This item is the archived peer-reviewed author-version of:

An improved two-stage variance balance approach for constructing partial profile designs for discrete choice experiments

Reference:

Kessels Roselinde, Jones Bradley, Goos Peter.- An improved two-stage variance balance approach for constructing partial profile designs for discrete choice experiments

Applied stochastic models in business and industry - ISSN 1524-1904 - 31:5(2015), p. 626-648

Full text (Publishers DOI): <http://dx.doi.org/doi:10.1002/asmb.2065>

Handle/Permalink: <http://hdl.handle.net/10067/1277320151162165141>

An improved two-stage variance balance approach for constructing partial profile designs for discrete choice experiments

Roselinde Kessels^{*†}

Universiteit Antwerpen

Bradley Jones[‡]

SAS Institute

Universiteit Antwerpen

Peter Goos[§]

Universiteit Antwerpen

Katholieke Universiteit Leuven

^{*}Roselinde Kessels is a postdoctoral researcher in econometrics at Universiteit Antwerpen, Faculty of Applied Economics, Department of Economics and StatUa Center for Statistics, Prinsstraat 13, 2000 Antwerpen, Belgium. Tel.: +32 (0)3 265 40 95. E-mail: roselinde.kessels@uantwerpen.be

[†]Corresponding author.

[‡]Bradley Jones is a principal research fellow at SAS Institute Inc., SAS Campus Drive, Cary, NC 27513, USA. He is also a guest professor at Universiteit Antwerpen, Faculty of Applied Economics and StatUa Center for Statistics, Prinsstraat 13, 2000 Antwerpen, Belgium. Tel.: +1 919 531 4161. E-mail: bradley.jones@jmp.com

[§]Peter Goos is a full professor in statistics at Universiteit Antwerpen, Faculty of Applied Economics and StatUa Center for Statistics, Prinsstraat 13, 2000 Antwerpen, Belgium. He is also a full professor in statistics at Katholieke Universiteit Leuven, Faculty of Bioscience Engineering, Department of Mechatronics, Biostatistics and Sensors, Kasteelpark Arenberg 30, 3001 Leuven. Tel.: +32 (0)3 220 40 59. E-mail: peter.goos@uantwerpen.be

An improved two-stage variance balance approach for constructing partial profile designs for discrete choice experiments

Abstract

In many discrete choice experiments set up for product innovation, the number of attributes is large, which results in a substantial cognitive burden for the respondents. To reduce the cognitive burden in such cases, Green (1974) suggested the use of partial profiles that vary only the levels of a subset of the attributes. In this paper, we present two new methods for constructing Bayesian D-optimal partial profile designs for estimating main-effects models. They involve alternative generalizations of Green’s approach that makes use of balanced incomplete block designs and take into account the fact that attributes may have differing numbers of levels. We refer to our methods as variance balance I and II because they vary an attribute with a larger number of levels more often than an attribute with fewer levels to stabilize the variances of the individual part-worth estimates. The two variance balance methods differ in the way attributes with differing numbers of levels are weighted. Both methods provide statistically more efficient partial profile designs for differing numbers of attribute levels than another generalization of Green’s approach that does not weigh the attributes. This method is called attribute balance. We show results from an actual experiment in software development demonstrating the usefulness of our methods.

Keywords: discrete choice experiments, Bayesian D-optimal design, partial profiles, lexicographic choice behavior, variance balance I and II, attribute balance

1 Introduction

Discrete choice experiments (DCEs), also called conjoint choice or stated choice experiments, are a standard tool for quantifying people’s preferences in marketing, transportation, and environmental, health and political studies (Louviere et al. 2003). Given a set of predefined attributes of a product or an item, DCEs identify those attributes that matter most and indicate the most appealing levels for these attributes. In a DCE, levels of different attributes are combined in a profile item or alternative. A select number of profiles are grouped into choice sets, for each of which respondents indicate the alternative they prefer. The statistical analysis of the respondents’ choices employs discrete choice modeling to estimate the preference parameters attached to each attribute.

DCEs generally work well when the number of attributes in a profile is not too large. Green (1974), Green and Srinivasan (1990), Schwabe et al. (2003), Bradlow (2005), Chrzan (2010) and Kessels et al. (2010), among others, argue that, in a DCE, respondents can process only a limited number of attributes, loosely speaking, four to seven depending on the application. At the same time, they recognize the demand for DCEs suitable for investigating larger numbers of attributes, especially in the case of high-tech durable products. To meet this demand, Green (1974) originated the idea of constructing profiles based on only a subset of the attributes. A choice set then consists of profiles that vary only in the levels of a subset of the attributes called partial profiles. By contrast, profiles that vary the levels of all attributes in a choice set are called full profiles. The profile strength is the number of attributes whose levels vary in the partial profiles. These varying attributes differ from choice set to choice set.

Assuming that the constant attributes do not interact with the varying attributes, the levels taken by the constant attributes in a choice set do not affect the profile that is chosen. That is why in a traditional partial profile design that displays the profiles as columns of words or sentences, the constant attributes are dropped from the choice sets. However, in many DCEs, especially those that display the profiles as images (e.g. images of products), it is at least preferable or even required to also show the levels of the constant attributes to the respondents. In this respect, Dellaert et al. (2012) argue that in the absence of the levels of the constant attributes, the parameter estimates of the choice model may not reflect real-world choice behavior. For this reason, we extend the use of the term “partial” profiles to simply indicate that the levels of some attributes remain constant in every profile in each choice set. In Section 2, we describe a DCE in software development where showing the levels of the constant attributes is required.

Some researchers have disagreed with the notion that the cognitive burden of a large number of attributes overwhelms respondents, which has served as the primary reason for using partial profiles. Mainly Louviere (2005) and Louviere et al. (2011) maintain that profiles with large numbers of attributes may not overburden respondents because they are faced with complex choices in everyday life. Hensher (2006) added some nuance

to this result by arguing that attribute relevancy may be more important than complexity. However, like many other researchers exploring the complexity and cognitive burden associated with DCEs (see, e.g., Swait and Adamowicz 2001; DeShazo and Fermo 2002; Caussade et al. 2005), he provided evidence that a large number of attributes does have a detrimental effect on the ability to choose, contributing to an increased error variance and increased parameter differences.

One consequence of presenting many attributes in a DCE is that respondents may make choices based on a hierarchical or lexicographic preference ordering of the attributes, in a way that no longer reflects their true preferences. Rather than trading off all the profiles' attribute levels, they focus on only one attribute or a small subset of the attributes, the levels of which drive the entire choice process (see, e.g., Scott 2002; Hensher 2006; Hensher and Rose 2009). This kind of choice behavior in a DCE then results in little or no information about the remaining attributes.

The fact that respondents resort to simplifying, non-compensatory decision rules when processing multi-attribute information in a DCE provides a compelling argument for using partial profiles. Holding the levels of some of the attributes constant in every choice set makes the choice tasks easier and prevents respondents from using non-compensatory decision strategies in a DCE for the purpose of avoiding complex trade-offs.

Note that we distinguish between two types of non-compensatory decision making processes, namely design-driven non-compensatory choice behavior and preference-driven non-compensatory choice behavior. By preference-driven non-compensatory behavior, we refer to situations in which the respondents mainly care about one attribute or a small subset of the attributes, whether in an artificial DCE setting or not. We view partial profiles primarily as a way to prevent design-driven non-compensatory choice behavior when many attributes are under study in a DCE. To some extent, partial profiles are also useful in the presence of preference-driven non-compensatory choice behavior. If a full profile discrete choice design is used and the level of a dominant attribute is varied in every choice set, then the DCE does not reveal any information concerning the other attributes. Holding the level of the dominant attribute constant in part of the choice sets, in a partial profile design, forces the respondents to evaluate the remaining attributes, resulting in useful information concerning these attributes. Obviously, information concerning the non-dominant attributes is as important as information concerning the dominant attribute, because successfully competing products and items often have the exact same key characteristics (i.e., the same level for the dominant attribute) and a wise choice of secondary features may result in a significant increase in general attractiveness.

So far, most of the literature for DCEs using partial profiles has focused on utility-neutral designs involving choice sets with two profiles (Grasshoff et al. 2003, 2004; Grossmann et al. 2006, 2009, 2013). These designs are optimal for one specific set of parameter values, namely the set of all zero values. The utility-neutral designs assume zero utilities

for any prospective profile, and, hence, an equal utility for each profile. The most recent utility-neutral partial profile designs created by Grossmann et al. (2009, 2013) only allow for two or three groups of attributes where the number of attribute levels is fixed in every group. There are also limitations on the allowable numbers of choice sets in all cases considered and on the allowable numbers of varying attributes when there are two groups of attributes. The designs developed are therefore not applicable for a wide variety of practical problems.

Recently, Kessels et al. (2011a) have adopted the Bayesian design approach to construct partial profile designs. Bayesian designs are optimized over an assumed prior distribution of likely parameter values, which makes them robust to misspecification of the parameters (Sándor and Wedel 2001). The authors used a Bayesian $\mathcal{D}$ -optimality criterion for the main-effects multinomial logit (MNL) model and called the resulting partial profile designs $\mathcal{D}_B$ -optimal to emphasize their Bayesian nature. To generate $\mathcal{D}_B$ -optimal partial profile designs, Kessels et al. (2011a) provided a generalization of Green’s (1974) approach to use a balanced incomplete block design (BIBD) for selecting the attributes whose levels vary in each choice set. Their generalized method, which they refer to as attribute balance, tries to vary all attributes the same number of times in the design and works for any number of choice sets, any number of attributes and any number of varying attributes.

A weakness of the attribute balance approach is that it disregards the numbers of levels of each attribute. As a result, for design situations that involve differing numbers of attribute levels, the part-worths of the attributes with a larger number of levels are estimated less precisely than the part-worths of the attributes with fewer levels. It is desirable, however, to obtain similar amounts of information on each part-worth estimate. Therefore, we present two new methods for selecting the varying attributes that account for the fact that attributes may have differing numbers of levels. We refer to these methods as variance balance I and II. Like attribute balance, they are generalizations of Green’s BIBD approach. We use the term variance balance because the methods are meant to estimate all part-worths with the same precision, i.e. with the same variance. As we show in the paper, the variance balance methods provide statistically more efficient designs for differing numbers of attribute levels than does attribute balance. Hence, the variance balance methods offer practitioners efficiency gains as well as full flexibility when designing practical DCEs.

The outline of the remainder of the paper is as follows. Section 2 revisits the attribute balance approach used to design an actual DCE in software development with partial profiles. Section 3 reviews the multinomial logit model and the $\mathcal{D}_B$ -optimality criterion used to construct partial profile designs. In Section 4, we present the variance balance algorithms for generating efficient partial profile designs for differing numbers of attribute levels. In Section 5, we compare different $\mathcal{D}_B$ -optimal partial profile designs for an illustrative problem involving differing numbers of attribute levels and demonstrate the

excellent performance of the variance balance techniques in terms of statistical efficiency. Section 6 continues the discussion on the DCE in software development with an overview of the analysis results and a comparison of follow-up design solutions showing the practical value of our techniques. Section 7 concludes the paper and highlights some further research possibilities.

