



Transportation Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Sampling of Alternatives in Random Regret Minimization Models

C. Angelo Guevara, Caspar G. Chorus, Moshe E. Ben-Akiva

To cite this article:

C. Angelo Guevara, Caspar G. Chorus, Moshe E. Ben-Akiva (2014) Sampling of Alternatives in Random Regret Minimization Models. Transportation Science

Published online in Articles in Advance 23 Dec 2014

. <http://dx.doi.org/10.1287/trsc.2014.0573>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2014, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Sampling of Alternatives in Random Regret Minimization Models

C. Angelo Guevara

Faculty of Engineering and Applied Sciences, Universidad de los Andes, Las Condes, Santiago, Chile 762001,
caguevara@miuandes.cl

Caspar G. Chorus

Faculty of Technology, Policy and Management, Delft University of Technology, 2628 BX Delft, The Netherlands,
c.g.chorus@tudelft.nl

Moshe E. Ben-Akiva

Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139,
mba@mit.edu

Sampling of alternatives is often required in discrete choice models to reduce the computational burden and to avoid describing a large number of attributes. This approach has been used in many areas, including modeling of route choice, vehicle ownership, trip destination, residential location, and activity scheduling. The need for sampling of alternatives is accentuated for random regret minimization (RRM) models because, unlike random utility models, the regret function for each alternative depends on all of the alternatives in the choice-set. In this paper we develop and test a method to achieve consistency, asymptotic normality, and relative efficiency of the estimators while sampling alternatives in a class of models that includes RRM. The proposed method can be seen as an extension of the approach used to address sampling of alternatives in multivariate extreme value models. We illustrate the methodology using Monte Carlo experimentation and a case study with real data. Experiments show that the proposed method is practical, performs better than a truncated model, and results in finite-sample estimates that provide a good approximation of those obtained with a model considering all of the alternatives.

Keywords: sampling of alternatives; random regret minimization; large choice-sets

History: Received: August 2012; revisions received: February 2014, April 2014; accepted: August 2014. Published online in *Articles in Advance*

1. Introduction

Various types of discrete choice models that are relevant in transport modeling involve huge choice-sets. This is the case, for example, for models of route choice, trip destination, residential location, or activity scheduling. Two types of difficulties may arise when the choice-set is too large. The first is the computational burden of managing a large number of alternatives and the second is the need to gather the data to describe them. Both difficulties may arise in estimation and forecasting. In this article we consider the former, proposing a solution method for random regret minimization (RRM) models.

In the context of the classical random utility maximization-based (RUM) logit model (McFadden 1974), a convenient method has been proposed (McFadden 1978) to obtain a consistent estimator for model parameters with a sample of alternatives. This estimator capitalizes on the fact that, because of its independently and identically distributed (i.i.d.) errors, the RUM-based logit model exhibits the independence of irrelevant alternatives (IIA) property. McFadden's (1978) result concerning the sampling of alternatives

for logit has been profusely used over the years. Examples abound in many areas such as route choice (see, e.g., Fosgerau, Frejinger, and Karlstrom 2013, Frejinger, Bierlaire, and Ben-Akiva 2009); vehicle ownership (Berkovec and Rust 1985); trip destination (Carrasco 2008); residential location (Lee and Waddell 2010); and activity based modeling (see, e.g., Daly, Hess, and Dekker 2014, Bradley, Bowman, and Griesenbeck 2010, Bowman and Ben-Akiva 2001).

Although very convenient from a modeler's perspective, this IIA property is often considered to be restrictive in terms of the implied behavior of decision makers. Since the 1970s, this observation has led to the development of a number of alternative discrete choice model forms whose errors are not i.i.d. Although still featuring closed form choice probabilities, these models do not exhibit the IIA property because they allow for correlation among the errors associated with different (subsets of) alternatives. A prominent example of this category is the nested logit model (Ben-Akiva 1973), which was shown a few years after its inception to belong, together with the logit, to the more

general family of closed form choice models based on a multivariate extreme value (MEV) distribution (McFadden 1978). More recently, MEV mixture models have been proposed that allow for even more flexibility in terms of the specification of the error term distribution and related behavioral implications and substitution patterns (e.g., McFadden and Train 2000).

The problem of sampling of alternatives in nonlogit models has been only recently studied. Guevara and Ben-Akiva (2013b) and Guevara (2010) proposed a method to achieve consistent estimation while sampling alternatives in MEV models, providing examples for the nested logit and the cross nested logit. The method consists in developing a proper correction of a term that gets truncated because of the sampling. Also, Guevara and Ben-Akiva (2013a) proposed a method for estimation while sampling alternatives in logit mixture models, showing that a naïve approach, in which the kernel of the mixture is replaced by McFadden's (1978) correction for logit, does achieve consistent estimation. With this, Guevara and Ben-Akiva (2013a) provide theoretical support for previous empirical results suggesting the suitability of the naïve approach for logit mixture models (McConnel and Tseng 2000, Nerella and Bhat 2004, Azaiez 2010, Lemp and Kockelman 2012).

Recently, a choice model has been proposed that does not exhibit the IIA-property even though (when written in logit form) its errors are i.i.d. This RRM model (Chorus 2010), which is the focus of this paper, is based on a regret minimization-based decision rule. The model postulates that when decision makers choose among alternatives, they try to avoid the situation where a nonchosen alternative is better than a chosen one in terms of one or more attributes. This translates into a regret function for a considered alternative that by definition features all attributes of all competing alternatives. Different from a model in which the attributes of all other alternatives are just included linearly in the utility (Ben-Akiva 1974), the cross elasticities of the RRM do not need to be the same for all alternatives. However, in some extreme cases, these cross elasticities may be ill behaved (Hensher, Greene, and Chorus 2013), a limitation of the RRM model that still needs to be addressed.

Since its introduction in 2010, the RRM model has been estimated and applied by various authors in the context of a variety of different choice contexts, involving—to name a few examples—travelers' choices among vehicle types, destinations, modes, routes, departure times, and driving maneuvers; politicians' choices among policy options; patients' choices among medical treatments; and tourists' choices among leisure activity locations. An overview of recent studies empirically comparing RRM with RUM models can be found in Chorus, van Cranenburgh, and Dekker (2014).

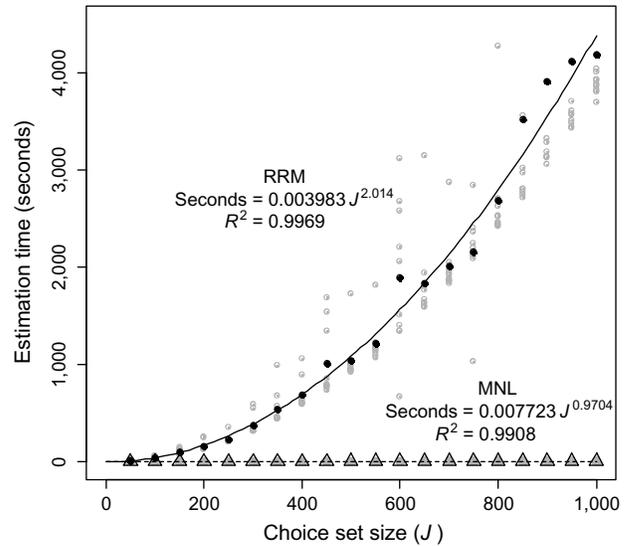


Figure 1 Estimation Time of RRM and MNL Models as a Function of the Number of Alternatives (J)

One disadvantage of the RRM model, which was highlighted in Chorus (2012), is that runtimes may suffer from combinatorial explosion when choice-sets become very large. This issue is a direct result from the behavioral postulate, incorporated in the regret function, that every alternative is compared with every other alternative in the choice-set in terms of every attribute.

The combinatorial explosion of RRM, compared to multinomial logit (MNL), is illustrated in Figure 1, which depicts estimation time (ordinates axis) as a function of the number of alternatives (J) in the choice-set (abscissas axis). The results were obtained from 10 Monte Carlo simulations for each value of J between 50 and 1,000, in steps of 50. The estimation time of each simulation is depicted in gray with a small symbol (a dot for RRM and a triangle for MNL) and the average within the 10 repetitions for each J is depicted with a larger dark symbol. Both the RRM and the MNL models consider only one attribute and 1,000 observations. Figure 1 shows that the average estimation time for RRM as a function of J is fitted almost perfectly by a quadratic function, reflecting the computational problems that arise with RRM models with large choice-sets. In turn estimation time for MNL is almost flat with J .

As a consequence, finding a proper way to estimate RRM models on sampled choice-sets is an important condition for the model to be useful in the context of choice situations involving very large numbers of alternatives. At this point it should be noted that since the RRM model does not exhibit the IIA-property, McFadden's (1978) result does not apply. As mentioned, this is the case even when—such as is the case for RRM-based logit models—errors are distributed i.i.d.

This paper extends the work of Guevara and Ben-Akiva (2013b) by presenting an estimator for the RRM-based logit model in the context of sampled choice-sets (§2). Furthermore, it analyzes the conditions required for consistency, asymptotic normality, and efficiency and determines the correct expansion factors required in some relevant examples (§3). Then it illustrates the methods and studies the finite sample properties of the estimators using Monte Carlo experimentation (§4) and real data (§5). The article finishes by summarizing the main results, their possible implications, and suggestions for future lines of research (§6).

2. Estimation and Sampling of Alternatives in Random Regret Minimization Models

We consider the RRM model proposed by Chorus (2010). The behavioral assumption behind the RRM model is that individual n chooses alternative i , within the choice-set C_n , if i minimizes the anticipated regret he or she may get from that decision. The regret is defined as a measure of how much worse is the chosen alternative i , regarding each attribute m , compared to all other alternatives $j \neq i$.

