DEPARTMENT OF ENVIRONMENT,
TECHNOLOGY AND TECHNOLOGY MANAGEMENT

A comparison of partial profile designs for discrete choice experiments with an application in software development

Roselinde Kessels, Bradley Jones & Peter Goos

UNIVERSITY OF ANTWERP
Faculty of Applied Economics

EARNED EXCELLENCE
AACSB ACCREDITED
EST. 1916
THE BEST BUSINESS SCHOOLS IN THE WORLD

Stadscampus
Prinsstraat 13, B.226
BE-2000 Antwerpen
Tel. +32 (0)3 265 40 32
Fax +32 (0)3 265 47 99
http://www.ua.ac.be/tew

FACULTY OF APPLIED ECONOMICS

DEPARTMENT OF ENVIRONMENT,
TECHNOLOGY AND TECHNOLOGY MANAGEMENT

A comparison of partial profile designs for discrete choice experiments with an application in software development

Roselinde Kessels, Bradley Jones & Peter Goos

RESEARCH PAPER 2012-004
FEBRUARY 2012

University of Antwerp, City Campus, Prinsstraat 13, B-2000 Antwerp, Belgium
Research Administration – room B.226
phone: (32) 3 265 40 32
fax: (32) 3 265 47 99
e-mail: joeri.nys@ua.ac.be

D/2012/1169/004

A comparison of partial profile designs for discrete choice experiments with an application in software development

Roselinde Kessels*
Bradley Jones†
Peter Goos‡§

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

†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. 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 professor of statistics at Erasmus Universiteit Rotterdam, Erasmus School of Economics, Postbus 1738, 3000 DR Rotterdam, The Netherlands. Tel.: +32 (0)3 220 40 59. E-mail: peter.goos@ua.ac.be

§Roselinde Kessels is grateful to the Fund for Scientific Research Flanders 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.

A comparison of partial profile designs for discrete choice experiments with an application in software development

Abstract

In a discrete choice experiment, each respondent chooses the best product or service sequentially from many groups or choice sets of alternative goods. The alternatives, called profiles, are described by level combinations from a set of predefined attributes. Respondents sometimes make their choices on the basis of only one dominant attribute rather than making trade-offs among all the attributes. For example, in studies involving price as an attribute, respondents may always choose the profile with the lowest price. Also, a choice task including many attributes may encourage respondent decisions that are not fully compensatory. To thwart these behaviors, the investigator can hold the levels of some of the attributes constant in every choice set. The resulting designs are called partial profile designs. In this paper, we construct D-optimal partial profile designs for estimating main-effects models. We use a Bayesian design algorithm that integrates the D-optimality criterion over a prior distribution of likely parameter values. To determine the constant attributes in each choice set, we provide three alternative generalizations of an approach that makes use of balanced incomplete block designs. Each of our three generalizations constructs partial profile designs accommodating attributes with any number of levels and allowing flexibility in the numbers of choice sets and constant attributes. We show results from an actual experiment in software development performed using one of these algorithms. Finally, we compare the algorithms with respect to their statistical efficiency and ability to avoid failures due to the presence of a dominant attribute.

Keywords: discrete choice experiments, Bayesian D-optimal design, partial profiles, lexicographic choice behavior, balanced incomplete block design, coordinate-exchange algorithm

1 Introduction

Discrete choice experiments (DCEs) are a standard tool in marketing for quantifying consumer preferences (Louviere et al. 2003). Given a set of predefined product attributes, DCEs identify those attributes that matter most and indicate the most appealing levels for these attributes. DCEs are also called conjoint choice or stated choice experiments. The term "conjoint" is *à propos* because levels of different attributes are combined or *con*sidered *joint*ly in a product profile or alternative. These profiles describe either existing products or hypothetical and possibly prospective products. A DCE consists of a select number of profiles grouped into choice sets. Each choice set represents a hypothetical market from which respondents indicate the product 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) and Chrzan (2010), among others, argue that 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 choice set. In Section 2, we describe a DCE in software development where showing the levels of the constant attributes is required.

Currently, there is some controversy around 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. For example, the cereal aisle in a supermarket confronts people with many product categories and plenty of attribute information. Still, most shoppers make their breakfast choices without undo stress. In opposition to this view, several studies exploring the complexity and cognitive burden associated with DCEs provide evidence that a large number of attributes has a detrimental effect on the ability to choose, contributing to an increased error variance and parameter differences (see, e.g., Swait and Adamowicz 2001; DeShazo and Fermo 2002; Caussade et al. 2005; Hensher 2006).

Discrete choice models treat the utility of a given profile as the sum of the individual utilities or part-worths of the attributes. Implicit in this formulation is the expectation that respondents make compensatory decisions. This means that an attractive level of one or more attributes can compensate for an unattractive level of another attribute in the respondent's mind. However, one consequence of presenting many attributes in a DCE is that respondents make choices based on a hierarchical or lexicographic preference ordering of the attributes. In other words, the levels of the most important attribute(s) may totally dictate the preferred profile in any choice set (see, e.g., Scott 2002; Hensher 2006; Hensher and Rose 2009). This sole focus on the levels of the most important attribute(s) shortcuts the compensatory decision process. This kind of non-compensatory decision behavior then yields little or no information about the secondary attributes. The part-worth estimates for the secondary attributes in such cases are rarely useful.

The fact that respondents may 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, rendering simplifying decision strategies unnecessary. Even better, a single attribute strategy is impossible if the levels of the key attribute remain constant in some choice sets. Thus, whether a DCE employs many attributes or just a few, partial profiles can prove useful in the presence of a dominant attribute. For choice sets in which the level of the 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.

We advocate using partial profiles to help prevent respondents from exploiting non-compensatory decision strategies, including attribute dominance, so that the compensatory discrete choice models remain valid. This motivation for partial profiles is stronger than the traditional motivation of reducing the cognitive burden associated with requiring respondents to evaluate a large number of attributes. There is, however, also a downside to partial profiles. In theory, if all the respondents employ compensatory decision making, then fielding DCEs with partial profiles instead of full profiles results in less precise parameter estimates (Kessels et al. 2010).

Early theoretical development of design for DCEs made use of orthogonal arrays commonly associated with models that are linear in the unknown parameters. The discrete choice model is nonlinear in the parameters leading to an apparent mismatch of design and model. However, if one sets all the parameters in a choice model to zero, this transforms the nonlinear design problem into one that is equivalent to a linear design problem. Zero parameters result in zero utilities for any prospective profile, which explains the use of the term, utility-neutral design, for this approach. In their book, Street and Burgess (2007) compile an extensive set of utility-neutral full profile designs for specific combinations of numbers of attribute levels and choice sets using optimal design theory for linear models (see, e.g., Atkinson et al. 2007).

