrix Σ_β.[13] In estimation, we use 30 starting values per model and 250 scrambled Halton draws per respondent. For the H-ROL, we fix $\sigma_1 = 1$.

For the HM-ROL, we consider three specifications. The first specification is equivalent to the latent class ROL in Fok et al. (2012), with the addition of allowing for individual-specific preference parameters. That is, we use two phases $\tilde{\sigma} = (1, \infty)$ and restrict the transition probabilities such that, between consecutive ranks, a respondent can only move from phase 1 to phase 2 and once in phase 2 stays there. In other words, a respondent can assign all ranks accurately (all choices based on phase 1), all ranks randomly (all choices based on phase 2), or the top j ranks accurately and the bottom $J - j$ randomly (in phase 1 for rank one until j, in phase 2 for ranks $j + 1$ and higher) for any j.

For the second specification, we also allow for bottom ranks to be assigned accurately. We use three phases $\tilde{\sigma} = (1, \infty, 1)$ and restrict the transition probabilities such that, between consecutive ranks, a respondent can only move from phase 1 to phase 2 or from phase 2 to phase 3, and once in phase 3 stays there. Also, a respondent is restricted to start (assign the top rank) in either phase 1 or 2. In the third specification, we include a middle phase with scale to be estimated: $\tilde{\sigma} = (1, \tilde{\sigma}_2, \infty, 1)$ to allow for a decrease in the ability of respondents to assign lower ranks. We restrict the transition and initial phase probabilities such that a respondent can only move one phase up and can only start in the first and third phase. For further parsimony, we restrict the transition probabilities to be equal to each other over ranks h, except for moving from the first to the second phase in the first two tasks (assign top ranks accurately) and from the third to the final phase in the last two tasks (assign bottom ranks accurately).

6.1. Results

The results for the political preferences ranking data are given in Table 5. With the standard ROL we find that individuals seem to attach most value to goals as 'maintain order', 'stable economy', 'fight crime', 'freedom of speech', and 'social security'. The estimated correlations across the individual-specific preference parameters in Σ_β indicate which goals are often

[13] For the political preferences ranking data, we have only one observation per individual and quite a number of free parameters in the (15×15) matrix Σ_β. Therefore, using a low-rank approximation of Σ_β or a Bayesian approach with informative priors might be useful to reduce the risk of overfitting.

ranked close by (results not shown).[14] We find that this holds most strongly for (i) 'maintain order' and 'fight crime', (ii) 'more say politics' and 'more say community', (iii) 'economic growth' and 'stable economy', (iv) 'defence forces' and 'cities and countryside', and (v) 'humane society' and 'ideas > money'.

Table 5: Parameter estimates and standard errors (in parentheses) for the political preferences ranking dataset. Baseline level is 'take good care of immigrants'.

	ROL		H-ROL		HM-ROL[a]		HM-ROL[b]		HM-ROL[c]	
maintain order	2.45	(0.05)	4.04	(0.10)	2.82	(0.07)	3.03	(0.08)	3.98	(0.15)
more say politics	1.20	(0.05)	2.32	(0.06)	1.51	(0.06)	1.67	(0.07)	2.29	(0.12)
fight rising prices	1.50	(0.05)	2.82	(0.08)	1.84	(0.07)	2.02	(0.08)	2.78	(0.13)
freedom of speech	2.31	(0.04)	3.84	(0.10)	2.64	(0.06)	2.86	(0.07)	3.77	(0.15)
economic growth	1.44	(0.05)	2.66	(0.07)	1.74	(0.07)	1.92	(0.08)	2.62	(0.13)
defence forces	-0.77	(0.05)	-2.22	(0.07)	-1.39	(0.11)	-1.74	(0.17)	-1.86	(0.15)
more say community	1.14	(0.05)	2.21	(0.06)	1.43	(0.06)	1.60	(0.07)	2.18	(0.12)
cities and countryside	0.17	(0.04)	0.44	(0.03)	0.29	(0.06)	0.36	(0.07)	0.46	(0.09)
stable economy	2.40	(0.05)	3.97	(0.11)	2.75	(0.06)	2.96	(0.08)	3.92	(0.15)
fight crime	2.40	(0.05)	3.99	(0.11)	2.77	(0.07)	2.99	(0.08)	3.92	(0.15)
humane society	1.67	(0.05)	3.04	(0.08)	1.97	(0.06)	2.17	(0.07)	2.99	(0.13)
ideas > money	0.91	(0.04)	1.82	(0.06)	1.15	(0.06)	1.31	(0.07)	1.81	(0.11)
fight unemployment	2.10	(0.05)	3.64	(0.10)	2.47	(0.06)	2.68	(0.08)	3.59	(0.15)
fight pollution	1.05	(0.04)	2.13	(0.06)	1.33	(0.05)	1.50	(0.07)	2.07	(0.11)
social security	2.31	(0.05)	3.85	(0.10)	2.65	(0.06)	2.88	(0.08)	3.79	(0.15)
$\tilde{\sigma}_2$									2.29	(0.17)
π_1					1.00	(0.02)	1.00	(0.02)	0.98	(0.02)
π_2					0.00	(0.01)	0.00	(0.01)		
π_3									0.02	(0.01)
π_4										
σ_1			1.00	-						
σ_2 or $q_{11,1}$			1.10	(0.04)	0.99	(0.01)	0.98	(0.01)		
σ_3 or $q_{11,2}$			1.23	(0.04)	1.00	(0.01)	0.99	(0.02)		
σ_4 or $q_{11,3}$			1.21	(0.04)	1.00	(0.01)	1.00	(0.01)		
σ_5 or $q_{11,4}$			1.36	(0.04)	0.99	(0.01)	0.98	(0.01)		
σ_6 or $q_{11,5}$			1.58	(0.05)	0.99	(0.01)	0.97	(0.02)		
σ_7 or $q_{11,6}$			1.63	(0.05)	0.98	(0.01)	0.96	(0.02)		
σ_8 or $q_{11,7}$			1.86	(0.05)	0.99	(0.02)	0.94	(0.03)		
σ_9 or $q_{11,8}$			1.90	(0.06)	0.98	(0.02)	0.92	(0.03)		
σ_{10} or $q_{11,9}$			1.99	(0.06)	0.94	(0.03)	0.89	(0.03)		
σ_{11} or $q_{11,10}$			2.44	(0.07)	0.96	(0.03)	0.91	(0.05)		
σ_{12} or $q_{11,11}$			2.78	(0.08)	0.94	(0.03)	0.83	(0.05)		
σ_{13} or $q_{11,12}$			2.59	(0.08)	0.97	(0.04)	0.94	(0.06)		
σ_{14} or $q_{11,13}$			3.27	(0.10)	0.87	(0.04)	0.78	(0.07)		
σ_{15} or $q_{11,14}$			4.54	(0.15)	0.77	(0.05)	0.61	(0.11)		
$q_{22,1}$							0.96	(5.63)		
$q_{22,2}$							0.96	(0.68)		

