Bayesian Approaches to Polyhedral Conjoint Analysis

Srikant Vadali
John Liechty
Arvind Rangaswamy

January 2009

Srikant Vadali is Assistant Professor at the University of Massachusetts, Lowell. His contact details are- Mail: FA207F, One University Avenue, Lowell, MA 01854, Email: Srikant_Vadali@uml.edu, Tel: (978) 934 2810, Fax (978) 934 2016.

John Liechty is Associate Professor of Marketing and Statistics at The Pennsylvania State University. His contact details are- Mail: 701 M BAB, University Park, PA 16803, Email: jcl12@psu.edu, Tel: (814) 865-0621

Arvind Rangaswamy is the Anchel Professor of Marketing at The Pennsylvania State University. His contact details are- 481 Business Building, The Smeal College of Business, Penn State University, University Park, PA 16802, arvindr@psu.edu, Tel: (814) 865 1907, Fax: (814) 865 3015.

Bayesian Approaches to Polyhedral Conjoint Analysis

Abstract

We develop probability models for FASTPACE (Toubia et al. 2003), an approach to Conjoint Analysis via polyhedral optimization techniques. First, we develop a theoretical foundation for FASTPACE by formulating a probability model which can be used to compute estimates identical to FASTPACE estimates. Next, via an appropriate prior for the probability model, we develop Bayesian FASTPACE and further extend it by incorporating shrinkage (i.e., hierarchical version of Bayesian FASTPACE). The resulting model called GENPACE nests FASTPACE, Bayesian FASTPACE, and Hierarchical Bayes models.

We investigate the performance of the competing models in the context of partial profile ratings-based Conjoint Analysis. Our simulations show that Bayesian FASTPACE results in 8% improvement in partworth recovery, on average, as compared to FASTPACE. Bayesian FASTPACE also has a smaller out-of-sample prediction error than FASTPACE in a real-world data set. Finally, our simulations show that GENPACE performance improves with the availability of more information, but it may exhibit poorer relative performance with inadequate levels of information (e.g., few questions asked of each respondent). We conclude by suggesting avenues for further research.

Keywords: Conjoint Analysis, Heterogeneity, Measurement and Inference, Probability Models, Simulation

INTRODUCTION

Conjoint analysis is a widely used method for obtaining empirical estimates of individual-level utility functions. For a review of conjoint analysis, see Green, Krieger, and Wind (2001). In recent years, there have been continuing efforts to enhance conjoint analysis methods and their impact. Such efforts include varying the experimental setting (see Ding, Grewal and Liechty 2005), developing methods to capture how consumers impute missing levels when presented with partial profiles (Bradlow, Hu and Ho, 2004), and accommodating heterogeneity with different hierarchical model structures (see Andrews, Ainslie and Currim, 2002 and Rossi, McCulloch and Allenby, 1996). In all of the different variations of conjoint analysis, the one constant is that potential consumers are asked to evaluate a set of product profiles that are described in terms of product attributes. Specifically, they are asked to either rate or rank a set of product profiles or to repeatedly choose between a small set of products.

The accuracy of data collected from a conjoint study would be high if individuals respond in a manner consistent with their true preferences over the entire set of products to which they respond; however, experience has demonstrated that is typically not the case. The process of preference construction during the course of a conjoint task suggests that an individual’s ratings may contain errors and that preferences may even evolve over the course of the study. Errors or inconsistent answers may in fact occur often (see Liechty, DeSarbo and Fong (2005) for a discussion). Models underlying conjoint analysis deal with inconsistent answers by incorporating an error structure when estimating partworths. For ratings-based conjoint data, which is the setting that we explore in this paper, two ways to incorporate error structures are the polyhedral-based models, which introduce the minimal amount of error required to resolve any inconsistencies, see Toubia et al. (2003), and the traditional Hierarchical Bayesian (HB) regression models, which minimize a squared error function (Lenk et al., 1996).

Toubia et al. (2003) proposed FASTPACE, a polyhedral optimization method for individual-level adaptive conjoint design and for individual partworth (utility) estimation in conjoint analysis. In adaptive conjoint, the product profiles provided to a respondent to rate are based on the ratings given by that respondent for the previous product profiles (questions). FASTPACE structures the questioning sequence such that the size of the feasible partworth space is reduced as rapidly as possible. The partworth estimate proposed by Toubia et al. (2003) is the analytic center of the space of the feasible partworths. When answers provided by respondents are consistent, the set of feasible partworths is non-empty, and the analytic center can be computed directly. However, if at any time during the questioning process the partworth space becomes infeasible (i.e., the partworth space is empty), FASTPACE introduces the minimal amount of error required to obtain a feasible partworth space which allows for the computation of the analytic center. The justification for the introduction of the error is to accommodate the preference inconsistencies of the respondent. A weakness of the FASTPACE model is the lack of an explicitly stated probability model for the error (see the discussion regarding “Error Theory” in Toubia et al (2003), pp 296-297)[1].

The major contribution of this paper is to present a probability model (and an associated response error distribution) which can be used to generate estimates consistent with the partworth estimates of FASTPACE. By using a probability model that is consistent with FASTPACE, researchers would be able to perform model selection, hypothesis testing and other statistically rigorous analysis of partworths. Using appropriate priors for the parameters of the proposed probability model, we develop Bayesian FASTPACE which offers the ability to estimate partworths via the mean of the posterior partworth distribution. Our results show that, on average, Bayesian FASTPACE outperforms FASTPACE in partworth recovery thus illustrating the superiority of the Bayesian estimator over the analytic center estimator in the context of FASTPACE.

FASTPACE offers the advantage of being able to estimate the complete partworth function at any stage of the question sequence, where the accuracy of the estimates improves as the number of questions increases. A natural alternative to FASTPACE is the HB regression model, which is based on a probability model with errors that are normally distributed. HB regression "pools information" across individuals allowing for estimates of population partworth estimates (along with an associated heterogeneity distribution to account for individual-level variations) even when there is a relatively small amount of data for each respondent.

A second contribution of this paper is that using the HB approach, we develop a unified hierarchical model, called Generalized Polyhedral Adaptive Conjoint Estimation (GENPACE), which nests Bayesian FASTPACE and HB regression under one model. In addition to nesting Bayesian FASTPACE and HB regression, GENPACE also nest other models including HB FASTPACE (which allows for pooling of information across individuals) and constrained HB regression (where we impose restrictions on partworths). From a theoretical perspective, GENPACE illustrates the relationships between the various models. From an empirical perspective, GENPACE can offer improved model performance in certain identifiable situations.

In summary, the main contributions of our research are the following: (a) we develop a probability model for FASTPACE thus embedding it within a sound theoretical foundation, (b) we demonstrate that Bayesian FASTPACE generally outperforms FASTPACE in partworth recovery, (c) we extend the Bayesian FASTPACE model to the GENPACE model, which nests

Bayesian FASTPACE and HB regression (and several variants of these models), and (d) we demonstrate that GENPACE performs significantly better than FASTPACE and the other models in certain specific situations. In this paper, we focus our discussion on estimation issues (i.e., obtaining estimates of partworths given a set of questions and participant responses) and leave for future research the issue of how to dynamically generate an adaptive sequence of questions using Bayesian FASTPACE and GENPACE.

In the next section, we identify the probability model which can be used to compute estimates consistent with that of FASTPACE and then develop the Bayesian FASTPACE model. Subsequently, we develop a unifying model, termed GENPACE, within which Bayesian FASTPACE and HB Regression can be nested. We then assess the performance of Bayesian FASTPACE and GENPACE relative to competing models using two simulations studies and using the experimental laptop data collected by Toubia et al. (2003). Finally, we summarize our main findings, note the limitations of our study, and point to some future directions for research.

BAYESIAN FASTPACE

In this section, we first identify a probability model called Bayesian FASTPACE that is consistent with the implicit assumptions in FASTPACE. Subsequently, we analyze how the maximum likelihood estimates of the proposed probability model relate to FASTPACE estimates.

