PSYCHOMETRIKA—VOL. 79, NO. 4, 647–674
OCTOBER 2014
DOI: 10.1007/S11336-013-9361-1

A BETA UNFOLDING MODEL FOR CONTINUOUS BOUNDED RESPONSES

YVONNICK NOEL

UNIVERSITY OF BRITTANY, RENNES 2

An unfolding model for continuous bounded responses is proposed, derived both from a hypothetical interpolation response mechanism and from the hypothesis of two opposite sources of item refusal being collapsed. These two sources of refusal are made explicit in a three-component Dirichlet response model and then collapsed to obtain a (two-component) beta response model. The two natural parameters of the beta are interpreted as acceptance and refusal parameters and expressed as functions of person-item distances on a latent continuum. The potentially bimodal shape of the beta is exploited to model chaotic response choices among ambivalent subjects.

Key words: item response theory, unfolding analysis, beta distribution, continuous bounded responses, bimodality.

1. Introduction

Though the origin of unfolding analysis for attitude measurement is often credited to Coombs (1950), who coined the term, its history could be traced back to Thurstone's (1927) works on attitude measurement. Despite such a long history, unfolding models have still been rarely used in applied studies until recently. One of the possible reasons is that the underlying response mechanism is slightly less intuitive than the cumulative one, which simply assumes that the higher the attitude level, the higher the response. The unfolding model of response assumes that an item rating will be higher around some ideal point on a latent attitudinal continuum, thus resulting in a single-peaked response function. The farther a person is from this ideal location, in either direction, the lower the item rating.

This response mechanism has been shown to be relevant in at least two main areas, evolutionary processes and conflictualized attitudes. Examples in the first field include: The development of moral judgements (Davison, Robins, & Swanson, 1978), of intellectual processes (DeMars & Erwin, 2003), of educational goals among students (Volet & Chalmers, 1992), and of processes of change in smoking cessation (Noel, 1999a, 1999b). Examples in the second area include: The structure of political opinions (van Schuur, 1993), attitude toward capital punishment (Andrich & Luo, 1993), abortion (Roberts & Laughlin, 1996), and environment (Andrich & Styles, 1998).

Though early developments in unfolding analysis already provided models for continuous data in the multidimensional scaling tradition (e.g., Ramsay, 1980; Greenacre & Browne, 1986), no probabilistic model exists (as far as we know) to unfold continuous bounded responses (CBR) in an item response theory framework. This paper aims at providing a model for CBR by which continuous responses with both a lower and an upper bound are designated. Such data may be collected by asking subjects to rate items by putting a mark at some point on a horizontal line segment with ends labeled, for instance, "0 % agreement" and "100 % agreement." The response is then measured as the distance from the left end of the segment to the subject's mark. But responses could also be proportions of time spent in one out of two tasks or, more generally, any kind of proportion data for which approximate continuity is arguable.

An R script for estimating the BUM parameters is available upon request from the author.

Requests for reprints should be sent to Yvonnick Noel, University of Brittany Rennes 2, Department of Psychology, Place du Recteur Henri Le Moal, CS 24307, 35043 Rennes Cedex, France. E-mail: yvonnick.noel@uhb.fr

Modeling the distribution of CBR has many advantages over modeling categorical data: Their graphical representation is often more suggestive, the number of parameters to be estimated is lower in general, and the quality of estimation is better with smaller sample sizes and shorter tests (Noel & Dauvier, 2007; Wang & Zeng, 1998). Moreover, CBR are now very easy to collect through sliders on a computer interface, and it has been reported that subjects find continuous scales more pleasing to use (McKelvie, 1978).

To model such data, one could think of dealing with the bounded nature of the response by using some ad hoc transformation (e.g., logit) that would change the response into an unbounded variable and that could then be modeled with some random-effects model (see, for instance, Samejima, 1973, 1974). By contrast, the approach adopted in the present paper is to directly model the response mechanism by which a bounded response is generated and to deduce a distributional model from this mechanism. This approach leads to a class of potentially more flexible models.

The paper is organized as follows: The interpolation response mechanism proposed by Noel and Dauvier (2007) to analyze CBR in a cumulative framework is first presented. The resulting Beta Response Model is then generalized to a Dirichlet Response Model, to deal with the situation where, besides item acceptability, two concurrent sources of refusal are potentially present in subjects' responses. An unfolding response model is derived by pooling the two sources of refusal in the expected response. An EM algorithm for estimating the model parameters is described, and a simulation study is conducted to test the quality of parameter recovery. An application to real data (attitudes toward abortion) is finally presented.

2. An Interpolation Response Mechanism

The situation where a person rates an item in a questionnaire on a continuous bounded response scale is now considered. Define the response variable X, measured as a distance from the left boundary and then arbitrarily rescaled to lie in [0; 1].

By contrast with Likert response scales, the continuous response format provides no semantic reference points except at the boundaries of the response segment. It is assumed that, in this situation, persons implicitly grant relevance values (or proximity judgments) to both extreme responses. Denote by $v^{(D)}$ and $v^{(A)}$ those psychological values, granted to the left and right ends, respectively (or Disagree and Agree extremities). By the interpolation response mechanism (Noel & Dauvier, 2007), the agreement response variable X on the unit-length segment is then assumed to be of the form

$$X = \frac{v^{(A)}}{v^{(D)} + v^{(A)}}. \tag{1}$$

This may be viewed as an interpolation process: Respondents are expected to put a check at a distance of the left boundary that is proportional to the relative value they give to the extreme positive answer. This simple mechanism closely resembles commonly used models in choice theory (Luce, 1959) or reinforcement learning (Herrnstein, 1961).

The random part of the model is introduced by defining a density for $v^{(D)}$ and $v^{(A)}$. Because those implicit values are thought of as nonnegative quantities, it is assumed for convenience that

$$\begin{aligned} v^{(A)} &\sim \Gamma(m, s), \\ v^{(D)} &\sim \Gamma(n, s). \end{aligned} \tag{2}$$

Latent values are assumed to follow a Gamma distribution with a common scale parameter. This assumption of a common scale parameter is reasonable given that respondents have

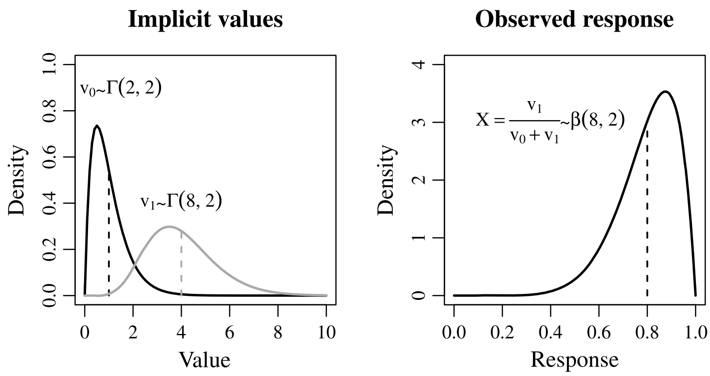

FIGURE 1.
Construction of an interpolated response. If the latent values for total disagreement and agreement are drawn from a $\Gamma(2, 2)$ (expected value 1) and a $\Gamma(8, 2)$ (expected value 4), respectively, the response will be generated from a $\beta(8, 2)$ (expected response $8/(8+2) = 0.8$).

to take both reference points simultaneously into account to construct a single response. From these assumptions, assuming conditional independence, it is known that (Kotz & Johnson, 1982, p. 229)

$$X = \frac{v^{(A)}}{v^{(D)} + v^{(A)}} \sim \beta(m, n). \tag{3}$$

The beta density, with two parameters m and n, is defined as

$$f(x; m, n) = \frac{\Gamma(m+n)}{\Gamma(m)\Gamma(n)} x^{m-1}(1-x)^{n-1} \quad \text{for } x \in [0; 1], \ m, n > 0, \tag{4}$$

with first moments

$$E(X; m, n) = \mu = \frac{m}{m+n}, \tag{5}$$

$$V(X; m, n) = \frac{mn}{(m+n)^2(m+n+1)} = \mu(1-\mu)\left[\frac{1}{m+n+1}\right]. \tag{6}$$

As the beta density is defined on [0; 1], it is well suited to the modeling of CBR, which may be easily rescaled to lie in [0; 1]. Expressions (5) and (6) show that it is relatively straightforward to define both mean and dispersion models within the beta, and this has already been derived in various flavors in the regression (Ferrari & Cribari-Neto, 2004; Smithson & Verkuilen, 2006; Verkuilen & Smithson, 2012) or IRT (Noel & Dauvier 2007; Tamhane, Ankenman, & Yang, 2002) fields. An example of this ratio construction is illustrated on Figure 1. In this example, implicit values for both extreme responses (disagreement and agreement) are assumed to be drawn from a $\Gamma(2, 2)$ (expected value 1) and a $\Gamma(8, 2)$ (expected value 4), respectively. It is thus expected that more weight be put on agreement than on disagreement in this situation. From (3), the observed response will be generated from a $\beta(8, 2)$. From (5), the expected response is simply calculated as $8/(8+2) = 0.8$, and high ratings are expected on average.

Depending on the values of m and n, the beta may show very different shapes, as illustrated on Figure 2. For $m = n = 1$, the density is uniform. Strong agreement becomes more likely as m is large relatively to n, and conversely. This legitimates an interpretation of m and n as "acceptance" and "refusal" parameters, respectively. For some parameter values ($m < 1$ and $n < 1$), it may also show a bimodal shape, which will be argued below to be especially relevant to model ambivalent attitudes.

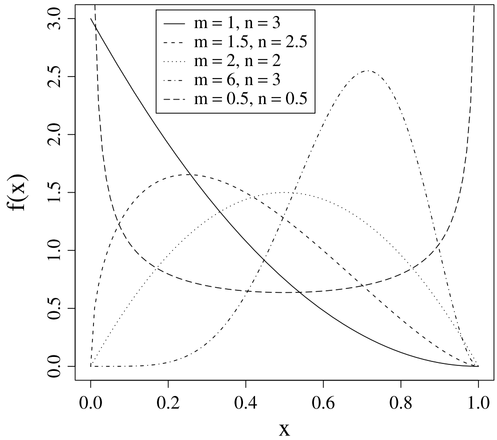

FIGURE 2.
The beta density for some parameter values.

Note that the choice of the Gamma for the value distribution is not the only one possible: Any density function defined on $\mathbb{R}^+$ could be used in a similar construction. For instance, the ratio of lognormal variates is known to follow a logistic-normal (Aitchison & Shen, 1980). But the Gamma has several advantages. First, it can have a mode at zero, which does not occur with other alternative distributions (lognormal, Weibull). Second, the sum of independent Gamma is still a Gamma, and this class-invariance property will be especially useful in the unfolding construction below, where two sources of refusal will be added. Alternatives like the lognormal, Weibull, or Generalized Exponential (Gupta & Kundu, 1999) do not enjoy this property. Third, the ratio construction above leads to a well-known distribution (the beta), with closed-form expressions both for the density and its first moments. This will help derive simple expressions for the expected rating curves, an important feature in the context of item response modeling. Fourth, the distribution of a ratio of Gamma variates can be bimodal, and bimodality will play a special role in the model to be presented below.

3. The Beta Unfolding Model

3.1. *A generic Dirichlet Response Model*

We consider the usual psychometric situation where N persons give ratings x_{ij} on a set of J items ($i = 1, \dots, N$, $j = 1, \dots, J$). Following Andrich and Luo (1993) and Verhelst and Verstralen (1993), an unfolding response process is generated each time a person is likely to refuse an item for two opposite reasons ("I cannot wholeheartedly support either side of the abortion debate"). In this context, the response is assumed to be determined by an implicit assessment of three possible extreme attitudes: Fully Agree (A), Fully Disagree for the first reason (D_1), and Fully Disagree for the second (opposite) reason (D_2). These are what Andrich and Luo (1993) named the "Agree", "Disagree from below", and "Agree from above" attitudes, respectively. Assume that for each item j, subject i implicitly assesses these three potential attitudes and grants them latent affinity values $v_{ij}^{(A)}$, $v_{ij}^{(D_1)}$, and $v_{ij}^{(D_2)}$.