Rose and Bliemer (2009) point out that the Street and Burgess designs may promote lexicographic choice behavior. This is because their utility-neutral full profile designs vary the levels of every attribute in each choice set allowing a dominant attribute to govern every choice. This underscores the potential benefit of using partial profile designs. So far, the literature for DCEs using partial profiles has almost exclusively focused on utility-neutral designs involving choice sets with two profiles (Grasshoff et al. 2003, 2004; Grossmann et al. 2006, 2009). Moreover, the most recent utility-neutral optimal designs created by Grossmann et al. (2009) only allow for two 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 and varying attributes. This approach is therefore not applicable for a wide variety of practical problems.

A fundamental aspect of nonlinear models in general and DCEs in particular is that the information content of any design depends on the unknown parameters (Atkinson and Haines 1996). DCEs constructed to be optimal for one specific set of parameter values are said to be locally optimal (Huber and Zwerina 1996). The utility-neutral optimal designs are optimal for one specific set of parameter values, namely the set of all zero values, so they belong to the class of locally optimal designs.

Bayesian design methodology provides a solution to the nonlinear design problem that is robust to misspecification of the parameters. Introduced in the choice design literature by Sándor and Wedel (2001), Bayesian design methodology assumes a prior distribution of likely parameter values and optimizes the design over that distribution. In this way, it incorporates uncertainty about the proposed parameters into the problem formulation, which is more realistic than any locally optimal design approach. Many researchers have implemented the Bayesian design approach to construct full profile designs for DCEs (see, e.g., Sándor and Wedel 2001, 2002, 2005; Kessels et al. 2006, 2008, 2009, 2011b; Ferrini and Scarpa 2007; Yu et al. 2009, 2011; Bliemer and Rose 2010). Our aim is to extend this work to partial profile designs, going beyond just blindly applying a full profile design method to the construction of partial profile designs, as advised by Kessels et al. (2011a). We adopt a Bayesian $\mathcal{D}$-optimality criterion for the main-effects multinomial logit (MNL) model and call the resulting partial profile designs $\mathcal{D}_B$-optimal to emphasize the fact that we are making use of the Bayesian design idea.

To generate $\mathcal{D}_B$-optimal partial profile designs, we propose three two-stage design algorithms. In the first stage, we generate a master design that determines the attributes whose levels vary in each choice set, while in the second stage, we generate an attribute level design that sets the levels of the varying attributes. The three algorithms differ in the methods for determining the varying attributes in each choice set. They provide alternative solutions that generalize Green's (1974) idea to use a balanced incomplete block design (BIBD) for selecting the varying attributes. As a matter of fact, the algorithms produce efficient partial profile designs accommodating attributes with any number of levels and allowing flexibility in the numbers of choice sets and varying attributes. This offers marketing researchers full flexibility when designing DCEs and breaks the status quo since Green's (1974) paper.

The outline of the remainder of the paper is as follows. Section 2 describes an actual DCE in software development with partial profiles created using one of our design algorithms. 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 partial profile design algorithms to generate the Bayesian designs. In Section 5, we construct and compare different $\mathcal{D}_B$-optimal partial profile designs for an illustrative problem. We study their performance relative to full profile testing and also use utility-neutral designs as benchmarks. 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. Section 7 concludes the paper and highlights some further research possibilities.

2 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.

<Insert Figure 1 about here>

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, so 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.

<Insert Tables 1-2 about here>

To obtain precise part-worth estimates, the JMP team generated a Bayesian $\mathcal{D}$-optimal partial profile design (see Section 3) containing 120 choice sets. For each group of eight respondents the 120 choice sets were randomly reshuffled into eight surveys of 15 choice sets each.

The selection of the three varying attributes in each choice set was based on Green's (1974) idea to use a balanced incomplete block design (BIBD) for specifying the attributes whose levels vary in each choice set. 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. 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. The goal of the JMP development team was to balance 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 partial profile design. The BIBD master design for these seven choice sets appears in Table 4. The varying attributes form a BIBD 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.

A problem with Green's BIBD approach is that BIBDs only exist for specific combinations of numbers of blocks and treatments. For instance, it is not possible to generate a BIBD with 120 blocks containing $t_v = 3$ out of $t = 7$ different treatments. One way to get around this problem is to adapt the choice design situation to fit an available BIBD master design. Fortunately, there is a better solution that adapts the master design to the choice design situation. Given that a BIBD is a special case of a $\mathcal{D}$-optimal design, the BIBD approach can be generalized by maximizing the $\mathcal{D}$-optimality criterion for a very specific ANOVA model. Since a BIBD is $\mathcal{D}$-optimal, this $\mathcal{D}$-optimal design method yields a BIBD if one exists for a given situation. Otherwise, the $\mathcal{D}$-optimal design approximates the balancing structure of a BIBD as closely as possible. Therefore, the JMP team used a $\mathcal{D}$-optimal design as master design to determine the varying attributes in the 120 choice sets. In this paper, we use the term attribute balance to describe this $\mathcal{D}$-optimal design method for determining the varying attributes in each choice set and provide construction details in Section 4.

<Insert Tables 3-4 about here>

3 The multinomial logit framework

The multinomial logit (MNL) model assumes that respondents to a DCE belong to a target group of consumers with homogeneous product preferences. The model employs random utility theory which describes the utility that a respondent attaches to profile j ($j = 1, ..., J$) in choice set s ($s = 1, ..., S$) as the sum of a systematic and a stochastic component:

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

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

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

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

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

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

with $\mathbf{X}_s = [\mathbf{x}'_{js}]_{j=1,\ldots,J}$ the submatrix of $\mathbf{X}$ corresponding to choice set s, $\mathbf{p}_s = [p_{1s}, \ldots, p_{Js}]'$ and $\mathbf{P}_s = \text{diag}\,[p_{1s}, \ldots, p_{Js}]$. Huber and Zwerina (1996), Sándor and Wedel (2001), Kanninen (2002), Kessels et al. (2006, 2009, 2011b) and Scarpa and Rose (2008), among others, implemented different design criteria or functions of the information matrix (3) for constructing optimal designs. This task is far from trivial since the information on $\boldsymbol{\beta}$ depends on the unknown part-worth values in $\boldsymbol{\beta}$ through the probabilities p_{js} so that part-worth values are required before it is possible to construct optimal designs. To deal with this dependency on $\boldsymbol{\beta}$, one can use a single prior guess, $\boldsymbol{\beta}_P$, in a locally optimal design approach. The most popular design criterion in this approach is the local $\mathcal{D}$-optimality criterion or $\mathcal{D}_P$-optimality criterion, which we define as

$$\mathcal{D}_P = \log\left|\mathbf{M}\left(\mathbf{X}, \boldsymbol{\beta}_P\right)\right|. \tag{4}$$

The design that maximizes the $\mathcal{D}_P$-criterion is the $\mathcal{D}_P$-optimal design for the MNL model (2).