For example, if m refers to a price attribute, then $\beta_m < 0$. Therefore, if an agent n chooses alternative i , he or she will perceive a price regret $\beta_m \{x_{jmn} - x_{imn}\}$ if $x_{jmn} < x_{imn}$, and zero otherwise, for each alternative j other than i . Formally, if m is the only attribute, the regret function R_{ijmn} can be summarized by the expression shown in Equation (1)

$$R_{ijmn} = \max[0, \beta_m \{x_{jmn} - x_{imn}\}]. \quad (1)$$

The regret function described in Equation (1) is difficult to implement in practice for estimation because it is not differentiable. For that, Chorus (2010) proposes to approximate R_{ijmn} by the expression shown in Equation (2)

$$R_{ijmn} = \max[0, \beta_m \{x_{jmn} - x_{imn}\}] \approx \ln(1 + \exp(\beta_m \{x_{jmn} - x_{imn}\})), \quad (2)$$

which can be seen either as a plausible approximation (see Figure 1 in Chorus 2010) or as the result of assuming unobserved heterogeneity in the regret function.

The regret function R_{in} of alternative i for agent n is completed by summing R_{ijmn} over all of the attributes m and alternatives $j \neq i$ in the choice-set C_n , as shown in Equation (3)

$$R_{in} = \sum_{\substack{j \in C_n \\ j \neq i}} \sum_{m=1}^M \ln[1 + \exp(\beta_m \{x_{jmn} - x_{imn}\})]. \quad (3)$$

Finally, it is considered that the individual seeks to minimize a random regret function $RR_{in} = R_{in} - \varepsilon_{in}$, where ε_{in} is a random term assumed to be i.i.d. extreme value $(0, \mu)$. Under those conditions, the probability that agent n will choose alternative i will correspond to the model shown in Equation (4)

$$P_n(i) = \frac{e^{-\mu R_{in}}}{\sum_{j \in C_n} e^{-\mu R_{jn}}}, \quad (4)$$

where the scale parameter μ is identifiable, but has been usually normalized equal to one. We will consider such normalization for the rest of the article.

Consider now that the researcher samples a subset D_n with \tilde{J}_n elements from the true choice-set C_n that is considered by the decision maker. As stated before, the sampling may be needed to reduce the computational burden and/or facilitate data collection. For estimation purposes, D_n must include the chosen alternative i , and then D_n is not independent of i . If i is not included in D_n , a probability measure constructed combining i and the elements in D_n may not be well defined since it may be larger than one. Also, if i is not included in D_n , the likelihood may be unbounded, precluding model estimation.

Define $\pi(i, D_n)$ the joint probability that agent n would choose alternative i and that the researcher would draw the set D_n . Using Bayes theorem, this joint probability can be rewritten as shown in Equation (5)

$$\pi(i, D_n) = \pi(D_n | i)P_n(i) = \pi(i | D_n)\pi(D_n), \quad (5)$$

where $\pi(i | D_n)$ is the conditional probability of choosing alternative i , given that the set D_n was drawn, and $\pi(D_n | i)$ is the conditional probability that the researcher drew the set D_n , given that alternative i was chosen by the agent.

Because the events of choosing each of the alternatives in C_n are mutually exclusive and totally exhaustive, we can write the probability $\pi(D_n)$ of constructing the set D_n as shown in Equation (6)

$$\pi(D_n) = \sum_{j \in C_n} \pi(D_n | j)P_n(j) = \sum_{j \in D_n} \pi(D_n | j)P_n(j), \quad (6)$$

where the second equality holds because $\pi(D_n | j) = 0 \forall j \notin D_n$.

Substituting Equation (6) and the choice probability $P_n(i)$ shown in Equation (4) into Equation (5), Equation (7) is obtained by canceling and rearranging terms

$$\pi(i | D_n) = \frac{e^{-R_{in} + \ln \pi(D_n | i)}}{\sum_{j \in D_n} e^{-R_{jn} + \ln \pi(D_n | j)}}. \quad (7)$$

The direct application of McFadden's (1978) result on sampling of alternatives for logit can be used to show that maximizing a conditional log-likelihood based

on the expression shown in Equation (7) would yield consistent estimators of the model parameters.

Equation (7) shows two things about the conditional probability $\pi(i | D_n)$. First, the form of the probability is very similar to Equation (4), except for the term $\ln \pi(D_n | j)$, which is known as the sampling correction. Second, the summation in the denominator is only over the alternatives in D_n .

However, Equation (7) does not yet offer a practical solution for the sampling of alternatives in random regret models. The problem is that even though the denominator of the choice probability depends only on D_n , the argument R_{in} still depends on the full choice-set C_n .

In this paper, we adapt Equation (7) to the problem of sampling of alternatives in RRM models by replacing R_i with an estimator that depends only on the subset D_n . We analyze the conditions required for consistency, asymptotic normality, and efficiency; determine the correct expansion factors required in some relevant examples; and illustrate the finite sample properties of the estimators using Monte Carlo experimentation and real data.

The results on consistency, asymptotic normality, and efficiency are summarized by the following theorem, which is a generalization of the result of Guevara and Ben-Akiva (2013b) to a class of models that includes the RRM model: models that can be written in a logit form, with nonlinear utility functions that depend on the full choice-set.

THEOREM 1. Consider N observations, a choice-set C_n of cardinality J_n , and two subsets $D_n \subseteq C_n$ and $\tilde{D}_n \subseteq C_n$ both with cardinality.¹ If

(a) the choice model is of the logit form, in the sense that it can be written as

$$P_n(i) = \frac{e^{W_{in}(C_n)}}{\sum_{j \in C_n} e^{W_{jn}(C_n)'}}$$

where $W_{in}(C_n)$ is any continuous and twice differentiable function of the attributes x_{jn} of all of the alternatives in C_n , and a set of parameters β^* .²

(b) $\hat{W}_{in}(\tilde{D}_n)$ is an unbiased estimator of W_{in} ;

(c) The variance of $\hat{W}_i(\tilde{D}_n)$ is bounded and decreases with \tilde{J}_n . Because $\hat{W}_i(\tilde{D}_n)$ is also unbiased, this means that $\hat{W}_i(\tilde{D}_n)$ is also consistent; and

(d) $\pi(D_n | j) > 0 \forall j \in D_n$ and $\pi(D_n | j) = 0 \forall j \notin D_n$, which holds when the chosen alternative is included in D_n .

Then the maximization of the conditional quasilog-likelihood (CQLL) function

$$\sum_{n=1}^N \ln \hat{\pi}(i | D_n, \tilde{D}_n) = \sum_{n=1}^N \ln \frac{e^{\hat{W}_{in}(\tilde{D}_n) + \ln \pi(D_n | i)}}{\sum_{j \in D_n} e^{\hat{W}_{jn}(\tilde{D}_n) + \ln \pi(D_n | j)}} \quad (8)$$

¹ This final assumption is just used to simplify the notation. It is not essential and can be generalized.

² Note that $W_{in}(C_n)$ includes, but is not limited to, the regret function defined in Equation (3).

yields, under general regularity conditions, consistent estimators of the model parameters β^* , as \tilde{J}_n increases with N at any rate. If \tilde{J}_n increases faster than \sqrt{N} does, the estimators of the model parameters will be consistent and asymptotically normal

$$\hat{\beta} \sim^a \text{Normal}(\beta^*, \mathbf{R}^{-1} \mathbf{\Omega} \mathbf{R}^{-1} / N) \quad (9)$$

where

$$\mathbf{\Omega} = \text{Var} \frac{\partial \ln \pi_n(\beta^* | D)}{\partial \beta} \quad \text{and} \quad \mathbf{R} = E \frac{\partial^2 \ln \pi_n(\beta^* | D)}{\partial \beta \partial \beta'}$$

where

$$\pi(i | D_n) = \frac{e^{W_{in}(C_n) + \ln \pi(D_n | i)}}{\sum_{j \in D_n} e^{W_{jn}(C_n) + \ln \pi(D_n | j)}}$$

This variance-covariance matrix can be approximated by the Berndt-Hall-Hausman (BHHH) estimator (Berndt, Hall, and Hall 1974), using $\hat{\pi}(i | D_n, \tilde{D}_n)$, evaluated at the optimal values.

Note that the variance-covariance matrix attained with Equation (8) is the same one attained by maximization of the impractical conditional log-likelihood function

$$\sum_{n=1}^N \ln \pi(i | D_n) = \sum_{n=1}^N \ln \frac{e^{W_{in}(C_n) + \ln \pi(D_n | i)}}{\sum_{j \in D_n} e^{W_{jn}(C_n) + \ln \pi(D_n | j)}} \quad (10)$$

This implies that the feasible estimator proposed in Equation (8) is relatively efficient, in the sense that it yields estimators that are as asymptotically efficient as the estimators obtained when considering the full choice-set to calculate $W_{in}(C_n)$. They are not globally efficient because some efficiency is lost when sampling alternatives in Equation (10).

Finally, if J_n is finite and the protocol is sampling without replacement, \tilde{J}_n needs to increase only up to $\tilde{J}_n = J_n$ to achieve the asymptotic distribution shown in Equation (9).

PROOF. The proof builds on the same approach used by Train (2009, pp. 247–257) to derive the statistical properties simulation-based estimators. In the appendix we provide a summarized demonstration, highlighting principal parts and including a justification for the main assumptions that are required.

The main difference between this theorem and the results of Guevara and Ben-Akiva (2013b) is in the term that gets truncated and the way this can be resolved in practice for RRM.

Two cases can be distinguished in the theorem: when the cardinality of the choice-set is finite and when it is infinite.

For finite J_n , if the protocol is sampling with replacement, \tilde{J}_n will have to grow infinitely with N to achieve consistency and grow faster than \sqrt{N} to achieve asymptotic normality. Instead, if the protocol is sampling without replacement, when \tilde{J}_n grows with N it will eventually reach J_n . At that point, the variance of $\hat{W}_{in}(\tilde{D}_n)$ will be zero because $W_{in} = \hat{W}_{in}$ and the CQLL

in Equation (8) would become the same as the conditional log-likelihood in Equation (10), achieving consistency and asymptotic normality.

When J_n is infinite, \tilde{J}_n could grow to infinity with N , even if the protocol is sampling without replacement. However, a relevant question in this case is whether or not an RRM model with an infinite number of alternatives will be well defined. Models with infinite choice-sets have been previously considered, among others, for spatial choice, labor demand, and route choice. Examples of those are the works of McFadden (1976), Ben-Akiva and Watanatada (1981), Dagsvik (1989) and, more recently, Fosgerau, Frejinger, and Karlstrom (2013). The validity of the RRM with an infinite number of alternatives can be stated for a fairly general case. Following the derivation of the continuous logit (McFadden 1976), RRM will be well defined for cases in which the alternatives are built as partitions of a space of elemental alternatives X , a space that can be finite or infinite. An example of a choice model of this type is the problem of residential location choice. Nevertheless, numerical limitations in the estimation and forecasting of a model with an infinite choice-set would make the model intractable, which is what motivates the need for sampling of alternatives.