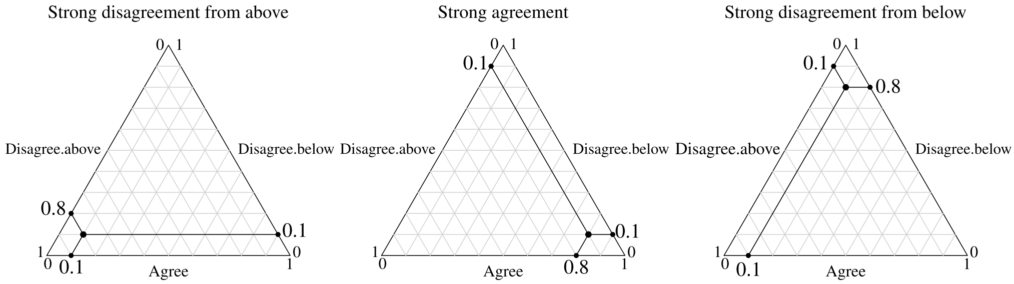

FIGURE 3.
Ternary plot of an unfolding response mechanism.

In the line of the preceding section, the following conditional, subject-specific, gamma random model is considered:

$$\begin{aligned} v_{ij}^{(A)} &\sim \Gamma(m_{ij}, s), \\ v_{ij}^{(D_1)} &\sim \Gamma(p_{ij}, s), \\ v_{ij}^{(D_2)} &\sim \Gamma(q_{ij}, s). \end{aligned} \tag{7}$$

As the scale parameter s will not play any role in the following, it is indifferent whether it is considered subject (or item) specific or not.

If made explicit in the response format (for instance, by instructing subjects to mark a dot in a unit-side ternary plot), these implicit assessments would lead subject i answering item j to give three observed responses, say x_{ij}, y_{ij}, and z_{ij} (or $1 - x_{ij} - y_{ij}$). In this situation, a straightforward generalization of the binary interpolation response mechanism presented above assumes that subjects would pick a position in the response triangle that is at the weighted average of the triangle vertices, using latent values as weights. On a unit-side ternary plot, this defines a random response vector of the form

$$(X_{ij}, Y_{ij}, Z_{ij})' = \left(\frac{v_{ij}^{(A)}}{v_{ij}^{(A)} + v_{ij}^{(D_1)} + v_{ij}^{(D_2)}}, \frac{v_{ij}^{(D_1)}}{v_{ij}^{(A)} + v_{ij}^{(D_1)} + v_{ij}^{(D_2)}}, \frac{v_{ij}^{(D_2)}}{v_{ij}^{(A)} + v_{ij}^{(D_1)} + v_{ij}^{(D_2)}} \right)'.$$

This is illustrated in Figure 3 for three contrasted attitudes. From the distributional assumptions in (7), a well-known result states that the three-dimensional response vector will be Dirichlet-distributed:

$$(X_{ij}, Y_{ij}, Z_{ij}) \sim \mathrm{Dir}(m_{ij}, p_{ij}, q_{ij}). \tag{8}$$

In the spirit of item response theory, a connection with a latent dimension and with person attitude (θ_i) and item location (δ_j) parameters is obtained by constructing the following structural model:

$$\begin{cases} m_j = \exp \lambda_j, \\ p_{ij} = \exp[\theta_i - \delta_j], \\ q_{ij} = \exp\bigl[-(\theta_i - \delta_j)\bigr], \end{cases} \tag{9}$$

where λ_j is an item-specific acceptability parameter. The exponential in these expressions simply ensures that m_j, p_{ij}, and q_{ij} parameters take on strictly positive values. The refusal parameters p_{ij} and q_{ij} are defined as increasing (p_{ij}) and decreasing (q_{ij}) functions of person-item distance

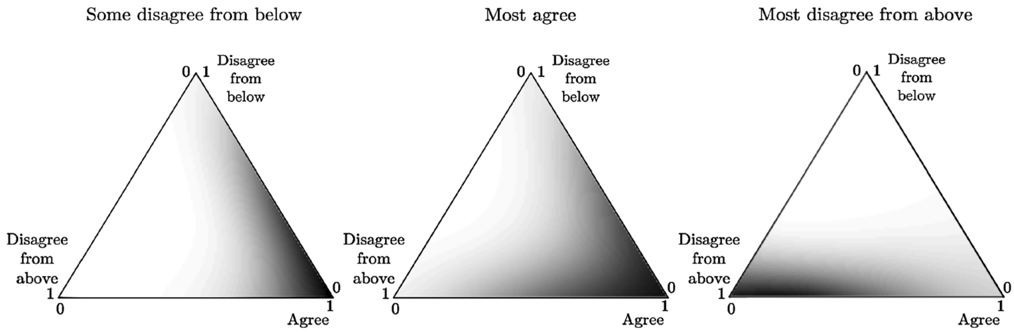

FIGURE 4.
Ternary density plots of sampled responses from the Dirichlet Response Model. *Middle panel*: When there is disagreement with the item, it is either on the "Disagree from below" or "Disagree from above" sides, but not both, hence the whiter area between both disagreement vertices.

on some latent continuum, modeling the interplay between two sources of refusal operating in opposite directions. In the line of Andrich and Luo (1993), the assumed response mechanism is essentially a double-refusal process. This takes into account that unfolding items are usually unlikely to be positively endorsed for themselves, but rather because refusability is low, for any one of two opposite reasons. But although acceptability is here modeled as a constant, item-specific characteristic, this will be shown below to lead to a subject-varying acceptance function, through the corresponding expectation function.

An illustrative plot of three-dimensional responses sampled from the Dirichlet Response Model is provided in Figure 4, for three items located at $\delta_1 = +1$, $\delta_2 = 0$, and $\delta_3 = -2$, drawing latent attitudes θ_i from a standard Gaussian. The contradictory nature of the two sources of refusal, as implied by this model, is apparent on the middle panel of Figure 4: When there is disagreement with the item, it is either on the "Disagree from below" or on the "Disagree from above" sides, but not on both (hence, the white area between bottom-left and top-right vertices).

3.2. *An Unfolding Model*

In the item rating context, where a simple Disagree–Agree response scale is provided, effects of the two sources of refusal are actually collapsed into a single observed response. From the properties of the Gamma, under the hypothesis of conditional independence across value assessments, the Dirichlet enjoys an interesting aggregation property: The response vector remains Dirichlet distributed upon aggregating some of the response components, the corresponding Dirichlet parameters simply being summed up.

Define the total disagreement value $v_{ij}^{(D)}$ as

$$v_{ij}^{(D)} = v_{ij}^{(D_1)} + v_{ij}^{(D_2)}.$$

Under the conditions mentioned above, we then have

$$v_{ij}^{(D)} \sim \Gamma(n_{ij}, s) \quad \text{with } n_{ij} = p_{ij} + q_{ij},$$

and the observed response X_{ij} on the unit response segment is such that $X_{ij}|\theta_i \sim \text{Dir}(m_j, p_{ij} + q_{ij})$, that is, a Beta distribution $\beta(m_j, n_{ij})$. Note that this aggregation property, inherited from the Gamma, does not exist in general for other distributions on the positive half-line, like the logistic-normal, as already mentioned by Aitchison and Shen (1980).

A Beta Unfolding Model (BUM) is thus fully specified as

$$\begin{cases} X_{ij}|\theta_i \sim \beta(m_j, n_{ij}), \\ m_j = \exp \lambda_j, \\ n_{ij} = \exp[\theta_i - \delta_j] + \exp\bigl[-(\theta_i - \delta_j)\bigr]. \end{cases} \tag{10}$$

From the properties of the beta density, the expected response function for item j is

$$\begin{aligned} E(X_{ij}|\theta_i) &= \mu_{ij} \\ &= \frac{m_j}{m_j + n_{ij}} \qquad (11) \\ &= \frac{\exp \lambda_j}{\exp \lambda_j + 2\cosh(\theta_i - \delta_j)}. \qquad (12) \end{aligned}$$

Interestingly, it is exactly of the hyperbolic cosine form independently derived for dichotomous data by Andrich and Luo (1993) and Verhelst and Verstralen (1993). Plots of the expected rating curve and the corresponding value functions are provided in Figure 6. The λ_j parameter here has the same interpretation of an item-specific acceptability parameter. But this quantity does not take the same importance for all values of θ_i in the expected response, and this leads to an attitude-varying model of acceptance. Note that, by contrast with Bernoulli models, the resulting response function in (12), although bounded by 0 and 1, is not a probability function, but an expectation function.

3.3. Generalized Forms of the Beta Unfolding Model

It may appear in practice desirable to model response expectation and variance in a (relatively) decoupled manner. Because the variance is directly related to the sum of the canonical parameters $m_j + n_{ij}$ (see Equation (6)) in a beta model, the model variance is easily modulated by introducing an additional scaling parameter $\phi_j = \exp \tau_j$ in m_j and n_{ij} such that

$$\begin{cases} m_j = \exp(\lambda_j + \tau_j), \\ n_{ij} = \exp\bigl[(\theta_i - \delta_j) + \tau_j\bigr] + \exp\bigl[-(\theta_i - \delta_j) + \tau_j\bigr]. \end{cases} \tag{13}$$

This will leave the expectation function unchanged:

$$\begin{aligned} E(X_{ij}|\theta_i) &= \frac{\exp(\lambda_j + \tau_j)}{\exp(\lambda_j + \tau_j) + \exp[(\theta_i - \delta_j) + \tau_j] + \exp[-(\theta_i - \delta_j) + \tau_j]} \\ &= \frac{\exp \lambda_j}{\exp \lambda_j + 2\cosh(\theta_i - \delta_j)}, \qquad (14) \end{aligned}$$

and the variance function now reads

$$V(X_{ij}|\theta_i) = \frac{\mu_{ij}(1 - \mu_{ij})}{1 + \phi_j[\exp \lambda_j + 2\cosh(\theta_i - \delta_j)]}. \tag{15}$$

The effects of varying λ_j and τ_j on response density and expectation are illustrated in Figure 5. Each of them mostly affects response expectation and variance, respectively. Higher values of λ correspond to higher expected responses (item acceptability increases), and higher values of τ correspond to a smaller response dispersion, given the true attitude. As will be further commented below, this parameterization suffices to create a family of flexible response functions with varying peakedness.

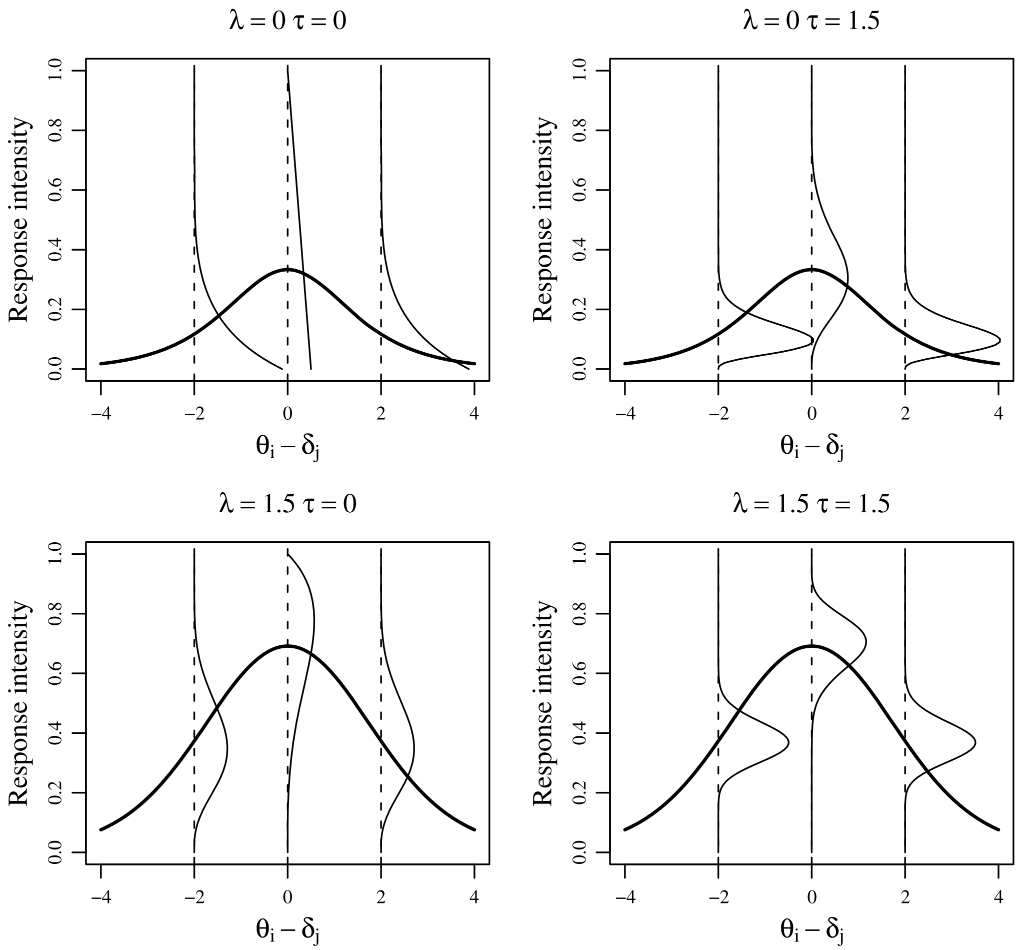

FIGURE 5.
Expectation rating curves and conditional response density functions. Varying λ (*from upper to lower panels*) and varying τ (*from left to right panels*) affects response expectation and conditional variance, respectively.

