ONLINE DISCRETE CHOICE MODELS: APPLICATIONS IN SMART MOBILITY

By

Mazen Danaf

B.E., Civil and Environmental Engineering, American University of Beirut (2011)
M.E., Transportation, American University of Beirut (2013)

Submitted to the Department of Civil and Environmental Engineering in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy in Transportation

at the

Massachusetts Institute of Technology

June 2019

Signature redacted

Signature of Author: ....
Department of Civil and Environmental Engineering
May 17, 2019

Signature redacted

Certified by .... Moshe E. Ben-Akiva
Edmund K. Turner Professor of Civil and Environmental Engineering, MIT
Thesis Supervisor

Signature redacted

Certified by .... Bilge Atasoy
Assistant Professor in Transportation Engineering and Logistics, TU Delft
Thesis Supervisor

Signature redacted

Accepted by .... Heidi Nepf
Donald and Martha Harleman Professor of Civil and Environmental Engineering
Chair, Graduate Program Committee

ONLINE DISCRETE CHOICE MODELS: APPLICATIONS IN SMART MOBILITY

By

Mazen Danaf

Submitted to the Department of Civil and Environmental Engineering on May 17, 2019 in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Transportation

ABSTRACT

Discrete choice models have been widely applied in different fields to better understand behavior and forecast market shares. Because of their ability to capture taste heterogeneity, logit mixture models have gained increasing interest among researchers and practitioners. However, since the estimation of these models is computationally expensive, their applications have been limited to offline contexts. On the other hand, online applications (such as recommender systems) require users' preferences to be updated frequently and dynamically.

The objective of this dissertation is to develop a methodology for estimating discrete choice models online, while accounting for inter- and intra-consumer heterogeneity. An offline-online framework is proposed to update individual-specific parameters after each choice using Bayesian estimation. The online estimator is computationally efficient, as it uses the data of the individual making the choice only in updating his/her individual preferences. Periodically, data from multiple individuals are pooled, and population parameters are updated offline.

Online estimation allows for new and innovative applications of discrete choice models such as personalized recommendations, dynamic personalized pricing, and real-time individual forecasting. This methodology subsumes the utility-based advantages of discrete choice models and the personalization capabilities of common recommendation techniques by making use of all the available data including user-specific, item specific, and contextual variables.

In order to enhance online learning, two extensions are proposed to the logit mixture model with inter- and intra-consumer heterogeneity. In the first extension, socio-demographic variables and contextual variables are used to model systematic inter- and intra-consumer taste heterogeneity respectively. In the second extension, a latent class model is used to allow for more flexibility in modeling the inter- and intra-consumer mixing distributions.

Finally, the online estimation methodology is applied to *Tripod*, an app-based travel advisor that aims to incentivize and shift travelers' behavior towards more sustainable alternatives. Stated preferences data are collected in the Greater Boston Area and used to estimate the population parameters, which are then used by the app in online estimation. Using the collected data, a large number of synthetic users is simulated, and the recommendation system is tested over several days, and under different scenarios. The results show that the average hit-rate generally increases over time as we learn individual preferences and population parameters.

Thesis Supervisors:

Moshe E. Ben-Akiva: Edmund K. Turner Professor of Civil and Environmental Engineering, MIT
Bilge Atasoy: Assistant Professor in Transportation Engineering and Logistics, TU Delft

ACKNOWLEDGEMENT

First, I am indebted to my thesis co-advisors, Professor Moshe Ben-Akvia and Professor Bilge Atasoy for accepting nothing less than excellence from me. Moshe's guidance and unmatched expertise in the areas of transportation and demand modeling were invaluable. On a personal level, he inspired me by his hardworking and passionate attitude. I will always live by and remember what he told me in the second year of my PhD, "We do not use cookbooks, we are the chefs!"

I am sincerely grateful to Prof. Bilge Atasoy for her patience, dedication, enthusiasm, and immense knowledge. The door to Bilge's office was always open whenever I ran into a trouble spot or had a question about my research or writing. She consistently allowed this dissertation to be my own work, but steered me in the right the direction whenever she thought I needed it.

Besides my co-advisors, I would like to thank my committee members: Prof. Angelo Guevara, Prof. Carolina Osorio, and Prof. Ricardo Daziano, for their insightful comments and encouragement, but also for their constructive feedback which incented me to widen my research from various perspectives. I am also grateful to Carlos Lima De Azevedo, Arun Prakash, Ravi Seshadri and Fang Zhao for sharing their expertise, guidance, and encouragement.

A very special gratitude goes to my mentors, Professor Isam Kaysi and Professor Maya Abou-Zeid who were my master's co-advisors at the American University of Beirut. Since my first day in graduate school, they believed in me like nobody else and gave me endless support. They also taught me the fundamentals of conducting scientific research from defining a research problem to publishing the results.

I would like to thank my friends and colleagues for their feedback, cooperation, and of course friendship. I am grateful to Felix, Xiang, Yifei, and Isabel whom I collaborated with on different parts of this dissertation. I am also grateful to the stimulating discussions and the countless nights that I spent in the lab with my colleagues Haizheng, Eytan, Youssef, Siyu, Ajinkya, Yundi, Peiyu, and Giacomo.

This dissertation would not have been possible without funding from the Advanced Research Projects Agency–Energy (ARPA-E), Ford, the Civil and Environmental Engineering Department at MIT, and the MIT-Singapore Alliance for Research and Technology (SMART).

I would also like to thank my friends, and the members of the Sidney Pacific community and the Lebanese Club at MIT for making my journey a pleasant and friendly experience. Last but not the least, I would like to thank my family. My late grandmother was the first teacher I ever had, and my academic aspirations are primarily attributed to her. I am also thankful to my parents and brother for their continuous support.

CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1: INTRODUCTION

1.1. MOTIVATION

The Random Utility Model (RUM) has been widely applied by researchers in different fields for a variety of reasons such as better understanding behavior, forecasting, and testing the validity of simpler model structures (Walker and Ben-Akiva, 2002). Applications of this model had proven successful in many areas, including transportation, energy, housing, e-business, and marketing (Train, 2009). Because of its ability to capture taste heterogeneity, the logit mixture model has gained increasing interest among researchers and practitioners. The ease of data collection and the development of estimation software made these models widely popular, especially because they allow for estimating individual level preferences (also known as individual part-worths).

Logit mixture models have been estimated using both classical and Bayesian methods. The most commonly used method is Maximum Simulated Likelihood (MSL) (Train, 2009), which uses numerical simulation to approximate integration over multidimensional distributions. In the Bayesian context, the Hierarchical Bayes (HB) estimator of logit mixture has been widely applied and documented (Allenby, 1997; Allenby and Rossi, 1998; Train, 2009, etc.) and implemented in different software packages (Sawtooth, 1999; 2003; 2009; Train, 2006; Dumont et al., 2015). Since both of these methods require a long time for estimation, applications of logit mixture models have been limited to offline contexts.

Overcoming computational limitations enables discrete choice models to be used in interactive decision support systems that aim to make optimal decisions in real-time, such as recommender systems (Guo and Sanner, 2010). These systems require individual preferences to be estimated and updated frequently; otherwise, the users' most recent choices might not be taken into consideration in generating the next recommendation. Chaptini (2005) identified scalability, user-friendliness, and automation as the three main challenges associated with using discrete choice models in personalized recommender systems.

The primary objective of this dissertation is to develop a methodology for estimating and updating individual preferences online using discrete choice models. This methodology can be utilized in real-time systems that learn individual preferences from repeated choices, since it overcomes the three limitations mentioned above (scalability, user-friendliness, and automation). First, it is scalable since online estimation can be performed in real-time, and to the individual making the choice only (which makes it possible to handle a large number of users). Second, it is user-friendly because preferences are estimated using the users' choices rather than revealed or stated preferences (RP/SP) data, or ratings as in traditional recommender systems. Finally, the whole process of estimating and updating preferences can be automated since online estimation can be triggered automatically after each choice.

1.2. PERSONALIZATION IN DISCRETE CHOICE MODELS

1.2.1. Discrete Choice Models

Discrete choice models are disaggregate behavioral models that predict or describe the process by which a decision maker (n) makes a choice among a set of discrete alternatives belonging to a choice set (C_{mn}) in a choice situation (m). The decision maker is characterized by his/her *characteristics*, and the alternatives are characterized by their *attributes*, representing their benefits and costs to the decision maker (Ben-Akiva and Lerman, 1985; Ben-Akiva and Bierlaire, 2003).

The *decision rule* is the process used by the decision maker to evaluate the alternatives and make a choice. Most discrete choice models assume *utility maximization*, which assumes that the decision maker selects the alternative with the highest utility. Using the *random utility framework*, individual n's utility of alternative j in choice situation m, U_{jmn}, is divided into a systematic component and a random component as shown in equation (1.1).

$$U_{jmn} = V_{jmn} + \epsilon_{jmn} \tag{1.1}$$

The systematic component (V_{jmn}) is specified as a function of the attributes of alternative j and the preferences and characteristics of decision maker n. The random component ϵ_{jmn} captures uncertainty, such as the modeler's imperfect knowledge, unobserved attributes, etc. Under utility maximization, the probability of individual n choosing alternative i in choice situation m is expressed as:

$$P(i|C_{mn}) = P\left[U_{imn} \geq U_{jmn}\ \forall\, j \in C_{mn}\right] = P[U_{in} = \max_{j \in C_{mn}} U_{jmn}] \tag{1.2}$$

The logit model assumes that the random components follow an Extreme Value (EV(0,1)) distribution, and that they are independently and identically distributed across individuals, choice situations, and alternatives, resulting in the closed form probability:

$$P(i|C_{mn}) = \frac{\exp(V_{imn})}{\sum_{j=1}^{J_{mn}} \exp(V_{jmn})} \tag{1.3}$$

While the *flat logit* model assumes that the preference parameters do not vary across individuals, the *logit mixture* model with random parameters can account for preference heterogeneity, by assuming that the preference parameters ζ_n follow a distribution in the population, $f(\zeta_n)$. The resulting probability of choosing alternative i can be expressed as:

$$P(i|C_{mn}) = \int_{\zeta_n} P(i|C_{mn}, \zeta_n) f(\zeta_n) d\zeta_n \tag{1.4}$$

The individual-specific parameters ζ_n can be learnt from the repeated choices of individual n (in the case of panel data), and used in personalization.

1.2.2. Personalization

Personalization is defined as the ability to offer content to consumers based on certain knowledge about their preferences and behavior, and the context in which this content is provided, with certain goal(s) in mind (Hagen, 1999; Adomavicius et al., 2008). Personalized recommendations are prevalent in e-commerce and web applications for various products, including books, movies, videos, travel, etc. (Aggarwal, 2016). In the context of smart mobility, personalization has been applied in trip and journey planners (Carreras et al., 2012; Nuzzolo et al., 2013; Chen et al., 2015; Chow et al., 2018, etc.) which provide real-time information and personalized recommendations to users that match their needs and preferences.

The conventional personalization and recommendation techniques (such as collaborative filtering and content-based methods), which mainly rely on item and user profiling, do not take full advantage of the available data, and suffer from a tradeoff between relevancy and diversity (Ziegler et al., 2005). On the other hand, discrete choice models, which have been widely pervasive in transportation, marketing, and demand forecasting, are rarely used in recommender systems despite their ability to integrate item specific, user specific, and contextual data in a single model (Chaptini, 2005).

According to Jiang et al. (2014), the use of discrete choice models in recommender systems can address the major limitations associated with standard recommendation approaches. Unlike typical recommendation techniques, discrete choice models are able to express user utility as a function of the attributes of the different alternatives, and not as a mere measure of similarity obtained by item and user profiling. In addition, these models can account for taste heterogeneity using individual-specific parameters in logit mixture models. Finally, discrete choice models can introduce product diversity in the proposed recommendations (Jiang et al., 2014; Teo et al., 2016).

1.2.3. Logit Mixture Models and Individual-Specific Parameters

According to Revelt and Train (1999; 2000) and Train (2009), researchers and firms are interested in estimating individual-specific parameters as well as their distribution in the population for different reasons such as market segmentation, targeted marketing, and individual forecasting. The most commonly used method for estimating individual level preferences is the HB estimator of logit mixture (Allenby, 1997; Allenby and Rossi, 1998; Train, 2009). This estimator assumes the utility specification shown in equation (1.1):

$$U_{jmn} = V_{jmn}(\zeta_n, X_{jmn}) + \epsilon_{jmn} \tag{1.4}$$

$$\zeta_n \sim N(\mu, \Omega) \tag{1.5}$$

where $\mathrm{U_{jmn}}$ is individual n's unobserved utility of alternative j in menu m, X_{jmn} represents a vector of individual characteristics and alternative attributes, ζ_n represents a vector of individual-specific parameters that is normally distributed in the population with mean μ and covariance matrix Ω, and ϵ_{jmn} is an error term following the extreme value distribution.

The HB estimator is a Gibbs sampler with an embedded Metropolis-Hastings (MH) algorithm which samples from the three conditional posteriors:

1. $\mu|\Omega, \zeta_n$: using a normal Bayesian update with unknown mean and known variance.
2. $\Omega|\mu, \zeta_n$: using a normal Bayesian update with known mean and unknown variance.
3. $\zeta_n|\mu, \Omega$: using a Metropolis-Hastings algorithm.

The estimated individual preferences correspond to the Monte Carlo Markov Chain (MCMC) draws of ζ_n.

Individual preferences can also be estimated using MSL (e.g. Revelt and Train, 2000; Chaptini, 2005), by estimating the distributions of tastes in the population, then conditioning on the observed data for each individual in order to derive the distributions of individual-specific parameters. Chaptini (2005) used this approach to develop an online academic advisor for MIT students that recommends academic courses based on the observed and latent attributes of the courses. However, in this application, student preferences were estimated using RP and SP surveys collected beforehand.

The abovementioned applications of logit mixture models were also limited to inter-consumer heterogeneity (defined as taste variations among different individuals). On the other hand, intra-consumer heterogeneity, which represents perturbations in preferences among choices done by the same individual, was treated as a nuisance factor. This unrealistically assumes a nearly neoclassical consumer with "permanent" individual preferences that do not change over time (Ben-Akiva et al., 2019). In the presence of multiple observations from each individual, it is possible to identify both inter- and intra-consumer heterogeneity. Several studies have found that accounting for intra-consumer heterogeneity improves the model fit and predictions (Becker et al., 2018; Song et al., 2018; Ben-Akiva et al., 2019, etc.).

Becker et al. (2018) extended the HB estimator described above to estimate the logit mixture model with inter- and intra-consumer heterogeneity (referred to as the *logit double mixture model*). This estimator serves as the backbone of the online estimation methodology proposed in this dissertation; it enables discrete choice models to be estimated online, while accounting for both levels of heterogeneity.

1.3. OBJECTIVES AND CONTRIBUTIONS

This dissertation addresses important gaps associated with using discrete choice models in personalization. Once these limitations are addressed, these models can provide a holistic framework for online personalization that makes use of all the available data, and overcomes several shortcomings of traditional personalization methods.

1.3.1. Online Estimation

Although discrete choice models have been deployed in some recommender systems (Chaptini, 2005; Jiang et al., 2014), the applications were mostly offline because updating individual preferences requires re-estimating the entire model (which becomes computationally burdensome as the sample size, number of attributes, or number of alternatives increases). In these few applications, individual preferences were assumed to be "*static*", and only updated periodically. On the other hand, many online applications require updating preferences in real-time so that they can be used in the next recommendation.

Chapter 2 presents a methodology for estimating discrete choice models online, which extends the HB estimator of the logit double mixture model proposed by Becker et al. (2018) and Ben-

Akiva et al. (2019). In the proposed methodology, individual-specific parameters are updated after each choice using an efficient Bayesian update, without the need to re-estimate the entire model. Periodically, data from multiple individuals are pooled, and population level parameters are updated by re-estimating the model with the new data. This is achieved using two interacting and repeated steps:

- An **offline estimation** which updates the individual and population parameters using data from all users, and is run periodically (e.g. weekly).

- An **online estimation** which uses the data of the individual making the choice only to update his/her individual preferences, and is triggered after each choice.

The up-to-date individual preferences can then be used to generate personalized recommendations to users through an optimization model, e.g., assortment optimization.

This novel contribution allows for new applications of discrete choice models, such as real-time individual forecasting, personalized recommendations, and dynamic personalized discounts. In such applications, discrete choice models can be more effective than traditional personalization techniques for several reasons. First, they can account for complex patterns of heterogeneity by making use of the estimated individual specific parameters. Second, they can account for all the available data including user-specific, item-specific, and contextual information (while some traditional methods rely mainly on item and user profiling) (Chaptini, 2005). Third, they express utility as a function of attributes. This enables us to recommend new items that have not been chosen or rated before, and applies to cases where the alternative attributes change over time. Finally, discrete choice models can integrate relevance and diversity in a unified utility-based framework (Jiang et al., 2014).

The online estimation methodology is validated in Chapter 5 using the app-based travel advisor, *Tripod* (Sustainable Travel Incentives with Prediction, Optimization and Personalization) (Azevedo et al., 2018; Song et al., 2018). Tripod aims to incentivize and shift travelers' behavior towards more sustainable travel alternatives (e.g. by changing their mode, route, or departure time choice) by offering them real-time incentives and information. A smartphone-based SP survey is used to collect data, and the online estimator is implemented in the smartphone-based app afterwards.

1.3.2. Advanced Level of Heterogeneity

Some online applications of discrete choice models in recommender systems were based on multinomial or nested logit/probit models, which do not account for preference heterogeneity. Such models can only be used in non-personalized recommendations. On the other hand, logit mixture models (which account for heterogeneity) could not be estimated in real-time because of computational constraints (estimation requires integration over multidimensional distributions in MSL, or drawing from complex posteriors in HB methods). The few applications of these models in recommender systems only accounted for inter-consumer heterogeneity, and preferences were estimated and updated periodically. Some studies have calibrated choice models on the individual level (e.g. Teo et al., 2016), however, this method also has limitations since it requires either a sufficiently large number of observations per individual, or good priors on the individual-specific parameters.

The online methodology proposed in Chapter 2 accounts for inter- and intra-consumer heterogeneity, both of which improve the quality of predictions and recommendations (Ben-Akiva et al., 2019, Song et al., 2018). The distributions of individual- and choice-specific parameters (pertaining to intra-consumer taste variations) are learnt from repeated choices done by the same individual. The subsequent predictions or recommendations are generated using the estimated individual distributions rather than point estimates of the individual parameters.

While the model proposed by Becker et al. (2018) restricts the distributions of preferences to normal and log-normal distributions, Chapter 4 extends this model and the HB estimator to incorporate a latent class model, which allows for more flexibility in modeling heterogeneity. The resulting model and estimator build on previous work on logit mixture models with inter-consumer heterogeneity (Rossi et al., 2005; Bujosa et al., 2010; Greene and Hensher, 2013; Keane and Wasi, 2013; and Krueger et al., 2018), in which the distribution of preferences is represented as a finite mixture of normals.

1.3.3. Endogeneity

The proposed online methodology allows for new applications of discrete choice models in real-time adaptive contexts (such as recommendation systems). However, in such contexts, the alternatives presented to a user are determined by the users' previous choices, which may result in endogeneity bias. This topic has not been investigated yet in the context of choice-based recommender systems. Several studies analyzed adaptive stated preferences (ASP) surveys, which suffer from the same limitation, however, they were mostly based on Monte Carlo simulations rather than theoretical analyses. In addition, those studies do not have a consensus; some of them found significant bias associated with ASP surveys while others did not. Addressing this issue is essential for the application of the offline-online methodology, as estimates will be biased otherwise.

Chapter 3 builds on the analysis of Liu et al. (2007), who investigated endogeneity bias in adaptive metric utility models. It presents a theoretical explanation of endogeneity bias in adaptive choice contexts, and outlines the practical implications on estimating behavioral models.

1.3.4. Context-Awareness and Cold Start

Context-awareness and "cold start" are two common issues in various personalization systems (and not only choice-based systems). Discrete choice models can be easily extended to introduce context-awareness and mitigate the cold start problem, while maintaining their online capability.

Context-aware systems generate more relevant recommendations by adapting them to the specific contextual situation of the user (Adomavicius and Tuzhilin, 2011). For example, an online e-commerce platform might recommend clothes to customers depending on the season and on their location (Aggarwal, 2016). In discrete choice models, we can express preferences as a function of contextual information (e.g. time, location, social data, etc.) to generate context-aware recommendations.

The cold start problem refers to the case where the number of initially available ratings or choices is relatively small. For example, we can only generate personalized recommendations to a given user if we know his/her choice history. In discrete choice models, this problem can be mitigated by expressing the user's preferences as a function of his/her characteristics, which exploits similarities between different users in a systematic way.

These extensions are presented in Chapter 4, which proposes a logit mixture model that accounts for random and systematic inter- and intra-consumer heterogeneity. Systematic inter-consumer heterogeneity is expressed as a function of socio-economic variables, while systematic intra-consumer heterogeneity is expressed as a function of contextual variables. Despite their complexity, these models can still be applied in online settings.

1.4. DISSERTATION ORGANIZATION

The remainder of this dissertation is organized as follows. Chapter 2 presents the methodology for estimating and updating individual preferences online, in addition to a proof of concept using SP data. The data used in Chapter 2 are exogenously generated; the attributes are independent of the previous choices of each individual. However, in some applications (such as recommender systems), endogeneity is expected because the attributes in a given menu or choice situation are determined by the individual's previous choices.

Chapter 3 investigates endogeneity in the context of adaptive choice situations using theoretical analyses and Monte Carlo simulations. Different experiments are presented representing cases where endogeneity causes the estimates to be inconsistent, and other cases where endogeneity can be ignorable.

The methodology presented in Chapter 2 builds on the logit double mixture model, which imposes strong assumptions on the inter- and intra-consumer mixing distributions. Chapter 4 presents extensions that relax these assumptions and enrich the behavioral models. In the first extension, random and systematic inter- and intra-consumer heterogeneity are modeled using socio-demographic variables and contextual variables respectively. In the second extension, the logit mixture model with inter- and intra-consumer heterogeneity is enhanced by using a latent class model (which is an extension to the mixture of normals logit, also known as the "mixed-mixed logit model").

Chapter 5 presents an application of the online estimation methodology to Tripod. SP data are collected in the Greater Boston Area (GBA) using a smartphone-based, context-aware SP method. The data are then used to implement the online estimator on the smartphone application. Using the collected data, a large number of synthetic Tripod users is simulated, and the recommendation system is tested over several days.

Finally, Chapter 6 concludes the dissertation by discussing the main contributions and limitations and suggesting directions for future research.

CHAPTER 2: ONLINE DISCRETE CHOICE MODELS: APPLICATIONS IN PERSONALIZED RECOMMENDATIONS

2.1. INTRODUCTION

Personalization has gained increasing interest among researchers and practitioners in the past two decades. Greater ease in data collection about users has made it possible for service providers to recommend items, services, and content in a non-intrusive way (Aggarwal, 2016) through online recommendations. Some of the conventional recommendation techniques, which mainly rely on item and user profiling, produce ratings that do not take full advantage of the available data. On the other hand, discrete choice models, which have been rarely used in online recommendations, integrate item specific, user specific, and contextual data in a single model (Chaptini, 2005).

According to Jiang et al. (2014), the use of discrete choice models in recommender systems can address some limitations associated with the standard recommendation approaches. The first limitation is the tradeoff between relevancy and diversity (which reflect how similar the recommendations are to what has been previously chosen by the user, and how different they are with respect to each other respectively) (Ziegler et al. 2005; Jiang et al., 2014). The second limitation is that both metrics (relevancy and diversity), which are commonly used to measure the degree of matching, do not necessarily explain user preferences. On the other hand, discrete choice models directly measure the individual-specific utility of an alternative (or a set of alternatives) as a function of its attributes (without the need to measure relevancy or diversity separately). Finally, and unlike most standard recommendation techniques, discrete choice models can be applied even when the universal set from which alternatives are recommended and the alternative attributes vary over time. For example, in travel recommendations, the travel time, cost, and availability of the different alternatives might vary over different choice situations.

This chapter presents a methodology for estimating discrete choice models online, which can be used to update user preferences continuously in an app-based setting (such as recommender systems). The framework presented in this chapter extends the Hierarchical Bayes (HB) estimator proposed by Becker et al. (2018) and Ben-Akiva et al. (2019) which accounts for inter- and intra-consumer heterogeneity. This estimator does not only allow for estimating double mixture models more efficiently, but also enables these models to be applied online. An offline-online methodology is proposed in which individual-specific parameters are updated after each choice without the need to re-estimate the whole model. Periodically, data from multiple individuals are pooled, and population level parameters are updated by re-estimating the model with the new data.

In order to validate this methodology, stated preferences (SP) data on the choice of transport mode in Switzerland (Bierlaire et al., 2001) are used. Individual preferences are estimated and updated using repeated observations, and then used in predicting the next choice and generating personalized recommendations. While our applications focus on personalized recommendations, this methodology allows discrete choice models to be applied online in various real-time applications and decision support systems such as personalized advertisement, real-time forecasting, personalized price discounts, and others.

The remainder of this chapter is organized as follows: Section 2.2 presents an overview of online recommendations and recent applications of discrete choice models in this domain. Section 2.3 presents the proposed methodology for estimating and updating user preferences online. Section 2.4 presents an application of this methodology to SP data. Section 2.5 presents a Monte Carlo experiment illustrating why both offline and online estimation are needed. A discussion of the modeling approach and its applications in online recommendations is presented in Section 2.6, and Section 2.7 concludes.

2.2. BACKGROUND

2.2.1. Online Recommendations

The goal of online recommendations is to suggest items of interest to a user from a much larger set in order to handle information overload (Chaptini, 2005; Jiang et al., 2014; Ricci et al., 2015). Personalized recommender systems must deliver relevant and precise recommendations based on each user's tastes and preferences, which should be determined with minimal involvement from the user. Recommendations must also be delivered in real-time so users are able to act immediately (Chaptini, 2005).

According to Ansari et al. (2000), online recommendations can make use of several information sources including the individual's expressed preferences or choices among different alternatives, preferences for product attributes, other people's preferences or choices, expert judgments, and individual characteristics that may predict these preferences and choices.

Collaborative filtering and content-based filtering are the two most popular recommendation techniques. Other techniques include knowledge-based and context-aware methods (Ricci et al., 2015). Collaborative filtering (Goldberg et al., 1992) provides recommendations to an individual based on overlapping interests with other individuals. In other words, it mimics 'word-of-mouth' recommendations. Content-based techniques match the attributes of the user profile against the attributes of an item (Ricci et al., 2015). These techniques make recommendations similar to those a given user has liked in the past (Chaptini, 2005). Knowledge- or utility-based recommender systems base their recommendations on the computation of the utility of each item for the user (Huang, 2011). These systems utilize previous knowledge about users, items, and the utility function (Ricci et al., 2015). Context-aware recommender systems (CARS) account for contextual information such as the user's knowledge level (e.g. expert user or beginner), the time a recommendation is requested, and the external context (e.g. proximity of restaurants to the user) (Ricci et al., 2015). Other techniques have been proposed that utilize traditional machine learning techniques such as support vector machines and latent class models (Cheung et al., 2003) and multi-armed bandit methods (Li et al, 2010; Song, 2016).

Despite the significant advances in online recommendations, several theoretical and practical challenges have been identified. For example, Ziegler et al. (2005) showed that the commonly used top-N lists do not necessarily map user satisfaction and utility. In some cases, measuring the (expected) utility of recommendations may be more important than measuring the accuracy of recommendations (Ziegler et al. 2005; Gunawardana and Shani, 2015). Another major challenge is that the commonly used recommendation techniques are designed to consider different configurations as different items (Ricci et al., 2015). Therefore, very few of these techniques can be applied when the universal set from which items are recommended varies over time (as in the case of travel advisors, where the attributes of alternatives such as time and cost vary over time).

Using discrete choice models in personalized recommendations overcomes many of the limitations mentioned above. First, these models represent utility as a function of the attributes of items (or alternatives), and the individual preferences towards each of these attributes. Therefore, utility is not inferred from measures of similarity obtained from item or user profiling. Second, since utility is modeled as a function of attributes, this method is able to handle cases where new items (with known attributes) could be recommended (e.g. items that have not been chosen or rated before), and cases where the attributes vary over time. The researcher decides on the specification of the utility functions, which may include alternative attributes, individual preferences, contextual variables, and individual characteristics, thus making use of all the available data. Third, since the users' preferences are inferred from their previous choices, this reduces the burden on users because they are not required to rate or evaluate any items. Finally, this method is able to deal with diversification and sparsity using the extensions discussed in Section 2.6.

2.2.2. Econometric and Discrete Choice Models in Recommender Systems

Discrete choice models are often used to predict choices on an aggregate level. More recently these models have been utilized in recommender systems due to their ability to predict individual choices (Chaptini, 2005, Polydoropoulou and Lambrou, 2012; Jiang et al., 2014).

Chaptini (2005) utilized discrete choice models to predict choices on the individual level and provide personalized recommendations. He developed an online academic advisor for MIT students that recommends academic courses based on observed and latent attributes of the courses (e.g. difficulty, workload, overall impression, etc.). These attributes were expressed as functions of the students' characteristics (such as gender, degree program, etc.). The model was estimated using maximum likelihood estimation with data collected via an online revealed preferences (RP)/stated preferences (SP) survey. He then conditioned on the individual choices to estimate individual level parameters that were used in generating course recommendations. In this study, preferences were estimated offline for each student and not updated as more choices were observed. In addition, the behavioral model accounted for inter-consumer heterogeneity only.

Jiang et al. (2014) used discrete choice models to measure users' preferences towards an entire recommendation list. The goal was to identify a recommendation list with the highest choice probability. A multi-level nested multinomial logit model was proposed, and the recommendation problem was formulated as a nonlinear binary integer programming problem. The authors noted that unlike typical recommender systems, discrete choice models introduce product diversity in the proposed recommendations. The main limitation of this approach was the lack of personalization, since a nested logit model was used (this model can be estimated at the individual level only if a large number of choices per individual is available).

Rubin and Steyvers (2009) introduced a probabilistic model of the process by which an individual selects and later rates an item. This model was applied to movie rating data collected by Netflix. A Latent Dirichlet Allocation (LDA) model was used to model the probability of selecting a movie given a set of movies classified by topic. An ordered logit model was used to model movie ratings. This model included an individual-specific bias term which determines the general tendency of a user to give favorable ratings, however, it did not account for heterogeneity in other parameters (preferences). In addition, since all parameters are learnt through Markov-Chain Monte Carlo methods, this model can only be estimated offline.

Ansari et al. (2000) proposed a Hierarchical Bayes approach for a recommender system that accounts for systematic (observed) and random (unobserved) heterogeneity in user preferences, unobserved product heterogeneity and attributes (such as holistic consumer judgements and product appeal structures), and expert judgements. Customer ratings were modeled as a function of product attributes, customer characteristics, and expert evaluations. User preferences were expressed as a function of fixed effects (i.e. observed customer and movie variables and their interactions) and random effects pertaining to the customer. The model was applied to movie recommendations on the internet and estimated using MCMC. The main advantage of this paper was accounting for various sources of information (i.e. movie genres, expert evaluations, and socio-demographic characteristics). The authors suggested two main extensions to their framework: (1) learning preferences from implicit rather than explicit information (i.e. revealed preferences or actual choices), and (2) accounting for more complex forms of heterogeneity. This model also cannot be estimated online because of the excessive estimation time.

The methodology presented in this chapter extends the abovementioned studies by considering models that account for taste heterogeneity, and yet can be estimated online after each choice. This methodology uses only implicit data (observed choices) in order to estimate and update user preferences (i.e. observed choices).

2.2.3. User Heterogeneity and Personalization

According to Castells et al. (2015), user preferences are complex, dynamic, context-dependent, heterogeneous, and even contradictory. Therefore, accounting for consumer heterogeneity is critical in recommender systems. Most of the methods mentioned earlier account for inter-consumer heterogeneity. On the other hand, limited research has been done on intra-consumer heterogeneity, representing taste variation among different choices done by the same individual. For example, in travel recommendations, the same user might be more or less sensitive to travel time depending on various unobserved factors specific to the particular choice situation, such as his/her schedule, the trip purpose, weather conditions, etc.

According to Ben-Akiva et al. (2019), ignoring intra-consumer heterogeneity assumes a nearly neoclassical consumer with "permanent" individual preferences that do not change over time. Perturbations in these preferences are treated as nuisance factors. In the presence of multiple observations from each individual, it is possible to identify inter- and intra-consumer heterogeneity. In the context of discrete choice models, excluding intra-consumer heterogeneity when its effect is significant results in biases due to a greater degree of unobserved effects (Ben-Akiva et al. 2019).

Models with intra-consumer heterogeneity have been estimated using MSL (Bhat and Castelar, 2002; Bhat and Sardesai, 2006; Hess and Rose; 2009, Hess and Train; 2011; and Yáñez et al., 2011), and using maximum approximate composite marginal likelihood (MACML) (Bhat and Sidharthan, 2011). These studies investigated taste variations among different choices done by the same individual, and demonstrated that accounting for such effects results in better estimates. However, these studies were mainly exploratory and limited to offline applications. Becker et al. (2018) and Ben-Akiva et al. (2019) introduced a Hierarchical Bayes (HB) estimator for such models by extending the standard HB procedure for logit mixture (Allenby and Rossi, 1998; Train, 2009). This model significantly reduces the computation time compared to the previously used MSL estimators, and can be extended to online applications.

In the following sections, a novel framework is proposed to estimate discrete choice models and update individual level preferences in an online setting, building on the HB estimator proposed by Becker et al. (2018) and Ben-Akiva et al. (2019). This framework can be used in various applications, but is particularly useful in recommender systems. The estimated preferences account for both inter- and intra-consumer heterogeneity, and are updated in real-time after each choice. They can serve as input to an assortment optimization algorithm, which recommends personalized menus to users by maximizing an objective function (e.g. the probability of choosing an alternative from the menu, the expected revenue of the menu, etc.).

2.3. METHODOLOGY

This section explains the methodology for estimating and continuously updating population and individual level preferences. The Hierarchical Bayes estimator of the logit mixture model with inter- and intra-consumer heterogeneity proposed by Becker et al. (2018) is used in order to estimate these preferences. Inter- and intra-consumer heterogeneity are used to improve the estimation results, and thus the predictive capabilities of the choice models (Ben-Akiva et al., 2019). Individual-specific parameters, which can be extracted from the estimation procedure, are used for personalization.