[14]The detailed results for Σ_β are available upon request.

Table 5: (Continued)

	ROL		H-ROL		HM-ROL[a]		HM-ROL[b]		HM-ROL[c]	
$q_{22,3}$							0.68	(0.28)		
$q_{22,4}$							0.78	(0.41)		
$q_{22,5}$							0.98	(0.35)		
$q_{22,6}$							0.96	(0.26)		
$q_{22,7}$							0.74	(0.19)		
$q_{22,8}$							0.41	(0.15)		
$q_{22,9}$							0.61	(0.17)		
$q_{22,10}$							1.00	(0.18)		
$q_{22,11}$							0.83	(0.12)		
$q_{22,12}$							0.72	(0.10)		
$q_{22,13}$							1.00	(0.15)		
$q_{22,14}$							0.99	(0.15)		
$q_{11,1}$									0.90	(0.05)
$q_{11,2}$									1.00	(0.06)
$q_{11,3:14}$									0.83	(0.02)
q_{22}									0.95	(0.01)
$q_{33,1:12}$									0.87	(0.04)
$q_{33,13}$									1.00	(0.15)
$q_{33,14}$									1.00	(0.16)
average σ_β	1.3		2.5		1.6		1.7		2.3	
# free parameters	135		149		150		164		144	
log-likelihood	-62,369		-62,193		-62,207		-62,172		-62,152	
BIC	125,781		125,537		125,573		125,610		125,417	
AIC3	125,144		124,833		124,865		124,835		124,736	
AIC2	125,009		124,684		124,715		124,671		124,592	

[a]HM-ROL with 2 phases $\tilde{\sigma} = (1, \infty)$: accurately assign top ranks.

[b]HM-ROL with 3 phases $\tilde{\sigma} = (1, \infty, 1)$: accurately assign top and bottom ranks.

[c]HM-ROL with 4 phases $\tilde{\sigma} = (1, \tilde{\sigma}_2, \infty, 1)$: accurately assign top and bottom ranks + decrease in accuracy.

We next consider the H-ROL which accounts for the behavior that individuals cannot rank all alternatives accurately in a homogeneous way across individuals. The estimates for the mean preference parameters b are further away from zero than for the ROL, and the scale parameters σ_h show a gradual increase as the rank h increases. Hence, individuals seem to be unable to assign all ranks accurately, rendering the estimator for the standard ROL biased. Moreover, individuals seem to most accurately assign the top ranks, followed by the middle and then the bottom ranks. The preference ordering of the political goals stays roughly the same.

The HM-ROL allows for individual differences in the ranking capabilities. The first HM-ROL specification is equivalent to the latent class ROL in Fok et al. (2012), with the addition of allowing for preference heterogeneity. We find that all individuals are able to rank the first alternative accurately ($\pi_1 = 1.0$). The probabilities of staying in the minimum variance phase 1 for consecutive ranks (q_{11h}) are mostly smaller than one, indicating that quite some individuals find it rather difficult to assign the middle and bottom ranks. Moreover, there seem to be individual differences in the number of top ranks that can be assigned accurately. The probabilities of staying in the minimum variance phase one are especially low for the final couple of ranks. This indicates that a large proportion of respondent who are able to accurately assign ranks 1 to 9, have more trouble assigning the lower ranks. These findings mostly agree with the findings Fok et al. (2012), with the exception that they find that 4% of individuals cannot rank the first alternative accurately. This difference suggests that it is important to allow for preference heterogeneity when allowing for differential capabilities in ranking.