However, as shown by Sándor and Wedel (2001), locally optimal designs may provide imprecise part-worth estimates if the prior part-worth vector is misspecified. The locally optimal designs are only guaranteed to work well if the true part-worth vector $\boldsymbol{\beta}$ is close to the one specified when constructing the design. A more robust design solution is a Bayesian strategy that averages the design criterion over a prior distribution of likely part-worth values, $\pi(\boldsymbol{\beta})$. 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 $\boldsymbol{\Sigma}_0$. As opposed to locally optimal designs, Bayesian optimal designs perform well for a broad range of part-worth vectors $\boldsymbol{\beta}$ (Kessels et al. 2011b,c; Rose 2011). To generate them, we use the Bayesian $\mathcal{D}$-optimality criterion or $\mathcal{D}_B$-optimality criterion, which we define as

$$\mathcal{D}_B = \int_{\mathcal{R}^k} \log\left|\mathbf{M}\left(\mathbf{X}, \boldsymbol{\beta}\right)\right| \pi(\boldsymbol{\beta}) \mathrm{d}\boldsymbol{\beta}. \tag{5}$$

The design that maximizes the $\mathcal{D}_B$-criterion is the $\mathcal{D}_B$-optimal design for the MNL model (2). Kessels et al. (2011b) were the first to use this definition of the $\mathcal{D}_B$-criterion to generate $\mathcal{D}_B$-optimal designs. It differs from the $\mathcal{D}_B$-optimality criterion used in most of the literature on optimal choice design (see, e.g., Sándor and Wedel 2001, 2005; Kessels et al. 2006, 2008, 2009; Bliemer and Rose 2010) because of the logarithmic transformation of the determinant. This transformation ensures that, in a Bayesian information theoretic sense, the design that maximizes the $\mathcal{D}_B$-optimality criterion (5) also maximizes the expected Shannon information (Chaloner and Verdinelli 1995; Atkinson et al. 2007). A practical advantage of the logarithmic transformation is that it makes the $\mathcal{D}_B$-criterion less sensitive to part-worth vectors resulting in very large determinant values.

To compare the estimation performance of a partial profile design to a full profile design, we compute the relative $\mathcal{D}_B$-efficiency. We define the $\mathcal{D}_B$-efficiency of a design $\mathbf{X}$ relative to a design $\mathbf{X}^*$ as

$$\mathit{Eff}_B\left(\mathbf{X}, \mathbf{X}^*\right) = \exp\left(\frac{\mathcal{D}_B(\mathbf{X}) - \mathcal{D}_B(\mathbf{X}^*)}{k}\right). \tag{6}$$

Also Holling and Schwabe (2011) proposed using Equation (6) to compute the relative $\mathcal{D}_B$-efficiency of two designs.

Rather than using a Bayesian design approach, some researchers have transformed the nonlinear design problem for the MNL model (2) into a linear one by creating locally optimal designs assuming that the probabilities for all J profiles in each choice set are equal to $1/J$, or equivalently that $\boldsymbol{\beta} = \mathbf{0}$ (see, e.g., Street and Burgess 2007; Grossmann et al. 2009). This assumption reflects the fact that respondents have no preference for any of the profiles in a choice set. Therefore, the designs are called utility-neutral designs. The information matrix (3) of such designs is, up to a proportionality constant, equal to

$$\mathbf{M}(\mathbf{X}) = \mathbf{X}'\mathbf{X} - \sum_{s=1}^{S} J^{-1} \left(\mathbf{X}_s' \mathbf{1}_J\right) \left(\mathbf{1}_J' \mathbf{X}_s\right), \tag{7}$$

where $\mathbf{1}_J$ is a J-dimensional vector of ones (see, e.g., Kessels et al. 2011b). To construct $\mathcal{D}$-optimal utility-neutral designs, we maximize the linear $\mathcal{D}$-optimality criterion,

$$\mathcal{D} = \left| \mathbf{X}'\mathbf{X} - \sum_{s=1}^{S} J^{-1} \left(\mathbf{X}_s' \mathbf{1}_J\right) \left(\mathbf{1}_J' \mathbf{X}_s\right) \right|. \tag{8}$$

In Section 5, we present a comparison study in which we use $\mathcal{D}$-optimal utility-neutral designs as benchmark designs.

4 Partial profile design algorithms

In this section, we provide the details of our algorithms for generating partial profile designs. These algorithms all have the same basic two-stage structure:

Stage 1: Creating a master design to determine the attributes whose levels vary in each choice set,

Stage 2: Creating an attribute level design to set the levels of these varying attributes.

We present three algorithms for the stage 1 design problem and one algorithm for the stage 2 design problem. We divide the algorithms for stage 1 into two types:

Attribute balance: A generalization of Green's BIBD approach for making all attributes vary the same number of times in the design,

Variance balance: A modification of Green's BIBD approach that ensures an equal amount of information on, or equivalently, an equal estimation precision of all part-worths. This modification is especially useful if not all attributes have the same number of levels. We implemented two slightly different modifications to achieve balance in the variances of the part-worth estimates.