2.3.1. Estimating Preferences

We consider the case whereby individual n ($n = 1, 2, \dots N$) is presented with a menu m ($m = 1, 2, \dots M_n$) and makes a choice among a set of alternatives ($j = 1, 2, \dots, J_{mn}$). Thus, each menu refers to a choice situation. The total number of individuals is N and the total number of menus presented to each individual is M_n.

In order to estimate user preferences, we use the HB estimator proposed by Becker et al. (2018), which extends the widely used 3-step HB estimator of logit mixture (Train, 2009) to a 5-step estimator in order to account for intra-consumer heterogeneity.

We assume the utility specification of alternative j in menu m presented in equation (2.1):

$$U_{jmn} = -P_{jmn} + X_{jmn}\eta_{mn} + \exp(\alpha_{mn})\,\epsilon_{jmn} \qquad (2.1)$$

Where U_{jmn} is individual n's unobserved utility of alternative j in menu m, P_{jmn} is the price of alternative *j* in menu *m* (with its coefficient fixed to -1), X_{jmn} represents a vector of individual characteristics and alternative attributes, η_{mn} represents a vector of parameters/preferences, α_{mn} is a scale parameter, and ϵ_{jmn} is an error term following the extreme value distribution EV(0, 1). The subscripts (mn) in η_{mn} and α_{mn} indicate that these parameters vary across individuals and across choice situations of the same individual respectively. The utility equation can also include interaction terms between the attributes and socio-economic characteristics (e.g. to model how high-income users have a higher willingness-to-pay for different attributes).

The model uses the money-metric utility specification as in Ben-Akiva et al. (2019), whereby the price coefficient is fixed to -1. This specification is equivalent to assuming that the scale is fixed and the price coefficient is distributed, but it is advantageous because all other coefficients represent the willingness-to-pay for the corresponding attributes (Train and Weeks, 2005). Since the price coefficient is fixed, the scale parameter α_{mn} can be estimated. The error term in equation (2.1) is $\exp(\alpha_{mn})\,\epsilon_{jmn}$, and its variance is given by $\exp(2\alpha_{mn})\frac{\pi^2}{6}$. However, we can divide all the elements of equation (2.1) by $\exp(\alpha_{mn})$ in order to obtain equation (2.2)[1], where the error term ϵ_{jmn} is distributed as EV(0,1) and its variance is $\frac{\pi^2}{6}$ as shown in equation 2.2:

$$U_{jmn} = \frac{1}{\exp(\alpha_{mn})}(-P_{jmn} + X_{jmn}\eta_{mn}) + \epsilon_{jmn} \qquad (2.2)$$

We start by defining three levels of parameters needed to account for both inter- and intra-consumer heterogeneity as proposed by Ben-Akiva et al. (2019):

1. Population distribution: the mean preferences (μ) in the population and the inter-consumer covariance matrix (Ω^b) respectively.

2. Individual distribution: the mean preferences (ζ_n) of a specific individual and the intra-consumer covariance matrix (Ω^w) respectively.

3. Menu-specific (choice-specific) preferences η_{mn}.

We assume that ζ_n and η_{mn} are normally distributed:

$$\eta_{mn} \sim \mathcal{N}_T(\zeta_n, \Omega^w) \qquad (2.3)$$

$$\zeta_n \sim \mathcal{N}_T(\mu, \Omega^b)$$

where T is the number of parameters. The probability of a sequence of choices (d_n) made by individual n can be expressed as:

$$P(d_n|\mu, \Omega^b, \Omega^w) = \int_{\zeta_n} \prod_{m=1}^{M_n} \left[\int_{\eta_{mn}} \prod_{j=1}^{J_{mn}} P_j(\eta_{mn})^{d_{jmn}} H(d\eta_{mn}|\zeta_n, \Omega^w)\right] F(d\zeta_n|\mu, \Omega^b) \qquad (2.4)$$

Where d_{jmn} is equal to one if individual n chooses alternative j in menu m and zero otherwise, and:

[1] The log-normally distributed scale parameter in the WTP space can be expressed as $1/\exp(\alpha_{mn})$ or $\exp(\alpha_{mn})$. Both formulations ensure that the menu specific parameter is always positive and result in similar model fit, however, the estimated population means have opposite signs. The latter formulation ($\exp(\alpha_{mn})$) assumes that the original error term in equation (2.1) is given by $\frac{1}{\exp(\alpha_{mn})}\epsilon_{jmn}$.

$$P_j(\eta_{mn}) = \frac{\exp(V_{jmn}(\eta_{mn}))}{\sum_{j'=1}^{J_{mn}} \exp(V_{j'mn}(\eta_{mn}))} \tag{2.5}$$

$$H(d\eta_{mn}|\zeta_n, \Omega^w) \sim \mathcal{N}_T(\zeta_n, \Omega^w) \tag{2.6}$$

$$F(d\zeta_n|\mu, \Omega^b) \sim \mathcal{N}_T(\mu, \Omega^b) \tag{2.7}$$

The Gibbs sampling procedure proposed by Becker et al. (2018) and Ben-Akiva et al. (2019) is used to estimate this model. Not only does this estimator offer superior computational performance compared to MSL estimators, but also it enables the online estimation of discrete choice models as shown in Section 2.3.2.

The posterior distribution is presented in equation (2.8):

$$K(\mu, \zeta_n\ \forall n, \eta_{mn}\ \forall mn, \Omega^b, \Omega^w|d) \propto \prod_{n=1}^{N}\left[\prod_{m=1}^{M_n}\left[\prod_{j=1}^{J_{mn}}\left[P_j(\eta_{mn})^{d_{jmn}}\right] h(\eta_{mn}|\zeta_n, \Omega^w)\right] f(\zeta_n|\mu, \Omega^b)\right] k(\Omega^w)k(\mu)k(\Omega^b) \tag{2.8}$$

Where:

$$k(\mu) \sim \mathcal{N}_T(\mu_0, A) \tag{2.9}$$

$$k(\Omega^b) \sim IW(T, I) \tag{2.10}$$

$$k(\Omega^w) \sim IW(T, I) \tag{2.11}$$

μ_0 represents a vector of prior means, A is a diagonal covariance matrix with diagonal values $\rightarrow \infty$ (uninformative prior), T is the number of unknown parameters in the utility equations, I is the T-dimensional identity matrix, and $IW(T, I)$ represent an Inverted Wishart distribution with T degrees of freedom and parameter I.

The model is estimated using the five-step Gibbs sampling procedure proposed by Ben-Akiva et al. (2019) and Becker et al. (2018). This procedure is explained below:

Step I: drawing from the conditional posterior of the population means:

$$K(\mu|\zeta_n\ \forall n, \eta_{mn}\ \forall mn, \Omega^w, \Omega^b) \propto f(\zeta_n \forall n|\mu, \Omega^b)k(\mu) \tag{2.12}$$

The conditional posterior on μ is $\mathcal{N}\left(\bar{\zeta}^{i-1}, \frac{\Omega^{b,i-1}}{N}\right)$ (where i is an iteration index) where:

$$\bar{\zeta} = \frac{1}{N}\sum_n \zeta_n^{i-1} \tag{2.13}$$

Step II: drawing from the conditional posterior of the inter-consumer covariance matrix:

$$K(\Omega^b|\mu, \zeta_n\ \forall n, \eta_{mn}\ \forall mn, \Omega^w) \propto f(\zeta_n \forall n|\mu, \Omega^b)k(\Omega^b) \tag{2.14}$$

The conditional posterior on Ω^b is Inverted Wishart with T+N degrees of freedom and parameter TI+N$\overline{V}_b$, where T is the number of unknown parameters in the utility equations, I is the T-dimensional identity matrix, and:

$$\overline{V}_b = \frac{1}{N}\sum_{n=1}^{N}(\zeta_n^{i-1} - \mu^i)(\zeta_n^{i-1} - \mu^i)' \tag{2.15}$$

Step III: drawing from the conditional posterior of the intra-consumer covariance matrix:

$$K(\Omega^w|\mu, \zeta_n\ \forall n, \eta_{mn}\ \forall mn, \Omega^b) \propto h(\eta_{mn}\ \forall mn|\zeta_n\ \forall n, \Omega^w)k(\Omega^w) \tag{2.16}$$

Given η_{mn}^{i-1} and ζ_n^{i-1} for all n, the conditional posterior on Ω^w is Inverted Wishart with degrees of freedom $T + M_t$ and parameter $\frac{TI+M_t\overline{V}_w}{T+M_t}$, where M_t represents the total number of menus faced by all individuals, and:

$$\overline{V}_w = \frac{1}{M_t}\sum_{n=1}^{N}\sum_{m=1}^{M_n}(\eta_{mn}^{i-1} - \zeta_n^{i-1})(\eta_{mn}^{i-1} - \zeta_n^{i-1})' \tag{2.17}$$

In this step, we assume a single covariance matrix for all individuals. Due to the potentially small number choice situations faced by each individual in a typical recommender system, it might not be possible to estimate an individual-specific covariance matrix.

Step IV: drawing from the conditional posterior of the individual level means:

$$K(\zeta_n|\mu, \eta_{mn}\ \forall mn, \Omega^b, \Omega^w) \propto h(\eta_{mn}\ \forall mn|\zeta_n\ \forall n, \Omega^w)f(\zeta_n|\mu, \Omega^b) \tag{2.18}$$

Using $N(\mu, \Omega^b)$ as a prior for ζ_n, the conditional posterior is $N(\overline{\zeta}_n, \Sigma_{\zeta n})$ where:

$$\overline{\zeta}_n = \left([\Omega^{b,i}]^{-1} + M_n[\Omega^{w,i}]^{-1}\right)^{-1}\left([\Omega^{b,i}]^{-1}\mu_i + M_n[\Omega^{w,i}]^{-1}\frac{1}{M_n}\sum_{m=1}^{M_n}\eta_{mn}^{i-1}\right) \tag{2.19}$$

And:

$$\Sigma_{\zeta n} = \left([\Omega^{b,i}]^{-1} + M_n[\Omega^{w,i}]^{-1}\right)^{-1} \tag{2.20}$$

Step V: drawing from the conditional posterior of the menu-specific parameters:

$$K(\eta_{mn}|\mu, \zeta_n, \Omega^b, \Omega^w) \propto \prod_{j=0}^{J_{mn}}[P_j(\eta_{mn})^{d_{jmn}}]\, h(\eta_{mn}\ \forall mn|\zeta_n, \Omega^w),\ n = 1, 2, \dots, N, m = 1, 2, \dots M_n \tag{2.21}$$

A draw of η_{mn}^i is obtained by the Metropolis-Hastings (MH) algorithm.

This five-step procedure assumes that all coefficients have inter- and intra-consumer heterogeneity. However, the estimator can account for coefficients with only inter-consumer heterogeneity, or coefficients without any heterogeneity by including two additional MH steps (Becker et al., 2018; Ben-Akiva et al., 2019).

The modified version of this Gibbs sampler uses the Hierarchical Inverted Wishart (HIW) prior[2] on the inter- and intra-consumer covariance matrices as recommended by Huang and Wand (2013), Akinc and Vandebroek (2018), and Song et al. (2019), who showed that the latter outperforms the commonly used IW priors in terms of noninformativity and parameter recovery.

2.3.2. Updating Preferences

These parameters are estimated and updated through two interacting and repeated steps: offline and online estimation.

Offline Estimation: The offline estimation updates all the population and individual level the parameters (μ, Ω^b, ζ_n, Ω^w , and η_{mn}). Data from multiple individuals are pooled and all the parameters are updated to reflect the effects of all choices made by all individuals since the last offline estimation. This is performed periodically (e.g. overnight or once a week) as it is computationally expensive. Updating population level parameters accounts for population trends, as shown in Section 2.6.

Online Estimation: The online estimation updates users' preferences in real-time as they make choices. The individual specific parameters (ζ_n and η_{mn}) are updated after each choice, assuming that the population parameters μ and Ω^b and the intra-consumer covariance matrix Ω^w are fixed. This update is computationally inexpensive, and it can be done for each individual at a time; when a given individual makes a choice, only his/her parameters are updated. The online procedure is executed by iterating steps IV and V of the 5-step Gibbs sampler only.

Ideally, if we ignore the computational constraints, the 5-step offline procedure would be used to update individual preferences after each choice. In this procedure, Steps IV and V update the individual- and menu-specific preferences for each individual using the intra-consumer covariance matrix $\Omega^{w,i}$ and the inter-consumer distribution $\mathcal{N}_T(\mu^i, \Omega^{b,i})$ as a prior as shown in equation (2.18). Conditional on the population level parameters (μ, Ω^b, and Ω^w), obtaining draws from ζ_n and η_{mn} for each individual is done independently from all other individuals. Therefore, if draws from the population level parameters were available, the individual- and menu-specific parameters could be updated separately for each individual by iterating steps IV and V. Since population level parameters are not expected to vary significantly between different offline estimations, these draws can be obtained from the last offline estimation. Section 2.4.2 illustrates that this method produces results that are very close to those of the full offline estimation.

Additionally, since an informative prior is used on the individual-specific parameters, the Markov Chains converge faster; stationarity is achieved quickly and a fewer number of draws is required in the online procedure (compared to the offline procedure). Finally, this procedure can also be implemented on the users' mobile phone in app-based settings.

[2] The HIW or Half-t distribution is defined by Huang and Wand (2013) as follows:
if $\sigma^2|a \sim Inverse-Gamma(\nu/2, \nu/a)$ and $a \sim Inverse-Gamma(1/2, 1/A^2)$, then $\sigma \sim Half-t(\nu, A)$
In the multivariate case: $\Sigma|a_1, \dots, a_p \sim Inverted\ Wishart\ (\nu + p - 1, 2\nu\ diag(1/a_1, \dots, 1/a_p))$;
$a_k \sim Inverse-Gamma\ (1/2, 1/A_k^2)\ , k = 1, \dots, p$ (where p is the dimensionality).

The key assumption in this procedure is that the population level preferences μ, Ω^{b}, and Ω^{w} do not vary significantly between successive offline estimations. The frequency of offline estimation depends on how fast the population level preferences change over time. This might vary from one application to another, and even between different parameters in the same application. This can be mitigated by observing the population level parameters obtained from successive offline estimations and deciding on the frequency of these estimations accordingly.

In addition, since μ and Ω^{b} are used as priors in Step IV, their effect diminishes as more observations per individual are observed, as the individual specific means ζ_n get closer to their true values. With few choice observations from each individual, the model suffers from "shrinkage", whereby individual level preferences are biased towards population means. However, if the number of observations per individual is large, deviations in the population level parameters from their true values will have a smaller effect on the individual level preferences.

2.3.3. Personalized Menu Generation

The online and offline estimations update individual and population level parameters respectively. These parameters are used as inputs to an online optimizer that performs menu optimization to present the user with a personalized list of alternatives to choose from. The system architecture is presented in Figure 1, which demonstrates how the online procedure uses the individual choices and the population level parameters obtained from the offline estimation (μ, Ω^{b}, and Ω^{w}) in order to update user preferences.

Personalized recommendations are generated using the menu optimization model proposed by Song et al. (2017, 2018). This model maximizes hit-rate or consumer-surplus (CS) in the form of a log-sum, subject to constraints specifying the maximum number of alternatives to be shown in a menu. Binary decision variables are defined for each alternative representing whether or not it is recommended. In the latter study, a Monte Carlo experiment representing a smart mobility service showed that models with intra-consumer heterogeneity provide better menus (i.e., higher hit-rates) compared to models with only inter-consumer heterogeneity.

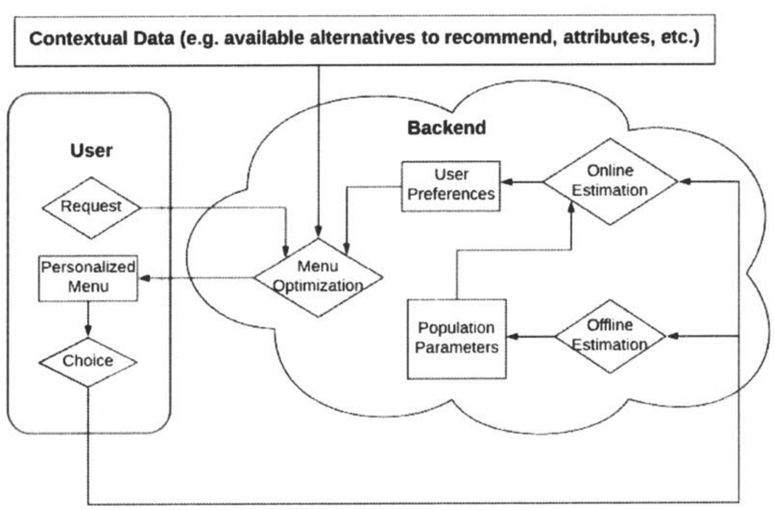

Figure 1: System architecture.

2.4. APPLICATION: SWISSMETRO DATA

2.4.1. Data and Model

The procedure described in Section 2.3 is applied to the Swissmetro dataset (Bierlaire et al., 2001), with the dependent variable being the transportation mode choice. The data was collected in Switzerland on the trains between St. Gallen and Geneva in 1998. Each survey respondent was presented with 9 hypothetical choice tasks, each having three alternatives (private car, Swissmetro (SM), and train). The attributes of these modes include the travel cost (fuel and parking costs for private car and fares for Swissmetro and train), travel time for all three modes, and Swissmetro and train headway. Since multiple observations are available from each respondent, we can use the offline-online procedure to demonstrate how preferences are learnt as more choices are observed.

In this application, we consider the simplified utility equations presented in equations (2.22-2.24)[3]. Since the cost coefficient is fixed to -1, the travel time coefficient $\exp(\eta_{mn})$ represents the value of time. Consequently, a scale parameter (α_{mn}) is estimated.

$$U_{Car,nm} = (ASC_{Car,mn} - \exp(\eta_{mn}) \times Time_{Car,mn} - Cost_{Car,mn})/\exp(\alpha_{mn}) + \epsilon_{Car,mn} \quad (2.22)$$

$$U_{SM,nm} = (ASC_{SM,mn} - \exp(\eta_{mn}) \times Time_{SM,mn} - Cost_{SM,mn})/\exp(\alpha_{mn}) + \epsilon_{SM,mn} \quad (2.23)$$

$$U_{Train,nm} = (\qquad - \exp(\eta_{mn}) \times Time_{Train,mn} - Cost_{Train,mn})/\exp(\alpha_{mn}) + \epsilon_{Train,mn} \quad (2.24)$$

Where:

- $U_{Car,mn}$, $U_{SM,mn}$, and $U_{Train,mn}$ represent the utilities of car, Swissmetro, and train in menu m for individual n, respectively.
- $Time_{j,mn}$ and $Cost_{j,mn}$ represent the total (door-to-door) travel time and travel cost of alternative j in menu m, respectively. The cost coefficient is fixed to -1.
- $ASC_{Car,mn}$ and $ASC_{SM,mn}$ represent alternative specific constants of car and train, respectively. The standard deviation of the train constant is normalized to zero.
- $\exp(\eta_{mn})$ and $\exp(\alpha_{mn})$ represent the coefficient of travel time and the scale parameter, respectively. Exponentiation is used in order model the log-normal distribution, which ensures that the travel time coefficient and the scale parameter are positive (and thus travel time and cost have a negative effect on utility to all individuals).
- $\epsilon_{Car,mn}$, $\epsilon_{SM,mn}$, and $\epsilon_{Train,mn}$ are error terms independently and identically distributed as Extreme Value type I.

[3] This simple specification is used only to validate the online estimation methodology. In Chapters 4 and 5, more elaborate choice models are presented which include socio-demographic variables and other interactions.

2.4.2. Results

In the following sections, the model with inter- and intra-consumer heterogeneity is estimated and the effect of personalization is analyzed. The base model is estimated using eight choices per individual, and the ninth choice is used for testing. Afterwards, models are estimated using fewer choices per individual (2, 5, or 7 choices out of 8), and then online estimation is applied to the remaining choices up to the eighth choice (and the ninth choice is also used for testing).

Estimation with Inter- and Intra-consumer Heterogeneity

The model is estimated using 400,000 Gibbs iterations, 200,000 of which are used as burn-in draws while the remaining 200,000 are used for sampling from the posterior distributions. The estimation results with menus 1-8 show significant inter-consumer heterogeneity in all coefficients. In addition, we find significant intra-consumer heterogeneity in the car and Swissmetro constants as shown in Table 1.

Table 1: Estimation results using the Swissmetro data.

	Population Mean		Inter-consumer Standard Deviation		Intra-consumer Standard deviation	
Parameter	Posterior Mean	Posterior Std. Dev.	Posterior Mean	Posterior Std. Dev.	Posterior Mean	Posterior Std. Dev.
ASC_{SM}	0.321	0.064	0.748	0.058	0.255	0.036
ASC_{Car}	0.574	0.071	1.286	0.050	0.091	0.045
Scale	-2.019	0.067	1.053	0.075	0.163	0.101
Travel time	0.179	0.037	0.912	0.029	0.024	0.016

The stationarity of the Markov chains is verified using the Heidelberg-Welch test (Heidelberger and Welch, 1983) and Gelman and Rubin's convergence diagnostic (Gelman and Rubin, 1992). All Markov chains pass the tests at the 95% level of confidence. All the convergence diagnostics are reported in Appendix A.

Predicting the Next Choice

In this section, similar models are estimated using fewer menus (e.g. 2, 5, or 7) and then the online procedure is applied to all individuals. Steps I through V (offline procedure) are iterated for menus 1-2, 1-5, or 1-7 and then steps IV and V are iterated for the remaining menus (online procedure).

This experiment mimics a scenario in which a few observations are initially observed from each individual. Individual and population level preferences are already estimated using these initial menus (menus 1-2, 1-5, or 1-7) by applying the five-step Gibbs sampler offline. Afterwards, new observations are available, and we need to update individual preferences to account for them using online estimation. The results of the offline-online estimation are compared with those of the full offline estimation as a benchmark.

In order to avoid overfitting, all the analyses are done using the test data (which include the ninth choice of each individual). The analyses are based on the mean of the posterior predictive distribution (PPD) given by equation (2.25) and the conditional log-likelihood of the estimated parameters.

$$P\left(d_{jm'n} = 1 \middle| d_n^*\right) = \int_{\eta_{m'n}} P_j(\eta_{m'n})\, K(\mathrm{d}\eta_{m'n} | d_n^*) \qquad (2.25)$$

where d_n^* denotes choices from recent menus and $K(\mathrm{d}\eta_{m'n}|d_n^*)$ is the posterior distribution of menu-specific parameters. The predicted probability of the chosen alternative is defined as the mean of the posterior predictive distribution across all individuals and draws. The conditional log-likelihood of the test data is calculated using individual-specific parameters and distributions, and therefore is conditioned on the choices made by individuals.

Since the five-step Gibbs sampler uses uninformative priors on μ, Ω^{b}, and Ω^{w}, the estimates and log-likelihood values obtained using this estimator are the same as those obtained using maximum simulated likelihood (MSL) since the posterior is dominated by the likelihood (Greene, 2004; Huber and Train, 2001; Ben-Akiva et al., 2019). Since uninformative priors are used, the conditional log-likelihood on the test data is used as a measure of performance.

Table 2 shows the log-likelihood and the predicted probability of the chosen alternative for the test menu (9th choice) using different estimation procedures. The results of the full offline estimation are presented in the last row (8 menus for each individual). This estimation achieves an average predicted probability of 0.717 and a log-likelihood of -400.9.

The results with a subset of the data (2, 5, and 7 choices per individual respectively) indicate that with fewer observations, the estimated models have lower average probabilities and log-likelihood values on the test data. However, the subsequent application of the online-procedure to the remaining menus recovers the drop in prediction accuracy as shown in the first two columns ("Double Mixture Personalized") in Table 2. This indicates that the online procedure produces results that are as accurate as those obtained from the full offline estimation.

Effect of Personalization

In order to test the effect of personalization, the results are compared to the unconditional log-likelihood and predicted probabilities. These are obtained by generating draws from the distributions $\zeta_n \sim \mathcal{N}\left(\hat{\mu}, \hat{\Omega}^{\mathrm{b}}\right)$ and $\eta_{mn} \sim \mathcal{N}(\zeta_n, \hat{\Omega}^{\mathrm{w}})$ respectively (where $\hat{\mu}, \hat{\Omega}^{\mathrm{b}}$, and $\hat{\Omega}^{\mathrm{w}}$ are the posterior means of μ, Ω^{b}, and Ω^{w}).

The unconditional predicted probability of choosing alternative j in the test menu m' of individual n can be calculated using equation (2.26).

$$\mathrm{P}\left(d_{jm'n} = 1 \middle| \hat{\mu}, \hat{\Omega}^{\mathrm{b}}, \hat{\Omega}^{\mathrm{w}}\right) = \int_{\zeta_n} \int_{\eta_{mn}} \mathrm{P}(d_{jm'n} = 1 | \eta_{mn}) \mathrm{H}\left(\mathrm{d}\eta_{mn} \middle| \zeta_n, \hat{\Omega}^{\mathrm{w}}\right) \mathrm{F}(\mathrm{d}\zeta_n | \hat{\mu}, \hat{\Omega}^{\mathrm{b}}) \qquad (2.26)$$

Table 2: Prediction results with the full offline, partial offline, and offline-online procedures.

	Double Mixture Personalized		Double Mixture		Flat Logit	
Estimation Procedure	Log-Likelihood	Probability	Log-Likelihood	Probability	Log-Likelihood	Probability
Partial Offline (1 – 2)	-550.5	0.668	-666.0	0.515	-689.2	0.505
Online (3 – 8)	-410.2	0.701	--	--	--	--
Partial Offline (1 – 5)	-436.7	0.699	-656.0	0.505	-650.9	0.492
Online (6 – 8)	-403.2	0.716	--	--	--	--
Partial Offline (1 – 7)	-414.8	0.711	-659.5	0.511	-649.7	0.496
Online (8)	-400.0	0.716	--	--	--	--
Full Offline (1 – 8)	-400.9	0.717	-656.9	0.510	-648.8	0.497

As shown in Table 2, the choice probabilities and log-likelihood values obtained from the flat logit model and the non-personalized double mixture model are inferior to those obtained using the posterior draws of the individual- and menu-specific parameters (personalized double mixture). In this example, using individual level parameters improves the predicted probabilities of the observed choices by 21% compared to non-personalized methods.

Generating Personalized Recommendations

In order to demonstrate the accuracy and robustness of the proposed method in personalized recommendations, it is compared to two different approaches: a simple content-based method (in which the most chosen alternative in the previous menus is recommended), and non-personalized discrete choice models (flat logit and double mixture model with inter- and intra-consumer heterogeneity). The first approach accounts for personalization by considering the choice history of each individual, however, it cannot account for the impact of changes in the attributes (travel cost and travel time) as these vary among different choices. On the other hand, the non-personalized choice models account for attributes, but do not make use of the choice history of each individual. The offline-online estimation methodology presented in Section 2.3 accounts for both the individual choice history and alternative attributes.

Personalized menu optimization is performed with the objective of maximizing the expected hit-rate (Song et al., 2017, 2018) on the 9th choice. Recommended menus are simulated with either one or two out of the three original alternatives. The hit-rate is defined as the fraction of individuals who choose an alternative that is included in the recommended menu. As shown in Table 3, the offline-online procedure can also approximate the full offline procedure in terms of hit-rate, and the observed effect of personalization is substantial.

Table 3 also indicates that the personalized double mixture model outperforms the content-based recommendation in all cases (by a margin of 1-2%), even when the online procedure is used. In addition, it is substantially better than the non-personalized flat logit and double mixture models.

Table 3: Hit-rates with the full offline, partial offline, offline-online, and content-based procedures.

	Content-Based (Most Chosen)		Flat Logit		Double Mixture		Double Mixture - Personalized	
Menu Size	1	2	1	2	1	2	1	2
Full Offline (1 – 8)	0.763	0.954	0.636	0.910	0.609	0.914	0.770	0.977
Partial Offline (1 – 2)	0.713	0.912	0.588	0.910	0.608	0.912	0.725	0.941
Partial Offline (1 – 5)	0.745	0.947	0.626	0.911	0.609	0.912	0.757	0.968
Online (3 – 8)	--	--	--	--	--	--	0.767	0.952
Online (6 – 8)	--	--	--	--	--	--	0.777	0.975

2.5. MONTE CARLO APPLICATION

After validating the online estimator using SP data, this section presents a Monte Carlo example demonstrating why intra-consumer heterogeneity, offline estimation, and online estimation are all needed.

2.5.1. The Data

The data mimics a route choice SP survey; we assume that 2,000 individuals are presented with successive choices between two alternatives: a faster/more expensive alternative, and a cheaper/slower alternative. The utility equation follows the money-metric specification, whereby a scale parameter is estimated and the cost parameter is fixed to -1. Therefore, the travel time parameter in this equation represents the value of time (VOT). The systematic utility equations of the two alternatives are given by equation (2.27).

$$V_{jmn} = \exp(\eta_{scale,mn})(-Cost_{jmn} - \exp(\eta_{time,mn}) \times Time_{jmn}) \quad j = 1,2 \qquad (2.27)$$

where $\eta_{mn} \sim N(\zeta_{mn}, \Omega_w)$ and $\zeta_{mn} \sim N(\mu_m, \Omega_b)$. We assume that there is a positive trend in the travel time parameter, indicating that preferences are changing over time:

$$\mu_m = 2 + 0.1 * m \qquad (2.28)$$

The inter-consumer variance of the travel time parameter is 0.5, and the intra-consumer variance is 0.25. The inter- and intra-consumer variances of the scale parameter are both 0.1, and the population mean is 0.5.

Ideally, we want to account for this trend in the model. However, this might not be possible because the trend might be difficult to detect in real applications. If we can detect the trend, we can model it explicitly using the random and systematic inter- and intra-consumer heterogeneity framework presented in Chapter 4. Otherwise, we would estimate a logit mixture model with inter- and intra-consumer heterogeneity as explained in Section 2.3.1, and perform offline estimation frequently in order to update the population parameters.

2.5.2. Static Predictions

In this application, two different models are estimated using 10 choices per individual, and used to forecast the 11th choice:

- **Model 1**: logit mixture with inter-consumer heterogeneity only.
- **Model 2**: logit mixture with inter- and intra-consumer heterogeneity.

The results are shown in Table 4. These results show that accounting for intra-consumer heterogeneity improves the predictions by about 1.3%. In addition, we can further improve our predictions by weighing the menu-specific parameters obtained from the model with intra-consumer heterogeneity differently. For example, we can forecast using the menu-specific parameters of the last 5 choices only, and improve the predictions by 1.1% (or 2.4% compared to the model with inter-consumer heterogeneity only).

Table 4: Monte Carlo results: predictions on the 11th choice.

	Model 1	Model 2	Model 2 (last 5 choices)
Predicted Probability	0.705	0.718	0.729
Log-Likelihood	-894.3	-868.4	-830.7

2.5.3. Dynamic Predictions

In the following application, we demonstrate the usefulness of both the offline and online estimation using the same Monte Carlo example. The goal is to predict the probability of the next choice given all the previous choices for $m = 1, 2, \ldots, 12$. Three different scenarios are compared:

1. **Offline-online estimation**: offline estimation is performed after menus 5 and 10, and online estimation is performed after each menu. The posterior draws of the last 5 choices are used in predictions.
2. **Online estimation only**: online estimation is applied after each choice using the population parameters estimated before the trend, $(m = 0)$. The population parameters are not updated at all throughout the simulation.
3. **Offline estimation only**: offline estimation is performed after menus 5 and 10, and individual- and menu-specific parameters are not updated in between.

The results of this simulation are shown in Figure 2. We can observe that the offline-online estimation performs consistently better than online estimation only. This is expected because the updated population parameters are used, which are closer to their true values (however, the estimated population mean of travel time will be lower than its true value because we are not accounting for the trend). On the other hand, applying offline estimation only results in a step function followed by a decreasing slope. The jumps indicate that our predictions improve after we update the population parameters, while the decreasing slopes are due to the fact that preferences change over time, but our estimates do not get updated to reflect this change.

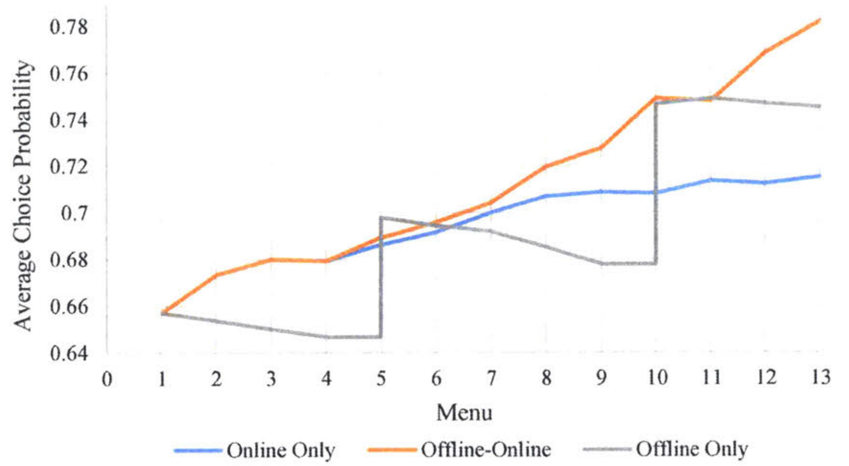

Figure 2: Comparison between offline-online estimation, offline estimation only, and online estimation only.

2.6. DISCUSSION

2.6.1. Model Estimation

This section presents some practical issues related to the effect of priors, identifiability, and applications in recommender systems.

Effect of priors

The basic HB procedure utilizes the Inverted Wishart (IW) prior, which has some undesirable properties, and thus can lead to biased estimates of standard deviations (Alvarez et al. 2014). Particularly, this prior tends to inflate standard deviations if their true values are small since it has a low density near zero. In order to overcome this issue, other priors can be used such as the Hierarchical Inverted Wishart (HIW), Scaled Inverted Wishart (SIW), and Separation Strategy (or BMM) (Huang and Wand, 2013; Akinc and Vandebroek, 2018; Song et al., 2019).