The second HM-ROL specification also allows for bottom ranks to be assigned accurately. The probability of staying in the first phase (q_{11h}) are close to one for the first six ranks, and quite a bit lower for the subsequent ranks. This indicates that quite some individuals can accurately assign ranks one to six, but find it more difficult to assign middle ranks from rank seven onwards. The probability of staying in the second phase (q_{22h}) differ quite a bit over ranks. These probabilities are often quite low, indicating that indeed some individuals are able to assign the bottom ranks accurately. Because of the uncertainty in these estimates, it might be sensible to add restrictions to the transition probabilities. For example, one can impose them equal across certain (middle) ranks.

In the third HM-ROL specification, we allow for a decrease in accuracy for assigning alternatives for consecutive ranks, as well as for bottom ranks to be assigned accurately. The estimated scale parameter for the second phase is equal to 2.3. We find that 10% of respondent in the minimum variance phase move to the second phase after rank one, 0% after rank two, and 17% of respondents after each remaining rank. Hence, the information content in the ranks assigned seems to highly differ across respondents.

The main difference between the three HM-ROL specifications is that the final specification allows for a decrease in accuracy in the alternatives assigned to consecutive ranks, whereas the first two specifications assume that an individual either completely accurately assigns a rank or completely randomly. According to the three information criteria, the third HM-ROL specification should be preferred. Hence, for this ranking dataset, it seems more likely that individuals do not completely randomly assign middle ranks, but that the choice is more random compared to top ranks.

When comparing all five models, the information criteria indicate that the third HM-

ROL specification should be most preferred. Even though this model has five parameters less to estimate than the H-ROL, the likelihood value indicates that it much better fits the data. The standard ROL should be the least preferred model, followed by the first HM-ROL specification. The H-ROL and second HM-ROL specification are at a shared second place.

7. Conclusion

The heteroscedastic logit model is useful to describe repeated choices of individuals when randomness in the choice-making varies over time. For example, due to fatigue, individuals may respond more randomly to survey questions as the survey progresses. Or when asked to give a complete ranking amongst multiple alternatives, individuals may more accurately assign top ranks than middle and bottom ranks.

In this paper, we generalize the standard heteroscedastic logit model to allow for individual differences in the dynamics in this randomness. In case individual differences exist, this individual-level approach has three main advantages: (i) it alleviates biases in the preference parameters, (ii) makes more efficient use of data, and (iii) allows for an analysis of individual behavior. The generalization amounts to adding an individual- and time/rank-specific scale parameter to the multinomial and rank-ordered logit model. We let the dynamics in the sequence of an individual's scale parameters be governed by a Markov process. Additionally, we allow for unobserved preference heterogeneity. For inference, we develop a simulated maximum likelihood estimation approach.

In a Monte Carlo study, we find that our proposed model works well and the proposed estimator seems unbiased in various settings. For the standard heteroscedastic logit model, the biases in the estimator for the preference parameters are clearly illustrated: the bias towards zero due to neglecting individual differences in the dynamics in the scale parameter, and the bias away from zero due to scaling the preference parameters based on the minimum of the estimated scale parameters. Depending on the data generating process, one of these biases may dominate the other. In case of heteroscedasticity, the estimator for the preference parameters of the standard MNL is clearly biased towards zero, because heteroscedasticity leads to more random-looking choice-making of respondents. Our proposed model and estimator eliminate these biases. Furthermore, when allowing for preference heterogeneity via a multivariate normal distribution, both the standard MNL and the heteroscedastic MNL tend to spuriously capture individual differences in the dynamics in the scale parameter in time-invariant correlations between preference parameters.

We also illustrate our model with two empirical applications: one using multinomial choice data from a discrete choice experiment on food choices to model learning and fatigue effects, and one on rank-ordered data from a survey to model differential capabilities in ranking. For the multinomial choices, we find that accounting for individual differences in learning and fatigue leads to a much better fit of the data, while needing less model parameters. The same holds for the rank-ordered data.

Our approach has one main limitation: each variable gets scaled with the same factor. Hence, the model cannot capture choice strategies where choices are made based on different subsets of attributes as time progresses, or where preferences change over time. The model could be extended to allow for a different scale parameter per variable, for example, by letting each scale parameter be governed by its own Markov process. However, for datasets with limited information per individual, such an approach would be susceptible to overfitting and estimation uncertainty can become problematic.

We provide three venues for future research. First, in case one wants to impose restrictions on the minimum number of tasks an individual should be in a phase, one can use a second- or higher-order Markov process. By using suitable restrictions on the transition probabilities, no extra parameters need to be estimated. Of course, if desired, one can also allow the transition probabilities to depend on the duration in a phase using such a higher-order Markov process. Second, for rank-ordered data, the Markov process over time and over ranks can be combined, to simultaneously allow for learning and fatigue and for differential capabilities in ranking.