Our proposed probability model formalizes the original FASTPACE model “wrapping it” within a probability framework. In a sense, our modeling approach continues a long tradition in the academic literature of proposing probabilistic frameworks for “ad hoc” estimation algorithms. In many research contexts, an optimization model is first proposed as something that works well in practice, but without clear explanation or theory as to why it works well, or with regards to how far the algorithm can be extended. Probably the most famous example goes back to when Gauss (1809) proposed the Normal (Gaussian) distribution in connection with his use of the ordinary least squares method which was originally proposed in Legendre (1805). The obvious impact of this connection was quickly felt and it helped contribute to the foundations of modern Probability Theory and Statistics. In a similar way (though not of the same significance), the proposed probability model provides a stronger foundation for the FASTPACE estimation algorithm.

We begin by reviewing the partworth estimation algorithm of FASTPACE. (See Toubia et al (2003) for a discussion of the FASTPACE question design algorithm and for further details about the associated partworth estimation process). FASTPACE starts by asking a respondent to state which of two levels within each attribute is the least preferred. As the model is individual-specific, without loss of generality, the partworth (utility) of the least preferred level of each attribute for a respondent is set to zero. To accommodate the optimization-based approach to inference and to fix the interval-scale of the partworth function, the model restricts each partworth to be less than 100. Thus, FASTPACE bounds partworths between 0 and 100. Unlike statistical approaches (such as HB Regression, Ordinary Least Squares), FASTPACE assumes there is no response error unless the answers of the respondent exhibit inconsistencies. In other words, FASTPACE starts by assuming there is no response error in the answers provided by a respondent. This is illustrated by the fact that in equation 1, there are no errors.

$$a_i = X_i u_i \tag{1}$$

where,

X_i: is the $n \times p$ matrix that corresponds to the n questions presented to the i^{th} respondent,

a_i: is a $n \times 1$ vector of the i^{th} respondent's answers.

u_i: is the partworth for respondent i.

When there are no errors, we can represent the set of partworths that is consistent with equation 1 and the fact that partworths are bounded between 0 and 100 as follows:

(2) $$S(a_i, X_i) = \{ u_i \mid a_i = X_i u_i, 0e \le u_i \le 100e \}$$

Where,

e: is a vector of ones of appropriate dimension[2].

However, if the set $S(a_i, X_i)$ is empty then FASTPACE estimates the least amount of response error that needs to be introduced to obtain a feasible partworth space. Specifically, the following optimization program allows one to obtain a feasible partworth space:

(3) min $$\delta_i$$

(3.1) subject to $$a_i - \delta_i e \le X_i u_i \le a_i + \delta_i e \quad ,$$

(3.2) $$0e \le u_i \le 100e$$

(3.3) $$\delta_i \ge 0 .$$

Where,

δ_i: is response error in the answers of the respondent.

The optimization program in (3) will yield a unique solution for δ_i. However**,** there *may* not be a unique solution for u_i because not *all* constraints in 3.1 and 3.2 may be binding. Since questions are partial profile, some of the partworth variables in u_i may not have a unique solution if they do not appear in constraints that bind at optimality. Therefore, the feasible set of partworths at optimality becomes:

(4) $$\tilde{S}(a_i, X_i) = \left\{ u_i \mid X_i u_i - \delta_i^{FP} \le a_i \le X_i u_i + \delta_i^{FP}, 0e \le u_i \le 100e \right\}$$

where,

δ_i^{FP}: is the optimal solution determined by FASTPACE using the optimization program in (3).

To obtain a unique point estimate for partworths (u_i), FASTPACE uses the analytic center of $\tilde{S}(a_i, X_i)$, which assumes that all points in $\tilde{S}(a_i, X_i)$ are equally likely (See Toubia et al (2003) p. 278). The probability model we develop for FASTPACE is based on this idea that all points in $\tilde{S}(a_i, X_i)$ are equally likely.

The intuition behind our proposed probability model is as follows: The optimization program in (3) and the resulting partworth set in (4) are together sufficient to also describe partworth estimation even when no response error is introduced into the estimation process. Specifically, if the partworth space, defined by $a_i = X_i u_i$, is non-empty then the optimization program (3) will yield $\delta_i^{FP} = 0$ and the set $\tilde{S}(a_i, X_i)$ is then identical to $S(a_i, X_i)$. Thus, the optimization program (3), the partworth set (4), and the choice of the analytic center as a point estimate together represent a complete description of the FASTPACE approach to partworth estimation. Therefore, if there exists a posterior distribution of partworths, it must be uniform on $\tilde{S}(a_i, X_i)$. If we assume a uniform prior for the partworths over the region $0e \le u_i \le 100e$, then the likelihood is a uniform distribution in order for the posterior distribution to be uniform. This suggests that the implied FASTPACE probability model has uniformly distributed errors with specific bounds.

The imposition of bounds on response errors is consistent with the notion that there is an overall tendency among humans to exhibit preference consistency. Research suggests there are strong evolutionary reasons for such consistency (Lee, Amir, and Ariely 2006; Gintis 2007, p. 10), although there are differences among individuals in the degree of preference consistency (Stanovich et al., 2003). Thus, it is reasonable to propose a utility model where the errors that people make in preference judgments are bounded, rather than unbounded. We now formalize our intuition. We assume that response errors in the answers of a respondent are distributed uniformly between symmetric lower and upper bounds. In other words let,

(5) $$a_i = X_i u_i + \varepsilon_i$$

where,

(6) $$\varepsilon_i \sim U(-\delta_i e, \delta_i e) \quad .$$

$\delta_i > 0$ is the ***maximum possible error*** made by the respondent.

(Note: To reduce notational clutter, we "re-use" the symbol δ_i from the linear program in (3) above with a new interpretation that it represents the maximum error.)

The resulting probability density is:

(7) $$f(a_i \mid u_i, \delta_i, X_i) = \frac{1}{k_i} I\left(X_i u_i - \delta_i e \le a_i \le X_i u_i + \delta_i e\right) \qquad ,$$

where,

$I(\cdot)$ is the indicator function and

$k_i = \int_{X_i u_i - \delta_i e}^{X_i u_i + \delta_i e} da_i = (2\delta_i)^n$ is the normalizing constant.

Given a set of questions and answers from a respondent, the likelihood equation for the model parameters is:

(8) $$L(u_i, \delta_i \mid a_i, X_i) = \frac{I\left(X_i u_i - \delta_i e \le a_i \le X_i u_i + \delta_i e\right)}{\left(2\,\delta_i\right)^n}$$

Maximizing the above likelihood is equivalent to minimizing δ_i subject to the constraints $I\left(X_i u_i - \delta_i e \le a_i \le X_i u_i + \delta_i e\right)$, $I\left(0e \le u_i \le 100e\right)$ and $\delta_i > 0$. Notice that the optimization program to find the maximum likelihood estimate is identical to the FASTPACE optimization algorithm outlined in (3). Therefore, $\delta_i^{MLE} = \delta_i^{FP}$ and the space of possible partworth MLEs at optimality is identical to $\tilde{S}(a_i, X_i)$. As we discussed in the context of the FASTPACE optimization algorithm, the set $\tilde{S}(a_i, X_i)$ may not be a singleton i.e., we may have multiple solutions for the partworth vector, all having the same likelihood value. Consequently, we need an additional mechanism to compute a point estimate for the partworth vector[3].