Identifiability of Individual level Preferences and Accounting for Uncertainty

The model with inter- and intra-consumer heterogeneity is only identifiable if multiple choice situations from each individual are available (unless informative priors are used). In addition, if we have a small number observations per individual, the model suffers from "shrinkage", whereby individual level preferences are biased towards their population means. HB is defined as a "data borrowing" technique that stabilizes individual level preferences for each individual using information not only from his/her past choices, but also from other individuals within the same dataset (Orme and Baker, 2000).

While Allenby and Rossi (1998) state that this procedure allows us to estimate the distributions of the population level parameters (μ and Ω^b) and yields exact finite-sample estimates of the posterior distribution of individual level parameters, Greene (2004) argues that these estimates are only "exact" for the assumed priors and the data used, and up to simulation variance. To account for uncertainty in individual level estimates, Allenby and Rossi (2006) indicate that these estimates are not precisely estimated, and the use of point-estimates leads to over-confident predictions of effect-sizes. To avoid this over-confidence, Allenby and Rossi (2006) suggest using all the posterior draws to make predictions instead of the point estimates (which is done in Sections 2.4 and 2.5).

Despite the fact that individual level preferences are not accurately estimated, the results show that we achieve significantly better predictions compared to those without any personalization. These preferences are "learnt" with more choices, which makes this procedure suitable for application in recommender systems.

In the case of sparse data, the posterior distributions of the individual- and menu-specific parameters can be used in Multi-armed bandit methods such as Thompson sampling and Upper Confidence Bounds as described in Teo et al. (2016) and Song (2018) to enhance learning via *exploration* and *exploitation*. For example, Thompson sampling uses individual draws from the posterior distributions rather than all the posterior draws. A distribution with large variance indicates high uncertainty in the estimated parameter. Therefore, attributes with uncertain parameter distributions become more likely to be recommended, which allows for learning these distributions more efficiently.

2.6.2. Application in Recommender Systems

Application to New Users

The online procedure can be applied to new users with observed choices using the estimates of μ, Ω^b and Ω^w obtained from the offline estimation with other users. For instance, these users might have joined the system and made choices after the last offline update. On the other hand, the procedure described in Section 2.4.2 can be applied to individuals with no previous choice history to generate unconditional predictions, and thus, these users will be first presented with non-personalized menus, i.e., population level parameters will be used for menu optimization. This is known as the "cold-start" problem (Chapter 4 addresses this issue by modeling random and systematic inter-consumer heterogeneity).

Data Collection and Endogeneity

In applications to recommender systems, the estimated models must account for endogeneity; the choice set presented to the user in each menu is based on this user's preferences, which are estimated based on his/her previous choices. Extensive research has been done on endogeneity corrections in discrete choice models, most of which falls into two categories: the BLP method (Berry, Levinsohn, and Pakes, 1995), and the control-function method (Heckman 1978; Hausman 1978).

Endogeneity is not a concern in the SP application presented in this chapter (since all the attributes used in the estimation of preferences were generated exogenously). However, this issue is investigated in Chapter 3, and practical implications are highlighted.

2.7. CONCLUSION AND FUTURE WORK

This chapter presented a methodology for estimating and updating consumer preferences online in the context of app-based recommender systems. I proposed an offline estimator, which estimates and updates individual and population level parameters periodically using a five-step Gibbs sampler, and an online estimator, which updates individual-specific parameters in real-time as more choices are made (assuming that the population parameters are fixed until the next offline estimation).

The proposed online estimator enables the use of discrete choice models in online decision support systems because it is (1) computationally efficient, (2) empirically accurate, and (3) theoretically justified. It is computationally efficient because it uses the data of the individual making the choice only, without the need to use data from other users. It is empirically accurate as it can achieve the same level of prediction accuracy as the offline estimator (which is computationally infeasible in real-time) as I have shown using real and Monte Carlo data. Finally, it is theoretically justified since it is equivalent to calibrating the model at the individual level, but with good priors representing the distribution of preferences in the population.

This methodology subsumes the utility-based advantages of discrete choice models and the personalization capabilities of standard recommendation techniques by making use of all the available data including user-specific characteristics and preferences, alternative-specific attributes, and contextual variables. Using the money-metric utility specification, the estimated distributions can be interpreted as the individual's "willingness-to-pay" for different features, which can be used in pricing, designing, and recommending new alternatives. In addition, the behavioral model is able to account for complex patterns of preference heterogeneity, namely intra-consumer heterogeneity which represents variations in preferences across different choices of the same individual. Therefore, we avoid the unrealistic assumption that preferences are static or stable over time. This has also been shown to improve the accuracy of recommendations and predictions (Song et al., 2018; Ben-Akiva et al., 2019).

Several limitations arise in the application of the proposed methodology. As in most Hierarchical models, individual- and menu-specific parameters might not be estimated precisely due to shrinkage (Liu et al., 2007). These preferences are "learnt" gradually from repeated choices.

So far, we have assumed that individual preferences and their corresponding distributions are constant over time. However, user preferences are dynamic and changing over time due to several factors. For example, an increase in a user's income might cause him/her to become less sensitive to travel cost and more sensitive to travel time. There are several approaches to deal with such cases. For example, we can base our recommendations on the user's most recent choices instead of his/her entire choice history (as in Section 2.5). We can also weigh more recent choices more than old choices both in estimation and in prediction. Future research must look into an optimal weighting model which would achieve the best predictions in terms of choice probabilities and menu hit-rates on real data.

Additionally, in recommender systems, the models estimated using the offline procedure might suffer from endogeneity bias, since personalized menus are generated based on the users' previous choices. Chapter 3 investigates this issue using theoretical and empirical approaches and outlines practical implications on such applications.

The underlying behavioral model with inter- and intra-consumer heterogeneity used in the offline-online estimation methodology imposes strong assumptions on the mixing distributions on both levels. First, only random (unobserved) sources of preference heterogeneity are considered. Second, the mixing distributions are restricted to normal (or log-normal) distributions, because of the desirable conjugacy property. Finally, the intra-consumer covariance matrix is assumed to be the same for all individuals. These assumptions are relaxed in Chapter 4, and more realistic behavioral models are proposed.

The framework presented in this Chapter is implemented in the app-based travel advisor Tripod (Sustainable Travel Incentives with Prediction, Optimization and Personalization) (Azevedo et al., 2018; Song et al., 2018) presented in Chapter 5. Tripod aims to incentivize and shift travelers' behavior towards more sustainable alternatives (e.g. changing mode, route, or departure time choice behavior), by offering them a personalized menu with incentives and information associated with each alternative.

CHAPTER 3: ENDOGENEITY IN ADAPTIVE CHOICE CONTEXTS

3.1. INTRODUCTION

Chapter 2 introduced a methodology for estimating and updating individual preferences online using discrete choice models. Preferences are learnt from repeated choices, and personalized menus are optimized based on the individual-specific preferences. However, when individuals are presented with personalized menus, the alternatives and attributes in each menu depend on the choices done by the same individual in the previous menus. Therefore, these attributes are endogenous. This chapter investigates potential endogeneity bias in these applications using both theoretical and empirical approaches, and presents practical implications on data collection and on estimating the behavioral models.

Endogeneity can arise in discrete choice models due to several factors including measurement errors, selection bias, omitted variables, and simultaneity, and results in inconsistent estimates of the model parameters (Guevara, 2015). The textbook definition of endogeneity is a correlation between the independent/observed variables in the model and the unobserved error term. A broader definition of endogeneity has been provided by Louviere et al. (2005), in which they defined "endogenous" as "all effects that are not exogenous". In the latter work, the authors attributed endogeneity to model misspecification.

In linear regression models, different methods have been proposed to address endogeneity, the most common of which are instrumental variables (IV) and Two-Stage Least Squares (TSLS). These methods use instruments that are correlated with the endogenous variables, but not with the error term. Other corrections include Heckman correction for selection bias (Heckman, 1977), and the estimation of simultaneous equation models.

In discrete choice models, several corrections have been proposed that mainly fall into two categories; the BLP method (Berry et al., 1995), and the control-function method (Hausman, 1978; Heckman, 1978). The BLP method (also known as the "product-market" control approach) suggests aggregating disaggregate consumer preferences obtained from discrete choice models into an aggregate market-level system, allowing for the application of standard instrumental variable methods. On the other hand, control-function methods use extra variables in the utility specification that are obtained using exogenous instruments. Different control-functions have been proposed, the most common of which are by Villas-Boas and Winer (1999), Blundell and Powell (2004), Guevara et al. (2006; 2012; 2016), and Petrin and Train (2010).

These methods are convenient when a few endogenous variables are identified, and relevant instruments are available. However, the cases considered in this chapter are of a different nature; in choice-based recommender systems, all of the attributes presented in a personalized menu depend on the choices done by the individual in the previous menus. A similar application is adaptive stated preferences (ASP) surveys. We refer to these two cases as adaptive choice contexts. In these contexts, endogeneity arises in all attributes, and instruments are generally difficult to obtain.

In this chapter, a theoretical analysis of endogeneity is presented, followed by a Monte Carlo experiment that mimics a choice-based recommender system that recommends Mobility as a Service (MaaS) plans. We demonstrate how endogeneity can cause inconsistent estimation results, and how this inconsistency can be avoided. The results show that when the system is initialized with exogenous attributes, and when all menus are included in the estimation, the estimates are consistent (i.e., the estimates are not significantly different from their true values). On the other hand, excluding data from the estimation leads to inconsistent estimates.

The remainder of this chapter is organized as follows. Section 3.2 presents an overview of endogeneity in adaptive contexts such as ASP surveys and adaptive metric utility experiments. Section 3.3 presents the methodology used in adaptive choice scenarios and a theoretical analysis of endogeneity. A Monte Carlo experiment mimicking a dynamic recommender system is presented in Section 3.4. Finally, Section 3.5 concludes. A similar application using a different recommender system and an example with adaptive linear regression are presented in Appendices B1 and B2.

3.2. BACKGROUND

This chapter primarily focuses on adaptive choice contexts whereby endogeneity might arise because the attributes of alternatives presented to a user are determined by the previous choices done by the same user. In such applications, all attributes are assumed to be endogenous. There are two common applications of this nature; choice based-recommender systems and adaptive stated preferences surveys. Endogeneity bias in recommender systems has not yet been studied in the literature. For example, Chaptini (2005) developed an online academic advisor for MIT students that recommends academic courses based on observed and latent attributes of the courses (such as difficulty, workload, overall impression, etc.). However, he calibrated the recommender systems exogenously, using revealed preferences (RP) and stated preferences (SP) data. Endogeneity has been studied in adaptive SP surveys (ASP) and adaptive metric utility experiments as discussed below.

3.2.1. Adaptive SP Surveys

In stated choice experiments, a respondent is presented with hypothetical alternatives and asked to choose one among those. Orthogonal designs have been commonly used to generate choice sets, either using a full factorial design, or a reduced version known as the fractional factorial design. More recently, efficient designs have been proposed to reduce the standard errors of the estimated parameters. However, these designs rely heavily on priors (Walker et al., 2018). To overcome this problem, Kanninen (2002) and Johnson et al. (2006; 2013) suggested updating the design during the data collection phase as knowledge of the true parameters increases. With increasing computational capabilities, adaptive stated preferences surveys (ASP) have been proposed in which the attributes presented to the users are determined based on the choices they have already made.

ASP surveys have been widely applied in marketing and transportation. Shinghal (1999) and Fowkes and Shinghal (2002) describe the development of the Leeds Adaptive Stated Preferences (LASP) survey, which was used by Shinghal and Fowkes (2002) to analyze freight mode choice. On the other hand, Toubia et al. (2004) developed a polyhedral method for survey design that "reduces the feasible set of parameters as rapidly as possible".

In other applications, Tilahun et al. (2007) used an ASP survey in order to evaluate individual preferences for different cycling environments in Minnesota. The authors argued that ASP surveys allow for measuring the exact values individuals attach to attributes of interest. Bolis and Maggi (2003) used an ASP survey to analyze freight transport demand in a logistics context. Similarly, Danielis et al. (2005) used an ASP to estimate logistics managers' preferences for freight service attributes.

Bradley and Daly (1993) analyzed several variations of ASP designs and concluded that endogenous SP designs might result in bias in the presence of taste variation in the sample. To counter this phenomenon, the authors suggested a few remedies such as (1) using market segmentation in order to ensure that each segment is as homogenous as possible, (2) using a Fixed/Adaptive approach (which uses a fixed basic design but avoids presenting certain choice options if they are deemed redundant based on prior choices), or (3) using an exogenous adaptive design in which an exogenous variable is used to adapt the design levels prior to the experiment.

Toubia et al. (2004) acknowledged the risk of endogeneity bias associated with the polyhedral method they proposed. Similarly, Abernethy et al. (2008), who proposed a similar method, but which is robust to response error, mentioned that such questionnaires are subject to endogeneity bias. However, they observed that bias decreases as more questions are observed.

On the other hand, Fowkes (2007) demonstrated that the LASP has no significant endogeneity bias since the models are calibrated at the level of the respondent. However, he also acknowledged that estimation over several respondents subjected to ASP could give rise to bias. In addition, similar to Abernethy et al. (2008), he concluded that the accuracy of an estimate improves and the apparent bias is reduced as more questions are added (due to alterations to the design), and concluded that LASP is "asymptotically unbiased". Similarly, Richardson (2002) developed a simulation study to estimate individual-specific values of time (VOT) using an ASP Survey and concluded that this method produces unbiased estimates of VOT.

Related work has been done on SP surveys constructed based on RP surveys by Train and Wilson (2008; 2009). One of the main findings was that in the presence of heterogeneity (e.g., random coefficients logit), the estimates based on combined RP and SP data are consistent, while estimates based on SP data alone are not. The authors attributed this inconsistency to the conditional distribution of coefficients, which differs over respondents and cannot be calculated without the RP data.

3.2.2. Adaptive Metric Utility Experiments

Although they are not choice based, adaptive metric utility conjoint experiments have the same endogeneity concerns as adaptive choice-based experiments (or ASP surveys). In this context, respondents are presented with paired-comparisons of two different items. The preference scale is metric, and questions are chosen based on previous responses to result in utilities that are nearly equal. Software packages such as Sawtooth (2003) use an adaptive conjoint analysis (ACA) method, which focuses on the attributes that are most relevant to the respondent, thus avoiding information overload. This method is referred to as "adaptive metric utility balance".

Hauser and Toubia (2005) argued that adaptive utility balance in metric conjoint analysis results in biases, inefficiencies, and higher response errors. The authors attributed this bias to endogeneity since new questions depend upon the errors made by respondents in their previous answers.

On the other hand, Liu et al. (2007) argued that in adaptive choice-based conjoint (CBC) experiments (such as the one implemented in Sawtooth software), endogeneity becomes ignorable for estimation once the data has been collected because of weak exogeneity and the "likelihood principle".

In linear regression models, weak exogeneity implies that the estimates might be biased, but are consistent. The likelihood principle states that "the likelihood function contains all the information in the data about the model parameters" (Fisher, 1922). This implies that the likelihood is meant to represent the true data generating mechanism. This has several practical implications on sequential sampling and pre-survey questions. According to Liu et al. (2007), in such cases one should simply condition on these questions instead of accounting for all the possible answers.

3.2.3. Summary of Background Literature

Based on the limited literature on adaptive SP surveys, it is undecided whether the estimation results are consistent in such applications or not. Research on discrete choice models applications in recommender systems is very limited, and thus endogeneity has not been discussed in these few applications.

The following sections build on the results by Liu et al. (2007), in order to demonstrate that the estimates are consistent in certain cases and inconsistent in others, depending on the data used and the likelihood function. A Monte Carlo simulation that replicates a recommender system context is then presented, however, the findings can be generalized to ASP surveys as well. The Monte Carlo results also show that biased estimates become closer to their true values as more questions are included in the estimation which is consistent with the literature on ASP surveys (Fowkes et al., 2007; Abernethy et al., 2008).

3.3. METHODOLOGY

In this section, adaptive choice contexts are analyzed using discrete choice models with random parameters, and theoretical explanations for consistent and inconsistent estimation strategies are provided.

In discrete choice models, researchers are primarily concerned with consistency rather than unbiasedness, as the latter cannot be established with standard estimators (e.g., Maximum Likelihood). Later in the Monte Carlo experiments (Section 3.4), we refer to "bias" as the difference between the true values of the model parameters and their corresponding estimates.

3.3.1. Model Specification

We consider a system (e.g., recommender system or ASP survey) in which an individual n is presented with a menu m with J_{mn} alternatives. The underlying behavioral model is a logit mixture with the linear utility specification shown in equation (3.1). We use the indices n for individuals ($n = 1,2, \dots, N$), m for menus ($m = 1,2, \dots M_n$), and j for alternatives ($j = 1,2, \dots, J_{mn}$).

$$U_{jmn} = \frac{1}{exp(\alpha_n)}(- P_{jmn} + X_{jmn}\beta_n) + \varepsilon_{jmn} \qquad (3.1)$$

U_{jmn} is individual n's unobserved utility of alternative j in menu m, P_{jmn} is the price associated with this alternative, X_{jmn} is a vector of alternative attributes, β_n is a vector of individual-specific coefficients/preferences, α_n is an individual-specific scale parameter, and ε_{jmn} is an error term following the extreme value distribution ($\text{EV}(0, 1)$). The model is in the willingness-to-pay (WTP) space; the price coefficient is fixed to -1 and a scale parameter is estimated. Therefore, all the parameters in β_n represent the willingness-to-pay for their corresponding attributes. This setting is used to be able to distinguish between the impact of endogeneity on the scale and on the ratio of the model coefficients (Guevara et al., 2012).

We define ζ_n as a vector of individual-specific parameters which includes both β_n and α_n. We assume that ζ_n is normally distributed in the population with mean μ and covariance matrix Ω.

$$\zeta_n \sim N(\mu, \Omega) \qquad (3.2)$$

We also define θ as the set of all model parameters (in this case, it includes μ, Ω, and ζ_n).

3.3.2. Menu Generation

In both ASP surveys and recommender systems, menus are generated using the previous choices and attributes. Without loss of generality, we assume that:

1. The system is initialized with one or more menus having exogenous attributes X_0 (which can also be pre-survey questions, or any exogenous screening questions). The choices/responses in these exogenous menus are denoted as d_0.

2. A known function (Q) is used for generating menu m with alternative attributes X_m based on the previous choices $(d_0, \dots, d_{m-1})$ and attributes $(X_0, \dots, X_{m-1})$ (where $Q(X_m | X_0, \dots, X_{m-1}, d_0, \dots, d_{m-1})$ represents the probability of menu X_m conditional on these choices and attributes).

In the case of recommender systems, this function is an assortment optimization method that generates a personalized menu of alternatives. A detailed example demonstrating this process is presented in Section 3.4. At each menu m, we have a universal set of alternatives to recommend from S_{mn}. This universal set is assumed to be exogenous; it is independent of any of the model parameters and the individual-specific parameters of individual n. The goal is to generate a recommended subset $S^*_{mn} \in S_{mn}$ with J_{mn} alternatives by maximizing a specific objective function given our estimates of the individual-specific parameters of individual n. Song et al. (2017; 2018) provide a methodology for generating recommendations with different objective functions (e.g., maximizing consumer surplus, hit-rate, etc.). In each menu m, individuals make the choice between any of the J_{mn} recommended alternatives and opting-out (i.e., not choosing any of the recommendations).

The recommended alternatives in menu m are generated as follows:

1. Individual-specific parameters are estimated using all menus up to $(m - 1)$ (e.g. using the HB procedure proposed by Train (2009)).

2. These parameters are fed into an assortment optimization to determine the J_{mn} alternatives to be included in the optimized menu (S^*_{mn}).

3. The choice among these J_{mn} alternatives and opting-out (i.e., not selecting any of the recommended alternatives) in menu m is observed.

4. Steps 1-3 are repeated after the new menu m is included in the next estimation.

In the case of ASP surveys, the menu generation function can be a deterministic (or even probabilistic) function which determines the attributes of menu m based on the previous choices. For example, a researcher might increase the price of an alternative if it was chosen before in order to see whether the subject will switch to another alternative or not.

3.3.3. Consistent Estimation

When a typical model estimation is done (e.g., through MSL or HB) using data from ASP surveys or recommender systems, the likelihood function that is being considered is presented in equation (3.3).

$$\mathrm{P}(d_0, \dots, d_m \,| X, \theta) = \mathrm{P}(d_0, d_1, \dots, d_m | X_0, X_1, \dots, X_m, \theta) \qquad (3.3)$$

Inconsistency can arise due to the misspecification of this likelihood function. Ideally, the likelihood function should include two components:

1. The probability of the observed choices conditional on the attributes $\mathrm{P}(d_m|X_m, \theta)$: This is conditional on the true model parameters θ because the choices are generated based on these parameters.

2. The probability of menu with attributes $\mathrm{X_m}$ conditional on the previous choices and attributes $\mathrm{Q}(X_m|d_1, \dots, d_{m-1}, X_1, \dots, X_{m-1}, \mathcal{S}_{mn})$: This is dependent on the true model parameters only through the choices.

In the case of recommender systems, the attributes are determined by the analysts' estimates of the model parameters $\hat{\theta}$, and not their true values. However, since $\hat{\theta}$ are estimated using the previous attributes and choices, the two expressions $\mathrm{Q}(X_m|\hat{\theta}, \mathcal{S}_{mn})$ and $\mathrm{Q}(X_m|d_1, \dots, d_{m-1}, X_1, \dots, X_{m-1}, \mathcal{S}_{mn})$ are equivalent.

The joint likelihood of the choices and attributes of the presented alternatives up to menu m is presented in equation (3.4).

$$\mathrm{P}(d_0, \dots, d_m, X_1, \dots, X_m|X_0, \theta) = \mathrm{P}(d_0|X_0, \theta)\mathrm{Q}(X_1|d_0, X_0, \mathcal{S}_1)\mathrm{P}(d_1|X_1, \theta) \dots$$

$$\dots \mathrm{Q}(X_m|d_0, \dots, d_{m-1}, X_0, \dots, X_{m-1}, \mathcal{S}_m)\mathrm{P}(d_m|X_m, \theta) \qquad (3.4)$$

Q can either be a deterministic or a non-deterministic function. If it is deterministic, then the outcome is either 0 or 1; it is 1 if the combination of $\{d_0, \dots, d_{m-1}, X_0, \dots, X_{m-1}, \mathcal{S}_m\}$ will result in presenting attributes X_m in the next menu and 0 otherwise. If it is non-deterministic, then the probabilities of presenting different attributes can be specified by the researcher.

Therefore, conditional on the previous choices, their attributes, and the universal set to recommend from $\mathcal{S}_{mn}$, the probability of the recommended alternatives in menu m is a constant. For example, in a recommender system with a deterministic function Q, given the previous choices and menus presented to a specific individual, and given the universal set of alternatives, the recommendation becomes deterministic:

$$\mathrm{Q}(X_1|d_0, X_0, \mathcal{S}_1) = 1$$

$$\dots$$

$$\mathrm{Q}(X_m|d_0, \dots, d_{m-1}, X_0, \dots, X_{m-1}, \mathcal{S}_m) = 1 \qquad (3.5)$$

Furthermore, if the universal sets of alternatives $\mathcal{S}_m$ are not known, but exogenous, even without conditioning on these universal sets, the menu generation probability is a constant and only depends on the distribution of attributes in the universal sets:

$$\mathrm{Q}(X_m|d_0, \dots, d_{m-1}, X_0, \dots, X_{m-1}) = \int_{\mathcal{S}_m} \mathrm{Q}(X_m|d_0, \dots, d_{m-1}, X_0, \dots, X_{m-1}, \mathcal{S}_m)\mathrm{f}(\mathcal{S}_m)d\mathcal{S}_m \qquad (3.6)$$

where $\mathrm{f}(\mathcal{S}_m)$ is the distribution of attributes in the universal sets. $\mathrm{Q}(X_m|d_0, \dots, d_{m-1}, X_0, \dots, X_{m-1}, \mathcal{S}_m)$ is equal to 1 for all universal sets that result in recommending X_m and 0 otherwise. Therefore, this probability is a constant and it is independent of the model parameters.

As a result, these constant terms can be removed from the likelihood without having any effect on the estimation results. Using the conditional independence of choices, we can conclude that the two likelihood functions presented in equations (3.3) and (3.4) are equivalent. In this case, the likelihood function fully explains the underlying data generation process, and therefore, endogeneity is ignorable.

3.3.4. Inconsistent Estimation

Inconsistent estimation results are obtained when the likelihood function does not reflect the data generation process. For example, if the exogenous menus are used in generating the first menu and ignored afterwards, the (misspecified) likelihood function that is being considered is:

$$\mathrm{P}\left(d_1, d_2, \ldots, d_m \mid X, \theta\right) = \mathrm{P}\left(d_1, d_2, \ldots, d_m \middle| X_1, X_2, \ldots, X_m, \theta\right) \qquad (3.7)$$

However, the correct likelihood function is presented in equation (3.8):

$$\mathrm{P}\left(d_1, \ldots d_m, X_1, \ldots, X_m \mid \theta\right) =$$

$$\mathrm{Q}(X_1 \mid \theta)\mathrm{P}(d_1 \mid X_1, \theta) \ldots \mathrm{Q}(X_m \mid d_1, \ldots, d_{m-1}, X_1, \ldots, X_{m-1})\mathrm{P}(d_m \mid X_m, \theta) \qquad (3.8)$$

$\mathrm{Q}(X_1 \mid \theta)$ is the marginal probability of observing menu X_1 (assuming that X_0 and d_0 are unknown). It can be obtained by integrating over the joint distribution of X_0, d_0, and X_1 given by equation (3.9), which is a function of θ because it depends on the choice probabilities $\mathrm{P}(d_0 \mid X_0, \theta)$, which are functions of θ.

$$\mathrm{f}(X_0, d_0, X_1 \mid \theta) = \mathrm{P}(d_0 \mid X_0, \theta)\mathrm{f}(X_0)\mathrm{Q}(X_1 \mid d_0, X_0) \qquad (3.9)$$

Where $\mathrm{f}(X_0)$ is the joint distribution of all the attributes in the exogenous menus, $\mathrm{P}(d_0 \mid X_0, \theta)$ is the probability of the choices given the exogenous attributes and model parameters, and $\mathrm{P}(X_1 \mid d_0, X_0)$ is the menu generation probability of the first menu.

$$\mathrm{Q}(X_1 \mid \theta) = \int_{X_0} \sum_{d_0} P(d_0 \mid X_0, \theta) f(X_0) Q(X_1 \mid d_0, X_0) dX_0 \qquad (3.10)$$

In equation (10), $\mathrm{Q}(X_1 \mid d_0, X_0)$ cannot be treated as a constant and dropped from the equation as before because the outcome is different for each combination of $\{d_0, X_0\}$. For example, if a deterministic function is used, this probability is either equal to zero or one, depending on the values of d_0 and X_0 (it is equal to one only for combinations of d_0 and X_0 that will result in recommending the first observed menu X_1).

Therefore, the probability of observing the first endogenous menu is a function of the model parameters θ, and excluding this expression from the likelihood function results in inconsistent estimates. Since the two likelihood functions are not equivalent, the one presented in equation

(3.7) does not account for the data generation process, and therefore, the estimates are inconsistent. In addition, the expression in equation (3.10) is difficult to evaluate because it requires multi-dimensional integration over all the possible attribute levels in the exogenous menus. The case presented above demonstrates an example of when inconsistency can arise, however, it is not the only one. Inconsistent estimates will be obtained if any of the following menus is excluded (if this menu was used in generating the subsequent menus).

3.4. MONTE CARLO EXPERIMENT

In this section, we simulate a dynamic recommender system following the methodology presented in Chapter 2 and in Song et al. (2017; 2018), mimicking the choice of Mobility as a Service (MaaS) plans. The experimental design is based on the Grapes Monte Carlo data in Ben-Akiva et al. (2019). Different model estimations are presented in order to demonstrate cases where the estimates are consistent and others where they are not, based on the previous analysis in Section 3.3.

3.4.1. Dataset Description

The Monte Carlo data assumes that 10,000 individuals are presented with eight successive menus. Each menu includes three alternatives (different plans of MaaS) with varying attributes and an opt-out alternative. Each plan has a different monthly price and three attributes: access to transit, access to bike sharing, and the number of on-demand trips (e.g., taxi, Uber/Lyft etc.) per month. Table 5 shows the distributions of the attributes in the universal sets (S_m). The dependent variable is the choice between the three different plans or not purchasing any of them (the opt-out alternative).

Table 5: Monte Carlo attributes and levels.

Attribute	Symbol	Levels
Monthly Price	P	\$100 to \$400
Transit	T	Available (1) or unavailable (0)
Bike Sharing	B	Available (1) or unavailable (0)
On-Demand Trips	OD	5 to 30 trips/month*

*On-demand trips are positively correlated with Price (r≈ **0.3**)

The utility equations (normalized to the opt-out alternative) are presented in equation (3.11):

$$U_{jmn} \equiv \frac{1}{exp(\alpha_n)}(-\ P_{jmn} + \beta_{T_n}T_{jmn} + \beta_{B_n}B_{jmn} + exp(\beta_{OD_n})OD_{jmn} + \beta_{q_n}) + \varepsilon_{jmn}$$

$$U_{opt-out,mn} \equiv 0 + \varepsilon_{opt-out,mn} \qquad (3.11)$$

where:

- n is an index for users ($n = 1,2, \dots, N$), m is an index for menus ($m = 1,2, \dots, M_n$), and j is an index for alternatives in the menu ($j = 1,2, \dots, J_{mn}$).

- U_{jmn} represents the utility of alternative j in menu m faced by individual n, and $U_{opt-out,mn}$ is the opt-out utility.

- P_{jmn} is the monthly price (in \$100's) of alternative j in menu m faced by individual n, with its coefficient normalized to -1.

- T_{jmn}, B_{jmn}, and OD_{jmn} represent access to transit, access to bike sharing, and the number of on-demand trips per month of alternative j with coefficients β_{T_n}, β_{B_n}, and $exp(\beta_{OD_n})$ respectively (exponentiation is used to obtain a log-normally distributed coefficient).

- β_{qn} is a constant term for choosing any plan (rather than opting out).

- α_n is a scale parameter (the log-normal distribution is also used to guarantee that the scale is positive for everyone).

- ε_{jmn} is an error component following the extreme value distribution (EV(0,1)).

All parameters are normally distributed in the sample. The true values of the population means and inter-consumer standard deviations are shown in Table 6. The true values of the individual-specific parameters are generated only once from their corresponding distributions, and then used in generating the choices in each menu. However, the attributes of the alternatives in the universal sets are different (both across different individuals and different menus).

The choices are simulated by estimating the systematic utilities (using the true individual-specific parameters and the attributes) and adding EV(0,1) error terms to these systematic utilities to obtain the total utilities. The alternative with the highest total utility is chosen.

Table 6: True values of the endogeneity simulation parameters.

Parameter	True mean	True standard deviation
α_n	-0.5	0.5
β_{T_n}	1.0	1.0
β_{B_n}	0.3	1.0
β_{OD_n}	-2.0	1.0
β_{q_n}	-1.0	1.0

3.4.2. Dynamic Recommendations

We assume that in each menu, three alternatives are recommended (and included in this menu) from a universal set (S_{mn}) of 100 alternatives. This universal set is different for each individual and for each menu, and the 100 alternatives in this set are generated by drawing the attributes from their corresponding distributions shown in Table 5.

In the first two menus, the recommendations are random; i.e., the three recommended alternatives are randomly chosen from the 100 alternatives in the universal set. Therefore, these two menus have exogenous attributes.

The following menus (3-8) are generated in a dynamic way based on each user's previous choices. The recommended alternatives in menu m are selected using the same recommendation procedure as in Section 3.3.2:

1. Estimating individual-specific parameters ($\hat{\zeta}_n$) using menus $1,2,\ldots,m-1$ (using the HB procedure in Train (2009)).

2. Calculating the systematic utilities of each alternative in the universal set using the estimated individual-specific parameters:

$$\hat{V}_{jmn} \equiv \frac{1}{exp(\hat{\alpha}_n)}\left(-\,P_{jmn} + \hat{\beta}_{S_n}S_{jmn} + \hat{\beta}_{C_n}C_{jmn} + exp(\hat{\beta}_{L_n})\,L_{jmn} + \hat{\beta}_{q_n}\right)$$

$$j = 1,2,\ldots,100 \qquad (3.12)$$

3. Selecting the three alternatives with the highest calculated utilities and including them in the recommended menus.

Since HB estimation is used, the systematic utilities are averaged over all the posterior draws of the individual-specific parameters.

In each menu, the choice among the three MaaS plans and opting-out is simulated using the true individual-specific parameters (ζ_n) (and not their estimates). The procedure is then repeated for the next menu ($m+1$).

3.4.3. Results

To estimate the models, the Hierarchical Bayes procedure of logit mixture (Allenby, 1997; Allenby and Rossi, 1998; Train, 2009) is used. The results presented in Table 7 show the estimates (posterior means) and standard errors (posterior standard deviations) of the population means and inter-consumer variances obtained with 200,000 Gibbs sampling draws (100,000 of which are burn-in draws and 100,000 are used for sampling). These results indicate that when all menus (up to menu m) are considered, there is no significant bias in any of the parameters at the 95% level of confidence (bias is defined as the difference between the estimates and their corresponding true values in the population).

Table 7: Estimation results with the first two menus included.