Of course, the λ_j and τ_j are both item-specific parameters that are not separable in (13). With no loss of generality, one may reparameterize as

$$\begin{cases} m_j = \exp \lambda_j^*, \\ n_{ij} = \exp\left[(\theta_i - \delta_j) + \tau_j\right] + \exp\left[-(\theta_i - \delta_j) + \tau_j\right], \end{cases} \tag{16}$$

with $\lambda_j^* = \lambda_j + \tau_j$. For convenience, λ_j^* will be simply written λ_j in what follows. The over-parameterization in n_{ij} is practically dealt with by imposing $\sum_i \theta_i = 0$.

It may also be convenient, for interpretation or comparison purposes, that the attitude distribution be scaled to a fixed unit, for instance, by imposing $\sum_i \theta_i^2 = N$. In this case, a common slope constant will play the role of a (positive) scaling factor:

$$\begin{cases} m_j = \exp \lambda_j, \\ n_{ij} = \exp\left[\alpha(\theta_i - \delta_j) + \tau_j\right] + \exp\left[-\alpha(\theta_i - \delta_j) + \tau_j\right] \\ \quad\; = 2\phi_j \cosh\left[\alpha(\theta_i - \delta_j)\right]. \end{cases} \tag{17}$$

The expected response function becomes

$$E(X_{ij}|\theta_i) = \frac{\exp\lambda_j}{\exp\lambda_j + 2\phi_j\cosh[\alpha(\theta_i - \delta_j)]}. \tag{18}$$

Although tempting, there is no need to free the scaling constant and make it item-specific (i.e., a discrimination parameter), as is common in Bernoulli response models. This would result in high parameter redundancy and lead to numerical estimation problems. Actually, the (λ, τ) parameterization already produces sufficiently flexible response curves (see Figure 5). Similar remarks may be found in Verhelst and Verstralen (1993) about a Bernoulli unfolding model with very similar equations.

3.4. Bimodality

An unusual feature of the BUM is the ability to model bimodal response density functions. This will occur each time one has simultaneously $m_j < 1$ and $n_{ij} < 1$ for an item j and some range of θ, that is,

$$\begin{cases} \lambda_j < 0, \\ \tau_j < -\ln\bigl[2\cosh\bigl(\alpha(\theta_i - \delta_j)\bigr)\bigr]. \end{cases} \tag{19}$$

In the following, such items will be referred to as bimodal response density (BRD) items, as opposed to unimodal response density (URD) items. It should be noted that this terminology characterizes the shape of the conditional response density for a fixed level of attitude, and not the shape of the expected response function, which is of course always single-peaked in an unfolding model. Value and response functions for an URD item ($\delta_j = 0$, $\lambda_j = 2$, $\alpha = 1$, $\tau_j = 0$) and a BRD one ($\delta_j = 0$, $\lambda_j = -0.8$, $\alpha = 1$, $\tau_j = -2$) are displayed on the top and bottom rows of Figure 6. These plots deserve some attention, as they summarize most of the BUM features. On each row, the left panel shows the value functions (i.e., acceptability and refusal parameters as functions of attitude), and the right panel displays the corresponding expected and modal response curves. Contour lines of the response density have been added to the response function graph (as thin lines on the right panels), to help see the difference between both cases.

In the URD case (top row), both value functions (m_j and n_{ij}) remain above one (left panel), and the conditional response density (right panel) is never bimodal: The response density is high around the modal response curve (curve of higher density responses) and decreases as more extreme responses (given θ) are considered above and below this curve. Note that, as a result of the dissymmetry of the beta, the modal response curve will always lie at some distance of the expectation, except at the two attitudes values where acceptability and refusal have equal intensities (see dots on the plots); the beta is perfectly symmetric at these points (along the vertical dotted lines). This is to keep in mind to correctly interpret real data plots, where the expected rating curve will seem not to pass "through the middle of the data" (but the modal response curve will). Although somewhat puzzling at first glance, it should be noted that the very same phenomenon exists with binomial models but the effect (i.e., the difference between $N\pi$ and $(N + 1)\pi$ in a $B(N, \pi)$ model) is so small that it is almost invisible. From one end of the attitudinal continuum to the other, the response density evolves from shapes where refusal dominates (for the first reason), then acceptance, and then refusal again (for the second reason). The acceptance region (marked as "A" on the plot), defined by $m_j > n_{ij}$, is bounded by the critical attitude values $\delta_j \pm \alpha^{-1}\text{acosh}[0.5\exp\lambda_j]$, for which agreement and refusal are equal and $E(X_{ij}) = 0.5$.

In the BRD case (bottom row), the conditional response density function is unimodal only for low or high values of $\theta_i - \delta_j$ (refusal regions) but becomes bimodal when both m_j and n_{ij} fall

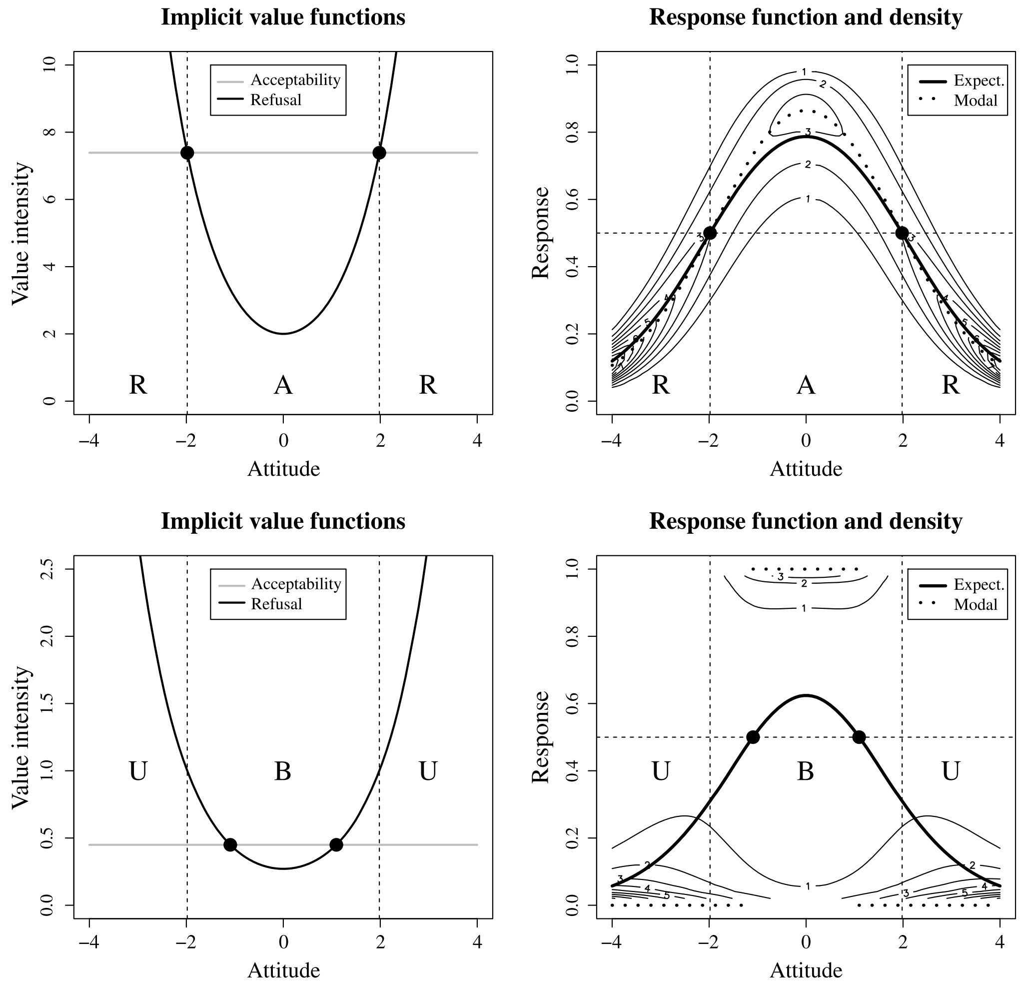

FIGURE 6.
Value functions and response isodensity plots of unimodal (*top row*) and bimodal (*bottom row*) response density items. Note: *Thick lines* are expected (*plain*) and modal (*dotted*) response curves, *thin plain lines* are response density contours. *Top row*: "R" and "A" mark regions where refusal or acceptation dominates. *Bottom row*: "U" and "B" mark unimodal and bimodal response density regions, respectively.

below 1 (left panel). This defines a central interval on θ (marked as "B" in Figure 6), bounded by the critical values $\delta_j \pm \alpha^{-1}\,\text{acosh}[0.5\exp(-\tau_j)]$. For any given θ_i in this interval, both extreme responses are more probable than the expected one. This may be viewed from contour plots of response density (right panel), for instance, for $\theta = 0$; the expected response at this point (plain curve) has a density lower than 1 when both extreme responses ($X_{ij} = 0$ and $X_{ij} = 1$) have densities greater than 1. Within this interval, the density is bimodal and perfectly symmetric when $X_{ij} = 0.5$. At these points (see dots on the right panel), extreme responses (close to 0 or 1) have the same (highest) density. In this example, the modal response curve (thick dotted line) becomes a step function, suddenly switching from one extreme to the other of the response scale as $\theta - \delta_j$ increases (this is why it could not fully replace the expected response function for item characteristics description).

Bimodality, be it symmetric or not, is clearly an unusual feature for an IRT model, but it may help model response choices among ambivalent subjects. The model implies that if you do not have strong reasons to either endorse or reject an item (i.e., both m_j and n_{ij} are low), then your response choice is likely to be chaotic and suddenly bifurcate from one extreme to the other of the response scale. This also means that, in the bimodal case, agreement is considered more unstable than refusal. This seems reasonable in the unfolding context, as intermediate attitudes are conflicting by nature, and probably more difficult to endorse.

If bimodality is not expected to occur, from theoretical considerations, or simply not desired for the data at hand, one may want to impose either one of $m_j \geq 1$ or $n_{ij} \geq 1$ $\forall i$ for a given item j, in the estimation process, that is,

$$\begin{cases} \lambda_j \geq 0, \\ \tau_j \geq -\ln\bigl[2\cosh\bigl(\alpha(\theta_i - \delta_j)\bigr)\bigr]. \end{cases} \tag{20}$$

Anyone of these conditions suffices. Note that the constraint on τ_j in (20) should hold for any value of θ_i. Because the hyperbolic cosine is U-shaped with a minimum of $2\exp\tau_j$ at $\theta_i - \delta_j = 0$, the minimal constraint on this second condition is to impose $\tau_j \geq -\ln 2$. In practice, we will look for the minimal scaling factor to apply simultaneously on m_j and n_{ij} (at $\theta_i = \delta_j$) to achieve at least one of these conditions. This approach affecting both m_j and n_{ij} by the same scaling factor is not the only choice possible for preserving unimodality, but has the advantage to leave the expectation function unchanged.

In any case, that the bimodality feature is present has to be tested on the dataset at hand. In practice, both the constrained and unconstrained versions of the model may be fitted and compared on the basis of some goodness-of-fit measure (see below). If bimodality is a feature of the data, the unconstrained BUM is likely to capture it and should show better fit.

4. Parameter Estimation

4.1. *Analysis of the Loglikelihood Surface*

For a given item j with parameters Δ_j and the response set X_j, the BUM loglikelihood reads

$$\ln L(\Theta, \Delta_j | X_j) = -\sum_{i=1}^{N} \ln B(m_j, n_{ij}) + (m_j - 1)\sum_{i=1}^{N} \ln x_{ij} + \sum_{i=1}^{N}(n_{ij} - 1)\ln(1 - x_{ij}), \tag{21}$$

where $B(\cdot)$ is the Eulerian Beta function.