There exist at least two approaches to compute a ‘reasonable’ point estimate from among all feasible partworth vectors in $\tilde{S}(a_i, X_i)$. The first approach is to mimic FASTPACE and use the analytic center and compute point estimates. Thus, we can recover FASTPACE estimates using our probability model, which validates our claim that the probability model identified in equations (5) and (6) represents an underlying probability framework for FASTPACE. The second approach is to extend the maximum likelihood framework to a Bayesian context and use the posterior mean of the partworth distribution as a point estimate. Extending the above framework to the Bayesian context would require a suitable assumption for the priors of u_i and δ_i. Equations (9) and (10) give the full conditional posterior distributions for u_i and δ_i if we choose very uninformative priors (e.g., an uniform distribution over $I\left(0e \le u_i \le 100e\right)$ for u_i and uninformative gamma prior for δ_i)

(9) $$f(u_i \mid -) \;\propto\; I\left(a_i - \delta_i e \le X_i u_i \le a_i + \delta_i e\right) I\left(0e \le u_i \le 100e\right)$$

$$f(\delta_i \mid -) \quad \propto \quad \left(\frac{1}{\delta_i}\right)^{sp_M - 1} \exp\left(\frac{-sc_\Delta}{\delta_i}\right) \qquad , \qquad (10)$$
$$I(\delta_i e \geq a_i - X_i u_i)\, I(\delta_i e \geq X_i u_i - a_i)\, I(\delta_i e \geq 0_i)$$

To see the difference between the analytic center and the Bayesian approaches to computing point estimates, consider the full conditional posterior distribution for δ_i as given in equation (10) above. Notice that the support of the posterior distribution is given by $I(\delta_i e \geq a_i - X_i u_i)\, I(\delta_i e \geq X_i u_i - a_i)\, I(\delta_i e \geq 0_i)$ conditioned on u_i. Therefore, the lower bound of the *unconditional* posterior distribution for δ_i can be computed by minimizing δ_i subject to the constraints $I1(\delta_i e \geq a_i - X_i u_i)\, I(\delta_i e \geq X_i u_i - a_i)\, I(\delta_i e \geq 0)$ and $I(0e \leq u_i \leq 100e)$. We include the restriction $I(0e \leq u_i \leq 100e)$ in the optimization program as the support of the posterior distribution in equation (10) is conditioned on u_i. The optimization program we just formulated to find the lower bound of the *unconditional* distribution for δ_i is identical to the FASTPACE optimization algorithm outlined in equation (3) above. Thus, the lower bound of the *unconditional* posterior distribution equals δ_i^{FP}. Therefore, the analytic center approach to computing point estimates computes the least possible error bound (i.e., δ_i^{FP}) and then uses the concept of the analytic center to compute a point estimate for the partworth vector given δ_i^{FP}. In contrast, the Bayesian approach to computing point estimates for the partworths involves sampling over the entire support of the posterior distribution for u_i and δ_i.

In our empirical analysis, we explore the impact of estimating u_i and δ_i via their respective posterior means, instead of using FASTPACE point estimates (i.e., estimates based on the analytic center) on partworth recovery and out-of-sample predictions.

GENPACE

In this section, we develop the GENPACE model by incorporating elements from both Bayesian FASTPACE and HB Regression models. We start by noting three points of difference between Bayesian FASTPACE and HB Regression models. First, Bayesian FASTPACE assumes that the response error distribution is a bounded uniform distribution, whereas HB Regression assumes that response error is an unbounded normal distribution. Second, Bayesian FASTPACE imposes the restriction that partworths are bounded between 0 and 100 whereas HB Regression does not impose any restrictions on partworths. Third, Bayesian FASTPACE does not incorporate partworth shrinkage whereas HB Regression shrinks partworths toward the population mean. We develop GENPACE by incorporating the appropriate assumptions that serve to provide a common framework within which the differences between the various models can be accommodated. We start by assuming that the response error has a bounded normal distribution:

$$(11) \qquad \varepsilon_i \sim N(0,\sigma_i^2 I)\; I(-\delta_i e \le \varepsilon_i \le \delta_i e)$$

We emphasize two points with respect to the above assumption: (1) The response error distribution in (11) reduces to a bounded uniform distribution if we set $\sigma_i^2 \to \infty$; a distribution that we used to develop Bayesian FASTPACE (see equation 6). (2) If we set $\delta_i \to \infty$ then the distribution reduces to an unbounded normal distribution, and it is consistent with HB regression. Thus, the bounded, normal response error distribution outlined above is a generalization that accommodates both Bayesian FASTPACE and HB regression in one framework. We accommodate the difference in partworth restrictions (“bounded between 0 and 100” vs “unrestricted”) by imposing restrictions on partworths as follows:

$$lb \le u_i \le ub \quad . \tag{12}$$

We can accommodate FASTPACE partworth restrictions (i.e., the restriction that partworths are bounded between 0 and 100) by setting $lb = 0e$ and $ub = 100e$. Setting $lb = -\infty$ and $ub = \infty$ makes the partworths unbounded thus accommodating HB regression. Finally, we accommodate the option to shrink partworths by assuming that

$$u_i \sim N(\bar{u}, \Sigma)\, I(lb \le u_i \le ub) \quad . \tag{13}$$

If we do not want to apply shrinkage to the partworths, then we can set $\Sigma \to \infty$. Incorporating assumptions (11), (12) and (13) and selecting suitable priors as given below completes the GENPACE specification:

Data Generating Process

$$a_i = X_i u_i + \varepsilon_i \tag{14}$$

$$\varepsilon_i \sim N(0, \sigma_i^2 I)\; I(-\delta_i e \le \varepsilon_i \le \delta_i e) \tag{15}$$

Shrinkage Specification

$$u_i \sim N(\bar{u}, \Sigma)\, I(lb \le u_i \le ub) \tag{16}$$

$$\sigma_i^{\,2} \sim N(\overline{\sigma^2}, \tau^2)\; I(\sigma_i^{\,2} > 0) \tag{17}$$

Priors

$$\Sigma^{-1} \sim Wishart(R^{-1}, d) \tag{18}$$

$$\bar{u} \sim N(\bar{u}_{pr}, v^2 I) \;\; , \tag{19}$$

$$\overline{\sigma^2} \sim N(\overline{\overline{\sigma^2}}, \tau^2) \tag{20}$$

$$\delta_i \sim Gamma(sc_{\delta_i}, sp_{\delta_i}) \quad , \tag{21}$$

The full conditional densities for the GENPACE model are:

$$\begin{gathered} f(u_i \mid -) \sim \quad N\left(\left[\sigma_i^{-2} X_i^T X_i + \Sigma^{-1}\right]^{-1}\left[\sigma_i^{-2} X_i^T a_i + \Sigma^{-1}\overline{u}\right], \left[\sigma_i^{-2} X_i^T X_i + \Sigma^{-1}\right]^{-1}\right) \\ I(a_i - \delta_i e \le X_i u_i \le a_i + \delta_i e)\, I(lb \le u_i \le ub), \end{gathered} \tag{22}$$

$$f(\overline{u} \mid -) \propto \exp\left(-\frac{1}{2v^2}(\overline{u} - \overline{u}_{pr})^T(\overline{u} - \overline{u}_{pr})\right)\prod_i K_1(\overline{u}, \Sigma)\exp\left(-\frac{1}{2}(u_i - \overline{u})^T \Sigma^{-1}(u_i - \overline{u})\right) \tag{23}$$

$$f\left(\Sigma^{-1} \mid -\right) \propto \left|\Sigma^{-1}\right|^{\frac{n+d-p-1}{2}} \exp\left(-tr\left(\left(R + \sum (u_i - \overline{u})(u_i - \overline{u})^T\right)\Sigma^{-1}\right)\right)\prod_i K_1(\overline{u}, \Sigma) \quad , \tag{24}$$

$$f\left(\sigma_i^2 \mid -\right) \propto K_2\left(\delta_i, \sigma_i^2\right)\exp\left(-\frac{\left(\sigma_i^2 - \overline{\sigma^2}\right)^2}{2\varpi^2}\right)\exp\left(-\frac{(a_i - X_i u_i)^T(a_i - X_i u_i)}{2\sigma_i^2} - \frac{sc_\sigma}{\sigma_i^2}\right) \tag{25}$$

$$f\left(\overline{\sigma^2} \mid -\right) \propto \prod_i K_3\left(\tau^2, \overline{\sigma^2}\right)\exp\left(-\frac{\left(\sigma_i^2 - \overline{\sigma^2}\right)^2}{2\varpi^2}\right) \tag{26}$$

$$f(\delta_i \mid -) \propto K_2(\delta_i, \sigma_i^2)\left(\frac{1}{\delta_i}\right)^{sp_M - 1}\exp\left(\frac{-sc_\Delta}{\delta_i}\right) I(a_i - \delta_i e \le X_i u_i \le a_i + \delta_i e) \quad , \tag{27}$$

Where,

$$K_1(\overline{u}, \Sigma) = \int_{\{u_i : I(lb \le u_i \le ub)\}} \frac{\exp\left(-0.5(u_i - \overline{u})^T \Sigma^{-1}(u_i - \overline{u})\right)}{(2\pi)^{0.5p}\left|\Sigma\right|^{0.5}} du_i \tag{28}$$

$$K_2\left(\delta_i, \sigma_i^2\right) = \left(\Phi\left(\frac{\delta_i}{\sigma_i}\right) - \Phi\left(-\frac{\delta_i}{\sigma_i}\right)\right)^{-n} \quad , \tag{29}$$

$$K_3\left(\tau^2, \overline{\sigma^2}\right) = \left(1 - \Phi\left(-\frac{\overline{\sigma^2}}{\tau}\right)\right)^{-1} \tag{30}$$

(Note: In equations (29) and (30), $\Phi(\ )$ represents the cdf of a univariate normal with mean 0 and variance 1. Also, the normalizing constant specified in (28) is a multivariate normal cumulative distribution function. We use an algorithm developed by Genz (1992) to calculate this constant).