Although the theoretical results hold asymptotically, the Monte Carlo experiments in §4 show that for finite N (1,000 in the example) and finite J_n (1,000), \tilde{J}_n as small as 30 can result in proper estimators. Moreover, depending on the behavior of $\text{Var}(\hat{W}_{in})$ for small \tilde{J}_n , the theorem sheds some light on the speed of convergence, which can be useful in practice. For example, if one would like to maintain the statistical properties attained when $\tilde{J}_n = 30$ and $N = 1,000$, but with $N = 2,000$, the theoretical result states that \tilde{J}_n would have to be at least 45 because

$$\tilde{J}_n = 45 \approx 30 \frac{\sqrt{2,000}}{\sqrt{1,000}}$$

3. Application of the Method in Practice

3.1. Introduction

For the application of the theorem to the RRM model in practice, it is convenient to note first what occurs if, in Equation (3), the incumbent alternative is included in the regret function

$$\begin{aligned} \tilde{R}_{in} &= \sum_{j \in C_n} \sum_{m=1}^M \ln[1 + \exp(\beta_m \{x_{jmn} - x_{imn}\})] \\ &= R_{in} + M \ln(2). \end{aligned} \quad (11)$$

Equation (11) suggests that the inclusion of the incumbent alternative implies the addition of the same constant $M \ln(2)$ to all of the alternatives, where M

is the number of attributes. Since this same constant cancels out, considering the incumbent alternative in the regret function has no impact in the choice probability shown in Equation (4). For the rest of the paper, we will consider the definition of the regret function including the incumbent alternative, as in Equation (11). This will facilitate the notation of the different versions of the practical application of the method. Also, going forward, we will remove the tilde from \tilde{R}_{in} to save notation.

We propose the following \hat{R}_{in} as a feasible approximation of R_{in} :

$$\hat{R}_{in} = \sum_{j \in \tilde{D}_n} w_{jn} \sum_{m=1}^M \ln[1 + \exp(\beta_m \{x_{jmn} - x_{imn}\})]. \quad (12)$$

The expansion factors w_{jn} in \hat{R}_{in} needed for attaining unbiasedness, as required by the theorem, would have to have the following form:

$$w_{jn} = \frac{\tilde{n}_{jn}}{E(\tilde{n}_{jn})}, \quad (13)$$

where \tilde{n}_{jn} corresponds to the number of times alternative j is included in the sample for agent n and $E(\tilde{n}_{jn})$ is its expected value (see Guevara and Ben-Akiva 2013b, Appendix B, for a proof of a similar case but in a different context of nests in RUM MEV models). Note that if the protocol used to draw alternatives is sampling without replacement, $\tilde{n}_{jn} = 1$ and $E(\tilde{n}_{jn})$ corresponds to the probability of sampling alternative j .

The expansion factors w_{jn} would depend on the sampling protocol used and, importantly, on whether or not the subset D_n used to write the sampling correction $\ln \pi(D_n | j)$ in Equation (7) is the same as the subset (\tilde{D}_n) used to build the expansion factors w_{jn} .

We will next explore three methods to construct in practice the expansion factors shown in Equation (13). These methods are analogous to some of the approaches explored by Guevara and Ben-Akiva (2013a, b) for the problem of sampling of alternatives in logit mixture and MEV, respectively.

3.2. Expansion Factors When Resampling Is Possible

Consider first the case when the researcher has full control of the data and is able to sample a set D_n from C_n to build the sampling correction $\ln \pi(D_n | i)$ and to sample a different set \tilde{D}_n from C_n to construct the expansion factors w_{jn} needed to build \hat{R}_{in} . To save notation we will consider that both D_n and \tilde{D}_n have the same cardinality \tilde{J} for all individuals, but this is not essential and can be generalized.

The expansion factors required depend on the protocol used for building \tilde{D}_n . We consider as an example that the protocol is a simple random sample without

replacement. Note that the chosen alternative does not need to necessarily be in \tilde{D}_n . The sampling in this case is random from all of the elements in C_n . This is crucial for the simplicity and practicality of applying this version of the method.

In such a case the expansion factors in \hat{R}_{in} that are needed to achieve an unbiased estimator of R_{in} are the following for each alternative j :

$$w_{jn} = \frac{\tilde{n}_{jn}}{E(\tilde{n}_{jn})} = \frac{1}{\tilde{J}/J} = \frac{J}{\tilde{J}}. \quad (14)$$

To describe the likelihood function required to estimate the model, we need to specify the sampling protocol used to build the set D_n in order to be able to determine McFadden's (1978) sampling correction. For example, consider that the protocol used in this case is the following. In the first step, the chosen alternative for each observation is included. Then nonchosen alternatives are randomly sampled without replacement to make a total of \tilde{J} . In this case the sampling correction will correspond to

$$\ln \pi_n(D | i) = \ln \binom{J-1}{\tilde{J}-1},$$

a term that, for this particular sampling protocol, is constant across alternatives and, therefore, cancels out in the calculation of the quasilog-likelihood function shown in Equation (8).

To summarize, given the particular sampling protocols for D_n and \tilde{D}_n described, the conditional probability of choosing alternative i , given that the sets D_n and \tilde{D}_n were drawn, can be approximated by

$$\hat{\pi}_n(i | D_n, \tilde{D}_n)^{\text{Resampling}} = \frac{e^{-\sum_{j \in \tilde{D}_n} (J/\tilde{J}) \sum_{m=1}^M \ln[1 + \exp(\beta_m(x_{jmn} - x_{imn}))]}}{\sum_{k \in D_n} e^{-\sum_{j \in \tilde{D}_n} (J/\tilde{J}) \sum_{m=1}^M \ln[1 + \exp(\beta_m(x_{jmn} - x_{kmn}))]}}. \quad (15)$$

Therefore, according to the theorem, a model estimated using the CQLL function built using Equation (15) will result in consistent and asymptotically normal estimators of the model parameters and the variance-covariance matrix of the estimators can be obtained using the BHHH estimator. This estimation tool is practical because it can be applied in canned estimation software such as BIOGEME (Bierlaire 2003) or ALOGIT (Daly 1992) with minor modifications, making it very attractive for practitioners.

Finally, note that the intuition behind Equation (15) is direct. If, for example, 10 out of 1,000 alternatives are sampled randomly to build \tilde{D}_n , the regret function has to be calculated with the 10 alternatives and then amplified by 100 to correct for regret being otherwise underestimated because of the smaller (sampled) choice-set. Note that since for this particular sampling

protocol the expansion factor is the same for all of the alternatives, the $w_{jn} = J/\tilde{J}$ term comes out of the sum and becomes indistinguishable with the overall utility scale.

Things become more troublesome when the researcher is forced to instead use the same set D_n to build the term \hat{R}_{in} . We will discuss this in §3.3.

3.3. Expansion Factors When Resampling Is Not Possible

Consider that the researcher does not have full control of the data and is not able to sample two sets D_n and \tilde{D}_n . This can occur, for example, when the researcher is using a database previously processed and for which he or she does not have access to the original source because of privacy concerns.

If the protocol used to build D_n (and therefore also \tilde{D}_n) was to draw first the chosen alternative and then to sample $\tilde{J} - 1$ alternatives randomly, the expansion factors required to attain unbiasedness are the following (see Guevara and Ben-Akiva 2013b, Appendix C, for a demonstration of an equivalent case):

$$w_{ij} = \frac{1}{P_n(j) + ((\tilde{J} - 1)/(J - 1))(1 - P_n(j))}. \quad (16)$$

There is a crucial difference between Equation (16) and Equation (14). The expression shown in Equation (16) depends on the choice probabilities, which are unknown beforehand in an application with real data. To avoid this limitation in practice, we postulate two methods called Pop.Shares and 1_0.

Method Pop.Shares. One way to approximate the choice probabilities needed for the calculation of the expansion factors is to use the population shares H_j of each alternative. Replacing choice probabilities by population shares in Equation (16), the expansion factors implied by this procedure become the following:

$$w_{jn} = \frac{1}{H_j + ((\tilde{J} - 1)/(J - 1))(1 - H_j)} \quad \forall n = 1, \dots, N; \quad \forall j \in C_n.$$

An advantage in this case is that the expansion factors w_{jn} can be directly calculated without incurring additional computational costs. Although the true population shares are not available in a real application, good approximations of them may be available from different sources (e.g., census or flow counts) or directly from the sample, provided it is random. In case the H_j has to be gathered from the sample, it could be calculated as

$$H_j \approx \frac{\sum_n y_{jn}}{N},$$

where y_{jn} equals one if individual n chooses alternative j and zero otherwise.

Given the particular sampling protocol described for D_n , the conditional probability of choosing alternative i , given that the set D_n was drawn, can be approximated by

$$\hat{\pi}_n(i | D_n)^{\text{Pop.Shares}} = \frac{e^{-Q(i)}}{\sum_{k \in mD_n} e^{-Q(k)}}, \quad (17)$$

with

$$Q(s) := \sum_{j \in D_n} \left[\frac{\sum_n y_{jn}}{N} + \frac{\tilde{J}-1}{\tilde{J}-1} \left(1 - \frac{\sum_n y_{jn}}{N} \right) \right]^{-1} \cdot \sum_{m=1}^M \ln[1 + \exp(\beta_m \{x_{jmn} - x_{smn}\})].$$

The Pop.Shares method could be easily implemented in canned estimation software with minor modifications, making it attractive for practitioners. The disadvantage is that the approximation $H_j \approx P_n(j)$ may be too rough, potentially causing large finite sample biases. This approach is studied using Monte Carlo experiments in §4.

Method 1_0. Another approach to avoid the need for the choice probabilities is to approximate them by considering that they take value 1 for the observed chosen alternative and 0 for the nonchosen ones. Replacing these assumptions in the example described in Equation (16) the expansion factors in this case would be the following:

$w_{jn} = 1$ if j is the chosen alternative and

$w_{jn} = (J_n - 1)/(\tilde{J}_n - 1)$ if j is not chosen.