The item log-likelihood function in (21) is basically a sum of conditionally independent beta log-likelihoods. As a member of the exponential family, the univariate beta has a globally concave log-likelihood with respect to its canonical parameters. The BUM log-likelihood is therefore globally concave with respect to m_j and n_{ij} parameters and has a unique maximum. A condition for the preservation of a unique maximum, when the log-likelihood is reexpressed as a function of BUM structural parameters, is that these parameters (up to an orthogonal transformation) all appear in one-to-one strictly monotone transforms. This is obviously the case for $m_j = \exp\lambda_j$. To see that this is also the case for δ_j and τ_j, it is convenient to reparameterize the (unscaled) model as

$$\begin{cases} m_j = \exp\lambda_j, \\ n_{ij} = \exp\bigl[\theta_i - \delta_{1j}\bigr] + \exp\bigl[-(\theta_i - \delta_{2j})\bigr], \end{cases}$$

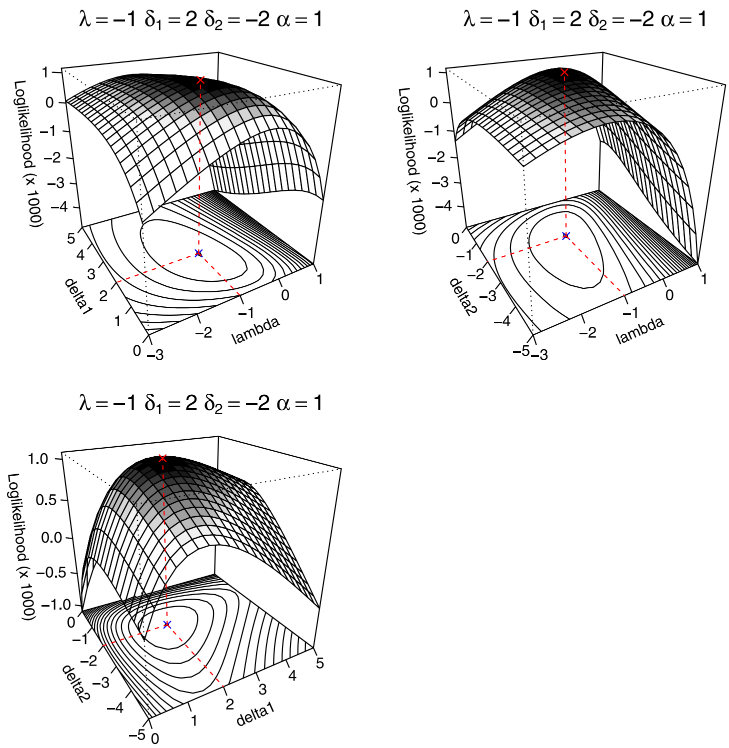

FIGURE 7.
Partial views of an item log-likelihood surface ($\lambda = -1$, $\delta_1 = 2$, $\delta_2 = -2$, $\alpha = 1$). Note: True parameter values and the corresponding log-likelihood values are plotted as *crosses* at the bottom and top of the graph.

where δ_{1j} and δ_{2j} threshold parameters are obtained by rotation/dilation in the (δ_j, τ_j) parameter subspace as

$$\begin{cases} \delta_{1j} = \delta_j - \tau_j, \\ \delta_{2j} = \delta_j + \tau_j. \end{cases}$$

These parameters do appear through a strictly monotonic (exponential) transform, so unimodality of the log-likelihood surface is preserved. Of course, this ensures that a global maximum exists but does not preserve concavity. This is illustrated in Figure 7, where partial views of the BUM log-likelihood surface are plotted for each possible pair of item parameters, the other two parameters being held fixed at their true values, for a sample item.

To construct these plots, a thousand attitude values have been generated, equally spaced between -4 and $+4$, and responses have been drawn from a BUM with $\lambda = -1$, $\delta_1 = 2$, $\delta_2 = -2$, and $\alpha = 1$. Contour lines for a number of fixed log-likelihood values have been added at the bottom of each plot, to help see that the corresponding parameter values do not necessarily form strictly convex sets, for instance, in the (δ_1, δ_2) subspace (the stretching effect of the exponential transforms in the δ_1 and δ_j directions is visible). However, concavity is not necessary to obtain good convergence properties, provided that the log-likelihood surface is unimodal, and the estimation procedure detailed below appears to work quite well in practice.

4.2. *Estimation Procedure*

As multimodal or nonsymmetric attitude distributions have sometimes been reported in previous unfolding studies (Andrich, 1995; Roberts, Donoghue, & Laughlin, 2000a; Roberts, Rost,

& Macready, 2000b), it seems unwise to assume a normal prior on θ in the general case. A finite mixture approach is adopted (Woodruff & Hanson, 1996), where the ability random variable is discretized to a set Θ of a priori defined θ_k ($k = 1, \dots, K$) values with probabilities π_k iteratively estimated along with item parameters (Mislevy, 1984). This approach may also be used with the π_k weights held fixed at values from some a priori defined probability model, if desired.

Let $f(\mathbf{x}_i|\Delta, \pi)$ be the density of an observed response vector $\mathbf{x}_i$, given the whole set Δ of item parameters and the vector π of latent scores probabilities. When the latent variable is discrete, this density is expressed as a finite mixture:

$$f(\mathbf{x}_i|\Delta, \pi) = \sum_k f(\mathbf{x}_i|\theta_k, \Delta)\pi_k. \tag{22}$$

An EM algorithm is used to estimate model parameters, which implies computing at each step s the expectation

$$Q\big(\Delta, \pi|\Delta^{(s)}, \pi^{(s)}\big) = E_\theta\left\{\ln\left[\prod_{i=1}^{N} f(\mathbf{x}_i, \theta_i|\Delta, \pi)\right]|\mathbf{X}, \Delta^{(s)}, \pi^{(s)}\right\} \tag{23}$$

in the E step and maximizing it with respect to item parameters and latent score probabilities in the M step.

Woodruf and Hanson (1996) show that, in the finite mixture reformulation, this expectation may be rewritten as the sum of two quantities,

$$\Phi_1(\Delta) = \sum_{k=1}^{K}\sum_{i=1}^{N}\left\{\ln\left[\prod_{j=1}^{p} f(x_{ij}|\theta_k, \Delta_j)\right] p\big(\theta_k|\mathbf{x}_i, \Delta^{(s)}, \pi^{(s)}\big)\right\} \tag{24}$$

and

$$\Phi_2(\pi) = \sum_{k=1}^{K}\ln(\pi_k)\sum_{i=1}^{N} p\big(\theta_k|\mathbf{x}_i, \Delta^{(s)}, \pi^{(s)}\big), \tag{25}$$

each of which depending only upon Δ and π, respectively.

In the E step, the posterior probability that the ith response vector results from a θ_k attitude level is computed as

$$p\big(\theta_k|\mathbf{x}_i, \Delta^{(s)}, \pi^{(s)}\big) = \frac{f(\mathbf{x}_i|\theta_k, \Delta^{(s)})\pi_k^{(s)}}{\sum_{k'=1}^{K} f(\mathbf{x}_i|\theta_{k'}, \Delta^{(s)})\pi_{k'}^{(s)}}. \tag{26}$$

In the M step, Φ_1 is maximized with respect to any item-specific η_j parameter by setting to zero the corresponding partial derivative

$$\frac{\partial \Phi_1(\Delta)}{\partial \eta_j} = \sum_{k=1}^{K}\sum_{i=1}^{N}\frac{\partial \ln f(x_{ij}|\theta_k, \Delta_j)}{\partial \eta_j} p\big(\theta_k|\mathbf{x}_i, \Delta^{(s)}, \pi^{(s)}\big), \tag{27}$$

where the quantity $p(\theta_k|\mathbf{x}_i, \Delta^{(s)}, \pi^{(s)})$ computed in the previous step is taken as fixed.

Finally, maximizing Φ_2 with respect to the π_k, subject to the constraint that $\sum_k \pi_k = 1$, leads to the updated estimates

$$\pi_k^{(s+1)} = \frac{1}{N}\sum_{i=1}^{N}\frac{f(\mathbf{x}_i|\theta_k, \Delta^{(s)})\pi_k^{(s)}}{\sum_{k'=1}^{K} f(\mathbf{x}_i|\theta_{k'}, \Delta^{(s)})\pi_{k'}^{(s)}}. \tag{28}$$

The partial derivatives of $\ln f(x_{ij}|\theta_k, \Delta_j)$ with respect to δ_j, λ_j, and τ_j are (see Appendix):

$$\frac{\partial \ln f(x_{ij}|\theta_k, \Delta_j)}{\partial \delta_j} = -2\alpha\phi_j \sinh\bigl[\alpha(\theta_k - \delta_j)\bigr] \times \bigl\{\Psi(m_j + n_{kj}) - \Psi(n_{kj}) + \ln(1 - x_{ij})\bigr\}, \tag{29}$$

$$\frac{\partial \ln f(x_{ij}|\theta_k, \Delta_j)}{\partial \lambda_j} = m_j\bigl\{\Psi(m_j + n_{kj}) - \Psi(m_j) + \ln x_{ij}\bigr\}, \tag{30}$$

$$\frac{\partial \ln f(x_{ij}|\theta_k, \Delta_j)}{\partial \tau_j} = n_{kj}\bigl\{\Psi(m_j + n_{kj}) - \Psi(n_{kj}) + \ln(1 - x_{ij})\bigr\}, \tag{31}$$

where $\Psi(x) = \partial \ln \Gamma(x)/\partial x$ is the digamma function, and $\phi_j = \exp \tau_j$.

Because log-transformed scores appear in the expressions above, responses that are exactly 0 or 1 cannot be dealt with as such. Moreover, it should be noted that in the unconstrained version of the BUM (i.e., m_j and n_{ij} are potentially lower than 1), exponents on x_{ij} and $1 - x_{ij}$ may be negative, and the density is no more defined for the extreme responses $x_{ij} = 0$ or $x_{ij} = 1$. Though it is not sure whether absolute agreement or disagreement with an item is psychologically plausible, extreme responses are of course observed in real data, and both issues are simply addressed by arbitrarily replacing 0 and 1 responses by numerical values arbitrarily closed to 0 and 1 (for instance, 10^{-3} and 1–10^{-3}). Although other proposals have been made to deal with extreme responses in beta models (Smithson & Verkuilen, 2005), we found in practice this slight correction to induce the smallest bias in the estimation of BUM parameters (see the simulation section).

Several iterative methods may be used to maximize ϕ_1 and ϕ_2. In the light of the discussion on the log-likelihood surface above, the Newton–Raphson algorithm does not seem appropriate. But the limited-memory quasi-Newton (L-BFGS-B) algorithm (Byrd, Lu, Nocedal, & Zhu, 1995) was found in practice to be very fast, numerically accurate, and flexible for managing parameter boundary constraints. In this algorithm, a negative definite approximation of the Hessian is numerically estimated at each iteration from the last t previous iterations ($t = 10$ in the following). This is implicitly done by maintaining in memory a history of the last t parameter vectors and gradients.

Although only gradient functions are needed for using the L-BFGS-B algorithm, exact expressions of the expected second-order cross-derivatives of the mixture log-likelihood with respect to item parameters are provided in the Appendix. The resulting 3×3 matrices are used at the end of the estimation process to compute standard errors of estimates for each item, by computing the inverse of minus these matrices and taking the square roots of their diagonal elements.

4.3. *Initial Values of Estimates*

Row and column scores from a correspondence analysis of the data matrix, centered and scaled with respect to row scores mean and variance, were found to provide a very good initial solution ($\theta_i^{(0)}$ and $\delta_j^{(0)}$) for θ_i and δ_j parameters. In the finite mixture formulation adopted here, the $\theta_i^{(0)}$ are then used to estimate initial values of the latent probabilities π_k, as relative frequencies of K fixed width intervals, over some a priori defined domain (e.g., $[-4.5; +4.5]$), outliers being included in the extreme intervals. The K interval midpoints are taken as the corresponding discrete θ_k attitude values. At step 0, all λ_j parameters are arbitrarily set to 1, and all τ_j are set to 0.

At the end of each estimation cycle over all items, the estimated distribution of the θ_k is centered around zero and scaled to unit variance (the corresponding observed standard deviation is used to fix the common slope α). The δ_j parameters are consequently adjusted.