Estimating the parameters of GENPACE is non-trivial. As pointed out by Boatwright et al. (1999), the parameters of the normal truncated distribution is technically identified but can be empirically unidentified, especially with a small dataset. Thus, the priors for the parameters in GENPACE have to be specified carefully in estimating GENPACE. The appendix summarizes the technical details covering such topics as specification of prior, reasons for shrinking response error variance, and estimation strategies for GENPACE. In the appendix, we explain the similarities and differences between our approach for dealing with issues of empirical identification with the approach proposed by Boatwright et al (1999). We also provide details regarding how we implemented GENPACE to assess its relative performance in simulations as well as in an empirical application using the data collected by Toubia et al (2003).

Table 1 and Figure 1 summarize the parameter restrictions that need to be imposed in GENPACE to recover estimates from different models: Table 1 provides a summary of the full conditional densities for the partworths and Figure 1 give the summaries in terms of probability model and priors[4].

[Insert Table 1 here]

[Insert Figure 1 here]

THE EMPIRICAL PERFORMANCE OF BAYESIAN FASTPACE AND GENPACE

We have two main goals in evaluating the empirical performance of the Bayesian FASTPACE and GENPACE models. First, we want to assess the potential benefits of using Bayesian posterior mean, instead of analytic center, for the partworth point estimates. Understanding the relative value of Bayesian approach versus the analytic center approach is valuable for both theory and practice. Second, we want to assess the conditions under which the added complexity of GENPACE is necessary to recover partworths accurately. While the theoretical nesting of various models such as FASTPACE, Bayesian FASTPACE, HB Regression etc within GENPACE is appealing, it is also important to assess the empirical performance of GENPACE as compared to simpler models.

We evaluated the empirical performance of FASTPACE, Bayesian FASTPACE, HB Regression, Constrained HB Regression and GENPACE using two simulations and also using the laptop data collected by Toubia et al (2003). Table 2 summarizes the key properties of the various models we evaluated:

[Insert Table 2 here]

In the first simulation, we explore the relative performance of the different models in recovering the true partworths. In the second simulation, we do a more detailed assessment of GENPACE's ability to recover the true parameters (including δ_i and partworths) under a broader set of conditions than in simulation 1, so that we can explore the conditions under which the added complexity of GENPACE is needed to recover partworths accurately. Finally, we present and discuss the results from applying the models to the laptop dataset.

Simulation 1: Relative Performance of FASTPACE, Bayesian FASTPACE, and The Other Competing Models

To compare the empirical performance of the various models summarized in Table 2, we used an experimental design that is similar to the one reported in Toubia et al. (2003). We generated questions for the simulations using FASTPACE[5] for 100 respondents for a product with 10 attributes and 2 levels per attribute. We varied the number of questions ($n = 8, 16, 32, \text{and } 100$), the partworth heterogeneity ($\Sigma = 100I$ or $900\ I$) and the standard deviation of response error distribution ($\sigma_i = 20 \text{ or } 40$). The size of the bound on response error was set to 1 standard deviation i.e., we set $\delta_i = \sigma$. Tables 3 and 4 summarize our experimental design.

[Insert Table 3 here]

[Insert Table 4 here]

Thus, we had 16 experimental conditions and estimated parameters for 20 sets of synthetic data for each experimental condition. Table 5 summarizes the simulation results. Specifically, we report the RMSE average (across the 20 replications) and the corresponding standard deviation for partworths for all the competing models.

[Insert Table 5 here]

We can draw the following conclusions based on the simulation results in Table 5.

Conclusion 1: Bayesian FASTPACE recovers partworths better than FASTPACE. The proposed Bayesian FASTPACE model offers, on average, a performance gain of about 8% compared to the original FASTPACE model.

There is a simple explanation for the superior results from Bayesian FASTPACE. As we discussed earlier, FASTPACE finds a corner solution for $\boldsymbol{\delta}_i$. Therefore, when the true response error is a bounded distribution, FASTPACE's estimate of $\boldsymbol{\delta}_i$ is typically an underestimate of the true bound which hurts partworth recovery. In contrast, Bayesian FASTPACE samples over the entire posterior distribution of $\boldsymbol{\delta}_i$ resulting in a better estimate of the bound on response error which helps partworth recovery. The above finding illustrates the superiority of Bayesian estimation as compared to analytic center estimates which is an interesting conclusion from a theoretical perspective. From a practice perspective, our results indicate that it is better to use the Bayesian FASTPACE to estimate partworths, instead of FASTPACE, at least when response errors are bounded.

Conclusion 2: GENPACE is better at partworth recovery than ***all the competing models*** when respondents are relatively homogeneous and response error is high, and this performance improvement occurs irrespective of the number of questions asked (low, moderate or high).

This finding suggests that, under specific conditions, there is an incremental benefit to using GENPACE instead of the competing models such as FASTPACE, Bayesian FASTPACE, HB Regression, or Constrained HB Regression. This result is useful for applications because in most situations we are able to ask each respondent only a few questions, due to limited time of surveys and potential respondent fatigue.

Conclusion 3: When respondents are heterogeneous, or when respondents are homogeneous and response error is low, GENPACE performs relatively poorly as compared to Bayesian FASTPACE in situations with fewer questions. However, GENPACE dominates Bayesian FASTPACE as the number of questions becomes large (100 in our simulation).

The reason for the relative poor performance of GENPACE when we have fewer questions is that often GENPACE needs a large number of responses to get a good estimate of $\boldsymbol{\delta_i}$, the maximum possible error made by a respondent *i*. This conclusion is mainly of theoretical interest because in practical contexts, we are interested in model performance where we are able to ask each respondent only a few questions.

Simulation 2: Recovery of GENPACE Parameters

The primary objective in the second simulation is to more closely evaluate the performance of GENPACE. Recall that GENPACE nests the other models, and its empirical performance should be superior to the other models if we have the luxury of asking respondents a large number of questions. Simulation 1 provided some insights on GENPACE's ability to recover true partworths. In the second simulation, we do a broader assessment of GENPACE's ability to recover both δ_i and the partworths. Following are the details of the steps used to generate our simulation data:

1. Generate a random vector of average partworths, $\bar{u} \sim N(0,25I)$
2. Generate $u_i \sim N(\bar{u}, 20I)$
3. Generate $\delta_i \sim Gamma(20, 0.5)$. Thus, $E(\delta_i) = 10$ and $V(\delta_i) = 5$.
4. Generate response error using $\varepsilon_i \sim N(0,100)\, I(-\delta_i \le \varepsilon_i \le \delta_i)$. In view of steps 3 and 4, the maximum response error was on average equal to one standard deviation (i.e., $E(\delta_i) = \sigma_i$)
5. Generate covariates (except for the covariate associated with the intercept term) from a $N(0, 1)$ distribution.
6. Compute, $y_i = X_i u_i + \varepsilon_i$

Table 6 summarizes the results from our simulations. We report the true partworth mean vector, the corresponding GENPACE estimate; RMSE for individual partworths, partworth means, the true mean and variance of δ and the corresponding GENPACE estimates. A key conclusion from this simulation is that as the amount of information available for estimation increases (i.e., more respondents and/or more questions), the performance of GENPACE improves. This is reassuring and shows that GENPACE is a well formed model that is identified.