It should be noted that the 1_0 approach to address the problem of estimation while sampling alternatives is the same as the one that was implicitly considered by Frejinger, Bierlaire, and Ben-Akiva (2009) and by Lee and Waddell (2010) in a different context.

Given the particular sampling protocol described for D_n , the conditional probability of choosing alternative i , given that the set D_n was drawn, can be approximated by

$$\hat{\pi}_n(i | D_n)^{1_0} = \frac{e^{-Q_1(i)}}{\sum_{k \in D_n} e^{-Q_1(k)}}, \quad (18)$$

where

$$Q_1(s) := \sum_{j \in D_n} \left[y_{jn} + (1 - y_{jn}) \frac{J_n - 1}{\tilde{J}_n - 1} \right] \cdot \sum_{m=1}^M \ln[1 + \exp(\beta_m \{x_{jmn} - x_{smn}\})].$$

The advantages and disadvantages of this procedure are similar to those of the Pop.Shares method: it can be directly implemented without using additional information and without incurring additional computational costs. Additionally, this method can be easily

implemented in canned estimation software with minor modifications, making it attractive for practitioners. The disadvantage is that this approximation may be too rough and may cause large finite sample biases. This approach is studied using Monte Carlo experiments in §4.

4. Monte Carlo Experiments

4.1. Introduction

In this section, we report three Monte Carlo experiments that illustrate the application of the different versions of the method outlined in the previous sections. These experiments also shed light on the relative performance of the variations of the method in finite samples; we also caution that, as with any Monte Carlo experiment, the results in this respect are only valid in the context of the experiments considered.

4.2. Assessment of Different Versions of the Method

In the first experiment we analyze the empirical finite sample properties of each version of the method in recovering the true parameters of the model, depending on the number of alternatives sampled \tilde{J} . The structure of this experiment is summarized in Figure 2. The true or underlying model is an RRM model with 1,000 alternatives and 1,000 observations, with a single attribute x distributed Uniform(-1, 1) and with parameter $\beta = 1$. The motivation for considering a single attribute in this experiment was to be able to estimate the true model considering a number of alternatives as large as 1,000 to be used as a benchmark.

The methodology used to implement the RRM model shown in Figure 2 for the Monte Carlo experimentation consists of several steps. First, the choice probability was calculated using the true value of the parameter ($\beta = 1$) in Equation (3). Then these choice probabilities were used to build a discrete cumulative distribution function by alternative. Afterward, a random number Uniform(0, 1) was generated for each observation. Finally, the chosen alternative was determined as the inverse of the cumulative distribution function, evaluated for each random number.

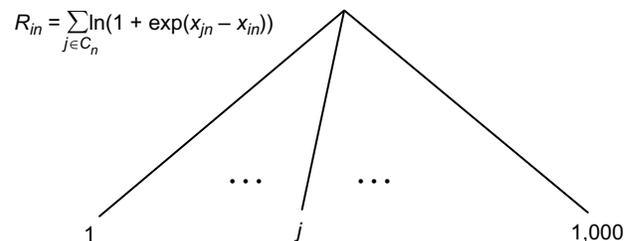


Figure 2 Structure of the Random Regret Model for the Monte Carlo Experiment

Note. $N = 1,000$; $J = 1,000$; $\tilde{J} = 5, 15, 30$, and 50.

The sampling protocol used to draw alternatives D_n from the choice-set C_n in this experiment was the following. First, the chosen alternative for each observation was included. Then nonchosen alternatives were randomly sampled, without replacement, to make a total of $\tilde{J} = 5, 15, 30,$ and 50 . The sampling protocol used to draw alternatives \tilde{D}_n , when it was considered to be different from D_n , was a simple random sample of \tilde{J} alternatives from C_n .

Under this setting we estimated the model using five different methods. The first method corresponds to the True model, a model where all alternatives are considered in the choice-set. This model acts as a benchmark, both in terms of the maximum quality that can be attained for the estimators and of the maximum estimation time.

The second estimation method corresponds to a Truncated version of the problem where only the elements in the subset D_n are used to build the term $\hat{R}_{in}^{\text{Truncated}} = \sum_{j \in D_n} \ln(1 + \exp(\beta(x_{jn} - x_{in})))$. This method acts as a benchmark in terms of the minimum quality that can be attained for the estimators.

The third estimation method considered is Resampling, a method in which an alternative set \tilde{D}_n is sampled to build the term \hat{R}_{in} . In this application, \tilde{D}_n was drawn as a random sample without replacement so that the expansion factors are calculated as $w_{jn} = J/\tilde{J}$. The CQLL considered in this case is the one shown in Equation (15).

The fourth estimation method considered is Pop.Shares. In this case $\tilde{D}_n = D_n$, the expansion factors are calculated using the sample shares as an approximation of the choice probabilities, and the CQLL is the one shown in Equation (17).

The final estimation method considered is 1_0. In this case $\tilde{D}_n = D_n$, the expansion factors are calculated using the observed choice as an approximation of the choice probabilities, and the CQLL considered in this case is the one shown in Equation (18).

The model was generated 100 times for different values for \tilde{J} . For each repetition of the model, we regenerated the attribute x , the choices, and the sets D_n and \tilde{D}_n . Estimation was performed using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) (Fletcher 1980) algorithm coded in the *optim* package of the open-source software R (R Development Core Team 2008) on an IBM eServer with a CPU Intel Xeon X5560 of 2.8 GHz and 12 GB RAM.

For each model estimated we report the following statistics to assess the empirical finite sample properties of each method in estimating the model coefficient β .

Bias. Difference between average estimator within the 100 repetitions and the true value of the parameter. The Bias should tend to zero if the mean of the sampling distribution is equal to the true value.

Root Mean Squared Error (RMSE). Square root of the sum of the sampling variance and the square of the bias. The smaller the RMSE, the better the method is in terms of small sample efficiency.

t-test. Ratio between the bias and the sampling standard deviation of the average of the estimators. This statistic can be used to test the null hypothesis that the mean of the sampling distribution is equal to its respective true value.

Count. Number of times, among the 100 repetitions, the estimator of each repetition is within a 75% confidence interval of the true value constructed using the sampling variance from all of the repetitions. This statistic is usually termed the empirical coverage. The larger this statistic is, the better the performance of the method. The closer to 75 this statistic is, the closer its empirical distribution is to its theoretical sampling distribution.

Together with these statistics, we report in Table 1 the respective \tilde{J} , the estimation time in minutes (Time), and the number of times—within the 100 repetitions—that the model was not estimable because of an error in the optimization procedure (Error).

The estimation results are also summarized in Figure 3. The abscissa corresponds to the \tilde{J} and the ordinate depicts the estimator $\hat{\beta}$ of the single model parameter. The value of $\tilde{J} = 1,000$ is not presented in scale, and the values of $\hat{\beta}$ are limited to those between 0.0 and 2.0. The true value of $\beta = 1.0$ is highlighted with a horizontal line. The estimators obtained for each method and repetition are drawn in gray with the respective symbols detailed in the legend of Figure 3 for each method. The average of the estimators, within the 100 repetitions, is marked with a larger symbol for each method.

The estimators of the True model, the one estimated using $\tilde{J} = 1,000$, are depicted with a dot in Figure 3. As expected, this model performs well. The average of the 100 repetitions is almost equal to the true value of $\beta = 1.0$, and each repetition is close and symmetrically around it. This is reaffirmed by the statistics deployed in Table 1. The Bias is about 0.5%. The RMSE is about 8%, and t-tests are far below the critical value of 1.984 to erroneously reject the null hypothesis that $\beta = 1.0$. Also, none of the 100 repetitions failed, and the empirical coverage was 76, which is almost equal to its nominal value of 75. Finally, the estimation time was about one hour per repetition on average.

The estimators of the Truncated method, the one estimated ignoring the impact of the truncation of the regret function caused by the sampling, are depicted with an x in Figure 3. As expected, although the Bias decreases with \tilde{J} (see Table 1), the results are still very poor for $\tilde{J} = 50$. Not even one estimator falls in the 0.0–2.0 range depicted in Figure 3. Table 1 shows that the Bias is above 25,000% compared to the true value and so is the RMSE. Also, for all values of \tilde{J} ,

Table 1 Statistical Analysis of $\hat{\beta}$ for Different Methods to Estimate RRM Model with Sampling of Alternatives, Varying \tilde{J}

Method	Bias	RMSE	<i>t</i> -test	Count	\tilde{J}	Time (min)	Error
True	0.005105	0.08092	0.06322	76	1,000	64.61	0
Truncated	347.0	390.4	1.939	12	5	0.002615	0
	284.4	288.6	5.794	0	15	0.01639	0
	270.8	273.1	7.703	0	30	0.05473	0
	259.0	260.2	10.15	0	50	0.1447	0
Resampling	0.3848	0.6276	0.7763	0	5	0.002875	68
	-0.01060	0.4702	0.02255	88	15	0.01984	0
	0.06234	0.5419	0.1158	94	30	0.07124	0
	0.02292	0.3844	0.05972	97	50	0.1841	0
Pop.Shares	0.06136	0.8994	0.06838	0	5	0.003164	4
	-0.06717	0.3184	0.2158	75	15	0.02044	0
	-0.01191	0.3216	0.03705	90	30	0.07146	0
	-0.01601	0.1993	0.08058	85	50	0.1813	0
1_0	363.9	414.1	1.841	0	5	0.004880	19
	292.5	296.7	5.923	0	15	0.02359	0
	287.8	290.0	8.155	0	30	0.09929	0
	287.9	289.0	11.30	0	50	0.3138	0

Note. $J = 1,000$; $N = 1,000$; $\beta = 1$; 100 repetitions; one attribute, distributed $U(-1, 1)$.

the *t*-test is above 1.984, the threshold for erroneously rejecting with 95% confidence the null hypothesis that the mean of the sampling distribution is equal to its respective true value. It is interesting that the best value of the *t*-test occurs for $\tilde{J} = 5$, which can be explained by noting that the sample variance is larger for such small \tilde{J} . Finally, none of the 100 repetitions of the estimation procedure failed and the estimation time was on average less than 10 seconds for $\tilde{J} = 50$. As a conclusion and completely in line with expectations, the Truncated method performs very poorly in all

aspects for small \tilde{J} , although it can be noted that results improve as \tilde{J} grows, slightly but steadily.