2 Revisiting attribute balance with a real-life DCE in software development

The JMP software product from SAS Institute is a general-purpose statistical analysis and graphical visualization tool for scientists and engineers. The JMP software group performed an online DCE to understand users' preferences for output displays generated by JMP. Typically, software developers make numerous product design decisions, but many of these decisions are not tested to find out whether the customer likes them. However, in this case, the JMP development team sought feedback from its customers by asking them to evaluate 15 choice sets of two output displays from a simple linear regression analysis. In each choice set, the two output displays differed in the levels of three attributes and respondents had to indicate the output display they liked better. Figure 1 shows an example of such a choice set.

The two output displays in Figure 1 differ as follows. First, the background color of the picture area around the graph is white in the left output display and creamish in the right output display. Second, the graph frame is bordered on two sides in the left output display and on all four sides in the right output display. Finally, an outer graph rectangle is absent in the left output display and present in the right output display.

The three attributes with different levels for the two output displays in a choice set are, however, not the same for all 15 choice sets. There were seven attributes under study, and the attributes whose levels varied in each choice set changed between choice sets. Table 1 shows the seven attributes in the DCE, together with their levels. The first attribute, the report background color, has four levels, whereas the remaining six attributes all have two levels. Note that the picture background color is either the same as or in contrast with the report background color.

The output displays shown in the 15 choice sets are partial profiles varying the levels of only three of the seven attributes in each choice set. For example, Table 2 shows the profiles with the attribute levels for the two output displays in Figure 1. Besides the three varying attributes, there are four constant attributes with levels that are depicted in gray. The JMP development team chose a design with partial profiles to keep the choice tasks manageable for the respondents.

To obtain precise part-worth estimates, the JMP team generated a Bayesian $\mathcal{D}$ -optimal partial profile design containing 120 choice sets using the design algorithm of Kessels et

al. (2011a). The 120 choice sets were divided into eight groups or surveys of 15 choice sets each such that each survey had a similar partial profile design structure (see below). Furthermore, the eight surveys were submitted to groups of eight respondents to make sure that each survey was filled out about an equal number of times.

The design algorithm of Kessels et al. (2011a) uses the attribute balance method to select the three varying attributes in each choice set. The method was inspired by Green’s (1974) idea to use a balanced incomplete block design (BIBD) for specifying the varying attributes. A BIBD describes how to arrange the t levels of a single qualitative factor, called treatments, in S blocks of size t_v , where $t_v < t$. In a BIBD, every treatment appears with the same frequency and the number of times each pair of treatments appear together in a block is the same for all pairs. We refer to Abel and Greig (2007) and Mathon and Rosa (2007) for more discussions and extant tables of BIBDs. By analogy, in the construction of partial profile designs, the choice sets represent the blocks and the varying attributes represent the treatments within the blocks. A BIBD can therefore be used as a master design to specify the varying attributes in each choice set. It balances the number of times an attribute varies over the entire design and the number of times pairs of varying attributes appear together within a choice set.

To illustrate the BIBD approach, Table 3 shows seven of the 120 choice sets of the software development partial profile design. The varying attributes form a BIBD master design because each attribute varies in three choice sets and each of the 21 possible pairs of varying attributes appears exactly once in the seven choice sets.

For the JMP study it is, however, not possible to generate a BIBD with 120 blocks containing $t_v = 3$ out of $t = 7$ different treatments. To get around this problem, Kessels et al. (2011a) present their attribute balance method as a flexible solution to adapt the master design of varying attributes to the specific choice design requirements. Motivated by the fact that a BIBD is a special case of a $\mathcal{D}$ -optimal design for a very specific ANOVA model, attribute balance generalizes the BIBD approach by maximizing the $\mathcal{D}$ -optimality criterion for that ANOVA model. Since a BIBD is $\mathcal{D}$ -optimal, attribute balance yields a BIBD in case such a design exists for a given situation. Otherwise, it approximates the balancing structure of a BIBD as closely as possible. This explains why the JMP team used the attribute balance approach to determine the varying attributes in the 120 choice sets.

3 The multinomial logit framework

The multinomial logit (MNL) model assumes that respondents to a DCE belong to a group of decision makers with homogeneous preferences. The model employs random utility theory which describes the utility that a respondent attaches to profile j ($j = 1, \dots, J$) in

choice set s ($s = 1, \dots, S$) as the sum of a systematic and a stochastic component:

$$U_{js} = \mathbf{x}'_{js}\boldsymbol{\beta} + \varepsilon_{js}. \quad (1)$$

In the systematic component $\mathbf{x}'_{js}\boldsymbol{\beta}$, $\mathbf{x}_{js}$ is a $k \times 1$ vector containing the attribute levels of profile j in choice set s . We assume all attributes are categorical. The vector $\boldsymbol{\beta}$ is a $k \times 1$ vector of parameter values representing the part-worths or main effects of the attribute levels on the utility. This part-worth vector is the same for every respondent. The stochastic component ε_{js} is the error term, which is assumed independently and identically extreme value distributed. Therefore, the MNL probability that a respondent chooses profile j in choice set s is the closed-form expression

$$p_{js} = \frac{\exp(\mathbf{x}'_{js}\boldsymbol{\beta})}{\sum_{t=1}^J \exp(\mathbf{x}'_{ts}\boldsymbol{\beta})}, \quad (2)$$

where $\boldsymbol{\beta}$ can be estimated using a maximum likelihood approach.

The construction of an optimal design $\mathbf{X} = [\mathbf{x}'_{js}]_{j=1, \dots, J; s=1, \dots, S}$ for estimating $\boldsymbol{\beta}$ in the MNL model (2) is based on the Fisher information matrix

$$\mathbf{M}(\mathbf{X}, \boldsymbol{\beta}) = \sum_{s=1}^S \mathbf{X}'_s (\mathbf{P}_s - \mathbf{p}_s \mathbf{p}'_s) \mathbf{X}_s, \quad (3)$$

with $\mathbf{X}_s = [\mathbf{x}'_{js}]_{j=1, \dots, J}$ the submatrix of $\mathbf{X}$ corresponding to choice set s , $\mathbf{p}_s = [p_{1s}, \dots, p_{Js}]'$ and $\mathbf{P}_s = \text{diag}[p_{1s}, \dots, p_{Js}]$. To deal with the dependency of the information matrix on $\boldsymbol{\beta}$, we generate Bayesian $\mathcal{D}$ -optimal or $\mathcal{D}_B$ -optimal designs that maximize the $\mathcal{D}_B$ -optimality criterion, which we define as

$$\mathcal{D}_B = \int_{\mathcal{R}^k} \log |\mathbf{M}(\mathbf{X}, \boldsymbol{\beta})| \pi(\boldsymbol{\beta}) d\boldsymbol{\beta}, \quad (4)$$

where $\pi(\boldsymbol{\beta})$ is a prior distribution of likely part-worth values over which a design is optimized. Often, this distribution is the multivariate normal distribution, $\mathcal{N}(\boldsymbol{\beta}|\boldsymbol{\beta}_0, \boldsymbol{\Sigma}_0)$, with prior mean $\boldsymbol{\beta}_0$ and prior variance-covariance matrix $\boldsymbol{\Sigma}_0$. Kessels et al. (2011a) were the first to use this criterion to generate $\mathcal{D}_B$ -optimal partial profile designs.

To compare the estimation performance of a partial profile design to a full profile design, we compute the relative $\mathcal{D}_B$ -efficiency. Following Holling and Schwabe (2011) and Kessels et al. (2011a), we define the $\mathcal{D}_B$ -efficiency of a design $\mathbf{X}$ relative to a design $\mathbf{X}^*$ as

$$Eff_B(\mathbf{X}, \mathbf{X}^*) = \exp\left(\frac{\mathcal{D}_B(\mathbf{X}) - \mathcal{D}_B(\mathbf{X}^*)}{k}\right). \quad (5)$$

Note that we have chosen the MNL framework as a basis for $\mathcal{D}_B$ -optimal partial profile design generation, even though we recognize that there are more sophisticated discrete choice models to date, the most well-known of which is the panel mixed logit

model (Train 2009). Unlike the MNL model, the panel mixed logit model captures the unobserved preference heterogeneity among respondents. As shown by Bliemer and Rose (2010), the generation of optimal designs for the panel mixed logit model requires a much larger computational effort than the generation of optimal designs for the MNL model. To avoid the long computing times, it is good practice to use an optimal design for the MNL model, since Bliemer and Rose (2010) also showed that an optimal design for the MNL model performs relatively well for the panel mixed logit model too. This result is due to the fact that the functional form of the MNL model is a building block of the panel mixed logit model and many other discrete choice models. As a result, using the MNL model for design generation is justified.

4 A two-stage variance balance approach for constructing partial profile designs

In this section, we introduce our new variance balance approach for generating efficient partial profile designs involving S choice sets of J profiles each. In every choice set of the designs, we vary the levels of t_v out of a total of t attributes, and keep the levels of the remaining $t - t_v$ attributes constant. The new approach is especially useful if not all attributes have the same number of levels.

Variance balance is a modification of Green’s BIBD approach that ensures an equal amount of information on, or equivalently, an equal estimation precision for all part-worths. We implemented two slightly different modifications to achieve balance in the variances of the part-worth estimates. The two variance balance algorithms have the same basic two-stage structure:

Stage 1: Create a master design to determine the t_v attributes whose levels vary in each of the S choice sets,

Stage 2: Create an attribute level design to set the levels of the t_v varying attributes in each choice set.

In Stage 1, we create a master design using one of the two variance balance methods by applying principles of the optimal design of experiments for linear models which we present in Section 4.1. These linear design principles also provided the basis for Kessels et al.’s (2011a) attribute balance method. In Stage 2, we create an attribute level design using the modified version of Meyer and Nachtsheim’s (1995) coordinate-exchange algorithm proposed by Kessels et al. (2011a) and reviewed in Section 4.2.

4.1 Stage 1: Creating the master design

4.1.1 An auxiliary two-way ANOVA model

In Section 2, we mentioned Green’s (1974) suggestion to use a balanced incomplete block design (BIBD) as a master design to determine the varying attributes in each choice set. To this end, the BIBD must have S blocks of t_v out of a total of t treatments each.

