Estimation of Cumulative Prospect Theory-based passenger behavioral models for dynamic pricing & transactive control of shared mobility on demand

by

Vineet Jagadeesan Nair

B.S., University of California, Berkeley (2018)
MPhil, University of Cambridge (2019)

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Computational Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2021

Author .....
Department of Mechanical Engineering
January 29, 2021

Certified by .....
Dr. Anuradha Annaswamy
Senior Research Scientist
Thesis Supervisor

Accepted by .....
Youseff Marzouk
Chairman, Department Committee on Graduate Theses

Estimation of Cumulative Prospect Theory-based passenger behavioral models for dynamic pricing & transactive control of shared mobility on demand

by

Vineet Jagadeesan Nair

Submitted to the Department of Mechanical Engineering
on January 29, 2021, in partial fulfillment of the
requirements for the degree of
Master of Science in Computational Science and Engineering

Abstract

This thesis studies the optimal design of large-scale shared mobility on demand services (SMoDS) in urban settings. Specifically, we build upon previous work done in the Active-Adaptive Control Lab lab on the dynamic pricing and routing of ride sharing services. We develop and characterize a novel passenger behavioral model based on Cumulative Prospect Theory (CPT) to more accurately represent decision making in the presence of significant risks and uncertainty associated with SMoDS' travel times. A comprehensive survey was designed to estimate both the mode-choice and CPT models. The mode choice section consisted of a series of discrete choice experiments created via factorial design, while the CPT section involved carefully constructed lottery questions and travel choice scenarios to elicit risk preferences. After conducting a pilot study and going through several iterations, the survey was launched via a panel firm. Data was collected from 1000+ respondents in the Greater Boston metro area. This was used to fully characterize the model and estimate parameters through methods including maximum simulated likelihood estimation, nonlinear least squares and global optimization tools. I also utilized other techniques like regularization and transfer learning to improve the quality of results obtained.

Beyond parameter estimation, the uncertainty associated with such behavioral models was quantified via well-established nonlinear programming methods. Sensitivity and robustness analyses were performed to assess the effects of CPT model parametrization errors on the performance of the SMoDS system and objectives like expected revenue, average waiting times etc. These insights were used to design and simulate a closed loop, feedback control mechanism for the SMoDS system to correct modelling errors in real-time, achieve setpoint tracking and enable parameter estimation. The scheme uses the dynamic tariff as a transactive control input to influence the passenger's behavior as desired. This was implemented via gradient-descent based control schemes to update the price signals, in order to drive the (i) drive the passengers' probabilities of accepting the SMoDS ride offer towards the desired value while

also (ii) learning the true passenger behavioral model parameters.

Thesis Supervisor: Dr. Anuradha Annaswamy
Title: Senior Research Scientist

Acknowledgments

Firstly, I would like to sincerely thank my advisor and thesis supervisor Dr. Anuradha Annaswamy for providing valuable mentorship and guidance throughout my time at MIT so far, and for the many insightful research discussions. In addition, I want to thank my fellow Active-Adaptive Control lab members, especially Dr. Yue Guan and Claudio Lombardi for several helpful discussions and feedback about this project and others.

I want to thank the Ford-MIT Alliance for funding this project and the entire team for their support and feedback over the past 1.5 years - including Dr. Eric Tseng, Dr. Baljeet Singh, Dr. James Fishelson, Dr. Sara Dadras and Dr. Dimitar Filev.I would like to acknowledge Dr. Shenhao Wang for lending his expertise in the area of Prospect theory and Dr. Joanna Moody for her valuable advice particularly during the survey design process. I would like to thank Sarah Humphrey and the rest of the Qualtrics team for all the assistance they provided in seamlessly launching the panel study and troubleshooting during survey data collection.

I would like to acknowledge the MIT SuperCloud and the team at the Lincoln Laboratory Supercomputing Center for providing HPC, database and consultation resources that have contributed to the research results reported within this thesis. I would also like to thank Leslie Regan and Kate Nelson for providing assistance and support with several administrative, academic and professional matters during graduate school.

I am grateful for the friends that I have gained during my time here and the events of the past year have only made me appreciate them even more. Finally, I want to thank my parents and my brothers for always unconditionally supporting me in my endeavors.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Shared Mobility on Demand Services (SMoDS)

In recent times, several ride sharing platforms have emerged, that offer the potential for increased affordability, efficiency and customizability [3] [4]. There is also a growing shift away from exclusive, door-to-door ride hailing services towards ride-pooling, which offers additional benefits including mitigating traffic congestion, reducing cumulative travel times and emissions while increasing fleet utilization rates [5]. This thesis pertains to such Shared Mobility on Demand Services (SMoDS) i.e. large-scale pooled ridesharing systems that operate primarily in urban and suburban settings. As can be seen in fig. 1-1, SMoDS offers a desirable balance of lower prices, medium carrying capacities, greater convenience and flexibility. Although mass transit options are cheaper and can accommodate more passengers, these are not as convenient or flexible due to their fixed, static routes that may not cater to all users. On the other hand, traditional ride hailing services are more convenient since they usually involve shorter waiting, walking and in-vehicle travel times. However, these are more expensive compared to SMoDS. Furthermore, a short walk by the passenger to/from the vehicle can significantly cut in-vehicle ride times in urban settings under certain traffic conditions.

Although SMoDS do offer numerous benefits, these come at the cost of increased uncertainty in travel times. For instance, current ride hailing services like Uber and

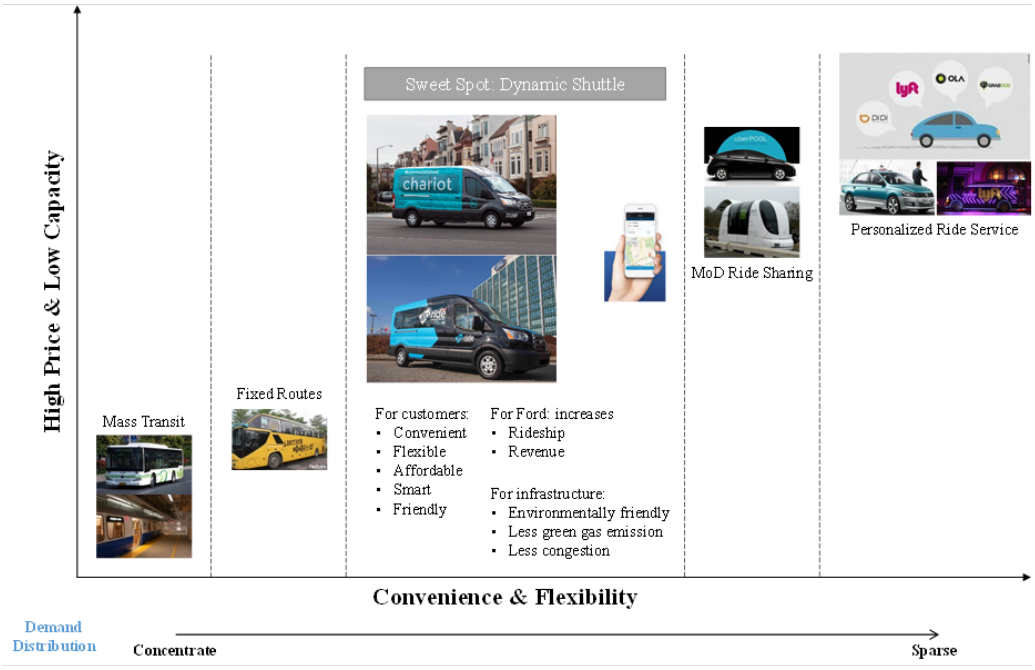

Figure 1-1: Comparison of different transport modes in terms of key performance metrics.

Lyft often provide passengers with a finite window for arrival times at their destination. In order to efficiently operate such SMoDS, we need to obtain a clear understanding of how passengers make tradeoffs between price and travel times while making mode choice decisions. This can then be used to plan and operate the SMoDS system in an optimal manner so as to match supply and demand in real-time, while also achieving desired objectives like maximizing expected revenue, minimizing average wait times, etc.

1.2 Motivation for Cumulative Prospect Theory (CPT)

Conventional Expected Utility Theory (EUT) postulates that consumers choose among travel options based on their respective expected utilities [6]. Cumulative Prospect Theory (CPT) is an alternative to EUT that better describes subjective human decision making in the presence of uncertainty and risk [7, 8]. This is necessary in the case of SMoDS due to the significant variability in travel times introduced by pooling rides. The behavioral model is described by the value function $V(\cdot)$ and probability

distortion $\pi(\cdot)$ given by [1], with $\pi(0) = 0$ and $\pi(1) = 1$ by definition. These non-linearities transform the objective utilities (u) and probabilities (p) of each possible outcome to subjective values, as perceived by the passengers.

$$V(u) = \begin{cases} (u - R)^{\beta^+} & \text{if } u \geq R \\ -\lambda(R - u)^{\beta^-} & \text{if } u < R \end{cases} \tag{1.1}$$

$$\pi(p) = e^{-(-ln(p))^{\alpha}} \tag{1.2}$$

The CPT parameters here describe loss aversion (λ), diminishing sensitivity in gains (β^+) and losses ($\beta-$) and probability distortion (α). The reference R is the baseline against which users compare the utilities of different prospects. These can vary across individuals and also depending on the particular set of alternatives that the customer is facing.

1.3 Summary of contributions

Prior work from our lab has described a comprehensive solution for the provision of SMoDS with (1) Dynamic routing via an alternating minimization (AltMin) optimization algorithm [9] and dynamic pricing using a passenger behavioral model based on Cumulative Prospect Theory (CPT) [1]. This work focuses on accurately estimating the proposed behavioral model in practical settings. Due to the lack of publicly available ride sharing data from transportation network companies (TNCs), we decided to collect our own data. I designed a comprehensive survey to assess passengers' travel preferences as well as their risk attitudes relating to travel time uncertainty. These were used to quantitatively estimate both the mode-choice and CPT models, where we experimented with several different solution methods in an effort to improve the quality of our results. In addition to conducting statistical tests to verify the validity and correctness of the model, a sensitivity analysis was conducted to assess its robustness to CPT parametrization errors. These results informed our efforts to develop a closed loop transactive control strategy to correct such errors in real time

using feedback and the dynamic tariff as a control input. To our knowledge, this is the first work that establishes a methodology for estimating CPT-based behavioral models specifically applied to pooled ride sharing services, and formulates a feedback control mechanism based on the notion of price-based transactive control.

Chapter 2

Literature review

2.1 Mode-choice modelling

Discrete choice experiments have long been established as a standard approach to estimate utility functions [10]. These experiments generally aim to infer *stated* preferences (SP) via stated choice experiments [11] in which the survey respondents are explicitly asked to make choices between several alternatives or options. Thus, SP data relies on what consumers say they will do in a hypothetical setting. This is in contrast to *revealed* preferences (RP) which relies on existing historical and observational data on what consumers actually do in reality. Most recent mode-choice studies have relied on SP methods since RP data may not always be available or feasible to collect due to budget and time constraints. Furthermore, SP methods are more useful to identify key explanatory variables, capture the effects of new attributes and features while also satisfying statistical model assumptions (since the SP study can be designed accordingly) [11].

Discrete choice methods have been extended to utility functions with both fixed and random coefficients, estimated via multinomial logit models. In particular, mixed logit models have been found to be especially useful since these allow for combinations of both fixed and random parameters as well as flexible choices for the underlying mixing distribution assumed for the parameter estimates [12, 13, 14]. Furthermore, these are more realistic and require fewer restrictive assumptions compared to stan-

dard logit models since they allow for differences in tastes and preferences across different people [15]. Finally, mixed logit is well suited to panel data involving a series of repeated choices, as is the case with our survey which has multiple mode choice experiments [16].

Given its numerous relative advantages, such methods have been widely used for mode choice studies in the transportation literature, to derive customers' utility functions from survey data [17, 18] and thus also forecast travel demand [19]. These approaches based on limited amounts of survey data have shown to be reasonably accurate and robust even in the presence of response errors such as sampling errors, self-selection biases, inconsistent answers etc. [20]. Such state of the art discrete choice methods are also related to studies in the transportation literature that aim to determine the value of time (VOT) and the value of reliability (VOR) from SP surveys of passengers [21, 22], as well as their willingness to pay for improvements in these metrics [23]. In particular, heterogeneity in the VOT across different legs of a trip (i.e. walking, waiting and riding/driving) are important considerations for the design of services [24] like SMoDS while VOR can be related to the travel time uncertainty introduced by such pooling services.

Numerous recent studies have applied such discrete choice SP methods to mobility-on-demand applications similar to the SMoDS such as taxi cabs [25, 26], shared autonomous vehicles [27, 28, 29], general Mobility as a service (MaaS) systems in urban areas [30], ride hailing [31] and ride pooling services [32]. Other related applications include predicting demand for clean-fuel (electric and hybrid) vehicles [33] and responsiveness to dynamic congestion pricing in cities [34]. The novelty of the current work primarily lies in the introduction and empirical evaluation of CPT as an alternative method to better understand consumers' risk attitudes towards travel time uncertainty and their willingness to pay for mitigating uncertainty.

2.2 Modelling risk preferences with CPT

The standard approach to model passengers' travel risk preferences has generally been to quantify measures like VOR and use the utility functions of each mode to compute their associated choice probabilities, using conventional expected utility theory (EUT). However, EUT has been shown to be insufficient to accurately predict consumer choices when there is significant uncertainty and risk involved in both the outcomes (or events) themselves and their associated probabilities. Cumulative Prospect Theory (CPT) is a framework introduced by Nobel prize-winning behavioral economists and psychologists that has been extensively shown to better represent decision making under uncertainty [7, 8]. It builds upon EUT by introducing additional nonlinear mappings as shown in eq. (1.1). These are necessary to model four characteristic behaviors displayed by agents, that are *irrational* and *subjective*, thus not accounted for under EUT [1]:

1. **Framing effect**: People perceive outcomes as gains and losses relative to a specific subjective reference point R (or baseline value) and not in absolute terms.