The estimators of the Resampling method, which is obtained by maximizing the CQLL shown in Equation (15), are depicted with a circumference in Figure 3. This estimation method performs acceptably with \tilde{J} as small as 30. From that point, the Bias is below 6% and the *t*-test is far below the critical value for rejecting the null hypothesis that β is equal to its true value. The RMSE is not as small as with the True model but is 600 times below the Truncated one. Also, it is interesting that 68 out of 100 repetitions failed for $\tilde{J} = 5$, but none failed for larger \tilde{J} . This may be explained because the fundamental part of the method is to gather a proper estimate of the regret function with a reduced number of alternatives, and maybe with $\tilde{J} = 5$ the estimator of R_{in} is so poor that it results in the estimation procedure becoming unbounded or undefined. Another possible explanation is that there might be a limitation of the estimation procedure BFGS in this context.

The estimation time of the Resampling method took about 11 seconds on average for $\tilde{J} = 50$, which is very similar to the Truncated method, and about 350 times smaller than that of the True model. Finally, the Count for the Resampling method is higher than the nominal value of 75. This may reflect that 100 repetitions may not be enough in this case for providing a proper account of the sampling distribution or that the finite sample distribution is not well behaved. As a conclusion, these results suggest that although the Resampling method works asymptotically, various finite sample properties, particularly the Bias, are below 6% with \tilde{J} as small as 30 out of 1,000. However, statistical testing with finite samples should be treated with care since results

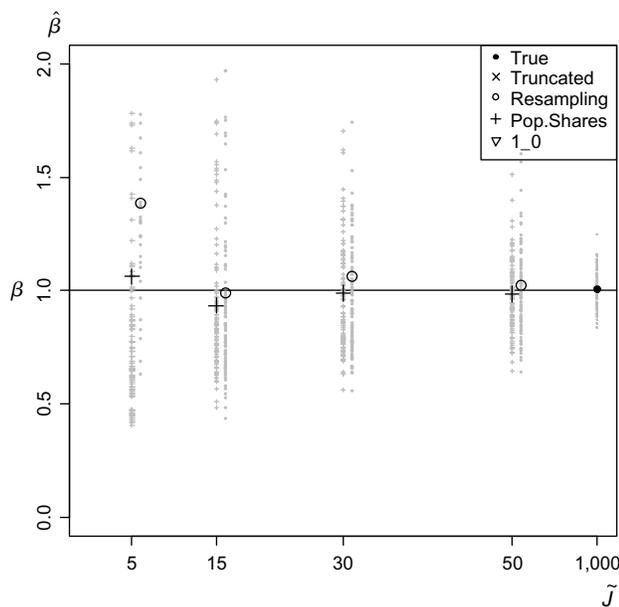


Figure 3 Estimators for Different Methods

Note. 100 repetitions for various (\tilde{J}).

suggest that the t -tests may have low power. Further investigation in this final issue is needed.

The estimators of the Pop.Shares method, the one obtained by maximizing the CQLL shown in Equation (17), are depicted with a cross in Figure 3. This method performs as well as the Resampling method. For some values of \tilde{J} , Pop.Shares is superior and for others Resampling is superior. Failed estimations also occur only for $\tilde{J} = 5$, but now in only four out of 100 repetitions, which suggests that this method is more robust with regard to this aspect. Estimation times are also of the same order of magnitude as for the Resampling method. As with the Resampling method, the Count in this case is larger than its nominal value. As a conclusion, the results suggest that the Pop.Shares method works as well as the Resampling method for finite samples.

Finally, the estimators of the 1_0 method, the one obtained by maximizing the CQLL based in the expression shown in Equation (18), are depicted with an inverted triangle in Figure 3. The results obtained with this method are very poor, in fact almost as poor as the results obtained with the Truncated method. As a conclusion, although the 1_0 method works asymptotically, the finite sample properties in this application are very poor, with Bias as large as those of the Truncated model. The Resampling and the Pop.Shares methods both showed substantially better results.

4.3. Sensitivity to the Variance of x

The second experiment is devised to analyze the relative performance of the methods when changing the variance of the data. The experiment is equivalent to the one described in the previous section in various aspects. There is also only one attribute x , the true parameter is $\beta = 1$, and there are 1,000 observations. In turn, the true choice-set in this case has 500 alternatives for

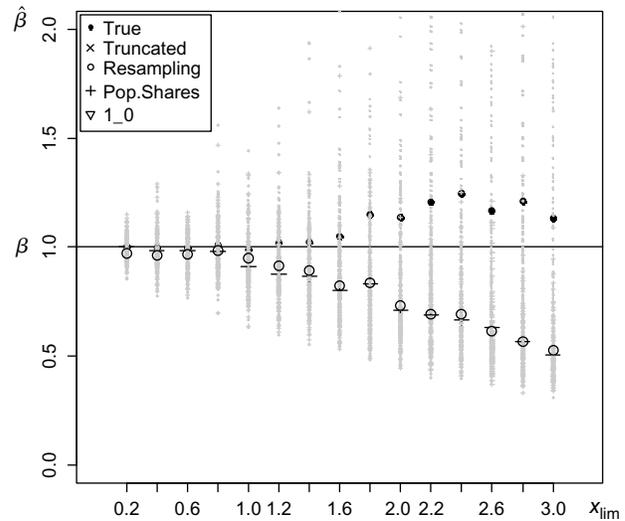


Figure 4 Estimators for Different Methods

Note. 100 repetitions, $J = 500$; $N = 1,000$; $\tilde{J} = 30$; X follows $U(-x_{lim}, x_{lim})$.

all individuals, and 30 alternatives are sampled. The attribute x is distributed $Uniform(-x_{lim}, +x_{lim})$, where x_{lim} varies from 0.2 to 3.0, in steps of 0.2. The data was generated 100 times and the estimators of the five methods are reported in Figure 4 and Table 2. In Table 2, only the results for $x_{lim} = 0.2, 1.0$, and 3.0 are reported.

The working hypotheses are fourfold. (1) A larger variance x will reflect a larger variance of the statistic that is being estimated (R_{in}) by the proposed method, which implies that a larger size J of the sampled choice-set would be needed to attain the same level of error at a given confidence level. (2) A larger variance of x would also imply a larger variance of the choice probability, implying that a larger N would be needed to maintain the statistical properties. This would impact

Table 2 Statistical Analysis of $\hat{\beta}$ for Different Methods to Estimate RRM

Method	x_{lim}	Bias	RMSE	t -test	Count	Time (min)	Error
True	0.2	0.001229	0.034632	0.035514	75	17.39	0
	1.0	-0.010732	0.071602	0.151598	77	17.21	0
	3.0	0.130163	0.504423	0.267089	79	18.54	0
Truncated	0.2	143.9	144.4	11.75	0	0.09732	0
	1.0	125.6	126.1	11.94	0	0.1113	0
	3.0	52.73	52.94	11.27	0	0.1108	0
Resampling	0.2	-0.02873	0.05970	0.5490	62	0.0686	0
	1.0	-0.05144	0.3231	0.1613	97	0.07934	0
	3.0	-0.4710	0.5009	2.765	4	0.1007	0
Pop.Shares	0.2	0.00008789	0.05354	0.001641	76	0.0677	0
	1.0	-0.08756	0.1569	0.6725	60	0.08052	0
	3.0	-0.4961	0.5148	3.605	3	0.09911	0
1_0	0.2	226.5	226.9	18.15	0	0.1136	0
	1.0	142.3	142.7	13.34	0	0.1374	0
	3.0	58.23	58.42	12.38	0	0.1355	0

Note. $J = 500$; $N = 1,000$; $\tilde{J} = 30$; X follows $U(-x_{lim}, x_{lim})$.

both the true model and the estimation with sampling of alternatives. (3) If the variance is too large, this may eventually cause numerical problems in the estimation methods. (4) Finally, an increase in the variance of the attribute will increase the level of information, improving the efficiency of the estimator.

Figure 4 and Table 2 show that just as for the experiments reported in §4.2, the methods Truncated and 1_0 have a poor performance. No realizations in the 0.0–2.0 range are observed for both methods. Furthermore, the Resampling and Pop.Shares methods show very similar performance as reported in the previous subsection. For $x_{lim}=1.0$ and smaller, both methods perform very well, with biases below 9%, small RMSE, and t -tests below the critical value to erroneously reject the null hypothesis that the coefficient is equal to its true value. Things become worse for larger x_{lim} , both for Resampling and for the Pop.Shares method. This can be explained by a mixture of Hypotheses 1 and 2. The fact that the Bias grows for $x_{lim} > 1.6$, even for the true model, suggests that from that point onward, the second hypothesis is more relevant, which means that a larger N is required to maintain good statistical properties.

It is interesting that the Bias is positive for $x_{lim} > 1.6$ for the True model and negative for Resampling and Pop.Shares. This suggests that the first and second hypotheses may be acting in different directions for Resampling and Pop.Shares methods. Further investigation is needed to properly explain this behavior.

There is no support for Hypothesis 3 for the range of values of x_{lim} analyzed because none of the estimations failed. There is also no support for Hypothesis 4 for the range of x_{lim} considered. For all cases the RMSE grows with x_{lim} , which suggests that the effect in terms of efficiency is offset by the other effects. For smaller x_{lim} (not reported) the adjustment slightly deteriorated in a similar way for both the Resampling and Pop.Shares method.

In conclusion, results suggest that the variance of the data impacts the \tilde{J} that is needed to attain a certain statistical quality of the estimators. This implies that it is not possible to suggest a proper \tilde{J} for all contexts—for example, as a given fraction of J . In §4.4 we propose a method to choose the number of alternatives to be sampled in practice.

4.4. Selection of \tilde{J} in Practice

The third Monte Carlo experiment was devised to illustrate how one may decide which \tilde{J} to use in a practical application. This procedure is not exclusive to the problem of sampling of alternatives in RRM but also applies to MEV, logit mixture, and logit.