The family of BIBDs is well-known because a BIBD is $\mathcal{D}$ - and $\mathcal{A}$ -optimal for estimating the treatment effects α_i in the two-way ANOVA model of the type

$$Y_{is} = \alpha_i + \gamma_s + \varepsilon_{is}, \quad (6)$$

where Y_{is} is the response of treatment i in block s , α_i represents the average response for treatment i , γ_s is the effect of block s and ε_{is} is the random error corresponding to the response for treatment i in block s (where all random errors are assumed to be independently normally distributed with zero mean and variance σ^2). In matrix notation, the two-way ANOVA model (6) can be written as

$$\mathbf{Y} = \mathbf{Q}\boldsymbol{\alpha} + \mathbf{Z}\boldsymbol{\gamma} + \boldsymbol{\epsilon}, \quad (7)$$

where $\mathbf{Y}$ is a vector of $r = St_v$ responses, $\mathbf{Q}$ is an $r \times t$ matrix containing the treatment design, $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_t]'$ is the vector containing the average responses for all treatments, $\mathbf{Z}$ is the $r \times (S - 1)$ matrix that assigns the treatments to the blocks, $\boldsymbol{\gamma} = [\gamma_1, \dots, \gamma_{S-1}]'$ is the vector containing the block effects, and $\boldsymbol{\epsilon}$ is the random error vector.

A key feature of the BIBD is that it maximizes the determinant of the information matrix on the ANOVA model’s parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$, which, up to a proportionality constant, is equal to

$$\mathbf{N} = \begin{bmatrix} \mathbf{Q}'\mathbf{Q} & \mathbf{Q}'\mathbf{Z} \\ \mathbf{Z}'\mathbf{Q} & \mathbf{Z}'\mathbf{Z} \end{bmatrix}. \quad (8)$$

The BIBD is therefore $\mathcal{D}$ -optimal for the ANOVA model. At the same time, the BIBD minimizes the trace of the inverse of the information matrix, $\mathbf{N}^{-1}$, as well as the sum of the variances of the ordinary least squares estimates of $\alpha_1, \dots, \alpha_t$, $\sum_{i=1}^t \text{var}(\hat{\alpha}_i)$. The BIBD is therefore also $\mathcal{A}$ -optimal for the ANOVA model.

A technical problem with BIBDs is that they only exist for a limited number of specific combinations of numbers of choice sets and attributes, which makes them impractical for general-purpose use. A more general approach to construct a master design for the varying attributes, which works for any number of choice sets and attributes, is to seek $\mathcal{D}$ -optimal or $\mathcal{A}$ -optimal designs for the two-way ANOVA model. This is the approach used by Kessels et al. (2011a), who seek a $\mathcal{D}$ -optimal master design maximizing the determinant of the information matrix (8) for a given number of choice sets, attributes and varying attributes.

Using the $\mathcal{D}$ -optimality criterion for creating the master design implies that an equal amount of attention is paid to each attribute. Therefore, this approach is called attribute

balance. In experimental scenarios where not all attributes have the same number of levels, however, it makes sense to pay more attention to attributes that have many levels. For these scenarios, we introduce two weighted $\mathcal{A}$ -optimal design approaches for creating the master design and refer to these approaches as variance balance.

4.1.2 Variance balance

A limitation of attribute balance is that it ignores the fact that attributes may have differing numbers of levels. In such cases, the part-worths of the attributes with a larger number of levels are estimated less precisely than the part-worths of the attributes with fewer levels. Therefore, it makes sense to vary an attribute with a larger number of levels in more choice sets than an attribute with fewer levels, to obtain similar amounts of information on each part-worth estimate. To decide how often each of the t attributes has to be varied in the S choice sets, we create weighted $\mathcal{A}$ -optimal designs for the two-way ANOVA model (6), which we refer to as $\mathcal{A}_w$ -optimal master designs. An $\mathcal{A}_w$ -optimal master design minimizes

$$\mathcal{A}_w = \sum_{i=1}^t w_i \text{var}(\hat{\alpha}_i), \quad (9)$$

where w_i is the weight attached to the variance of the individual parameter estimator $\hat{\alpha}_i$ of the ANOVA model. The variances are the diagonal elements of the $t \times t$ -dimensional upper left hand submatrix of $\mathbf{N}^{-1}$, which, according to Theorem 8.5.11 in Harville (1997), can be computed as

$$\{\mathbf{Q}'(\mathbf{I}_r - \mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}')\mathbf{Q}\}^{-1}, \quad (10)$$

where $\mathbf{I}_r$ denotes the r -dimensional identity matrix.

The weights w_i in the weighted $\mathcal{A}$ -optimality criterion increase with the number of levels, d_i , of attribute i . By increasing that weight, we explicitize that we prefer having a larger amount of information on attribute i . In terms of the two-way ANOVA model, the larger weight expresses that we want to estimate the treatment effect α_i more precisely than other treatment effects. The resulting master design then uses treatment i more often than other treatments, which implies that the level of attribute i in the choice design is varied more often. In general, increasing an attribute's weight in the $\mathcal{A}_w$ -optimality criterion causes that attribute to vary in more choice sets.

We generate two classes of $\mathcal{A}_w$ -optimal master designs and refer to them as variance balance I and variance balance II designs. The variance balance I and II approaches differ in the exact weights utilized. The two sorts of weights we use play a crucial role in optimal design theory for linear models used to construct utility-neutral optimal designs for DCEs, as was done by Grasshoff et al. (2004) and Grossmann et al. (2006, 2009, 2013). The mathematical justification for the two types of weights is given in the Appendix. Note that in case all attributes have the same number of levels, the variance balance I and II

approaches reduce to the attribute balance approach.

Variance balance I

In the variance balance I approach, the weight w_i is the number of part-worth values for attribute i in the MNL model relative to the total number of part-worth values:

$$w_i = \frac{d_i - 1}{\sum_{i=1}^t d_i - t} = \frac{d_i - 1}{k}. \quad (11)$$

The intuition behind this weight formula is that this way of selecting the varying attributes spreads the information losses from the constant attributes equally over all k part-worth values. In the Appendix, we show that this weight is required to compute the information matrix of utility-neutral optimal partial profile designs.

Variance balance II

In the variance balance II approach, each weight is computed as

$$w_i = \frac{(d_i - 1)^2}{2d_i}. \quad (12)$$

In the Appendix, we show that this weight formula has been derived from the information matrix of utility-neutral optimal full profile designs.

4.1.3 Illustration of master designs

To compare the new stage 1 algorithms, variance balance I and II, with Kessels et al.'s (2011a) attribute balance algorithm, we generated master designs conforming the illustrative problem presented in Kessels et al. (2011a). The partial profile designs in this illustration are main-effects designs which consist of a single survey involving 15 choice sets with two profiles, three 2-level attributes, two 3-level attributes and one 5-level attribute. Because these six attributes have heterogeneous numbers of levels, we constructed $\mathcal{A}_w$ -optimal master designs for the selection of four varying attributes in each choice set and compared them to the $\mathcal{D}$ -optimal master design. All master designs appear in Table 4. The four varying attributes in each choice set are indicated by check marks. The remaining two attributes, shown in gray, are constant. The last line of the table contains the numbers of times each attribute is varying in a choice set.

The first master design in the table is the $\mathcal{D}$ -optimal master design generated by Kessels et al. (2011a) using the attribute balance approach. Each attribute varies in exactly ten choice sets and each of the 15 possible pairs of varying attributes appears in exactly six choice sets. The $\mathcal{D}$ -optimal master design is therefore a BIBD in this example.

The remaining two designs are the $\mathcal{A}_w$ -optimal master designs generated using variance balance I and II, respectively. As opposed to the $\mathcal{D}$ -optimal master design, the $\mathcal{A}_w$ -optimal

master designs take into account the numbers of levels of each attribute. Both $\mathcal{A}_w$ -optimal master designs hold pairs of 2-level and/or 3-level attributes constant in the choice sets. The level of the 5-level attribute varies all the time. Compared to variance balance II, variance balance I varies the level of a 2-level attribute in two more choice sets at the expense of a 3-level attribute.

We revisit this illustrative example in Section 5. To complete the partial profile design construction, we now discuss the stage 2 design algorithm to set the levels of the varying attributes.

4.2 Stage 2: Creating the attribute level design

In the last stage of the partial profile design construction, we determine the levels of the varying attributes in each choice set. We choose the levels of the constant attributes randomly because, for main-effects models, they have no effect on the information acquired from the experiment. That is, the information matrix (3) does not depend on the levels of the constant attributes in a choice set.

The algorithm we use to create the attribute level design for maximizing the $\mathcal{D}_B$ -optimality criterion (4) in Bayesian partial profile designs is a modified version of Meyer and Nachtsheim’s (1995) coordinate-exchange algorithm. Kessels et al. (2011a) created this modification because they identified a weakness in the coordinate-exchange algorithm when it comes to generating Bayesian partial profile designs. Because the algorithm does not require the levels of the varying attributes to actually change, they showed that the levels of one or more of these attributes remain fixed in certain choice sets of some partial profile designs. In these choice sets, the resulting number of varying attributes is smaller than t_v . Kessels et al. (2011a) modified the coordinate-exchange algorithm to avoid this undesirable behavior of the algorithm. Because the modification involves a restriction in the exchange procedure of the algorithm, Kessels et al. (2011a) referred to the entire two-stage design algorithm as the restricted partial profile design algorithm.

In this paper, we follow Kessels et al. (2011a) and also use the restricted coordinate-exchange algorithm. In our view, the most important justification for doing so is that, if the number of varying attributes is not the same in each choice set, then this results in some choice sets being easier to evaluate than others so that the error variance is unlikely to be constant across choice sets. Also, in DCEs where the profiles are displayed as columns of words or sentences and the constant attributes are dropped, the columns do not all have the same length if the number of varying attributes differs between choice sets. This is aesthetically jarring. Finally, Kessels et al. (2011a) showed that the restricted coordinate-exchange algorithm performs virtually equally well as the original algorithm.

5 A comparison of partial profile designs for an illustrative problem

We revisit the illustrative study by Kessels et al. (2011a) in which they compare a series of $\mathcal{D}_B$ -optimal partial profile designs to the corresponding $\mathcal{D}_B$ -optimal full profile design. Other benchmark designs in the study are the $\mathcal{D}$ -optimal utility-neutral designs with full and partial profiles. The partial profile designs constructed by Kessels et al. (2011a) are attribute balance designs. We extend the comparison study by constructing variance balance I and II designs for the Bayesian and utility-neutral design cases. We compute the losses in $\mathcal{D}_B$ -efficiency of the variance balance designs relative to the Bayesian full profile design and quantify to what extent the variance balance designs outperform the attribute balance designs.

As mentioned in Section 4.1.3, all designs are main-effects designs which consist of a single survey involving 15 choice sets with two profiles, three 2-level attributes, two 3-level attributes and one 5-level attribute. The partial profile designs hold the levels of one or two attributes constant. We review the setup of the comparison study in Section 5.1, and discuss the optimal designs and their $\mathcal{D}_B$ -efficiencies in Sections 5.2 and 5.3, respectively.