For creating the master design, both the attribute balance and the two variance balance methods make use of principles of the optimal design of experiments for linear models. We introduce these linear design principles and describe their application for creating the master design in Section 4.1. We discuss the algorithm for creating the attribute level design 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 as many blocks as there are choice sets, and the number of treatments within each block of the BIBD has to be equal to the number of varying attributes in a choice set. In other words, 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 in the two-way ANOVA model of the type

$$Y_{is} = \alpha_i + \gamma_s + \varepsilon_{is}, \tag{9}$$

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 (9) can be written as

$$\mathbf{Y} = \mathbf{Q}\boldsymbol{\alpha} + \mathbf{Z}\boldsymbol{\gamma} + \boldsymbol{\epsilon}, \tag{10}$$

where $\mathbf{Y}$ is a vector of $r = St_v$ responses, $\mathbf{Q}$ is an $r \times t$ matrix containing the treatment design, $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_t]'$ 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, $\boldsymbol{\gamma} = [\gamma_1, \ldots, \gamma_{S-1}]'$ is the vector containing the block effects, and $\boldsymbol{\epsilon}$ is a 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

$$\mathbf{N} = \begin{bmatrix} \mathbf{Q}'\mathbf{Q} & \mathbf{Q}'\mathbf{Z} \\ \mathbf{Z}'\mathbf{Q} & \mathbf{Z}'\mathbf{Z} \end{bmatrix}. \tag{11}$$

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 $\alpha_1, ..., \alpha_t$, $\sum_{i=1}^{t} \text{var}(\hat{\alpha}_i)$. 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. We use the former type of optimal designs for the two-way ANOVA model in scenarios where we want to pay an equal amount of attention to each attribute. We refer to this situation as attribute balance. In experimental scenarios where not all attributes have the same number of levels, it makes sense, however, to pay more attention to attributes that have many levels. For these scenarios, we use a weighted $\mathcal{A}$-optimal design for the two-way ANOVA model as a master design. We refer to this situation as variance balance.

4.1.2 Attribute balance

Our attribute balance approach generalizes Green's BIBD approach in the sense that it works for any number of choice sets, any number of attributes and any number of varying attributes, while paying an equal amount of attention to all attributes, independent of the numbers of levels of each attribute. To this end, we use a computerized search to find the master design that maximizes the $\mathcal{D}$-optimality criterion for the two-way ANOVA model. Such a $\mathcal{D}$-optimal design can always be found, even in situations where a BIBD does not exist. Hence, by using a $\mathcal{D}$-optimal master design, we are no longer hindered by the restrictive nature of BIBDs.

The attractive feature of a BIBD is that it balances the number of times an attribute varies in the choice design (first-order balance). In addition, it also balances the number of times an attribute varies in the same choice set as another attribute (second-order balance or pairwise balance). The $\mathcal{D}$-optimal master design which we advocate achieves the same kinds of balance for suitable choices of the numbers of choice sets (S), attributes (t) and varying attributes (t_v), i.e. in situations where a BIBD exists. If no BIBD exists, the $\mathcal{D}$-optimal master designs come as close to providing both first- and second-order balance as possible.

The $\mathcal{D}$-optimal master design maximizes the determinant of the $(r+S-1)\times(r+S-1)$ information matrix (11), $|\mathbf{N}|$, for given values of S, t_v and t. To save computing time when searching for the $\mathcal{D}$-optimal master design, we apply Theorem 13.3.8 in Harville (1997) to find that

$$\begin{aligned} |\mathbf{N}| &= \begin{vmatrix} \mathbf{Q}'\mathbf{Q} & \mathbf{Q}'\mathbf{Z} \\ \mathbf{Z}'\mathbf{Q} & \mathbf{Z}'\mathbf{Z} \end{vmatrix}, \\ &= |\mathbf{Z}'\mathbf{Z}|\,|\mathbf{Q}'\mathbf{Q} - \mathbf{Q}'\mathbf{Z}\,(\mathbf{Z}'\mathbf{Z})^{-1}\,\mathbf{Z}'\mathbf{Q}|\,. \end{aligned} \tag{12}$$

The matrix $\mathbf{Z}$ only depends on S and t_v, so that it is constant for a given design problem. Also, $|\mathbf{Z}'\mathbf{Z}|$ and $\mathbf{Z}\,(\mathbf{Z}'\mathbf{Z})^{-1}\,\mathbf{Z}'$ are constant, so that we only need to compute them once. If we denote the r-dimensional identity matrix by $\mathbf{I}_r$, then, to achieve attribute balance, we can maximize the t-dimensional determinant

$$|\mathbf{Q}'(\mathbf{I}_r - \mathbf{Z}\,(\mathbf{Z}'\mathbf{Z})^{-1}\,\mathbf{Z}')\mathbf{Q}| \tag{13}$$

instead of the $(r+S-1)$-dimensional determinant $|\mathbf{N}|$. This approach leads to computational time savings that increase with the number of blocks, S.

4.1.3 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 start again from the two-way ANOVA model (9). This time, we create a weighted $\mathcal{A}$-optimal design for the ANOVA model, which we refer to as an $\mathcal{A}_w$-optimal master design. The $\mathcal{A}_w$-optimal master design minimizes

$$\mathcal{A}_w = \sum_{i=1}^{t} w_i \text{var}(\hat{\alpha}_i), \tag{14}$$

where w_i are the weights attached to the variances of the individual parameter estimators, $\hat{\alpha}_i$. These 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 most efficiently as

$$\{\mathbf{Q}'(\mathbf{I}_r - \mathbf{Z}\,(\mathbf{Z}'\mathbf{Z})^{-1}\,\mathbf{Z}')\mathbf{Q}\}^{-1}. \tag{15}$$

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). The mathematical justification for the two types of weights is given in the Appendix.

Variance balance I

In the variance balance I approach, the weight w_i is proportional to the number of part-worth values associated with attribute i:

$$w_i = \frac{d_i - 1}{\sum_{i=1}^{t} d_i - t} = \frac{d_i - 1}{k}. \tag{16}$$

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

$$w_i = \frac{(d_i - 1)^2}{2d_i}. \tag{17}$$

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

4.1.4 Illustration of master designs

To provide insight in the stage 1 algorithms, attribute balance and variance balance I and II, we used them to generate master designs for an illustrative problem. We constructed the $\mathcal{D}$-optimal and $\mathcal{A}_w$-optimal master designs for a partial profile design involving 15 choice sets, three 2-level attributes, two 3-level attributes and one 5-level attribute. At this stage of the choice design process, the number of profiles, J, in every choice set is irrelevant, though any value is applicable. The master designs appear in Table 5. Given six attributes under study, they have four varying attributes in each choice set 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 constant in a choice set.

<Insert Table 5 about here>

The first master design is the $\mathcal{D}$-optimal master design generated using the attribute balance approach. It perfectly satisfies first-order balance because each attribute varies in ten choice sets. Also, it perfectly satisfies second-order balance because 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 will revisit this illustrative example in Section 5 where we present a comparison study of partial profile designs. To complete the partial profile design construction, we now describe 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 an attribute level design. 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) is not a function of the constant attributes in a choice set.