Third, for the applications, we recommend to use more flexible forms of preference heterogeneity then the used multivariate normal. This especially holds when one wants to include scale parameters between one and infinity. A mixture of multivariate normal distributions might be able to capture more realistically the differences in preferences across individuals. One way in which individuals may differ is that some individuals may answer more randomly throughout the observed period than others, also known as time-invariant scale heterogeneity. In the multivariate normal, such behavior is partly captured in the correlations in the covariance matrix. Using a more flexible form than one multivariate normal could further reduce the way in which the Markov process can capture part of the time-invariant scale heterogeneity in case one of the scale parameters is allowed to be between one and infinity.

References

Beggs, S., Cardell, S., & Hausman, J. (1981). Assessing the potential demand for electric cars. *Journal of Econometrics*, *17*(1), 1–19.

Bhat, C. R. (2001). Quasi-random maximum simulated likelihood estimation of the mixed multinomial logit model. *Transportation Research Part B: Methodological*, *35*(7), 677–693.

Bhat, C. R. (2003). Simulation estimation of mixed discrete choice models using randomized and scrambled Halton sequences. *Transportation Research Part B: Methodological*, *37*(9), 837–855.

Bhat, C. R., & Castelar, S. (2002). A unified mixed logit framework for modeling revealed and stated preferences: Formulation and application to congestion pricing analysis in the San Francisco Bay area. *Transportation Research Part B: Methodological*, *36*(7), 593–616.

Bhat, C. R., & Sidharthan, R. (2011). A simulation evaluation of the maximum approximate composite marginal likelihood (MACML) estimator for mixed multinomial probit models. *Transportation Research Part B: Methodological*, *45*(7), 940–953.

Braaten, E., & Weller, G. (1979). An improved low-discrepancy sequence for multidimensional quasi-Monte Carlo integration. *Journal of Computational Physics*, *33*(2), 249–258.

Bradley, M., & Daly, A. (1994). Use of the logit scaling approach to test for rank-order and fatigue effects in stated preference data. *Transportation*, *21*(2), 167–184.

Braga, J., & Starmer, C. (2005). Preference anomalies, preference elicitation and the discovered preference hypothesis. *Environmental and Resource Economics*, *32*(1), 55–89.

Campbell, D., Boeri, M., Doherty, E., & Hutchinson, W. G. (2015). Learning, fatigue and preference formation in discrete choice experiments. *Journal of Economic Behavior & Organization*, *119*, 345–363.

Chapman, R. G., & Staelin, R. (1982). Exploiting rank ordered choice set data within the stochastic utility model. *Journal of Marketing Research*, *19*(3), 288–301.

Chib, S. (1998). Estimation and comparison of multiple change-point models. *Journal of Econometrics*, *86*(2), 221–241.

Czajkowski, M., Giergiczny, M., & Greene, W. H. (2014). Learning and fatigue effects revisited: Investigating the effects of accounting for unobservable preference and scale heterogeneity. *Land Economics*, *90*(2), 324–351.

Danaf, M., Atasoy, B., & Ben-Akiva, M. (2020). Logit mixture with inter and intra-consumer heterogeneity and flexible mixing distributions. *Journal of Choice Modelling*, *35*, 100188.

Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, *39*(1), 1–22.

DeSarbo, W. S., Lehmann, D. R., & Hollman, F. G. (2004). Modeling dynamic effects in repeated-measures experiments involving preference/choice: An illustration involving stated preference analysis. *Applied Psychological Measurement*, *28*(3), 186–209.

Eddelbuettel, D., & François, R. (2011). Rcpp: Seamless R and C++ integration. *Journal of Statistical Software*, *40*(8), 1–18. https://doi.org/10.18637/jss.v040.i08

Fiebig, D. G., Keane, M. P., Louviere, J., & Wasi, N. (2010). The generalized multinomial logit model: Accounting for scale and coefficient heterogeneity. *Marketing Science*, *29*(3), 393–421.

Fok, D., Paap, R., & Van Dijk, B. (2012). A rank-ordered logit model with unobserved heterogeneity in ranking capabilities. *Journal of Applied Econometrics*, *27*(5), 831–846.

Goldfeld, S. M., & Quandt, R. E. (1973). A Markov model for switching regressions. *Journal of Econometrics*, *1*(1), 3–15.

Green, P. E. (1974). On the design of choice experiments involving multifactor alternatives. *Journal of Consumer Research*, *1*(2), 61–68.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the Econometric Society*, 357–384.

Hamilton, J. D. (1990). Analysis of time series subject to changes in regime. *Journal of Econometrics*, *45*(1-2), 39–70.

Hausman, J. A., & Ruud, P. A. (1987). Specifying and testing econometric models for rank-ordered data. *Journal of econometrics*, *34*(1-2), 83–104.

Hess, S., Hensher, D. A., & Daly, A. (2012). Not bored yet–revisiting respondent fatigue in stated choice experiments. *Transportation Research Part A: Policy and Practice*, *46*(3), 626–644.

Hess, S., & Rose, J. M. (2009). Allowing for intra-respondent variations in coefficients estimated on repeated choice data. *Transportation Research Part B: Methodological*, *43*(6), 708–719.

Hess, S., & Rose, J. M. (2012). Can scale and coefficient heterogeneity be separated in random coefficients models? *Transportation*, *39*(6), 1225–1239.

Holmes, T. P., & Boyle, K. J. (2005). Dynamic learning and context-dependence in sequential, attribute-based, stated-preference valuation questions. *Land Economics*, *81*(1), 114–126.