5.1 Setup of the comparison study

For the sake of direct comparison, we adopted the same setup of the comparison study as described in Kessels et al. (2011a). First, we used 2000 random starts of the restricted coordinate-exchange algorithm to generate the variance balance I and II designs. Second, we modeled the attribute levels using effects-type coding which constrains the part-worths associated with any given attribute to sum to zero. If d_i denotes the number of levels for attribute i , $i = 1, \dots, t$, effects-type coding implies that only the part-worths attached to the first $d_i - 1$ levels of that attribute are estimated explicitly. Therefore, the vector $\boldsymbol{\beta}$ contains $k = 11$ unknown part-worth values in the study. Third, for the construction of the Bayesian designs, we used the multivariate normal prior distribution $\pi(\boldsymbol{\beta}) = \mathcal{N}(\boldsymbol{\beta}|\boldsymbol{\beta}_0, \boldsymbol{\Sigma}_0)$, with prior mean

$$\boldsymbol{\beta}_0 = [-0.5, -0.5, -0.5, -0.5, 0, -0.5, 0, -0.5, -0.25, 0, 0.25]' \quad (13)$$

and prior variance-covariance matrix

$$\Sigma_0 = \begin{bmatrix} 0.16 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.16 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.16 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.16 & -0.08 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -0.08 & 0.16 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.16 & -0.08 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -0.08 & 0.16 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.16 & -0.04 & -0.04 & -0.04 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.04 & 0.16 & -0.04 & -0.04 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.04 & -0.04 & 0.16 & -0.04 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.04 & -0.04 & -0.04 & 0.16 \end{bmatrix}. \quad (14)$$

For the reason why this prior mean β_0 and prior variance-covariance matrix Σ_0 were chosen, we refer to Kessels et al. (2011a).

5.2 Optimal designs

5.2.1 Full profile design

Table 5 shows the $\mathcal{D}_B$ -optimal full profile design. This design has seven choice sets with one constant 2-level attribute. These constant attributes are marked in gray and can assume any possible attribute level. There are 6 attributes and 15 choice sets giving 90 possibilities for level overlap. Hence, there is about 8% of level overlap.

5.2.2 One constant attribute

Table 6 shows the $\mathcal{D}_B$ -optimal partial profile designs involving one constant attribute. It includes the (restricted) attribute balance design from Kessels et al. (2011a) and our newly generated variance balance design. We report only one variance balance design because the variance balance I and II weighting schemes produced the same master design in case one constant attribute was desired. The attribute balance design varies the levels of each attribute in either 12 or 13 choice sets. Hence, the $\mathcal{D}$ -optimal master design used to select the varying attributes is inevitably not a BIBD. The variance balance design varies the levels of each of the 2-level attributes in 10 choice sets. The levels of the two 3-level attributes and the 5-level attribute vary all the time. This causes some concern if one of the latter attributes is a dominant attribute, because it implies the risk that nothing will be learnt about the levels of the non-dominant attributes. The attribute balance design would be a better design option then.

5.2.3 Two constant attributes

Table 7 shows the $\mathcal{D}_B$ -optimal partial profile designs involving two constant attributes. The first partial profile design is the attribute balance design from Kessels et al. (2011a). The remaining designs are the variance balance I and II designs. The three designs are based on the master designs discussed in Section 4.1.3. Note that the variance balance designs do not hold the 5-level attribute constant in any choice set. If decision making is based on this attribute only, the variance balance designs are therefore not preferable.

5.2.4 Utility-neutral designs

Table 8 shows the $\mathcal{D}$ -optimal utility-neutral full profile design and Tables 9 and 10 show the $\mathcal{D}$ -optimal utility-neutral partial profile designs involving one and two constant attributes. The full profile design as well as the attribute balance designs are provided by Kessels et al. (2011a), but the variance balance designs are new. All utility-neutral designs are level balanced in the varying attributes within and across choice sets.

5.3 $\mathcal{D}_B$ -efficiencies

To evaluate the estimation performance of the variance balance I and II designs for the Bayesian and utility-neutral design cases, we computed the $\mathcal{D}_B$ -efficiencies of the designs relative to the Bayesian full profile design. We contrast these $\mathcal{D}_B$ -efficiencies with those of the (restricted) attribute balance designs provided by Kessels et al. (2011a). Table 11 shows the $\mathcal{D}_B$ -optimality criterion values and $\mathcal{D}_B$ -efficiencies of all designs.

Examining the $\mathcal{D}_B$ -optimal partial profile designs, we observe that the variance balance designs perform much better than the attribute balance designs in terms of $\mathcal{D}_B$ -efficiency. In the case of one constant attribute, the variance balance design hardly loses 2% in $\mathcal{D}_B$ -efficiency, while the attribute balance design loses about 10%. This means that, in order to obtain the same amount of information as from full profile testing, the variance balance design requires about 2% more respondents (computed as $1/0.9867 - 1$) and the attribute balance design requires about 10% more respondents. In the case of two constant attributes, the attribute balance design loses about 20% in $\mathcal{D}_B$ -efficiency, requiring 27% more respondents compared to full profile testing. The efficiency losses of the variance balance I and II designs are only half as large. There is a small difference in performance between the variance balance I and II designs, in favor of variance balance II.

Examining the $\mathcal{D}$ -optimal utility-neutral designs, the utility-neutral full profile design loses about 20% in $\mathcal{D}_B$ -efficiency compared to the Bayesian full profile design. When evaluating this efficiency loss, we need to take into account the fact that the outperformance of Bayesian designs over utility-neutral designs depends on the prior distribution used. The results of Kessels et al. (2011b-c), who performed an in-depth comparison of Bayesian and utility-neutral full profile designs, indicate that the efficiency loss of 20% of

the $\mathcal{D}$ -optimal utility-neutral full profile design is not extreme, but rather average.

For the utility-neutral partial profile designs with one and two constant attributes, we again observe that the variance balance designs outperform the attribute balance designs. Also, the difference in performance between the variance balance I and II designs is small. Like the utility-neutral full profile design, the utility-neutral variance balance designs yield efficiency losses of about 20% compared to the $\mathcal{D}_B$ -optimal full profile design. As a result, the estimation performance is unaffected by the constant attributes in the utility-neutral variance balance designs. Overall, the $\mathcal{D}_B$ -efficiencies of the utility-neutral full and partial profile designs are much smaller than those of the Bayesian designs, making the utility-neutral design options the worst to consider.

For this illustrative design problem, the use of variance balance II weights for constructing Bayesian partial profile designs results in the highest statistical efficiency in the case of two constant attributes. In the case of one constant attribute, there is no difference in efficiency between the designs produced by the two variance balance methods.

Given there are 15 choice sets in the $\mathcal{D}_B$ -optimal partial profile designs, the variance balance designs suffer from the drawback that they do not hold the 5-level attribute constant in any choice set using one or two constant attributes. Using one constant attribute, they also do not hold the 3-level attribute constant in any choice set. We therefore only recommend the variance balance designs if the attributes that vary all the time are not dominant. Otherwise, an attribute balance design should be used. The next section revisits the software development study in which the most statistically efficient design, obtained using the variance balance II weights, can be used without any problem because none of the attributes is dominant.

6 Follow-up designs for the software development experiment

We present three follow-up partial profile designs for the JMP graphics choice experiment discussed in Section 2. Table 12 shows the likelihood ratio tests of the attribute effects based on choice data collected from 239 respondents in the original experiment. The picture color and the frame line color attributes are not significant so we drop them from the follow-up experiment. The report color attribute is marginally significant ($p = 0.0924$). For this reason, we include it in the follow-up experiment. From the part-worth estimates of the simplified model in Table 13, we can compute the difference in part-worths between the bluish report color (estimated part-worth -0.112) and the white report color (estimated part-worth $0.125 = -(-0.112 - 0.018 + 0.005)$) to be 0.237. This difference was large enough to matter to the JMP development team assuming the effect is real. A decision to change the report color to white hung in the balance. Thus, a follow-up experiment was necessary to reduce the uncertainty about the part-worth estimates of

the report color.

The construction of $\mathcal{D}_B$ -optimal designs for the follow-up experiment requires the input of a prior mean and variance-covariance matrix for the parameters. We used the part-worth estimates in Table 13 and their variance-covariance matrix in Table 14 as the mean and variance of the multivariate normal prior distribution $\pi(\boldsymbol{\beta}) = \mathcal{N}(\boldsymbol{\beta}|\boldsymbol{\beta}_0, \boldsymbol{\Sigma}_0)$. The prior mean vector of part-worth estimates, $\boldsymbol{\beta}_0$, and its covariance matrix, $\boldsymbol{\Sigma}_0$, are 7-dimensional.

For the sake of illustration, we show three partial profile designs each generated using a different algorithm and consisting of one survey of 15 choice sets of two profiles. The designs have three varying attributes and two constant attributes. They appear in Table 15. The first attribute in the designs represents the report color having four levels. The remaining four attributes concern the graph color, frame all sides, outer graph rectangle and Y-axis title, each having two levels. The first partial profile design is the attribute balance design for which a $\mathcal{D}$ -optimal master design has been generated to select the varying attributes. The $\mathcal{D}$ -optimal master design varies each attribute in exactly nine choice sets, but not all pairs of varying attributes appear equally often. Therefore, the $\mathcal{D}$ -optimal master design is not a BIBD.

The remaining partial profile designs in Table 15 are the variance balance designs. In these designs, the 4-level attribute varies in many more choice sets than each of the 2-level attributes. The variance balance I design varies the levels of the 4-level attribute in 13 choice sets, while the variance balance II design varies them in all choice sets.

To evaluate the estimation performance of the $\mathcal{D}_B$ -optimal partial profile designs, we also generated the $\mathcal{D}_B$ -optimal full profile design, which is shown in Table 16. The full profile design does not hold any of the attributes constant. Table 17 shows the estimation performances of the Bayesian partial profile designs. Like in the illustrative study of Section 5, the variance balance II design has the highest $\mathcal{D}_B$ -efficiency, closely followed by the variance balance I design. The attribute balance design is the least efficient.

Although the variance balance II design has the highest statistical efficiency, as mentioned previously, it does not hold the 4-level attribute constant in any choice set. We would recommend the variance balance I design or the attribute balance design if this was a preliminary study. In that case, we would advise against varying any attribute in every choice set to avoid a failed study due to a dominant attribute. However, this is a follow-up experiment and the 4-level attribute is one that was least significant in the original study. It makes sense, therefore, to allow this attribute to vary in every choice set of the follow-up study. For the software development case study, we thus recommended using the most statistically efficient design using the variance balance II weights.