The algorithm we use to create the attribute level design is Meyer and Nachtsheim's (1995) coordinate-exchange algorithm. It starts by randomly generating a level for each varying attribute in each profile of the design. For each of these levels or "coordinates", the algorithm then tries all possible levels and chooses the level corresponding to the best value of the optimality criterion selected for partial profile design evaluation. We use the $\mathcal{D}_B$-optimality criterion (5) to create Bayesian partial profile designs and the linear $\mathcal{D}$-optimality criterion (8) to create utility-neutral partial profile designs. The algorithm runs several times through the attribute level design and restarts for a given number of times.

Kessels et al. (2011b) also used the coordinate-exchange algorithm with the $\mathcal{D}_B$-optimality criterion (5) as objective function to construct Bayesian full profile designs. We follow their methods to compute the $\mathcal{D}_B$-criterion value of a partial profile design. We use a multivariate normal prior distribution $\pi(\boldsymbol{\beta}) = \mathcal{N}(\boldsymbol{\beta}|\boldsymbol{\beta}_0, \boldsymbol{\Sigma}_0)$ and approximate the k-dimensional integral related to this distribution using the quadrature scheme of Gotwalt et al. (2009) and Gotwalt (2010). Yu et al. (2010) showed that this quadrature method outperforms any other methods for Bayesian objective function evaluation.

Kessels et al. (2011a), however, pointed out a weakness of using the coordinate-exchange algorithm to generate 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 that case, more than the required number of attributes remain constant in such choice sets. We modified the coordinate-exchange algorithm so that it must vary the levels of the varying attributes in the choice sets. This affects the $\mathcal{D}_B$-optimal design outcome in some design situations.

The main reason why we do not allow additional constant attributes selected by the $\mathcal{D}_B$-optimality criterion (5) in Bayesian partial profile designs 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. Lastly, it turns out that in cases where $\mathcal{D}_B$-efficiency losses are incurred because all varying attributes have to change, these losses are most often negligible. This is because the partial profile designs already have $t_c = t - t_v$ attributes constant so that the number of additional constant attributes selected by the $\mathcal{D}_B$-optimality criterion (5) is rarely large.

In the next section, we evaluate Bayesian partial profile designs against utility-neutral partial profile designs. For main-effects utility-neutral designs, the linear $\mathcal{D}$-optimality criterion (8) always varies the levels of the varying attributes in the choice sets so that the partial profiles have exactly t_c constant attributes. As a result, there is no need for a restriction in the coordinate-exchange algorithm that excludes additional constant attributes from the choice sets. Moreover, the main-effects utility-neutral designs are level balanced in the varying attributes within and over all choice sets given appropriate numbers of attribute levels, choice sets and profiles per choice set. For other design situations, the main-effects designs are level balanced to the largest possible extent.

5 A comparison of partial profile designs for an illustrative problem

This section compares a series of $\mathcal{D}_B$-optimal partial profile designs to the corresponding $\mathcal{D}_B$-optimal full profile design. Other benchmark designs are the $\mathcal{D}$-optimal utility-neutral designs with full and partial profiles. 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 examine the loss in $\mathcal{D}_B$-efficiency due to using partial profiles instead of full profiles. We describe 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

We generated the Bayesian and utility-neutral optimal designs using 2000 random starts of the coordinate-exchange algorithm. 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, $i = 1, ..., t$, effects-type coding requires that only the part-worths attached to the first $d_i - 1$ levels of that attribute need to be estimated as the part-worth attached to the last level d_i automatically results. Therefore, the vector $\boldsymbol{\beta}$ contains $k = 11$ unknown part-worth values in our illustrative study.

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

$$\boldsymbol{\beta}_0 = [-0.5, -0.5, -0.5, -0.5, 0, -0.5, 0, -0.5, -0.25, 0, 0.25]' \qquad (18)$$

and prior variance-covariance matrix

$$\mathbf{\Sigma}_0 = \begin{bmatrix} 0.16 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.16 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.16 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.16 & -0.08 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -0.08 & 0.16 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.16 & -0.08 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -0.08 & 0.16 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.16 & -0.04 & -0.04 & -0.04 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.04 & 0.16 & -0.04 & -0.04 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.04 & -0.04 & 0.16 & -0.04 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.04 & -0.04 & -0.04 & 0.16 \end{bmatrix}. \tag{19}$$

For the specification of the prior mean $\boldsymbol{\beta}_0$, we chose to equally space the part-worth values between -0.5 and 0.5 for each attribute. The first three part-worth values in $\boldsymbol{\beta}_0$ are the prior mean utilities associated with the first level of each of the three 2-level attributes. The next two sets of two part-worth values reflect the prior mean utilities of the first and second level of each of the two 3-level attributes. Finally, the last four part-worth values correspond to the prior mean utilities of the first four levels of the 5-level attribute. Due to the effects-type coding, a prior mean part-worth value of 0.5 automatically results for the last level of each attribute. Using this prior mean specification, we assume that all attributes are equally important and that they are ordinal, where the levels of each attribute are ordered from least preferred to most preferred.

For the prior variance-covariance matrix $\mathbf{\Sigma}_0$, we specified $k = 11$ variances that are all equal to 0.16 and negative covariances between the $d_i - 1$ part-worths of each attribute i. We computed these covariances using a correlation coefficient of $-1/(d_i - 1)$. As explained by Kessels et al. (2008), this ensures that the variances of all part-worths corresponding to a given attribute are the same, meaning that the variance associated with the last implied level d_i of attribute i also equals 0.16.

5.2 Optimal designs

5.2.1 Full profile design

Table 6 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 such level overlap. Here, we see about 8% of level overlap.

5.2.2 One constant attribute

Table 7 shows the $\mathcal{D}_B$-optimal partial profile designs involving one constant attribute. It includes the attribute balance design and the variance balance design. We report only one variance balance design because the variance balance I and II weighting schemes produced the same master design. The attribute balance design holds each attribute constant in either two or three choice sets. It does not satisfy attribute balance perfectly because the number of choice sets, $S = 15$, is not a multiple of the number of attributes, $t = 6$. Hence, the $\mathcal{D}$-optimal master design used to select the varying attributes is inevitably not a BIBD. The variance balance design holds each of the 2-level attributes constant in five 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. In that case, the variance balance design does not prevent the selection of profiles governed by such attribute. The attribute balance design would be a better design option then.

5.2.3 Two constant attributes

Table 8 shows the $\mathcal{D}_B$-optimal partial profile designs involving two constant attributes. The first partial profile design is the attribute balance design. The remaining designs are the variance balance I and II designs. The three designs are based on the master designs discussed in Section 4.1.4. 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 not preferable.