Population Means					
	Transit	Bike Sharing	On-Demand	Constant	Scale
True Value	1.000	0.300	-2.000	-1.000	-0.500
Value in Sample	0.992	0.319	-1.994	-0.996	-0.497
Estimation with m = 1, 2	0.946 (0.030)	0.270 (0.028)	-1.986 (0.022)	-0.958 (0.056)	-0.534 (0.040)
Estimation with m = 1, 2, 3	0.953 (0.026)	0.306 (0.020)	-1.987 (0.021)	-0.977 (0.052)	-0.517 (0.022)
Estimation with m = 1, 2 , 3 ,4	0.961 (0.024)	0.310 (0.017)	-1.983 (0.019)	-0.988 (0.046)	-0.522 (0.017)
Estimation with m = 1, 2,...,5	0.964 (0.023)	0.328 (0.017)	-1.991 (0.017)	-0.980 (0.040)	-0.522 (0.015)
Estimation with m = 1, 2,...,6	0.962 (0.021)	0.327 (0.016)	-1.994 (0.017)	-0.972 (0.039)	-0.519 (0.014)
Estimation with m = 1, 2,...,7	0.953 (0.022)	0.326 (0.015)	-1.999 (0.016)	-0.957 (0.038)	-0.522 (0.012)
Estimation with m = 1, 2,...,8	0.961 (0.021)	0.323 (0.014)	-1.996 (0.015)	-0.965 (0.033)	-0.515 (0.012)
Variances					
	Transit	Bike Sharing	On-Demand	Constant	Scale
True Value	1.000	1.000	1.000	1.000	0.250
Value in Sample	1.010	0.998	0.979	0.996	0.247
Estimation with m = 1, 2	1.004 (0.095)	1.033 (0.096)	1.001 (0.057)	0.803 (0.204)	0.279 (0.071)
Estimation with m = 1, 2, 3	0.999 (0.070)	1.005 (0.067)	0.994 (0.044)	0.857 (0.109)	0.241 (0.040)
Estimation with m = 1, 2 , 3 ,4	1.000 (0.058)	1.012 (0.057)	0.970 (0.036)	0.869 (0.112)	0.223 (0.028)
Estimation with m = 1, 2,...,5	0.967 (0.051)	0.956 (0.048)	0.964 (0.032)	0.948 (0.107)	0.224 (0.023)
Estimation with m = 1, 2,...,6	0.967 (0.050)	0.959 (0.045)	0.963 (0.031)	0.976 (0.097)	0.237 (0.020)
Estimation with m = 1, 2,...,7	0.951 (0.044)	0.948 (0.042)	0.970 (0.027)	0.935 (0.086)	0.234 (0.019)
Estimation with m = 1, 2,...,8	0.951 (0.043)	0.944 (0.038)	0.969 (0.026)	0.968 (0.081)	0.230 (0.015)

On the other hand, Table 8 presents similar estimation results, but the first two menus are excluded from the estimation and are only used in generating the first recommendation ($m = 3$). In this case, significant differences are observed in most of the parameters (at the 95% level of confidence). Such biases may have negative effects on transport policies. For example, underestimating the coefficients of transit in the MaaS example may result in designing plans that do not include this option as frequently. As a result, this may be less convenient to both MaaS users and operators.

Table 8: Estimation results with the first two menus excluded.

Population Means					
	Transit	Bike Sharing	On -Demand	Constant	Scale
True Value	1.000	0.300	-2.000	-1.000	-0.500
Value in Sample	0.992	0.319	-1.994	-0.996	-0.497
Estimation with m = 3	0.715 (0.066)	0.266 (0.035)	-2.303 (0.110)	7.223 (3.529)	-0.243 (0.079)
Estimation with m = 3, 4	0.917 (0.048)	0.327 (0.027)	-1.964 (0.048)	-0.224 (0.199)	-0.220 (0.047)
Estimation with m = 3, 4, 5	0.969 (0.043)	0.318 (0.021)	-1.991 (0.046)	-0.628 (0.130)	-0.351 (0.040)
Estimation with m = 3, 4,..., 6	0.960 (0.031)	0.316 (0.018)	-1.993 (0.033)	-0.857 (0.110)	-0.420 (0.023)
Estimation with m = 3, 4,..., 7	0.972 (0.029)	0.324 (0.017)	-1.970 (0.028)	-0.955 (0.086)	-0.423 (0.022)
Estimation with m = 3, 4,..., 8	0.968 (0.028)	0.330 (0.015)	-1.981 (0.024)	-0.962 (0.079)	-0.440 (0.019)
Variances					
	Transit	Bike Sharing	On-Demand	Constant	Scale
True Value	1.000	1.000	1.000	1.000	0.250
Value in Sample	1.010	0.998	0.979	0.996	0.247
Estimation with m = 3	1.583 (0.538)	1.386 (0.286)	1.427 (0.328)	77.530 (49.269)	0.584 (0.243)
Estimation with m = 3, 4	1.106 (0.169)	1.016 (0.129)	0.932 (0.122)	4.500 (0.997)	0.723 (0.123)
Estimation with m = 3, 4, 5	1.171 (0.129)	0.977 (0.081)	1.000 (0.090)	2.199 (0.392)	0.429 (0.086)
Estimation with m = 3, 4,..., 6	1.088 (0.082)	0.986 (0.060)	1.012 (0.065)	1.247 (0.198)	0.230 (0.043)
Estimation with m = 3, 4,..., 7	1.067 (0.069)	1.014 (0.055)	0.972 (0.051)	1.261 (0.207)	0.231 (0.039)
Estimation with m = 3, 4,..., 8	1.037 (0.063)	1.015 (0.049)	0.985 (0.045)	1.126 (0.169)	0.238 (0.030)

Similar to the findings of Fowkes (2007) and Abernethy et al. (2008), the bias decreases as more choices are observed. In this example, the largest bias is observed in the estimates and standard deviations of the constant and the scale parameter. However, in some cases, other parameters can be significantly different from their true values as well.

Finally, in the latter case, the observed bias decreases as more choices are included in the estimation because the effect of the omitted term ($Q(X_3|\theta)$) on the likelihood decreases. The attributes in any menu m after the third one are determined partially by the attributes and choices of menus $3, \ldots, m$. Therefore, the bias decreases since the misspecified likelihood function accounts partially for the data generation process. The results are also consistent with the findings of Abernethy et al. (2008), who concluded that LASP is asymptotically unbiased, and Toubia et al. (2004) who concluded that bias decreases as more questions are observed.

These results are also in accordance with the findings of Train and Wilson (2008), who concluded that when SP is experiments are based on the RP choices, estimation is only consistent when RP choices are included in the estimation. This case can be interpreted as a special case of the framework presented above. To demonstrate this, we denote by d_{RP} and X_{RP} the choices and attributes in the RP data, and d_{SP} and X_{SP} the choices and attributes in the SP data respectively. Assuming the RP attributes are exogenous and the function Q is used to generate SP attributes, the joint Likelihood of observing the SP attributes and the RP and SP choices is expressed as:

$$P(d_{RP}, X_{SP}, d_{SP}|X_{RP}, \theta) = P(d_{RP}|X_{RP}, \theta)Q(X_{SP}|d_{RP}, X_{RP})P(d_{SP}|X_{SP}, \theta) \tag{3.13}$$

In this case, $Q(X_{SP}|d_{RP}, X_{RP})$ is constant (independent of θ), and thus can be dropped. Therefore, joint estimation with SP and RP data ensures that the likelihood function is correctly specified.

However, if we do not include the RP data (d_{RP} and X_{RP}) in the estimation, we have to integrate over their distributions:

$$\begin{aligned} P(X_{SP}, d_{SP}|\theta) &= P(d_{SP}|X_{SP}, \theta)P(X_{SP}|\theta) \\ &= P(d_{SP}|X_{SP}, \theta) \int_{X_{RP}} \sum_{d_{RP}} P(d_{RP}|X_{RP}, \theta)f(X_{RP})Q(X_{SP}|d_{RP}, X_{RP})dX_{RP} \end{aligned} \tag{3.14}$$

Where $f(X_{RP})$ denotes the distribution of the RP attributes.

These results also hold for different menu generation procedures. Appendix B1 presents a similar Monte Carlo experiment, however, the three alternatives presented in the menu are the three nearest neighbors of the user's choice in the previous menu. We also observe significant biases when the first menu is excluded, and the magnitude of the bias also decreases as more choices are observed.

3.5. CONCLUSION

This chapter investigated endogeneity in adaptive choice contexts, such as choice-based recommender systems and adaptive stated preferences surveys. Traditional correction methods for endogeneity (e.g., control-functions) are difficult to apply in these cases since all attributes are endogenous and finding relevant instruments seems infeasible.

A Monte Carlo experiment was used to demonstrate cases where the estimation results are consistent and cases where they are not, in a choice-based recommender system in which alternatives are recommended based on the estimates of the individual-specific preferences. The results indicate that starting with exogenous menus and including all the data in the model estimation will result in consistent estimates. In the latter case, the estimates were not significantly different from their true values.

These findings have implications on the design and analysis of ASP surveys and choice-based recommender systems. Researchers and practitioners should make sure that the likelihood function accounts for the data generation process, which is achieved by conditioning on the previous choices and attributes. It is also important that the system is initialized exogenously, and that this initialization is accounted for in the estimation. For example, if the menus are generated based on pre-survey questions, the likelihood function should include the probability of the responses to these pre-survey questions.

In some cases, it might be infeasible to include the entire data in estimation due to computational limitations. Future research should address endogeneity bias in adaptive choice contexts without the need to use the entire data in estimation, either using traditional methods (such as the Wooldridge method and control-functions) or using other techniques.

Another important finding is that the magnitude of bias decreases as more choices per individual are observed. This was also observed in Nearest Neighbor recommendations (Appendix B1), and in other studies on ASP surveys (Abernethy et al., 2008; Fowkes, 2007). This indicates that when a sufficient number of choices per individual is included in estimation, endogeneity bias is not a concern even when only part of the data is used.

CHAPTER 4: EXTENSIONS TO THE BEHAVIORAL MODEL

4.1. INTRODUCTION

Chapter 2 introduced a methodology for estimating discrete choice models online, while accounting for random taste heterogeneity on two levels: inter- and intra-consumer heterogeneity. However, the underlying behavioral model (Becker et al., 2018; Ben-Akiva et al., 2019) relies on strong assumptions on the mixing distributions of preferences on both levels. This chapter relaxes some of these assumptions and proposes more realistic and better predictive behavioral models. The proposed extensions can also be integrated in the online estimation framework and applied in real-time personalization.

The first assumption is that taste heterogeneity is characterized by random inter- and intra-consumer mixing distributions (e.g. normal, log-normal, mixture of normals, etc.). This assumption is restrictive because it assumes that preference heterogeneity is completely random, while in fact it, part of it can be explained by many variables available to the researcher.

Preferences can be affected by several factors, including the individual's characteristics, attitudes, perceptions, and previous experiences, in addition to the context specific to the choice situation. For example, an individual's sensitivity to cost or price can be partially explained by the individual's income (people with higher income are expected to be less sensitive to cost). It might also be explained by other socio-demographic characteristics, such as age, gender, employment status, etc. Preferences might also be affected by context-specific variables that vary among choices. For example, in transportation mode choice, travel time sensitivity might be affected by weather conditions, trip purpose, and other contextual variables.

The second assumption is that both mixing distributions are assumed to be normal (or log-normal), i.e. the menu-specific parameters are normally distributed around the individual-specific means, which are in turn normally distributed around the population level means. Although this assumption facilitates the estimation of logit mixture models using the HB estimators, it might not be very realistic.

Finally, the behavioral model used in Chapter 2 assumes a single intra-consumer covariance matrix for all individuals. This assumption is also unrealistic, because different individuals might have different levels of variation in their preferences across choices.

This chapter introduces two extensions to the double mixture model which overcome the main limitations mentioned above. The first extension presents an estimator that accounts for both random and systematic preference heterogeneity by including individual characteristics and contextual variables. The second extension introduces a latent class extension, which results in more flexibility in modeling inter-consumer heterogeneity, and allows for different intra-consumer covariance matrices across users. Monte Carlo experiments are used to demonstrate these two extensions.

4.2. RANDOM AND SYSTEMATIC HETEROGENEITY

4.2.1. Motivation

This section presents a holistic framework for modeling complex dynamics in taste heterogeneity using advanced mixture models. The proposed framework accounts for the following elements:

1. Systematic inter-consumer heterogeneity: representing taste variations among different individuals that can be explained by the individuals' characteristics and socio-economic variables.
2. Systematic intra-consumer heterogeneity: representing taste variations among different choices done by the same individual, which can be explained by contextual variables specific to each choice.
3. Random inter- and intra-consumer heterogeneity: representing taste variations among different individuals and different choices of the same individual respectively, which are affected by factors that are unobserved to the modeler.

While models with systematic inter-consumer heterogeneity have been estimated before (Ansari et al., 2000; Bhat, 2000; Small et al., 2005; Horsky et al., 2006), as well as models with inter- and intra-consumer heterogeneity (Hess and Rose, 2009; Bhat and Sidharthan, 2011; Hess and Train, 2011; Yáñez et al., 2011, Becker et al., 2019; etc.), systematic intra-consumer heterogeneity has not yet been addressed in the literature. In addition, none of these studies has combined all sources of heterogeneity in a general framework. This section assembles all of the abovementioned elements in a single model, which results in better behavioral interpretability and more accurate predictions.

In addition, this section presents an HB estimator of the proposed model consisting of a Gibbs sampler with an embedded Metropolis-Hastings (MH) algorithm. Such models are usually infeasible to estimate using traditional methods such as MSL (Bhat and Sidharthan, 2011; Becker et al., 2019), as the dimensionality explodes quickly. The proposed HB estimator builds on the work of Becker et al. (2019), which was limited to random inter- and intra-consumer heterogeneity.

4.2.2. Background

Using socio-demographics to explain the parameters in mixed logit models (or the class probabilities in a latent class model) enables us to obtain insights into the likely location of an individual on the sample distribution (Hess, 2014). Therefore, several studies in the literature have modeled random and systematic sources of heterogeneity by extending the random coefficients model to incorporate observed individual characteristics.

Bhat (2000) formulated a multinomial-logit-based model of travel mode choice which includes both random and systematic taste variations. The model was used to represent mode preferences and responsiveness to level-of-service as a function of the individuals' socio-demographic characteristics. For estimation, Bhat used a simulated maximum likelihood procedure in which the multi-dimensional integral was evaluated using Monte Carlo techniques. A similar approach was deployed by Small et al. (2005) who modeled motorists' lane choice and sensitivity to travel time reliability. They used individual characteristics to model preferences and estimated the model using simulated maximum likelihood.

Similarly, Horsky et al. (2006) modeled individual level brand preferences by accounting for random and systematic sources of heterogeneity. A similar approach to Bhat (2000) was used in which preferences were modeled as functions of socio-demographic and household characteristics. However, unlike the aforementioned studies, (Bhat et al., 2000 and Small et al., 2005), this model was estimated using Bayesian Monte Carlo methods. The authors indicated that accounting for systematic heterogeneity in modeling preferences can more accurately reflect household cross-sectional differences and exploit the "panel" nature of the data, which does not only improve the model fit and prediction but also result in cleaner, more precise estimates of how intrinsic brand value, brand loyalty, and price influence choice. Both studies (Bhat, 2000 and Horsky et al., 2006) have found a significantly better fit of models that account for both random and systematic heterogeneity compared to models that do not account for both sources.

Ansari et al. (2000) also used a Bayesian estimation method in order to model random and systematic heterogeneity towards movie ratings. The authors applied their model to an internet-based recommendation system that recommends movies to users based on their characteristics and choice history.

While all of the abovementioned studies included random and systematic inter-consumer heterogeneity, none of them accounted for systematic intra-consumer heterogeneity. Inter-consumer heterogeneity is typically explained by individual characteristics and socio-demographics. On the other hand, intra-consumer heterogeneity can be explained by variables that vary among choices, such as contextual variables.

The following section presents an extension to these models and to the double mixture model proposed by Becker et al. (2018) and Ben-Akiva et al. (2019) in order to represent random and systematic inter- and intra-consumer heterogeneity. Socio-demographic characteristics are used to partially explain inter-consumer heterogeneity, and contextual variables are used to partially explain intra-consumer heterogeneity. Random taste heterogeneity is still present at both levels by specifying normally distributed error components. This extension allows for incorporating additional information into the model which results in better behavioral insights and more accurate predictions (Bhat, 2000; Horsky et al., 2006). The model is estimated using Gibbs sampling and applied to Monte Carlo data and to the Swissmetro dataset used in Chapter 2.

4.2.3. The Model

The proposed model is an innovative extension to the double mixture model with inter- and intra-consumer heterogeneity presented in Chapter 2. The main assumptions are that:

1. Menu specific parameters (η_{mn}) - which vary among choices, are affected by contextual variables that also vary from one choice to another.
2. Individual specific means (ζ_n) – which vary among individuals, but not across choices, are affected by socio-demographic characteristics that are assumed to be stable across choices.

The utility of alternative j in menu m is given by $V_{jmn}(\eta_{mn}, X_{jmn})$ where X_{jmn} is a vector of attributes of alternative j in menu m. Here we assume that the menu-specific scale parameter is included in η_{mn}.

We assume that η_{mn} is distributed according to the parametrized distribution:

$$\eta_{mn} \sim N(\zeta_n W_{mn}, \Omega^w) \tag{4.1}$$

Where ζ_n is a matrix of individual-specific parameters and W_{mn} is a vector of contextual variables. For simplicity, we assume that the intra-consumer covariance matrix Ω^w is diagonal. This assumption can be relaxed, but at the cost of significant computational costs. In addition, correlations among the different terms can be captured through the explanatory variables W_{mn}.

The menu-specific parameter of the k^{th} coefficient can be expressed as:

$$\eta_{mnk} = \zeta_{n0k} + \zeta_{nk} W_{mnk} + \omega_{mnk} \quad ; \qquad k = 1, \dots, K \tag{4.2}$$

Where ζ_{n0k} is the individual-specific intercept, ζ_{nk} is a vector of individual-specific parameters to the vector of contextual variables W_{mnk}, and $\omega_{mnk} \sim N(0, \sigma_k^{w2})$ where σ_k^{w2} is the k^{th} diagonal term of Ω^w.

Similarly, we assume that the vector of intercepts ζ_{n0} follows a parametrized distribution:

$$\zeta_{n0} \sim N(\mu S_n, \Omega^b) \tag{4.3}$$

Where μ is a matrix of population level parameters, S_n is a vector of individual characteristics and socio-demographic variables of individual n, and Ω^b is the inter-consumer covariance matrix. This matrix is also assumed to be diagonal. For simplicity, we assume that all the other individual-specific parameters in ζ_{nk} are normally distributed in the population as shown in equation (4.4) (without any covariates). This assumption can be easily relaxed in order to model the effects of socio-demographic variables on the sensitivities to contextual variables.

$$\zeta_{nk} \sim N(\mu'_k, \Omega_k^{b'}) \tag{4.4}$$

μ_k' and $\Omega_k^{b\prime}$ represent the population means of the individual-specific sensitivities to contextual variables, and the inter-consumer covariance matrix respectively. This distribution accounts for the fact that contextual variables might affect individuals differently (for example, in transportation mode choice, rainy weather can have a different effect on the walking time parameters of different individuals).

Equation (4.3) can be expressed as:

$$\zeta_{n0k} = \mu_{0k} + \mu_{1k}S_{n1} + \mu_{2k}S_{n2} + \cdots + \mu_{Lk}S_{nL} + \delta_{nk} \quad (4.5)$$

Where $\delta_{nk} \sim N(0, \sigma_k^{b^2})$, $\sigma_k^{b^2}$ is the k^{th} diagonal element of Ω^b, and L is the number of socio-demographic variables.

The resulting probability of individual n choosing alternative j in menu m can be expressed as:

$$\mathrm{P}(d_{jmn} = 1 | \mu, \mu_1', \dots, \mu_K', \Omega^b, \Omega_1^{b\prime}, \dots, \Omega_K^{b\prime}, \Omega^w) =$$

$$\int_{\zeta_n} \int_{\eta_{mn}} \mathrm{P_j}(\eta_{mn})\, \mathrm{H}(d\eta_{mn} | \zeta_n, \Omega^w; W_{mn})\, \mathrm{F}(d\zeta_n | \mu, \mu_1', \dots, \mu_K', \Omega^b, \Omega_1^{b\prime}, \dots, \Omega_K^{b\prime}; \mathrm{S_n}) \quad (4.6)$$

Where ζ_n represents the matrix of individual-specific parameters (including ζ_{n0k} and ζ_{nk} for each k).

This model can be estimated by modifying the 5-steps Gibbs Sampler as follows: In steps 1 and 4 of the Gibbs sampler, Bayesian linear regressions are used instead of Bayesian normal updates in order to draw from the distributions of ζ_{n0k} and μ_k. These regressions preserve the normal and Hierarchical Inverted Wishart conjugacy properties and are easy to implement.

The modified Gibbs sampler is summarized below:

Step 1a: drawing from the conditional posterior of population level parameters (μ), using a linear regression with diffuse (uninformative) priors, with known variance of the error term and using (ζ_{n0}, S_n) as the data.

Step 1b: drawing from the conditional posterior of the means of individual-specific sensitivities to contextual variables (μ_k') using multiple normal Bayesian updates with known variances and unknown means, and using ζ_{nk} as the data.

Step 2a: drawing from the conditional posterior of the inter-consumer variances of the individual-specific intercepts (Ω^b) using normal Bayesian updates with known means (draws from the fitted values of ζ_{0n}, obtained from Step 1a) and unknown variances, and using ζ_{0n} as the data.

Step 2b: drawing from the conditional posterior of the inter-consumer variances of the remaining individual-specific parameters ($\Omega_k^{b\prime}$) using normal Bayesian updates with known means (μ_k') and unknown variances, and using ζ_{nk} as the data.

Step 3: drawing from the conditional posterior of the intra-consumer covariance matrix (Ω^w) using a normal Bayesian update with known means and unknown variances, and using η_{mn} as the data. The means are the fitted values of η_{mn} obtained from Step 4.

Step 4: drawing from the conditional posterior of the individual-specific parameters ζ_{n0k} and ζ_{nk} for each individual n using a linear regression with η_{mn} as the data and informative priors. The mean of the prior on ζ_{n0k} is the fitted value from Step 1a ($\mu_k^i S_n$) and the mean of the prior on ζ_{nk} is the population mean obtained from Step 1b.

Step 5: drawing from the conditional posterior of the menu-specific parameters η_{mn} using the MH algorithm.

4.2.4. Monte Carlo Example

The Data

The estimator described above is applied to Monte Carlo data mimicking a route choice SP survey. We assume that 5,000 individuals are presented with 10 choice tasks, each including two alternatives (a faster/more expensive alternative, and a cheaper/slower alternative). The utility equation follows the money-metric specification, whereby a scale parameter is estimated and the cost parameter is fixed to -1. Therefore, the travel time parameter in this equation represents the value of time (VOT).

The systematic utility equations of the two alternatives are given by equation (4.7)

$$V_{jmn} = \exp(\eta_{scale,mn})(-Cost_{jmn} - \exp(\eta_{time,mn}) \times Time_{jmn}) \qquad j = 1,2 \tag{4.7}$$

For the sake of simplicity, we assume that the travel time parameter is partially explained by one socio-economic variable (employment status) and one contextual variable (day of the week). Following the above formulation, the menu-specific parameters are distributed as:

$$\eta_{scale,mn} = \zeta_{n0,scale} + \omega_{mn,scale} \tag{4.8}$$
$$\eta_{time,mn} = \zeta_{n0,time} + \zeta_{n,time} \times weekday_{mn} + \omega_{mn,time}$$

Where $weekday_{mn}$ is a dummy variable equal to 1 if the trip is on a weekday and 0 otherwise, and $\omega_{mn,scale}$ and $\omega_{mn.time}$ are independently and identically normally distributed with mean 0 and variance to be estimated.

Similarly, the inter-consumer distributions are given by:

$$\zeta_{n0,scale} = \mu_{0,scale} + \delta_{n,scale} \tag{4.9}$$
$$\zeta_{n0time} = \mu_{0,time} + \mu_{1,time} \times employed_n + \delta_{n,time}$$

$$\zeta_{n,time} = \mu'_{time} + \delta'_{n,time}$$

Where $employed_n$ is a binary indicator of being employed and $\delta_{n,scale}$, $\delta_{n,time}$, and $\delta'_{n,time}$ are independently and identically normally distributed with zero means and variances to be estimated. The true values of the model parameters are presented in Table 9. The underlying assumptions are that (1) employed people have a higher VOT, and (2) VOT is higher on weekday trips compared to weekend trips. The VOT distribution across individuals and choices has mean of 34.3 USD/hr and a median of 19.4 USD/hr.

Table 9: True values of the model parameters.

Parameter	True Value	Parameter	True Value
$\mu_{0,scale}$	0.00	$\mu_{0,time}$	2.00
$\mu_{1,time}$	1.00	μ'_{time}	1.00
$\mathrm{Var}(\delta_{n,time})$	0.25	$\mathrm{Var}(\delta_{n,scale})$	0.10
$\mathrm{Var}(\delta'_{n,time})$	0.25	$\mathrm{Var}(\omega_{mn,scale})$	0.10
$\mathrm{Var}(\omega_{mn,time})$	0.25		

Estimation Results

The model is estimated using R with 100,000 Gibbs sampling iterations, half of which were burn-in iterations and the other half were used for sampling from the posterior distributions. The estimation results are presented in Table 10, showing the true values, the estimates (posterior means), and standard errors in parentheses (posterior standard deviations).

Table 10: Estimation results.

Parameter	True Value	Estimate	Parameter	True Value	Estimate
$\mu_{0,scale}$	0.00	-0.049 (0.028)	$\mu_{0,time}$	2.00	2.021 (0.015)
$\mu_{1,time}$	1.00	0.976 (0.021)	μ'_{time}	1.00	0.975 (0.019)
$\mathrm{Var}(\delta_{n,time})$	0.25	0.245 (0.012)	$\mathrm{Var}(\delta_{n,scale})$	0.10	0.091 (0.014)
$\mathrm{Var}(\delta'_{n,time})$	0.25	0.235 (0.020)	$\mathrm{Var}(\omega_{mn,scale})$	0.10	0.081 (0.013)
$\mathrm{Var}(\omega_{mn,time})$	0.25	0.241 (0.017)			

The results indicate that all of the model parameters are not statistically different from their true values, which validates the estimator presented in Section 4.2.3.

4.2.5. Benchmarking Against Simpler Models

In order to demonstrate the benefits of accounting for systematic inter- and intra-consumer heterogeneity, we compare the proposed model to three simpler structures. Therefore, we have the following four models:

1. **Model 1**: the full model with systematic and random inter- and intra-consumer heterogeneity (which is the model proposed in Section 3.1, and which reflects the true data generation process).

2. **Model 2**: a similar model with systematic and random intra-consumer heterogeneity, but only random inter-consumer heterogeneity.
3. **Model 3**: a similar model with systematic and random inter-consumer heterogeneity, but only random intra-consumer heterogeneity.
4. **Model 4**: the logit mixture model proposed by Becker et al. (2018) with only random inter- and intra-consumer heterogeneity.

The four models are compared using two metrics: the mean of the posterior predictive distribution (PPD), representing the average predicted probability of the chosen alternative, and the log-likelihood. For both metrics, the conditional and unconditional values are calculated using two hold-out samples. The first sample consists of the same 5,000 users used in estimation (whose user-specific parameters are available from the estimated posterior distributions). The second sample consists of 5,000 new users (i.e. users with no choice history) whose preferences follow the same distributions as the existing users. Each sample includes two choices per user (a weekday trip and a weekend trip).

The conditional values of the two metrics are obtained using the estimated posterior distributions of the individual- and choice-specific parameters. These are conditional on the previous choices observed by each user. Therefore, the first sample is used (*"existing users"*). The unconditional values of the two metrics are obtained using the posterior distributions of the estimated population parameters. These are not conditional on the previous choices, and therefore, the second sample is used (*"new users"*).

The results of this comparison are shown in Table 11, which shows the mean and confidence interval of the posterior predictive distribution (PPD) on the two samples, as well as the log-likelihood of the two samples. The conditional values of PPD and log-likelihood are higher than the unconditional values as expected, indicating that we could learn the preferences of existing users.

The results indicate that Model 1 (with both systematic inter- and intra-consumer heterogeneity) outperforms Models 2 and 3 for the sample of new users. Both models 2 and 3 (with either systematic inter- or systematic intra-consumer heterogeneity) outperform Model 4, which only has random taste heterogeneity.

On the other hand, Models 1 and 2 outperform Models 3 and 4 for the sample of existing users. These results show that accounting for systematic inter-consumer heterogeneity could improve predictions significantly for new users, but only marginally for users with previous choice history. This is expected because after observing 10 choices, we do not expect any gains from modeling the user-specific parameters as a function of user characteristics; we simply need to condition on the previous choices. On the other hand, accounting for systematic intra-consumer heterogeneity improves the predictions for users with and without previous choice history.

Table 11: Comparison between the double mixture model and models with random and systematic heterogeneity.

	Model 1	Model 2	Model 3	Model 4
Unconditional PPD	0.725 [0.720, 0.729]	0.691 [0.686, 0.696]	0.696 [0.691, 0.700]	0.665 [0.659, 0.670]
Unconditional log-likelihood	-4265.4	-4718.8	-4703.9	-5149.1
Conditional PPD	0.751 [0.747, 0.754]	0.746 [0.743, 0.750]	0.712 [0.708, 0.716]	0.708 [0.704, 0.712]
Conditional log-likelihood	-3975.3	-4041.9	-4483.2	-4547.2

In this example, accounting for random intra-consumer heterogeneity results in small improvements in predictions for both samples. This is shown in Table 12 below, showing the predictions of Multinomial logit and two models with only inter-consumer heterogeneity: the first includes random and systematic inter-consumer heterogeneity, and the second only includes random inter-consumer heterogeneity.

We observe from this table that Model 3 (with systematic and random inter-consumer heterogeneity, and only random intra-consumer heterogeneity) performs slightly better than the model with only random and systematic inter-consumer heterogeneity and no intra-consumer heterogeneity shown below. Similarly, Model 4 above performs slightly better than the model with only random inter-consumer heterogeneity.

Table 12: Similar results for models with only inter-consumer heterogeneity.

	Radom and Systematic Inter-Consumer Heterogeneity	Radom Inter-Consumer Heterogeneity	Flat Logit
Unconditional PPD	0.690 [0.687, 0.693]	0.660 [0.656, 0.664]	0.654 [0.653, 0.655]
Unconditional log-likelihood	-4727.8	-5177.5	-5283.3
Conditional PPD	0.707 [0.705, 0.709]	0.703 [0.701, 0.705]	-
Conditional log-likelihood	-4568.9	-4618.8	-

From the above simulations, we can conclude the following:

1. Accounting for random inter- and intra-consumer heterogeneity improves the unconditional predictions slightly.
2. Accounting for systematic inter- and intra-consumer heterogeneity improves the unconditional predictions (i.e. for users without previous choice history) significantly.
3. Accounting for random inter-consumer heterogeneity improves the conditional predictions (i.e. for users with previous choice history) significantly.

4. Accounting for systematic intra-consumer heterogeneity improves the conditional predictions significantly, but systematic inter-consumer heterogeneity has a smaller effect.

While the above conclusions are in line with our apriori expectations, the effect of random and systematic inter- and intra-consumer heterogeneity also depends on several factors, including the number of choice situations per user, the levels of inter- and intra-consumer heterogeneity, and the degree to which contextual variables and individual characteristics affect both the inter- and intra-consumer distributions. For example, Section 2.5 showed that random intra-consumer heterogeneity improves predictions significantly when there is a positive trend in preferences.

4.2.6. Benchmarking Against Traditional Model Specifications

In typical model specifications, individual characteristics and contextual variables are specified linearly in the utility equations, without being interacted with the taste parameters, and without modeling systematic inter- or intra-consumer heterogeneity.

In this section, we compare Model 1 (with random and systematic inter- and intra-consumer heterogeneity) to the logit mixture model with only random inter- and intra-consumer heterogeneity, and with the dummy variables of "employment" and "weekday" included linearly in the utility equation of the slower alternative (i.e., similar to the common specifications of choice models). The results are presented in Table 13.

Table 13: Comparison with the traditional model specification.

	Random and systematic inter- and heterogeneity	Traditional specification with inter- and intra-consumer heterogeneity.
Unconditional PPD	0.725 [0.720, 0.729]	0.710 [0.707, 0.713]
Unconditional log-likelihood	-4265.4	-4492.2
Conditional PPD	0.751 [0.747, 0.754]	0.730 [0.728, 0.733]
Conditional log-likelihood	-3975.3	-4178.1

As expected, the correctly specified model (with random and systematic heterogeneity) outperforms the misspecified model where the explanatory variables are included linearly in the utility equations. However, in this case, the latter model performs better than models with only random taste heterogeneity (e.g. Model 4 in Table 11), because it accounts for these explanatory variables (even though it is misspecified).

4.2.7. Real Application: Swissmetro Data

The model described above is also applied to the Swissmetro dataset (used in Chapter 2). The utility equations are specified as before. The money-metric utility is used in which a scale parameter is estimated and the cost coefficient is fixed to -1. The constants for car and Swissmetro enter the utility equations, while the train constant is normalized to 0.

Exponentiation is used for the travel time coefficient and the scale parameter in order to ensure that the effects of cost and time on the utility are negative. The systematic utility equations are presented in equation (4.10).

$$V_{Car,\,nm} = (ASC_{Car,mn} - exp(\beta_{mn}) \times Time_{Car,mn} - Cost_{Car,mn}) \times exp(\alpha_{mn}) \quad (4.10)$$

$$V_{SM,\,nm} = (ASC_{SM,mn} - exp(\beta_{mn}) \times Time_{SM,mn} - Cost_{SM,mn}) \times exp(\alpha_{mn})$$

$$V_{Train,nm} = (\qquad - exp(\beta_{mn}) \times Time_{Train,mn} - Cost_{Train,mn}) \times exp(\alpha_{mn})$$

The individual-specific means (ζ_n) for the two constants and the travel time parameter are specified to be functions of the socio-economic variables available in the dataset: gender and income, as shown in equation (4.11).

$$\zeta_{n,Car} = \mu_{0,Car} + \mu_{1,Car} * Male_n + \mu_{2,Car} * LowIncome_n + \mu_{3,Car} * HighIncome_n + \omega_{n,Car}$$

$$\zeta_{n,SM} = \mu_{0,SM} + \mu_{1,SM} * Male_n + \mu_{2,SM} * LowIncome_n + \mu_{3,SM} * HighIncome_n + \omega_{n,SM}$$

$$\zeta_{n,Time} = \mu_{0,Time} + \mu_{1,Time} * Male_n + \mu_{2,Time} * LowIncome_n + \mu_{3,Time} * HighIncome_n + \omega_{n,Time} \quad (4.11)$$

where $Male_n$ is a binary indicator of being a male, $LowIncome_n$ and $HighIncome_n$ are indicators of the low and high income categories in the Swissmetro dataset, and $\omega_{n,Car}, \omega_{n,SM},$ and $\omega_{n,Time}$ are normally distributed error terms.