In general, as was highlighted in §4.3, it is not possible to provide a recommendation for \tilde{J} as a fraction of J . The \tilde{J} needed will depend, among other things, on the

distribution of the data, the number of attributes, the true value of the parameters, the number of observations N , the optimization procedure, and the computing capabilities. The choice of a proper value for \tilde{J} involves a trade-off between estimation time and quality of the estimators. The larger \tilde{J} is, the longer it will take to estimate the model, but the estimates will be better.

In a practical application, the researcher will have a single database. To assess the fit of the model for a given \tilde{J} , the researcher can sample K sets $D_k(\tilde{J})$ and $\bar{D}_k(\tilde{J})$, obtaining a respective series of $\hat{\beta}_k$. With this, the following two statistics can be calculated:

$$\bar{\beta} = \frac{1}{K} \sum_{k=1}^K \hat{\beta}_k(\tilde{J}) \quad \text{and} \quad \hat{\sigma}_{\hat{\beta}} = \sqrt{\frac{1}{K-1} \sum_{k=1}^K (\hat{\beta}_k - \bar{\beta})^2}.$$

The statistic $\hat{\sigma}_{\hat{\beta}}$ can be seen as an estimator of the noise of the estimation parameter, which is equivalent to the concept of noise described for the estimation of the score in the demonstration shown in the appendix, but also includes the finite sample bias caused by using the impractical conditional maximum likelihood shown in Equation (10), instead of the log-likelihood of the true model.

The larger the \tilde{J} , the smaller the $\hat{\sigma}_{\hat{\beta}}$ will be. Eventually, when $\tilde{J} = J$, $\hat{\sigma}_{\hat{\beta}}^2 = 0$. The noise $\hat{\sigma}_{\hat{\beta}}$ is a measure that the researcher may want to constrain and trade-off with estimation time, when choosing the \tilde{J} to be used in practice.

If the researcher is able to estimate the model with the full choice-set to obtain $\hat{\beta}(C)$, $\bar{\beta}$ can be used to estimate what can be defined as the sampling bias $\bar{\beta} - \hat{\beta}(C)$. Note that this bias is not the same as the one we considered in the context of the previous experiments. In those experiments, the bias was calculated with respect to the true value of the parameter. In this case, the bias is calculated with respect to the estimator obtained when considering the full choice-set and for a given data set. This notion of bias is equivalent to the concept of bias described for the estimation of the score in the appendix. The larger \tilde{J} , the smaller the sampling bias will be. Eventually, when $\tilde{J} = J$, the sampling bias will be zero. The statistic $\bar{\beta} - \hat{\beta}(C)$ is thus a measure that the researcher may want to constrain and trade off with estimation time when choosing the \tilde{J} to be used in practice.

If the researcher is not able to estimate the model with the full choice-set to obtain $\hat{\beta}(C)$, $\bar{\beta}$ can still be used directly to choose the proper \tilde{J} by checking its stability. This is analog to the way that the number of draws has to be chosen when estimating a model by maximum simulated likelihood, as suggested by Chiou and Walker (2007).

To illustrate this procedure we report a Monte Carlo experiment in which the true model is an RRM that

Table 3 Practical Determination of \tilde{J} in an RRM

\tilde{J}	Sampling bias	$\tilde{\beta}$	$\hat{\sigma}_{\tilde{\beta}}$	Time (seconds)	Error
5	-0.08610	0.8461	0.3090	0.1663	4
15	0.09320	1.025	0.6257	1.180	0
30	-0.03381	0.8984	0.1526	3.792	0
50	-0.02142	0.9107	0.1278	10.29	0
100	-0.001991	0.9302	0.0893	45.40	0

Note. $J = 1,000$; $N = 1,000$; 30 repetitions; population shares method $\hat{\beta}(C) = 0.9322$; time (C) = 57 minutes.

has 1,000 alternatives; there are 1,000 observations and a single attribute x distributed $\text{Uniform}(-1, 1)$ with parameter $\beta = 1$. The estimation of the true model for a particular realization of x results in $\hat{\beta}(C) = 0.9322$ and is performed in about 57 minutes.

Table 3 summarizes the statistics obtained using the Pop.Shares method for various \tilde{J} , considering $R = 30$ repetitions for each \tilde{J} . It should be noted that the repetitions in this case are not the same as in the previous experiments where x , the choices, and the choice-sets were regenerated each time. In this case, the only thing that changes across repetitions is $D_r(\tilde{J}) = \tilde{D}_r(\tilde{J})$. For completeness, we also include in Table 3 the number of times, within the 30 repetitions, that the optimization procedure failed.

To choose a proper \tilde{J} , the researcher would have to make a decision on the desirable sampling bias, shift in $\tilde{\beta}$, noise $\hat{\sigma}_{\tilde{\beta}}$, and estimation time. An additional criterion would also be to consider not having errors in the estimation procedure. These criteria could be accomplished, for example, for $\tilde{J} = 30$ because it has a sampling bias below 5%, $\hat{\sigma}_{\tilde{\beta}}$ of about 15%, estimation time below five seconds, the shift in $\tilde{\beta}$ about 1%, and not a single failed estimation within the 30 repetitions.

4.5. Conclusion

The Monte Carlo experiments show that the method we proposed for sampling of alternatives in the context of RRM models is practical. Results suggest that Pop.Shares and Resampling versions of the method seem to provide suitable results for samples of alternatives as small as 30 out of 1,000 alternatives.

Results also illustrate that the sample size that is needed for obtaining a given level of quality of the estimates depends on many features, including the distribution of the data. This means that it is not possible to provide a simple criterion for the choice of the proper sample size, such as that \tilde{J} has to be some fraction of the true J . In turn, the approach described in §4.4 is recommended to decide based on the trade-off between estimation time and goodness of fit. This approach is valid not only for RRM but also for sampling of alternatives in MEV, logit mixture, and logit.

The experiments also suggest that although the finite sample bias can be small for small values of \tilde{J} , the

power of the t -tests with finite samples may be low. This can be inferred from the observation that empirical coverage tended to be larger than its nominal value for most cases. This issue should be analyzed in further research.

Finally, it should be noted that, as with any Monte Carlo experiment, the relative assessment of the methods presented in this section are only valid for the experiments reported and does not represent a complete description of the finite sampling properties of the estimators.

5. Application with Real Data

In this section we revisit a real data experiment used by Chorus (2010) to demonstrate the RRM model. The data concerns parking choices and was collected by Van der Waerden, Borgers, and Timmermans (2008) on the campus of Eindhoven University of Technology. The choice-set consists of 14 parking lot alternatives and 350 cases (which is the sample used for estimation by Chorus 2010).

The choice model considers four attributes of the parking lots. The first is NR_SPACES, which corresponds to the number of spaces available in each parking lot. The second is ROOM_MANEUV, which is a dummy that takes value 1 if the parking lot has extra space for making maneuvers. The third attribute is RIGHT_OF_WAY, which is a dummy that takes value 1 if the driver has right-of-way when leaving the parking lot. The final attribute is DISTANCE, which corresponds to the distance between the agent's workplace and the parking lot and is discretized as follows: equals 1 when the distance is approximately 100 meters, equals 2 when the distance is approximately 300 meters, and equals 3 when it is approximately 500 meters.

Although the choice-set may not seem particularly large ($J = 14$), as a proof of concept of the method, we preferred not to generate a pseudo-synthetic experiment with a larger choice-set (as in Bierlaire, Bolduc, and McFadden 2008) and instead used the real data as they were. The reason is that a pseudo-real data set will not offer fundamentally new insights compared to the experiments described in §4; additionally, our real data will illustrate the behavior of the method with various attributes and provide additional support to the statement that the choice of the proper \tilde{J} cannot be specified as a given fraction of J .

Table 4 summarizes the estimators obtained for the true RRM model of parking lot choices using all 14 alternatives available. These results are the same as the ones reported by Chorus (2010). We will use the Resampling method to estimate the model with sampling of alternatives, varying \tilde{J} from 2 to 14. The model was estimated 30 times for each \tilde{J} .

We report in Figure 5 the average estimators $\tilde{\beta}$ within the 30 repetitions. The values of $\beta(C)$, which

Table 4 RRM True Model of Parking Lot Choices

	$\hat{\beta}$	S.E.
NR_SPACES	0.08671	(0.01430)
ROOM_MANEUV	0.09066	(0.02750)
RIGHT_OF_WAY	0.03387	(0.02763)
DISTANCE	-1.444	(0.4517)
$L(0)$		-923.7
$L(\hat{\beta})$		-404.7
ρ^2		0.5619
N		350
J		14

is the respective parameter attained with the true model reported in Table 4, are depicted with a dashed line. We also report a bandwidth of 10% deviation from each $\beta(C)$. As expected, all $\hat{\beta}$ get closer to $\beta(C)$ as \tilde{J} grows. However, the speed of convergence is heterogeneous. On one side $\hat{\beta}$ for NR_SPACES is within the 10% bandwidth for $\tilde{J} = 2$. On the other side, for DISTANCE, this occurs only as $\tilde{J} = 13$ out of 14.

In Figure 6 we report the standard deviation $\hat{\sigma}_{\hat{\beta}}$ for each \tilde{J} . Note that $\hat{\sigma}_{\hat{\beta}} = 0$ for $\tilde{J} = 14$. This value is depicted with a dashed line. As expected, each $\hat{\sigma}_{\hat{\beta}}$ shrinks as \tilde{J} grows. However, as with $\hat{\beta}$, the behavior is heterogeneous. The noise $\hat{\sigma}_{\hat{\beta}}$ for NR_SPACES is always

below 0.04, whereas for DISTANCE, it only occurs for $\tilde{J} = 14$.

The heterogeneity in $\hat{\beta}$ and $\hat{\sigma}_{\hat{\beta}}$ illustrates that when choosing \tilde{J} in a model with various attributes the researcher would have to consider some type of norm to account for the degree of convergence of the full vector of parameters. A robust strategy could be to consider the convergence of the worst behaved parameter. In addition, because for one of the parameters a somehow reasonable convergence is attained only for $\tilde{J} = 93\%$ of J serves to illustrate that the choice of \tilde{J} cannot be settled as a fixed fraction of J .