7 Summary and discussion

We presented two new methods for constructing $\mathcal{D}_B$ -optimal main-effects partial profile designs for DCEs. They provide alternative generalizations of Green’s (1974) balanced incomplete block design (BIBD) approach and Kessels et al.’s (2011a) attribute balance approach to select the varying attributes, and they account for the fact that attributes may have differing numbers of levels. A strength of our new methods is that they pay an equal amount of attention to each attribute level or each part-worth. We refer to them as variance balance I and II because they vary an attribute with a larger number of levels more often than an attribute with fewer levels to stabilize the variances of the individual part-worth estimates. The two variance balance methods differ in the way attributes with differing numbers of levels are weighted and reduce to the attribute balance method in case all attributes have the same number of levels. Both the attribute balance and the variance balance methods are used in the first stage of a two-stage partial profile design generation algorithm introduced by Kessels et al. (2011a). The methods provide a master design that determines the varying attributes and are complemented with a second-stage algorithm that sets the levels of these varying attributes.

The motivation for partial profile designs is twofold. First, by reducing the complexity of the choices in a DCE, partial profiles have the potential to prevent respondents from resorting to non-compensatory decision rules when the number of attributes under study is large. We refer to these rules as design-driven non-compensatory decision rules. Second, whether choices are complex or not, partial profiles can prove useful in situations in which respondents mainly care about one attribute or a small subset of the attributes. We refer to the resulting decision rules as preference-driven non-compensatory decision rules. For choice sets in which the level of a dominant attribute is varied, respondents will always choose the profile with the most attractive level for that attribute. If the level of the dominant attribute is, however, constant, then respondents have to trade off the remaining attributes. This yields useful information about the trade-offs between the levels of the non-dominant attributes.

Keeping certain attributes constant in each choice set also has a cost. It reduces the theoretical information content of each choice set compared to full profile testing. Using two comparison studies that involve differing numbers of attribute levels, we investigated the losses in estimation precision from using partial profiles. The results show that the losses in estimation precision are the smallest using variance balance II, closely followed by variance balance I. Attribute balance results in considerably higher losses in estimation precision. Note that this theoretical drawback of using partial profiles vanishes if using full profiles causes respondents to short cut the compensatory choice process and resort to design-driven non-compensatory decision rules. Also note that one should be cautious when using the variance balance methods because it is possible that they do not hold every attribute constant in at least a few choice sets. In that case, the investigator is

not protected against the impact of a potentially dominant attribute among the set of attributes that vary all the time. We verified that the results obtained from the two comparison studies hold for other design problems as well.

The partial profile design algorithms are flexible in that they can accommodate arbitrarily many attributes and varying attributes. Choice sets may have any number of profiles and, though the number of choice sets must be adequate to fit the underlying model, there can be as many choice sets as desired. These can be divided into separate surveys that are assigned to different respondents so that the respondents are not overburdened by having to make too many choices. Sándor and Wedel (2005) call such a design with different surveys a heterogeneous design and showed that it is statistically more efficient than a single homogeneous design that is assigned to every respondent. The efficiency gain from using a heterogeneous design accrues from different respondents being given different surveys, which causes more variation in the attributes, which in turn enables the variation in the part-worths to be captured more effectively. For the sake of illustration, the partial profile designs discussed in the comparison studies in this paper are small, consisting of 15 choice sets only, and are therefore homogeneous. For real-life DCEs, however, we recommend using heterogeneous designs. The design used in the original software development experiment discussed in this paper is a heterogeneous design.

Another reason for using heterogeneous choice designs with partial profiles is that, in the case of attribute dominance, a homogeneous design with partial profiles may not be adequate to estimate all part-worths of the remaining attributes. Consider, for example, the $\mathcal{D}_B$ -optimal attribute balance design with one constant attribute in Table 6. In this design, the first attribute is kept constant in only two of the 15 choice sets. If this attribute is dominant, then the two choice sets will not be sufficient to estimate the part-worths of the remaining attributes. It is therefore necessary to use a number of surveys that each consist of 15 different choice sets so that the part-worths of the remaining attributes can be estimated when the first attribute turns out to be dominant.

The research in this paper can be extended in the following direction. The two variance balance methods provide more efficient $\mathcal{D}_B$ -optimal partial profile designs for differing numbers of attribute levels than the attribute balance method. It would be interesting to investigate whether we can still improve the estimation performance of the partial profile designs using other types of master designs. Grossmann et al. (2009, 2013) presented master designs that are $\mathcal{D}$ -optimal for very specific utility-neutral design situations. We could use these master designs as a basis to find partial profile designs that perform better in terms of the $\mathcal{D}_B$ -optimality criterion. Eventually, the goal is then to find $\mathcal{D}_B$ -optimal master designs, which would depend on the prior distributions for each of the model’s parameters. Instead of the two-step algorithm we have used in this paper for the construction of partial profile designs, such a goal requires a one-step algorithm in which the master design and the attribute level design are optimized simultaneously, at the

expense of a large computing time. In addition to that, instead of using the $\mathcal{D}_B$ -optimality criterion, we could use another Bayesian optimality criterion derived from a choice model that explicitly accounts for the choice task complexity such as that in Sándor and Franses (2009) or Danthurebandara et al. (2011). This would increase the computing time even further, though it has the potential to demonstrate that partial profile designs are more efficient than full profile designs, after correcting for choice task complexity. We could also use alternative information matrices studied in Yu et al. (2012) who recommend using the generalized Fisher information matrix as a basis for determining Bayesian optimal designs for small sample DCEs.

To conclude, however, in our consulting sessions we noticed that practitioners consider the choice of a master design to be an important decision. Given the low computational cost for generating master designs using attribute balance or variance balance, the two-step algorithm offers the advantage that several master designs can be computed, compared and discussed on the fly during a consulting session. In any sensible one-step algorithm, the generation of the master design would be computationally very intensive because of its dependency on the prior distributions for each of the model’s parameters. Hence, it might not be possible to generate master designs on the fly during a consulting session with a one-step algorithm, and leave the practitioner with the uneasy feeling that the entire design generation is a black box.

Appendix. Motivation of variance balance methods

We show that the weight formulas in (11) and (12) for variance balance I and variance balance II, respectively, are derived from utility-neutral $\mathcal{D}$ -optimal design theory for DCEs with $J = 2$ profiles per choice set. We verified, however, that the same weight formulas apply regardless of the value of J .

Variance balance I

Using variance balance I, each weight, $w_i, i = 1, \dots, t$, increases linearly with the number of part-worth values associated with attribute i , as defined in (11). In a study using approximate design theory, Grossmann et al. (2006, 2009, 2013) found that the information matrix of a utility-neutral $\mathcal{D}$ -optimal partial profile design with choice sets of size two is based on a similar weight formula. In approximate design theory, the number of times a given choice set appears in a design is allowed to be fractional rather than being required to be an integer. This aids in proofs but limits applicability to very restrictive cases where the relative frequencies of different choice sets exactly match their optimal fractions. The information matrix of a utility-neutral $\mathcal{D}$ -optimal approximate design with two profiles

per choice set is block diagonal and equal to

$$\mathbf{M} = \frac{t_v}{k} \begin{bmatrix} (d_1 - 1)\mathbf{M}_1 & & 0 \\ & \ddots & \\ 0 & & (d_t - 1)\mathbf{M}_t \end{bmatrix}, \quad (\text{A.1})$$

where

$$\mathbf{M}_i = \frac{2}{d_i - 1} \begin{bmatrix} 2 & 1 & \cdots & 1 \\ 1 & 2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 1 \\ 1 & \cdots & 1 & 2 \end{bmatrix}. \quad (\text{A.2})$$

The matrix $\mathbf{M}_i$ has $d_i - 1$ rows and columns, where d_i is the number of levels of attribute i . Ignoring the constant, t_v , the multiplier of each $\mathbf{M}_i$ is $(d_i - 1)/k$, which is the weight defined in (11).

Variance balance II

Using variance balance II, each weight, $w_i, i = 1, \dots, t$, is computed as

$$w_i = \frac{(d_i - 1)^2}{2d_i}.$$

This weight formula is based on the work of Grasshoff et al. (2004) who, again using approximate design theory, provided the information matrix of a utility-neutral $\mathcal{D}$ -optimal full profile design with choice sets of size two. The information matrix is block diagonal and given by

$$\mathbf{M} = \frac{1}{4} \begin{bmatrix} \mathbf{M}_1 & & 0 \\ & \ddots & \\ 0 & & \mathbf{M}_t \end{bmatrix}. \quad (\text{A.3})$$

The weight formula is derived from the diagonal elements of the inverse of $\mathbf{M}_i$:

$$\mathbf{M}_i^{-1} = \frac{(d_i - 1)}{2d_i} \begin{bmatrix} d_i - 1 & -1 & \cdots & -1 \\ -1 & d_i - 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & -1 \\ -1 & \cdots & -1 & d_i - 1 \end{bmatrix}. \quad (\text{A.4})$$

Each weight, w_i , corresponds to the diagonal elements of matrix $\mathbf{M}_i^{-1}$, which are all identical.

Acknowledgements

Roselinde Kessels is grateful to the Research Foundation Flanders (FWO) for her postdoctoral research grant. The authors would like to thank Heiko Grossmann for providing valuable comments on the variance balance methods for partial profile design construction.

References

- Abel, R. J. R. and Greig, M. (2007). BIBDs with small block size, in *The CRC Handbook of Combinatorial Designs*, 2nd ed., Ch. II.3, Colbourn, C. J. and Dinitz, J. H., eds. Boca Raton (Florida): CRC Press, 72–79.
- Bliemer, M. C. J. and Rose, J. M. (2010). Construction of experimental designs for mixed logit models allowing for correlation across choice observations, *Transportation Research B* **44**: 720–734.
- Bradlow, E. T. (2005). Current issues and a ‘wish list’ for conjoint analysis, *Applied Stochastic Models in Business and Industry* **21**: 319–323.
- Caussade, S., Ortúzar, J. de D., Rizzi, L. I. and Hensher, D. A. (2005). Assessing the influence of design dimensions on stated choice experiment estimates, *Transportation Research B* **39**: 621–640.
- Chrzan, K. (2010). Using partial profile choice experiments to handle large numbers of attributes, *International Journal of Market Research* **52**: 827–840.
- Danthurebandara, V. M., Yu, J. and Vandebroek, M. (2011). Effect of choice complexity on design efficiency in conjoint choice experiments, *Journal of Statistical Planning and Inference* **141**: 2276–2286.
- Dellaert, B. G. C., Donkers, B. and van Soest, A. (2012). Complexity effects in choice experiment-based models, *Journal of Marketing Research* **49**: 424–434.
- DeShazo, J. R. and Fermo, G. (2002). Designing choice sets for stated preference methods: the effects of complexity on choice consistency, *Journal of Environmental Economics and Management* **44**: 123–143.
- Grasshoff, U., Grossmann, H., Holling, H. and Schwabe, R. (2003). Optimal paired comparison designs for first-order interactions, *Statistics* **37**: 373–386.
- Grasshoff, U., Grossmann, H., Holling, H. and Schwabe, R. (2004). Optimal designs for main effects in linear paired comparison models, *Journal of Statistical Planning and Inference* **126**: 361–376.
- Green, P. E. (1974). On the design of choice experiments involving multi-factor alternatives, *Journal of Consumer Research* **1**: 61–68.
- Green, P. E. and Srinivasan, V. (1990). Conjoint analysis in marketing: new developments