5.2.4 Utility-neutral designs

Table 9 shows the $\mathcal{D}$-optimal utility-neutral full profile design and Tables 10 and 11 show the $\mathcal{D}$-optimal utility-neutral partial profile designs involving one and two constant attributes. These designs vary all possible attribute levels in the choice sets. They are level balanced in the varying attributes within and over all choice sets.

<Insert Tables 6-11 about here>

5.3 $\mathcal{D}_B$-efficiencies

To evaluate the estimation performance of the Bayesian and utility-neutral optimal designs, we compared them to the $\mathcal{D}_B$-optimal full profile design in terms of $\mathcal{D}_B$-efficiency. By doing so, we learn how much we lose in estimation precision by using a partial profile design and/or a utility-neutral design. Table 12 shows the $\mathcal{D}_B$-optimality criterion values and $\mathcal{D}_B$-efficiencies of the designs.

<Insert Table 12 about here>

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 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 favoring the latter design.

A comparison of the $\mathcal{D}_B$-optimal attribute balance designs involving one and two constant attributes reveals that the efficiency losses roughly double when using two constant attributes instead of one. Examining the $\mathcal{D}_B$-optimal variance balance designs on this matter, we observe that the efficiency losses increase by a factor of five when increasing the number of constant attributes from one to two. The efficiency losses thus increase relatively more rapidly for the variance balance designs. This is due to the fact that in the case of one constant attribute, the variance balance design only holds 2-level attributes constant, while in the case of two constant attributes, the two variance balance designs also hold 3-level attributes constant (see Tables 7 and 8). The use of 3-level attributes as constant attributes has a detrimental effect on the efficiency, illustrating that efficiency losses are larger when attributes with a larger number of levels are kept constant.

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), 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.

To conclude, for this illustrative design problem, the use of variance balance II weights for constructing Bayesian partial profile designs results in the highest statistical efficiency. The difference in efficiency from using variance balance I weights is small or nonexistent depending on the number of varying attributes.

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 discusses a design study in which the most statistically efficient design using the variance balance II weights can be used without any problem.

6 Follow-up designs for the software development experiment

In this section, we present three follow-up partial profile designs for the JMP graphics choice experiment discussed in Section 2. Table 13 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. However, the report color attribute is marginally significant ($p = 0.0924$). From the part-worth estimates of the simplified model in Table 14, 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. While this result is not statistically significant, it 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 14 and their variance-covariance matrix in Table 15 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.

<Insert Tables 13-15 about here>

For the sake of illustration, we show three partial profile designs generated using our algorithms each 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 16. 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 perfectly satisfies first-order balance because each attribute is constant in six choice sets. It, however, does not satisfy second-order balance perfectly, because the number of choice sets, $S = 15$, is not a multiple of the number of possible attribute pairings, which is 10. The $\mathcal{D}$-optimal master design is, therefore, inevitably not a BIBD.

The remaining partial profile designs in Table 16 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 holds the 4-level attribute constant in two choice sets, while the variance balance II design does not hold the 4-level attribute constant in any choice set.

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 17. The full profile design does not hold any of the attributes constant. Table 18 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.

<Insert Tables 16-18 about here>

Though 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 future research

We presented and compared three flexible algorithms for constructing $\mathcal{D}_B$-optimal main-effects partial profile designs for DCEs. Green (1974) originated the idea of partial profile designs to reduce the complexity of the choices that a respondent must make by varying the levels of only a subset of the attributes within each choice set. Using verbal descriptions of the attributes, the attributes whose levels are assumed constant, were dropped from the choice sets. Dellaert et al. (2012), however, 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. Also by construction, many DCEs use images or real-life prototypes as profiles requiring the levels of the constant attributes. For these reasons, we use the term "partial" profiles in this paper to simply indicate that the levels of some attributes remain constant in every choice set.

The motivation for partial profile designs is twofold. First, by reducing the complexity of the choices, partial profiles have the potential to prevent respondents from resorting to non-compensatory approaches towards making their choices, which would violate the implicit assumption of compensatory decision making in discrete choice models. Second, whether choices are complex or not, partial profiles can prove useful in the presence of a dominant attribute. For choice sets in which the level of the 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.

The three partial profile design algorithms we propose all have the same basic two-stage structure. In the first stage, we generate a master design that determines the attributes whose levels vary in each choice set, while in the second stage, we generate an attribute level design that sets the levels of the varying attributes. The three algorithms differ in the methods for constructing the master design. They provide alternative generalizations of Green's (1974) approach that makes use of balanced incomplete block designs (BIBDs) to select the varying attributes. We call our approaches attribute balance, variance balance I and variance balance II. The attribute balance approach pays an equal amount of attention to each attribute. It generates a BIBD in the event a BIBD exists for the design situation under study. Otherwise, it produces a master design that reflects the balancing structure of a BIBD as much as possible. The variance balance approaches, on the other hand, pay an equal amount of attention to each attribute level, which makes sense in experimental scenarios where not all attributes have the same number of levels. 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 approaches differ in the way attributes with different numbers of levels are weighted.

The partial profile design algorithms are flexible in that they can accommodate arbitrarily many attributes, each with any number of levels. 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 7. 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.

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. The results of the two comparison studies discussed in this paper revealed 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. However, this theoretical drawback vanishes if using full profile testing causes respondents to short cut the compensatory choice process. We should, however, be cautious when using the variance balance approaches because it is possible that they do not hold every attribute constant in at least one choice set. 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. Because the attribute balance approach holds every attribute constant the same number of times, this approach can be used in any design situation. We verified that the results obtained from the two comparison studies hold for other design problems as well.

A number of extensions of this work are worth investigating. First, the partial profile design algorithms consider main-effects models only. Examining the use of partial profiles when possible two-attribute interactions are present would be an interesting research topic. Second, the Bayesian design algorithms only allow for a multivariate normal prior distribution of the parameters. Extending the method to support other distributions would be a useful contribution. Finally, the algorithms only apply to the multinomial logit model assuming homogeneous preferences of the respondents. Exploring the use of more sophisticated nonlinear models that take into account respondent heterogeneity would be a challenging topic for future research. For example, the work by Bliemer and Rose (2010) and Yu et al. (2011) on full profile design construction methods for the panel mixed logit model could be extended to partial profiles.

Appendix. Motivation of variance balance methods

We show that the weight formulas in (16) and (17) 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, $w_i, i = 1, ..., t$, is proportional to the number of part-worth values associated with attribute i, as defined in (16). In a study using approximate design theory, Grossmann et al. (2006, 2009) 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 numbers of different choice sets exactly match their optimal fractions. Given that the design matrix of an experiment with two profiles per choice set contains effects-type coded difference vectors, that is, $\mathbf{X} = [\mathbf{x}'_{1s} - \mathbf{x}'_{2s}]_{s=1,\ldots,S}$, the information matrix of the $\mathcal{D}$-optimal approximate design is block diagonal and equal to

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

where

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

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 (16).