6. Conclusion

This article proposes a method to obtain consistent, asymptotically normal, and efficient estimators (i.e., efficient relative to any other estimator using the same sample) for the problem of sampling of alternatives in the context of random regret minimization models. In light of the fact that runtimes of RRM models increase almost quadratically with choice-set size, finding a proper way to estimate RRM-models on sampled choice-sets is a crucial condition to ensure that the RRM approach remains a feasible and attractive alternative

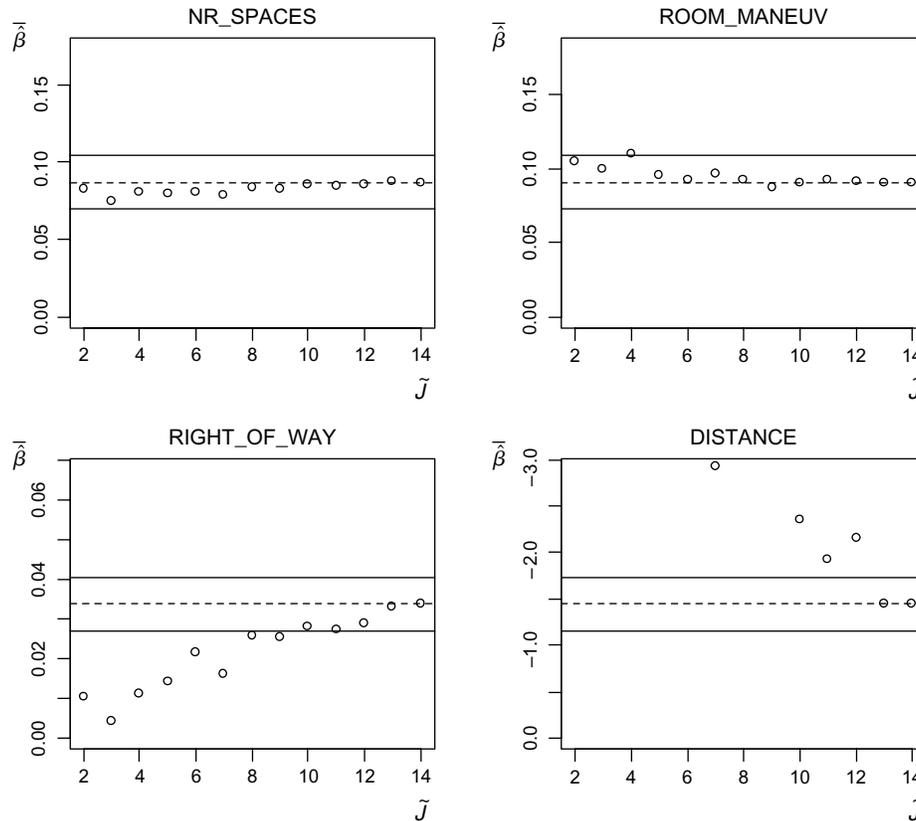


Figure 5 Average $\hat{\beta}$ as a Function of \tilde{J} for the RRM of Parking Lot Choice

Note. Thirty repetitions; resampling method.

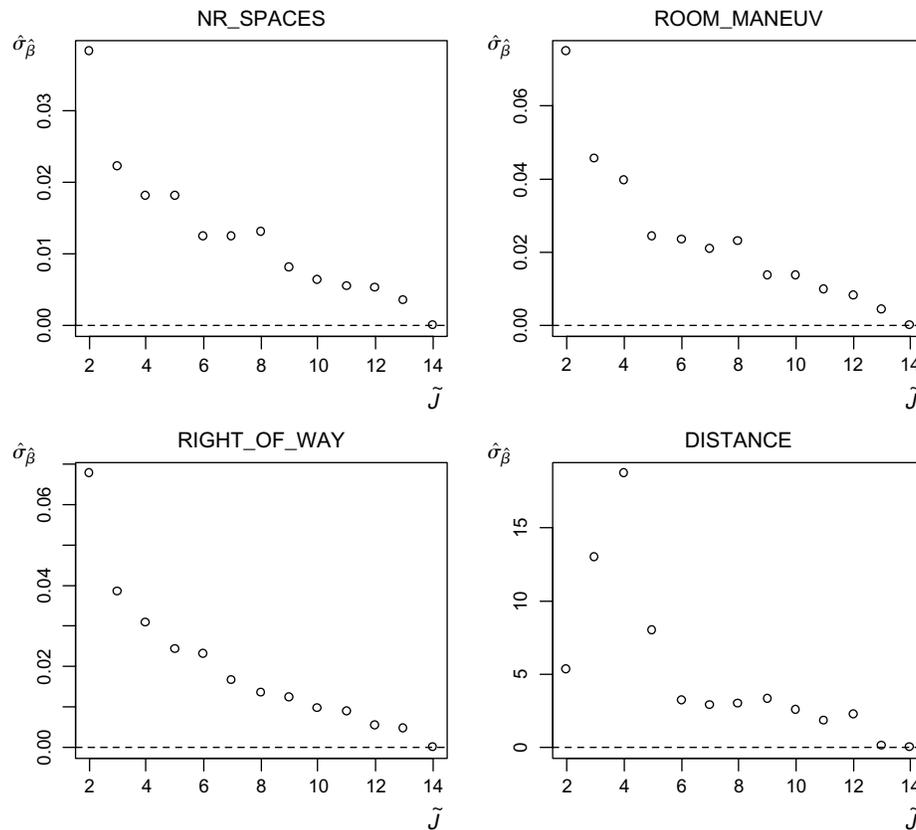


Figure 6 Average noise $\hat{\sigma}_\beta$ as a Function of \tilde{J} for the RRM of Parking Lot Choice

Note. Thirty repetitions; resampling method.

for random utility maximization models in the context of (very) large choice-sets. Given that the RRM model, even when written in logit form (i.e., with i.i.d. errors), does not exhibit the IIA property, McFadden's (1978) result cannot be applied to obtain a proper correction term when choice-sets are sampled. To overcome this situation, a tailor-made correction approach for RRM models is presented in this paper, which is a direct extension of the approach developed by Guevara and Ben-Akiva (2013b) to address a similar problem in RUM-based MEV models.

In line with expectations, Monte Carlo experiments showed that sampling of alternatives causes a significant bias in the estimators of the RRM-model parameters when no correction is applied. In addition, these experiments as well as an application on real data show that the proposed method for correcting the terms that get truncated because of the sampling performed reasonably well. In cases where the researcher has full control of the data and it is possible to obtain an additional sample to expand the truncated regret function, the method proposed is easily applicable. When it is not possible to resample, the method requires knowledge of the choice probabilities in order to build the expansion factors. In this final case, one practical approximation method showed reasonably good results.

The sample size required to obtain good estimators while sampling alternatives (in general, not only for RRM models) will vary on a case-by-case basis and cannot be expressed as a percentage of the cardinality of the true choice-set. Using synthetic and real data, we show that an appropriate strategy to determine if the size of the sample of alternatives is large enough is to test the stability of the estimators with different numbers of alternatives sampled and to analyze the sampling bias and noise.

Acknowledgments

Funding for this research came from Fondecyt, Chile [Grant 11110131], and the Singapore National Research Foundation under the Future Urban Mobility research group of the Singapore-MIT Alliance for Research and Technology. Support from The Netherlands Organisation for Scientific Research (NWO), in the form of VIDI [Grant 016-125-305], is gratefully acknowledged by the second author. The authors are also grateful for the valuable comments of two anonymous referees and of the participants of the 92nd Annual Meeting of the Transportation Research Board (TRB) and the Eighth Triennial Symposium on Transportation Analysis (TRISTAN VIII), where preliminary versions of this paper were presented. All Monte Carlo and real data experiments were generated and/or estimated using the open-source software R (R Development Core Team 2008). The authors would also like to thank the editorial contributions of Katherine Rosa.

Appendix

The demonstration of the theorem is analog to the two step procedure used by Train (2009, pp. 247–257) to derive the asymptotic distribution of simulation-based estimators. The first step consists in the derivation of the distribution of the approximated score

$$\begin{aligned} \hat{g}(\beta) &= \frac{1}{N} \sum_{n=1}^N \frac{\partial \ln \hat{\pi}_n(\beta | D_n)}{\partial \beta} \\ &= \frac{1}{N} \sum_{n=1}^N \frac{\partial}{\partial \beta} \ln \frac{e^{\hat{W}_{in}(\tilde{D}_n) + \ln \pi(D_n | i)}}{\sum_{j \in D_n} e^{\hat{W}_{jn}(\tilde{D}_n) + \ln \pi_n(D_n | j)}} \end{aligned}$$

relative to the true score

$$\begin{aligned} g(\beta) &= \frac{1}{N} \sum_{n=1}^N \frac{\partial \ln \pi_n(\beta | D_n)}{\partial \beta} \\ &= \frac{1}{N} \sum_{n=1}^N \frac{\partial}{\partial \beta} \ln \frac{e^{W_{in}(C_n) + \ln \pi(D_n | i)}}{\sum_{j \in D_n} e^{W_{jn}(C_n) + \ln \pi_n(D_n | j)}} \end{aligned}$$

The second step is to derive the distribution of $\hat{\beta}$, noting that $\hat{\beta}$ is the root of the equation $\hat{g}(\hat{\beta}) = 0$.

In what follows we provide a summarized account of the first step, in order to highlight why $\hat{W}_{in}(\tilde{D}_n)$ needs to be an unbiased estimator of W_{in} and why the variance of \hat{W}_{in} needs to be bounded and decrease with \tilde{J} , which also means that $\hat{W}_{in}(\tilde{D}_n)$ is a consistent estimator of W_{in} . The reader is referred to Train (2009, pp. 247–257) or Guevara and Ben-Akiva (2013b) for further details.