Since this dataset does not include contextual variables that vary among choices, only random (unobserved) intra-consumer heterogeneity is modeled (using normal distributions as before). Therefore, the menu-specific coefficients in equation (4.10) ($\mathrm{ASC}_{\mathrm{Car,mn}}$, $\mathrm{ASC}_{\mathrm{SM,mn}}$, and β_{mn}) concatenated as η_{mn} are distributed as $\eta_{mn} \sim N(\zeta_n, \Omega^{\mathrm{w}})$, where ζ_n is a concatenation of $\zeta_{n,Car}$, $\zeta_{n,SM}$, and $\zeta_{n,Time}$ and Ω^{w} is the inter-consumer covariance matrix.

However, systematic intra-consumer heterogeneity in the scale parameter is modeled as a polynomial function of the menu number as shown in equation (4.12).

$$\alpha_{mn} = \zeta_{n,0} + \zeta_{n,1} \times m + \zeta_{n,2} \times m^2 + \omega_{mn,scale} \qquad (4.12)$$

Where $\zeta_{n,0}$, $\zeta_{n,1}$, and $\zeta_{n,2}$ are normally distributed in the population, and $\omega_{mn,scale}$ is a normally distributed error term representing random intra-consumer heterogeneity in the scale parameter.

The estimation results are presented in Tables 14 and 15 showing random and systematic heterogeneity in the constants and the travel time parameter, and in the scale parameter respectively. The standard deviations of the posterior draws are presented in parentheses.

Table 14 : Random and systematic heterogeneity in the constants and travel time parameter.

	ASC_{Car}	ASC_{SM}	β_{time}
Intercept	-0.469 (0.152)	0.062 (0.108)	-0.801 (0.206)
Male	0.701 (0.155)	0.045 (0.106)	0.474 (0.187)
Low Income	-0.131 (0.151)	-0.417 (0.108)	-0.054 (0.149)
High Income	0.145 (0.132)	-0.046 (0.101)	0.395 (0.122)
Inter-Consumer Std. Dev.	1.031 (0.017)	0.304 (0.021)	0.608 (0.014)
Intra-Consumer Std. Dev.	0.133 (0.077)	0.060 (0.093)	0.083 (0.059)

The estimation results indicate that respondents with higher income are more sensitive to travel time as expected, and have a slightly higher preference towards car. In addition, males have a higher preference towards car, and a higher sensitivity to travel time. Respondents with a lower income have a lower preference to SM compared to other categories.

The values presented in Table 14 represent the corresponding parameters in Equation 4.11. For example, we can express the individual-specific parameter of travel time as follows:

$$\zeta_{n,Time} = -0.801 + 0.474 * Male_n - 0.054 * LowIncome_n + 0.395 * HighIncome_n + \omega_{n,Time} \tag{4.13}$$

Where $\omega_{n,Time} \sim N(0, 0.608)$. Having an expression of $\zeta_{n,Time}$, we can represent the menu-specific travel time parameter shown in equation (4.10) as $\beta_{mn} \sim N(\zeta_{n,Time}, 0.083)$. Using these relationships, the average values of time by gender and income category are calculated and shown in Table 16.

The values of intra-consumer heterogeneity are close to those obtained from the double mixture model in Chapter 2. However, the inter-consumer variances have lower magnitudes than those obtained in Chapter 2 because part of this heterogeneity is explained using the sociodemographic variables.

Table 15: Random and systematic heterogeneity in the scale parameter.

	Intercept	Menu	Menu2
Population Mean	1.744 (0.294)	3.776 (1.494)	-3.698 (1.430)
Inter-Consumer Variance	0.256 (0.107)	0.299 (0.149)	0.394 (0.214)
Intra-Consumer Variance		0.160 (0.058)	

Table 16: Average value of time (CHF/hr)

	Low Income	Medium Income	High Income
Males	96.51	101.87	151.21
Females	60.08	63.41	94.13

A second degree polynomial is used to model systematic intra-consumer heterogeneity in the scale parameter (higher order polynomial terms are not significant). The values of m and m^2 were scaled by 10.0 and 100.0 respectively. We can express the menu-specific scale parameter as in equation 4.12 using the results shown in Table 15, where:

$$\zeta_{n,0} \sim N(1.744, 0.256) \qquad (4.14)$$

$$\zeta_{n,1} \sim N(3.776, 0.299)$$

$$\zeta_{n,2} \sim N(-3.698\ 0.394)$$

$$\omega_{mn,scale} \sim N(0, 0.160)$$

Figure 3 presents the fitted expectation of the scale parameter as a function of the menu number (for menus 1-9). This figure indicates an increase in scale up to menu 5, and a decrease afterwards. A higher scale parameter indicates a higher explanatory power of the model variables compared to the error component and vice versa. The increase might be due to the fact that respondents are becoming more familiar with the survey (e.g. presentation of the attributes, framing, etc.). The following decrease indicates the presence of "fatigue" which results in inattentiveness to the presented attributes. These results have implications on survey design, suggesting that having a large number of choice tasks per individual might lead to biases due to loss of attentiveness and fatigue.

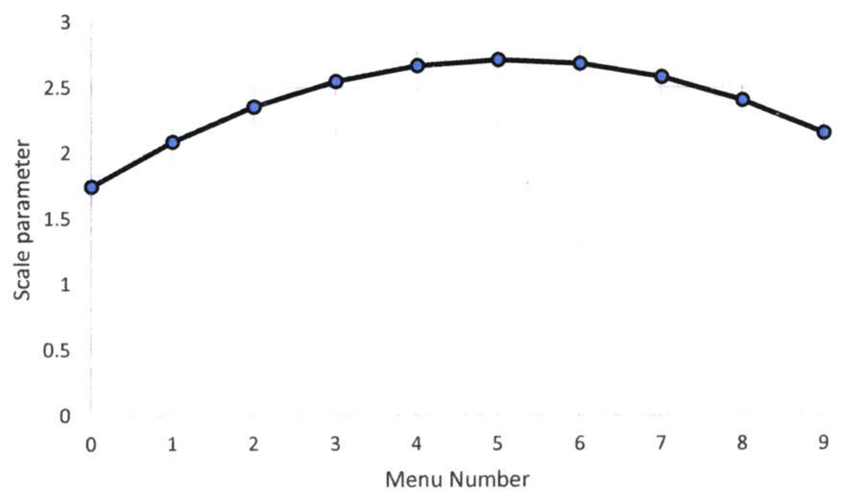

Figure 3: Change in the scale parameter by menu number.

4.2.8. Applications in Personalization

In the context of personalization and recommender systems, accounting for random and systematic heterogeneity has three advantages. First, accounting for systematic inter-consumer heterogeneity partially mitigates the "cold start" problem. The preferences of new users (i.e. with no choice history) can be inferred from their socio-demographics instead of presenting them with generic (non-personalized) recommendations. These socio-demographics can also be used to improve predictions for existing users with limited choice history as well.

Second, accounting for intra-consumer systematic heterogeneity enables us to provide context-aware recommendations, and thus make use of all the available data. Different recommendations can be proposed depending on the contextual variables that are used to model the menu/choice specific parameters. The individual-specific sensitivities to contextual variables are also learnt from repeated choices (in the same manner as the individual-specific means ζ_n are learnt in the double mixture model).

Third, accounting for systematic heterogeneity relaxes the assumptions on the mixing distributions; only assumptions on the distributions of the random part of heterogeneity are needed, while the distributions of the systematic parts are determined by the data.

Finally, the online procedure for updating preferences can be applied in the same manner as before; by iterating Steps 4 and 5 of the Gibbs sampler. This procedure is feasible in real-time, as it only uses the data of the individual making the choice.

4.3. LATENT CLASS EXTENSION

4.3.1. Motivation

The behavioral model used in Chapter 2 relies on two strong assumptions on the mixing distributions of inter- and intra-consumer heterogeneity. The first assumption is that the mixing distributions are assumed to be normal (or log-normal), i.e. the menu-specific parameters are normally distributed around the individual-specific means, which are in turn normally distributed around the population means. The key limitation of these parametric distributions is the assumption of uni-modality (Hess, 2014). Several studies have shown that nonparametric mixture models (i.e. latent class) can sometimes outperform continuous mixture models (e.g. Vij and Krueger, 2017), especially when the number of observations per individual is small (Andrews et al., 2002).

The second assumption is that the intra-consumer covariance matrix is assumed to be the same for all individuals, which means that the level of heterogeneity across different choices of a given individual is the same across the population. Becker et al. (2018) justified this assumption because it is not possible to estimate individual-specific covariance matrices given the small number of observations per individual.

This section extends the logit double mixture model and the estimator proposed by Becker et al. (2018) by incorporating a latent class model, which:

1. Relaxes the normality assumptions on the mixing distribution of inter-consumer heterogeneity and allows for approximating any shape. For example, a distribution with S modes can be approximated by a normal mixture of S classes.

2. Allows different individuals to have different intra-consumer covariance matrices, depending on the class to which they belong.

Semi-parametric distributions have not been used before in logit mixture models with inter- and intra-consumer heterogeneity, even though they have proven useful in models with inter-consumer heterogeneity only (Rossi et al., 2005; Bujosa et al., 2010; Greene and Hensher, 2013; Keane and Wasi, 2013; Krueger et al., 2018). In this section, I propose an efficient Hierarchical Bayes (HB) estimator for these models, and demonstrate that they can outperform logit mixture models with unimodal (normal or lognormal) distributions of inter- and intra-consumer heterogeneity.

4.3.2. Background

The majority of studies in discrete choice models have used normal and log-normal distributions, with a few using Johnson's S_b gamma, and triangular distributions (Train, 2016). Because of the limitations associated with these distributions, several extensions have been proposed in the literature including latent class models, models that account for scale heterogeneity, and discrete mixtures of continuous distributions (also known as mixed-mixed logit models).

Latent class models are semi-parametric models that assume a finite number of market segments, with preferences being homogenous within each segment. These models are advantageous as they do not require numerical simulation. However, they have been found to understate the extent of heterogeneity in the data (Elrod and Keane, 1995; Allenby and Rossi, 1998).

Several procedures for flexible mixing distributions were proposed by Bajari et al. (2007), Fosgerau and Bierlaire (2007), Train (2008), Fox et al. (2011), Burda et al. (2008), Train (2016), and others. Train (2016) proposed a logit-mixed-logit (LML) model in which any continuous distribution can be approximated by a discrete distribution defined with a logit kernel. Burda et al. (2008) suggested a convolution of a normal kernel with a skewing function.

Nevertheless, the most popular procedure to represent flexible mixing distributions is the mixture of normals. Different methods have been proposed for logit (Rossi et al., 2005; Bujosa et al., 2010; Greene and Hensher, 2013; Keane and Wasi, 2013; and Krueger et al., 2018) and probit (Geweke and Keane, 2001; 2007). These models subsume the advantages of continuous mixture models and latent class models. In addition, they are convenient for estimation using simulated maximum likelihood or Bayesian estimation methods (Rossi et al., 2005; Sarrias and Daziano, 2017). According to Vij and Kreuger (2017), semi-parametric mixture models can asymptotically approximate any shape, which allows for modeling complex patterns of heterogeneity that cannot be captured using unimodal distributions.

The following section presents an extension to the logit double mixture estimator using a similar latent class extension.

4.3.3. The Model

The model extends the double mixture logit model by assuming that the population is composed of S segments where $S \geq 2$. Each segment has its own mean (μ_s) and inter- and intra-consumer covariance matrices (Ω_s^b and Ω_s^w respectively). We also define s_n as a segment membership indicator for each individual n:

$$s_n \in \{1, 2., \dots, S\} \tag{4.15}$$

We assume that ζ_n and η_{mn} are normally distributed within each segment:

$$\zeta_{n|s_n=s} \sim \mathcal{N}_T(\mu_s, \Omega_s^b) \tag{4.16}$$

$$\eta_{mn|s_n=s} \sim \mathcal{N}_T(\zeta_n, \Omega_s^w) \tag{4.17}$$

Finally, we define the marginal probability of belonging to a segment s as π_s, and π is the vector $(\pi_1, \dots, \pi_s)$:

$$\pi_s = P(s_n = s) \tag{4.18}$$

Conditional on the segment membership, the probability of individual n choosing alternative j in menu m is:

$$P(d_{jmn} = 1 | s_n = s) = P(d_{jmn} = 1 | \mu_s, \Omega_s^b, \Omega_s^w)$$

$$= \int_{\zeta_n} \int_{\eta_{mn}} P_j(\eta_{mn})\, H(d\eta_{mn} | \zeta_n, \Omega_s^w)\, F(d\zeta_n | \mu_s, \Omega_s^b) \tag{4.19}$$

Where:

$$H(d\eta_{mn} | \zeta_n, \Omega_s^w) \sim \mathcal{N}_T(\zeta_n, \Omega_s^w) \tag{4.20}$$

$$F(d\zeta_n | \mu_s, \Omega_s^b) \sim \mathcal{N}_T(\mu_s, \Omega_s^b) \tag{4.21}$$

However, since the class membership is unknown, the unconditional probability is obtained as a weighted sum over all the possible segments (weighted by π_s):

$$P(d_{jmn} = 1 | \mu_s, \Omega_s^b, \Omega_s^w, \pi_s\ \forall\, s \in \{\mathbb{Z}: [1, S]\}) = \sum_{s=1}^{S} \pi_s P(d_{jmn} = 1 | \mu_s, \Omega_s^b, \Omega_s^w)$$

$$= \sum_{s=1}^{S} \pi_s \int_{\zeta_n} \int_{\eta_{mn}} P_j(\eta_{mn})\, H(d\eta_{mn} | \zeta_n, \Omega_s^w)\, F(d\zeta_n | \mu_s, \Omega_s^b) \tag{4.22}$$

This model can be estimated by extending the 5-steps Gibbs Sampler as follows. Hierarchical Bayes (HB) and data augmentation are needed, where we assume that s_n, ζ_n, and η_{mn} are unknowns. We also use a diffuse prior on π_s defined as $k(\pi)$.

The posterior on $\zeta_n, s_n, \eta_{mn}, \mu_s, \Omega_s^w, \Omega_s^b\ \forall\ s$, and π is given by equation (4.23):

$$K(\zeta_n, s_n\ \forall n, \eta_{mn}\ \forall mn, \mu_s, \Omega_s^w, \Omega_s^b, \pi_s\ \forall\ s|d)$$
$$\propto \prod_{n=1}^{N}\left[\prod_{m=1}^{M_n}\left[\prod_{j=1}^{J_{mn}}[P_j(\eta_{mn})^{d_{jmn}}]\, h(\eta_{mn}|\zeta_n, \Omega_{s_n}^w)\right] f(\zeta_n|\mu_{s_n}, \Omega_{s_n}^b) P(s_n|\pi)\right] k(\pi) \prod_{s=1}^{S} k(\Omega_s^w) k(\mu_s) k(\Omega_s^b)$$

(4.23)

Where:

$$k(\mu_s) \sim \mathcal{N}_T(\mu_{0,s}, A_s) \qquad (4.24)$$

$$k(\Omega_s^b) \sim \mathrm{HIW}(\nu_{bs}, A_{bs}) \qquad (4.25)$$

$$k(\Omega_s^w) \sim \mathrm{HIW}(\nu_{ws}, A_{ws}) \qquad (4.26)$$

$$k(\pi) \sim \mathrm{DIR}(\alpha) \qquad (4.27)$$

where $\mu_{0,s}$ represents a vector of prior means, A_s is a diagonal covariance matrix with diagonal values $\rightarrow \infty$ (uninformative prior), HIW is the Hierarchical Inverted Wishart (Half-t) prior with parameters $\nu_{bs}, \nu_{ws}, A_{bs}$, and A_{ws} (Huang and Wand; 2013), and α is a vector of concentration parameters (e.g. $\alpha_s = 1\ \forall\ s$).

This model is estimated using a 7-step Gibbs sampler which draws from the conditional posteriors shown below:

Step I: drawing from the conditional posterior of the population means for each class s:

$$K(\mu_s|\zeta_n, s_n \forall n, \eta_{mn}\ \forall mn, \Omega_s^w, \Omega_s^b, \pi_s) \propto f(\zeta_n \forall n_{|s_n=s}|\mu_s, \Omega_s^b) k(\mu_s) \qquad (4.28)$$

This includes S Normal Bayesian updates with known variances and unknown means. The conditional posterior on each μ_s is $\mathcal{N}\left(\bar{\zeta}_s, \frac{\Omega_s^b}{N_s}\right)$, $\bar{\zeta}_s$ is the mean of ζ_n draws over all individuals who belong to segment s (in the same Gibbs iteration) ($s_n = s$), and N_s is the number of individuals belonging to segment s in this iteration.

Step II: drawing from the conditional posterior of the population level covariance matrix for each class:

$$K(\Omega_s^b|\zeta_n, s_n\ \forall n, \eta_{mn}\ \forall mn, \mu_s, \Omega_s^w, \pi_s) \propto f(\zeta_n \forall n_{|s_n=s}|\mu_s, \Omega_s^b) k(\Omega_s^b) \qquad (4.29)$$

This includes S normal Bayesian updates with unknown variances and known means, using HIW priors. Each step is the same as Step 2 in the original Gibbs sampler, however, it uses only observations belonging to class s in the previous iteration.

Step III: drawing from the conditional posterior of the segment probabilities π_s:

$$\mathrm{K}(\pi_1, \ldots, \pi_S | \zeta_\mathrm{n}, s_n\ \forall \mathrm{n}, \eta_{\mathrm{mn}}\ \forall \mathrm{mn}, \mu_\mathrm{s}, \Omega_\mathrm{s}^\mathrm{w}, \Omega_\mathrm{s}^\mathrm{b}) \propto \mathrm{P}(s_n|\pi)\mathrm{k}(\pi) \tag{4.30}$$

This step is simply drawing from a Dirichlet distribution with parameters $N_s + \alpha_s$ for each segment s.

Step IV: drawing from the conditional posterior of the class membership indicators $s_n\ \forall\ n$:

$$P(s_n|\mu_s, \Omega_s^\mathrm{b}, \Omega_s^\mathrm{w}\ \forall\ s, \pi_s, \zeta_n\ \forall\ n, \eta_{mn} \forall\ mn) \propto \tag{4.31}$$

$$P(s_n|\pi) f(\zeta_n|\mu_{s_n}, \Omega_{s_n}^\mathrm{b}) \prod_{m=1}^{M_n} h(\eta_{mn}|\zeta_n, \Omega_{s_n}^\mathrm{w})$$

This can be performed by calculating the segment probabilities for each individual and simulating segment membership:

$$P(s_n = s|\mu_s, \Omega_s^\mathrm{b}, \Omega_s^\mathrm{w}, \pi_s \forall s, \zeta_n\ \forall\ n, \eta_{mn} \forall\ mn) \tag{4.32}$$

$$= \frac{\pi_s f(\zeta_n|\mu_s, \Omega_s^\mathrm{b}) \prod_{m=1}^{M_n} h(\eta_{mn}|\zeta_n, \Omega_s^\mathrm{w})}{\sum_{k=1}^{S} \left[\pi_k f(\zeta_n|\mu_k, \Omega_k^\mathrm{b}) \prod_{m=1}^{M_n} h(\eta_{mn}|\zeta_n, \Omega_\mathrm{k}^\mathrm{w})\right]}$$

Step V: drawing from the conditional posterior of the intra-consumer covariance matrix for each segment:

$$\mathrm{K}(\Omega_\mathrm{s}^\mathrm{w}|s_n, \zeta_\mathrm{n}\ \forall \mathrm{n}, \eta_{\mathrm{mn}}\ \forall \mathrm{mn}, \mu_\mathrm{s}, \Omega_\mathrm{s}^\mathrm{b}, \pi_s) \propto \mathrm{h}(\eta_{\mathrm{mn}}\ \forall \mathrm{mn}_{|s_\mathrm{n}=s}|\zeta_\mathrm{n}\ \forall \mathrm{n}, \Omega_\mathrm{s}^\mathrm{w})\mathrm{k}(\Omega_\mathrm{s}^\mathrm{w}) \tag{4.33}$$

This includes S normal Bayesian updates with unknown variances and known means, using HIW priors. Each step is the same as Step 3 in the original Gibbs sampler, however, it uses only observations belonging to class s in the previous iteration.

Step VI: drawing from the conditional posterior of the individual level means:

$$\mathrm{K}(\zeta_\mathrm{n}|\mu_\mathrm{s}, \eta_{\mathrm{mn}}, \Omega_\mathrm{s}^\mathrm{b}, \Omega_\mathrm{s}^\mathrm{w}, s_n, \pi) \propto \mathrm{h}(\eta_{\mathrm{mn}}|\zeta_\mathrm{n}\ \forall \mathrm{n}, \Omega_{\mathrm{s_n}}^\mathrm{w})\mathrm{f}(\zeta_\mathrm{n}|\mu_{\mathrm{s_n}}, \Omega_{\mathrm{s_n}}^\mathrm{b}) \tag{4.34}$$

This is performed using a normal Bayesian update with η_{mn} as the data and $N(\mu_{s_n}, \Omega_{\mathrm{s_n}}^\mathrm{b})$ as a prior.

Step VII: drawing from the conditional posterior of the individual- and menu-specific parameters:

$$K(\eta_{mn}|\mu_s, \zeta_n, \Omega_s^b, \Omega_s^w, \pi, s_n) \propto \prod_{j=0}^{J_{mn}} [P_j(\eta_{mn})^{d_{jmn}}] h(\eta_{mn}|\zeta_n, \Omega_{s_n}^w)$$

$$n = 1, 2, \dots, N, m = 1, \dots M_n \tag{4.35}$$

A draw of η_{mn} is obtained by using the Metropolis-Hastings procedure, where the jumping distribution is $N(\zeta_n, \Omega_{s_n}^w)$.

4.3.4. Monte Carlo Experiment

Monte Carlo data are used in order to demonstrate the estimation procedure presented above. A simpler version of the Mobility as a Service (MaaS) experiment described in Chapter 3 is used, but with the addition of intra-consumer heterogeneity. We assume that individuals are presented with 8 menus, and that each menu includes three MaaS plans (with no opt-out alternative). Each plan has two binary attributes: access to transit (T) and access to on-demand services (OD).

The utility functions are given by equation (4.36):

$$V_{jmn} = \exp(\alpha_{mn})(-P_{jmn} + \beta_{T,mn} T_{jmn} + \beta_{OD,mn} OD_{jmn}) \quad , \quad j = 1, 2, 3 \tag{4.36}$$

The population consists of two segments with probabilities 0.6 and 0.4 (π_1 and π_2 respectively). All parameters (α_{mn}, $\beta_{T,mn}$, and $\beta_{OD,mn}$) are normally distributed within each class with inter- and intra-consumer heterogeneity. Class 1 has high inter-consumer heterogeneity and low intra-consumer heterogeneity, and class 2 has high intra-consumer heterogeneity and low inter-consumer heterogeneity. Class 1 has a higher preference to on-demand services on average, while class 2 has a higher preference to transit. The true values of the parameters (population means and inter- and intra-consumer standard deviations) are presented in Table 17.

The sample size used in the experiment is 20,000 individuals. Estimation is done using 40,000 Gibbs iterations, the first 20,000 of which are burn-in iterations, and the remaining 20,000 are used for sampling from the posterior distributions. The estimation results are also presented in Table 17. The results indicate that the estimator is able to recover the true values of the model parameters.

4.3.5. Benchmarking with a Simpler Model

In order to demonstrate the benefits of the latent class extension, we compare the predictions and the conditional log-likelihood of the estimated model to those obtained from the logit double mixture model (presented in Chapter 2). A hold-out sample is generated from the same users that were used in estimation (but with different menus). Table 18 presents the mean and 95% confidence interval of the posterior predictive distribution (PPD) and the conditional log-likelihood of the two models. The predictions of the model with the latent class extension are approximately 1.8% better than those of the logit double mixture model, and the conditional log-likelihood is higher by 2911 points.

Table 17: Estimation results with the latent class extension for inter- and intra-consumer heterogeneity.

Class 1						
	Population Mean		Inter-consumer var.		Intra-consumer var.	
Parameter	Estimate	True value	Estimate	True value	Estimate	True value
Scale	1.004 (0.016)	1.000	0.428 (0.035)	0.500	0.275 (0.035)	0.250
Transit	0.003 (0.014)	0.000	0.524 (0.025)	0.500	0.218 (0.025)	0.250
On-Demand	3.001 (0.026)	3.000	0.489 (0.028)	0.500	0.280 (0.028)	0.250
Class Prob.	0.599 (0.004)					
Class 2						
	Population Mean		Inter-consumer var.		Intra-consumer var.	
Parameter	Estimate	True value	Estimate	True value	Estimate	True value
Scale	1.018 (0.022)	1.000	0.257 (0.051)	0.250	0.542 (0.051)	0.500
Transit	3.046 (0.035)	3.000	0.302 (0.042)	0.250	0.531 (0.042)	0.500
On-Demand	-0.028 (0.020)	0.000	0.236 (0.038)	0.250	0.545 (0.038)	0.500
Class Prob.	0.401 (0.004)					

Table 18: Comparison between the standard double mixture model and the latent class extension.

	Conditional Log-Likelihood	PPD
Standard double mixture model	-55233.5	0.684 [0.682, 0.686]
Latent class extension	-52322.6	0.701 [0.700, 0.704]

4.3.6. Modified Online Procedure

Similar to the offline-online methodology proposed in Chapter 2, an online estimator of individual- and menu-specific parameters can be used based on the 7-step Gibbs sampler described above. The offline procedure iterates steps 1-7 and updates all of the model parameters (μ_s, Ω_s^{b}, Ω_s^{w}, π_s, ζ_n, η_{mn}, and s_n $\forall \mathrm{m}, n$).

On the other hand, the online procedure assumes that the population level parameters (π_s, μ_s, Ω_s^{b}, and Ω_s^{w}) are fixed, and iterates over steps 4, 6, and 7 only to update the class membership (s_n) and the individual- and menu-specific parameters ζ_n and η_{mn}.

4.3.7. Discussion

The model described above enhances the double mixture model with inter- and intra-consumer heterogeneity by incorporating a latent class model. Since the individual- and menu-specific parameters are assumed to be normally distributed conditional on the class, the HB estimator retains the desirable properties of the normal distribution in Gibbs sampling (conjugate priors on the population means and covariance matrices), and uses the Dirichlet and Multinomial distributions (which also have the convenient conjugacy property).

In estimating such models, careful consideration is needed in deciding on the optimal number of classes. One approach is to use the Akaike Information Criterion (AIC) which is common in finite mixture models (e.g. the Consistent Akaike Information Criterion, CAIC (Bozdogan, 1987) and the Bayesian Information Criterion, BIC (Schwartz, 1978)). These approaches require estimating models with a different number of classes beforehand and deciding on the best model. Another possible approach is to use out-of-sample cross validation as in Section 4.3.5.

This model can also be extended to include a "Chinese Restaurant Process", in which the number of classes varies at each iteration of the Gibbs sampler. Such models have been proposed by Burda et al. (2008) for logit and probit mixture models (with inter-consumer heterogeneity only). However, since the number of classes varies at each iteration, these models might not be convenient in online estimation.

Another relevant extension is to enrich the class membership model by including covariates, such as socio-economic variables and individual characteristics. For example, a multinomial logit model can be used instead of the fixed probabilities π. This extension is particularly useful in improving the predictions of new users (i.e., unconditional predictions). As for existing users (with a long choice history), their class membership can be learnt from repeated choices (in a similar manner to learning ζ_n). This extension includes substituting Step 3 (drawing from the Dirichlet distribution) with a Metropolis-Hastings step that draws from the parameters of the class membership model.

4.4. CONCLUSION

In this chapter, two extensions to the logit double mixture model are proposed. The first extension allows for modeling random and systematic inter- and intra-consumer heterogeneity. This relaxes the assumption that heterogeneity is completely random, and uses socio-demographic characteristics and contextual variables in order to partially explain individual- and menu-specific preferences. The second extension incorporates a latent class model, and extends the widely used mixture of normals model (mixed-mixed logit) to inter- and intra-consumer heterogeneity. Both extensions entail fewer assumptions on the inter- and intra-consumer mixing distributions of preferences, and allow the shapes of these distribution(s), to be partially determined by the data.

These models are mainly constrained by computation. In both applications, all the inter- and intra-consumer covariance matrices are diagonal, and therefore the computation time is significantly reduced. However, if full covariance matrices need to be estimated, the running times will be substantially exacerbated. For example, in the first application, multivariate regressions need to be estimated in Steps 1 and 4 (and the correlations need to be accounted for in Steps 2 and 3 as well). Finally, due to the model complexity, both models required large Monte Carlo datasets for estimation.

Both models can be easily applied in a recommendation system setting similar to the one proposed in Chapter 2, by simple extensions to the online estimation procedure. Future research should focus on such applications and quantify the benefits of these two extensions on real datasets.

CHAPTER 5: CASE STUDY: SUSTAINABLE TRAVEL INCENTIVES WITH PREDICTION, OPTIMIZATION, AND PERSONALIZATION

5.1. INTRODUCTION

The online estimation methodology presented in Chapter 2 was validated using "static" SP data on transportation mode choice. Consumer behavior may differ significantly between SP experiments and app-based choices. For instance, the time intervals between successive choices may vary considerably between the app-based systems and SP experiments, resulting in higher magnitudes of intra-consumer heterogeneity. In addition, unlike adaptive recommender systems, all of the attributes in the SP data are exogenous.

This chapter presents a real application to an online travel advisor, where menus are generated based on previous choices, and preferences are updated in real-time. Scalability is demonstrated by simulating a large synthetic population of users. Models with inter- and intra-consumer heterogeneity are estimated using choices obtained from personalized menus, and preferences are updated after each choice using online estimation.

This methodology is applied to *Tripod*, an app-based on-demand system that influences individuals' real-time travel decisions by offering them information and incentives with the objective of achieving system-wide energy savings (Azevedo et al., 2018). Revealed and stated preferences (RP/SP) data are collected from a sample of 202 users in Greater Boston Area (GBA). A discrete choice model with inter- and intra-consumer heterogeneity is estimated using the SP data to model the choice among different alternatives. A subset of these alternatives is recommended by Tripod, and incentives are allocated to the alternatives that result in positive energy savings. The estimated parameters are used in a pilot of the Tripod app, where user preferences are updated using the online estimation procedure described in Chapter 2.

In addition, a simulation of Tripod is presented in which the offline-online estimation framework is applied to a larger sample. The simulation uses synthetic users whose trips are generated based on the trip patterns of the survey respondents, and whose preference parameters are generated based on the estimated parameters from the SP data collected from the same respondents. Online estimation procedure is performed after each choice, and offline estimation is performed periodically. Learning individual preferences is examined by analyzing the evolution of the average hit-rate, defined as the average probability of selecting one of the recommended alternatives.

The remainder of this chapter is organized as follows: section 5.2 presents a brief overview of Tripod. Section 5.3 presents the data collection methodology and the experimental design. Section 5.4 presents the estimation results and their deployment in the Tripod application. Section 5.5 presents the Tripod simulation and the application of the offline-online estimation. Finally, Section 5.6 concludes this chapter.

5.2. TRIPOD OVERVIEW

Tripod (Sustainable Travel Incentives with Prediction, Optimization and Personalization) is an app-based trip planner offering a personalized menu with real-time predicted information and optimized incentives in the form of tokens that can be exchanged for goods and services (Azevedo et al., 2018). The travel decisions of interest are mode choice, route choice, departure time, trip-making, and driving style. In response to changes in any of those dimensions, users receive incentives in the form of *tokens* that can then be redeemed in a market place for a variety of goods and services offered by third party providers.

Tripod maximizes system-wide energy savings taking into account system-wide supply and demand interactions as well as individual specific preferences. The optimization problem is decomposed into two tractable, loosely coupled problems: System Optimization (SO) and User Experience (UE) (Azevedo et al., 2018). Tripod's framework is shown in Figure 4 below.

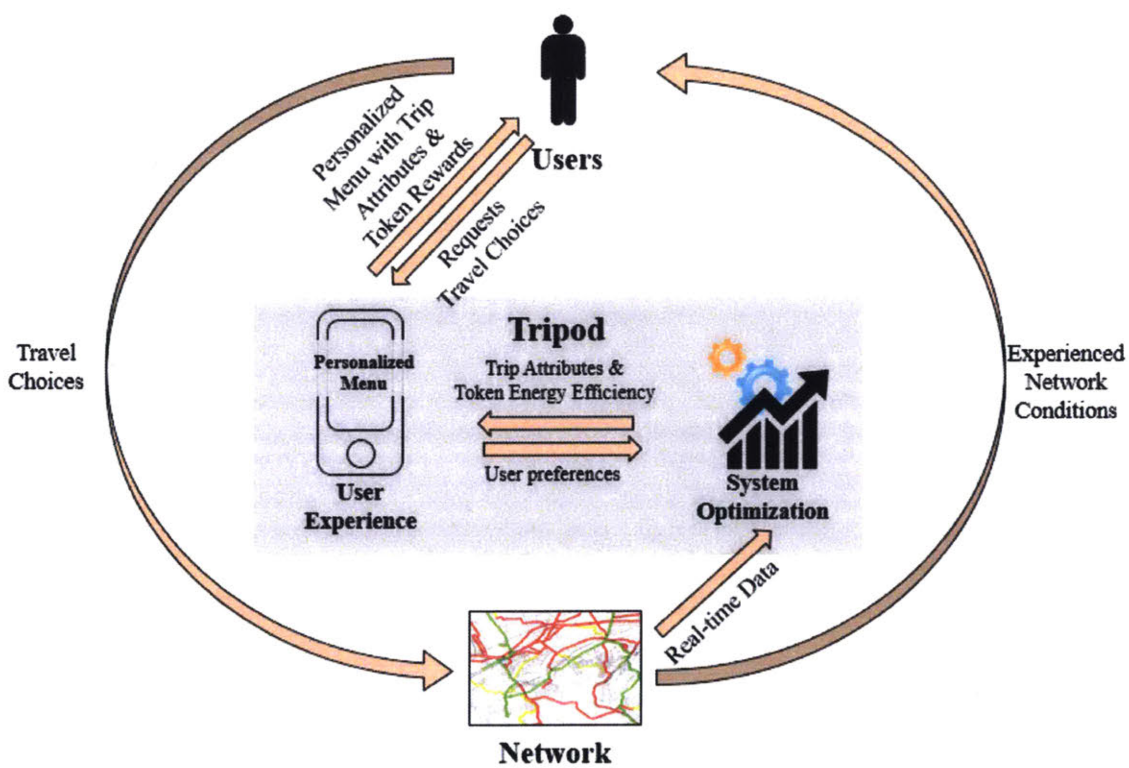

Figure 4: Tripod Framework.