[Insert Table 6 here]

Laptop Data: Empirical Performance of GENPACE With the Laptop Data

Toubia et al. (2003) collected data for a conjoint analysis study to predict consumers' preferences for various laptop bags. Their data were extensive and covered a range of question design algorithms (e.g., FASTPACE question design, random design) and data collection enhancements (e.g., self-explicated data). For testing Bayesian FASTPACE and GENPACE, we restricted our attention to data that were generated using the FASTPACE question design procedure, and we chose not to use the Self-Explicated data for our estimation.

In their study, Toubia et al. (2003) collected data from 88 respondents for a laptop design which had 10 attributes, each with two levels (*e.g.* price, color, size). FASTPACE sequentially generated partial-profile paired-comparison questions, based on the respondent's ratings. These questions and ratings were then used to estimate individual-level parameter estimates. After rating 20 paired comparisons, there was a hold out task in which the respondents were asked to rank-order five laptop bags that they had not seen before. Using either 8 or 16 questions, Toubia et al. (2003) assessed model performance based on the rank-order correlation between the actual ranks provided by the respondents and the ranks predicted by the respective model. We replicated the above analysis with the laptop dataset and computed the rank order correlations for the five models we chose to compare i.e., FASTPACE, HB Regression, GENPACE, Constrained

HB Regression, Bayesian FASTPACE. Table 7 reports the relative performance of various models using the hold-out sample of the laptop dataset.

[Insert Table 7 here]

Two points emerge from the above analysis. One, Bayesian FASTPACE is at least as good as FASTPACE with respect to the hold-out sample, with its predictions being more accurate than FASTPACE when we have 8 questions. At the same time, hold out performance of GENPACE is worse than the other three models (FASTPACE, Bayesian FASTPACE and HB). Also, the performance of GENPACE relative to FASTPACE and Bayesian FASTPACE degrades as the number of questions increases from 8 to 16. The relative inferior performance of GENPACE suggests that for the laptop data, a uniform response error distribution is likely to be the appropriate choice. Overall, the above results justify the value of using Bayesian FASTPACE, instead of FASTPACE, to estimate partworths for ratings-based conjoint data.

The results from our two simulations and the laptop data attest to the value of using a Bayesian approach to estimating partworths when we use the polyhedral Conjoint Analysis for generating question sequences. Recall that our focus is on partworth estimation, and the question sequence is still being generated by FASTPACE. All that we do is to add a Bayesian approach to estimating partworths from the responses obtained from respondents. The fact that Bayesian FASTPACE performs better than the traditional FASTPACE estimation in our study shows that estimates based on posterior mean, instead of analytic center, generates more accurate partworth estimates from the FASTPACE model. Our results for the GENPACE model attest to the value of the Bayesian approach in a broader set of contexts than just FASTPACE. By incorporating shrinkage, we can get more accurate estimates of partworths from Conjoint ratings data. However, this improvement in performance from using GENPACE comes with a cost – the number of questions to be asked of the respondent increases, especially when response errors are likely to be small.

SUMMARY AND CONCLUSIONS

For ratings-based conjoint data, which is the setting explored in this paper, there are currently three ways to incorporate error structures: (1) the polyhedral-based models, which introduce the minimal amount of error required to resolve any inconsistencies (see Toubia et al. 2003), (2) the traditional Hierarchical Bayesian (HB) regression models, which minimize a squared error loss function (Lenk et al., 1996), and (3) machine-learning based methods that offer the ability to estimate nonlinear preference functions while accommodating response error (Toubia et al, 2008). We extend our understanding of the polyhedral-based approaches by showing that we can obtain superior point estimates by using a probability model with a bounded response error distribution, where the bound is symmetric and the error is uniform between the bounds. As we had indicated earlier, from a theoretical perspective, there is good reason to impose bounds on the errors (although the researcher, or even the respondent, may not know the exact value of the bound) to accommodate the notion that people tend to exhibit preference consistency.

We accommodate the differing assumptions underlying Bayesian FASTPACE and HB Regression by proposing a unified model called GENPACE. We demonstrate how GENPACE nests various models, including Bayesian FASTPACE, HB Regression, and Constrained HB Regression.

For the laptop data, originally collected by Toubia et al. (2003), we find that Bayesian FASTPACE performs at least as well as FASTPACE and is in fact better than FASTPACE when the number of questions is low, a situation of particular interest in practice. Based on our simulation studies, we conclude that Bayesian FASTPACE outperforms FASTPACE and that GENPACE outperforms all models including Bayesian FASTPACE when respondents are relatively homogeneous and when response error is high. However, when we do not have adequate number of responses GENPACE performs relatively poorly as compared to Bayesian FASTPACE; *e.g.* when respondents are heterogeneous, or when respondents are homogeneous and response error is low. In summary, our results show that Bayesian approaches are superior to maximum likelihood approaches to estimating partworths in the context of FASTPACE, where response errors are bounded.

There are several limitations of this study which may offer opportunities for further research. From an empirical perspective, Bayesian FASTPACE can only be used for ratings-based conjoint analysis. A natural extension would be to consider appropriate link functions for choice-based polyhedral conjoint analysis. Note, however, that both ratings-based and choice-based conjoint models predict equally well on holdout data, but ratings-based conjoint models may be particularly appropriate if the researcher is looking to segment customers based on their stated preferences (Elrod, Louviere, and Davey 1992).

From the perspective of experimental design and sequential question generation, we only focus on the estimation aspects. It is also important to explore the relationship between question design using Bayesian FASTPACE and the question design algorithm of FASTPACE. Following this line of investigation, enhancing GENPACE to allow for generating adaptive questions is a potentially interesting and useful avenue for future research. For example, one possible approach to selecting the next question presented to the respondent is by minimizing expected preposterior predictive error. More generally, there may be opportunities to develop adaptive question design and estimation approaches that minimize endogeneity bias (Liu, Otter, and Allenby 2007).

From an empirical perspective, estimating partworths accurately using GENPACE is a challenge. It would be useful to assess the potential value of incorporating empirically derived priors and/or self-explicated priors to improve GENPACE estimates, especially with respect to out-of-sample predictions. For example, as a study progresses, the ratings provided by the initial set of participants could be used to derive the empirical priors for later respondents. Another empirical analysis worth considering is a more comprehensive simulation-based assessment to specify finer-grained conditions under which it would be valuable to use GENPACE. In particular, it would be worthwhile to explore the performance of GENPACE over a broader range of experimental factors and factor levels, over a broader set of performance measures, and over a broader set of distributions for generating simulation parameters. Such a simulation experiment would not invalidate the findings of this paper; rather, it would add texture to the results we have already reported.

In sum, we hope the paper has helped generate insights about the potential value of combining two modeling strategies for Conjoint Analysis, namely the optimization based FASTPACE algorithm and the statistically-oriented Hierarchical Bayes regression model. More generally, we hope that our approach points the way to combining optimization-based and probability-based modeling approaches in future research.

REFERENCES

Andrews, R., Ainslie, A. and Currim, Imran. (2002), "An Empirical Comparison of Logit Choice Models with Discrete Versus Continuous Representations of Heterogeneity." *Journal of Marketing Research*. **39**, 479.

Boatwright, Peter, Robert McCulloch, Peter Rossi (1999), "Account-Level Modeling for Trade Promotion: An Application of a Constrained Parameter Hierarchical Model" *Journal of the American Statistical Association*, **94**, 448.

Bradlow, Eric T., Ye Hu and Teck-Hua Ho (2004), "A Learning-Based Model for Imputing Missing Levels in Partial Conjoint Profiles" *Journal of Marketing Research*, **41**, 369.

Damien, Paul, Jon Wakefield and Stephen Walker (1999), "Gibbs Sampling for Bayesian Non-Conjugate and Hierarchical Models by Using Auxiliary Variables", *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, **61**, 2.

Ding, Min, Raj Grewal and John Liechty (2005), "Incentive Aligned Conjoint Analysis." *Journal of Marketing Research*. **42**, 67.