Variance balance II

Using variance balance II, each weight, $w_i, i = 1, ..., t$, 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. Assuming $\mathbf{X} = [\mathbf{x}'_{1s} - \mathbf{x}'_{2s}]_{s=1,\ldots,S}$, the information matrix is block diagonal and given by

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

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

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

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

References

Atkinson, A. C., Donev, A. N. and Tobias, R. D. (2007). *Optimum Experimental Designs, with SAS*, Oxford, U.K.: Clarendon Press.

Atkinson, A. C. and Haines, L. M. (1996). Designs for nonlinear and generalized linear models, in *Handbook of Statistics: Design and Analysis of Experiments*, Vol. 13, Ghosh, S. and Rao, C. R., eds. Amsterdam: Elsevier, 437–475.

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.

Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: a review, *Statistical Science* **10**: 273–304.

Chrzan, K. (2010). Using partial profile choice experiments to handle large numbers of attributes, *International Journal of Market Research* **52**: 827–840.

Dellaert, B. G. C., Donkers, B. and van Soest, A. (2012). Complexity effects in choice experiment-based models, *Journal of Marketing Research*, in press.

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.

Ferrini, S. and Scarpa, R. (2007). Designs with a priori information for nonmarket valuation with choice experiments: a Monte Carlo study, *Journal of Environmental Economics and Management* **53**: 342–363.

Gotwalt, C. M. (2010). Addendum to "Fast computation of designs robust to parameter uncertainty for nonlinear settings", *Technometrics* **52**: 137.

Gotwalt, C. M., Jones, B. A. and Steinberg, D. M. (2009). Fast computation of designs robust to parameter uncertainty for nonlinear settings, *Technometrics* **51**: 88–105.

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., 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.

Huber, J. and Zwerina, K. (1996). The importance of utility balance in efficient choice designs, *Journal of Marketing Research* **33**: 307–317.

Kanninen, B. J. (2002). Optimal design for multinomial choice experiments, *Journal of Marketing Research* **39**: 214–217.

Kessels, R., Goos, P. and Vandebroek, M. (2006). A comparison of criteria to design efficient choice experiments, *Journal of Marketing Research* **43**: 409–419.

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., Goos, P. and Vandebroek, M. (2008). Recommendations on the use of Bayesian optimal designs for choice experiments, *Quality and Reliability Engineering International* **24**: 737–744.

Kessels, R., Jones, B., Goos, P. and Vandebroek, M. (2009). An efficient algorithm for constructing Bayesian optimal choice designs, *Journal of Business and Economic Statistics* **27**: 279–291.

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.

Meyer, R. K. and Nachtsheim, C. J. (1995). The coordinate-exchange algorithm for constructing exact optimal experimental designs, *Technometrics* **37**: 60–69.

Rose (2011). Discussion on "The usefulness of Bayesian optimal designs for discrete choice experiments", *Applied Stochastic Models in Business and Industry* **27**: 193–196.

Rose, J. M. and Bliemer, M. C. J. (2009). Constructing efficient stated choice experimental designs, *Transport Reviews* **29**: 587–617.

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. (2002). Profile construction in experimental choice designs for mixed logit models, *Marketing Science* **21**: 455–475.

Sándor, Z. and Wedel, M. (2005). Heterogeneous conjoint choice designs, *Journal of Marketing Research* **42**: 210–218.

Scarpa, R. and Rose, J. M. (2008). Design efficiency for non-market valuation with choice modelling: how to measure it, what to report and why, *The Australian Journal of Agricultural and Resource Economics* **52**: 253–282.

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.

Street, D. J. and Burgess, L. (2007). *The Construction of Optimal Stated Choice Experiments: Theory and Methods*, Hoboken, New Jersey: Wiley.

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.

Yu, J., Goos, P. and Vandebroek, M. (2009). Efficient conjoint choice designs in the presence of respondent heterogeneity, *Marketing Science* **28**: 122–135.

Yu, J., Goos, P. and Vandebroek, M. (2010). Comparing different sampling schemes for approximating the integrals involved in the efficient design of stated choice experiments, *Transportation Research B* **44**: 1268–1289.

Yu, J., Goos, P. and Vandebroek, M. (2011). Individually adapted sequential Bayesian conjoint-choice designs in the presence of consumer heterogeneity, *International Journal of Research in Marketing* **28**: 378–388.

Tables

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

Attributes	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.

Attributes	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: BIBD master design for the seven choice sets in Table 3. The three varying attributes in each choice set are indicated by check marks and form a BIBD. The four constant attributes are shown in gray.

Choice set	Report color	Picture color	Graph color	Frame line color	Frame all sides	Graph rectangle	Y-axis title
1			✓		✓		✓
2			✓	✓		✓	
3		✓			✓	✓	
4		✓		✓			✓
5	✓					✓	✓
6	✓			✓	✓		
7	✓	✓	✓				

Table 5: 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	Attributes with levels						Attributes with levels						Attributes with levels					
set	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	✓	✓	✓	✓			✓	✓	✓			✓	✓	✓		✓		✓
# cst	*5*	*5*	*5*	*5*	*5*	*5*	*7*	*7*	*8*	*4*	*4*	*0*	*8*	*8*	*8*	*3*	*3*	*0*

Table 6: $\mathcal{D}_B$-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	Attributes					
set	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

# cst	*2*	*2*	*3*	*0*	*0*	*0*

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 five varying attributes and one constant attribute. Variance balance I and II lead to the same design in this scenario.

	Attribute balance						Variance balance					
Choice	Attributes						Attributes					
set	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
# cst	*2*	*3*	*3*	*3*	*2*	*2*	*5*	*5*	*5*	*0*	*0*	*0*

Table 8: $\mathcal{D}_B$-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.

	Attribute balance						Variance balance I						Variance balance II					
Choice	Attributes						Attributes						Attributes					
set	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
# cst	*5*	*5*	*5*	*5*	*5*	*5*	*7*	*7*	*8*	*4*	*4*	*0*	*8*	*8*	*8*	*3*	*3*	*0*

Table 9: $\mathcal{D}$-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

# cst	0	0	0	0	0	0

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 five varying attributes and one constant attribute. Variance balance I and II lead to the same design in this scenario.