4.4. *Estimation of Latent Attitudes*

Once item parameters have been estimated, they are taken as fixed in the estimation of attitude parameters. First (and often very good) estimates are already provided by the expected value of the Θ posterior distribution:

$$\hat{\theta}_i = \sum_{k=1}^{K} \theta_k p(\theta_k|\mathbf{x}_i, \Delta, \pi). \tag{32}$$

These can be still improved by maximizing the joint likelihood with respect to all θ_i, all item parameters being fixed. Denoting by $L(\Theta, \Delta|\mathbf{X})$ the joint likelihood of the BUM, we have:

$$\ln L(\Theta, \Delta|\mathbf{X}) = \sum_{i=1}^{N} \sum_{j=1}^{p} \ln f(x_{ij}|\theta_i, \Delta_j). \tag{33}$$

Person parameters are obtained by setting to zero the first derivative:

$$\begin{aligned} \frac{\partial \ln L(\Theta, \Delta|\mathbf{X})}{\partial \theta_i} &= 2\sum_{j=1}^{p} \alpha\phi_j \sinh\bigl[\alpha(\theta_i - \delta_j)\bigr] \\ &\quad \times \bigl\{\Psi(m_j + n_{ij}) - \Psi(n_{ij}) + \ln(1 - x_{ij})\bigr\}. \end{aligned} \tag{34}$$

A Fisher scoring scheme is obtained by using

$$E\left[\frac{\partial^2 \ln L(\Theta, \Delta|\mathbf{X})}{\partial \theta_i^2}\right] = 4\sum_{j=1}^{p} \alpha^2\phi_j^2 \sinh^2\bigl[\alpha(\theta_i - \delta_j)\bigr]\bigl\{\Psi'(m_j + n_{ij}) - \Psi'(n_{ij})\bigr\} \tag{35}$$

in an iterative procedure updating the θ_i as follows:

$$\hat{\theta}_i^{(t)} = \hat{\theta}_i^{(t-1)} - E\left[\frac{\partial^2 \ln L(\Theta, \Delta|\mathbf{X})}{\partial \theta_i^2}\right]^{-1} \left[\frac{\partial \ln L(\Theta, \Delta|\mathbf{X})}{\partial \theta_i}\right]. \tag{36}$$

4.5. *Information Functions*

The test information function reads

$$\begin{aligned} I(\theta_i) &= -E\left[\frac{\partial^2 \ln L(\Theta, \Delta|\mathbf{X})}{\partial \theta_i^2}\right] \\ &= 4\sum_{j=1}^{p} \phi_j^2\alpha^2 \sinh^2\bigl[\alpha(\theta_i - \delta_j)\bigr]\bigl\{\Psi'(n_{ij}) - \Psi'(m_j + n_{ij})\bigr\}. \end{aligned} \tag{37}$$

The item information function

$$\begin{aligned} I_j(\theta_i) &= 4\phi_j^2\alpha^2 \sinh^2\bigl[\alpha(\theta_i - \delta_j)\bigr]\bigl\{\Psi'(n_{ij}) - \Psi'(m_j + n_{ij})\bigr\} \\ &= I_j^{(1)}(\theta_i) \times \bigl\{I_j^{(2)}(\theta_i)\bigr\} \end{aligned} \tag{38}$$

is displayed in Figure 8 for various values of λ_j and τ_j (item location is fixed to $\delta = 0$). From the properties of the hyperbolic sine and the trigamma, the functions

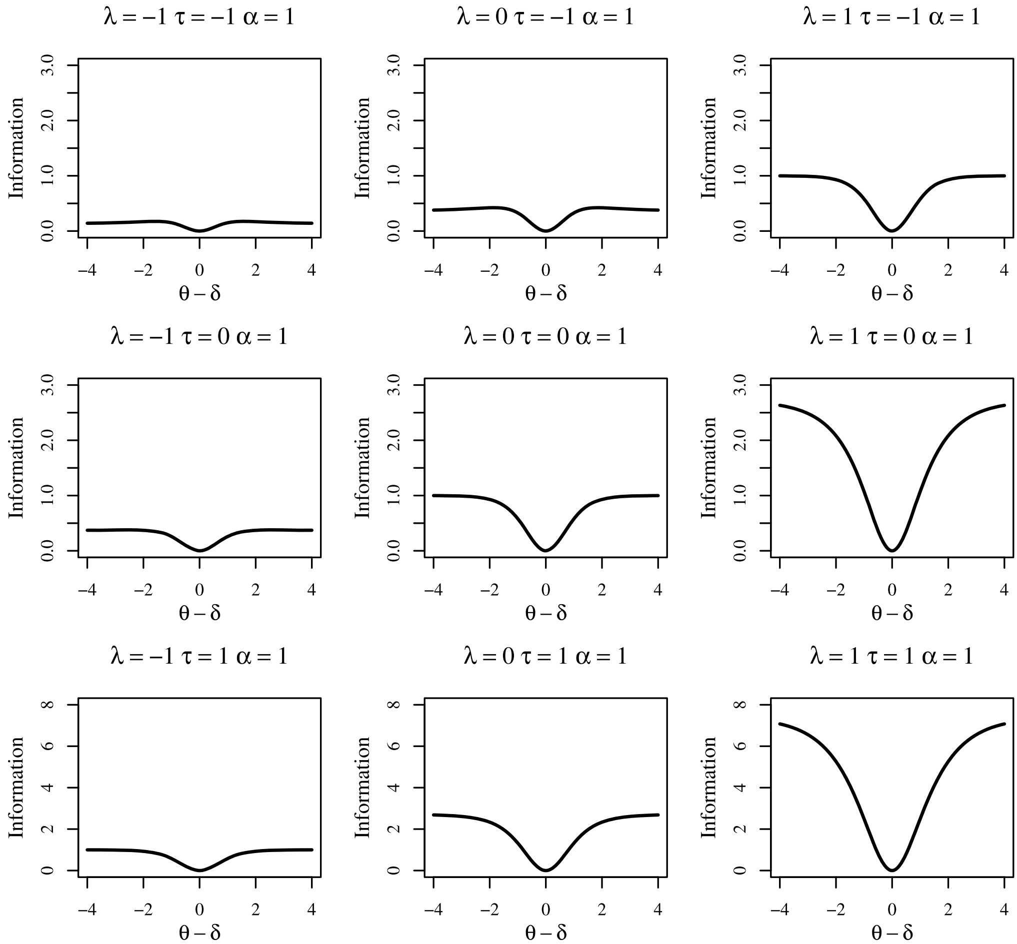

FIGURE 8.
Theoretical item information functions for varying values of λ and τ parameters. Note: λ_j values are varied from left to right, and τ_j values from top to bottom (note the different y-scale on the bottom row).

$$I_j^{(1)}(\theta_i) = 4\phi_j^2\alpha^2 \sinh^2\big[\alpha(\theta_i - \delta_j)\big],$$
$$I_j^{(2)}(\theta_i) = \Psi'(n_{ij}) - \Psi'(m_j + n_{ij})$$

are both second-order functions of θ_i, and so $I_j(\theta_i)$ is of fourth order, as has been observed with categorical unfolding models (Andrich & Luo, 1993; Roberts & Laughlin, 1996). It is null for $\theta_i - \delta_j = 0$ and increases, up to some optimum, with increasing values of $|\theta_i - \delta_j|$. Because the trigamma is strictly monotone decreasing on $\mathbb{R}^+$, higher values of ϕ_j (or τ_j) and λ_j are essentially associated with higher information levels.

5. A Simulation Study

A simulation study was conducted to test the accuracy of parameter estimates obtained by the algorithm described in the previous section, which was implemented in the R programming language (R Core Team, 2013). Three hundreds simulated datasets were generated for each of the 16 conditions resulting from factorially combining four levels of sample size (100, 250, 500, and 1000 observations) with four levels of test length (10, 15, 20, and 30 items). For each replication, subject parameters where drawn from a normal distribution and then exactly centered on zero and scaled to unit variance. Item locations were fixed equally spaced between -2 and $+2$. Dispersion and acceptance parameters were drawn at random from the uniform distributions in the $[-3; 1]$, and $[-1.5; 1.5]$ intervals, respectively. The scale parameter was fixed to 1.5. These choices resulted from preliminary experience with available datasets (see next section) but also ensured that bimodal response density was possible (τ_j and λ_j are potentially lower than $-\ln 2$ and 0, respectively). For the estimation of latent attitudes, the number of latent scores θ_k was fixed to $K = 30$, over the $[-4.5; 4.5]$ domain.

Several measures of estimation quality have been computed. Correlations (R) between true and estimated parameters are reported in Table 1. For a given item parameter (η), the root mean square error (*RMSE*), the signed bias (*BIAS*), and the standard error (*SE*) of estimates have been computed as follows:

$$RMSE = \left[\frac{1}{300}\sum_{r=1}^{300}(\eta - \hat{\eta}_r)^2\right]^{\frac{1}{2}},$$

$$BIAS = \frac{\sum_{r=1}^{300}(\eta - \hat{\eta}_r)}{300},$$

$$SE = \left[\frac{1}{300}\sum_{r=1}^{300}(\hat{\eta}_r - \bar{\eta})^2\right]^{\frac{1}{2}},$$

where $\hat{\eta}_r$ is the parameter estimate obtained at the rth replicate, and $\bar{\eta}$ is the averaged estimate over all 300 samples. These quantities are reported in Table 2. The squared bias is easily recovered from Table 2 by computing $BIAS^2 = RMSE^2 - SE^2$.

On the whole, the estimation algorithm performed very well, larger sample sizes being associated with increased accuracy in item parameter estimation, and larger test lengths with higher correlations between true and estimated attitude values. The average correlation between true and estimated attitude values was never lower than 0.96 (Table 1).

As to item locations, the *RMSE* values reported in the present analysis are lower than those usually reported for categorical unfolding models for the same sample size and test length. For comparison, Roberts and Laughlin (1996) report an $RMSE = 0.217$ for item location estimates obtained in the Graded Unfolding Model for $N = 500$ ($J = 20$, 6 response categories). This is in the order of magnitude of what is observed for $N = 100$ in the present simulation. Roberts, Donoghue, and Laughlin (1998) report *RMSE* values for δ estimates in the Generalized Graded Unfolding Model that fall below 0.10 for at least $N = 2000$. This is achieved for $N = 500$ in the present study. Wang and Zeng (1998) report *RMSE* values around 0.08 for item location recovery in Samejima's Continuous Response Model for $N = 500$, which is quite comparable to what is obtained here, for an unfolding model. So the algorithm seems to perform satisfactorily. We also note that the bias is quite negligible as far as the item locations are concerned.

The acceptability parameters had the smallest standard errors in the estimation but also the highest (positive) bias. This remains acceptable however (the larger bias value was 0.038 for $N = 100$ and $J = 10$) and seems to be well compensated by the sample size, the bias values falling below 0.01 as soon as $N \geq 500$. Conversely, dispersion parameters had the highest standard errors and a slight positive bias (which appears negligible with regard to the standard errors). These *SE* values fall below 0.10 only for $N = 1000$, so τ_j are the parameters in this model for which there is most uncertainty.

TABLE 1.
Correlations between true and estimated parameters.

J	N	Attitude	Location	Acceptability	Dispersion
10	100	0.979	0.992	0.989	0.969
10	250	0.966	0.997	0.992	0.995
10	500	0.967	0.998	0.998	0.995
10	1000	0.970	0.999	0.998	0.997
15	100	0.979	0.988	0.990	0.970
15	250	0.980	0.996	0.995	0.994
15	500	0.984	0.998	0.998	0.996
15	1000	0.980	0.999	0.999	0.997
20	100	0.984	0.988	0.987	0.973
20	250	0.989	0.997	0.996	0.991
20	500	0.986	0.998	0.998	0.996
20	1000	0.986	0.999	0.999	0.998
30	100	0.991	0.988	0.984	0.975
30	250	0.990	0.996	0.994	0.991
30	500	0.991	0.998	0.998	0.995
30	1000	0.991	0.999	0.999	0.998

Note: These figures are averages over 300 replications.

TABLE 2.
Root mean square error, signed bias, and standard errors of item parameter estimates.