2. **Reflection effect**: In both the gain and loss regimes, the agent's sensitivity diminishes as they move farther away from the reference. This implies that the perceived value function is concave in gains and convex in losses ($0 < \beta^-, \beta^+ < 1$).

3. **Loss aversion**: Losses hurt much more than equivalent gains, resulting in an asymmetric value function.

4. **Probability distortion**: People tend to subjectively overweight the likelihood of occurrence of less likely, small probability events and underweight more common, large probability events ($0 < \alpha < 1$).

These effects are clearly illustrated in the plots shown in fig. 2-1 below.

CPT has widely been applied in the literature to study and describe human behavior in situations involving significant uncertainty. Some examples include the

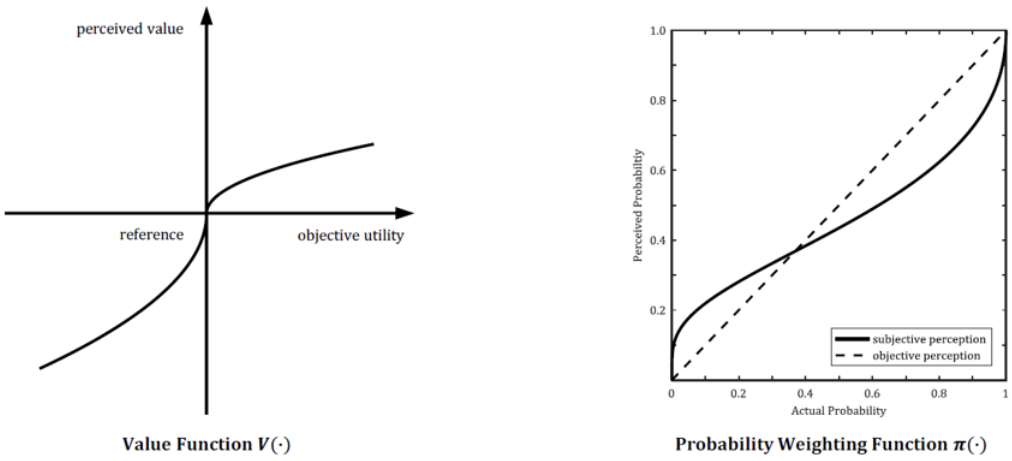

Figure 2-1: Characteristic examples of possible value function and probability weights [1].

insurance industry, financial lotteries [35, 36, 2] and optimal portfolio selection by investors considering both risky as well as risk-free assets [37]. Past studies have also examined the relationships between risk attitudes towards time and money, and socioeconomic indicators like household income [38]. These studies illustrate how CPT is more accurate and versatile in its ability to capture decision making, compared to simpler models based on just expected utility theory, loss aversion and exponential time discounting. A few recent studies have applied CPT specifically to the transportation setting, to model the route choice behavior of drivers on highways or freeways [39, 40] and the adoption of autonomous vehicles [28]. While [40] used real-time traffic information and GPS records to build the CPT-based route choice model, [39] utilized SP questionnaires presenting different scenarios to develop a structural behavioral model. To our knowledge, CPT has not yet been applied specifically for shared mobility services, to understand the implications of the high degree of travel time variability. However, almost all of the papers surveyed attempt to infer risk attitudes and estimate the CPT parameters using simplified lottery questions (for an example see the appendix). In reality though, passengers' risk preferences towards monetary gains and losses may differ significantly from their perceptions of travel time uncertainty. Furthermore, past studies like [28] used simpler model parametrizations based on prospect theory (PT) which does not use cumulative probability weighting

[8]. The survey developed as part of this thesis is unique in that we attempt to directly estimate travel time risk parameters using discrete mode choice games, rather than using simplistic lottery questions as an analog.

2.3 Sensitivity and robustness analysis

There are several sources of error and uncertainty involved in the parameter estimation process that need to be addressed. These include sampling errors, survey design issues, response biases etc. Population level mode-choice models require very large sample sizes to obtain relatively accurate distributions for utility function coefficients and still involve some finite levels of uncertainty. Furthermore, travel preferences of consumers are fluid and can vary significantly over time and among different individuals in a population, but surveys capture only a static snapshot for a subset. Inaccuracies in model parameters can result in setting sub-optimal dynamic tariffs that reduce the operational efficiency of the SMoDS, leading to decreases in revenue, ridership and fleet utilization rates. Additionally, CPT parameters describing risk attitudes are often specific to each user and determined from a much smaller set of responses, thus lacking statistical properties like asymptotic normality [28]. Since the novel aspect of this behavioral model is the incorporation of CPT and because the parameters associated with it involve a greater deal of uncertainty compared to the mode-specific utility functions, this work will focus on analyzing the sensitivity only with respect to these CPT parameters. To our knowledge, CPT has not been explored for SMoDS in the literature and neither has such a sensitivity analysis been considered to date.

There has been significant past work conducted in analyzing the sensitivity of parametric linear and nonlinear optimization problems. The early work in this field analyzed the sensitivity and stability of nonlinear programs (NLP) [41, 42]. These established the mathematical foundation including basic theorems related to the smoothness, continuity and differentiability properties of the optimal solution and value function [43], performing both first and second order sensitivity analyses as well

as determining asymptotic bounds on sensitivity derivatives [44]. Several different approaches and solution techniques have also been explored, including penalty-function methods [45], generalized perturbation approaches [46] and directional derivatives [47]. In addition to localized analyses that focus on varying one parameter at a time, methods for global sensitivity analysis have also been studied, often using Monte Carlo techniques [48].

The main contribution of this paper lies in applying methods from sensitivity and robustness analysis specifically to behavioral models based on prospect theory, thereby developing tools to manage the significant uncertainty and estimation errors associated with such models in the SMoDS context.

Chapter 3

Methodology

In this chapter, we describe the methodology used for survey design, data collection, estimation of the mode choice and CPT models, statistical testing and validation, and the sensitivity analysis. The different sub-components of our system are shown in fig. 3-1 below. The model proposed here provides a complete end-to-end solution for the overall provision of SMoDS. The key performance indicator (KPI) we consider here is the expected waiting time (EWT) for passengers over the whole SMoDS system or region being served, which in turn depends on the passengers' probability of accepting the SMoDS. This acceptance probability is the main output that we would like to predict and regulate in this study. In addition, we could also have secondary objectives like expected revenue, total ridership, fleet utilization rates etc.

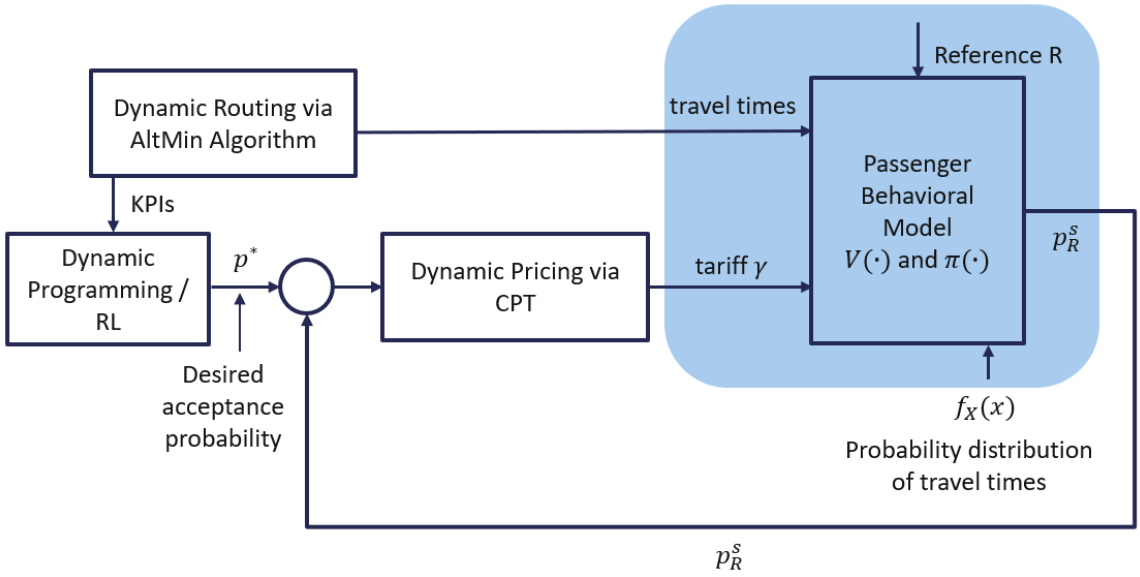

Figure 3-1: Illustration of our complete integrated solution for the SMoDS.

3.1 Mode choice modelling

3.1.1 Specification of utility functions

In this study, we constructed a model that considered three different modes of transport: Public transit, Exclusive ride hailing (similar to UberX or Lyft) and Pooled ridesharing i.e., the SMoDS (similar to UberPool or Lyft Line). Our model was developed specifically for the Boston greater metropolitan area. Thus, there are three possible types within the public transit mode: Bus, Subway (T) and Commuter rail. A different probabilistic discrete choice model was fitted for each mode using random, objective utility functions, assuming utility-maximizing behavior by all decision-makers. Here the observed utility U_i of a travel alternative (or mode) i is comprised of two distinct parts [49]:

$$U_i = V_i + \epsilon_i \tag{3.1}$$

1. U_i is the true, observed utility of option i.

2. V_i is the systematic (mean) or deterministic portion of the utility that can be

described in terms of the attributes or features of the i^{th} alternative.

3. ϵ_i are the random error terms, i.e., the portion of the utility that's unknown or unobservable.

Here, we assume that the systematic utility is a function of two key trip attributes, travel time and price. The travel time is further segmented into the times spent on each leg of the trip: walking, waiting and in-vehicle riding times. Thus, the deterministic mode-specific utility function is given by the following multinomial or conditional logit (logistic regression) model:

$$V_i = a_{walk,i}\, t_{walk,i} + a_{wait,i}\, t_{wait,i} + a_{ride,i}\, t_{ride,i} + b\gamma_i + ASC_i \tag{3.2}$$

$$i \in \{\text{Public transit, Exclusive ride hailing, Ride pooling}\} \tag{3.3}$$

where $a_{walk,i}$, $a_{wait,i}$ and $a_{ride,i}$ are the travel time coefficients while b is the coefficient on the tariff γ_i charged for that mode. Since increases in travel times and price would cause disutility to passengers, we expect all of these coefficients to be negative in sign. ASC_i is a bias/offset referred to as the alternative specific constant, which accounts for all other observable trip attributes which affect its utility, other than time and price. These could include factors like convenience, safety, carbon footprint, privacy, other externalities etc. Formally, ASC is defined as the average impact of all factors not included in the model, similar to the constant term in linear regression [10]. An important thing to note is that the mode choice model is estimated for the population as a whole, i.e. the coefficients are not specific to each respondent. Thus, we obtain a consistent estimator and the distributions for the mode choice parameters will be approximately Gaussian, as long as the sample size is large enough (by the Central Limit Theorem).

Since utility is a relative (and not absolute) measure of well-being, we need to define a reference point with respect to which it is measured. For this purpose, we set public transit as the *baseline* alternative for which $ASC_{transit} = 0$. Then, we can interpret the other two $ASCs$ for exclusive ride hailing and pooled ridesharing

relative to it. Once the utilities of each of the modes are known, the probability of choosing each mode is given by it's respective marginal utility:

$$Pr(i) = \frac{e^{V_i}}{\sum_j e^{V_j}} = \frac{1}{1 + \sum_{j \neq i} e^{V_j - V_i}} \tag{3.4}$$

The error terms ϵ_i cannot be observed or measured. These can be represented using several possible distributions over the alternatives i and all the individuals in the population. Different assumptions on the distribution of error terms lead to different kinds of logit models. The most widely used model is standard logit, which assumes that the error terms are independent and identically distributed (i.i.d) according to an extreme value (Gumbel) distribution across all alternatives and individuals/observations [10]. This implies that the errors and unobserved factors are uncorrelated and also have the same variance across individuals, choices and over time. This leads to several nice properties such as the independence from irrelevant alternatives (IIA) which means that while deciding between any two given options i and j, the agent does not need to take into consideration any of the other options in the choice set:

$$\frac{Pr(i)}{Pr(j)} = \frac{e^{V_i}}{e^{V_j}} = e^{V_i - V_j} \tag{3.5}$$

Such a formulation allows us to directly relate the choice probabilities of each mode to their systematic utilities:

$$log\left(\frac{P_i}{1 - P_i}\right) = V_i \tag{3.6}$$

Although this considerably simplifies the analysis, it is not realistic in our setting since the features are likely going to be correlated across different modes (or alternatives) as well as for sequential choices made by the same person over time. Furthermore, standard logit is restrictive in that it enforces identical preferences across all the respondents i.e., we would only be able to estimate the utility function coefficients as fixed, static parameters. For these reasons, mixed logit models are preferred since these do not require the restrictive i.i.d assumption, allowing for correlated errors and variations in tastes across respondents. This enables us to use random parameter

logit models, estimating distributions for the parameters rather than fixed values.

3.1.2 Logit models

If the errors and unobserved factors are assumed to be distributed jointly normal over alternatives, respondents and time, we get the Multinomial Probit (MNP) probabilistic choice model. Although Probit offers flexibility in modeling any type of correlation pattern, its reliance on the normal distribution makes it unsuitable for most mode choice models. For example, we know from intuition that the tariff and travel time coefficients are negative. Assuming a normal distribution would contradict this since it has non-zero density on both sides of zero [10]. Furthermore, MNP models offer practical challenges for discrete choice analysis since they are less tractable and thus harder to numerically estimate, predict and interpret results [49]. For instance, MNP models can only be estimated by evaluating expensive, high-dimensional multivariate integrals (e.g. to compute normalization factors) which causes numerical issues. Thus, MNL models are preferred since they use Gumbel (Type I) extreme value distributions that roughly resembles normals (as seen in fig. 3-2 below) but offers several computational benefits for maximizing the likelihood function, resulting in a closed form model.