System Optimization estimates the current state of the transportation network, predicts the state of the network given different token awarding strategies, and optimizes this strategy. The output of this optimization is the system-wide token energy efficiency value in terms of energy savings per token, which is used as an input to User Experience.

User Experience includes two components: the Preference Updater and the User Optimization (UO). The Preference Updater uses the offline-online estimation framework presented in Chapter 2 in order to update individual level preferences after each choice. UO uses the estimated individual level preferences to generate a personalized menu of travel options to Tripod users upon request. It uses the system-wide token energy efficiency as well as the transportation performance predictions and the energy impacts generated by SO (Azevedo et al, 2018).

A Tripod user has to subscribe to the app and, before each trip, decide whether to request a Tripod menu. The personalized menu is presented to the user (see Figure 5) with information about the recommended alternatives and the associated tokens. The user may select an option from the menu and use the Tripod app to navigate to the destination or opt-out and travel without Tripod's guidance and rewards. In the first case, the app monitors the travel of the user and rewards her/him at the end of the trip if the guidance was followed.

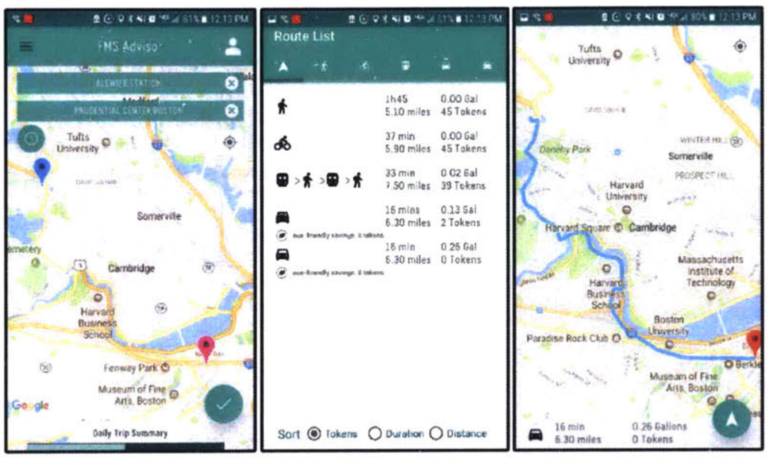

Figure 5: User interface of the Tripod app.

5.3. DATA COLLECTION

The parameters in Tripod UE are obtained from a behavioral model estimated using SP data collected in GBA. SP data can be used to initialize the recommender system described in Chapter 2 when data are not available. They can also be used to enrich the existing data, especially because they do not suffer from the endogeneity problem described in Chapter 3.

5.3.1. Methodology and Experimental Design

The SP data collection method used in this chapter leverages RP data and pre-survey data in generating SP choice tasks. The RP data includes the departure and arrival times, origin and destination, trip mode and purpose, and activity duration. These are obtained using the smartphone-based sensing app, FMS (Zhao et al., 2015a; 2015b). External contextual data are obtained from a trip planner (e.g. Google Maps). These are used together with the user's personal characteristics obtained from a pre-survey, e.g., car ownership, usage of car/bike sharing services, etc. The experimental design uses those data in order to generate SP choice experiments with reasonable alternatives and attributes.

Profile generation is based on the random design, however, constraints are used in order to enhance realism and remove dominant or inferior alternatives as recommended by Walker et al. (2018). After selecting a reference trip from the activity diary, the trip context (activity, origin, destination, and departure time) is fed into the Google Maps API in order to obtain some baseline attributes, such as travel times using different modes and transit access and egress times. These attributes are then pivoted using random "*design parameters*" in order to construct the attributes shown in the profiles. Design parameters are chosen from a pre-defined set of levels. For example, travel time by car is calculated by multiplying the travel time obtained from Google Maps by a design parameter, the *travel time ratio*, with levels of 0.7, 0.85, 0.95, 1.00, 1.05, 1.15, 1.20, or 1.50. This means that the pivoted travel time can range between 70% and 150% of the travel time obtained from Google Maps for the same trip. A uniform distribution is used such that each level of the design parameter has an equal probability of being selected (i.e., the travel time ratio is equally likely to be chosen from any of the values above). More information on the experimental design is presented in Danaf et al. (2019).

The user flow and choice task generation process are explained below:

1. Upon registering to the platform, the user's socio-demographic data as well as attitudes and perceptions are collected using a pre-survey.
2. The app starts to track the user, who then has to validate the activity diaries with trip and activity information every day, directly on the smartphone app.
3. After a full day is validated, a trip context (activity, origin/destination, and departure time) is selected from the trips that the user has performed as a reference trip.
4. The alternatives and attributes (e.g., travel times) corresponding to the selected trip context are retrieved from Google Maps.
5. SP profiles are generated using the random design explained above.
6. The choice tasks are checked afterwards for realism and internal validity.
7. A validated choice task is presented to the user in the FMS platform, who is asked about his/her choice if he/she were to repeat the trip under different hypothetical scenarios (including Tripod alternatives).

5.3.2. Data Collection

Using the methodology explained above, data collection for Tripod was carried out in GBA where 1940 observations were obtained from 202 participants, 154 of which have finished the required 14 days of responses and exited the survey.

In the pre-survey, users were asked about their socio-demographic characteristics such as age, gender, working status, income, car ownership, bike ownership, and how frequently they use different transportation modes. Descriptive statistics of the sample and screenshots of the survey are presented in Appendix C1. Since the survey is smartphone-based, the sample is biased towards young respondents.

The universal set of SP alternatives includes non-motorized modes (walking, biking, and bike-sharing), private motorized modes (car and carpooling), on-demand modes (e.g. Uber/UberPool, Lyft/Lyft Line, car sharing, and taxi), and transit (with walk, bike, or car access). The attributes of these alternatives are presented in Appendix C2.

The Tripod menu includes a subset of the existing alternatives with changes across multiple dimensions that generate energy savings, e.g., the departure time may be delayed (between 15 and 90 minutes), a different route may be suggested, or the driver may be asked to drive in an eco-friendly way. Information on energy savings (relative to the RP choice) and tokens assigned to alternatives are also presented in the Tripod menu. Energy consumption values are obtained using the TripEnergy software (Needell et al., 2016). Only alternatives with positive energy savings are included in this menu in the SP.

In the first SP choice task, users are presented with a "marketplace" showing the items that can be traded in for tokens. The redeem value of tokens is randomly generated for each individual but then kept the same for that individual during the data collection period. The marketplace can be accessed to view the available items and the redeem values anytime throughout the SP.

5.4. TRIPOD IMPLEMENTATION

The SP data were used to implement a pilot of the Tripod app in order to demonstrate how individual preferences are estimated and updated in real-time. The population parameters used by the online estimation were obtained from a behavioral model estimated using SP data. The Preference Updater and User Optimization were implemented as discussed in Section 5.2 (Peaks, 2018; Song et al., 2018).

5.4.1. Model Estimation

A discrete choice model with inter- and intra-consumer heterogeneity was estimated using the collected SP data. Only observations in which the respondent accessed the Tripod menu were included in estimation. The final sample consisted of 772 observations collected from 135 individuals.

The model assumes log-normal distributions on the coefficients of in-vehicle and out-of-vehicle travel time, tokens and scale, and normal distributions on the coefficients of car, transit, walking, and on-demand modes. All of these coefficients have inter- and intra-consumer heterogeneity. The mean and standard deviation of biking are normalized to 0. In addition, individual characteristics such as age, income, and dummy variables for being a student or having zero household vehicles are included without any heterogeneity. The systematic utility is presented in equation (5.1).

$$V_{jmn} = \exp(\alpha_{mn})(-Cost_{jmn} + \beta_{car,mn} IsCar_{jmn} + \beta_{walk,mn} IsWalk_{jmn} + \beta_{transit,mn} IsTransit_{jmn} + \beta_{OnDemand,mn} IsOnDemand_{jmn} - \exp(\beta_{IVTT,mn})\, IVTT_{jmn} - \exp(\beta_{OVTT,mn})\, OVTT_{jmn} + \exp(\beta_{token,mn})\, Tokens_{jmn} + \beta_{Xj} X_n) \quad (5.1)$$

Where:

- α_{mn} is a scale parameter.
- $IsCar_{jmn}$, $IsWalk_{jmn}$, $IsTransit_{jmn}$, and $IsOnDemand_{jmn}$ are binary variables equal to 1 if the mode in alternative j is car, walk, transit, or on-demand and zero otherwise.
- $Cost_{jmn}$ is the total cost of alternative j in 10's of USD (scaling is needed to obtain coefficients in the same order of magnitude).
- $IVTT_{jmn}$ and $IVTT_{jmn}$ are the in-vehicle and out-of-vehicle travel time of alternative j respectively in minutes.
- $Tokens_{jmn}$ is the Dollar equivalent of tokens offered for alternative j, obtained from the individual-specific marketplace (as different individuals were presented with different marketplace menus in the SP).

- X_n is a vector of socio-demographics of individual n including age, log of income, and dummy variables for having zero household vehicles and being a student.

The model was estimated using the 5-step Gibbs sampler proposed by Becker et al. (2018) and Ben-Akiva et al. (2019). The estimation results are presented in Table 19, and the posterior standard deviations are shown in parentheses. The mean values of in-vehicle and out-of-vehicle travel time across all observations (mean of $\exp(\beta_{IVTT,mn})$ and $\exp(\beta_{OVTT,mn})$ draws) are 42.1 USD/hr and 73.2 USD/hr respectively, and the medians are 23.0 USD/hr and 50.3 USD/hr. The mean value of 1.0 USD-equivalent of tokens is 0.69 USD. It is also observed that intra-consumer heterogeneity generally has a lower magnitude than inter-consumer heterogeneity (which was also observed in the Swissmetro data), indicating more variability in preferences between individuals than between different choices of the same individual.

The age parameter in the utility equations of on-demand modes is negative indicating that older respondents are less likely to choose these alternatives. The student dummy parameter in the non-motorized utility equations is positive indicating the students are more likely to choose walking or biking. The logarithm of income has a positive coefficient in the car utility equation indicating that people with higher income are more likely to choose car alternatives. Finally, respondents with zero household vehicles are more likely to choose transit alternatives since the corresponding parameter has a positive sign in the transit utility equations.

Table 19: Model estimation results for Tripod (posterior standard deviations in parentheses).

	Mean	Inter-consumer variance	Intra-consumer variance
Walk constant	1.236 (0.225)	0.689 (0.334)	0.347 (0.174)
Car constant	-0.749 (0.440)	0.690 (0.399)	0.345 (0.184)
On-demand constant	0.886 (0.424)	0.533 (0.263)	0.302 (0.230)
Transit constant	-0.658 (0.289)	0.686 (0.369)	0.402 (0.148)
Tokens	-0.900 (0.295)	0.713 (0.486)	0.344 (0.156)
Scale	0.355 (0.157)	0.659 (0.267)	0.217 (0.129)
In-vehicle travel time	0.836 (0.493)	0.901 (0.630)	0.297 (0.138)
Out-of-vehicle travel time	1.814 (0.148)	0.279 (0.091)	0.173 (0.057)
Age (on-demand)	-0.269 (0.101)	-	-
Student dummy (non-motorized)	0.756 (0.459)	-	-
Log income (car)	0.329 (0.104)	-	-
Zero Household vehicles (transit)	0.941 (0.420)	-	-

5.4.2. Application of the Online Estimation

The results in Table 19 are deployed in a pilot of the Tripod app implemented on the FMS platform (*FMS Advisor*, Peaks (2018)). Using the estimation results of the full 5-step Gibbs sampler, the population level parameters are used to update preferences after each choice using online estimation. Menu optimization is performed using the optimization method proposed by Song et al. (2017; 2018). A constraint is included in the optimization to guarantee that at least one alternative per mode is included in each menu.

Due to the limited amount of data available, the full offline-online estimation framework described in Chapter 2 was not tested on real data. However, the effect of personalization was tested with two real users labelled as a "balanced user" and a "car lover": the balanced user evaluates all of the alternatives in the menu and selects the best one, while the car lover always selects a car alternative. The two users were asked to request menus for 20 trips having the same origins and destinations, and make choices accordingly. Figure 6 shows the first and 20th menus for the two users (this figure was generated using an earlier version of the app that did not include on-demand modes). The first menu for both users is shown in 6(a), and the twentieth menus for the "balanced user" and the "car lover" are shown in 6(b) and 6(c) respectively.

This experiment shows that the balanced user still obtains balanced menus after 20 choices including multiple alternatives of walking, biking, and car, while the car lover obtains one alternative of each of walking, biking, and transit, and 7 car and carpooling alternatives.

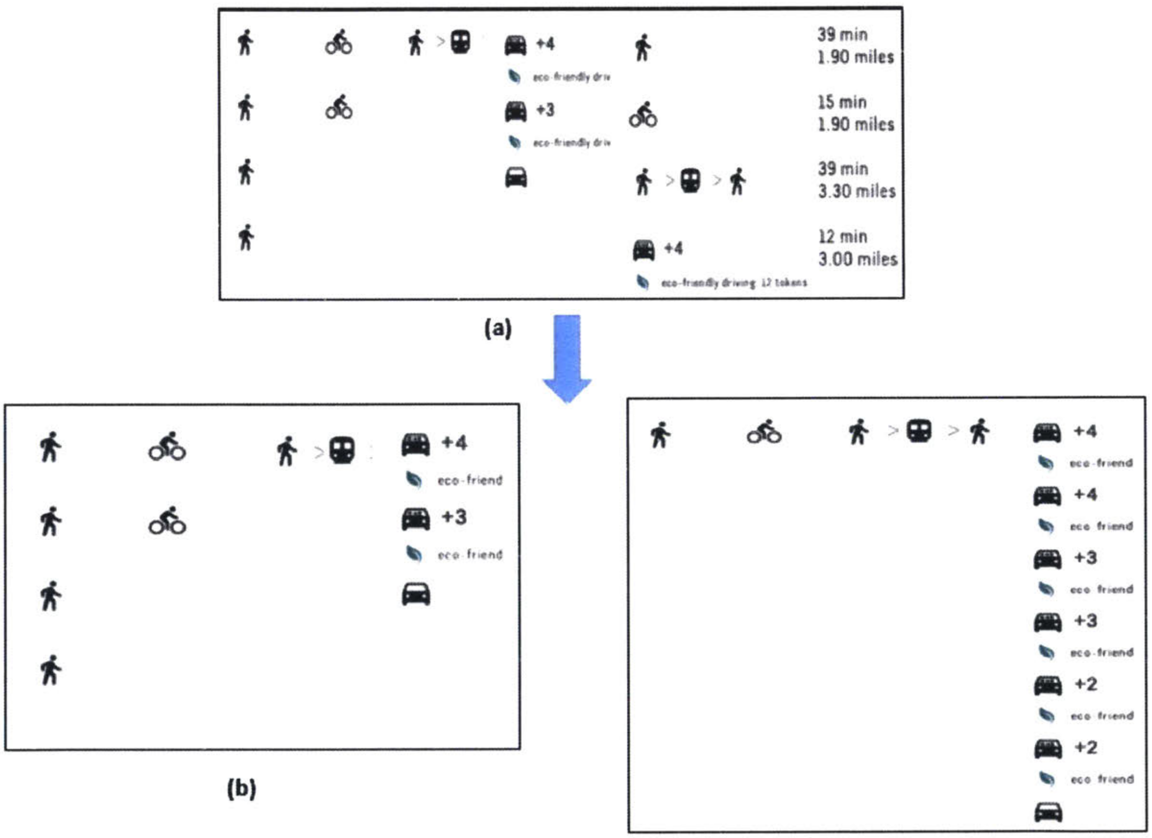

(c)

Figure 6: Tripod menus for the two users: (a) first menu for both users, (b) twentieth menu for the "balanced" user, and (c) twentieth menu for the "car lover".

5.5. TRIPOD SIMULATION

Using the same estimated parameters, Tripod is simulated with a larger synthetic population of users in order to demonstrate the process of generating personalized recommendations and the application of the offline-online estimation framework to update individual preferences. The following sections present the simulation setup, synthetic population generation, and results.

The objectives of the simulation are:

1. Validating the offline-online estimation methodology using realistic data, and a large number of users.
2. Benchmarking against a simpler method that does not apply online estimation and does not account for intra-consumer heterogeneity.
3. Analyzing the effect of personalization by comparing personalized and non-personalized recommendations.
4. Testing the effect of the frequency of offline estimation under different circumstances.
5. Testing the effect of inter- and intra-consumer heterogeneity levels.

5.5.1. Simulation Setup

This simulation represents a realistic application of Tripod, in which three alternatives are recommended from a universal set of 15 alternatives obtained from a trip planner (e.g. Google Maps). Users are presented with personalized menus including different modes (car, walking, biking, on-demand, and transit). The attributes associated with these modes are travel time, travel cost, and tokens, which are all shown to users.

Personalized recommendations are generated for each trip once a Tripod menu is requested. Draws of the menu-specific parameters (η_{mn}) of the previous menus are used in order to predict the probability of choosing each alternative from the universal choice set (without accounting for tokens). Based on the predicted probabilities, three alternatives are recommended. The tokens of the recommended alternatives are calculated using the optimization algorithm (Peaks, 2018). Non-recommended alternatives are assigned zero tokens.

Following the framework described in Chapter 2, online estimation is used to update individual preferences after each choice, and offline estimation is performed periodically. The frequency of offline estimation is varied and the results are analyzed.

5.5.2. Synthetic Population

Generating the synthetic population involves generating the individual preferences, trip patterns, and Tripod usage probabilities.

<u>Individual Preferences</u>

The true values of the preference parameters (which include the population parameters shown in Table 20 and the individual-specific means ζ_n) are assumed to be unknown, and they are only used in simulating the choices. Individual means (ζ_n) are generated for each user by drawing from the distribution:

$$\zeta_n \sim N(\mu, \Omega^{\mathrm{b}}) \qquad (5.2)$$

Where μ and Ω^{b} are the true values of the population means and the inter-consumer covariance matrix respectively (shown in Table 20, based on Table 19). The values of ζ_n are stored for each individual and assumed to be unchanged throughout the simulation. On the other hand, menu-specific parameters are simulated before each choice by drawing from the distribution:

$$\eta_{mn} \sim N(\zeta_n, \Omega^{\mathrm{w}}) \qquad (5.3)$$

Where Ω^w is the true intra-consumer covariance matrix.

In order to initialize the system, we need estimates of the population parameters (μ, Ω^b, and Ω^w). We simulate a pre-deployment data collection process using a smaller sub-sample of the synthetic population (1000 users), but with non-personalized incentives. This can either mimic an initial warm-up period of the system or an SP survey. The estimates obtained from this dataset are used as starting values for the population parameters, and are updated after the first offline estimation. The true values of the population parameters and the starting values are shown in Table 20.

Table 20: Starting values and true values of the population parameters in the Tripod simulation.

Starting Values								
	Tokens	IVTT	OVTT	Walking	Car	On-Demand	Transit	Scale
Population Mean	-0.794	0.759	1.657	1.310	-0.724	-0.943	-0.700	0.314
Inter-Consumer Std. Dev.	0.650	1.472	0.602	0.575	0.741	0.960	0.556	0.776
Intra-Consumer Std. Dev.	0.170	0.421	0.253	0.356	0.250	0.529	0.515	0.300
True Values (Boston SP estimates)								
	Tokens	IVTT	OVTT	Walking	Car	On-Demand	Transit	Scale
Population Mean	-0.900	0.836	1.614	1.236	-0.749	-0.886	-0.658	0.355
Inter-Consumer Std. Dev.	0.713	0.901	0.579	0.689	0.690	0.533	0.686	0.659
Intra-Consumer Std. Dev.	0.344	0.297	0.173	0.347	0.345	0.302	0.402	0.217

<u>Individual Trips</u>

The distribution of trips in the FMS data is used to generate users' trips in the synthetic population. Each individual in the synthetic population is matched with a random survey respondent from the FMS data (sampling with replacement). The daily trips of the simulated individual are also drawn with replacement from those of the corresponding FMS respondent.

The contexts of the simulated trips (origins, destinations, and departure times) are obtained from those of the real trips. Tokens are generated for each alternative by calculating the expected energy savings using TripEnergy, and the user optimization algorithm in Song (2018) and Peaks (2018).

<u>Tripod Usage</u>

The choice whether to request a Tripod menu or not for each trip is simulated using the FMS respondents' response to the post-survey question: "*I would use Tripod if it were available today*". The response variable belonging to the Likert scale "*very unlikely, unlikely, somewhat likely, likely, and very likely*" is used as the dependent variable in an ordered logit model that models the likelihood of using Tripod as a function of socio-demographic characteristics.

The model is then applied to the population of synthetic users to predict their probability of using Tripod before each trip. This ensures that individuals might perform multiple trips in a given day, and that the number of daily trips and Tripod requests vary among users in the simulation.

The simulation framework is shown in Figure 7.

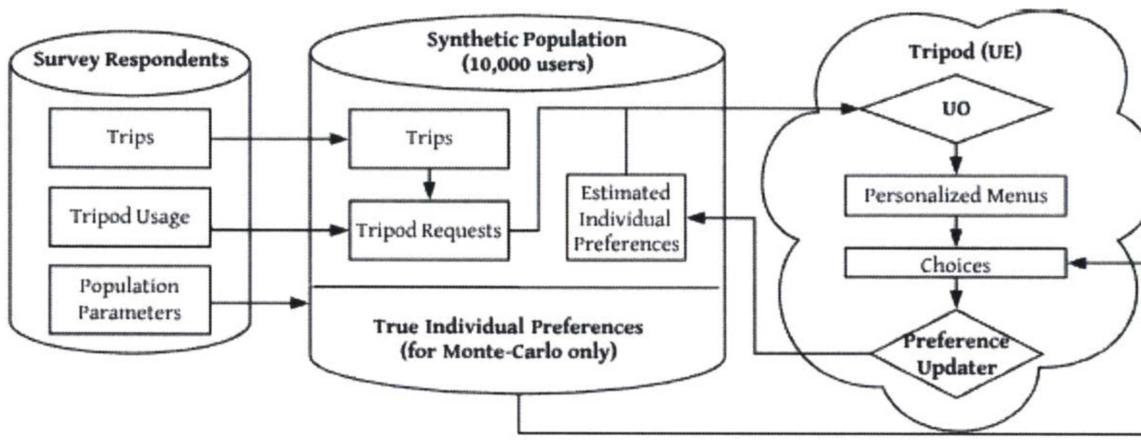

Figure 7: Tripod simulation framework.

5.5.3. Simulation Results

Base Case

In the base case, the sample size is 10,000 users. Two weeks are simulated and offline estimation is performed every fifth day. As the simulation evolves, the average hit-rate by day is observed. This rate is defined as the percentage of cases in which one of the recommended alternatives is chosen.

Using the starting values presented in Table 20 as initial population parameters, online estimation is performed after each choice on days 1-4. The first offline estimation is performed on day 5. After that, the online estimations on days 6-9 use the draws of μ, Ω^b, and Ω^w obtained from the first offline estimation. The second offline estimation is performed on day 10, and the population parameters are updated again and used in the online estimations on days 11-14.

Before observing any choices, recommendations are generated using the unconditional η_{mn} draws, which are obtained from the distributions shown in equations (5.2) and (5.3) using the starting values of μ, Ω^b, and Ω^w. These are non-personalized recommendations because they are not conditional on the users' previous choices.

The average hit-rate of the 14 simulated days is presented in Figure 8. The hit-rate is generally increasing from 76.8% to 82.4% because we are learning the individual preferences. However, it is not monotonically increasing due to intra-day variability in trips and simulation randomness. The increasing trend is observed on most days, and not only when offline estimation is performed. The simulation results are also benchmarked against a simpler method that does not apply online estimation, and uses a behavioral model with only inter-consumer heterogeneity. The average hit-rate of this method is also shown in Figure 8.

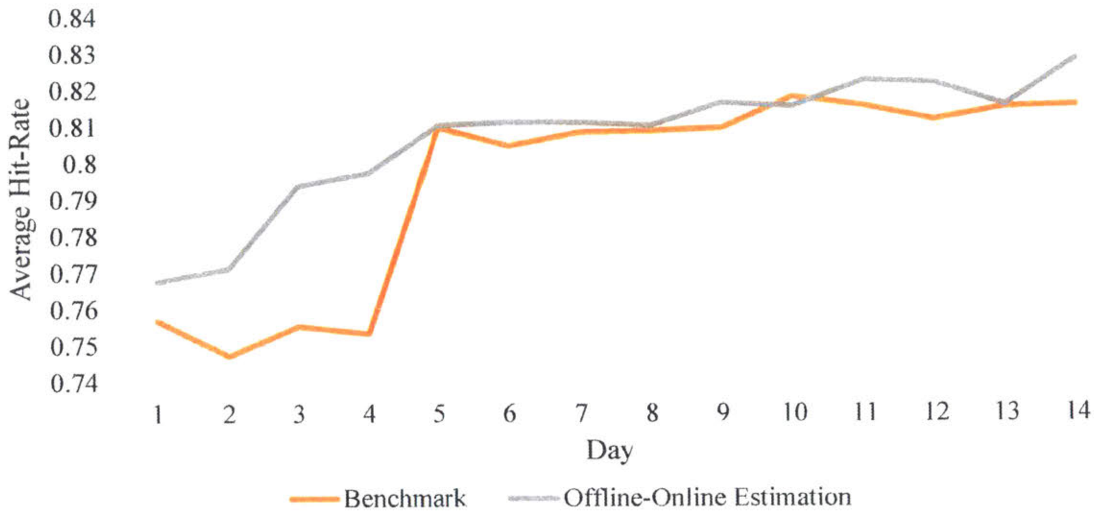

Figure 8: Average hit-rate over two weeks of simulation.

In order to eliminate the effects of day-to-day variability, we calculate the average hit-rate on a randomly generated day (i.e. testing data) using the parameters estimated on days 1-14. For this test day, we can calculate an upper bound on the expected average hit-rate, by assuming that the true individual means (ζ_n) and the true intra-consumer covariance matrix (Ω^w) are known, and generating draws from the true individual-specific distributions $\eta_{mn} \sim N(\zeta_n, \Omega^w)$ to be used in recommendations. This achieves an average hit-rate of 84.2%.

The evolution of the average hit-rate on the test day is presented in Figure 9, showing the average hit-rate by day, the approximate upper bound, and the benchmark method hit-rate. As expected, the hit-rate is consistently and almost monotonically increasing, indicating that we learn individual preferences from repeated choices. Substantial increases are observed in the first few menus, and the learning rate diminishes afterwards. We can also observe that for the benchmark method, the average hit-rate only increases when offline estimation is performed (as preferences are not updated otherwise, which is similar to the result obtained in Section 2.5.3).

The population parameters are also estimated accurately in the offline estimations on days 5 and 10. These are presented in Table 21. Since we use all the data in estimation, the results are close to their true values and no endogeneity bias is observed.

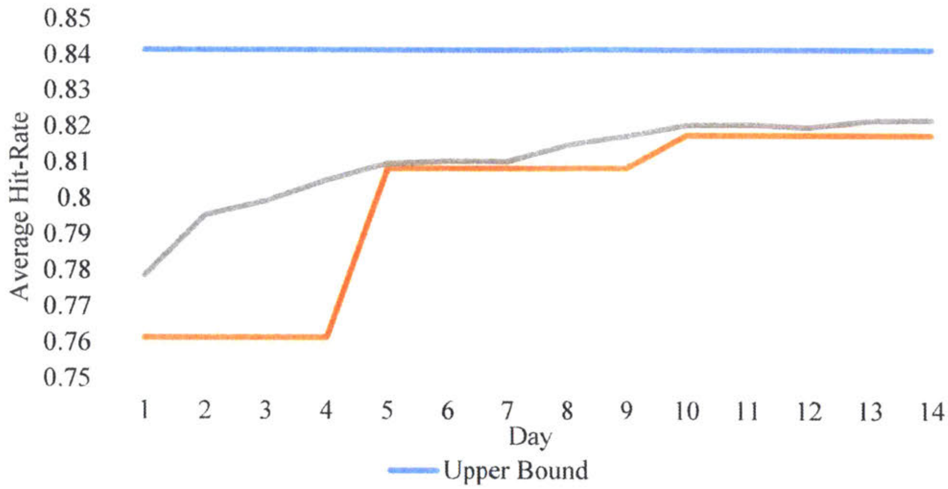

Figure 9: Evolution of the average hit-rate in the base case on the test day.

Table 21: Offline estimation results on days 5 and 10.

True Values (SP estimates)								
	Tokens	IVTT	OVTT	Walking	Car	On-Demand	Transit	Scale
Population Mean	-0.900	0.836	1.614	1.236	-0.749	-0.886	-0.658	0.355
Inter-Consumer Std. Dev.	0.713	0.901	0.579	0.689	0.690	0.533	0.686	0.659
Intra-Consumer Std. Dev.	0.344	0.297	0.173	0.347	0.345	0.302	0.402	0.217
Offline (Day 5)								
	Tokens	IVTT	OVTT	Walking	Car	On-Demand	Transit	Scale
Population Mean	-0.949 (0.036)	0.835 (0.045)	1.620 (0.022)	1.227 (0.034)	-0.738 (0.035)	-0.946 (0.090)	-0.664 (0.032)	0.374 (0.024)
Inter-Consumer Std. Dev.	0.738 (0.102)	0.836 (0.050)	0.590 (0.020)	0.751 (0.086)	0.691 (0.060)	0.589 (0.067)	0.710 (0.054)	0.648 (0.030)
Intra-Consumer Std. Dev.	0.320 (0.056)	0.286 (0.039)	0.172 (0.013)	0.472 (0.075)	0.278 (0.039)	0.422 (0.070)	0.450 (0.069)	0.245 (0.025)
Offline (Day 10)								
	Tokens	IVTT	OVTT	Walking	Car	On-Demand	Transit	Scale
Population Mean	-0.941 (0.025)	0.822 (0.029)	1.609 (0.016)	1.242 (0.027)	-0.722 (0.027)	-0.919 (0.02)	-0.648 (0.024)	0.375 (0.017)
Inter-Consumer Std. Dev.	0.786 (0.043)	0.846 (0.051)	0.597 (0.018)	0.670 (0.043)	0.710 (0.039)	0.593 (0.057)	0.703 (0.041)	0.703 (0.021)
Intra-Consumer Std. Dev.	0.282 (0.037)	0.298 (0.032)	0.187 (0.008)	0.385 (0.023)	0.413 (0.057)	0.359 (0.081)	0.439 (0.041)	0.206 (0.019)

Benefits of personalization

In order to assess the benefits of personalization, we compare the hit-rates obtained with personalized and non-personalized recommendations. The results are presented in Table 22.

On day 0, we can only calculate non-personalized hit-rates because we assume that we have not collected any choices from users yet (we use the estimated population parameters (μ, Ω^b, and Ω^w) in order to generate recommendations). On days 5 and 10, we repeat the same procedure, but with the updated population parameters following the offline estimations. Finally, we repeat the same procedure with the true population parameters shown in Table 20. The results indicate that the average hit-rate for non-personalized recommendations increases slightly from 75.9% to 77.2% even when the population parameters are perfectly learnt.

On the other hand, the average hit-rate for personalized recommendations is substantially higher, and it increases substantially as we observe more choices (in this case, up to 82% on day 10). For users with a sufficiently long choice history, we expect average hit-rate to be close to that of the true parameters, which is 84.2%.

Table 22: Personalized vs. Non-Personalized Hit-rates.

	Non-personalized	Personalized
Day 0	75.9%	75.9%*
Day 5	77.1%	81.0%
Day 10	77.2%	82.0%
True Parameters	77.2%	84.2%

*On day 0, there are no observed choices, and we start with the population distribution, so the personalized and non-personalized hit-rates are the same.

Learning Individual Distributions

In order to illustrate how individual-specific distributions are learnt, we visualize the estimated distributions of one user in the synthetic population. These are estimated using online estimation, with the population parameters obtained from offline estimation on day 10. They are then compared to:

- the true distributions of this individual, obtained by generating draws of η_{mn} using the true values of ζ_n of this user and the true intra-consumer covariance matrix, Ω^w; and
- the population distributions, obtained by sampling from the distributions shown in equations 5.2 and 5.3 respectively.

The results are presented in Figure 10, showing the distributions of the in-vehicle travel time and the tokens parameters. The results show that the estimated distributions fall between the true distributions of this user and the population distributions. This is expected because we start with the population distributions and gradually learn the true individual distributions as we observe more choices from this user. This does not hold for every user and every parameter in the data, but it is generally observed in many cases.

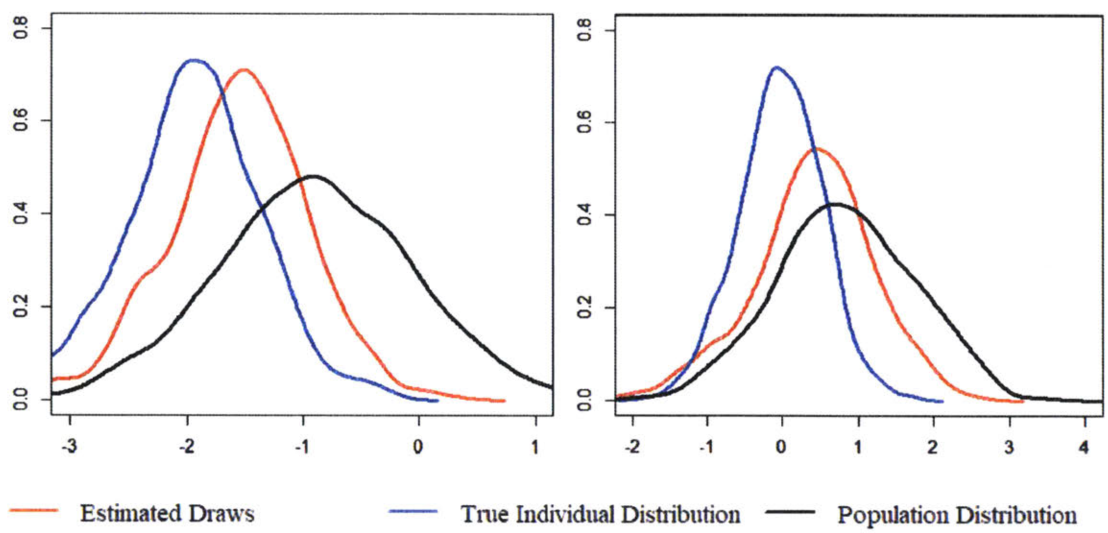

Figure 10: True and estimated distributions of the travel time (left) and walking (right) parameters.