Kim, C.-J. (1994). Dynamic linear models with Markov-switching. *Journal of Econometrics*, *60*(1-2), 1–22.

Koç, H., & van Kippersluis, H. (2017). Thought for food: Nutritional information and educational disparities in diet. *Journal of Human Capital*, *11*(4), 508–552.

Koppelman, F. S., & Sethi, V. (2005). Incorporating variance and covariance heterogeneity in the generalized nested logit model: An application to modeling long distance travel choice behavior. *Transportation Research Part B: Methodological*, *39*(9), 825–853.

Lavrakas, P. J. (2008). *Encyclopedia of survey research methods*. Sage Publications.

Louviere, J. J., & Woodworth, G. (1983). Design and analysis of simulated consumer choice or allocation experiments: An approach based on aggregate data. *Journal of Marketing Research*, *20*(4), 350–367.

Manski, C. F. (1977). The structure of random utility models. *Theory and Decision*, *8*(3), 229–254.

McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in Econometrics* (pp. 105–142). Academic Press: New York.

Plott, C. R. (1993). Rational individual behavior in markets and social choice processes.

R Core Team. (2013). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria. http://www.R-project.org/

Rao, V. R. (2014). *Applied conjoint analysis*. Springer.

Savage, S. J., & Waldman, D. M. (2008). Learning and fatigue during choice experiments: A comparison of online and mail survey modes. *Journal of Applied Econometrics*, *23*(3), 351–371.

Sociaal en Cultureel Planbureau. (2004). Culturele veranderingen in Nederland 2004 - CV'04 [data retrieved from DANS, https://doi.org/10.17026/dans-xtu-d36b].

A. Maximum simulated likelihood estimation

We use maximum simulated likelihood estimation to estimate the parameters of the hidden Markov multinomial and rank-ordered logit model. For this purpose, we maximize the (approximated) likelihood function directly with respect to θ, q, π, and $\tilde{\sigma}$ using a quasi Newton-Raphson algorithm. During optimization, we use analytic gradients of the simulated log-likelihood function and approximate the Hessian with the BFGS algorithm. In this appendix, we provide more details for the estimation approaches including the explicit likelihood functions for multinomial choices in Section A.1 and for rank-ordered choices in Section A.2.

A.1. Hidden Markov multinomial logit model

The likelihood function of the HM-MNL can be written as

$$\begin{aligned}
p(y|\theta, q, \pi, \tilde{\sigma}) &= \prod_{i=1}^{N} \left[\int \left(\sum_{s_i^* \in \mathcal{S}} \Pr[S_i = s_i^* | q, \pi] p(y_i|\beta_i, \tilde{\sigma}, s_i^*) \right) f(\beta_i|\theta) d\beta_i \right] \\
&= \prod_{i=1}^{N} \left[\int \left\{ \sum_{s_i^* \in \mathcal{S}} \left(\Pr[S_i = s_i^* | q, \pi] \prod_{t=1}^{T} \frac{\exp\left(\frac{1}{\tilde{\sigma}_{s_{it}}}(x_{itj}'\beta_i)\right)}{\sum_{l=1}^{J} \exp\left(\frac{1}{\tilde{\sigma}_{s_{it}}}(x_{itl}'\beta_i)\right)} \right) \right\} f(\beta_i|\theta) d\beta_i \right] \\
&= \prod_{i=1}^{N} \left[\int \left\{ \pi' f_{i1} \left(\prod_{t=2}^{T-1} Q_{t-1} f_{it} \right) Q_{T-1} \tilde{f}_{iT} \right\} f(\beta_i|\theta) d\beta_i \right], \qquad (12)
\end{aligned}$$

where Q_t is a $(M \times M)$ transition probability matrix with element (m, n) equal to q_{mnt}, and $\tilde{f}_{it}$ is a $(M \times 1)$ vector with the likelihood contribution of task t of respondent i given β_i and s_{it}, with element m equal to

$$\tilde{f}_{itm} \equiv \Pr[Y_{it} = y_{it}|\beta_i, \tilde{\sigma}, s_{it} = m] = \frac{\exp\left(\frac{1}{\tilde{\sigma}_m}(x_{itj}'\beta_i)\right)}{\sum_{l=1}^{J} \exp\left(\frac{1}{\tilde{\sigma}_m}(x_{itl}'\beta_i)\right)}.$$

Furthermore, f_{it} is a diagonal $(M \times M)$ matrix with the diagonal equal to $\tilde{f}_{it}$.