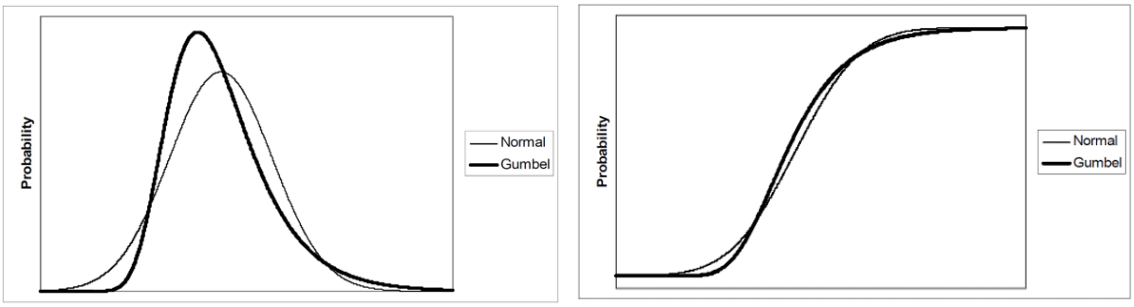

(a) Probability density function. (b) Probability density function.

Figure 3-2: Illustration of PDFs and CDFs for a Normal and Gumbel distribution with identical mean and variance.

Mixed logit models are the most general form of logit since they can estimate any type of discrete choice model following an arbitrary mixing distribution. Choice

probabilities in mixed logit can be represented as [10]:

$$P_{ni} = \int L_{ni}(\beta) f(\beta \mid \theta) d\beta \tag{3.7}$$

$$L_{ni}(\beta) = \frac{e^{V_{it}(\beta)}}{\sum_j e^{V_{jt}(\beta)}} \tag{3.8}$$

where $L_{ni}(\beta)$ are the logit probabilities calculated at various values of the model parameters β, which follow a mixture distribution specified by the density function $f(\beta \mid \theta)$, generally taken to be continuous. θ represents the parameters that fully characterize the probability density of β. I experimented with several different distributions for the mode choice model, including normal, lognormal, truncated normal, uniform and exponential distributions. The lognormal, truncated normal and exponential forms are particularly useful for the travel time and price coefficients since they always give coefficients of the same sign for every decision maker, and these parameters are expected to be negative. Normal distributions are more appropriate for the ASC terms since ride hailing and ride pooling could have either positive or negative values for these relative to public transit, for which we set $ASC_{transit} = 0$ by definition. For example, for normal distributions β $N(\mu, \sigma^2)$ or lognormal $ln(\beta)$ $N(\mu, \sigma^2)$, the mean and standard deviation parameters $\theta = (\mu, \sigma)$ are estimated to describe the mixing distribution. However, one can occasionally run into numerical estimation issues with lognormals, when the logit probabilities become unbounded (approach $+\infty$ or $-\infty$) for certain combinations of parameters. The mode choice model from eq. (3.3) is also modified slightly to account for the fact that the survey produces panel data with several choices made by each respondent i at each time instant t:

$$V_{it} = a_{walk,i}\, t_{walk,it} + a_{wait,i}\, t_{wait,it} + a_{ride,i}\, t_{ride,it} + b\gamma_{it} + ASC_{it} \tag{3.9}$$

$$i \in \{\text{Public transit, Exclusive ride hailing, Ride pooling}\} \tag{3.10}$$

3.1.3 Value of time (VOT)

Once the mode-specific utility functions in eq. (3.10) are determined, we can then use the estimated travel time and price coefficients to determine the passengers' value of time (VOT) spent on different modes. The value of time is defined as the extra tariff that a person would be willing to pay or cost incurred to save an additional unit of time, i.e., it measures the willingness to pay (WTP) for extra time savings. In absolute terms, the VOT spent on mode i can be calculated as the ratio between the marginal utilities of travel time and trip cost:

$$VOT_i = \frac{\frac{\partial U_i}{\partial t}}{\frac{\partial U_i}{\partial \gamma}} = \frac{a_i}{b_i} \tag{3.11}$$

The VOT is often presented in the literature as some proportion of the hourly wage rate. This can be easily calculated by normalizing the price term γ_i in the utility function for each respondent by their hourly wage (which would also need to be collected in the survey). Then the VOT for mode i is given by an $\frac{a_i}{b_i}$ proportion of their hourly wages. Past studies in the literature have estimated absolute VOT to be \$ 0.22/min for SMoDS and \$ 0.77/min for UberX, and \$ 0.40/min for business travelers in the US [34]. On the other hand, [50] estimated the VOT of commute trips to be $\approx 50 - 60\%$ of the average wage rate.

3.2 CPT model

We will now delve into the elements of the cumulative prospect theory (CPT) model in greater detail. Firstly, CPT postulates that when presented with risk or uncertainty, decision makers perceive outcomes by applying a nonlinear transformation on their objective utilities. This is known as the value function which maps objective utilities of individual outcomes to their subjective values [7].

$$V(u) = \begin{cases} (u-R)^{\beta^+} & \text{if } u \geq R \\ -\lambda(R-u)^{\beta^-} & \text{if } u < R \end{cases} \tag{3.12}$$

where u is the objective utility of an uncertain prospect and R is the reference relative to which the person measures outcomes as being gains or losses (framing effect). In this thesis, I experimented with several different types of reference points, as will be outlined in section 4.2 later. $\lambda > 1$ is the loss aversion parameter while $0 < \beta^{+} < 1$ and $0 < \beta^{-} < 1$ are the diminishing sensitivity parameters in the gain and loss regimes.

In the discrete case, suppose the SMoDS travel option has n possible outcomes $u_i \in \mathbb{R},\ \forall\ i\ \in\ \{1, \ldots, n\}$ with $u_1 < u_2 \ldots < u_n$, each occurring with probability $p_i\ \in (0, 1)$ and $\sum_{i=1}^{n} p_i = 1$. Then, under EUT, the total objective utility of this option would simply be given by the expectation of the random variable:

$$U^o = \sum_{i=1}^{n} p_i u_i \tag{3.13}$$

which results in the following objective probability of acceptance:

$$p_R^o = \frac{e^{U^o_{SMoDS}}}{\sum_{i\,\in\,\mathcal{I}} e^{U_i^o}} \tag{3.14}$$

where $\mathcal{I}$ is the set of all travel alternatives available to the decision maker for the current trip, including the SMoDS. However, under CPT, the total *subjective* utility is given by a weighted average of the subjective values:

$$U_R^s = \sum_{i=1}^{n} w_i V(u_i) \tag{3.15}$$

where w_i refers to the probability weights assigned by the person to each outcome according to their subjective perception, calculated using:

$$w_i = \begin{cases} \pi[F_U(u_i)] - \pi[F_U(u_{i-1})] & u_i < R \\ \pi[1 - F_U(u_{i-1})] - \pi[1 - F_U(u_i)] & u_i\ \geq R \end{cases} \tag{3.16}$$

where $F_U(u)$ is the cumulative distribution function (CDF) of the objective utility of the uncertain prospect (SMoDS). This can be modelled as being either discrete or

continuous, and can be extracted from historical travel data and demand patterns [1]. $\pi(\cdot)$ is the probability weighting function due to the distortion effect, described by the parameter $0 < \alpha < 1$. We propose using the form for the weighting function proposed by [51].

$$\pi(p) = e^{-(-ln(p))^{\alpha}} \tag{3.17}$$

If F_U is continuous, U_R^s can be calculated in a similar fashion [1]:

$$U_R^s = \int_{-\infty}^{R} V(u) \frac{d}{du} \left\{ \pi \left[F_U(u) \right] \right\} du + \int_{R}^{\infty} V(u) \frac{d}{du} \left\{ -\pi \left[1 - F_U(u) \right] \right\} du \tag{3.18}$$

The subjective probability of acceptance for the SMoDS can then be computed as:

$$p_R^s = \frac{e^{U^s_{R,\,SMoDS}}}{\sum_{i \in \mathcal{I}} e^{U^s_{R,\,i}}} \tag{3.19}$$

3.3 Survey design

3.3.1 Overall information flow

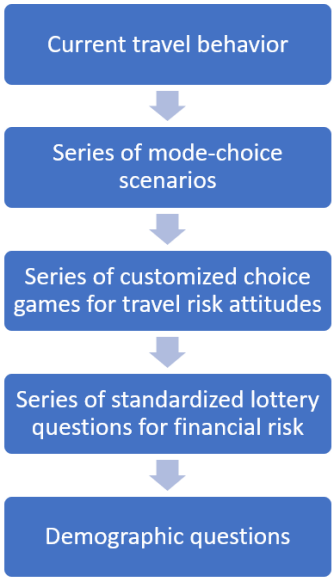

Figure 3-3: Different sections of the final version of the survey.

The survey begins by collecting information on the respondent's current travel behavior. This is used to customize the latter survey sections specifically to each passenger by identifying their most commonly used current mode, preferred type of public transit, commuting distances, current trip costs, CPT reference point (R) etc. The second section consists of several discrete choice scenarios used to estimate the mode choice model. Here, respondents are asked to imagine hypothetical travel scenarios that are as similar as possible to their present daily commutes, in order to ground them and make realistic decisions. In order to control for other effects like weather, time of day, trip purposes etc., these were standardized across all survey respondents. Thus, all respondents were given scenarios involving their commute to work/school between 8-10 AM, with the weather outside being sunny and pleasant. The third section consists of a series of scenarios where the user is asked to choose between SMoDS and their current mode of travel, in order to estimate their CPT travel risk parameters. In the fourth section, we ask ten lottery questions to also determine their financial risk attitudes. Finally, the last sections asks a few demographic questions to explore the heterogeneity of individual-level risk parameters across our sample by income, age, gender, occupation etc.

We went through several iterations of the survey before arriving at this final version. After a small initial pilot with members of our lab, we significantly altered our discrete choice experimental design. This was followed by a soft launch to around 100 respondents from the Mechanical Engineering department at MIT. We made a large number of changes based on feedback and preliminary results from this round. For example, we removed driving as a mode option since there seemed to be a heavy bias towards it among our target audience. Furthermore, we realized that it couldn't be fairly compared against other travel *services* like transit or ride sharing, and there were some complications arising from parking costs. We also added several attention and quality checks throughout. The third iteration of the survey was then widely distributed in the full launch. The latest version of the survey can be found online at https://mit.co1.qualtrics.com/jfe/form/SV_cAU6s5ZxDROOPmC.

3.3.2 Sample size and eligibility

The survey was restricted to the Boston designated market area (DMA) shown in fig. 3-4, with a population of $\approx$ 4.8 million people. Only people between the ages of 18-65 with a valid driver's license were allowed to complete and they had to take it on a desktop or laptop computer. We also enforced age and gender quotas in order to get a nationally representative sample and added in a disclaimer due to COVID-19: Asking the respondents to answer the survey questions keeping in mind their travel behavior *before* the pandemic. Finally, we restricted our respondents to those who used ride sharing services at least a few times a month, in order to ensure that they would be familiar and comfortable with the idea of taking SMoDS.

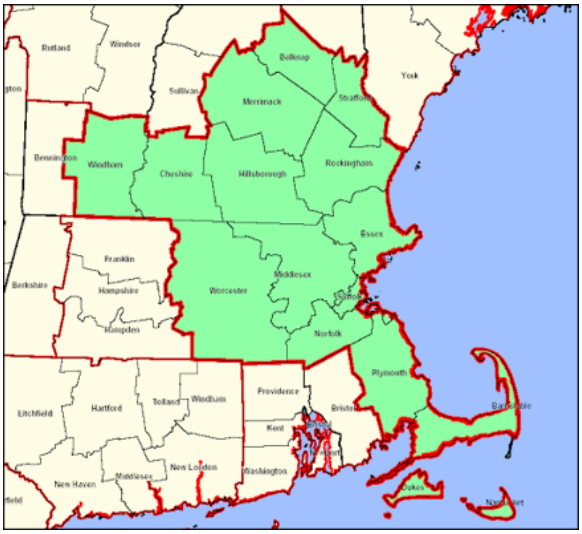

Figure 3-4: Map showing the Boston DMA.

We calculate the desired sample size as follows [1]:

$$N = Z^2 \cdot \frac{\sigma(1-\sigma)}{e^2} \tag{3.20}$$

[1] https://www.qualtrics.com/experience-management/research/determine-sample-size/

where

$$Z = 1.96 \text{ (95\% confidence level)} \tag{3.21}$$

$$e = \pm 3\% \text{ (Margin of error)} \tag{3.22}$$

$$\sigma = 0.5 \text{ (Conservatively assume a large standard deviation)} \tag{3.23}$$

This formula is valid for a large population and we get $N \approx 1000$. We enlisted Qualtrics as our panel firm and collected 955 quality-complete responses over the course of one month, with the average completion time being around 8 minutes.

3.3.3 Discrete choice experimental design

The mode-choice section of the survey consisted of a series of choice scenarios or games where respondents had to choose among the three presented travel options: transit (bus, subway or rail), exclusive ride hailing and pooled ride sharing. These scenarios were constructed by varying the tariffs and travel times per leg (walking, waiting, riding) for each mode, in a systematic manner. Such stated preference studies need to be planned by carefully varying the different attributes (or features) of interest, and their levels so as to accurately estimate the desired effects and test specific hypotheses.

The most straightforward approach is a full factorial design where each level of each attribute is combined with all possible levels of every other attribute being considered in the model. According to eq. (3.10), we have a total of 12 features that can be varied, i.e., 3 travel times (walk, wait, ride) and 1 tariff, for each of the 3 modes. If we considered 3 levels (low, medium, high) for each of these, a factorial enumeration of all these would result in a total of $3^{12} = 531441$ possible combinations or choice situations. Obviously, designing a survey with such a large number of mode choice scenarios is infeasible. Thus, we explored options to reduce the number of choice situations needed.