<u>Frequency of Offline Estimation</u>

In this section, the base case is modified by changing the frequency of the offline estimation. We consider three scenarios:

- Scenario 1: offline estimation is performed every 5 days (base case).
- Scenario 2: offline estimation is performed more frequently (every 2 days)
- Scenario 3: offline estimation is performed less frequently (every 10 days).

Ten days are simulated starting with the same population parameters as before. The results are presented in Figure 11, showing the average hit-rate by day for the three scenarios. The results indicate that all three scenarios achieve similar hit-rates for the entire simulation period. This is because the starting values of the population level parameters are close to their true values. Therefore, offline estimation is not needed in this case.

However, offline estimation can be beneficial if the system is initialized with bad population parameters. To illustrate this, the starting values of the population parameters of IVTT and OVTT (including population means and inter- and intra-consumer variances) are halved. The estimation results are shown in Figure 12. These results show that Scenario 2 achieves the highest hit-rate on days 2-5 since the population parameters are updated quickly. On the other hand, Scenario 3 results in the worst hit-rate on days 2-9, that is before the offline estimation. All three scenarios achieve similar hit-rates by day 10.

This illustrates the importance of the offline estimation in the beginning of operation, where population parameters might not be accurately estimated. However, after this initial phase, assuming that the population parameters do not vary over time, frequent offline estimations are not needed.

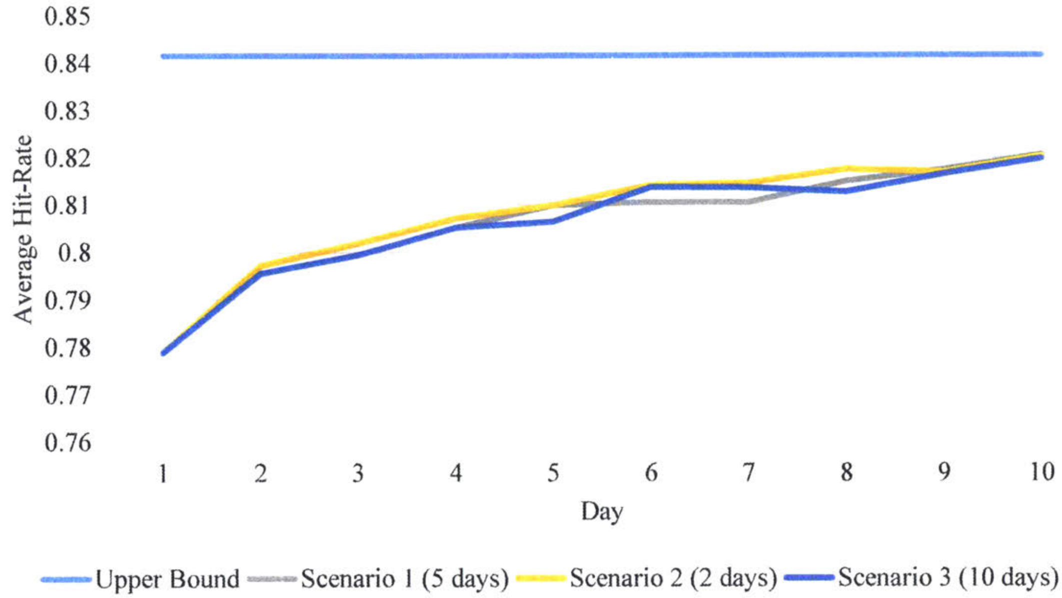

Figure 11: Simulation results with varied frequency of offline estimation.

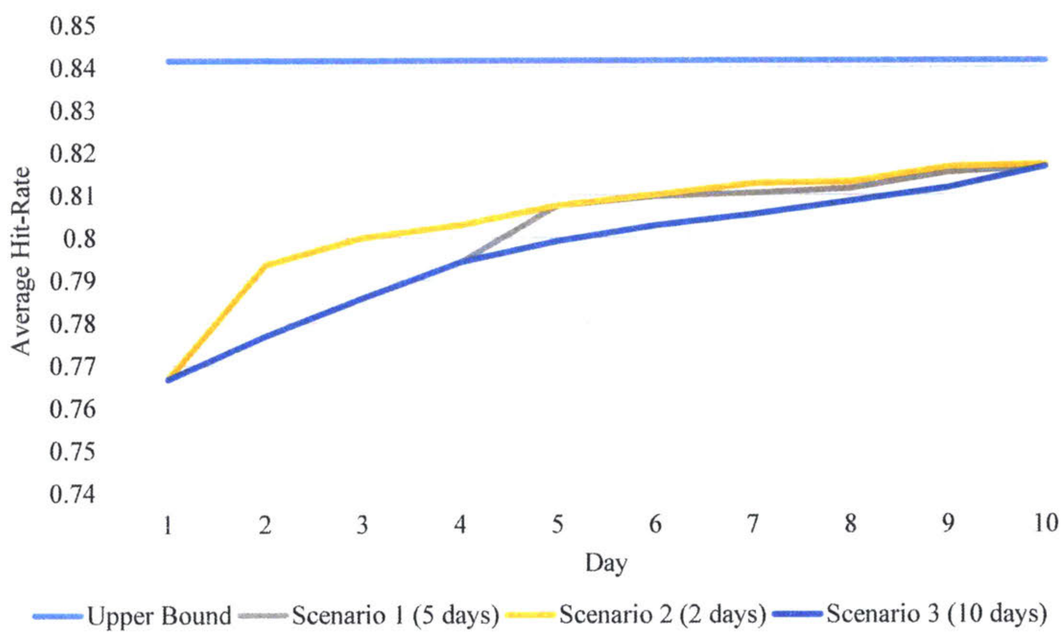

Figure 12: Simulation results with varied frequency of offline estimation and bad starting values.

Figure 12 shows that in Scenario 3, even though we do not perform the offline estimation until day 10, we still obtain results that are close to those obtained in Scenarios 1 and 2. This is because the frequency of trip requests is high (about 1 trip per person per day), and therefore, we are able to learn preferences from the online estimation. Since the misspecified population parameters have the same effect as priors in this estimation, their effect diminishes as we observe more choices from each individual.

This might not be the case if the data is sparse (i.e. request frequency is low). To demonstrate this, the base case is modified by reducing the frequency of Tripod requests to one third of its original value. The starting values of the population parameters are misspecified as before, and three scenarios are compared:

- Scenario 1: offline estimation is performed every 5 days (base case).
- Scenario 2: offline estimation is performed every 10 days.
- Scenario 3: offline estimation is not performed at all.

Figure 13 shows the simulation results for a period of four weeks. We can clearly observe a slower learning rate compared to figures 11 and 12 as expected. The average hit-rates on day 10 is approximately 1.7% lower than those observed in Figure 12. The results also show larger improvements in hit-rates after offline estimation (on day 5 for Scenario 1, and day 10 for scenario 2). The average hit-rates of Scenario 3 are consistently lower than those of Scenarios 1 and 2, and only approache them after 24 days. Therefore, when we have a small number of observations per individual, improving the population parameters (i.e. using the offline estimation) results in substantial improvements in recommendations.

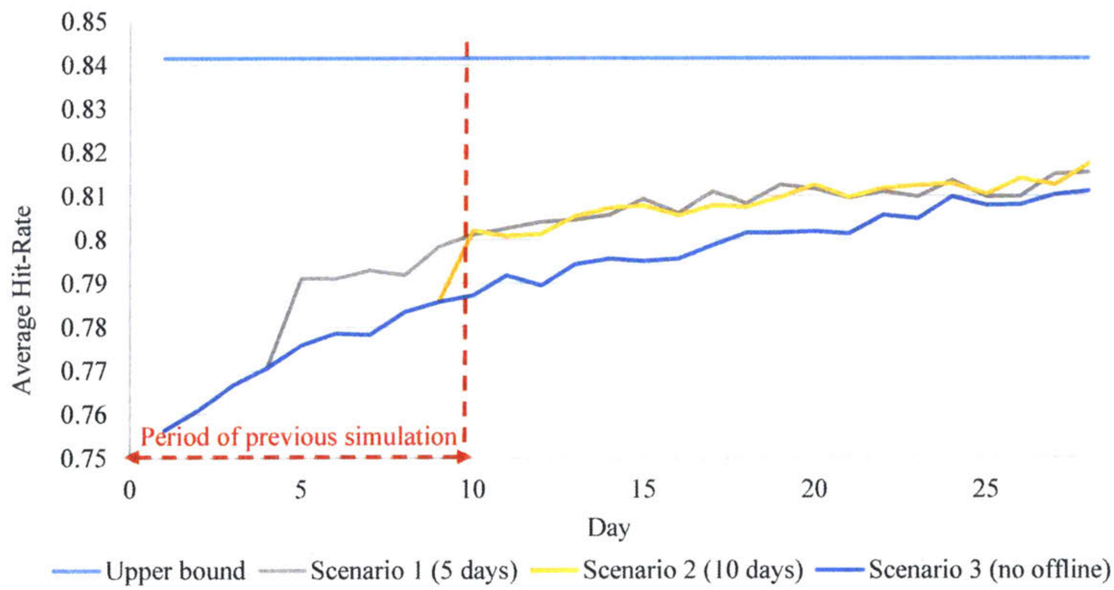

Figure 13: Simulation results with bad starting values and sparse data.

5.5.4. Varying Levels of Heterogeneity

In this section, the levels of inter- and intra-consumer heterogeneity in the synthetic population are varied, and the effects on the average hit-rate are analyzed. Four different cases are considered separately:

1. High inter-consumer heterogeneity: double the variances of inter-consumer heterogeneity in the base case.
2. Low inter-consumer heterogeneity: half the variances of inter-consumer heterogeneity in the base case.
3. High intra-consumer heterogeneity: double the variances of intra-consumer heterogeneity in the base case.
4. Low intra-consumer heterogeneity: half the variances of intra-consumer heterogeneity in the base case.

In each of these cases, the variances of all parameters are changed except for the scale parameter. The results of the four simulations, in addition to the base case, are shown in Figures 13 and 14.

These results show that with high inter-consumer heterogeneity, the learning rate is higher as indicated by the steep slope of the average hit-rate in Figure 14. On the other hand, low inter-consumer heterogeneity results in a slow learning rate, and the plot is almost flat. This result is expected; in the extreme case preferences do not vary across users, we do not expect to learn at all from the users' previous choices (as all users are identical, and perturbations are only due to intra-consumer heterogeneity).

We also notice that in the beginning of the simulation, the case with low inter-consumer heterogeneity achieves better hit-rates than that with high inter-consumer heterogeneity. This is also expected, because in the beginning of the simulation, the recommendations are non-personalized (i.e. not conditional on previous choices).

On the other hand, Figure 15 shows that higher levels of intra-consumer heterogeneity result in lower hit-rates. This trend is consistent throughout the simulation. This result is also expected; we achieve higher hit-rates if consumers' preferences are stable over time and do not very considerably across choices.

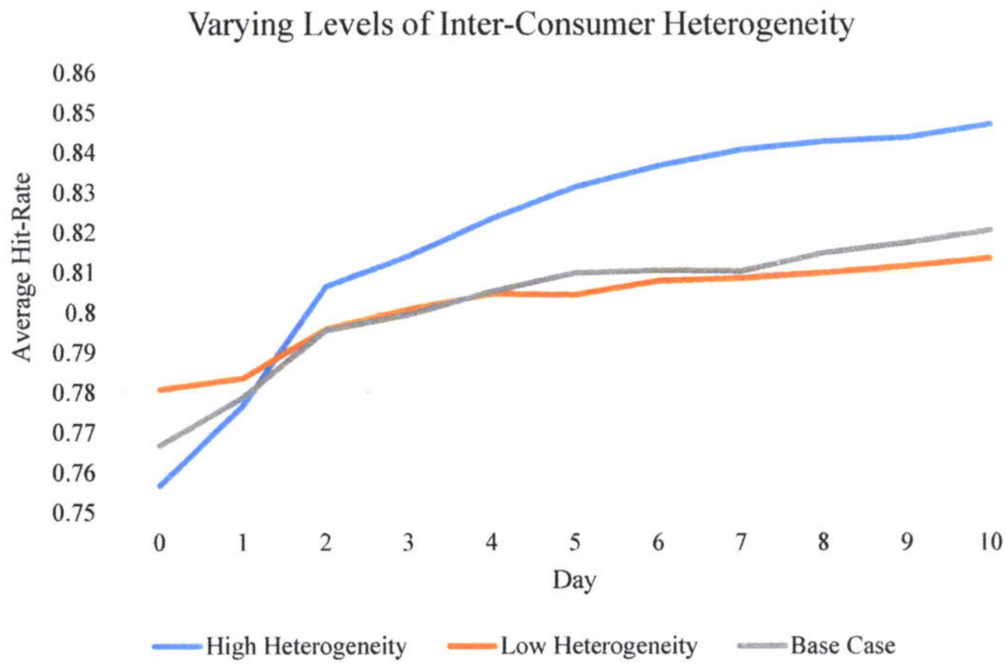

Figure 14: Simulation results with varying levels of inter-consumer heterogeneity.

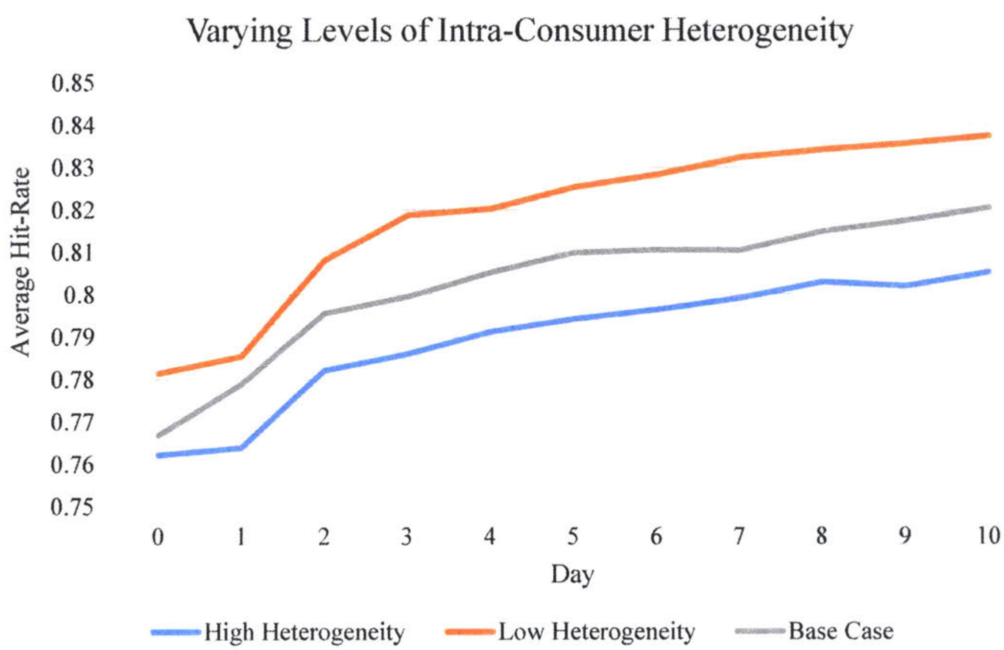

Figure 15: Simulation results with varying levels of intra-consumer heterogeneity.

5.6. CONCLUSION

This chapter presented an application of the offline-online estimation framework to the personalized travel advisor, Tripod. RP and SP data were collected in GBA and used to estimate a discrete choice model with inter- and intra-consumer heterogeneity. The estimates were used in the online preference updater of the Tripod smartphone app.

This chapter also presented a simulation of Tripod using a synthetic population of 10,000 users. The results show that the average hit-rate increases as more choices are observed, and approaches the approximate upper bound after 14 simulated days. In addition, the benefit of the offline estimation was demonstrated by varying the frequency of these estimations in the simulation and the starting values of the population parameters. The results show that offline estimation is needed when the population parameters are poorly estimated, for example, when the system is poorly initialized.

One of the main limitations was the small sample size used in estimating the initial model from the SP data, and therefore, a parsimonious model was used. Once more data becomes available, the behavioral model can be further developed to include coefficients for other variables such as schedule delay, carpooling, energy savings, and eco-friendly driving in addition to various socio-economic and demographic variables. In addition, the population parameters obtained from the SP data can be replaced by those obtained from the offline estimation.

Future research should focus on analyzing trends and structural changes in the model parameters. For example, the transit parameter might increase or decrease over time due to improvements in the service quality, or due to seasonal variations (such as weather). While the results in Section 5.5.3 show that offline estimation might not be needed frequently after the first few days of operation, it can be necessary in cases where preferences change over time (as illustrated in Section 2.5.3).

In addition, we can also apply the framework presented in Chapter 4 to model random and systematic heterogeneity in order to represent structural changes explicitly. The regression equations of the individual or menu specific parameters might include time variables to account for trends or trigonometric variables to account for seasonality.

While this simulation assumed that the model parameters are unknown, the true model specification was assumed to be known. In order to adequately assess the performance of the offline-online personalization framework, future research should use real data obtained from Tripod users. This would also allow us to objectively compare this recommendation methodology to other methods, such as content-based recommendations and other state-of-the-practice methods.

The simulation also assumed that the request frequency is exogenous and independent of the quality of recommendations. This might not be realistic, as users are more likely to return to the app if they had found it useful in their last usage. Ongoing research is focused on modeling the app-usage frequency as a function of the expected maximum utility of the recommended menus using the framework proposed by Xie et al. (2019).

CHAPTER 6: CONCLUSION

6.1. SUMMARY OF CONTRIBUTIONS

The Random Utility Model (RUM) has been widely applied in various fields such as transportation, marketing, e-commerce, and others. However, due to computational limitations, most of its applications were offline. This dissertation introduced a method for estimating and updating individual-specific preferences online using discrete choice models. The Hierarchical Bayes (HB) estimator proposed by Becker et al. (2018) was used to estimate logit mixture models with inter- and intra- consumer heterogeneity, which allows for estimating individual- and population level parameters (offline). Building on this estimator, an online estimator was proposed to update individual-specific parameters in real-time (i.e. after each choice). This approach is the first to apply logit mixture models in online applications, while accounting for complex patterns of heterogeneity.

The online estimation procedure is computationally efficient and as accurate as offline estimation as demonstrated using real and Monte Carlo data. This technique enables discrete choice models to be applied in real-time decision support systems (such as recommender systems). The resulting recommendations make use of all available data including user-specific characteristics and preferences, alternative-specific attributes, and contextual data.

This framework is implemented in the app-based travel advisor Tripod, which aims to incentivize and shift travelers' behavior towards more sustainable alternatives (e.g. by changing their mode, route, or departure time choice). The smartphone app offers users personalized menus of alternatives with incentives and information associated with each alternative.

Chapter 3 addressed a main limitation of adaptive choice contexts, which is endogeneity. In choice-based recommender systems, the alternatives are generated using individual-specific parameters, which are estimated using the users' previous choices. Traditional correction methods for endogeneity (e.g., control-functions) are difficult to apply in these cases since all attributes are endogenous and finding relevant instruments seems infeasible. The literature has not yet addressed endogeneity bias in choice-based recommender systems, however, different Monte Carlo experiments were developed for a similar application; adaptive stated preferences (ASP) surveys. Some of these studies found significant endogeneity bias (deviation from the true values), while others did not.

Chapter 3 built on the findings of Liu et al. (2007), and identified cases where endogeneity can be ignorable and cases where it cannot. The results indicated that including all the data in estimation results in consistent estimates if the initialization is exogenous. These findings have implications on the design and analysis of ASP surveys and choice-based recommender systems. Researchers and practitioners should make sure that the likelihood function accounts for the data generation process, which is achieved by conditioning on all the previous choices. It is also important that the system is initialized exogenously, and that this initialization is accounted for in estimation.

Chapter 4 presented methodological extensions to the logit mixture model with inter- and intra-consumer heterogeneity. The first extension entails random and systematic inter- and intra-consumer heterogeneity, by incorporating socio-demographic data and contextual variables to model the individual- and menu-specific preferences respectively.

In the context of recommender systems, accounting for systematic heterogeneity has three advantages. First, accounting for systematic inter-consumer heterogeneity partially mitigates the "cold start" problem, where the number of initially available choices is relatively small. For example, we can only generate personalized recommendations to a given user if we have his/her choice history. Using the proposed model, this problem can be mitigated by expressing the user's preferences as a function of his/her characteristics and socio-economic variables. Second, accounting for systematic intra-consumer heterogeneity enables us to provide context-aware recommendations, and thus make use of all the available data. For example, an online e-commerce platform could recommend clothes to customers using contextual data, such as the season and their location (Aggarwal, 2016). Finally, the online estimator proposed in Chapter 2 can be applied to this model, and is feasible in real-time as it only uses the data of the individual making the choice.

The second extension incorporated a latent class model into the logit double mixture model, which results in more flexibility in the modeling the mixing distributions, and allows different individuals to have different intra-consumer covariance matrices. The proposed model is easily interpretable, as we can distinguish between different classes in the population; each class has its own population means and inter- and intra-consumer covariance matrices. An HB estimator was proposed to estimate this model and validated using Monte Carlo data.

Chapter 5 presented an application of the offline-online estimation methodology to the app-based travel advisor, Tripod, which incentivizes users towards more sustainable travel behavior. Context-aware SP data were collected using the state-of-the-art smartphone based platform, FMS. The data were used to estimate a discrete choice model with inter- and intra-consumer heterogeneity. The estimation results were then used in a simulation which demonstrated that preferences are learnt from repeated choices and hit-rates improve over time.

6.2. FUTURE RESEARCH DIRECTIONS

The online discrete choice methodology presented in this dissertation can be applied primarily in recommender systems. In this context, this methodology can be further enhanced using "exploration" and "exploitation". Such applications were proposed by Teo et al. (2016) and Song (2018) in the context of choice-based recommender systems, who used Multi-armed Bandit techniques such as Thompson sampling (Thompson, 1933) to accelerate the learning process. Exploration can be particularly useful when the data is sparse (i.e., the number of choices per individual is small).

In the literature on choice-based recommender systems, different methods were proposed to introduce diversity in the recommended menus. These methods include the "submodular optimization problem" proposed by Teo et al. (2016), and the multi-level nesting approach proposed by Jiang et al. (2014). Future research can extend this methodology to achieve better diversity in recommendations.

While online estimation is fast and computationally efficient, offline estimation is computationally expensive especially as the number of individuals and observations per individual increases. Future research should focus on developing sequential Bayesian methods in which the posteriors obtained from past estimations can be used as priors (while accounting for endogeneity). Although Chapter 3 demonstrated that endogeneity is ignorable when the likelihood function is correctly specified (i.e. when it accounts for the data generation process), this might not be feasible in some applications, as it requires using all the available data in estimation. Future research should develop remedies for endogeneity which allow us to obtain consistent estimates without the need to use the entire data.

Alternatively, parallel MCMC algorithms (such as the one proposed by Neiswanger et al. (2014)) can be developed to account for correlated/panel data. In the latter algorithm, data is partitioned into different batches, which are then processed independently (with very little communication) on multiple machines (using any classical MCMC method such as Gibbs sampling). Samples from the sub-posteriors are then combined to form samples from the full posterior. This algorithm can still be used in offline estimation if the data is partitioned cross-sectionally (i.e. observations of the same individual belong to the same batch). Future research should develop similar parallel MCMC algorithms in which the data can be partitioned across different time intervals rather than individuals. This allows sub-posteriors from sequential estimations to be combined efficiently and accurately, avoiding the need to include all the historical data in a single estimation.

Future research should also focus on developing other semi-parametric and non-parametric mixing distributions that can be efficiently applied in online settings. For example, the latent class model described in Chapter 4 can be extended to a "Chinese Restaurant Process", in which the number of classes is not determined beforehand, but updated in each step of the Gibbs sampler (Burda et al., 2008). In addition, the class membership model can be enriched by using covariates instead of assuming constant probabilities.

Finally, the estimators of the behavioral models can be enhanced in several ways. For example, more efficient Metropolis-Hastings algorithms can be developed depending on the application. Alternatively, variational Bayesian inference (Krueger et al., 2019) can be used in offline estimation instead of HB to obtain draws from the population parameters.

REFERENCES

Abernethy, J., Evgeniou, T., Toubia, O., & Vert, J. P. (2008). Eliciting consumer preferences using robust adaptive choice questionnaires. IEEE Transactions on Knowledge and Data Engineering, 20(2), 145-155.

Adomavicius, G., Huang, Z., & Tuzhilin, A. (2008). Personalization and recommender systems. State-of-the-art decision making tools in the information-intensive age, 55-100.

Adomavicius, G., & Tuzhilin, A. (2011). Context-aware recommender systems. In Recommender systems handbook (pp. 217-253). Springer, Boston, MA.

Aggarwal, C. C. (2016). Recommender systems. Springer International Publishing.

Akinc, D., & Vandebroek, M. (2018). Bayesian estimation of mixed logit models: Selecting an appropriate prior for the covariance matrix. Journal of choice modelling, 29, 133-151.

Allenby, G. (1997). An introduction to hierarchical Bayesian modeling. In Tutorial notes, Advanced Research Techniques Forum, American Marketing Association.

Allenby, G. M., & Rossi, P. E. (1998). Marketing models of consumer heterogeneity. Journal of econometrics, 89(1-2), 57-78.

Allenby, G. M., & Rossi, P. E. (2006). Hierarchical Bayes models. The handbook of marketing research: Uses, misuses, and future advances, 418-440.

Alvarez, I., Niemi, J., & Simpson, M. (2014). Bayesian inference for a covariance matrix. arXiv preprint arXiv:1408.4050.

Andrews, R. L., Ainslie, A., & Currim, I. S. (2002). An empirical comparison of logit choice models with discrete versus continuous representations of heterogeneity. Journal of Marketing Research, 39(4), 479-487.

Ansari, A., Essegaier, S., and Kohli, R. (2000). Internet recommendation systems. Journal of Marketing Research, 37(3), 363 – 375.

Azevedo, C. L., Seshadri, R., Gao, S., Atasoy, B., Akkinepally, A. P., Christofa, E., Zhao, F., Trancik, J., & Ben-Akiva, M. (2018). Tripod: Sustainable Travel Incentives with Prediction, Optimization, and Personalization. Transportation Research Board Annual Meeting, Washington, D.C.

Bajari, P., Fox, J. T., & Ryan, S. P. (2007). Linear regression estimation of discrete choice models with nonparametric distributions of random coefficients. American Economic Review, 97(2), 459-463.

Becker, F., Danaf, M., Song, X., Atasoy, B., & Ben-Akiva, M. (2018). Hierarchical Bayes estimator of a logit mixture with inter- and intra-consumer heterogeneity. Transportation Research Part B: Methodological 117, 1-17.

Ben-Akiva, M., & Bierlaire, M. (2003). Discrete choice models with applications to departure time and route choice. In Handbook of transportation science (pp. 7-37). Springer, Boston, MA.

Ben-Akiva, M., & Lerman, S. R. (1985). Discrete choice analysis: theory and application to travel demand (Vol. 9). MIT press.

Ben-Akiva, M., McFadden, D., & Train, K. (2019). Foundations of stated preference elicitation: Consumer behavior and choice-based conjoint analysis. Foundations and Trends in Econometrics, 10(1-2), 1-144.

Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile Prices in Market Equilibrium. Econometrica: Journal of the Econometric Society, 841-890.

Bhat, C. R. (2000). Incorporating observed and unobserved heterogeneity in urban work travel mode choice modeling. Transportation science, 34(2), 228-238.

Bhat, C. R. (2011). The maximum approximate composite marginal likelihood (MACML) estimation of multinomial probit-based unordered response choice models. Transportation Research Part B: Methodological, 45(7), 923-939.

Bhat, C. R., & Castelar, S. (2002). A unified mixed logit framework for modeling revealed and stated preferences: formulation and application to congestion pricing analysis in the San Francisco Bay area. Transportation Research Part B: Methodological, 36(7), 593-616.

Bhat, C. R., & Sardesai, R. (2006). The impact of stop-making and travel time reliability on commute mode choice. Transportation Research Part B: Methodological, 40(9), 709-730.

Bhat, C. R., & Sidharthan, R. (2011). A simulation evaluation of the maximum approximate composite marginal likelihood (MACML) estimator for mixed multinomial probit models. Transportation Research Part B: Methodological, 45(7), 940-953.

Bierlaire, M., Axhausen, K. & Abay, G. (2001). Acceptance of modal innovation: the case of the Swissmetro, Proceedings of the 1st Swiss Transportation Research Conference, Ascona, Switzerland.

Blundell, R. W., & Powell, J. L. (2004). Endogeneity in semiparametric binary response models. The Review of Economic Studies, 71(3), 655-679.

Bolis, S., & Maggi, R. (2003). Logistics strategy and transport service choices: an adaptive stated preference experiment. Growth and Change, 34(4), 490-504.

Bozdogan, H. (1987). Model selection and Akaike's information criterion (AIC): The general theory and its analytical extensions. Psychometrika, 52(3), 345-370.

Bradley, M. A., & Daly, A. J. (1993). New analysis issues in stated preference research. PTRC-PUBLICATIONS-P, 75-75.

Bujosa, A., Riera, A., & Hicks, R. L. (2010). Combining discrete and continuous representations of preference heterogeneity: a latent class approach. Environmental and Resource Economics, 47(4), 477-493.

Burda, M., Harding, M., & Hausman, J. (2008). A Bayesian mixed logit–probit model for multinomial choice. Journal of econometrics, 147(2), 232-246.

Carreras, I., Gabrielli, S., Miorandi, D., Tamilin, A., Cartolano, F., Jakob, M., & Marzorati, S. (2012). SUPERHUB: a user-centric perspective on sustainable urban mobility. In Proceedings of the 6th ACM workshop on Next generation mobile computing for dynamic personalised travel planning (pp. 9-10). ACM.

Castells, P., Hurley, N. J., & Vargas, S. (2015). Novelty and diversity in recommender systems. In Ricci, F., Rokach, L., and Shapira, B., Recommender Systems Handbook, Second Edition, Springer Science + Business Media, New York, 2015.

Chaptini, B. (2005). Use of discrete choice models with recommender systems. PhD dissertation, Department of Civil and Environmental Engineering, MIT.

Chen, C., Zhang, D., Guo, B., Ma, X., Pan, G., & Wu, Z. (2015). TripPlanner: Personalized trip planning leveraging heterogeneous crowdsourced digital footprints. IEEE Transactions on Intelligent Transportation Systems, 16(3), 1259-1273.

Cheung, K.W., Kwok, J.T., Law, M.H. & Tsui, K.C. (2003). Mining customer product ratings for personalized marketing. Decision Support Systems, 35(2), pp.231-243.

Chow, J. Y., Dmitriyeva, A., & Fay, D. (2018). City-scalable Destination Recommender System for On-demand Senior Mobility.

Danaf, M., Atasoy, B., Azevedo, C., Mastera, J., Abou-Zeid, M., Cox, N., Zhao, F., Ben-Akiva, M. (2019). Smartphone-based stated preferences methods for mobility solutions. Journal of Choice Modelling, 31, 35-50.

Danielis, R., Marcucci, E., & Rotaris, L. (2005). Logistics managers' stated preferences for freight service attributes. Transportation Research Part E: Logistics and Transportation Review, 41(3), 201-215.

Dumont, J., Keller, J., Carpenter, C., & Dumont, M. J. (2015). Package 'RSGHB'.

Elrod, T., & Keane, M. P. (1995). A factor-analytic probit model for representing the market structure in panel data. Journal of Marketing Research, 1-16.

Fosgerau, M., & Bierlaire, M. (2007). A practical test for the choice of mixing distribution in discrete choice models. Transportation Research Part B: Methodological, 41(7), 784-794.

Fowkes, A. S., & Shinghal, N. (2002). The Leeds adaptive stated preference methodology, working paper, University of Leeds.

Fowkes, T. (2007). The design and interpretation of freight stated preference experiments seeking to elicit behavioural valuations of journey attributes. Transportation Research Part B: Methodological, 41(9), 966-980.

Fox, J. T., Kim, K. I., Ryan, S. P., & Bajari, P. (2011). A simple estimator for the distribution of random coefficients. Quantitative Economics, 2(3), 381-418.

Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. Statistical science, 7(4), 457-472.

Geweke, J., & Keane, M. (2001). Computationally intensive methods for integration in econometrics. In Handbook of econometrics (Vol. 5, pp. 3463-3568). Elsevier.

Geweke, J., & Keane, M. (2007). Smoothly mixing regressions. Journal of Econometrics, 138(1), 252-290.

Goldberg, D., Nichols, D., Oki, B. M., & Terry, D. (1992). Using collaborative filtering to weave an information tapestry. Communications of the Association for Computing Machinery, 35 (12), 61–70.

Greene, W. H., & Hensher, D. A. (2013). Revealing additional dimensions of preference heterogeneity in a latent class mixed multinomial logit model. Applied Economics, 45(14), 1897-1902.

Greene, W. H. (2004). Interpreting estimated parameters and measuring individual heterogeneity in random coefficient models, Department of Economics, Stern School of Business, New York University.

Guevara, C. A. (2015). Critical assessment of five methods to correct for endogeneity in discrete-choice models. Transportation Research Part A: Policy and Practice, 82, 240-254.

Guevara, C. A., & Ben-Akiva, M. (2006). Endogeneity in residential location choice models. Transportation Research Record, 1977(1), 60-66.

Guevara, C. A., & Ben-Akiva, M. E. (2012). Change of scale and forecasting with the control-function method in logit models. Transportation Science, 46(3), 425-437.

Guevara, C. A., & Polanco, D. (2016). Correcting for endogeneity due to omitted attributes in discrete-choice models: the multiple indicator solution. Transportmetrica A: Transport Science, 12(5), 458-478.

Gunawardana, A., & Shani, G. (2015). Evaluating recommender systems. In Ricci, F., Rokach, L., and Shapira, B., Recommender Systems Handbook, Second Edition, Springer Science + Business Media, New York, 2015.