- with implications for research and practice, *Journal of Marketing* **54**: 3–19.
- Grossmann, H., Grasshoff, U. and Schwabe, R. (2009). Approximate and exact optimal designs for paired comparisons of partial profiles when there are two groups of factors, *Journal of Statistical Planning and Inference* **139**: 1171–1179.
- Grossmann, H., Grasshoff, U. and Schwabe, R. (2013). A catalogue of designs for partial profiles in paired comparison experiments with three groups of factors, *Statistics*, published online: <http://dx.doi.org/10.1080/02331888.2013.844146>.
- Grossmann, H., Holling, H., Grasshoff, U. and Schwabe, R. (2006). Optimal designs for asymmetric linear paired comparisons with a profile strength constraint, *Metrika* **64**: 109–119.
- Harville, D. A. (1997). *Matrix Algebra from a Statistician's Perspective*, New York: Springer.
- Hensher, D. A. (2006). How do respondents process stated choice experiments? Attribute consideration under varying information load, *Journal of Applied Econometrics* **21**: 861–878.
- Hensher, D. A. and Rose, J. M. (2009). Simplifying choice through attribute preservation or non-attendance: Implications for willingness to pay, *Transportation Research E* **45**: 583–590.
- Holling, H. and Schwabe, R. (2011). Discussion on “The usefulness of Bayesian optimal designs for discrete choice experiments”, *Applied Stochastic Models in Business and Industry* **27**: 189–192.
- Kessels, R., Goos, P. and Vandebroek, M. (2010). Optimal two-level conjoint designs with constant attributes in the profile sets, *Journal of Statistical Planning and Inference* **140**: 3035–3046.
- Kessels, R., Jones, B. and Goos, P. (2011a). Bayesian optimal designs for discrete choice experiments with partial profiles, *Journal of Choice Modelling* **4**: 52–74.
- Kessels, R., Jones, B., Goos, P. and Vandebroek, M. (2011b). The usefulness of Bayesian optimal designs for discrete choice experiments, *Applied Stochastic Models in Business and Industry* **27**: 173–188.
- Kessels, R., Jones, B., Goos, P. and Vandebroek, M. (2011c). Rejoinder: The usefulness of Bayesian optimal designs for discrete choice experiments, *Applied Stochastic Models in Business and Industry* **27**: 197–203.
- Louviere, J. J. (2005). Comment on “Current issues and a ‘wish list’ for conjoint analysis”, *Applied Stochastic Models in Business and Industry* **21**: 325–326.
- Louviere, J. J., Pihlens, D. and Carson, R. (2011). Design of discrete choice experiments: A discussion of issues that matter in future applied research, *Journal of Choice Modelling* **4**: 1–8.

- Louviere, J. J., Street, D. J. and Burgess, L. (2003). A 20+ years' retrospective on choice experiments, in *Marketing Research and Modeling: Progress and Prospects*, Ch. 8, Wind, Y. and Green, P. E., eds. New York: Kluwer Academic Press, 201–214.
- Mathon, R. and Rosa, A. (2007). $2-(v, k, \lambda)$ designs of small order, in *The CRC Handbook of Combinatorial Designs*, 2nd ed., Ch. II.1, Colbourn, C. J. and Dinitz, J. H., eds. Boca Raton (Florida): CRC Press, 25–58.
- Meyer, R. K. and Nachtsheim, C. J. (1995). The coordinate-exchange algorithm for constructing exact optimal experimental designs, *Technometrics* **37**: 60–69.
- Sándor, Z. and Franses, P. H. (2009). Consumer price evaluations through choice experiments, *Journal of Applied Econometrics* **24**: 517–535.
- Sándor, Z. and Wedel, M. (2001). Designing conjoint choice experiments using managers' prior beliefs, *Journal of Marketing Research* **38**: 430–444.
- Sándor, Z. and Wedel, M. (2005). Heterogeneous conjoint choice designs, *Journal of Marketing Research* **42**: 210–218.
- Schwabe, R., Grasshoff, U., Grossmann, H. and Holling, H. (2003). Optimal 2^K paired comparison designs for partial profiles, in *PROBASTAT2002, Proceedings of the Fourth International Conference on Mathematical Statistics*, Smolenice 2002, Tatra Mountains Mathematical Publications, Vol. 26, Stulajter, F. and Wimmer, G., eds., 79–86.
- Scott, A. (2002). Identifying and analysing dominant preferences in discrete choice experiments: An application in health care, *Journal of Economic Psychology* **23**: 383–398.
- Swait, J. and Adamowicz, W. (2001). Choice environment, market complexity and consumer behavior: A theoretical and empirical approach for incorporating decision complexity into models of consumer choice, *Organizational Behaviour and Human Decision Processes* **86**: 141–167.
- Train, K. E. (2009). *Discrete Choice Models with Simulation*, 2nd ed., Cambridge, U.K.: Cambridge University Press.
- Yu, J., Goos, P. and Vandebroek, M. (2012). A comparison of different Bayesian design criteria for setting up stated preference studies, *Transportation Research B* **46**: 789–807.

Tables

Table 1: Attributes and attribute levels for the JMP graphics choice experiment.

Attribute	Attribute levels			
1. Report background color	bluish	creamish	light gray	white
2. Picture background color	contrast	same		
3. Graph background color	light gray	white		
4. Frame line color	black	gray		
5. Frame all sides	no	yes		
6. Outer graph rectangle	no	yes		
7. Y-axis title	horizontal above	vertical left		

Table 2: Partial profiles with attribute levels for the two output displays of the example choice set in Figure 1.

Attribute	Profile 1	Profile 2
1. Report background color	creamish	creamish
2. Picture background color	contrast	same
3. Graph background color	light gray	light gray
4. Frame line color	black	black
5. Frame all sides	no	yes
6. Outer graph rectangle	no	yes
7. Y-axis title	vertical left	vertical left

Table 3: Subset of seven choice sets from the 120 choice sets of the partial profile design used in the JMP graphics choice experiment.

Choice set	Report color	Picture color	Graph color	Frame line color	Frame all sides	Graph rectangle	Y-axis title
1	bluish	same	white	gray	no	no	vertical
1	bluish	same	light gray	gray	yes	no	horizontal
2	light gray	contrast	light gray	gray	yes	no	horizontal
2	light gray	contrast	white	black	yes	yes	horizontal
3	creamish	contrast	light gray	black	no	no	vertical
3	creamish	same	light gray	black	yes	yes	vertical
4	light gray	same	white	gray	yes	no	horizontal
4	light gray	contrast	white	black	yes	no	vertical
5	creamish	contrast	white	black	no	no	vertical
5	bluish	contrast	white	black	no	yes	horizontal
6	light gray	contrast	white	gray	yes	no	horizontal
6	creamish	contrast	white	black	no	no	horizontal
7	light gray	same	white	black	no	yes	horizontal
7	white	contrast	light gray	black	no	yes	horizontal

Table 4: Master designs generated with attribute balance and variance balance I and II. They involve 15 choice sets, three 2-level attributes, two 3-level attributes and one 5-level attribute. They have four varying attributes, indicated by check marks, and two constant attributes, shown in gray.

	Attribute balance						Variance balance I						Variance balance II					
Choice set	Attributes with levels						Attributes with levels						Attributes with levels					
	2	2	2	3	3	5	2	2	2	3	3	5	2	2	2	3	3	5
1				✓	✓	✓				✓	✓	✓				✓	✓	✓
2			✓		✓	✓				✓	✓	✓				✓	✓	✓
3			✓	✓		✓			✓		✓	✓				✓	✓	✓
4			✓	✓	✓				✓		✓	✓			✓	✓	✓	✓
5			✓	✓	✓	✓			✓		✓	✓			✓	✓	✓	✓
6	✓				✓	✓			✓	✓		✓			✓	✓	✓	✓
7	✓			✓		✓			✓	✓		✓			✓	✓	✓	✓
8	✓			✓	✓		✓				✓	✓			✓		✓	✓
9	✓			✓	✓	✓					✓	✓			✓	✓	✓	✓
10	✓	✓				✓					✓	✓			✓	✓	✓	✓
11	✓	✓			✓				✓			✓			✓		✓	✓
12	✓	✓			✓	✓			✓	✓		✓			✓	✓		✓
13	✓	✓	✓								✓	✓			✓			✓
14	✓	✓	✓			✓				✓		✓			✓		✓	✓
15	✓	✓	✓	✓					✓			✓			✓			✓
<i>Freq</i>	10	10	10	10	10	10	8	8	7	11	11	15	7	7	7	12	12	15

Table 5: $\mathcal{D}_B$ -optimal full profile design involving 15 choice sets with 2 profiles, three 2-level attributes, two 3-level attributes and one 5-level attribute.

Choice set	Attributes					
	1	2	3	4	5	6
1	1	2	*	2	3	1
1	2	1	*	3	1	5
2	*	1	2	3	2	2
2	*	2	1	2	1	5
3	*	1	2	2	3	2
3	*	2	1	3	2	1
4	2	1	1	1	3	3
4	1	2	2	3	1	2
5	1	1	*	2	2	5
5	2	2	*	1	3	1
6	1	2	1	3	3	2
6	2	1	2	2	2	4
7	1	2	2	1	2	5
7	2	1	1	3	1	3
8	2	1	1	2	3	5
8	1	2	2	1	2	3
9	1	2	*	1	1	4
9	2	1	*	3	2	1
10	1	1	2	1	3	5
10	2	2	1	2	1	4
11	2	*	1	1	3	4
11	1	*	2	2	1	1
12	2	1	2	1	1	1
12	1	2	1	3	2	3
13	2	2	1	2	2	2
13	1	1	2	3	1	4
14	1	1	1	3	3	4
14	2	2	2	1	1	3
15	2	*	1	1	2	2
15	1	*	2	2	3	3
<i>Freq</i>	<i>13</i>	<i>13</i>	<i>12</i>	<i>15</i>	<i>15</i>	<i>15</i>

Table 6: $\mathcal{D}_B$ -optimal partial profile designs involving 15 choice sets with 2 profiles, three 2-level attributes, two 3-level attributes and one 5-level attribute. The designs have five varying attributes and one constant attribute. Variance balance I and II lead to the same design in this scenario.