Firstly, we assumed that travel time coefficients for walking and waiting modes were identical across modes i.e. $a_{walk,i} = a_{walk}$ and $a_{wait,i} = a_{wait}\ \forall\ i$. This was done to simplify the model and reduce the number of independent (explanatory) variables

involved, thereby also decreasing the number of choice situations that we needed to consider in order to fully characterize our discrete choice logit model. In reality, this may not hold exactly since for example, time spent waiting at home for an Uber differs markedly from waiting at a bus stop or subway station. However, we decided to instead focus on differences in the in-vehicle travel time coefficients across modes, since ride time often makes up bulk of the total trip duration and also contributes the most to the overall travel time uncertainty of the SMoDS. This reduced the number of features involved by 4. We also assumed that the disutility coefficient for tariff was equal for all modes, i.e., $b_{transit} = b_{exclusive} = b_{pooled} = b$ since this study is mainly focused on the uncertainty and disutility caused by travel times rather than price. Moreover, the heterogeneity in travel time coefficients across modes would still provide insights on how passengers make tradeoffs between time and price. We further restricted ourselves to varying the prices for only the exclusive and pooled ride sharing options in the design of our mode choice scenarios, and not transit. This is because the costs of taking public transit are generally fixed beforehand and do not vary with time in most cities. Thus, we no longer needed to vary $t_{ride,transit}$ either in our study. Under these assumptions, we have 7 features that need to be changed. However, this still requires $3^7 = 2187$ choice scenarios and it's not practical for survey respondents to make such a large number of decisions. In order to further narrow the choice situations, we used a fractional factorial design.

Fractional factorial design

These involve sampling a specific subset from the complete factorials. Rather than simply sampling in a random, clever sampling techniques can be used to achieve certain desired properties [11]. Since our proposed utility model eq. (3.10) is linear in all the terms (or features) and does not involve any two-way interactions or higher order terms, we decided to use a *main effects screening design*. This is reasonable since main effects generally account for around $70 - 90\%$ of all explained variance in most models [11]. Main effects screening designs have good statistical efficiency properties, in terms of [52]:

1. **Orthogonality**: They are generally either orthogonal or near orthogonal. An orthogonal design is one where every pair of levels occurs equally often across all pairs of factors.

2. **Well-balanced**: Balance is a measure of how close the experimental design is to one where every attribute level occurs equally often within each factor or feature.

The main effects screening design for the mode choice models with 7 factors (having 3 levels each) was constructed using SAS `JMP Pro` software [2]. The resulting design had a total of 18 choice situations and is shown in table 3.1 below.

Choice situation	$\gamma_{exclusive}$	γ_{pooled}	t_{wait}	t_{walk}	$t_{ride,\ transit}$	$t_{ride,\ exclusive}$	$t_{ride,\ pooled}$
1	1	1	3	1	1	1	2
2	1	1	2	3	3	1	1
3	1	3	1	2	1	3	1
4	1	3	3	1	2	3	3
5	1	2	1	3	3	2	3
6	1	2	2	2	2	2	2
7	2	1	1	1	3	3	2
8	2	3	3	2	3	2	2
9	2	1	2	2	2	3	3
10	2	3	2	1	1	2	1
11	2	2	3	3	2	1	1
12	3	2	1	2	1	1	3
13	3	2	3	2	3	3	1
14	3	3	1	3	2	1	2
15	3	1	3	3	1	2	3
16	3	2	2	3	1	3	2
17	3	1	1	1	2	2	1
18	3	3	2	1	3	1	3

Table 3.1: Factor level combinations used in main effects screening design, where the factor levels are coded as 1 = low, 2 = medium and 3 = high.

Following this, I also excluded certain factor combinations that were unrealistic or impossible to achieve in real-life situations e.g. situations where ride pooling was significantly cheaper than exclusive ride hailing but also had far shorter travel times.

[2] https://www.jmp.com/en_us/software/predictive-analytics-software.html

Removing other such scenarios resulted in a final set of 11 combinations of factor levels that were used for the mode choice section of the survey. An illustrative example of one mode choice question can be seen in fig. 3-5.

Travel option	Total trip cost ($)	Walking time (min)	Waiting time (min)	In-vehicle riding time (min)
Public transit: Subway (T)	2.4	14	5	8
Exclusive rideshare	6.4	0	4	12
Pooled rideshare	3.23	4	3	15

◯ Public transit: Subway (T)

◯ Exclusive rideshare

◯ Pooled rideshare

Figure 3-5: Example of mode choice scenario.

3.3.4 Design of travel risk scenarios for SMoDS

The CPT risk preference parameters were determined through the method of *certainty equivalents*. A certainty equivalent (CE) is defined as the guaranteed return that someone would accept now, rather than taking an uncertain (but potentially higher) return in the future. Thus, it can be viewed as being analogous to a risk premium. This can be used to elicit risk preferences when users compare the risky SMoDS option (due to its uncertain travel times) against the mode that they currently use most frequently. We assume that any uncertainty in the transit and exclusive ride hailing options (due to traffic conditions, weather etc.) is negligible in comparison to the SMoDS, where the travel times could vary significantly depending on route changes and other passengers being added to the ride pooling trip in real time. Although the SMoDS could still offer some time and/or cost savings, it depends on various factors like ridership and demand patterns etc. Thus, the current mode is a certain or sure

prospect. The respondent was asked to choose between the uncertain, risky SMoDS and their current baseline travel method which is certain. The key idea behind this survey section is to determine several pairs of prospects (i.e. the baseline current mode vs SMoDS) between which the respondent is indifferent. This allows us to infer the CE and thus set up a system of nonlinear equations with the CPT parameters (α, β^+, β^-, λ) being the unknowns. We designed 6 such choice scenarios which were generated randomly for each respondent, resulting in 6 nonlinear equations per passenger in the sample. An illustrative example of one SMoDS scenario used for CPT risk preferences, can be seen in fig. 3-6.

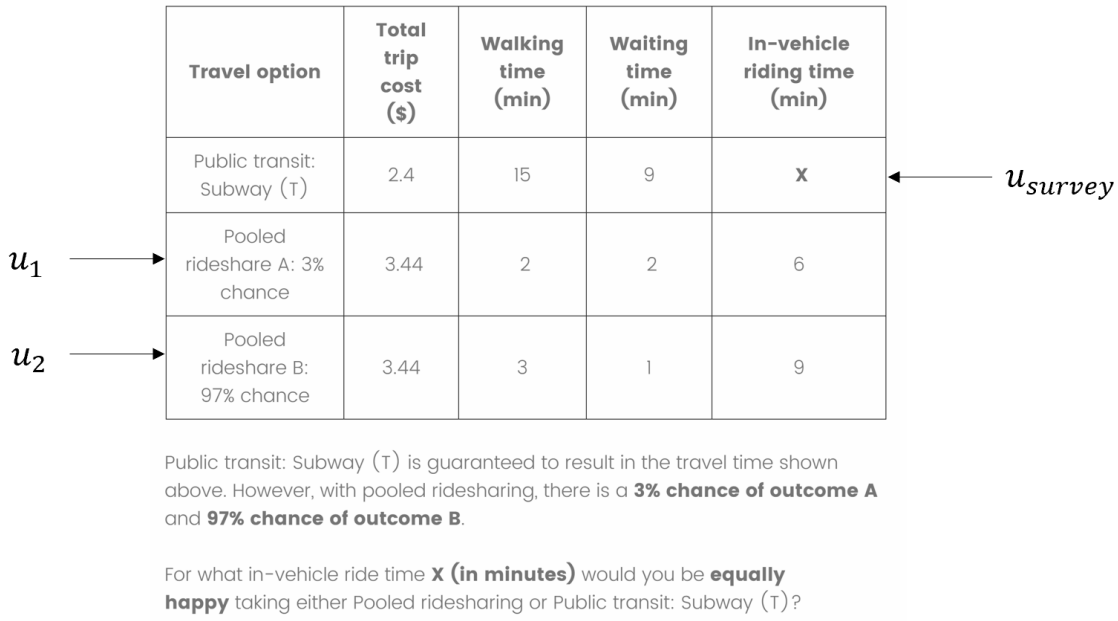

Figure 3-6: Example of CPT risky choice scenario for SMoDS.

Financial risk: Lottery questions

In addition to travel risk, we also added section with 10 standard lottery questions from a prior study [2]. These are also based on a similar idea to certainty equivalents involving two uncertain outcomes, with the 10 questions spanning various possible outcome pairs in the pure gain (both are gains), pure loss (both are losses) and mixed regimes (one gain and one loss). An example from the survey is shown infig. 3-7

below.

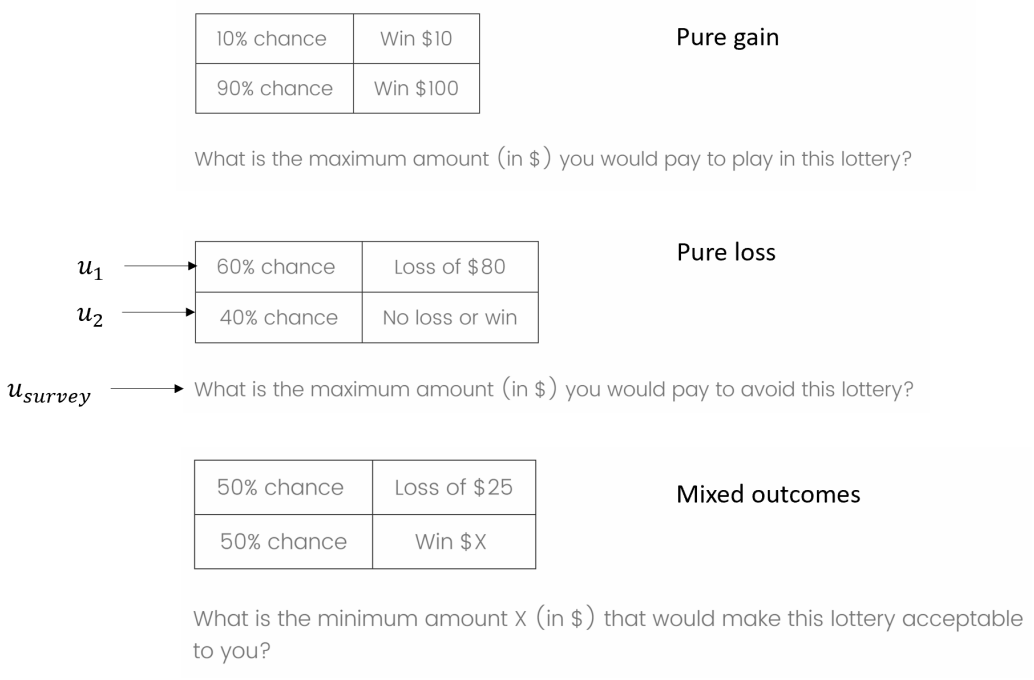

Figure 3-7: Illustrative examples for types of lottery questions used.

This provided a benchmark or ground truth against which we could validate our own SMoDS results, since lotteries are a well established method in the literature to study risk attitudes. Furthermore, it would be interesting to see how the respondents' financial risk preferences compare with their attitudes towards risk in terms of travel times. The lottery scenarios used in our survey are shown in fig. 3-8 below.

3.4 Estimation of mode-choice model

Mixed logit models are generally estimated via maximum likelihood estimation (MLE), which determine the population parameters from which the observed sample of data is most likely to have been generated [11]. The idea behind MLE is to maximum the likelihood function for the observed data that's available. The likelihood function of observing some data $\mathbf{y} = [y_1,\ y_2, \ldots y_n]$ given model parameters $\boldsymbol{\beta}$ is the joint probability density function (pdf), which can be decomposed into a product of the PDFs

Lottery	Outcome A($)	Prob(A)	Outcome B($)	Prob(B)
1	10	0.1	100	0.9
2	0	0.4	100	0.6
3	0	0.1	100	0.9
4	0	0.4	10,000	0.6
5	0	0.9	100	0.1
6	0	0.4	400	0.6
7	−80	0.6	0	0.4
8	−100	0.6	0	0.4
9	−25	0.5	–	0.5
10	−100	0.5	–	0.5

Figure 3-8: Scenarios used for the lottery questions [2].

for individual observations, assuming that they were all draw independently [53]:

$$L(\boldsymbol{\beta}) = P(\mathbf{y} \mid \boldsymbol{\beta}) = P(y_1 \mid \boldsymbol{\beta}) \cdot P(y_2 \mid \boldsymbol{\beta}) \dots P(y_n \mid \boldsymbol{\beta}) \tag{3.24}$$

In the case of our mode choice data, this is the joint PDF of the observed sample of individuals, decisions and alternatives, which can be derived from the choice probabilities [49]:

$$L(\boldsymbol{\beta}) = \prod_k f(y_k \mid \mathbf{x}_k, \boldsymbol{\beta}) = \prod_{\forall i} \prod_{\forall j} (P_{ij}(\boldsymbol{\beta}))^{\delta_{ij}} \tag{3.25}$$

$$\text{where} \quad \delta_{ij} = \begin{cases} 1 & \text{if option } j \text{ is chosen by individual } i \\ 0 & \text{otherwise} \end{cases} \tag{3.26}$$

$$P_{ij}(\boldsymbol{\beta}) = \frac{e^{V_{ij}(\boldsymbol{\beta})}}{\sum_{\forall k} e^{V_{ik}(\boldsymbol{\beta})}} \tag{3.27}$$

where $f(y_k \mid \boldsymbol{\beta})$ is the pdf for the k^{th} observation $(\mathbf{x}_k, y_k)$ in the sample. Here $P_{ij}(\boldsymbol{\beta})$ is the probability that individual i chooses mode j where V_{ij} is objective utility calculated using mode j's utility function for passenger j's trip. Thus, it is a function of the mode choice model parameters represented by $\boldsymbol{\beta}$. We take the natural logarithm of this since it's easier to differentiate and has its maximum at the same point due to

concavity. The log-likelihood function is then given by:

$$LL(\boldsymbol{\beta}) = \sum_{\forall\, i} \sum_{\forall\, j} \delta_{ij}\ ln(P_{ij}(\boldsymbol{\beta})) \tag{3.28}$$