Elrod, Terry, Jordan J. Louviere, and Krishnakumar S. Davey (1992), "An Empirical Comparison of Ratings-Based and Choice-Based Conjoint Models, " Journal of Marketing Research, Vol. XXIX (August), p. 368-277.

Gauss, C.F. (1809) *Theoria Motus Corporum Coelestium*, Perthes & Besser, Hamburgi. (English edn., *The Heavenly Bodies Moving around the sun in Conic Sections,* Dover, New York, 1963.)

Genz, Alan (1992), "Numerical Computation of Multivariate Normal Probabilities." *Journal of Computational and Graphical Statistics*, **1**, 141.

Gintis, Herbert (2007), Game Theory Evolving, Second Edition, Princeton University Press, NJ: Princeton. (Forthcoming).

Green, Paul E., Krieger, Abba M., & Wind, Yoram (2001), "Thirty years of conjoint analysis: Reflections and prospects." *Interfaces,* (May/June), S56.

Lee, Leonard , Amir, On and Ariely, Dan, "In Search of Homo Economicus: Preference Consistency, Emotions, and Cognition" (August 2006). Available at SSRN: http://ssrn.com/abstract=925978.

Legendre, Adrien Marie (1805) *Nouvelles methods pour la determination des orbites des cometes*. Paris: Courcier.

Lenk, Peter J., DeSarbo, Wayne S., Green, Paul E., & Young, Martin R. (1996), "Hierarchical Bayes conjoint analysis: Recovery of partworth heterogeneity from reduced experimental designs." *Marketing Science*, **15**, 173.

Liechty, John C., Wayne S. DeSarbo and Duncan K. H. Fong (2005), "Dynamic Models Incorporating Individual Level Heterogeneity: Utility Evolution in Conjoint Analysis," *Marketing Science*, **24**, 285.

Liu, Qing, Thomas Otter, and Greg M Allenby (2007), "Investigating Endogeneity Bias in Marketing" Marketing Science, 26, 642.

Rossi, Peter E., Robert E. McCulloch, and Greg H. Allenby (1996), "The Value of Purchase History Data in Target Marketing," *Marketing Science*, **15**, 321.

Stanovich, Kieth. E., Sá, Walter. C., & West, Richard. F. (2003), "Individual Differences in Thinking, Reasoning, and Decision Making." In Jacqueline P. Leighton and Robert J. Sternberg (Eds.), *The Nature of Reasoning* (pp. 375-409). New York: Cambridge University Press.

Toubia, Olivier, Duncan I. Simester, John R. Hauser, and Ely Dahan (2003), "Fast Polyhedral Adaptive Conjoint Estimation," *Marketing Science,* **22**, 273.

Toubia, Olivier, John R.Hauser, and Rosanna Garcia (2007), "Probabilistic Polyhedral Methods for Adaptive Choice-Based Conjoint Analysis: Theory and Application," Marketing Science, 26(5).

Toubia, Olivier, Theodoros Evgeniou, and John R. Hauser (2007), "Optimization-Based and Machine-Learning Methods for Conjoint Analysis: Estimation and Question Design," in *Conjoint Measurement: Methods and Applications* (4th edition) Springer Berlin Heidelberg

FOOTNOTES

---

[1] Toubia et al. (2007) incorporate an error structure in the context of polyhedral approaches to discrete choice conjoint. They accommodated response errors in polyhedral discrete choice conjoint by assuming a mixture distribution over the choices that are available to a respondent. This approach is reasonable for accommodating response error in discrete choice conjoint, but is not appropriate in the context of ratings-based conjoint where ratings are assumed to be a continuous measure of a consumer’s preference. Our focus in this paper is on ratings-based conjoint with continuous measures of preference.

[2] We use $\boldsymbol{\epsilon}$ to denote a vector of ones throughout the paper with the understanding that its dimension is appropriate for the context in which it is used.

[3] The issue of multiple maximum likelihood estimates for a parameter is similar to the problem of multiple Nash equilibria in game theory. To select a specific point estimate from all possible MLE estimates, we need an additional criteria to select one of the MLEs as a point estimate, just as game theorists need additional selection rules to select a specific equilibrium from all possible equilibria.

[4] In the context of adaptive designs there has been some recent discussion in the literature about endogeneity bias and potential violation of the likelihood principle by statistical models (Liu, Otter, and Allenby 2007). Bayesian FASTPACE and GENPACE do not violate the likelihood principle as the specifics of the question generation algorithm are not part of their respective likelihood functions.

[5] We used the software code developed by Toubia et al. (2003), suitably modified, to implement FASTPACE question design for our simulations.

Tables

Model Name	Partworths Density	Parameter Restrictions
Individual Level OLS	$N\left(\left[X_i^T X_i\right]^{-1}\left[X_i^T a_i\right], \sigma_i^2\left[X_i^T X_i\right]^{-1}\right)$	$\delta_i \rightarrow \infty$ $lb \rightarrow -\infty$ $ub \rightarrow \infty$ $\Sigma \rightarrow \infty$
Bayesian FASTPACE	$I(a_i - \delta_i e \leq X_i u_i \leq a_i + \delta_i e) \quad I(0e \leq u_i \leq 100e)$	$lb = 0e$ $ub = 100e$ $\sigma_i^2 \rightarrow \infty$ $\Sigma \rightarrow \infty$
HB FASTPACE	$f(u_i \mid -) \propto I(a_i - \delta_i e \leq X_i u_i \leq a_i + \delta_i e) \quad I(0e \leq u_i \leq 100e)$ $\exp\left(-\frac{\left(u_i - \overline{u}\right)' \Sigma^{-1}\left(u_i - \overline{u}\right)}{2}\right)$	$lb = 0e$ $ub = 100e$ $\sigma_i^2 \rightarrow \infty$
HB regression	$N\left(\left[\sigma_i^{-2} X_i^T X_i + \Sigma^{-1}\right]^{-1}\left[X_i^T a_i + \Sigma^{-1}\overline{u}\right], \left[\sigma_i^{-2} X_i^T X_i + \Sigma^{-1}\right]^{-1}\right)$	$\delta_i \rightarrow \infty$ $lb \rightarrow -\infty$ $ub \rightarrow \infty$
Constrained HB regression	$N\left(\left[\sigma_i^{-2} X_i^T X_i + \Sigma^{-1}\right]^{-1}\left[X_i^T a_i + \Sigma^{-1}\overline{u}\right], \left[\sigma_i^{-2} X_i^T X_i + \Sigma^{-1}\right]^{-1}\right)$ $I(0e \leq u_i \leq 100e)$	$\delta_i \rightarrow \infty$ $lb = 0e$ $ub = 100e$
GENPACE (our implementation)	$N\left(\left[\sigma_i^{-2} X_i^T X_i + \Sigma^{-1}\right]^{-1}\left[X_i^T a_i + \Sigma^{-1}\overline{u}\right], \left[\sigma_i^{-2} X_i^T X_i + \Sigma^{-1}\right]^{-1}\right)$ $I(a_i - \delta_i e \leq X_i u_i \leq a_i + \delta_i e) \quad I(0e \leq u_i \leq 100e)$	$lb = 0e$ $ub = 100e$
GENPACE	$N\left(\left[\sigma_i^{-2} X_i^T X_i + \Sigma^{-1}\right]^{-1}\left[X_i^T a_i + \Sigma^{-1}\overline{u}\right], \left[\sigma_i^{-2} X_i^T X_i + \Sigma^{-1}\right]^{-1}\right)$ $I(a_i - \delta_i e \leq X_i u_i \leq a_i + \delta_i e) \quad I(lb \leq u_i \leq ub)$	Not Applicable

Table 1: This table summarizes the partworth full conditional densities for various probability models that can be specified using the GENPACE model.