Guo, S., & Sanner, S. (2010). Real-time multiattribute Bayesian preference elicitation with pairwise comparison queries. In Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (pp. 289-296).

Hagen, P., 1999. Smart personalization. Forrester Report, Cambridge, MA, 1–18.

Hauser, J. R., & Toubia, O. (2005). The impact of utility balance and endogeneity in conjoint analysis. Marketing Science, 24(3), 498-507.

Hausman, J. A. (1978). Specification tests in econometrics. Econometrica: Journal of the econometric society, 1251-1271.

Heckman, J. J. (1977). Sample selection bias as a specification error (with an application to the estimation of labor supply functions). Center for Economic Analysis of Human Behavior Aid Social Institutions National Bureau of Economic Research.

Heckman, J.J. (1978). Dummy endogenous variables in a simultaneous equation system. Econometrica 46(4) 931–959.

Heidelberger, P., & Welch, P. D. (1983). Simulation run length control in the presence of an initial transient. Operations Research, 31(6), 1109-1144.

Hess, S. (2014). Latent class structures: taste heterogeneity and beyond. In Handbook of choice modelling, Chapter 14, 311-329. Edward Elgar Publishing Cheltenham.

Hess, S, & Train, K. (2011). Recovery of inter- and intra-consumer heterogeneity using mixed logit models. Transportation Research Part B: Methodological, 45(7), 973 – 990.

Hess, S., & Rose, J. (2009). Allowing for intra-respondent variations in coefficients estimated on repeated choice data. Transportation Research Part B: Methodological, 43(6), 708–719.

Horsky, D., Misra, S., & Nelson, P. (2006). Observed and unobserved preference heterogeneity in brand-choice models. Marketing Science, 25(4), 322-335.

Huang, S. (2011). Designing utility-based recommender systems for e-commerce: Evaluation of preference-elicitation methods. Electronic Commerce Research and Applications, 10(4), 398–407.

Huang, A., & Wand, M. P. (2013). Simple marginally noninformative prior distributions for covariance matrices. Bayesian Analysis, 8(2), 439-452.

Huber, J., & Train, K. (2001). On the similarity of classical and Bayesian estimates of individual mean partworths. Marketing Letters, 12(3), 259-269.

Jiang, H., Qi, X., & Sun, H. (2014). Choice-based recommender systems: a unified approach to achieving relevancy and diversity. Operations Research, 62(5), 973-993.

Johnson, F. R., Kanninen, B., Bingham, M., & Özdemir, S. (2006). Experimental design for stated-choice studies. In Valuing Environmental Amenities Using Stated Choice Studies (pp. 159-202). Springer, Dordrecht.

Johnson, F. R., Lancsar, E., Marshall, D., Kilambi, V., Mühlbacher, A., Regier, D. A., Bresnahan, B. W., Kanninen, B., & Bridges, J. F. (2013). Constructing experimental designs for discrete-choice experiments: report of the ISPOR conjoint analysis experimental design good research practices task force. Value in Health, 16(1), 3-13.

Kanninen, B. J. (2002). Optimal design for multinomial choice experiments. Journal of Marketing Research, 39(2), 214-227.

Keane, M. and Wasi, N. (2013). Comparing alternative models of heterogeneity in consumer choice behavior. Journal of Applied Econometrics, 28(6):1018–1045.

Krueger, R., Bansal, P., Bierlaire, M., Daziano, R. A., & Rashidi, T. H. (2019). Variational Bayesian Inference for Mixed Logit Models with Unobserved Inter-and Intra-Individual Heterogeneity. arXiv preprint arXiv:1905.00419.

Krueger, R., Vij, A., & Rashidi, T. H. (2018). A Dirichlet Process Mixture Model of Discrete Choice. arXiv preprint arXiv:1801.06296.

Li, L., Chu, W., Langford, J., & Schapire, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. In Proceedings of the 19th International Conference on World Wide Web (pp. 661-670). ACM.

Liu, Q., Otter, T. & Allenby, G.M. (2007). Investigating endogeneity bias in marketing. Marketing Science, 26(5), 642-650.

Louviere, J., Train, K., Ben-Akiva, M., Bhat, C., Brownstone, D., Cameron, T. A., ... & Hensher, D. (2005). Recent progress on endogeneity in choice modeling. Marketing Letters, 16(3-4), 255-265.

Needell, Z.A., McNerney, J., Chang, M.T., & Trancik, J.E. (2016). Potential for widespread electrification of personal vehicle travel in the United States. Nature Energy, 1, 16112.

Neiswanger, W., Wang, C., & Xing, E. (2013). Asymptotically exact, embarrassingly parallel MCMC. arXiv preprint arXiv:1311.4780.

Nuzzolo, A., Comi, A., Crisalli, U., & Rosati, L. (2013, October). An advanced pre-trip planner with personalized information on transit networks with ATIS. In Intelligent Transportation Systems-(ITSC), 2013 16th International IEEE Conference on (pp. 2146-2151). IEEE.

Orme, B., & Baker, G. (2000). Comparing Hierarchical Bayes Draws and Randomized First Choice for Conjoint Simulations. Sawtooth Software Research Paper Series.

Peaks, N. (2018). Smartphone interface for incentivized energy optimization system: FMS Advisor (Doctoral dissertation, Massachusetts Institute of Technology).

Petrin, A., & Train, K. (2010). A control function approach to endogeneity in consumer choice models. Journal of marketing research, 47(1), 3-13.

Polydoropoulou, A., & Lambrou, M. (2012). Development of an e-learning recommender system using discrete choice models and Bayesian theory: a pilot case in the shipping industry. In Security Enhanced Applications for Information Systems, ISBN: 978-953-51-0643.

RA Fisher, M. A. (1922). On the mathematical foundations of theoretical statistics. Phil. Trans. R. Soc. Lond. A, 222(594-604), 309-368.

Revelt, D. & Train, K. (1999). Mixed logit with repeated choices. Review of Economics and Statistics 80, 647–657.

Revelt, D., & Train, K. (2000). Customer-specific taste parameters and mixed logit: Households' choice of electricity supplier. Department of Economics, UCB.

Ricci, F., Rokach, L., & Shapira, B. (2015). Recommender systems: introduction and challenges. In Ricci, F., Rokach, L., and Shapira, B., Recommender Systems Handbook, Second Edition, Springer Science + Business Media, New York, 2015.

Richardson, A. (2002). Simulation study of estimation of individual specific values of time by using adaptive stated-preference survey. Transportation Research Record: Journal of the Transportation Research Board, (1804), 117-125.

Rossi, P., Allenby, G. and R. McCulloch (2005), Bayesian Statistics and Marketing, John Wiley and Sons, Hoboken, N.J.

Rubin, T., & Steyvers, M. (2009). A topic model for movie choices and ratings. In Proceedings of the 9th International Conference on Cognitive Modeling, Manchester.

Sarrias, M., & Daziano, R. A. (2017). Multinomial logit models with continuous and discrete individual heterogeneity in R: the gmnl package. Journal of Statistical Software, 79(2), 1-46.

Sawtooth Software (1999), The CBC/HB Technical Paper: http://www.sawtoothsoftware.com/TechPap.htm.

Sawtooth Software (2003). ACA/Hierarchical Bayes v2.0, Technical paper, http://www.sawtoothsoftware.com/technicaldownloads.

Sawtooth Software (2009). The CBC/HB system for hierarchical Bayes estimation version 5.0 Technical Paper. Technical Paper Series, Sawtooth Software, Orem, UT.

Schwarz, G. (1978). Estimating the dimension of a model. The annals of statistics, 6(2), 461-464.

Shinghal, N. (1999). An application of stated preference methods to the study of intermodal freight transport services in India (Doctoral dissertation, University of Leeds).

Shinghal, N., & Fowkes, T. (2002). Freight mode choice and adaptive stated preferences. Transportation Research Part E: Logistics and Transportation Review, 38(5), 367-378.

Small, K. A., Winston, C., & Yan, J. (2005). Uncovering the distribution of motorists' preferences for travel time and reliability. Econometrica, 73(4), 1367-1382.

Song, X. (2016). A Bayesian bandit approach to personalized online coupon recommendation. Master's Thesis, MIT.

Song, X. (2018). Personalization of Future Urban Mobility (Doctoral dissertation, Massachusetts Institute of Technology).

Song, X., Atasoy, B., & Ben-Akiva, M. (2017). Smart Mobility through Personalized Menu Optimization. Transportation Research Board Annual Meeting, Washington, DC, USA.

Song, X., Becker, F., Danaf, M., Atasoy, B., & Ben-Akiva, M. (2019). Enhancement of Hierarchical Bayes for Logit Mixture. Working paper, MIT.

Song, X., Danaf, M., Atasoy, B., & Ben-Akiva, M. (2018). Personalized menu optimization with preference updater: a Boston case study. Transportation Research Record, Journal of the Transportation Research Board.

Teo, C. H., Nassif, H., Hill, D., Srinivasan, S., Goodman, M., Mohan, V., & Vishwanathan, S. V. N. (2016). Adaptive, personalized diversity for visual discovery. In Proceedings of the 10th ACM Conference on Recommender Systems (pp. 35-38). ACM.

Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika, 25(3/4), 285-294.

Tilahun, N. Y., Levinson, D. M., & Krizek, K. J. (2007). Trails, lanes, or traffic: Valuing bicycle facilities with an adaptive stated preference survey. Transportation Research Part A: Policy and Practice, 41(4), 287-301.

Toubia, O., Hauser, J. R., & Simester, D. I. (2004). Polyhedral methods for adaptive choice-based conjoint analysis. Journal of Marketing Research, 41(1), 116-131.

Train, K. E. (2006). Mixed Logit Estimation by Hierarchical Bayes. http://elsa.berkeley.edu/Software/abstracts/train1006mxlhb.html. (accessed Aug. 28, 2011).

Train, K. E. (2008). EM algorithms for nonparametric estimation of mixing distributions. Journal of Choice Modelling, 1(1), 40-69.

Train, K. E. (2009). Discrete choice methods with simulation. Cambridge university press.

Train, K. E. (2016). Mixed logit with a flexible mixing distribution. Journal of choice modelling, 19, 40-53.

Train, K. E. & Weeks, M. (2005). Discrete choice models in preference space and willingness-to-pay space. In Applications of simulation methods in environmental and resource economics (pp. 1-16). Springer, Dordrecht.

Train, K. E., & Wilson, W. W. (2008). Estimation on stated-preference experiments constructed from revealed-preference choices. Transportation Research Part B: Methodological, 42(3), 191-203.

Train, K. E., & Wilson, W. W. (2009). Monte Carlo analysis of SP-off-RP data. Journal of Choice Modelling, 2(1), 101-117.

Vij, A., & Krueger, R. (2017). Random taste heterogeneity in discrete choice models: Flexible nonparametric finite mixture distributions. Transportation Research Part B: Methodological, 106, 76-101.

Villas-Boas, J. M., & Winer, R. S. (1999). Endogeneity in brand choice models. Management science, 45(10), 1324-1338.

Walker, J. L., Wang, Y., Thorhauge, M., & Ben-Akiva, M. (2018). D-efficient or deficient? A robustness analysis of stated choice experimental designs. Theory and Decision, 84(2), 215-238.

Walker, J., & Ben-Akiva, M. (2002). Generalized random utility model. Mathematical social sciences, 43(3), 303-343.

Xie, Y., Danaf, M., Azevedo, C.L., Akkinepally, A., A., Atasoy, B., Jeong, K., Seshadri, R., & Ben-Akiva, M. (2019). Behavioral Modeling of On-Demand Mobility Services: General Framework and Application to Sustainable Travel Incentives. 97th Annual Conference of the Transportation Research Board, Washington D.C.

Yáñez, M., Cherchi, E., Heydecker, B., & Dios-Ortuzar, J. (2011). On the treatment of repeated observations in panel data: efficiency of mixed logit parameter estimates. Networks and Spatial Economics, 11(3), 393-418.

Zhao, F., Ghorpade, A., Pereira, F.C., Zegras, C., & Ben-Akiva, M. (2015a). Quantifying mobility: pervasive technologies for transport modeling. In Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers (pp. 1039-1044). ACM.

Zhao, F., Pereira, F. C., Ball, R., Kim, Y., Han, Y., Zegras, C., & Ben-Akiva, M. (2015b). Exploratory analysis of a smartphone-based travel survey in Singapore. Transportation Research Record: Journal of the Transportation Research Board, 2(2494), 45-56.

Ziegler, C. N., McNee, S. M., Konstan, J. A., & Lausen, G. (2005, May). Improving recommendation lists through topic diversification. In Proceedings of the 14th international conference on World Wide Web (pp. 22-32). ACM.

APPENDIX A: CONVERGENCE DIAGNOSTICS

The stationarity of the Monte Carlo Markov Chains was verified using the Gelman and Rubin and the Heidelberg-Welch diagnostics. For the Gelman and Rubin diagnostic, estimation was repeated twice. The R-hat values for the model with inter- and intra-consumer heterogeneity using the Swissmetro data are presented below.

Table A1: R-hat values for the model parameters with inter- and intra-consumer heterogeneity using the Swissmetro data.

	Population Mean	Inter-consumer Standard Deviation	Intra-consumer Standard deviation
ASC_{SM}	1.01	1.01	1.04
ASC_{Car}	1.01	1.05	1.14
Scale	1.02	1.01	1.02
Travel time	1.01	1.00	1.01

Table A2: Heidelberg-Welch diagnostics (p-value) for stationarity.

	Population Mean	Inter-consumer Standard Deviation	Intra-consumer Standard deviation
ASC_{SM}	0.919	0.369	0.744
ASC_{Car}	0.301	0.693	0.336
Scale	0.497	0.064	0.711
Travel time	0.583	0.542	0.275

The R-hat values for the model with random and systematic inter- and intra-consumer heterogeneity using the Swissmetro data are presented below.

Table A3 : R-hat diagnostics for the model with random and systematic heterogeneity in the constants and travel time parameter.

	ASC_{Car}	ASC_{SM}	β_{time}
Intercept	1.02	1.02	1.08
Male	1.01	1.05	1.08
Low Income	1.00	1.03	1.04
High Income	1.00	1.01	1.00
Inter-Consumer Std. Dev.	1.00	1.05	1.04
Intra-Consumer Std. Dev.	1.01	1.01	1.01

Table A4 : R-hat diagnostics for the model with random and systematic heterogeneity in the scale parameter.

	Intercept	Menu	Menu2
Population Mean	1.08	1.16	1.20
Inter-Consumer Variance	1.07	1.01	1.01
Intra-Consumer Variance	1.03		

Finally, the R-hat values for the model with inter- and intra-consumer heterogeneity using the Tripod SP data are presented below. This model was estimated with 800K iterations, 200K of which were used for estimation and 600K for sampling.

Table A5: R-hat diagnostics for the model with inter- and intra-consumer heterogeneity using the Tripod SP data.

	Population Mean	Inter-consumer variance	Intra-consumer variance
Walk constant	1.02	1.17	1.01
Car constant	1.03	1.05	1.02
On-demand constant	1.01	1.02	1.04
Transit constant	1.00	1.02	1.00
Tokens	1.03	1.00	1.00
Scale	1.00	1.00	1.00
In-vehicle travel time	1.03	1.01	1.02
Out-of-vehicle travel time	1.00	1.00	1.02
Age (on-demand)	1.00	-	-
Student dummy (non-motorized)	1.00	-	-
Log income (car)	1.00	-	-
Zero Household vehicles (transit)	1.02	-	-

APPENDIX B1: ENDOGENEITY RESULTS USING A DIFFERENT MENU GENERATION PROCEDURE

This section presents similar results to Section 3.4, but with a different menu generation procedure. For generating each menu, only the previous choice is taken into consideration; first we generate a universal dataset of 10 alternatives, and recommend the three nearest neighbors to the alternative that was chosen in the previous menu.

The model is first estimated with 5 menus per individual. Afterwards, the first menu is excluded, and different estimations are performed with the remaining menus. The results are shown in Table B1. The results indicate similar trends to those observed in Section 3.4.3: (1) the estimates obtained from all 5 menus are very close to their true values, (2) excluding the first menu results in significant biases, and (3) the magnitude of bias decreases as more menus per individual are included in the estimation.

Significant biases are observed in the constant, which is similar to the results obtained in Section 3.4.3. However, unlike Section 3.4.3, biases in the scale parameter are negligible, while biases in the on-demand parameter are not. This is because a higher weight is used for on-demand compared to transit and bike sharing in the nearest neighbor algorithm.

Table B1: Endogeneity Results using KNN recommendations.

	Population Means				
	Transit	**Bike Sharing**	**On-Demand**	**Constant**	**Scale**
True Values	1.000	0.300	-2.000	-1.000	-0.500
Menu 2	1.012 (0.036)	0.338 (0.035)	-2.263 (0.060)	0.232 (0.204)	-0.539 (0.028)
Menus 2, 3	1.048 (0.027)	0.344 (0.024)	-2.228 (0.036)	-0.430 (0.062)	-0.503 (0.022)
Menus 2, 3, 4	1.027 (0.024)	0.335 (0.020)	-2.165 (0.031)	-0.557 (0.059)	-0.484 (0.017)
Menus 2, 3, 4, 5	1.030 (0.021)	0.334 (0.017)	-2.122 (0.026)	-0.67 (0.051)	-0.475 (0.014)
Menus 1,2, .., 5	1.030 (0.019)	0.328 (0.016)	-1.998 (0.019)	-1.022 (0.043)	-0.484 (0.013)
	Variances				
	Transit	**Bike Sharing**	**On-Demand**	**Constant**	**Scale**
True Values	1.000	1.000	1.000	1.000	0.250
Menu 2	0.697 (0.193)	0.859 (0.115)	1.874 (0.198)	10.309 (3.396)	0.173 (0.021)
Menus 2, 3	0.940 (0.110)	0.968 (0.081)	1.751 (0.104)	1.324 (0.143)	0.284 (0.036)
Menus 2, 3, 4	0.995 (0.068)	0.955 (0.060)	1.500 (0.079)	1.235 (0.112)	0.254 (0.026)
Menus 2, 3, 4, 5	1.042 (0.053)	0.955 (0.048)	1.365 (0.064)	1.089 (0.109)	0.218 (0.020)
Menus 1,2, .., 5	1.065 (0.047)	0.980 (0.043)	1.021 (0.040)	0.961 (0.098)	0.229 (0.015)

APPENDIX B2: ENDOGENEITY IN ADAPTIVE IN LINEAR REGRESSION

Liu et al. (2007) showed that in adaptive choice-based conjoint (CBC) experiments (such as the one implemented in Sawtooth software), endogeneity becomes ignorable for estimation once the data has been collected because of strict exogeneity and the "likelihood principle". To demonstrate this, they consider the panel data regression shown in equation (B1):

$$y_t = \beta_1 x_{1t} + \beta_2 x_{2t} + \epsilon_t \qquad t = 1, 2, 3 \qquad \text{(B1)}$$

All the explanatory variables are exogenous (drawn randomly from $\{-1,1\}$ except x_{23} which is given by:

$$x_{23} = \begin{cases} -1 \; if \; y_1 y_2 > 0 \\ +1 \; if \; y_1 y_2 \leq 0 \end{cases} \qquad \text{(B2)}$$

The true values of the coefficients β_1 and β_2 are 1 and 2 respectively, and the sample size is 10,000. If a regression model is estimated using the entire data (30,000 observations), the estimates are consistent and close to the true values. However, if a separate regression is estimated for each individual (3 observations each), the estimates are inconsistent and far from the true values (see Liu et al., 2007 for details).

The authors indicate that the likelihood of the data $\pi(y_1, y_2, x_3, y_3 | \beta, \sigma^2)$ can be expressed as:

$$\pi(y_1, y_2, x_3, y_3 | \beta, \sigma^2) = \pi_1(y_1, y_2 | \beta, \sigma^2) \times \pi_2(x_3 | y_1, y_2, \beta, \sigma^2) \times \pi(y_3 | x_3, \beta, \sigma^2) \qquad \text{(B3)}$$

$\pi_2(x_3 | y_1, y_2, \beta, \sigma^2)$ is always equal to 1, since given y_1 and y_2, the selection of x_3 is deterministic. Therefore, the likelihood of model parameters is unchanged if π_2 is excluded.

However, this experiment does not consider any agent effects (or heterogeneity). We modify the above experiment by adding an individual-specific normally distributed term (i.e. agent effect) as shown in equation (B4):

$$y_{nt} = \beta_1 x_{1nt} + \beta_2 x_{2nt} + \delta_n + \epsilon_{nt} \quad t = 1, 2, 3, \quad \delta_n \sim N(0, \sigma_\delta^2) \quad \text{(B4)}$$

In addition, we try different scenarios where:

- Agent effect is included or excluded in the estimation (using random effects);
- x_{23} is endogenous (as per equation 6) or exogenous (random).

Using OLS regression, bias is only observed when both agent effect and endogeneity are present. When either one is absent, OLS results in unbiased estimates. Additionally, the bias (when both endogeneity and agent effect are present) is eliminated by using random effects regression, which accounts for heterogeneity. The results are presented in Table B2.

Table B2: Estimation results of the regression Monte Carlo experiment

	Intercept	β_1	β_2
True Values	0.0	1.0	2.0
No endogeneity (all variables are exogenous)	-0.032	1.085	2.035
One endogenous variable, No agent effect (same as in Liu et al., 2017)	-0.001	1.024	2.005
One endogenous variable + agent effect (Linear regression)	0.721	0.179	1.388
One endogenous variable + agent effect (Random effects regression)	-0.056	1.016	2.025
No endogeneity + agent effect	0.022	1.039	1.997

I also present an extension of this example to a dynamic setting with 8 observations per individual and an individual/agent effect $\delta_n \sim N(0,9)$ as shown in equation B5. The endogenous explanatory variable x_{2t} is given by equation (A6). The used sample size is 100,000.

$$y_{nt} = \beta_1 x_{1nt} + \beta_2 x_{2nt} + \delta_n + \epsilon_{nt} \quad t = 1, 2, \dots 8 \quad \delta_n \sim N(0,\ 9),\ \epsilon_{nt} \sim N(0,25) \tag{B5}$$

x_{1nt} is randomly drawn from {0, 1} in all time periods ($t = 1, 2, \dots, 8$). In the first two time periods, x_{2n1} and x_{2n2} are randomly drawn from {0, 1}, while in the subsequent time periods, x_{2nt} is generated as follows:

$$x_{2nt} = \begin{cases} -1 \text{ if } y_{n(t-2)} + y_{n(t-1)} \leq 3 \\ +1 \text{ if } y_{n(t-2)} + y_{n(t-1)} > 3 \end{cases} \qquad t = 3, 4, \dots, 8 \tag{B6}$$

The model is estimated using random effects regression, which accounts for the added agent effect. When estimation is done over all 8 time periods, the estimates are consistent as shown in Table B3.

Table B3: Estimation Results Using 8 Time Periods

	True Value	Estimate	Std. Error	t value
(Intercept)	0.0	-0.011	0.012	-0.850
x1	1.0	0.987	0.012	84.090
x2	2.0	2.013	0.007	281.700

However, when the first two (exogenous) observations are excluded, the estimates of β_2 are far from their true values. It is also observed that the bias decreases as more time periods are included in the estimation as shown in Table B4.

Table B4: Estimation Results with Different Time Periods

	Estimate of β_2	Standard Error of β_2
Estimation with t=3, 4	3.425	0.013
Estimation with t=3,4, 5	3.284	0.011
Estimation with t=3,4,5, 6	3.070	0.010
Estimation with t=3,4,..., 7	2.810	0.009
Estimation with t=3,4,..., 8	2.602	0.008
Estimation with t=1,2,..., 8	2.013	0.007

Note that this bias is not observed if the agent effect was not present. It requires both the agent effect and the endogeneity of x_{2t}.

APPENDIX C1: TRIPOD SP DESCREPTIVE STATISTICS AND USER INTERFACE

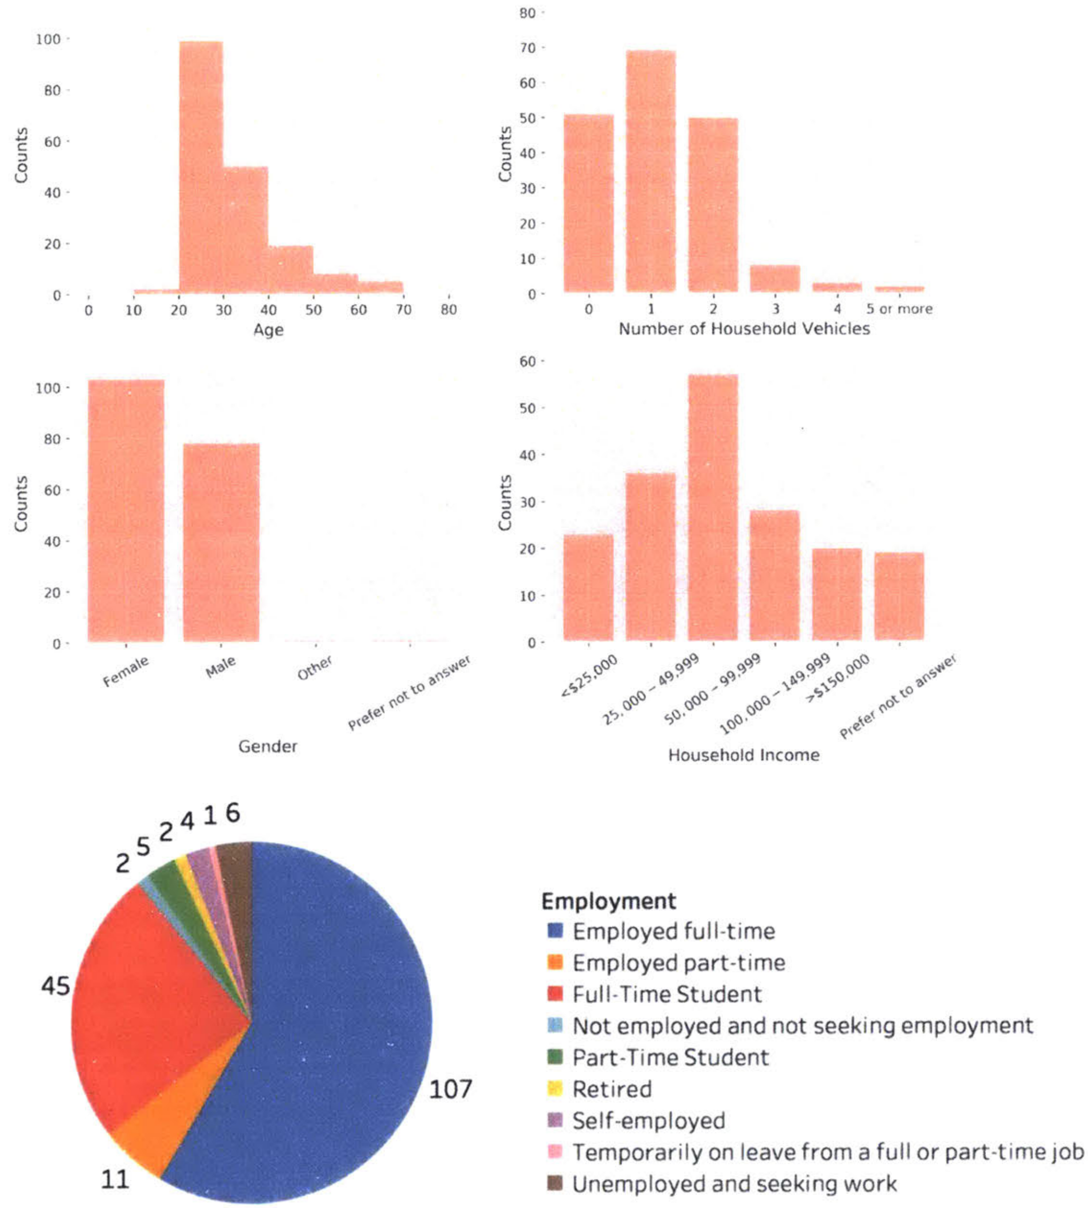

Figure C1: Sample characteristics.

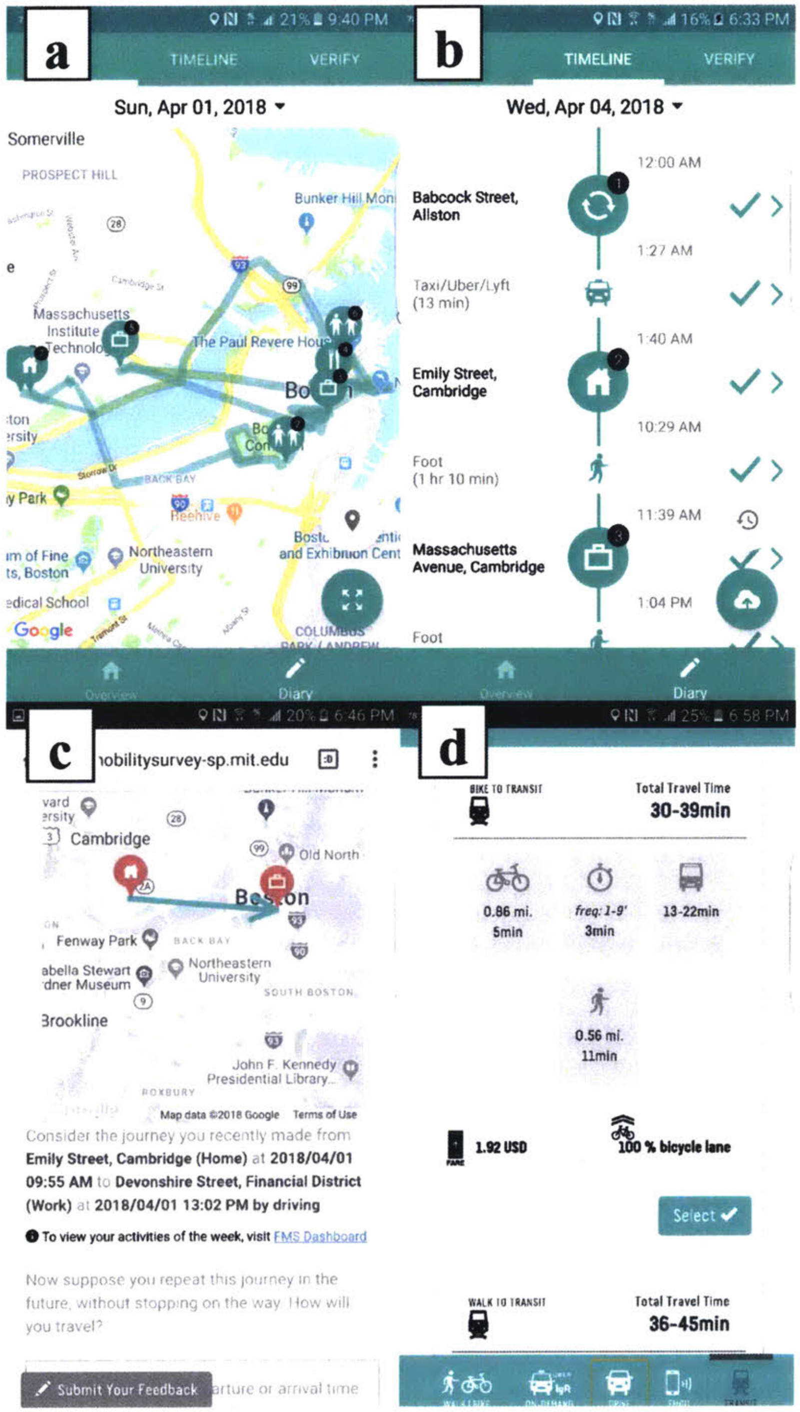

Figure C2: Screenshots of (a) day trajectories, (b) trip/activity diary validation, (c) reference trip for SP, and (d) SP profiles (in the same order in which the user sees them).

APPENDIX C2: TRIPOD SP ATTRIBUTES

This section presents the attributes of the basic alternatives presented in Tripod SP (walking, biking, bike-sharing, drive alone, carpooling, car-sharing, taxi, on-demand services, and transit). The attributes across all modes are listed in Table C1.

For drive alone and carpool alternatives, the existence of toll and parking costs is randomly decided using Boolean design variables. If there is a non-zero parking cost, it is calculated based on the activity duration. The experimental design ensures that the drive alone travel time is not greater than carpooling travel time. Drive alone availability is determined by car ownership and a valid driving license. If these two conditions are satisfied, the user might also choose carpooling "as a driver" or "as a passenger". Otherwise, only carpooling "as a passenger" is available.

For on-demand services (e.g. Uber, Lyft, etc.), the design ensures that travel time for shared alternatives (e.g., UberPool) is never less than that of the private ones (e.g. UberX). These alternatives are assumed to be always available except car-sharing which necessitates a valid driving license. For subscription-based services (e.g., car-sharing or bike-sharing), users who are not subscribed to these services are presented with annual subscription fees.

Under non-motorized modes, bike-sharing has additional attributes of access and egress times (as users pick up/drop off bikes at the stations), annual subscription, and time-based rental costs. The availability of these modes is determined by pre-specified maximum walking and biking distances. Bike availability is contingent on bike ownership; however, bike-sharing is always displayed given that maximum biking distance is not violated.

Finally, transit modes include bus and train, with walk or car access (i.e., park-and-ride). The number of transfers is randomly determined in a time-based manner, e.g., trips shorter than 10 minutes have no transfers, those between 10 and 20 minutes may have up to one transfer, and trips longer than 20 minutes may have up to two transfers. The fares include a fixed and a distance-based component for flexibility (e.g., setting the distance-based rate to zero results in a flat fare). A transit alternative may include bus, train, or a combination of the two. The availability is either based on the existing conditions (referring to Google Maps), or defined by the researcher in order to test hypothetical scenarios.

Table C1: Considered attributes across alternatives

	Non-motorized			Motorized		On-demand			Transit
	Walk	Bike	Bike-sharing	Drive alone	Carpool	Taxi	Uber/Lyft	Car-sharing	
Walking time	x								
Biking time		x	x						
Waiting time						x	x		
Schedule delay									
Access/egress time			x	x	x			x	x
In-vehicle travel time				x	x	x	x	x	x
Parking time				x	x				
% Bike lane		x	x						
Annual subscription cost			x					x	
Distance/time-based variable cost			x			x	x	x	x
Fuel cost				x	x				
Toll cost				x	x			x	
Parking cost				x	x			x	
Transfers									x
Headway									x