The MLE estimator $\hat{\boldsymbol{\beta}}$ is given by the parameter values that maximize this log-likelihood function, or equivalently minimize the negative of the log-likelihood. For linear utility functions, this function has shown to be globally concave with a unique maximum [10], which can be solved by setting the gradient to zero:

$$\nabla_{\boldsymbol{\beta}} LL(\boldsymbol{\beta}) = 0 \tag{3.29}$$

With linear utility functions of the form $V_{ij} = \boldsymbol{\beta}^{\intercal}\mathbf{x}_{ij}$ as in our mode choice model (eq. (3.10)), this first order condition can be further simplified:

$$\begin{aligned}
LL(\boldsymbol{\beta}) &= \sum_{\forall\, i} \sum_{\forall\, j} \delta_{ij}\ ln(P_{ij}(\boldsymbol{\beta})) && (3.30)\\
&= \sum_{\forall\, i} \sum_{\forall\, j} \delta_{ij}\ ln\left(\frac{e^{\boldsymbol{\beta}^{\intercal}\mathbf{x}_{ij}}}{\sum_{\forall\, k} e^{\boldsymbol{\beta}^{\intercal}\mathbf{x}_{ik}}}\right) && (3.31)\\
&= \sum_{\forall\, i} \sum_{\forall\, j} \delta_{ij}\boldsymbol{\beta}^{\intercal}\mathbf{x}_{ij} - \sum_{\forall\, i} \sum_{\forall\, j} \delta_{ij} ln\left(\sum_{\forall\, k} e^{\boldsymbol{\beta}^{\intercal}\mathbf{x}_{ik}}\right) && (3.32)\\
\nabla_{\boldsymbol{\beta}} LL(\boldsymbol{\beta}) &= \sum_{\forall\, i} \sum_{\forall\, j} \delta_{ij}\mathbf{x}_{ij} - \sum_{\forall\, i} \sum_{\forall\, j} \delta_{ij} \sum_{\forall\, k} P_{ik}\mathbf{x}_{ik} && (3.33)\\
&= \sum_{\forall\, i} \sum_{\forall\, j} \delta_{ij}\mathbf{x}_{ij} - \sum_{\forall\, i} \left(\sum_{\forall\, k} P_{ik}\mathbf{x}_{ik}\right) \sum_{\forall\, j} \delta_{ij} && (3.34)\\
&= \sum_{\forall\, i} \sum_{\forall\, j} \delta_{ij}\mathbf{x}_{ij} - \sum_{\forall\, i} \left(\sum_{\forall\, k} P_{ik}\mathbf{x}_{ik}\right) && (3.35)\\
&= \sum_{\forall\, i} \sum_{\forall\, j} (\delta_{ij} - P_{ij})\mathbf{x}_{ij} && (3.36)
\end{aligned}$$

In addition to the above analytical methods, the log-likelihood function can also be maximized numerically via a standard Newton-Raphson procedure, with the pa-

rameter update step given by:

$$\boldsymbol{\beta}_{t+1} = \boldsymbol{\beta}_t - H_t^{-1} \cdot \nabla_{\boldsymbol{\beta}} LL(\boldsymbol{\beta})\big|_{\boldsymbol{\beta}_t} \tag{3.37}$$

where H_t^{-1} is the inverse of the Hessian matrix of second derivatives of $LL(\boldsymbol{\beta})$ evaluated using the parameters at iteration t.

$$H_t = \left.\frac{\partial^2 LL(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}^2}\right|_{\boldsymbol{\beta}_t} \tag{3.38}$$

This was the method implemented for our study, using the `quasi-newton` algorithm for unconstrained function minimization (`fminunc`) in MATLAB's optimization toolbox.

3.4.1 Maximum simulated likelihood estimation

Standard maximum likelihood estimation requires the use of the exact probabilities of occurrence of the observed choice P_{ij}. These choice probabilities in mixed logit models involve high dimensional integrals (e.g. to compute the normalization factor) which cannot be analytically or exactly solved. Thus, the exact PDFs and log likelihood function do not have a closed form solution in such cases, and must instead be approximately numerically. Moreover, in practical situations it may not be possible to compute this accurately for finite sample sizes with only a limited number of observations being available. MLE estimators on such smaller sample sizes may not have desirable statistical properties like consistency, efficiency or asymptotic normality [10]. An alternative approach that is preferred in such cases is maximum *simulated* likelihood estimation (MSL), where the integral for choice probabilities are calculated via simulation by random draws from the specified mixing distribution [54]. This follows the same procedure as MLE, with the only difference being that simulated probabilities are used instead of the exact ones, resulting in a simulated log likelihood function:

$$SLL(\boldsymbol{\beta}) = \sum_{\forall\, k} ln\ \tilde{f}(\boldsymbol{\beta} \mid \mathbf{x}_k, y_k) = \sum_{\forall\, i} \sum_{\forall\, j} \delta_{ij}\ ln(\tilde{P}_{ij}(\boldsymbol{\beta})) \tag{3.39}$$

where $(\mathbf{x}_k, y_k)$ are the data (input features) and observations in the sample for passengers i and alternatives j, while $\tilde{P}(\boldsymbol{\beta})$ are the simulated approximations to the true, exact probabilities $P(\boldsymbol{\beta})$. $\tilde{f}(y_k|\mathbf{x}_k, \boldsymbol{\beta})$, the simulated approximations to the exact PDFs $f(y_k|\mathbf{x}_k, \boldsymbol{\beta})$ are obtained by taking a large number of repeated, random draws of the model parameters from the specified mixing distribution of $\boldsymbol{\beta}$, for each person (decision maker) in the sample. For a simulation using R draws, this is given by:

$$\tilde{f}(y_k \mid \mathbf{x}_k, \boldsymbol{\beta}) = \frac{1}{S} \sum_{s=1}^{S} g(y_i \mid \mathbf{x}_i, \boldsymbol{\beta}, w^{(s)}) \tag{3.40}$$

where $g(\cdot)$ is the conditional density computed using the drawn simulations $w^{(s)}$ from the mixing distribution. Thus, the MSL estimators describing the optimal mixing distribution for the parameters are obtained by maximizing this simulated log likelihood:

$$\tilde{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta}}{\operatorname{argmax}} \frac{1}{N} \sum_{i=1}^{N} SLL(\boldsymbol{\beta} \mid y_i, \mathbf{x}_i, \mathbf{w}_i) \tag{3.41}$$

$$= \underset{\boldsymbol{\beta}}{\operatorname{argmax}} \frac{1}{N} \sum_{i=1}^{N} \tilde{f}(y_k \mid \mathbf{x}_k, \boldsymbol{\beta}) \tag{3.42}$$

In addition to offering computational benefits, MSL also has nice statistical properties which make it preferable to other simulated-assisted estimation methods such as the method of simulated moments (MSM) and the method of simulated scores (MSS) [10]. When the number of simulations or draws R (per respondent) rises faster than $\sqrt{N}$, where N is the number of observations (samples), then the MSL estimator can be shown to be consistent, asymptotically normal, efficient and equivalent to exact MLE. Thus, as $N \to \infty$ and $R \to \infty$, we need $\sqrt{N}/R \to 0$. However, MSL is inconsistent if R is held constant (even with increasing sample size) and if R increases at a slower rate relative to $\sqrt{N}$, then it is consistent but not asymptotically normal [10]. The software developed by Kenneth Train was used as a starting point to perform the MSL estimation [3].

[3] https://eml.berkeley.edu/~train/software.html

3.4.2 Validation of mode choice results

The estimated mode choice parameters were validated using several methods.

Hypothesis testing

Firstly, a standard one-sample, asymptotic t-test was performed under the null hypothesis that the mean values of all parameters are equal to zero:

$$t = \frac{\hat{\mu} - \mu_0}{SE} \tag{3.43}$$

where $\hat{\mu}$ is the estimated parameter mean (of the chosen mixing distribution) and the hypothesized mean $\mu_0 = 0$. SE is the standard error and n is the total number of respondents in the sample. Since the two sample sets being compared are identical, we use a paired or dependent (correlated) t-test with $n - 1$ degrees of freedom [4]. The value of this t-statistic is used to measure the significance of the corresponding regressor - whether these coefficient means are away from zero in a statistically significant manner. If the p-values are sufficiently small (or equivalently, the t-values are sufficiently large) at the desired confidence level, we can reject this null hypothesis, indicating that the estimated main effects are significant. For example, if $|t| > 1.96$, we can have 95% or greater confidence that the mean is significantly different from zero [11].

Since the sample size is quite large ($n = 955$) we can also use a Z-test, which is often more convenient than the Student's t-test. Here, the assumption is that the test statistic can be approximated by a normal distribution under the null hypothesis. This approximation is reasonable in our case due to the central limit theorem. We can thus perform a two-tailed Z-test to test the alternative hypothesis that the parameter estimate means are non zeros (i.e. $\mu_0 \neq 0$) by computing the Z-score:

$$Z = \frac{\hat{\mu} - \mu_0}{s} = \frac{\hat{\mu}}{\sigma/\sqrt{n}} \tag{3.44}$$

[4] https://www.investopedia.com/terms/t/t-test.asp

where the sample variance σ^2 is substituted in place of the population variance, which is unknown. The variance σ^2 and standard errors for the parameters can be readily computed as part of the estimation process by inverting the Hessian matrix of the log-likelihood function at its optimal, maximizing value (L^*). This gives the variance-covariance matrix of the MLE estimates, with the σ^2 values given by the diagonal elements.

Likelihood ratio test

The likelihood ratio index is another useful measure of the goodness of fit of discrete choice models, to the observed data. It compares how much better a model with the optimal estimated MLE parameters performs, when compared to one where all the parameters (coefficients and constant terms) are set to zero. It is defined as:

$$\rho = 1 - \frac{LL(\hat{\boldsymbol{\beta}})}{LL(\mathbf{0})} \tag{3.45}$$

where $LL(\hat{\boldsymbol{\beta}})$ and $LL(\mathbf{0})$ are the values of the log likelihood function evaluated at the estimated parameters and at zero, respectively. The likelihood ratio index ($0 \leq \rho \leq 1$) ranges from 0 when the MLE parameter estimates do not perform any better than all parameters being zero, to 1 when the estimated model perfectly predict the mode choices of all passengers in the sample [10]. Thus, for models that are estimated on identical samples of data and using the same choice set, we can conclude that the parameters with higher ρ (i.e., closer to 1) has a better fit.

3.5 CPT model estimation

3.5.1 Data pre-processing

Prior to estimation, the data was processed in order to screen out potentially erroneous respondents from the final data set. This was necessary because in spite of the quality and attention checks incorporated in the survey, it became evident that

some respondents either fully understand the questions or did not put enough effort into giving thoughtful responses. We used the responses to the financial risk lottery questions from fig. 3-8 in order to remove the invalid responses and prevent them from corrupting the estimates. In order to deem responses as valid, we used the procedure outlined in a previous study [2] to see if they satisfied basic axioms of rational choice:

1. **Internal validity**: This is violated if the certainty equivalent (CE) of a lottery is chosen to be less than the smallest possible outcome or larger than the largest possible outcome. In our survey design, such errors are automatically avoided by enforcing lower/upper limits on the slider bar used.

2. **Monotonicity of probabilities**: For pairs of lotteries that have the same outcomes but assign different probabilities, the CE of the lottery with better winning chances should be higher (or at least equal). Thus, the lotteries of interest from fig. 3-8 are 2, 3 and 5.

3. **Monotonicity of outcomes**: For pairs of lotteries that have the same probabilities but different associated outcomes, the CE of the lottery with the better expected outcome should be higher (or at least equal). Thus, the lotteries of interest here are 7 and 8.

If any one of these conditions does not hold true, first-order stochastic dominance is violated. Then, the individualized CPT risk parameters were only estimated for the subset of respondents who satisfied all three of the above conditions.

3.5.2 Detection of CPT-like behaviors

Before quantitatively estimating the model parameters, a qualitative analysis was performed to see if the respondents actually exhibited CPT-like behaviors that cannot be explained by expected utility theory (EUT) alone. This was also done using their answers to the lottery questions, following a similar approach to [2].

Reflection effect

This effect refers to the person being risk-averse in gains and risk-seeking in the losses regime, leading to a convex-concave structure for the value function $V(\cdot)$. This effect can be tested by comparing *symmetric* gain and loss lotteries, such as questions 2 and 8 from fig. 3-8. These have the same probabilities and magnitudes for both outcomes but lottery 2 is a pure gain scenario while lottery 8 is pure losses. Thus, respondents who were found to be risk averse in lottery 2 but behaved in a risk seeking manner in lottery 8, were considered to be displaying the reflection effect. In our specific scenario, this implies that the maximum amount the respondent was willing to pay to either (i) play the gain lottery or (ii) not participate in the loss lottery, was less than the lottery's expected value or mean outcome of \$60.

Probability distortion effect

Finally, we explore whether the respondents' actually show probability distortion, which is the probably the most novel aspect of CPT and PT based models. For this study, we focus on probability weighting effects in the gain regime since 6 out of the 10 lottery questions were gains, but similar analyses can also be conducted for losses. The following procedure from [55] was used to detect *over weighting of small probabilities* at an individual level, by studying the lotteries 2, 3 and 5 from fig. 3-8.