		Location (δ)			Acceptability (λ)			Dispersion (τ)		
J	N	RMSE	BIAS	SE	RMSE	BIAS	SE	RMSE	BIAS	SE
10	100	0.220	0.002	0.220	0.158	0.038	0.153	0.320	0.014	0.319
10	250	0.118	−0.008	0.118	0.093	0.014	0.092	0.167	0.019	0.166
10	500	0.086	0.004	0.086	0.064	0.009	0.064	0.120	0.005	0.120
10	1000	0.061	0.000	0.061	0.044	0.004	0.044	0.086	0.001	0.086
15	100	0.217	−0.009	0.217	0.136	0.030	0.133	0.309	0.015	0.308
15	250	0.113	−0.004	0.113	0.095	0.018	0.093	0.154	0.016	0.153
15	500	0.078	−0.007	0.078	0.063	0.008	0.062	0.111	0.008	0.111
15	1000	0.068	−0.003	0.068	0.044	0.005	0.043	0.095	0.003	0.095
20	100	0.210	0.009	0.209	0.147	0.034	0.143	0.299	0.017	0.299
20	250	0.105	0.001	0.105	0.088	0.015	0.087	0.151	0.011	0.151
20	500	0.083	−0.001	0.083	0.059	0.006	0.059	0.119	0.006	0.119
20	1000	0.051	0.002	0.051	0.044	0.004	0.043	0.072	0.003	0.072
30	100	0.200	−0.001	0.200	0.142	0.029	0.139	0.287	0.015	0.286
30	250	0.112	0.000	0.112	0.085	0.012	0.084	0.162	0.012	0.161
30	500	0.077	0.000	0.077	0.060	0.006	0.059	0.110	0.005	0.110
30	1000	0.056	0.000	0.056	0.042	0.004	0.042	0.080	0.004	0.079

Note: These figures are averages over 300 replications. Squared biases may be recovered as $BIAS^2 = RMSE^2 - SE^2$.

Finally, if for practical purposes, one mainly considers the quality of estimation of attitude and item location parameters, then in the light of this simulation study, the minimum require-

TABLE 3.
Goodness-of-fit measures for several BUM models.

Unimodality	Latent distribution	Log-likelihood	#Parameters	#Constraints	AIC
Constrained	Gaussian	18961.89	99	80	−37725.79
Constrained	Estimated	19009.23	128	51	−37762.46
Unconstrained	Gaussian	24340.01	149	30	−48382.03
Unconstrained	Estimated	24406.64	178	1	−48457.29

ments would be to take $N \geq 500$ and $J \geq 10$ to get good estimates. This is an encouraging result, in conformity with what Noel and Dauvier (2007) already observed with the cumulative Beta Response Model (BRM): Working on continuous data do help to get good levels of accuracy.

6. Application on Real Data

A fifty-item questionnaire on attitude toward abortion (Roberts et al., 2000a; Roberts, Lin, & Laughlin, 2001) was translated into French and submitted to 443 subjects. For these sample size and test length, results in the previous section suggest that the quality of estimation should be acceptable. The subjects had to report their responses on a web interface written in PHP/Javascript. The response scale appeared as a 480-pixel horizontal slider. On page loading, the initial positions of the sliders were drawn at random, as was the item ordering. The recorded responses were divided by 480 to lie in the [0; 1] interval, and the responses that were exactly 0 or 1 were replaced by 10^{-3} and $1 - 10^{-3}$.

A succession of four BUM models were then fitted to the resulting dataset, constraining or not the response density to be everywhere unimodal and estimating the latent distribution or fixing it to be Gaussian. Goodness-of-fit measures (log-likelihood and Akaike Information Criterion) for each of these models are reported in Table 3, along with the numbers of estimated and fixed parameters. All parameters that had been constrained during the estimation process were considered as fixed in the computation of the *AIC*.

Estimating the latent distribution gave better fit than fixing it to be normal, in both conditions of constraint on the response density: A plot of the estimated score distribution (Figure 9) suggests a slight positive skewness. This dissymmetry is made apparent in Figure 9 by superimposing a kernel smoothed version of the estimated distribution with a standard Gaussian curve.

Relaxing the unimodality constraint on the response conditional density also resulted in better fits. This strongly suggests that some kind of bimodal response process does underlie the data, where for a given attitude (essentially in the ambivalence region illustrated in Figure 6, bottom panels), two opposite manifest responses are highly (and potentially equally) likely (this is further discussed below). The estimated item parameters are reported in Table 4 for a subset of 15 items, approximately evenly spaced along the dimension. Asymptotic standard errors are also reported, and they appear of the order of magnitude of what was observed in the simulation study (for $N = 500$). As is common in IRT estimation, the standard errors slightly increase for extreme item locations, which might be related to the fact that fewer subjects lie in this region of the latent dimension, and so less information is available. For information, classical item INFIT and OUTFIT measures (Wright & Masters, 1982) are also reported in the table, but more work is needed to explore the properties and behavior of these goodness-of-fit measures with the BUM.

Items are clearly scaled from positive ("43. Restrictions should never be placed on a woman's right to an abortion.") to negative ("13. Abortion is unacceptable under any circumstances.") attitudes toward abortion, ambivalent items lying in the middle range ("34. Abortion should be a woman's choice, but should never be used simply due to its convenience."). Observed

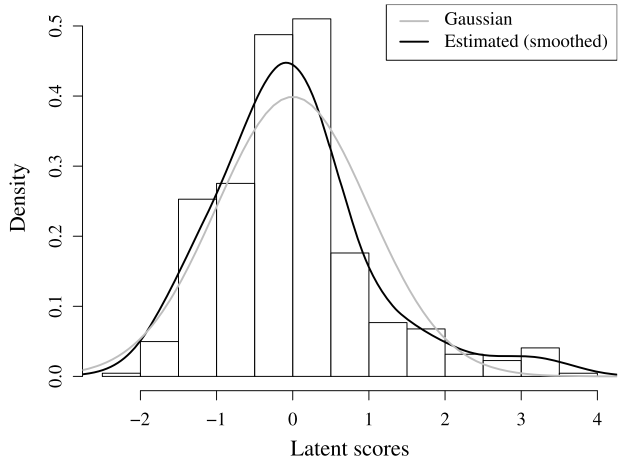

FIGURE 9.
Observed and smoothed distribution of estimated latent scores. Note: A Gaussian curve is added for comparison purposes.

TABLE 4.
Item parameter estimates for the abortion data.

Item	$\hat{\delta}_j$	A.S.E	$\hat{\lambda}_j$	A.S.E	$\hat{\tau}_j$	A.S.E	INFIT	OUTFIT
Item 43	−1.010	0.083	−0.536	0.063	−2.222	0.086	0.969	0.834
Item 47	−0.934	0.078	−0.626	0.063	−2.259	0.083	0.819	0.828
Item 48	−0.916	0.078	−0.646	0.062	−2.198	0.082	0.925	0.868
Item 40	−0.868	0.071	−0.323	0.069	−2.502	0.076	0.707	0.656
Item 50	−0.792	0.069	−0.569	0.066	−2.458	0.075	0.876	0.794
Item 29	0.535	0.073	−1.055	0.055	−1.390	0.103	0.928	0.884
Item 30	0.583	0.073	−1.032	0.056	−1.477	0.104	0.780	0.866
Item 34	0.685	0.062	−0.735	0.063	−2.465	0.093	1.016	0.901
Item 27	0.716	0.070	−0.889	0.057	−1.653	0.105	0.809	0.826
Item 16	1.968	0.107	−0.906	0.056	−2.522	0.156	0.881	1.084
Item 10	2.125	0.112	−1.451	0.055	−3.421	0.163	1.119	0.936
Item 22	2.176	0.108	−0.954	0.057	−3.354	0.157	1.016	0.924
Item 1	2.503	0.130	−0.979	0.056	−3.149	0.174	0.725	0.907
Item 6	2.609	0.137	−0.868	0.056	−3.057	0.179	0.699	0.769
Item 13	2.716	0.154	−1.022	0.055	−2.803	0.191	0.984	1.035

data (as dots), response density contours (as thin plain lines), and expectation functions (thick gray lines) for these items are plotted on Figure 10. In interpreting these plots, one should keep in mind that the BUM expected response function in (18) is not supposed to go "through the middle" of the data: Because the beta is not symmetric in general, the distribution expectation always lies at some distance from the mode (see Figure 6, thick lines).

The response density contours (thin lines) help appreciate how response density varies both with response value and latent attitude. For example, a subject whose attitude θ on the latent dimension is close to the location of item 47 ("Abortion is a reasonable alternative if a woman feels that having a baby might ruin her life", $\delta_{47} = -0.934$) is more likely to provide a response near the boundaries of the response scale (0 or 1), where the response density is above 1, than in the middle range, where the response density is lower than 1. Another view of this phenomenon is provided by the response density surface plots in Figure 11 for items 48 ("Abortion should be an accepted mechanism for family planning."), 29 ("Sometimes I am in favor of a woman's right

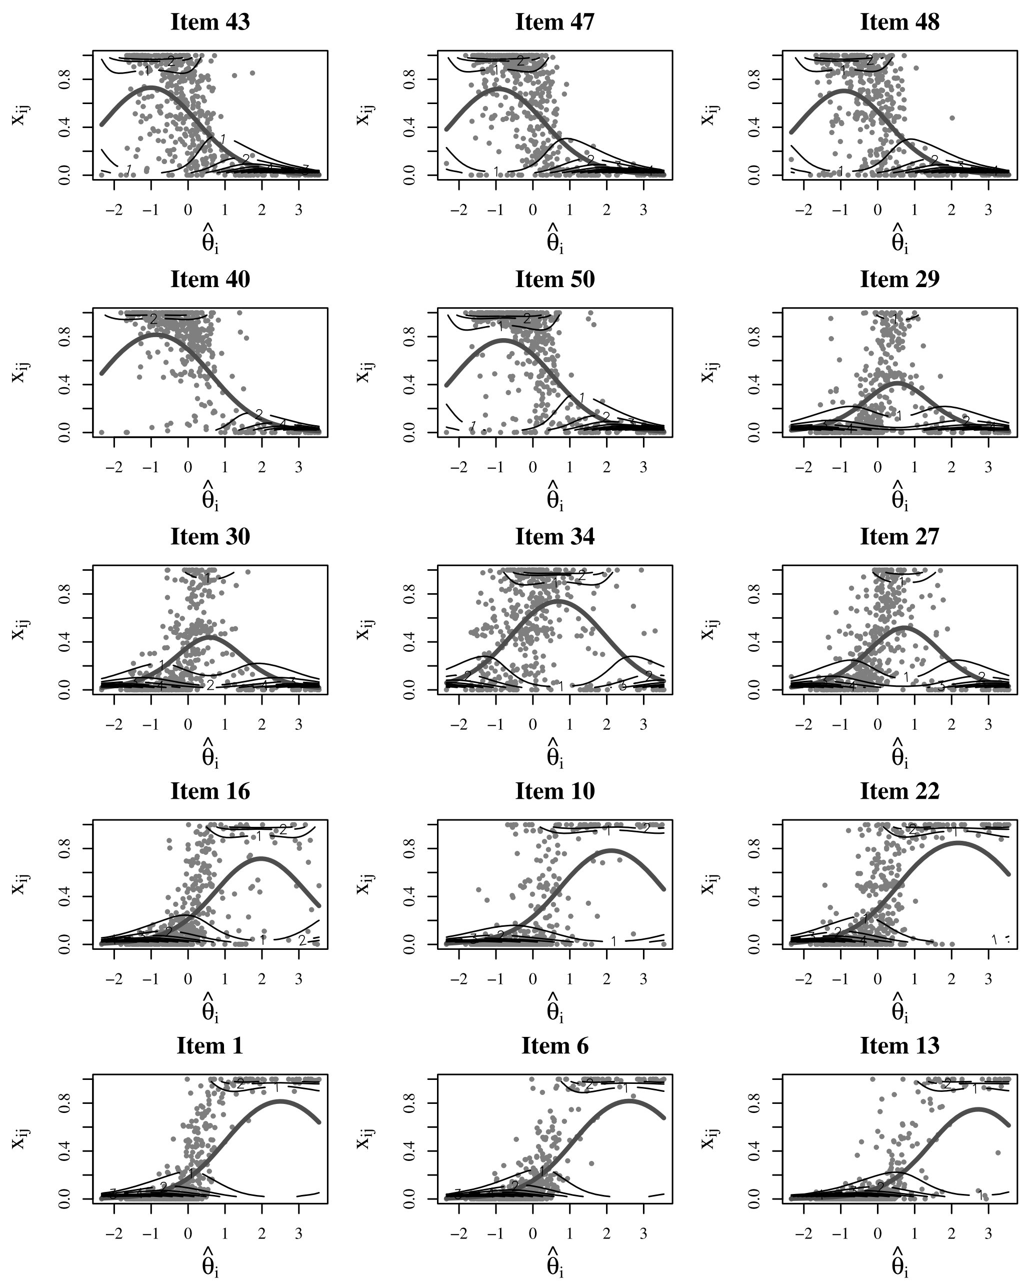

FIGURE 10.
Observed data, response density contours, and expectation function for the abortion data. The *thick plain lines* display the expected rating curve, the *thin lines* the response density contours as functions of attitude. Figures in the contour lines are the density values at that particular level.