Choice set	Attribute balance						Variance balance					
	Attributes						Attributes					
	1	2	3	4	5	6	1	2	3	4	5	6
1	*	1	1	2	3	4	*	1	1	3	3	2
1	*	2	2	1	1	2	*	2	2	1	2	5
2	*	2	1	3	1	1	*	1	2	2	2	3
2	*	1	2	1	2	3	*	2	1	1	1	5
3	1	*	1	3	3	2	*	1	2	3	1	2
3	2	*	2	1	1	1	*	2	1	1	2	4
4	1	*	2	1	3	2	*	2	1	2	1	3
4	2	*	1	3	2	3	*	1	2	3	3	1
5	2	*	1	1	1	4	*	1	2	3	1	4
5	1	*	2	2	2	1	*	2	1	2	3	2
6	1	2	*	2	3	3	2	*	1	3	1	5
6	2	1	*	3	2	4	1	*	2	1	3	1
7	1	1	*	3	3	5	1	*	2	1	3	3
7	2	2	*	2	1	4	2	*	1	2	2	1
8	1	2	*	2	1	5	1	*	2	2	1	5
8	2	1	*	1	3	1	2	*	1	3	2	4
9	1	2	1	*	2	4	1	*	1	3	3	3
9	2	1	2	*	1	5	2	*	2	1	2	2
10	1	1	2	*	3	4	1	*	2	2	1	4
10	2	2	1	*	2	2	2	*	1	1	3	5
11	2	1	1	*	1	3	1	2	*	3	2	1
11	1	2	2	*	2	4	2	1	*	1	3	4
12	1	2	1	1	*	5	2	1	*	2	3	1
12	2	1	2	2	*	2	1	2	*	3	2	2
13	2	1	1	2	*	5	1	1	*	2	2	5
13	1	2	2	3	*	3	2	2	*	1	1	1
14	1	1	2	2	2	*	2	1	*	1	2	3
14	2	2	1	3	3	*	1	2	*	2	3	4
15	1	1	2	3	1	*	2	2	*	2	1	3
15	2	2	1	1	3	*	1	1	*	1	2	5

<i>Freq</i>	13	12	12	12	13	13	10	10	10	15	15	15
-------------	----	----	----	----	----	----	----	----	----	----	----	----

Table 7: $\mathcal{D}_B$ -optimal partial profile designs involving 15 choice sets with 2 profiles, three 2-level attributes, two 3-level attributes and one 5-level attribute. The designs have four varying attributes and two constant attributes.

Choice set	Attribute balance						Variance balance I						Variance balance II					
	Attributes						Attributes						Attributes					
	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6
1	*	*	2	1	1	5	*	*	1	2	3	5	*	*	2	2	2	1
1	*	*	1	3	3	3	*	*	2	1	1	4	*	*	1	1	3	3
2	*	1	*	2	3	5	*	*	1	3	1	3	*	*	2	1	1	5
2	*	2	*	3	1	1	*	*	2	1	3	1	*	*	1	3	2	4
3	*	2	1	*	2	1	*	2	*	1	1	5	*	*	2	2	3	2
3	*	1	2	*	1	2	*	1	*	2	2	3	*	*	1	3	1	5
4	*	2	2	2	*	2	*	2	*	1	2	1	*	1	*	2	3	4
4	*	1	1	3	*	5	*	1	*	3	3	2	*	2	*	3	2	3
5	*	2	1	2	1	*	*	1	*	3	3	1	*	2	*	3	1	1
5	*	1	2	3	3	*	*	2	*	2	2	4	*	1	*	2	2	5
6	1	*	*	2	3	1	*	2	1	*	3	1	*	2	*	1	2	2
6	2	*	*	1	2	3	*	1	2	*	2	5	*	1	*	3	1	4
7	2	*	1	*	1	3	*	1	1	3	*	4	*	2	*	2	1	4
7	1	*	2	*	2	4	*	2	2	2	*	2	*	1	*	3	3	1
8	1	*	2	2	*	3	1	*	*	2	3	3	*	1	2	*	2	3
8	2	*	1	3	*	2	2	*	*	3	2	2	*	2	1	*	3	2
9	2	*	1	1	3	*	2	*	*	2	1	1	2	*	*	1	2	4
9	1	*	2	3	1	*	1	*	*	1	3	4	1	*	*	2	3	3
10	2	1	*	*	1	4	1	*	*	2	1	4	2	*	*	1	3	1
10	1	2	*	*	2	2	2	*	*	1	2	3	1	*	*	3	1	2
11	1	2	*	3	*	5	1	*	1	*	2	5	1	*	1	*	2	5
11	2	1	*	1	*	4	2	*	2	*	1	3	2	*	2	*	1	3
12	2	1	*	2	2	*	2	*	1	1	*	2	2	*	1	2	*	5
12	1	2	*	1	3	*	1	*	2	3	*	1	1	*	2	1	*	4
13	1	2	1	*	*	4	2	1	*	*	1	5	1	*	2	3	*	5
13	2	1	2	*	*	1	1	2	*	*	2	2	2	*	1	2	*	3
14	2	2	2	*	1	*	1	2	*	3	*	3	2	1	*	*	1	2
14	1	1	1	*	2	*	2	1	*	2	*	4	1	2	*	*	3	4
15	1	1	2	1	*	*	1	1	2	*	*	2	2	1	*	3	*	2
15	2	2	1	2	*	*	2	2	1	*	*	4	1	2	*	2	*	1
<i>Freq</i>	10	10	10	10	10	10	8	8	7	11	11	15	7	7	7	12	12	15

Table 8: $\mathcal{D}$ -optimal utility-neutral full profile design involving 15 choice sets with 2 profiles, three 2-level attributes, two 3-level attributes and one 5-level attribute.

Choice set	Attributes					
	1	2	3	4	5	6
1	2	1	2	1	2	3
1	1	2	1	3	1	5
2	1	1	1	3	1	2
2	2	2	2	2	2	1
3	2	2	1	1	1	1
3	1	1	2	3	2	3
4	1	1	1	2	3	3
4	2	2	2	3	2	2
5	1	2	1	1	3	2
5	2	1	2	3	1	5
6	1	1	2	1	3	4
6	2	2	1	2	1	3
7	1	1	2	2	1	1
7	2	2	1	1	2	5
8	1	2	1	2	2	4
8	2	1	2	3	3	1
9	1	2	2	1	3	5
9	2	1	1	2	2	2
10	1	1	2	2	2	5
10	2	2	1	3	3	3
11	1	1	1	1	2	1
11	2	2	2	2	3	4
12	2	1	2	1	1	2
12	1	2	1	3	2	1
13	1	2	2	3	1	4
13	2	1	1	2	3	5
14	1	2	2	1	1	3
14	2	1	1	3	3	4
15	1	2	2	2	3	2
15	2	1	1	1	1	4
<i>Freq</i>	<i>15</i>	<i>15</i>	<i>15</i>	<i>15</i>	<i>15</i>	<i>15</i>

Table 9: $\mathcal{D}$ -optimal utility-neutral partial profile designs involving 15 choice sets with 2 profiles, three 2-level attributes, two 3-level attributes and one 5-level attribute. The designs have five varying attributes and one constant attribute. Variance balance I and II lead to the same design in this scenario.

Choice set	Attribute balance						Variance balance					
	Attributes						Attributes					
	1	2	3	4	5	6	1	2	3	4	5	6
1	*	2	2	2	3	1	*	1	2	3	1	2
1	*	1	1	1	1	3	*	2	1	1	3	1
2	*	1	1	2	3	5	*	2	2	2	3	2
2	*	2	2	1	1	2	*	1	1	3	1	5
3	2	*	2	2	2	2	*	1	2	1	3	3
3	1	*	1	1	1	1	*	2	1	2	2	5
4	2	*	2	1	3	3	*	1	2	1	3	5
4	1	*	1	3	2	2	*	2	1	2	2	3
5	1	*	2	2	1	4	*	2	2	2	1	5
5	2	*	1	3	2	5	*	1	1	3	2	1
6	2	1	*	2	1	1	1	*	1	1	3	2
6	1	2	*	3	3	4	2	*	2	2	1	1
7	2	1	*	1	3	2	1	*	1	1	1	4
7	1	2	*	3	2	3	2	*	2	3	2	3
8	1	2	*	2	1	5	1	*	2	1	2	4
8	2	1	*	3	3	4	2	*	1	2	3	2
9	1	2	1	*	3	2	2	*	1	1	1	3
9	2	1	2	*	2	4	1	*	2	3	2	4
10	1	1	2	*	3	3	2	*	1	3	3	4
10	2	2	1	*	1	4	1	*	2	1	2	1
11	2	2	2	*	1	5	2	2	*	1	1	4
11	1	1	1	*	2	1	1	1	*	2	3	5
12	2	2	2	3	*	1	2	2	*	1	2	5
12	1	1	1	2	*	4	1	1	*	2	1	1
13	1	1	2	1	*	5	1	2	*	3	1	2
13	2	2	1	2	*	3	2	1	*	2	3	4
14	2	2	1	1	2	*	2	2	*	3	3	1
14	1	1	2	3	1	*	1	1	*	2	2	3
15	1	2	2	1	2	*	1	2	*	3	3	3
15	2	1	1	3	1	*	2	1	*	1	2	2

<i>Freq</i>	13	12	12	12	13	13	10	10	10	15	15	15
-------------	----	----	----	----	----	----	----	----	----	----	----	----

Table 10: $\mathcal{D}$ -optimal utility-neutral partial profile designs involving 15 choice sets with 2 profiles, three 2-level attributes, two 3-level attributes and one 5-level attribute. The designs have four varying attributes and two constant attributes.