Given two lotteries A and B, both with binary outcomes 0 and $X > 0$. A provides a payoff of \$X with a probability of p_A while B does so with probability p_B. Let's suppose that the survey respondent can be described by CPT but does not distort probabilities, i.e., they only have the reflection effect. Then, by definition of the certainty equivalent (CE), the user is indifferent between a guaranteed payoff of CE

(i.e. with probability 1) versus the expected payoff from the lottery:

$$\begin{aligned} & V(CE_A) = p_A V(X),\ V(CE_B) = p_B V(X) \\ \implies & \frac{V(CE_A)}{p_A} = \frac{V(CE_B)}{p_B} = V(X) \\ \implies & \frac{V(CE_A)}{V(CE_A)} = \frac{p_A}{p_B} \end{aligned}$$

Thus, without probability weighting, we would expect the ratio of the certainty equivalences for such lottery pairs to be a fixed value. If this is not the case, it indicates that the respondent has at least some degree of probability weighting. Specifically, for the lotteries under consideration, we have $X = \$100$, $p_2 = 0.6$, $p_3 = 0.9$ and $p_5 = 0.1$. Thus, probability overweighting occurs if one or more the below conditions are satisfied:

$$\frac{V(CE_5)}{V(CE_2)} > \frac{p_5}{p_2} = \frac{1}{6} \implies \text{Overweighting between 10 and 60\% probability} \quad (3.46)$$

$$\frac{V(CE_2)}{V(CE_3)} > \frac{p_2}{p_3} = \frac{2}{3} \implies \text{Overweighting between 60 and 90\% probability} \quad (3.47)$$

$$\frac{V(CE_5)}{V(CE_3)} > \frac{p_5}{p_3} = \frac{1}{9} \implies \text{Overweighting between 10 and 90\% probability} \quad (3.48)$$

Loss aversion

Although risk averse behavior is also one of the features of conventional utility theory, the extent of it observed in situations with uncertainty (such as lotteries and SMoDS) can often not be justified with even fairly high levels of loss aversion in EUT model [2]. By inspecting the responses to lotteries 9 and 10 where both losses (outcome A) and gains (outcome B) are equally likely, a large value of the the ratio A/B (i.e. significantly higher than 1) implies that the respondents require proportionately much higher gain outcomes to compensate for losses even they both occur with equal probability. Thus, prospect theory is likely a more accurate and realistic representation of such high degrees of loss aversion in the their preferences.

3.5.3 Quantitative estimation of parameters

We now proceed to determine the individual level CPT parameters that describe the risk preferences of each individual in the sample. Note that only valid responses as determined in section 3.5.1 were used for this analysis. The survey questions in the travel risk section determine, for each respondent, several pairs of prospects (i.e. their baseline current travel mode vs. the SMoDS) between which they are indifferent, by eliciting the certainty equivalent. This allows us to equate the subjective utilities derived from both alternatives and the resulting nonlinear system can be solved to estimate the model parameters.

The true certainty equivalent of the respondent can be computed using their answers u_{survey} as shown in the example questions from fig. 3-6 and fig. 3-7:

$$CE_{true} = U^s_{R,\,SMoDS} = V(u_{survey}) \tag{3.49}$$

Since the user is indifferent between the certain alternative (denoted as A) and the SMoDS at this point, and if they do indeed follow the proposed CPT model, we would expect this to be equal to the subjective utility offered by SMoDS. This can be viewed as the predicted certainty equivalent for that respondent:

$$CE_{pred} = U^s_{R,\,A} = \widehat{U}^s_{R,\,SMoDS} \tag{3.50}$$

$$= w_1 V(u_1) + w_2 V(u_2) \tag{3.51}$$

where we applied the CPT model equations as described in eq. (3.12) - eq. (3.17), which have simplified considerably here since these survey scenarios only consider two alternatives and the SMoDS is assumed to have only two possible outcomes u_1 and u_2 (i.e., a Bernoulli distribution) as shown in fig. 3-6. The error or residual for the i^{th} scenario is then defined as the difference between the predicted certainty equivalent

and the actual value determined from the survey:

$$e_i = (CE_{true})_i - (CE_{pred})_i = \left(U^s_{R,\,SMoDS}\right)_i - \left(\widehat{U}^s_{R,\,SMoDS}\right)_i \tag{3.52}$$

$$= f_i(\alpha, \beta^+, \beta^-, \lambda) \tag{3.53}$$

Each of the six scenarios results in an error equation as in eq. (3.53), producing a nonlinear system for each respondent with the unknowns being the four CPT parameters of interest. The exact form of the equations will depend on several factors specific to each scenario such as the choice of reference R, the CDF $F_U(u)$ of the SMoDS' travel times and objective utilities, etc. To see an example of the complex nonlinearities involved, consider a simplified special case with only two outcomes for the SMoDS: one gain and one loss i.e., $(u_1,\ u_2) \sim (p,\ 1-p)$ with $u_1 < u_2$. Further, assume that we use the expected value of the SMoDS as our reference:

$$R = pu_1 + (1-p)u_2 \implies u_1 < R < u_2 \tag{3.54}$$

Finally suppose that the objective utility of the user's response $u_{survey} > R$, thus perceived as a gain and this is a certain prospect. Then, the nonlinear equations for this scenario are:

$$w_1 = \pi(F_U(u_1)) = \pi(p) = e^{-[-ln(p)]^\alpha} \tag{3.55}$$

$$w_2 = \pi(1 - F_U(u_1)) - \pi(1 - F_U(u_2)) = \pi\left(1 - e^{-[-ln(p)]^\alpha}\right) \tag{3.56}$$

$$= e^{-\left[-ln\left(1-e^{-[-ln(p)]^\alpha}\right)\right]^\alpha} \tag{3.57}$$

$$U^s_R = (u_{survey} - R)^{\beta^+} \tag{3.58}$$

$$\widehat{U}^s_R = -\lambda e^{-[-ln(p)]^\alpha}(R - u_1)^{\beta^-} + e^{-\left[-ln\left(1-e^{-[-ln(p)]^\alpha}\right)\right]^\alpha}(u_2 - R)^{\beta^+} \tag{3.59}$$

The nonlinear error equations can be similarly derived for the lottery questions as well, using the survey responses to the constructed scenarios as illustrated in fig. 3-7. This results in a analogous nonlinear system in four unknown, but having 10 equations per respondent. In the case of the travel scenarios for the SMoDS, it is not obvious

what the correct reference point R should be and this could vary depending on the respondent as well as the trip characteristics. Thus, we have flexibility in the choice of R. However, in the case of the lottery questions, $R = \$\ 0$ (i.e. neither gain nor loss) is the natural reference choice that also makes intuitive sense. Furthermore, the lottery survey responses and monetary gains/losses can directly be used as a measure of objective utility without requiring any intermediate transformations. However, for the travel questions, the answers need to first be converted to utilities using the mode-specific utility functions. For example, if the survey question ask for the in-vehicle ride time of public transit X (min) at certainty equivalence as in fig. 3-6, then the utility corresponding to the response is given by:

$$u_{survey} = a_{walk,\,tr}\; t_{walk,\,tr} + a_{wait,\,tr}\; t_{wait,\,tr} + a_{ride,\,tr}\; X + b\gamma_{tr} + ASC_{tr} \tag{3.60}$$

Nonlinear least squares

After setting up the nonlinear system as in eq. (3.53) for each passenger, we can estimate the individual-level CPT parameters by minimizing the error (i.e. the sum of squared residuals) across all the scenarios considered subject to the upper and lower bound constraints on the risk parameters. This can be formulated as a constrained optimization problem for passenger j in the sample:

$$\min_{\alpha_j,\,\beta_j^+,\,\beta_j^-,\,\lambda_j} \|\mathbf{e}_j\|_2^2 = \sum_{\forall\, i} \left(f_i(\alpha, \beta^+, \beta^-, \lambda)\right)^2 \tag{3.61}$$

$$\text{s.t. } 0 < \alpha, \beta^+, \beta^- < 1,\ \lambda > 1 \tag{3.62}$$

where $\mathbf{e}_j$ is the vector of error equations across all the questions for respondent j. The index i loops over all 6 scenarios per respondent for the travel scenarios, and over all 10 questions for the lotteries. This is an example of a nonlinear least squares data-fitting problem and can be efficiently solved using the `lsqnonlin` function from MATLAB's optimization toolbox. However, it does not guarantee global convergence. Thus, I adjusted several hyperparameters of this solver in order to avoid getting stuck at local minima, as detailed in section 4.2.4. I also tried a few other methods to solve

the above minimization problem.

- Grid search: An exhaustive grid search was performed over the entire parameter space being considered, i.e. $\alpha, \beta^+, \beta^- \in (0, 1)$ and $\lambda \in (1, 100)$. Even with a very fine grid spacing, we were not able to reduce the squared norm of the estimation significantly below the values obtained by nonlinear least squares. Furthermore, this brute-force method if much more computationally expensive and takes far longer to complete.

- Multidimensional Newton's method: This root-finding method was also implemented to solve the nonlinear system. However, it is much slower to converge, mainly due to the expensive step of computing the Jacobian matrix at each iteration. Furthermore, the algorithm also fails if the Jacobian is close to being singular near the desired solution.

- Local solvers: Other local MATLAB solvers like `fmincon` (for minimizing constrained nonlinear multivariable functions) provided either comparable or inferior results to `lsqnonlin`.

- Global solvers: I experimented with several solvers from MATLAB's Global Optimization toolbox, such as global search, multiple starting point search, particle swarm optimization, genetic (evolutionary) and pattern search algorithms. Our objective function is highly nonlinear, non convex and also likely to be non smooth in the variable space of the CPT parameters. For minimizing such functions subject to bound constraints, the genetic, particle swarm and direct pattern search algorithms are generally expected to be most well-suited to reach a single global solution. However, these were found to be much slower to run than nonlinear least squares, especially when extended to the full sample of 955 respondents. Although these do theoretically guarantee convergence to global minima, the results obtained were not significantly any better than those from `lsqnonlin` in terms of reducing the squared prediction error in the certainty equivalents. In fact in some cases, the results obtained by these solvers were

even worse than those obtained using `lsqnonlin`, in terms of the distributions of CPT parameters across the sample. Thus, the marginal improvements in the quality of results (if any) were outweighed by the increased computational burden.

Thus, going forward, I primarily decided to use the nonlinear least squares solver instead of any of the alternative methods listed above. It offered a good balance between good quality of results, relatively low estimation errors and faster runtimes. A pictorial summary of our overall estimation process, going from raw survey data to final mode choice utility functions and CPT parameters, is shown in fig. 3-9.

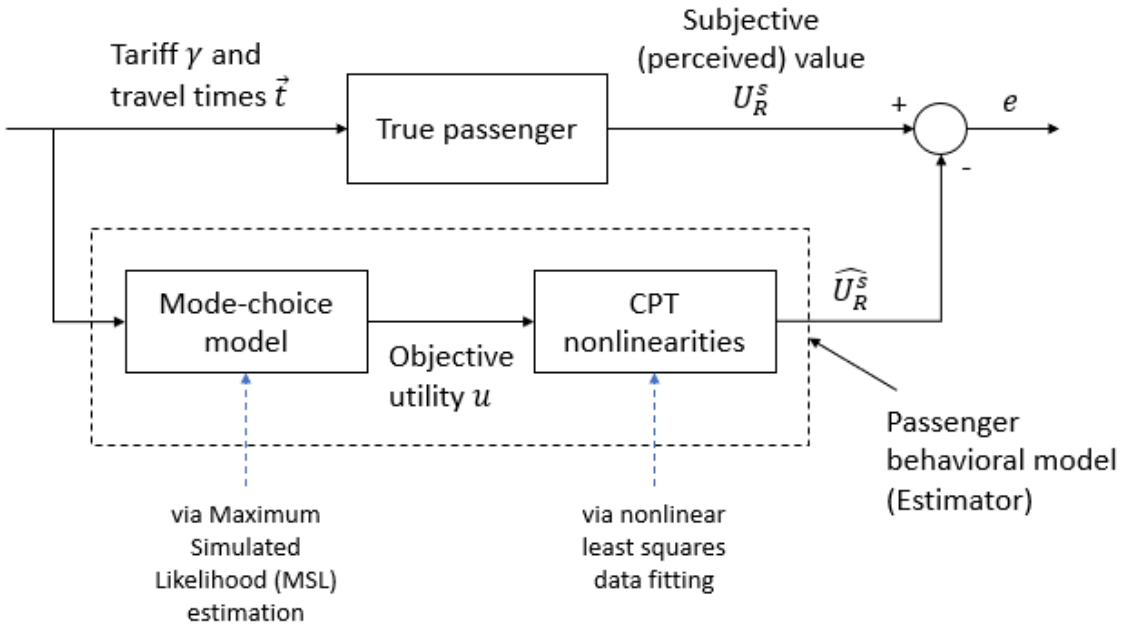

Figure 3-9: Illustration of our proposed estimation process for the SMoDS CPT model.

3.6 Parallelization and high performance computing

During both the mode choice and CPT model estimation steps, I conducted a large number of computational experiments to optimize various design choices and hyperparameters. Since both the maximum simulated likelihood and nonlinear least

squares tasks involve relatively high computational effort, I also attempted to use parallel and high performance computing techniques in order to efficiently perform these parameter sweep tasks. As a starting point, I modified my programs to be able to run in parallel on multiple cores (or workers). For the mode choice task, I executed independent MSL estimation iterations in parallel each using different hyperparameters such as the number of draws or simulations used, types of random draws made, mixing distributions and starting values for the simulations. In addition, the solver settings used for unconstrained minimization of the negative log-likelihood function were set to compute the loss gradients in parallel.

For the CPT estimation task, I was able to parallelize the nonlinear least squares estimation over all the respondents in the sample since these are all independent programs using different data to obtain estimates specific to each passenger. Furthermore, the gradient computations in `lsqnonlin` could also be parallelized similar to above. Finally, I was able to also use parallel computing to more quickly run an outer loop where we modified certain hyperparameters such as the regularization term used in the nonlinear optimization.