We approximate the likelihood function using Monte Carlo integration:

$$\begin{aligned}
p(y|\theta, q, \pi, \tilde{\sigma}) &\approx \prod_{i=1}^{N} \left[\frac{1}{R} \sum_{r=1}^{R} p(y_i|\beta_i^{(r)}, q, \pi, \tilde{\sigma}) \right] \\
&= \prod_{i=1}^{N} \left[\frac{1}{R} \sum_{r=1}^{R} \left(\pi' f_{i1}^{(r)} \left(\prod_{t=2}^{T-1} Q_{t-1} f_{it}^{(r)} \right) Q_{T-1} \tilde{f}_{iT}^{(r)} \right) \right],
\end{aligned}$$

where $\beta_i^{(r)}$ is a draw from a distribution with density $f(\beta_i|\theta)$ and $f_{it}^{(r)}$ has m^{th} element $\Pr[Y_{it} = y_{it}|\beta_i^{(r)}, \tilde{\sigma}, s_{it} = m]$ for $r = 1, ..., R$. The corresponding simulated log-likelihood function is given by

$$\log p(y|\theta, q, \pi, \tilde{\sigma}) \approx \sum_{i=1}^{N} \log \left[\frac{1}{R} \sum_{r=1}^{R} \left(\pi' f_{i1}^{(r)} \left(\prod_{t=2}^{T-1} Q_{t-1} f_{it}^{(r)} \right) Q_{T-1} \tilde{f}_{iT}^{(r)} \right) \right].$$

A.2. Hidden Markov rank-ordered logit model

For the HM-ROL, the likelihood function can be written as

$$\begin{aligned}
p(y|\theta, q, \pi, \tilde{\sigma}) &= \prod_{i=1}^{N} \left[\int \left(\sum_{s_i^* \in \mathcal{S}} \Pr[S_i = s_i^* | q, \pi] p(y_i | \beta_i, \tilde{\sigma}, s_i^*) \right) f(\beta_i|\theta) d\beta_i \right] \\
&= \prod_{i=1}^{N} \left[\int \left\{ \sum_{s_i^* \in \mathcal{S}} \left(\Pr[S_i = s_i^* | q, \pi] \prod_{t=1}^{T} \prod_{h=1}^{J-1} \frac{\exp\left(\frac{1}{\tilde{\sigma}_{s_{ih}}} (x'_{ity_{ith}} \beta_i) \right)}{\sum_{l=h}^{J} \exp\left(\frac{1}{\tilde{\sigma}_{s_{ih}}} (x'_{ity_{itl}} \beta_i) \right)} \right) \right\} f(\beta_i|\theta) d\beta_i \right] \\
&= \prod_{i=1}^{N} \left[\int \left\{ \sum_{s_i^* \in \mathcal{S}} \left(\Pr[S_i = s_i^* | q, \pi] \prod_{h=1}^{J-1} \prod_{t=1}^{T} \frac{\exp\left(\frac{1}{\tilde{\sigma}_{s_{ih}}} (x'_{ity_{ith}} \beta_i) \right)}{\sum_{l=h}^{J} \exp\left(\frac{1}{\tilde{\sigma}_{s_{ih}}} (x'_{ity_{itl}} \beta_i) \right)} \right) \right\} f(\beta_i|\theta) d\beta_i \right] \\
&= \prod_{i=1}^{N} \left[\int \left\{ \pi' f_{i1} \left(\prod_{h=2}^{J-2} Q_{h-1} f_{ih} \right) Q_{J-2} \tilde{f}_{i,J-1} \right\} f(\beta_i|\theta) d\beta_i \right], \qquad (13)
\end{aligned}$$

where Q_h is a $(M \times M)$ transition probability matrix, and $\tilde{f}_{ih}$ is a $(M \times 1)$ vector with the likelihood contribution of respondent i at rank h given β_i and s_{ih} with m^{th} element equal to

$$\tilde{f}_{ihm} \equiv \prod_{t=1}^{T} \Pr[Y_{ith} = y_{ith} | y_{it1}, ..., y_{it,h-1}, \beta_i, \tilde{\sigma}, s_{it} = m] = \prod_{t=1}^{T} \frac{\exp\left(\frac{1}{\tilde{\sigma}_m} (x'_{ity_{ith}} \beta_i) \right)}{\sum_{l=h}^{J} \exp\left(\frac{1}{\tilde{\sigma}_m} (x'_{ity_{itl}} \beta_i) \right)}$$

Furthermore, f_{ih} is a diagonal $(M \times M)$ matrix with the diagonal equal to $\tilde{f}_{ih}$.

We approximate the likelihood function using Monte Carlo integration:

$$\begin{aligned}
p(y|\theta, q, \pi, \tilde{\sigma}) &\approx \prod_{i=1}^{N} \left[\frac{1}{R} \sum_{r=1}^{R} p(y_i | \beta_i^{(r)}, q, \pi, \tilde{\sigma}) \right] \\
&= \prod_{i=1}^{N} \left[\frac{1}{R} \sum_{r=1}^{R} \left(\pi' f_{i1}^{(r)} \left(\prod_{h=2}^{J-2} Q_{h-1} f_{ih}^{(r)} \right) Q_{J-2} \tilde{f}_{i,J-1}^{(r)} \right) \right],
\end{aligned}$$

where $\beta_i^{(r)}$ is a draw from a distribution with density $f(\beta_i|\theta)$ and $f_{ih}^{(r)}$ has m^{th} element $\prod_{t=1}^{T} \Pr[Y_{ith} = y_{ith} | y_{it1}, ..., y_{it,h-1}, \beta_i^{(r)}, \tilde{\sigma}, s_{it} = m]$ for $r = 1, ..., R$. The corresponding simulated log-likelihood function is given by

$$\log p(y|\theta, q, \pi, \tilde{\sigma}) \approx \sum_{i=1}^{N} \log \left[\frac{1}{R} \sum_{r=1}^{R} \left(\pi' f_{i1}^{(r)} \left(\prod_{h=2}^{J-2} Q_{h-1} f_{ih}^{(r)} \right) Q_{J-2} \tilde{f}_{i,J-1}^{(r)} \right) \right].$$

A.3. Miscellaneous details

The parameters q, π, and $\tilde{\sigma}$ are constrained, as are possibly several parameters in θ. To ensure unconstrained optimization of the simulated log-likelihood function, we reparametrize the constrained parameters in terms of parameters that are unconstrained and optimize over these unconstrained parameters. Furthermore, to increase the probability of finding a global maximum, we recommend using multiple starting values and picking the solution with gives the highest log-likelihood value.