Choice set	Attribute balance						Variance balance I						Variance balance II					
	Attributes						Attributes						Attributes					
	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6
1	*	*	2	1	3	3	*	*	2	1	1	1	*	*	2	1	2	1
1	*	*	1	2	1	2	*	*	1	3	3	2	*	*	1	2	1	4
2	*	1	*	2	3	4	*	*	2	1	3	4	*	*	2	1	2	4
2	*	2	*	1	1	1	*	*	1	3	1	3	*	*	1	3	3	1
3	*	1	2	*	2	2	*	*	2	3	2	3	*	*	1	3	2	3
3	*	2	1	*	1	4	*	*	1	2	1	2	*	*	2	2	3	2
4	*	1	1	2	*	3	*	2	*	3	3	1	*	1	*	3	2	2
4	*	2	2	1	*	4	*	1	*	2	2	3	*	2	*	2	1	1
5	*	1	1	3	1	*	*	1	*	2	3	1	*	1	*	3	1	4
5	*	2	2	2	2	*	*	2	*	1	1	4	*	2	*	2	3	3
6	1	*	*	3	2	5	*	1	*	3	2	1	*	1	*	2	2	1
6	2	*	*	1	3	2	*	2	*	2	3	3	*	2	*	3	1	5
7	2	*	2	*	1	3	*	2	1	*	2	5	*	1	2	*	3	4
7	1	*	1	*	3	1	*	1	2	*	1	3	*	2	1	*	2	2
8	1	*	1	1	*	5	2	*	*	2	1	5	*	2	2	3	*	2
8	2	*	2	3	*	1	1	*	*	1	3	3	*	1	1	1	*	5
9	2	*	1	1	2	*	2	*	*	3	1	4	1	*	*	1	3	3
9	1	*	2	2	1	*	1	*	*	1	2	5	2	*	*	2	2	5
10	1	1	*	*	2	4	2	*	*	2	2	4	2	*	*	1	1	3
10	2	2	*	*	3	5	1	*	*	1	1	5	1	*	*	3	3	5
11	2	1	*	2	*	1	2	*	1	1	*	1	2	*	*	1	3	5
11	1	2	*	3	*	2	1	*	2	3	*	5	1	*	*	3	1	1
12	1	1	*	1	1	*	2	1	*	*	3	5	2	*	1	*	3	4
12	2	2	*	3	3	*	1	2	*	*	2	2	1	*	2	*	1	5
13	2	1	2	*	*	5	1	2	*	2	*	1	1	*	1	1	*	2
13	1	2	1	*	*	3	2	1	*	1	*	2	2	*	2	3	*	3
14	1	1	2	*	3	*	1	1	2	*	*	2	1	2	*	*	2	4
14	2	2	1	*	2	*	2	2	1	*	*	3	2	1	*	*	1	2
15	2	1	1	3	*	*	1	1	1	*	*	4	1	1	*	2	*	3
15	1	2	2	2	*	*	2	2	2	*	*	2	2	2	*	1	*	1
<i>Freq</i>	10	10	10	10	10	10	8	8	7	11	11	15	7	7	7	12	12	15

Table 11: $\mathcal{D}_B$ -optimality criterion values and $\mathcal{D}_B$ -efficiencies of the optimal designs.

Criterion	Number of cst attributes	Balance	$\mathcal{D}_B$ -value	$\mathcal{D}_B$ -efficiency
$\mathcal{D}_B$ -optimal	0		11.24378	100.00%
$\mathcal{D}_B$ -optimal	1	attribute	10.20127	90.96%
$\mathcal{D}_B$ -optimal	1	variance	11.09641	98.67%
$\mathcal{D}_B$ -optimal	2	attribute	8.64183	78.94%
$\mathcal{D}_B$ -optimal	2	variance I	10.03448	89.59%
$\mathcal{D}_B$ -optimal	2	variance II	10.24672	91.33%
$\mathcal{D}$ -optimal utility-neutral	0		8.77159	79.87%
$\mathcal{D}$ -optimal utility-neutral	1	attribute	7.16088	68.99%
$\mathcal{D}$ -optimal utility-neutral	1	variance	8.59738	78.62%
$\mathcal{D}$ -optimal utility-neutral	2	attribute	6.36122	64.15%
$\mathcal{D}$ -optimal utility-neutral	2	variance I	8.33708	76.78%
$\mathcal{D}$ -optimal utility-neutral	2	variance II	8.63521	78.89%

Table 12: Likelihood ratio tests of the attribute effects in the JMP graphics choice experiment.

Attribute	L-R ChiSquare	DF	Prob>ChiSq
Graph color	133.339	1	<.0001*
Frame all sides	110.715	1	<.0001*
Y-axis title	34.173	1	<.0001*
Graph rectangle	12.707	1	0.0004*
Report color	6.432	3	0.0924
Frame line color	0.641	1	0.4234
Picture color	0.001	1	0.9779

Table 13: Part-worth estimates and standard errors for the JMP graphics choice experiment obtained using a maximum likelihood approach for the simplified model.

Model term	Estimate	Standard error
Report color[bluish]	-0.112	0.057
Report color[creamish]	-0.018	0.054
Report color[light gray]	0.005	0.067
Graph color[light gray]	-0.313	0.028
Frame all sides[no]	-0.306	0.030
Graph rectangle[no]	-0.098	0.028
Y-axis title[horizontal]	-0.164	0.027

Table 14: Variance-covariance matrix of the part-worth estimates in Table 13 for the JMP graphics choice experiment.

	Report color [bluish]	Report color [creamish]	Report color [light gray]	Graph color [light gray]	Frame all sides[no]	Graph rectangle[no]	Y-axis title [horizontal]
Report color[bluish]	0.00329	-0.00083	-0.00168	-0.00004	0.00022	0.00000	-0.00010
Report color[creamish]	-0.00083	0.00290	-0.00097	0.00013	-0.00022	0.00012	0.00002
Report color[light gray]	-0.00168	-0.00097	0.00453	-0.00005	-0.00001	0.00003	0.00004
Graph color[light gray]	-0.00004	0.00013	-0.00005	0.00077	-0.00009	0.00002	-0.00005
Frame all sides[no]	0.00022	-0.00022	-0.00001	-0.00009	0.00088	-0.00009	-0.00004
Graph rectangle[no]	0.00000	0.00012	0.00003	0.00002	-0.00009	0.00076	-0.00008
Y-axis title[horizontal]	-0.00010	0.00002	0.00004	-0.00005	-0.00004	-0.00008	0.00072

Table 15: $\mathcal{D}_B$ -optimal follow-up partial profile designs for the JMP graphics choice experiment.

Choice set	Attribute balance				Variance balance I				Variance balance II						
	Report color	Graph color	Frame all sides	Graph rectangle	Y-axis title	Report color	Graph color	Frame all sides	Graph rectangle	Y-axis title	Report color	Graph color	Frame all sides	Graph rectangle	Y-axis title
1	*	*	yes	yes	vertical	*	*	yes	yes	vertical	creamish	*	*	no	vertical
1	*	*	no	no	horizontal	*	*	no	no	horizontal	white	*	*	yes	horizontal
2	*	white	*	yes	vertical	*	light gray	no	no	*	creamish	*	*	yes	vertical
2	*	light gray	*	no	horizontal	*	white	yes	yes	*	bluish	*	*	no	horizontal
3	*	light gray	no	*	vertical	light gray	*	*	yes	horizontal	white	*	no	*	vertical
3	*	white	yes	*	horizontal	bluish	*	*	no	vertical	creamish	*	yes	*	horizontal
4	*	light gray	no	*	horizontal	light gray	*	*	yes	vertical	light gray	*	yes	*	horizontal
4	*	white	yes	*	vertical	white	*	*	no	horizontal	bluish	*	no	*	vertical
5	*	white	no	yes	*	white	*	no	no	vertical	light gray	*	no	*	horizontal
5	*	light gray	yes	no	*	creamish	*	yes	*	horizontal	white	*	yes	*	vertical
6	*	light gray	yes	yes	*	bluish	*	yes	*	horizontal	creamish	*	no	yes	*
6	*	white	no	no	*	creamish	*	no	*	vertical	bluish	*	yes	no	*
7	white	*	*	yes	horizontal	creamish	*	no	yes	*	creamish	*	no	no	*
7	creamish	*	*	no	vertical	light gray	*	yes	no	*	bluish	*	yes	yes	*
8	white	*	no	no	vertical	white	*	no	yes	*	bluish	light gray	*	*	vertical
8	light gray	*	yes	yes	horizontal	creamish	*	yes	no	*	white	white	*	*	horizontal
9	creamish	*	yes	*	vertical	white	white	*	*	horizontal	light gray	light gray	*	*	vertical
9	bluish	*	no	*	vertical	bluish	light gray	*	*	vertical	creamish	white	*	*	horizontal
10	light gray	*	no	yes	*	bluish	white	*	*	horizontal	light gray	white	*	*	vertical
10	bluish	*	yes	no	*	light gray	light gray	*	*	vertical	bluish	light gray	*	*	horizontal
11	light gray	*	yes	no	*	white	white	*	*	vertical	white	light gray	*	yes	*
11	bluish	*	no	yes	*	light gray	light gray	*	*	horizontal	light gray	white	*	no	*
12	bluish	white	*	*	horizontal	light gray	white	*	no	*	white	white	*	no	*
12	white	light gray	*	*	vertical	creamish	light gray	*	yes	*	light gray	light gray	*	yes	*
13	light gray	light gray	*	*	vertical	creamish	white	*	no	*	white	light gray	*	no	*
13	white	white	*	*	horizontal	bluish	light gray	*	yes	*	bluish	white	*	yes	*
14	creamish	light gray	*	yes	*	bluish	white	no	*	*	light gray	light gray	yes	*	*
14	light gray	white	*	no	*	white	light gray	yes	*	*	bluish	white	no	*	*
15	creamish	white	no	*	*	light gray	white	no	*	*	light gray	white	no	*	*
15	bluish	light gray	yes	*	*	white	light gray	yes	*	*	creamish	light gray	yes	*	*
Freq	9	9	9	9	9	13	8	8	8	8	15	8	7	7	8

Table 16: $\mathcal{D}_B$ -optimal follow-up full profile design for the JMP graphics choice experiment.

Choice set	Report color	Graph color	Frame all sides	Graph rectangle	Y-axis title
1	creamish	white	no	yes	horizontal
1	white	light gray	yes	no	vertical
2	creamish	white	no	no	vertical
2	white	light gray	yes	yes	horizontal
3	white	white	no	yes	vertical
3	light gray	light gray	yes	no	horizontal
4	creamish	white	no	no	horizontal
4	light gray	light gray	yes	yes	vertical
5	creamish	light gray	yes	yes	vertical
5	light gray	white	no	no	horizontal
6	bluish	white	yes	no	vertical
6	white	light gray	no	yes	horizontal
7	light gray	white	yes	yes	horizontal
7	white	light gray	no	no	vertical
8	bluish	light gray	no	yes	vertical
8	light gray	white	yes	no	horizontal
9	bluish	light gray	yes	no	horizontal
9	white	white	no	yes	vertical
10	bluish	white	no	yes	horizontal
10	creamish	light gray	yes	no	vertical
11	bluish	light gray	yes	no	horizontal
11	light gray	white	no	yes	vertical
12	white	white	no	no	horizontal
12	bluish	light gray	yes	yes	vertical
13	light gray	light gray	no	yes	vertical
13	white	white	yes	no	horizontal
14	bluish	white	yes	yes	horizontal
14	light gray	light gray	no	no	vertical
15	bluish	white	no	no	vertical
15	creamish	light gray	yes	yes	horizontal
<i>Freq</i>	<i>15</i>	<i>15</i>	<i>15</i>	<i>15</i>	<i>15</i>

Table 17: $\mathcal{D}_B$ -optimality criterion values and $\mathcal{D}_B$ -efficiencies of the optimal designs.

Number of cst attributes	Balance	$\mathcal{D}_B$ -value	$\mathcal{D}_B$ -efficiency
0		13.82160	100.00%
2	attribute	10.65022	63.57%
2	variance I	11.29531	69.70%
2	variance II	11.46489	71.41%

Figures

Figure 1: Example choice set used in the JMP graphics choice experiment.