	Attribute balance						Variance balance					
Choice	Attributes						Attributes					
set	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
# cst	*2*	*3*	*3*	*3*	*2*	*2*	*5*	*5*	*5*	*0*	*0*	*0*

Table 11: $\mathcal{D}$-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.

	Attribute balance						Variance balance I						Variance balance II					
Choice	Attributes						Attributes						Attributes					
set	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
# cst	*5*	*5*	*5*	*5*	*5*	*5*	*7*	*7*	*8*	*4*	*4*	*0*	*8*	*8*	*8*	*3*	*3*	*0*

Table 12: $\mathcal{D}_B$-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 13: Likelihood ratio tests of the attribute effects.

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 14: 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 15: Variance-covariance matrix of the part-worth estimates in Table 14 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 16: $\mathcal{D}_B$-optimal follow-up partial profile designs for the JMP graphics choice experiment.

	Attribute balance					Variance balance I					Variance balance II				
Choice set	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	*	white	*	yes	vertical	light gray	*	*	yes	horizontal	white	*	no	*	vertical
3	*	light gray	*	no	horizontal	bluish	*	*	no	vertical	creamish	*	yes	*	horizontal
4	*	white	yes	*	horizontal	light gray	*	*	yes	vertical	light gray	*	yes	*	horizontal
4	*	light gray	no	*	vertical	white	*	*	no	horizontal	bluish	*	no	*	vertical
5	*	white	yes	*	horizontal	white	*	no	*	vertical	light gray	*	no	*	horizontal
5	*	light gray	no	*	vertical	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	bluish	*	*	no	vertical	creamish	*	no	yes	*	creamish	*	no	no	*
7	white	*	*	yes	horizontal	light gray	*	yes	no	*	bluish	*	yes	yes	*
8	white	*	*	no	vertical	white	*	no	yes	*	bluish	light gray	*	*	vertical
8	creamish	*	*	yes	horizontal	creamish	*	yes	no	*	white	white	*	*	horizontal
9	bluish	*	yes	*	horizontal	white	white	*	*	horizontal	light gray	light gray	*	*	vertical
9	light gray	*	no	*	vertical	bluish	light gray	*	*	vertical	creamish	white	*	*	horizontal
10	light gray	*	yes	no	*	bluish	white	*	*	horizontal	light gray	white	*	*	vertical
10	white	*	no	yes	*	light gray	light gray	*	*	vertical	bluish	light gray	*	*	horizontal
11	light gray	*	no	yes	*	white	white	*	*	vertical	white	light gray	*	yes	*
11	white	*	yes	no	*	light gray	light gray	*	*	horizontal	light gray	white	*	no	*
12	creamish	light gray	*	*	vertical	light gray	white	*	no	*	white	white	*	no	*
12	light gray	white	*	*	horizontal	creamish	light gray	*	yes	*	light gray	light gray	*	yes	*
13	creamish	white	*	no	*	creamish	white	*	no	*	white	light gray	*	no	*
13	bluish	light gray	*	yes	*	bluish	light gray	*	yes	*	bluish	white	*	yes	*
14	bluish	white	no	*	*	bluish	white	no	*	*	light gray	light gray	yes	*	*
14	creamish	light gray	yes	*	*	white	light gray	yes	*	*	bluish	white	no	*	*
15	light gray	light gray	yes	*	*	light gray	white	no	*	*	light gray	white	no	*	*
15	creamish	white	no	*	*	white	light gray	yes	*	*	creamish	light gray	yes	*	*
# cst	*6*	*6*	*6*	*6*	*6*	*2*	*7*	*7*	*7*	*7*	*0*	*7*	*8*	*8*	*7*

Table 17: $\mathcal{D}_B$-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
# cst	*0*	*0*	*0*	*0*	*0*

Table 18: $\mathcal{D}_B$-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	10.65028	63.57%
2	variance I	11.29531	69.70%
2	variance II	11.46489	71.41%

Figures

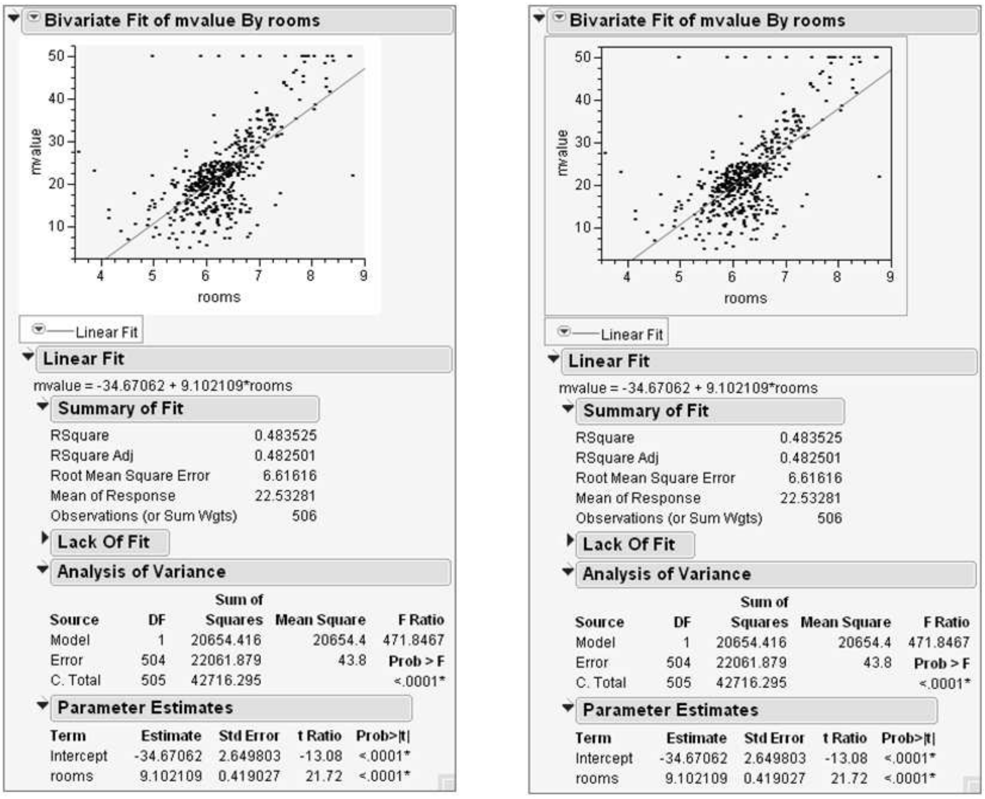

Figure 1: Example choice set used in the JMP graphics choice experiment.