to abortion, but at other times I am not."), 30 ("I cannot whole-heartedly support either side of the abortion debate."), and 6 ("Abortion is a threat to our society."). Along the response axis, for attitudes in the neighborhood of the item location (and especially for $\theta = \delta$), the density displays a characteristic "saddle" profile, the surface going upwards for extreme responses. This bimodality feature was in fact also present in Noel and Dauvier's (2007) BRM but judged unrealistic or

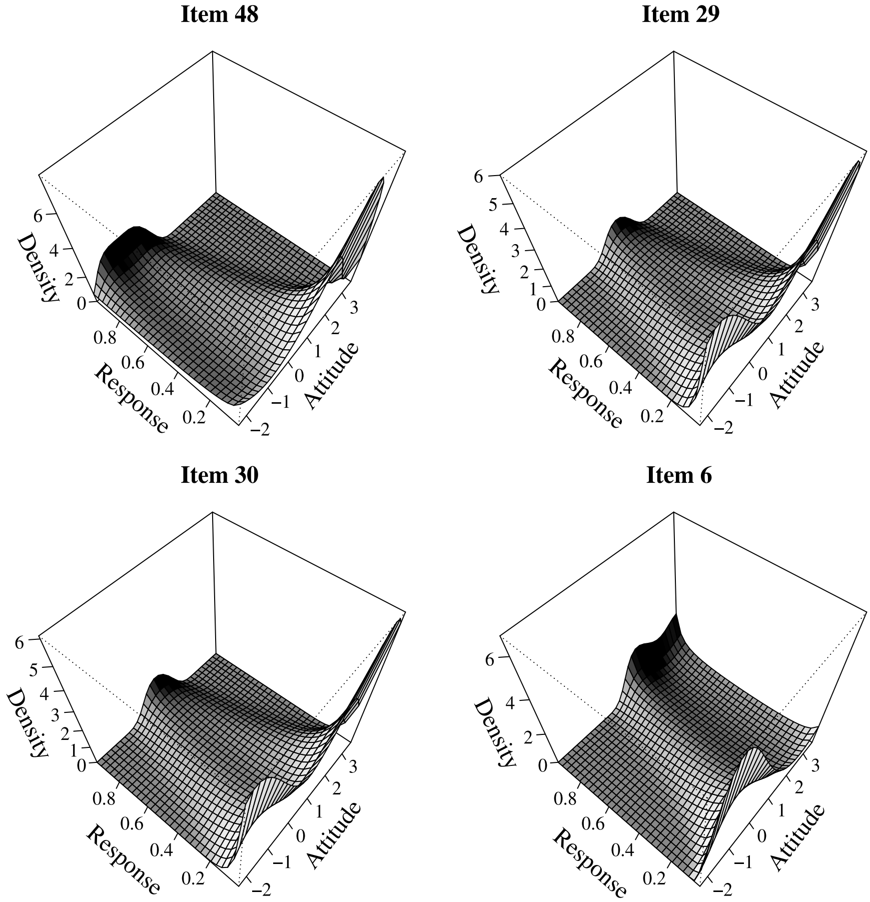

FIGURE 11.
Response density surfaces.

nonpsychologically plausible by these authors (and not apparent on their mood data). The results reported above show that it may indeed appear in attitude data.

7. Discussion

The use of continuous bounded responses (CBR) in psychological measurement may have a number of advantages. A first aspect is parsimony. The number of parameters to be estimated under the BUM favorably compares to the requirements of categorical models in general, which depend upon the chosen number of response categories. At a maximum, $3p-1$ item parameters are to be estimated under the BUM, which is comparable to a Graded Unfolding Model with fixed slopes and four response categories (Roberts et al., 2000a). Item and test information levels in categorical response models are known to increase with the number of response categories available. But the advantages of arbitrarily increasing this number are rapidly neutralized by the need of larger sample sizes to get acceptable levels of estimation quality. This is no more an issue when the set of possible responses is potentially infinite and statistically modeled with a continuous density.

The Beta Unfolding Model for CBR proposed in this paper is a direct extension of Noel and Dauvier's (2007) cumulative BRM. The model is derived both from a hypothetical interpolation response mechanism and from the hypothesis of two opposite sources of item refusal being collapsed. These two sources of refusal have been made explicit in a three-component Dirichlet model and then collapsed to obtain a (two-component) beta response model, taking advantage of the aggregation property of the Gamma. Because its parameters are readily interpretable in terms of acceptance and refusal, the beta is well suited to the modeling of Agree–Disagree responses on a continuous response scale. The reexpression of these distribution parameters in terms of structural parameters (person attitude and item location, dispersion, and acceptability) is also straightforward, through exponential transforms. The resulting expected response function has the same shape as the ones already proposed for binary data by Verhelst and Verstralen (1993) and Andrich and Luo (1993), and this model may be viewed as a generalization of these earlier proposals to the continuous case.

In theory, other beta unfolding models could be constructed in a similar way by defining a suitable U-shaped function of the person-item distance in the n_{ij} refusal parameter. The square (Andrich, 1989) or any power (Hoijtink, 1993) of the (absolute) person-item distance could be used, and the resulting expected response function would have a ratio form quite similar to the one proposed by these authors. However, the choice made in the present paper was to define separate and explicit components for both refusal sources. Besides the fact that it is easily done through a Dirichlet model, given its additivity property, this also leaves open the possibility to define potentially different submodels on the two refusal terms (for instance, one source of refusal may have a stronger impact than the other). This possibility is still to be explored.

The most salient and intriguing aspect of the BUM is the ability to model bimodal response distributions, given θ. More research is needed to see how useful this feature can be in applied studies. But this did happen on all items in the analysis of abortion data reported in this paper: Persons located near the item position tended, for the same attitude value, to provide opposite responses. From the properties of the model, we interpret this result as indicating that subjects in the neighborhood of an item location did not find so strong reasons to either agree or disagree with it (although agreement dominates). For those person and item parameter values, the expected response appears as the least likely. Although counterintuitive, this model feature allows a new distinction between firm and unstable acceptance for the same attitude and observed response. This is potentially measured by the response conditional variance, given the attitude (Equation (15)), and is made possible (by contrast with Bernoulli models) by the fact that the beta has two parameters. We believe these features to be especially meaningful in the field of attitude measurement and, in particular, in the modeling of sudden bifurcation in the choice of a response among ambivalent subjects.

It could be a matter of debate in general whether a bimodal response distribution reveals some kind of chaotic response process, or simply an extreme (dichotomous) response process, the middle range of the response scale simply not being used. In the present case, examination of the abortion items plots in Figure 10 shows that, for the same item, subjects did use the whole range of the response scale at some points of the attitudinal continuum (see Item 47 for example, around $\theta = 0$), and not at others (around $\theta = -1$), so that this phenomenon is attitude-dependent, as predicted by the model. Whatever the interpretation, would this binary response process appear, possibly for some items and not for others, the model would be flexible enough to accommodate both cases simultaneously.

It would probably be easier to illustrate bifurcative response processes in the field of behavior change, where sudden changes are sometimes expected to occur, and explicitly modeled as such (van der Maas & Molenaar, 1992). Preliminary results in the analysis of smoking cessation data (Noel, 2009) strongly suggest that bimodality may be given a substantial interpretation, at least in some contexts. In this research, bimodality appeared specifically for those smokers that were about to quit. This is a strong incentive to explore the response bimodality issue further, in particular, in the domains of behavior change and conflictual attitude measurement.

Appendix

A.1. First Derivatives of the BUM Logdensity Function

The logdensity of a response x_{ij}, given the latent score θ_i and the set Δ_j of item j parameters, is

$$\begin{aligned}\ln f(x_{ij}|\theta_i,\Delta_j) &= \ln\Gamma(m_j+n_{ij}) - \ln\Gamma(m_j) - \ln\Gamma(n_{ij})\\ &\quad + (m_j-1)\ln x_{ij} + (n_{ij}-1)\ln(1-x_{ij})\end{aligned}\tag{A.1}$$

with

$$\begin{cases} m_j = \exp\lambda_j, \\ n_{ij} = \exp\{\alpha(\theta_i-\delta_j)+\tau_j\} + \exp\{-\alpha(\theta_i-\delta_j)+\tau_j\}. \end{cases}$$

For any item-specific parameter η_j, the first derivative is given by

$$\begin{aligned}\frac{\partial \ln f(x_{ij}|\theta_i,\Delta_j)}{\partial\eta_j} &= \Psi(m_j+n_{ij})\left[\frac{\partial m_j}{\partial\eta_j}+\frac{\partial n_{ij}}{\partial\eta_j}\right] - \Psi(m_j)\frac{\partial m_j}{\partial\eta_j} - \Psi(n_{ij})\frac{\partial n_{ij}}{\partial\eta_j}\\ &\quad + \ln x_{ij}\frac{\partial m_j}{\partial\eta_j} + \ln(1-x_{ij})\frac{\partial n_{ij}}{\partial\eta_j}\\ &= \frac{\partial m_j}{\partial\eta_j}\left[\Psi(m_j+n_{ij}) - \Psi(m_j) + \ln x_{ij}\right]\\ &\quad + \frac{\partial n_{ij}}{\partial\eta_j}\left[\Psi(m_j+n_{ij}) - \Psi(n_{ij}) + \ln(1-x_{ij})\right],\end{aligned}\tag{A.2}$$

where $\Psi(x) = \partial\ln\Gamma(x)/\partial x$ is the digamma function.

From (17) we have for the present model:

$$\frac{\partial m_j}{\partial\delta_j} = 0, \qquad \frac{\partial m_j}{\partial\lambda_j} = m_j, \qquad \frac{\partial m_j}{\partial\tau_j} = 0\tag{A.3}$$

and

$$\frac{\partial n_{ij}}{\partial\delta_j} = -2\alpha\phi_j\sinh\left[\alpha(\theta_i-\delta_j)\right], \qquad \frac{\partial n_{ij}}{\partial\lambda_j} = 0, \qquad \frac{\partial n_{ij}}{\partial\tau_j} = n_{ij},\tag{A.4}$$

where $\phi_j = \exp\tau_j$.

Replacing these expressions in (A.2) leads to (29), (30), and (31).

A.2. Second Order Cross Derivatives of the BUM Logdensity Function

The second-order cross derivative of a response logdensity with respect to any pair of item-specific parameters η_j and η'_j is given by

$$\begin{aligned}\frac{\partial^2 \ln f(x_{ij}|\theta_i, \Delta_j)}{\partial\eta_j\partial\eta'_j} &= \frac{\partial^2 m_j}{\partial\eta_j\partial\eta'_j}\big[\Psi(m_j+n_{ij}) - \Psi(m_j) + \ln x_{ij}\big] \\ &\quad + \frac{\partial m_j}{\partial\eta_j}\left[\Psi'(m_j+n_{ij})\left[\frac{\partial m_j}{\partial\eta'_j} + \frac{\partial n_{ij}}{\partial\eta'_j}\right] - \Psi'(m_j)\frac{\partial m_j}{\partial\eta'_j}\right] \\ &\quad + \frac{\partial^2 n_{ij}}{\partial\eta_j\partial\eta'_j}\big[\Psi(m_j+n_{ij}) - \Psi(n_{ij}) + \ln(1-x_{ij})\big] \\ &\quad + \frac{\partial n_{ij}}{\partial\eta_j}\left[\Psi'(m_j+n_{ij})\left[\frac{\partial m_j}{\partial\eta'_j} + \frac{\partial n_{ij}}{\partial\eta'_j}\right] - \Psi'(n_{ij})\frac{\partial n_{ij}}{\partial\eta'_j}\right], \end{aligned} \tag{A.5}$$

where $\Psi'(x) = \partial\Psi(x)/\partial x$ is the trigamma function.