To simplify the notation we will assume that $\tilde{D}_n = \tilde{D}$ for all n . Consider $\hat{g}(\beta)$ in the vicinity of the true values β^* in the following form:

$$\hat{g}(\beta^*) = \underbrace{g(\beta^*)}_{A_1} + \underbrace{[E(\hat{g}(\beta^*)) - g(\beta^*)]}_{A_2} + \underbrace{[\hat{g}(\beta^*) - E(\hat{g}(\beta^*))]}_{A_3}$$

The first term $A_1 = g(\beta^*)$ is the statistic that is being approximated by $\hat{g}(\beta^*)$. The second term A_2 corresponds to the bias of the estimator of $g(\beta^*)$, and the third term A_3 is the noise of the approximation.

The noise (A_3) corresponds to the deviation of $\hat{g}(\beta^*)$ from its expected value, which will depend on a particular draw of the alternatives to construct the choice-set \tilde{D} . Because $\hat{W}_{in}(\tilde{D}_n)$ is bounded and decreases with \tilde{J}_n , we can claim that the same occurs with the variance of the noise. This can be expressed as $\text{Var}(A_{3n}) = S_n / \tilde{J}$, where S_n is the variance of A_3 for a given n when $\tilde{J} = 1$. Then by the generalized version of the central limit theorem (see, e.g., Train 2009, p. 246), the noise A_3 will have the following limiting distribution:

$$\sqrt{N} A_3 \xrightarrow{d} \text{Normal}(0, \mathbf{S} / \tilde{J}),$$

where \mathbf{S} is the population mean of S_n . Consequently, the asymptotic distribution of the noise A_3 will be

$$A_3 \overset{d}{\sim} \text{Normal}(0, \mathbf{S} / (\tilde{J}N)),$$

and the noise will vanish as N increases, even if \tilde{J} is fixed.

The bias term A_2 can be studied by taking a second order Taylor's approximation of $\hat{W}_{in}(\tilde{D}_n)$ around $\hat{W}_{in}(\tilde{D}_n) = W_{in}$. Noting that $\hat{g}_n(\beta, W_{in}) = g_n(\beta)$, it follows that

$$\begin{aligned} \hat{g}_n(\beta) &= g_n(\beta) + \frac{\partial \hat{g}_n}{\partial \hat{B}_n} [\hat{W}_n(\beta) - W_n(\beta)] \\ &\quad + \frac{1}{2} \frac{\partial^2 \hat{g}_n}{\partial \hat{B}_n^2} [\hat{W}_n(\beta) - W_n(\beta)]^2 + o_n. \end{aligned}$$

Then taking expectations (over possible realizations of the set \tilde{D}_n), recalling that $\hat{W}_{in}(\tilde{D}_n)$ is an unbiased estimator of W_{in} , and considering that the discrepancy o_n has zero mean, this Taylor's approximation can be rewritten as

$$E(\hat{g}_n(\beta)) - g_n(\beta) = \frac{1}{2} \frac{\partial^2 \hat{g}_n(\beta)}{\partial \hat{W}_n^2} \text{Var}(\hat{W}_n(\beta)).$$

The fact that $\text{Var}(W_n(\beta))$ is bounded and decreases with \tilde{J} can be captured by the expression $\text{Var}(\hat{W}_n(\beta)) = K_n / \tilde{J}$, where K_n is scalar. Then the expected value of the bias A_2 can be rewritten as $A_2 = Z / \tilde{J}$, where Z is the sample average of $(K_n / 2) (\partial^2 \hat{g}_n / \partial \hat{B}_n^2)$.

The bias A_2 will vanish as N increases, if and only if \tilde{J} increases also with N . Otherwise, $\hat{g}(\beta)$ will be an inconsistent estimator of $g(\beta)$. Instead, an even stronger assumption is required to achieve asymptotic normality. To understand why, consider the bias A_2 normalized for sample size N

$$\sqrt{N} A_2 = \frac{\sqrt{N}}{\tilde{J}} Z.$$

This term will vanish as N increases, if and only if \tilde{J} increases faster than does \sqrt{N} . Otherwise, the estimator $\hat{g}(\beta)$ will have neither a limiting nor an asymptotic distribution.

In summary, if \tilde{J} increases with N at any rate, $\hat{g}(\beta) \xrightarrow{p} g(\beta)$, and when \tilde{J} increases faster than \sqrt{N} , $\hat{g}(\beta)$ will be asymptotically normal. Given that $\hat{g}(\beta) \xrightarrow{p} g(\beta)$, the limiting and asymptotic distributions of $\hat{g}(\beta)$ will be the same as those of $g(\beta)$.

References

- Azaiez I (2010) Sampling of alternatives for logit mixture models. Unpublished masters thesis. Transport and Mobility Laboratory, École Polytechnique Fédérale de Lausanne, Switzerland.
- Ben-Akiva M (1973) Structure of passenger travel demand models. Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA.
- Ben-Akiva M (1974) Note on the specification of a logit model with utility functions that include attributes of competing alternatives. Working paper, Department of Civil Engineering, Massachusetts Institute of Technology, Cambridge, MA.
- Ben-Akiva M, Watanatada T (1981) Application of a continuous spatial choice logit model. Manski CF, McFadden D, eds. *Structural Analysis of Discrete Data with Econometric Applications* (MIT Press, Cambridge, MA), 320–343.
- Berkovec J, Rust J (1985) A nested logit model of automobile holdings for one vehicle households. *Transportation Res. Part B: Methodological* 19(4):275–285.
- Berndt EK, Hall BH, Hall RE (1974) Estimation and inference in nonlinear structural models. *Ann. Econom. Soc. Measurement* 3(4):103–116.

- Bierlaire M (2003) BIOGEME: A free package for the estimation of discrete choice models. *Proc. 3rd Swiss Transportation Res. Conf. Ascona, Switzerland*. <http://infoscience.epfl.ch/record/117133/files/bierlaire.pdf>.
- Bierlaire M, Bolduc D, McFadden D (2008) The estimation of generalized extreme value models from choice-based samples. *Transportation Res. Part B: Methodological* 42(4):381–394.
- Bowman JL, Ben-Akiva ME (2001) Activity-based disaggregate travel demand model system with activity schedules. *Transportation Res. Part A: Policy Practice* 35(1):1–28.
- Bradley M, Bowman JL, Griesenbeck B (2010) SACSIM: An applied activity-based model system with fine-level spatial and temporal resolution. *J. Choice Modelling* 3(1):5–31.
- Carrasco N (2008) Deciding where to shop: Disaggregate random utility destination choice modeling of grocery shopping in Canton Zurich. Doctoral dissertation, IVT, ETH Zurich.
- Chiou L, Walker J (2007) Masking identification of discrete choice models under simulation methods. *J. Econometrics* 141(2): 683–703.
- Chorus CG (2010) A new model of random regret minimization. *Eur. J. Transport Infrastructure Res.* 10(2):181–196.
- Chorus CG (2012) *Random Regret-Based Discrete Choice Modeling: A Tutorial* (Springer, Heidelberg, Germany).
- Chorus CG, van Cranenburgh S, Dekker T (2014) Random regret minimization for consumer choice modeling: Assessment of empirical evidence. *J. Bus. Res.* 67(11):2428–2436.
- Dagsvik JK (1989) The generalized extreme value random utility model for continuous choice, 1989–41. Report, Center for Economic Research, Tilburg University, The Netherlands.
- Daly A, Hess S, Dekker T (2014) Practical solutions for sampling alternatives in large scale models. *Transportation Res. Record* 2429:148–156.
- Daly AJ (1992) *ALOGIT 3.2 User's Guide* (Hague Consulting Group, The Hague, The Netherlands).
- Fletcher R (1980) *Practical Methods of Optimization*, Vol. 1 (John Wiley & Sons, New York).
- Fosgerau M, Frejinger E, Karlstrom A (2013) A link based network route choice model with unrestricted choice-set. *Transportation Res. Part B: Methodological* 56:70–80.
- Frejinger E, Bierlaire M, Ben-Akiva M (2009) Sampling of alternatives for route choice modeling. *Transportation Res. Part B: Methodological* 43(10):984–994.
- Guevara CA (2010) Endogeneity and sampling of alternatives in spatial choice models. Unpublished doctoral dissertation, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA.
- Guevara CA, Ben-Akiva ME (2013a) Sampling of alternatives in logit mixture models. *Transportation Res. Part B: Methodological* 58:185–198.
- Guevara CA, Ben-Akiva ME (2013b) Sampling of alternatives in multivariate extreme value (MEV) models. *Transportation Res. Part B: Methodological* 48:31–52.
- Hensher DA, Greene WH, Chorus CG (2013) Random regret minimization or random utility maximization: An exploratory analysis in the context of automobile fuel choice. *J. Advanced Transportation* 47(7):667–678.
- Lee BH, Waddell P (2010) Residential mobility and location choice: A nested logit model with sampling of alternatives. *Transportation* 37(4):587–601.
- Lemp J, Kockelman K (2012) Strategic sampling for large choice-sets in estimation and application. *Transportation Res. Part A: Policy Practice* 46(3):602–661.
- McConnel K, Tseng W (2000) Some preliminary evidence on sampling of alternatives with the random parameters logit. *Marine Resource Econom.* 14(4):317–332.
- McFadden D (1974) Conditional logit analysis of qualitative choice behavior. Zarembka P, ed. *Frontiers in Econometrics* (Academic Press, New York), 105–142.
- McFadden D (1976) The mathematical theory of demand models. Stopher P, Meyburg A, eds. *Behavioral Travel Demand Models* (Lexington Books, Lexington, MA), 305–314.
- McFadden D (1978) Modeling the choice of residential location. Karlquist A, Lundqvist L, Snickers F, Weibull JW, eds. *Spatial Interaction Theory and Residential Location* (North Holland, Amsterdam), 75–96.
- McFadden D, Train KE (2000) Mixed MNL models for discrete response. *J. Appl. Econometrics* 15(5):447–470.
- Nerella S, Bhat C (2004) A numerical analysis of the effect of sampling of alternatives in discrete choice models. *Transportation Res. Record* 1894:11–19.
- R Development Core Team (2008) *R: A Language and Environment for Statistical Computing* (R Foundation for Statistical Computing, Vienna).
- Train K (2009) *Discrete Choice Methods with Simulation*, 2nd ed. (Cambridge University Press, New York).
- Van der Waerden P, Borgers A, Timmermans HJP (2008) Modeling parking choice behavior in business areas. *87th Annual Meeting of the Transportation Res. Board, Washington, DC*.