Model Name	**Partworth Estimation**	**Partworth Restrictions**	**Response Error Distribution**	**Partworth Shrinkage**
FASTPACE	Analytic Center	Bounded between 0 and 100	N/A	
Bayesian FASTPACE	Posterior Mean		Bounded, Uniform	No
Constrained HB Regression			Unbounded, Normal	Yes
HB Regression		None		
GENPACE		Bounded between 0 and 100	Bounded, Normal	

Table 2: Properties of alternative models

Simulation Variable	Value
No of Attributes	10
No of Respondents	100
Partworth Mean	50
No of replicates per experimental cell	20
Bound on response error distribution	1

Table 3: Variables whose values were fixed in the experimental design

Simulation Variable	Values Investigated
Standard Deviation of Partworths	10, 30
Response Error Standard Deviation	20, 40
No of Questions	8, 16, 32 and 100

Table 4: Experimental variables whose values were manipulated (as shown in the table)

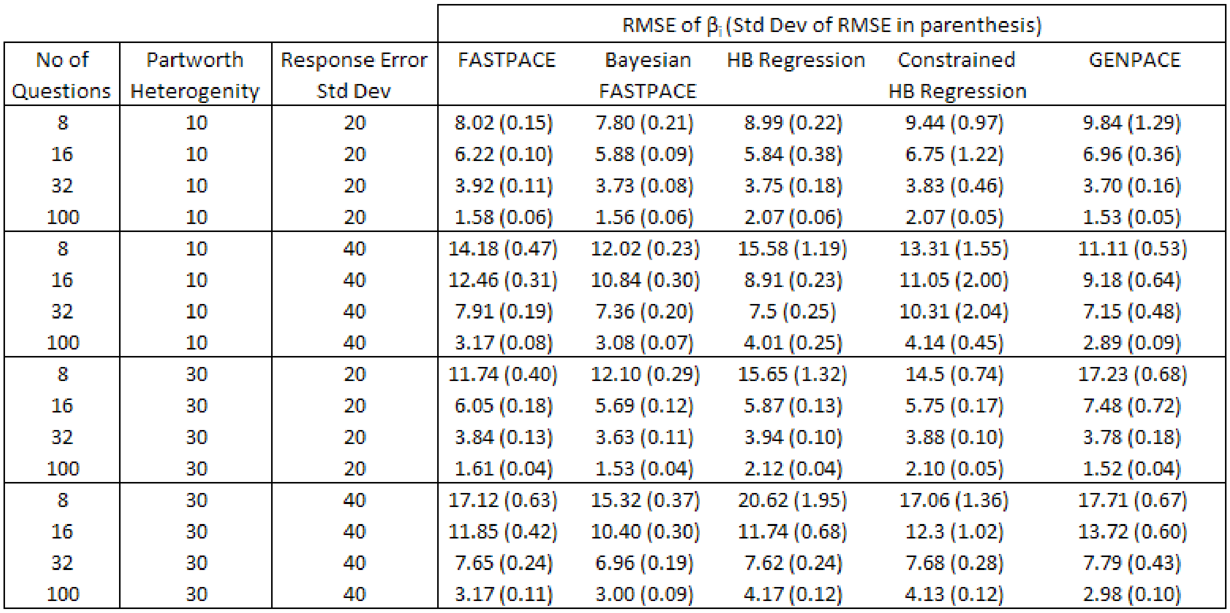

			RMSE of β_i (Std Dev of RMSE in parenthesis)				
No of Questions	Partworth Heterogenity	Response Error Std Dev	FASTPACE	Bayesian FASTPACE	HB Regression	Constrained HB Regression	GENPACE
8	10	20	8.02 (0.15)	7.80 (0.21)	8.99 (0.22)	9.44 (0.97)	9.84 (1.29)
16	10	20	6.22 (0.10)	5.88 (0.09)	5.84 (0.38)	6.75 (1.22)	6.96 (0.36)
32	10	20	3.92 (0.11)	3.73 (0.08)	3.75 (0.18)	3.83 (0.46)	3.70 (0.16)
100	10	20	1.58 (0.06)	1.56 (0.06)	2.07 (0.06)	2.07 (0.05)	1.53 (0.05)
8	10	40	14.18 (0.47)	12.02 (0.23)	15.58 (1.19)	13.31 (1.55)	11.11 (0.53)
16	10	40	12.46 (0.31)	10.84 (0.30)	8.91 (0.23)	11.05 (2.00)	9.18 (0.64)
32	10	40	7.91 (0.19)	7.36 (0.20)	7.5 (0.25)	10.31 (2.04)	7.15 (0.48)
100	10	40	3.17 (0.08)	3.08 (0.07)	4.01 (0.25)	4.14 (0.45)	2.89 (0.09)
8	30	20	11.74 (0.40)	12.10 (0.29)	15.65 (1.32)	14.5 (0.74)	17.23 (0.68)
16	30	20	6.05 (0.18)	5.69 (0.12)	5.87 (0.13)	5.75 (0.17)	7.48 (0.72)
32	30	20	3.84 (0.13)	3.63 (0.11)	3.94 (0.10)	3.88 (0.10)	3.78 (0.18)
100	30	20	1.61 (0.04)	1.53 (0.04)	2.12 (0.04)	2.10 (0.05)	1.52 (0.04)
8	30	40	17.12 (0.63)	15.32 (0.37)	20.62 (1.95)	17.06 (1.36)	17.71 (0.67)
16	30	40	11.85 (0.42)	10.40 (0.30)	11.74 (0.68)	12.3 (1.02)	13.72 (0.60)
32	30	40	7.65 (0.24)	6.96 (0.19)	7.62 (0.24)	7.68 (0.28)	7.79 (0.43)
100	30	40	3.17 (0.11)	3.00 (0.09)	4.17 (0.12)	4.13 (0.12)	2.98 (0.10)

Table 5: Simulation results comparing relative performance of competing models.

No of Respondents	No of questions	True Partworth Mean	GENPACE Estimate of Partworth Mean	RMSE of Partworth Mean	RMSE of Parworths	RMSE of Delta	True Mean of Delta	Mean of Genpace Estimates of Delta	True Variance of Delta	Variance of Genpace Estimates of Delta
10	10	3.76 -1.43 -6.33 4.25 1.19	4.27 0.64 -10.69 5.02 -0.43	5.17	2.69	4.32	9.82	5.69	2.96	2.34
10	25	-4.30 1.17 -2.73 1.07 2.81	-6.04 -1.80 -3.56 1.34 2.89	3.55	1.22	2.53	10.57	8.20	3.39	2.34
50	10	3.23 10.02 4.32 -0.22 -9.10	2.59 10.78 3.23 -0.89 -8.82	1.64	2.65	5.08	9.89	5.07	4.24	2.52
50	25	0.32 -2.05 -7.87 -0.66 6.95	1.13 -2.17 -7.21 0.05 6.94	1.27	1.13	2.56	9.98	7.61	3.95	2.35
50	50	5.52 -0.57 -3.05 0.79 1.35	5.36 0.18 -2.91 0.70 0.32	1.30	0.74	1.46	10.41	9.14	6.39	4.71
150	10	-0.74 2.50 -1.89 -4.37 -4.12	-0.88 2.32 -2.30 -4.83 -3.23	1.11	2.81	5.46	10.21	5.11	5.60	3.47
150	25	-8.59 5.04 1.63 10.53 0.51	-8.29 4.55 1.53 10.65 0.42	0.59	1.24	2.51	9.94	7.69	4.88	2.75
150	100	5.76 4.14 -4.76 -4.64 -5.10	6.12 3.78 -4.19 -4.60 -4.97	0.78	0.43	0.76	9.95	9.31	4.57	3.49
250	10	-3.42 -12.06 -2.71 4.94 3.09	-2.75 -12.44 -2.36 4.70 3.02	0.88	2.81	5.28	10.10	5.13	4.57	3.38
250	25	-1.83 1.13 -2.17 2.25 3.99	-1.75 1.31 -1.93 2.27 3.86	0.34	1.28	2.52	10.01	7.74	4.99	3.07
250	100	-9.92 2.73 4.33 -9.99 -0.97	-9.91 2.69 4.44 -10.07 -1.18	0.25	0.41	0.75	9.90	9.23	3.79	3.13

Table 6 – Summary of GENPACE parameter recovery

No of Questions	FASTPACE	Bayesian FASTPACE	HB Regression	Constrained HB Regression	GENPACE
8	0.59	0.62	0.61	0.55	0.59
16	0.68	0.68	0.66	0.58	0.65