I then tested its performance both with and without enabling parallelization on my local machine - a Dell XPS 15 9560 laptop with a 2.8GHz Intel Core i7-7700HQ (3.8GHz boost) having 4 cores and 8 threads, alongside an Nvidia GTX 1050 GPU with 4GB RAM. This was done seamlessly using MATLAB's parallel computing toolbox. After verifying that a speedup was in fact obtained, I submitted these *high-throughput computing* tasks as batch jobs to the MIT Supercloud[5], a remote HPC cluster operated by the MIT Lincoln Laboratory [56]. The code had to be slightly modified to run on this cluster since it uses `pMatlab`, an MPI-based parallel toolbox for multicore and multinode systems [6].

[5] https://supercloud.mit.edu/
[6] https://www.mit.edu/~kepner/pMatlab/

3.7 Sensitivity and robustness analysis of CPT model

3.7.1 Problem formulation

This study will focus on a single passenger taking a trip using the SMoDS. Without loss of generality, we consider a case where the passenger chooses between only two modes of transport: the uncertain SMoDS sm against a baseline travel alternative A (e.g. public transit, driving or exclusive door-to-door ridesharing) that can be treated as a certain prospect, offering a fixed objective utility u_o.

The objective utility of a trip with a certain travel option is calculated using a linear multinomial logit choice mode, based on the travel times spent on each leg, tariff γ and the alternative-specific constant (ASC) c of the service:

$$u = \mathbf{a}^\intercal \, \mathbf{t} + b\gamma + c \tag{3.63}$$

$$= x + b\gamma \tag{3.64}$$

where $\mathbf{t} = [t_{walk}, t_{wait}, t_{ride}]^\intercal$ denote the walking, waiting, and riding times, respectively and $\mathbf{a} = [a_{walk}, a_{wait}, a_{ride}]^\intercal$ are the travel time coefficients for each leg. Here, x is used to compactly represent the component of objective utility due to all travel times on different legs combined along with the ASC of the travel mode. The coefficients on travel times ($\mathbf{a}$) and tariff (b) are negative since these represent disutilities to the consumer, while c can be either positive or negative depending on the characteristics of the given travel option.

For simplicity, for any given ride offer, the possible outcomes with the SMoDS are modelled as following a Bernoulli distribution, i.e., it is assumed to have only two possible travel time outcomes $\overline{\mathbf{t}}$ and $\underline{\mathbf{t}}$ ($\underline{\mathbf{t}} \leq \overline{\mathbf{t}}$) having corresponding utilities $\underline{u}$ and $\overline{u}$ ($\underline{u} \leq \overline{u}$), occurring with probabilities of $p \in [0, 1]$ and $1 - p$ respectively. This choice of distribution is a reasonable starting point and makes the problem more tractable. Moreover, accurately estimating probability distributions for travel times expected by each passenger would require a very large number of draws for each respondent in the population. Thus, any distribution fitted using data from a reasonably large

sample size will still involve some finite error.

We also assume that both outcomes offer the same price since most ridesharing services guarantee trip tariff at the time of ride offer. The analysis we present below can readily be extended to situations with multiple travel alternatives. This framework can also be applied to model more complex probability distributions for SMoDS travel times, having more than two outcomes. These include both continuous (e.g. Gaussian, extreme value) and discrete (e.g. Poisson) distributions. The specific assumptions made here are primarily for simplicity and tractability while deriving analytical results, the broader insights and trends also hold for more general cases. The utilities of the SMoDS and alternative are then given by:

$$\underline{u} = \mathbf{a}_{sm}^{\intercal}\, \overline{\mathbf{t}} + b_{sm}\gamma_{sm} + c_{sm} = \underline{x} + b\gamma \tag{3.65}$$

$$\overline{u} = \mathbf{a}_{sm}^{\intercal}\, \underline{\mathbf{t}} + b_{sm}\gamma_{sm} + c_{sm} = \overline{x} + b\gamma \tag{3.66}$$

$$u_o = \mathbf{a}_o^{\intercal}\, \mathbf{t}_o + b_o\gamma_o + c_o \tag{3.67}$$

If $u_o \leq \underline{u}$, the customer would always choose the SMoDS since it offers strictly better outcomes and conversely if $u_o \geq \overline{u}$, they would always choose option A. Thus, the only cases considered are where $\underline{u} \leq u_o \leq \overline{u}$ are considered (note: u_o can still be either a gain or loss) such that the consumer's choice (of accepting or rejecting the SMoDS ride offer) is non-trivial. Given that the SMoDS outcomes follow a Bernoulli distribution, its cumulative distribution function (CDF) is defined on the support $[\underline{u}, \overline{u}]$:

$$F_U(u) = \begin{cases} 0 & \text{if } u < \underline{u} \\ p & \text{if } \underline{u} \leq u < \overline{u} \\ 1 & \text{if } u \geq \overline{u} \end{cases} \tag{3.68}$$

3.7.2 CPT model overview

The study focuses on analyzing the model's sensitivity to the CPT model parameters as follows. The reader is referred to [1] for more details regarding the CPT based passenger behavioral model in the SMoDS context:

1. **CPT parameters**: It assumed for simplicity that $\beta^+ = \beta^- = \beta$ i.e. the passenger displays similar reductions in sensitivity while moving away from the reference value, in both the gain and loss regime. The sensitivity with respect to these parameters is computed in the standard sense by allowing continuous variations in their values.

2. **Reference (R)**: Treated as a hyper-parameter and case studies using a few different reference types are considered:

   - Static: Fixed for each customer, independent of the SMoDS ride offer. This could be set as the objective utility of the most frequently used travel alternative (excluding SMoDS) i.e. $R = u_o$.

   - Dynamic: R depends on the uncertain prospect itself i.e. it varies with the SMoDS offer. Some examples considered here are:

     – Expected utility of SMoDS

     $$R = p\underline{u} + (1-p)\overline{u} \tag{3.69}$$

     – Best ($R = \overline{u}$) or worst-case ($R = \underline{u}$) utilities corresponding to the shortest ($\underline{\mathbf{t}}$) and longest ($\overline{\mathbf{t}}$) travel times, respectively.

3. **Probability distributions**: These are the probabilistic distributions of expected travel times spent on different legs (i.e. walking, waiting and riding), as perceived by the users themselves. In the current study, we model the objective utility of the SMoDS as a binary random variable. Thus, this distribution can be varied by altering the parameter p which is the probability of the worst-case SMoDS outcome ($\underline{u}$) occurring.

The subjective utilities of the SMoDS (i.e. U_R^s) and A (i.e. A_s^R) can then be calculated as:

$$A_R^s = \pi(1) \cdot V(u_o) = V(u_o) \tag{3.70}$$

$$U_R^s = w_1 \cdot V(\underline{u}) + w_2 \cdot V(\overline{u}) \tag{3.71}$$

where w_1 and w_2 are subjective probability weights calculated using:

$$w_i = \begin{cases} \pi[F_U(u_i)] - \pi[F_U(u_{i-1})] & u_i < R \\ \pi[1 - F_U(u_{i-1})] - \pi[1 - F_U(u_i)] & u_i \geq R \end{cases} \tag{3.72}$$

where $u_1 = \underline{u} \leq u_2 = \overline{u}$ and $F_U(\cdot)$ is the cumulative distribution function (CDF) of the SMoDS utility. The subjective probability of acceptance can be calculated using:

$$p_s^R(\gamma; \vec{\theta}) = \frac{e^{U_R^s}}{e^{U_R^s} + e^{A_R^s}} \tag{3.73}$$

where $\vec{\theta} = [\alpha, \beta, \lambda, p, R]^\intercal$ consists of all the parameters of interest, assembled together. From eq. (1.1) and equations (3.64)-(3.73), it is easy to see that p_s^R, the main output of the CPT-based passenger behavioral model is a function of the SMoDS tariff γ parametrized by $\vec{\theta}$. In the following, we will evaluate this model's sensitivity with respect to these key parameters by formulating the problem as one of nonlinear optimization.

3.7.3 Optimization

The dynamic tariff γ is set by solving a constrained, nonlinear parametric optimization problem. We can consider several possible objective functions. For example, maximizing expected ridership for the fleet would be equivalent to directly maximizing acceptance probability p_s^R itself. If we maximize expected revenue, the below NLP

results:

$$\min_{\gamma} \quad -f(\gamma;\vec{\theta}) \triangleq -\gamma \cdot p_s^R(\gamma;\vec{\theta}) \tag{3.74}$$

$$\text{s.t. } g^1: \quad \underline{\gamma} - \gamma \leq 0 \tag{3.75}$$

$$g^2: \quad \gamma - \overline{\gamma} \leq 0 \tag{3.76}$$

eq. (3.74) reflects the expected revenue because the revenue per passenger per trip is γ with a probability p_R^s and 0 with probability $1-p_R^s$. Yet another option is to perform a weighted multiobjective optimization considering multiple objectives like revenue, ridership and utilization. The sensitivity analysis is then aimed at understanding how changes in parameters affect the optimal dynamic tariff γ^* and value of the objective function $f(\gamma;\vec{\theta})$, where γ is the decision variable and $\vec{\theta}$ represents all the model parameters. Note that the constraints in eq. (3.75) and eq. (3.76) only include upper and lower bound on the dynamic tariff charged such that it is within a reasonable range, i.e., $\gamma \in [\underline{\gamma}, \overline{\gamma}]$. All other constraints related to travel times have already been accounted for by the routing algorithm in generating the SMoDS ride offer and its possible outcomes. In practice, the lower bound could be the minimum break-even price per trip and the upper bound could be some sensible limit e.g. SMoDS tariff cannot be higher than that of exclusive ridesharing.

3.7.4 Optimality conditions

In the following, subscripts indicate partial derivatives w.r.t. that variable. The Lagrangian dual for the NLP formulated in eq. (3.74)-(3.76) is [57]:

$$\mathcal{L}(\gamma;\vec{\theta}) = -f(\gamma;\vec{\theta}) + \mu_1 \cdot (\gamma - \overline{\gamma}) + \mu_2 \cdot (\underline{\gamma} - \gamma) \tag{3.77}$$

The Karush-Kuhn-Tucker (KKT) conditions for the optimal point γ^* are:

1^{st} order necessary conditions

$$\frac{\partial \mathcal{L}}{\partial \gamma} = -f_{\gamma}(\gamma^*; \vec{\theta}) + \mu_1^* - \mu_2^* = 0 \tag{3.78}$$

Complementary slackness conditions

$$\mu_1^* \cdot (\gamma^* - \overline{\gamma}) = 0 \ , \ \mu_2^* \cdot (\underline{\gamma} - \gamma^*) = 0 \tag{3.79}$$

Primal problem feasibility

$$\mu_1^* \ , \mu_2^* \geq 0 \tag{3.80}$$

If one of the constraints g^i is active at the nominal optimum, then its corresponding multiplier $\mu_i > 0$ by strict complementarity, and if inactive, then $\mu_i = 0$. In addition to the above necessary conditions, the following Strong 2^{nd} Order Sufficient Condition (SSC) guarantees that γ^* is a local minimum of the NLP, even if it is non-convex. The Hessian of the Lagrangian must be positive definite on the null space of the Jacobian of active constraints g^a [58]:

$$\nu^{\intercal} \, \mathcal{L}_{\gamma\gamma} \, \nu > 0 \quad \forall \, \nu \neq 0 \quad \text{s.t.} \quad g_{\gamma}^1 \, \nu = 0 \quad \text{or} \quad g_{\gamma}^2 \, \nu = 0 \tag{3.81}$$

Since at most one constraint can be active, this implies that either $\nu = 0$ or $-\nu = 0$ and thus the SSC holds automatically if either the lower (eq. (3.75)) or upper (eq. (3.76)) bound is active. If neither constraint is active, then the SSC requires positive definiteness of $\mathcal{L}_{\gamma\gamma} = -f_{\gamma\gamma}$ over all possible values of γ in the domain. It can be shown that the objective function $f(\gamma; \vec{\theta})$ is concave in γ as long as a specific condition holds on the parameters, price and travel times. For instance, if $R = \overline{u}$, this condition turns out to be:

$$\begin{aligned} e^{-\lambda(\overline{u}-u_0)^{\beta}} \left(e^{-e^{-\lambda(\overline{u}-\underline{u})^{\beta}(-ln(p))^{\alpha}}} + \gamma\lambda\beta b_{sm}(\overline{u}-u_0)^{\beta-1} \right) \\ \leq -\left(e^{-e^{-\lambda(\overline{u}-\underline{u})^{\beta}(-ln(p))^{\alpha}}} \right)^2 \end{aligned} \tag{3.82}$$

This implies $\mathcal{L}_{\gamma\gamma} \geq 0$ and the NLP reduces to a convex optimization problem for which KKT conditions are sufficient for both local and global minima.