Distinguishing between the three cases where (i) both η and η' only depend upon the acceptability parameter m_j (i.e., only λ_j under this model), (ii) both of them only depend upon the refusal parameter n_{ij} (δ_j and τ_j), and (iii) one of them depends upon m_j and the other one on n_{ij}, the corresponding expressions respectively simplify to

$$\begin{aligned}\frac{\partial^2 \ln f(x_{ij}|\theta_i, \Delta_j)}{\partial\eta_j\partial\eta'_j} &= \frac{\partial^2 m_j}{\partial\eta_j\partial\eta'_j}\big[\Psi(m_j+n_{ij}) - \Psi(m_j) + \ln x_{ij}\big] \\ &\quad + \left(\frac{\partial m_j}{\partial\eta_j}\right)\left(\frac{\partial m_j}{\partial\eta'_j}\right)\big[\Psi'(m_j+n_{ij}) - \Psi'(m_j)\big], \end{aligned} \tag{A.6}$$

$$\begin{aligned}\frac{\partial^2 \ln f(x_{ij}|\theta_i, \Delta_j)}{\partial\eta_j\partial\eta'_j} &= \frac{\partial^2 n_{ij}}{\partial\eta_j\partial\eta'_j}\big[\Psi(m_j+n_{ij}) - \Psi(n_{ij}) + \ln(1-x_{ij})\big] \\ &\quad + \left(\frac{\partial n_{ij}}{\partial\eta_j}\right)\left(\frac{\partial n_{ij}}{\partial\eta'_j}\right)\big[\Psi'(m_j+n_{ij}) - \Psi'(n_{ij})\big], \end{aligned} \tag{A.7}$$

$$\frac{\partial^2 \ln f(x_{ij}|\theta_i, \Delta_j)}{\partial\eta_j\partial\eta'_j} = \left(\frac{\partial m_j}{\partial\eta_j}\right)\left(\frac{\partial n_{ij}}{\partial\eta'_j}\right)\Psi'(m_j+n_{ij}). \tag{A.8}$$

A.3. Expected Values of the Second-Order Cross Derivatives

Expectations of the second derivatives are useful to get approximate standard errors and for estimating the parameters in a Fisher scoring approach. Considering the response variable $X_{ij} \sim \beta(m_j, n_{ij})$ and a known θ_i attitude level for subject i and denoting by L_{ij} the model likelihood function, we know from likelihood theory that

$$E\left[\frac{\partial \ln L_{ij}}{\partial m_j}\right] = 0,$$

$$E\left[\frac{\partial \ln L_{ij}}{\partial n_{ij}}\right] = 0,$$

from which we get:

$$E\big[\Psi(m_j+n_{ij})-\Psi(m_j)+\ln X_{ij}\big]=0,$$
$$E\big[\Psi(m_j+n_{ij})-\Psi(n_{ij})+\ln(1-X_{ij})\big]=0,$$

or

$$E[\ln X_{ij}]=\Psi(m_j)-\Psi(m_j+n_{ij}),$$
$$E\big[\ln(1-X_{ij})\big]=\Psi(n_{ij})-\Psi(m_j+n_{ij}),$$

where the expectation is taken over X_{ij}.

So finally, conditional on the true θ_i attitude value, taking the expectation on both sides of (A.5), we get somewhat simpler expressions for each of the three cases distinguished above:

$$E\left[\frac{\partial^2 \ln f(x_{ij}|\theta_i,\Delta_j)}{\partial\eta_j\partial\eta'_j}\right]=\left(\frac{\partial m_j}{\partial\eta_j}\right)\left(\frac{\partial m_j}{\partial\eta'_j}\right)\big[\Psi'(m_j+n_{ij})-\Psi'(m_j)\big], \tag{A.9}$$

$$E\left[\frac{\partial^2 \ln f(x_{ij}|\theta_i,\Delta_j)}{\partial\eta_j\partial\eta'_j}\right]=\left(\frac{\partial n_{ij}}{\partial\eta_j}\right)\left(\frac{\partial n_{ij}}{\partial\eta'_j}\right)\big[\Psi'(m_j+n_{ij})-\Psi'(n_{ij})\big], \tag{A.10}$$

$$E\left[\frac{\partial^2 \ln f(x_{ij}|\theta_i,\Delta_j)}{\partial\eta_j\partial\eta'_j}\right]=\left(\frac{\partial m_j}{\partial\eta_j}\right)\left(\frac{\partial n_{ij}}{\partial\eta'_j}\right)\Psi'(m_j+n_{ij}). \tag{A.11}$$

Replacing the first derivatives of m_j and n_{ij} with respect to the parameters δ_j, λ_j, and τ_j in the above formulas by their values in (A.3) and (A.4), we get, for the first case,

$$E\left[\frac{\partial^2 \ln f(x_{ij}|\theta_k,\Delta_j)}{\partial\lambda_j^2}\right]=m_j^2\big[\Psi'(m_j+n_{kj})-\Psi'(m_j)\big].$$

For the second case:

$$E\left[\frac{\partial^2 \ln f(x_{ij}|\theta_k,\Delta_j)}{\partial\tau_j^2}\right]=n_{kj}^2\big[\Psi'(m_j+n_{kj})-\Psi'(n_{kj})\big],$$

$$E\left[\frac{\partial^2 \ln f(x_{ij}|\theta_k,\Delta_j)}{\partial\tau_j\partial\delta_j}\right]=-2\alpha\phi_j n_{kj}\sinh\big[\alpha(\theta_k-\delta_j)\big]\big[\Psi'(m_j+n_{kj})-\Psi'(n_{kj})\big],$$

$$E\left[\frac{\partial^2 \ln f(x_{ij}|\theta_k,\Delta_j)}{\partial\delta_j^2}\right]=4\alpha^2\phi_j^2\sinh^2\big[\alpha(\theta_k-\delta_j)\big]\big[\Psi'(m_j+n_{kj})-\Psi'(n_{kj})\big].$$

For the third case:

$$E\left[\frac{\partial^2 \ln f(x_{ij}|\theta_k,\Delta_j)}{\partial\lambda_j\partial\tau_j}\right]=m_j n_{kj}\Psi'(m_j+n_{kj}),$$

$$E\left[\frac{\partial^2 \ln f(x_{ij}|\theta_k,\Delta_j)}{\partial\lambda_j\partial\delta_j}\right]=-2m_j\alpha\phi_j\sinh\big[\alpha(\theta_k-\delta_j)\big]\Psi'(m_j+n_{kj}).$$

These provide (minus) the elements of the information matrix. The standard errors of estimates may be obtained by taking the square root of the diagonal elements of the inverse information matrix.

References

Aitchison, J., & Shen, S. (1980). Logistic-normal distributions: some properties and uses. *Biometrika*, *2*, 261–272.

Andrich, D. (1989). A probabilistic IRT model for unfolding preference data. *Applied Psychological Measurement*, *13*, 193–216.

Andrich, D. (1995). Hyperbolic cosine latent trait models for unfolding direct responses and pairwise preferences. *Applied Psychological Measurement*, *19*, 269–290.

Andrich, D., & Luo, G. (1993). A hyperbolic cosine latent trait model for unfolding dichotomous single-stimulus responses. *Applied Psychological Measurement*, *17*, 253–276.

Andrich, D., & Styles, I.M. (1998). The structural relationship between attitude and behaviour statements from the unfolding perspective. *Psychological Methods*, *3*, 454–469.

Byrd, R.H., Lu, P., Nocedal, J., & Zhu, C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, *16*, 1190–1208.

Coombs, C. (1950). Psychological scaling without a unit of measurement. *Psychological Review*, *57*, 145–158.

Davison, M., Robbins, S., & Swanson, D. (1978). Stage structure in objective moral judgments. *Developmental Psychology*, *14*, 137–146.

DeMars, C.E., & Erwin, T.D. (2003). Revising the scale of intellectual development: application of an unfolding model. *Journal of College Student Development*, *44*, 168–184.

Ferrari, S.L.P., & Cribari-Neto, F. (2004). Beta regression for modeling rates and proportions. *Journal of Applied Statistics*, *10*, 1–18.

Greenacre, M., & Browne, M. (1986). An efficient alternating least-squares algorithm to perform multidimensional unfolding. *Psychometrika*, *51*, 241–250.

Gupta, R.D., & Kundu, D.K. (1999). Generalized exponential distributions. *Australian & New Zealand Journal of Statistics*, *41*, 173–188.

Herrnstein, R. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, *4*, 267–272.

Hoijtink, H. (1993). PARELLA: a parametric approach to parallelogram analysis. *Quantitative Methods*, *43*, 93–107.

Kotz, S., & Johnson, N. (1982). Beta distributions. In S. Kotz & N. Johnson (Eds.), *Encyclopedia of statistical sciences* (Vol. 1, pp. 228–229). New York: Wiley.

Luce, R.D. (1959). *Individual choice behavior*. New York: Wiley.

McKelvie, S. (1978). Graphic rating scales: how many categories? *British Journal of Psychology*, *69*, 185–202.

Mislevy, R. (1984). Estimating latent distribution. *Psychometrika*, *49*, 359–381.

Noel, Y. (1999a). An ordination approach to the quantification of smokers' readiness to change. *Swiss Journal of Psychology*, *58*, 123–133.

Noel, Y. (1999b). Recovering latent unimodal patterns of change by unfolding analysis: application to smoking cessation. *Psychological Methods*, *4*, 173–191.

Noel, Y. (2009, August). *When extreme responses are substantial: a generalized beta response model of behavior change*. Paper presented at the 42nd Annual Meeting of the Society for Mathematical Psychology, Amsterdam, The Netherlands.

Noel, Y., & Dauvier, B. (2007). A beta item response model for continuous bounded responses. *Applied Psychological Measurement*, *31*, 47–73.

R Core Team (2013). *R: a language and environment for statistical computing* [Computer software manual]. Vienna, Austria. Available from http://www.R-project.org/ (ISBN 3-900051-07-0).

Ramsay, J. (1980). The joint analysis of direct ratings, pairwise preferences, and dissimilarities. *Psychometrika*, *45*, 149–165.

Roberts, J.S., Donoghue, J.R., & Laughlin, J.E. (1998). *The generalized graded unfolding model: a general parametric item response model for unfolding graded responses* (RR-98-32). Princeton: Educational Testing Service.

Roberts, J.S., Donoghue, J.R., & Laughlin, J.E. (2000a). A general item response theory model for unfolding unidimensional polytomous responses. *Applied Psychological Measurement*, *24*, 3–32.

Roberts, J.S., & Laughlin, J.E. (1996). A unidimensional item response model for unfolding responses from a graded disagree-agree response scale. *Applied Psychological Measurement*, *20*, 231–255.

Roberts, J.S., Lin, Y., & Laughlin, J.E. (2001). Computerized adaptive testing with the generalized graded unfolding model. *Applied Psychological Measurement*, *25*, 177–196.

Roberts, J.S., Rost, J., & Macready, G.B. (2000b). *An unfolding mixture model to explore the latitude of acceptance concept in attitude measurement*. Invited paper presented at the Fifth International Conference on Social Science Methodology, Cologne, Germany.

Samejima, F. (1973). Homogeneous case of the continuous response model. *Psychometrika*, *38*, 203–219.

Samejima, F. (1974). Normal ogive model on the continuous response level in the multidimensional latent space. *Psychometrika*, *39*, 111–121.

van Schuur, W.H. (1993). Nonparametric unidimensional unfolding for multicategory data. *Political Analysis*, *4*, 41–74.

Smithson, M., & Verkuilen, J. (2005). Beta regression: practical issues in estimation (supplementary material). Retrieved from http://supp.apa.org/psycarticles/supplemental/met_11_1_54/met_smithson1631_supplementalmaterial.doc.

Smithson, M., & Verkuilen, J. (2006). A better Lemon-Squeezer? Maximum likelihood regression with beta-distributed dependent variables. *Psychological Methods*, *11*, 54–71.

Tamhane, A.C., Ankenman, B., & Yang, Y. (2002). The beta distribution as a latent response model for ordinal data (I): estimation of location and dispersion parameters. *Journal of Statistical Computation and Simulation*, *72*, 473–494.

Thurstone, L. (1927). A law of comparative judgement. *Psychological Review*, *34*, 278–286.
van der Maas, H., & Molenaar, P. (1992). Stagewise cognitive development: an application of catastrophe theory. *Psychological Review*, *99*, 395–417.
Verhelst, H., & Verstralen, H. (1993). A stochastic unfolding model derived from the partial credit model. *Quantitative Methods*, *43*, 73–92.
Verkuilen, J., & Smithson, M. (2012). Mixed and mixture regression models for continuous bounded responses using the beta distribution. *Journal of Educational and Behavioral Statistics*, *37*, 82–113.
Volet, S.E., & Chalmers, D. (1992). Investigation of qualitative differences in university students' learning goals based on an unfolding model of stage development. *British Journal of Educational Psychology*, *62*, 17–34.
Wang, T., & Zeng, L. (1998). Item parameter estimation for a continuous response model using an EM algorithm. *Applied Psychological Measurement*, *22*, 333–344.
Woodruff, D., & Hanson, B. (1996). *Estimation of item response models using the EM algorithm for finite mixtures*. (ACT research report series No. 96-6). Iowa City, IA: ACT, Inc.
Wright, B., & Masters, G. (1982). *Rating scale analysis*. Chicago: MESA Press.

Manuscript Received: 11 MAR 2008
Published Online Date: 11 DEC 2013