Finally, we compute standard errors using the square root of the diagonal elements of the inverse of the negative Hessian of the log-likelihood function. We approximate the Hessian using the outer-product-of-gradients approximation based on the analytic gradient of the log-likelihood function. For this purpose, we consider the Hessian with respect to the untransformed, (possibly) constrained parameters in θ, q, π, and $\tilde{\sigma}$. Moreover, the log-likelihood function is again approximated using the same draws as used for the optimization.

B. Conditional distribution of S_{it}

For the hidden Markov multinomial logit model, the distribution of S_{it} conditional on the choices y_i of individual i is a multinomial distribution with outcomes $1, ..., M$ with corresponding probabilities that can be computed as follows. It holds that

$$\begin{aligned}
\Pr[S_{it} = m|y_i, \theta, q, \pi, \tilde{\sigma}] &= \int \Pr[S_{it} = m, \beta_i|y_i, \theta, q, \pi, \tilde{\sigma}] d\beta_i \\
&= \int \Pr[S_{it} = m|y_i, \beta_i, \theta, q, \pi, \tilde{\sigma}] f(\beta_i|y_i, \theta, q, \pi, \tilde{\sigma}) d\beta_i \\
&= \int \Pr[S_{it} = m|y_i, \beta_i, q, \pi, \tilde{\sigma}] \frac{p(y_i|\beta_i, q, \pi, \tilde{\sigma})}{p(y_i|\theta, q, \pi, \tilde{\sigma})} f(\beta_i|\theta) d\beta_i \\
&= \frac{1}{p(y_i|\theta, q, \pi, \tilde{\sigma})} \int \Pr[S_{it} = m|y_i, \beta_i, q, \pi, \tilde{\sigma}] p(y_i|\beta_i, q, \pi, \tilde{\sigma}) f(\beta_i|\theta) d\beta_i,
\end{aligned}$$

which can be approximated by

$$
\begin{aligned}
\Pr[S_{it}=m|y_i,\theta,q,\pi,\tilde{\sigma}] &\approx \frac{1}{p(y_i|\theta,q,\pi,\tilde{\sigma})} \times \frac{1}{R}\sum_{r=1}^{R} \Pr[S_{it}=m|y_i,\beta_i^{(r)},q,\pi,\tilde{\sigma}] \times p(y_i|\beta_i^{(r)},q,\pi,\tilde{\sigma}) \\
&= \frac{\sum_{r=1}^{R} \Pr[S_{it}=m|y_i,\beta_i^{(r)},q,\pi,\tilde{\sigma}] \times p(y_i|\beta_i^{(r)},q,\pi,\tilde{\sigma})}{\sum_{n=1}^{M}\sum_{r=1}^{R} \Pr[S_{it}=n|y_i,\beta_i^{(r)},q,\pi,\tilde{\sigma}] \times p(y_i|\beta_i^{(r)},q,\pi,\tilde{\sigma})},
\end{aligned}
$$

where for the second equality we use that $\sum_m \Pr[S_{it}=m|y_i,\theta,q,\pi,\tilde{\sigma}]=1$, and we let $\beta_i^{(r)}$ be a draw from a distribution with density $f(\beta_i|\theta)$ for $r=1,...,R$. The probability $\Pr[S_{it}=m|y_i,\beta_i^{(r)},q,\pi,\tilde{\sigma}]$ can be computed with the Hamilton filter (Hamilton, 1989) and a smoother (Kim, 1994).

The Hamilton filter sequentially computes the filtered probabilities ($\xi_{itm|t} \equiv \Pr[S_{it}=m|\{y_{il}\}_{l=1}^{t},\beta_i,q,\pi,\tilde{\sigma}]$) and predicted probabilities ($\xi_{i,t+1,m|t} \equiv \Pr[S_{i,t+1}=m|\{y_{il}\}_{l=1}^{t},\beta_i,q,\pi,\tilde{\sigma}]$) using

$$
\begin{aligned}
\xi_{itm|t} &= \frac{\xi_{itm|t-1}\Pr[Y_{it}=y_{it}|\beta_i,\sigma_{it}=\tilde{\sigma}_m]}{\sum_{n=1}^{M}\xi_{itn|t-1}\Pr[Y_{it}=y_{it}|\beta_i,\sigma_{it}=\tilde{\sigma}_n]}, \\
\xi_{i,t+1,m|t} &= \sum_{n=1}^{M} Q_{nm}\xi_{itn|t},
\end{aligned}
$$

for $m=1,...,M$ and $t=1,...,T$. The filter is initialised by $\xi_{i1m|0}=\Pr[S_{i1}=m|\beta_i,q,\pi,\tilde{\sigma}]=\pi_m$. Given the filtered and predicted probabilities up to $t=T$, the required smoothed estimates can be computed sequentially using (Kim, 1994)

$$
\Pr[S_{it}=m|y_i,\beta_i,q,\pi,\tilde{\sigma}] = \sum_{n=1}^{M} \xi_{i,t+1,n|T}\frac{Q_{m,n}\xi_{itm|t}}{\xi_{i,t+1,n|t}},
$$

for $t=T-1,T-2,...,1$.