Table 7: Rank order correlations between the actual ranks provided by respondents and predictions by various models for the laptop dataset

Figures

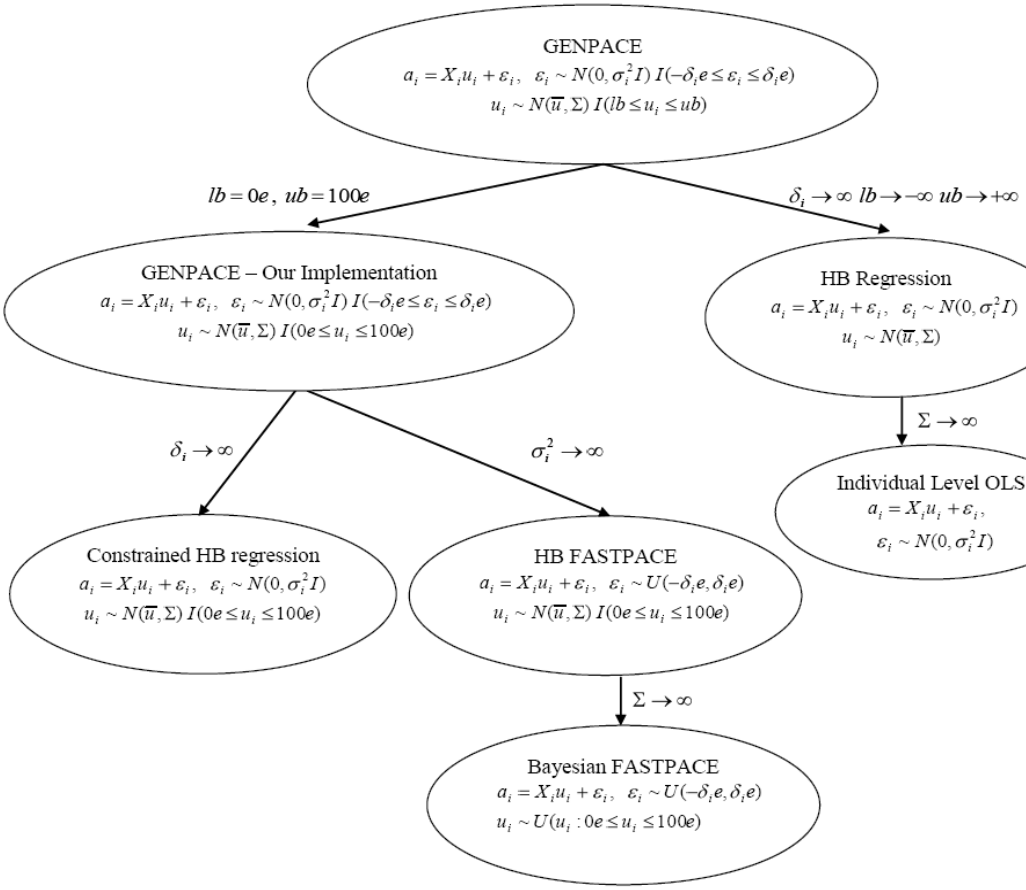

Figure 1: Graphical representation showing how GENPACE nests the other derivative models with respect to parameter restrictions. This figure shows the additional parameter restrictions that need to be imposed to go from a given model to its corresponding derivative model(s).

Appendix

1. Prior Structures for GENPACE

Our experience implementing GENPACE indicates that unless the priors for σ_i^2 are chosen carefully, the individual variance terms σ_i^2 may go to infinity, something that occurs because of the very flat tails of the full conditional distribution of σ_i^2 when δ_i is finite. Our experience is consistent with that of Boatwright et al (1999) where they encounter situations in which parameters of the truncated normal distribution are empirically unidentified with small data sets. The solution adopted by Boatwright et al. was to impose restrictions on the normalizing constant such that the MCMC sampler does not accept very high values for the variance parameter σ_i^2. We adopt a two-pronged methodology to address the identification issue with respect to σ_i^2. First, we shrink σ_i^2 and second, we impose upper bound restrictions on acceptable draws of σ_i^2, which is philosophically consistent with the approach of Boatwright et al., who implicitly impose constraints on σ_i^2 by imposing restrictions on the normalizing constant (see Boatwright et al, 1999 for details of their procedure). Details of our restrictions are outlined in the next paragraph.

To avoid excursions of individual σ_i^2 towards infinity, we have to choose the priors for σ_i^2 carefully. We found that by placing a rather large bound on how spread out the σ_i^2s can become, or by including strong subjective beliefs on the variance of the variances, the model was able to mix and perform well. To simplify calculations, we assumed that σ_i^2 followed a truncated normal (strictly positive), with a mean that has vague priors and a variance that is equal to a point mass. (Given the tendency for the σ_i^2s to head to infinity, this would be equivalent to having the variance follow a uniform on 0 to an upper bound, where the upper bound equaled the point mass). For example, we found that by setting the point mass at 10 times the variance of posterior estimates of variance of the σ_i^2s, (which came from individual HB models with vague priors) we were able to stop the excursions of individual σ_i^2 towards infinity and recover GENPACE model parameters. Hence, to ensure that our model can recover the true parameters with vague (uninformative priors) we used the following prior structures:

$$\sigma_i^2 \sim N\left(\overline{\sigma^2}, 10 Var\left(\sigma_{i,HB}^2\right)\right) I\left(\sigma_i^2 > 0\right) \quad \text{, with } \overline{\sigma^2} \sim N(Mean\left(\sigma_{i,HB}^2\right), 10{,}000) \quad .$$

The prior structures for the remaining parameters of the model are given in table A1.

Prior	Parameter	Value
$\bar{u} \sim Normal(\bar{u}_{pr}, v^2 I)$	$\bar{u}_{pr}$	**0**
	v^2	10^{10}
$\Sigma^{-1} \sim Wishart(R^{-1}, d)$	R^{-1}	**I**
	d	11
$\delta_i \sim Gamma(sc_{\delta_i}, sp_{\delta_i})$	sc_{δ_i}	10^5
	sp_{δ_i}	1

Table A1: Priors for $\overline{u}$, Σ^{-1} and δ_i

2. <u>Rationale for Shrinking Response Error Variance</u>

From a theoretical perspective, shrinking response error variance is not required for unifying Bayesian FASTPACE and HB Regression within GENPACE formulation. However, our experience with estimating GENPACE indicates that without shrinking response error variance we would need a large number of questions to recover δ_i. Therefore, we include an option for shrinking response error variance in GENPACE to obtain robust estimates of δ_i.

3. <u>Proposal Distributions</u>

We used the slice sampler (see Damien, Wakefield and Walker, 1999 for details about the slice sampler) to generate draws from the full conditional distributions of u_i and σ_i^2 and used the Metropolis-Hastings algorithm to sample from the full conditional distributions of $\overline{\boldsymbol{u}}$, Σ^{-1}, $\overline{\sigma^2}$ and δ_i. We drew proposed samples for the Metropolis-Hastings algorithm by means of random-walk proposals as shown by the table A2 below:

Parameter	Proposal
$\overline{\boldsymbol{u}}$	$\bar{u}_{\text{Proposed}} \sim N(\bar{u}_{\text{Current}}, \tau^2 I)$
Σ^{-1}	$\Sigma^{-1}_{\text{Proposed}} \sim W\left(\frac{\Sigma^{-1}_{\text{Current}}}{\rho+11}, \rho+11\right)$
$\overline{\sigma^2}$	$\overline{\sigma^2}_{\text{Proposed}} \sim G\left(\theta+2, \frac{\overline{\sigma^2}_{\text{Current}}}{\theta+2}\right)$
δ_i	$\delta_{i,\text{Proposed}} \sim N\left(\delta_{i,\text{Current}}, \eta^2\right) I(\delta_{i,\text{Proposed}} e \geq X_i u_i - a_i) I(\delta_{i,\text{Proposed}} e \geq a_i - X_i u_i)$

Table A2: Random walk proposals for GENPACE