For the hidden Markov rank-ordered logit model, the conditional probabilities that S_{ih} is equal to a phase m can be computed in a similar fashion.

C. Monte Carlo study: results DGPs 4-6

Table 6: Mean and RMSE (in parentheses) of the parameter estimates for the Monte Carlo study for DGPs 4 to 6. Based on 1,000 Monte Carlo replications per DGP.

		DGP 4[a]			DGP 5[b]			DGP 6[c]		
Parameter	True	MNL	H-MNL	HM-MNL	MNL	H-MNL	HM-MNL	MNL	H-MNL	HM-MNL
b_1	1.00	0.65	0.84	0.97	0.67	0.86	0.91	0.43	0.67	0.86
		(0.35)	(0.17)	(0.07)	(0.33)	(0.15)	(0.10)	(0.57)	(0.34)	(0.16)
b_2	0.30	0.20	0.25	0.29	0.20	0.26	0.27	0.13	0.20	0.26
		(0.11)	(0.05)	(0.03)	(0.10)	(0.05)	(0.04)	(0.17)	(0.10)	(0.05)
b_3	-0.50	-0.33	-0.43	-0.48	-0.34	-0.43	-0.46	-0.22	-0.34	-0.43
		(0.17)	(0.09)	(0.06)	(0.16)	(0.08)	(0.07)	(0.28)	(0.17)	(0.09)
$\sigma_{\beta,1}$	0.50	0.49	0.64	0.48	0.46	0.59	0.46	0.35	0.52	0.41
		(0.03)	(0.15)	(0.07)	(0.05)	(0.10)	(0.06)	(0.16)	(0.06)	(0.12)
$\sigma_{\beta,2}$	0.40	0.30	0.39	0.39	0.30	0.38	0.37	0.21	0.32	0.34
		(0.11)	(0.04)	(0.04)	(0.11)	(0.05)	(0.05)	(0.19)	(0.10)	(0.08)
$\sigma_{\beta,3}$	0.70	0.52	0.68	0.68	0.52	0.66	0.64	0.37	0.55	0.60
		(0.19)	(0.08)	(0.08)	(0.19)	(0.08)	(0.09)	(0.34)	(0.17)	(0.14)
$\rho_{\beta,12}$	0.00	0.25	0.26	-0.02	0.20	0.20	0.00	0.26	0.26	-0.05
		(0.26)	(0.27)	(0.17)	(0.22)	(0.22)	(0.13)	(0.29)	(0.29)	(0.27)
$\rho_{\beta,13}$	0.00	-0.25	-0.25	0.01	-0.19	-0.20	0.00	-0.26	-0.26	0.03
		(0.26)	(0.27)	(0.17)	(0.21)	(0.22)	(0.15)	(0.30)	(0.30)	(0.26)
$\rho_{\beta,23}$	0.00	-0.12	-0.12	0.01	-0.09	-0.09	0.00	-0.13	-0.13	0.01
		(0.17)	(0.17)	(0.14)	(0.15)	(0.16)	(0.14)	(0.23)	(0.23)	(0.20)
π_1				0.280			0.257			0.344
π_2				0.593			0.664			0.600
π_3				0.127			0.079			0.056
q_{11}				0.378			0.397			0.586
q_{22}				0.981			0.984			0.938

For DGPs 4 to 6, we simulate data from the HM-MNL in Equations (3)-(7) with $\beta_i \sim MVN(b, \Sigma_\beta)$. In the DGPs, the transition probabilities are set equal over tasks, and respondents can only move one phase up. More details:
[a] DGP 4: 5 phases, $\tilde{\sigma} = (\infty, 2, 1, 2, \infty)$, $\pi = (0.2, 0.2, 0.4, 0.1, 0.1)$, $q_{11} = 0.20$, $q_{22} = 0.35$, $q_{33} = 0.98$, and $q_{44} = 0.80$.
[b] DGP 5: 5 phases, $\tilde{\sigma} = (10, 2, 1, 2, 10)$, $\pi = (0.2, 0.2, 0.4, 0.1, 0.1)$, $q_{11} = 0.20$, $q_{22} = 0.35$, $q_{33} = 0.98$, and $q_{44} = 0.80$.
[c] DGP 6: 9 phases, $\tilde{\sigma} = (\infty, 10, 5, 2, 1, 2, 5, 10, \infty)$, $\pi = (0.1, 0.1, 0.1, 0.1, 0.4, 0.05, 0.05, 0.05, 0.05)$, $q_{11} = q_{22} = q_{33} = q_{44} = 0.20$, $q_{55} = 0.90$, and $q_{66} = q_{77} = q_{88} = 0.70$