3.7.5 Local sensitivity analysis

Suppose the *nominal* problem with assumed parameters θ_0 has optimal tariff γ_0^* and optimal objective value $f(\gamma_0^*; \theta_0) = f_0^*$. If the actual parameters turn out to be $\hat{\theta}$, we consider how this impacts the optimal tariff and optimal objective for the perturbed problem. Here, we consider cases where only one of the parameters is perturbed at a time while keeping the others fixed. This doesn't account for how interactions between parameters may influence the objective function. Local sensitivity analysis considers relatively small perturbations or uncertainties in the parameters for which the active set remains constant. Following [58], local sensitivity differentials can be derived analytically in the neighbourhood of the nominal optimum operating point $(\gamma_0^*, \vec{\theta_0})$, considering variations in a single parameter θ:

$$\begin{bmatrix} \frac{d\gamma^*}{d\theta} \\ \frac{d\mu^a}{d\theta} \end{bmatrix} = - \begin{bmatrix} \mathcal{L}_{\gamma\gamma} & g^{a\intercal} \\ g^a & 0 \end{bmatrix}^{-1} \begin{bmatrix} \mathcal{L}_{\gamma\theta} \\ g^a \end{bmatrix} \tag{3.83}$$

where all the quantities are evaluated at the nominal values γ_0^*, θ_0 and μ_0^a. This gives us local sensitivity derivatives of both the optimal solution $\gamma^*(\theta)$ and multipliers μ^a corresponding to inequality constraints g^a active at the nominal optimum. Furthermore, $\frac{d\mu^{ina}}{d\theta} = 0$ for all inactive constraint multipliers. If neither g^1 nor g^2 is active, we get:

$$\frac{d\gamma^*}{d\theta} = -\mathcal{L}_{\gamma\gamma}^{-1} \mathcal{L}_{\gamma\theta} \tag{3.84}$$

The 1^{st} order sensitivity of the optimal objective yields [58]:

$$\frac{df^*}{d\theta}(\gamma(\theta); \theta)|_{\theta=\theta_0} = \mathcal{L}_\theta(\gamma_0^*, \mu_0^a, \theta_0) \tag{3.85}$$

3.7.6 Real-time approximations by Taylor expansions

We can also approximate the perturbed optimal solution $\gamma^*(\theta)$ and objective function f^* in the neighbourhood of the nominal optimum, using 1^{st} order Taylor expansions about this operating point:

$$\gamma^*(\theta) = \gamma_0^* + \frac{d\gamma^*}{d\theta}(\theta_0)(\theta - \theta_0) \tag{3.86}$$

$$f^*(\theta) = f_0^* + \frac{df^*}{d\theta}(\theta_0)(\theta - \theta_0) \tag{3.87}$$

We can measure the quality of this approximation by comparing it with the 2^{nd} order expansion [58]:

$$\begin{aligned} f^*(\theta) &= f_0^* + \frac{df^*}{d\theta}(\theta_0)(\theta - \theta_0) + \frac{d^2 f^*}{d\theta^2}(\theta_0)(\theta - \theta_0)^2 \\ \frac{d^2 f^*}{d\theta^2}(\theta_0) &= \mathcal{L}_{\gamma\gamma}(\theta_0)\left(\frac{d\gamma^*}{d\theta}(\theta_0)\right)^2 + 2\mathcal{L}_{\gamma\theta}\frac{d\gamma^*}{d\theta}(\theta_0) + \mathcal{L}_{\theta\theta} \end{aligned} \tag{3.88}$$

3.7.7 Prediction of local domain

The analysis in section 3.7.5 assumes that the perturbation does not alter the set of active constraints. We can estimate the largest allowable magnitude of such parameter changes that still preserves the active set. 1^{st} order estimates for the Lagrange multipliers are used to determine when their corresponding constraints become either slack or tight, as in eq. (3.89) and eq. (3.90) respectively.

1. When a constraint leaves the active set, its non-zero multiplier becomes zero:

$$\begin{aligned} \mu^a(\theta) &\approx \mu^a(\theta_0) + \frac{d\mu^a}{d\theta}(\theta_0)(\theta - \theta_0) = 0 \\ (\hat{\theta} - \theta_0)_{max} &= \Delta\theta_{max} = -\frac{\mu^a(\theta_0)}{\frac{d\mu^a}{d\theta}(\theta_0)} \end{aligned} \tag{3.89}$$

2. When a constraint enters the active set, it becomes tight and equal to zero.

$$g_{ina}(\gamma;\theta) \approx g_{ina}(\gamma_0^*;\theta_0) + \frac{dg_{ina}}{d\theta}(\gamma_0^*,\theta_0)(\theta-\theta_0) = 0$$
$$(\hat{\theta}-\theta_0)_{max} = \Delta\theta_{max} = -\frac{g_{ina}(\gamma_0;\theta_0)}{\frac{dg_{ina}}{d\theta}(\gamma_0;\theta_0)} \tag{3.90}$$

3.7.8 Global sensitivity analysis

Global methods study the effects of varying multiple parameters simultaneously and relatively larger perturbations that cause the active set to change. In our CPT behavioral model, this could mean large errors in parameters or fundamentally misclassifying users, for example, assuming a given passenger to be loss averse ($\lambda > 1$) when in fact they are not ($\lambda < 1$). It is generally not possible to obtain explicit sensitivity derivatives at such points. However, an iterative scheme can be used to calculate directional derivatives for the optimal tariff and value function [58, 59]:

1. Calculate the initial optimum (γ^*, μ^a) and sensitivity differentials $\left(\frac{d\gamma^*}{d\theta}, \frac{d\mu^a}{d\theta}\right)$ at the nominal value $\theta_0 = \theta_0^1$.

2. Compute the local domain as in section 3.7.7 and the perturbed parameter θ_0^2 that disturbs the active set.

3. Calculate sensitivity differentials at θ_0^2 and update the active set to calculate 1^{st} order changes:

$$\Delta\gamma^* = \frac{d\gamma^*}{d\theta}(\theta_0^1)(\theta_0^2-\theta_0^1) + \frac{d\gamma^*}{d\theta}(\theta_0^2)(\theta-\theta_0^2) \tag{3.91}$$
$$\Delta\mu^a = \frac{d\mu^a}{d\theta}(\theta_0^1)(\theta_0^2-\theta_0^1) + \frac{d\mu^a}{d\theta}(\theta_0^2)(\theta-\theta_0^2) \tag{3.92}$$

4. Compute new optimal solutions and multipliers, as well as 1^{st} and 2^{nd} order approximations of f^*.

5. Repeat steps (1)-(4) whenever the active set is detected to change with incrementally larger perturbations.

Thus, while local sensitivity analysis constructs a linear approximation around the nominal operating point, the global sensitivity analysis creates a *piecewise* linear approximation with discontinuities at points where the active set changes.

3.7.9 Mismatch loss

Parameter estimation errors cause losses in the objective function, resulting from a mismatch between parameters assumed while designing the dynamic price ($\widetilde{\theta}$) versus the true but unknown, behavioral model of passengers (θ_{true}).

$$\Delta f = f(\gamma^*_{true}; \theta_{true}) - f(\widetilde{\gamma}^*; \theta_{true}) \tag{3.93}$$

$$\gamma^*_{true} = \underset{\gamma}{\text{argmax}} \;\; f(\gamma; \theta_{true}) \tag{3.94}$$

$$\widetilde{\gamma}^* = \gamma^*(\widetilde{\theta}) = \underset{\gamma}{\text{argmax}} \;\; f(\gamma; \widetilde{\theta}) \tag{3.95}$$

which implies that $f(\gamma^*_{true}; \theta_{true}) \geq f(\widetilde{\gamma}^*; \theta_{true})$.

3.7.10 Numerical simulations

Several assumptions need to be made to obtain reasonably accurate analytical solutions, for e.g, regarding the largest magnitude of allowed perturbations for a localized analysis to be valid, accounting for possible changes in the active set (section 3.7.7), checking curvature and convexity (eq. (3.82)) etc. These can be quite restrictive especially if we wish to consider larger perturbations and uncertainties in parameters. However, when the above mentioned assumptions do hold, analytical methods can be advantageous and much faster since most of the calculations can be done offline.

In addition, numerical approaches can also be used. The updated optimal solutions, value functions and mismatch losses can be computed using solvers in MATLAB's Global Optimization Toolbox, such as `Global Search` - which provides fast, proven quadratic convergence to local optima for such smooth problems using gradient-based methods. Results can be obtained by artificially constructing sensible travel scenarios (for both the SMoDS and the alternative) and repeatedly solving the NLP

in eq. (3.74)-(3.76) under both nominal and perturbed conditions. Numerical results can be applied more generally and provide a benchmark against which we can measure the accuracy of analytical approximations. On the other hand, such simulation-based methods are much more computationally expensive. This brute-force method is feasible here since the problem size and associated computational burden are relatively small. However, it may not be practical for larger, higher dimensional problems with more constraints. Furthermore, such solvers generally do not provide proven, theoretical guarantees for convergence to global optima. In our simulations, the termination criteria were tweaked to reach globally optimally solutions.

Chapter 4

Results and Discussion

4.1 Mode choice model results

4.1.1 Maximum simulated likelihood (MSL) estimation

The initial baseline results for the mode choice parameters, using both fixed and random coefficients are shown in table 4.1 and table 4.2, respectively. The likelihood ratio index was computed as explained in section 3.4.2, using the log-likelihood value at convergence and the log-likelihood when all the model parameters are set to zero which was estimated as $\approx -1.1541 \times 10^4$. The random parameters model presented here in table 4.2 were obtained using a normal mixing distributions for all the parameters, with $N = 500$ modified Latin hypercube draws.

Parameter	Estimate	Standard error (SE)
a_{walk}	-0.0220	0.0008
a_{wait}	0.0089	0.0067
$a_{ride,\ transit}$	-0.0027	0.0004
$a_{ride,\ exclusive}$	-0.0058	0.0007
$a_{ride,\ pooled}$	-0.0053	0.0004
b	-0.0060	0.0011
$ASC_{exclusive}$	-1.0960	0.0382
ASC_{pooled}	-1.2601	0.0423
Log-likelihood value at convergence	-1.0126 $\times 10^4$	
Likelihood ratio index ρ	0.1225	

Table 4.1: Results for mode choice utility functions estimated using *fixed* parameters for all the coefficients.

Parameter	Mean μ	SE	SD σ	SE
a_{walk}	-0.0586	0.0053	-0.1412	0.0079
a_{wait}	0.0113	0.0182	0.1491	0.0356
$a_{ride,\ transit}$	-0.0105	0.0013	0.0284	0.0017
$a_{ride,\ exclusive}$	-0.0086	0.0014	-0.0058	0.0010
$a_{ride,\ pooled}$	-0.0186	0.0013	-0.0095	0.0007
b	-0.0518	0.0050	0.0597	0.0042
$ASC_{exclusive}$	-2.5926	0.1800	2.3034	0.1558
ASC_{pooled}	-2.2230	0.1497	1.8175	0.1530
Log-likelihood value at convergence	-6.5350×10^3			
Likelihood ratio index ρ	0.4338			

Table 4.2: Mode choice utility functions estimated using random (normal) distributions for all coefficients.

Fixed vs. Random coefficients

Even from these initial results, we see that the random parameters model performs much better than the one with all fixed coefficients. This is evident since the random logit model has a significantly higher likelihood ratio index (ρ is close to 1) as well as the higher log-likelihood value at convergence. Thus, all of the subsequent analysis focuses on models where are all the parameters are allowed to belong to random distributions since fixed coefficients are inadequate to capture the variation in our sample well. It's important to note that all of the estimates presented are relative to setting the alternative specific constant of public transit as zero, i.e. $ASC_{transit} = 0$. Thus, it is the relative values of the terms and utilities that we should pay attention to, rather than their absolute magnitudes.

The tariff coefficient and all of the travel time coefficients except that for waiting time are negative, as expected since passengers suffer a disutility from losing either time or money. However, the coefficient on waiting time has a positive mean value. This non-intuitive sign of a_{wait} may be due to some of the simplifying assumptions we made, e.g. implicitly taking the disutility caused by waiting to be the same across different modes. Other possible reasons for this inconsistency could be to sampling or respondent biases, incorrect or inaccurate responses, other types of survey errors etc. Furthermore, the high standard deviation associated with this term implies that it still has non-zero probability density in the negative range, for a sizeable portion of respondents.

Another surprising observation is that the mean values of the constant terms (ASC) for both exclusive ride hailing and pooled ride sharing options are negative. This seems to imply that these modes offer less utility than transit, if we exclude price and time from consideration. This may be due to factors like the traffic congestion, increased carbon emissions, safety or privacy concerns etc. associated with ride sharing options, which may outweigh their other benefits like convenience and comfort. However, the standard deviation for these terms is also very large, indicating that there is a great deal of heterogeneity in how passengers perceive these external

factors apart from travel time and price.

4.1.2 Value of time (VOT) insights

Trip leg or mode	**VOT** (in \$/h)
Walking	67.8702
Waiting	13.1480
Transit ride	12.1703
Exclusive ride hailing	9.9466
Pooled ride sharing	21.5549

Table 4.3: Value of time spent on different modes, obtained from the random parameters logit model.

The value of time spent on different transport modes and trip legs can be calculated from the mean values of the estimated utility function coefficients from table 4.2, using the method described in section 3.1.3. We find that the orders of magnitude of the VOT values are comparable to the average hourly wage of ≈ \$32.40 in the Boston-Cambridge-Nashua metro area[1]. Furthermore, the relative trends also generally make intuitive sense. A higher value of time indicates that the passenger is willing to pay more to reduce the time spent on that leg. Thus, a higher VOT implies that the mode causes more disutility to them. Thus, it makes sense that respondents find walking to be more cumbersome than waiting, which in turn causes greater inconvenience compared to riding or in-vehicle travel time. Exclusive door-to-door ride hailing is more convenient for users than both pooled ridesharing and public transit. The only trend that's a bit surprising is that transit has a lower VOT than ride pooling. This could perhaps be due to the extra uncertainty in travel times and risk of delayed arrival associated with pooling.

[1] https://www.bls.gov/regions/new-england/news-release/occupationalemploymentandwages_boston.htm

4.1.3 Statistical significance of mode choice model

Parameter mean	t-stat	p-value	Z-score	p-value
a_{walk}	-11.0141	0.0000	12.8238	2.000
a_{wait}	0.6244	0.7338	2.3518	1.9813
$a_{ride,\, transit}$	-7.9241	0.0000	-11.4393	0.0000
$a_{ride,\, exclusive}$	-6.1996	0.0000	45.7569	2.0000
$a_{ride,\, pooled}$	-14.2325	0.0000	60.7051	2.0000
b	-10.4452	0.0000	-26.8046	0.0000
$ASC_{exclusive}$	-14.4046	0.0000	-34.7825	0.0000
ASC_{pooled}	-14.8511	0.0000	-37.7979	0.0000

Table 4.4: Results from conducting a hypothesis test on the mean values using the Student's t-distribution, for the statistical significance of the coefficients estimated by MSL.

We find that all of the estimates are statistically significant at the 5% significance level but at the 1% level, the null hypothesis cannot be rejected for some of the parameters. The p-values for the Z and t-test statistics are also very similar, which makes sense given that our sample size is relatively large.

4.1.4 Effects of MSL hyperparameters

Number of draws or simulation runs

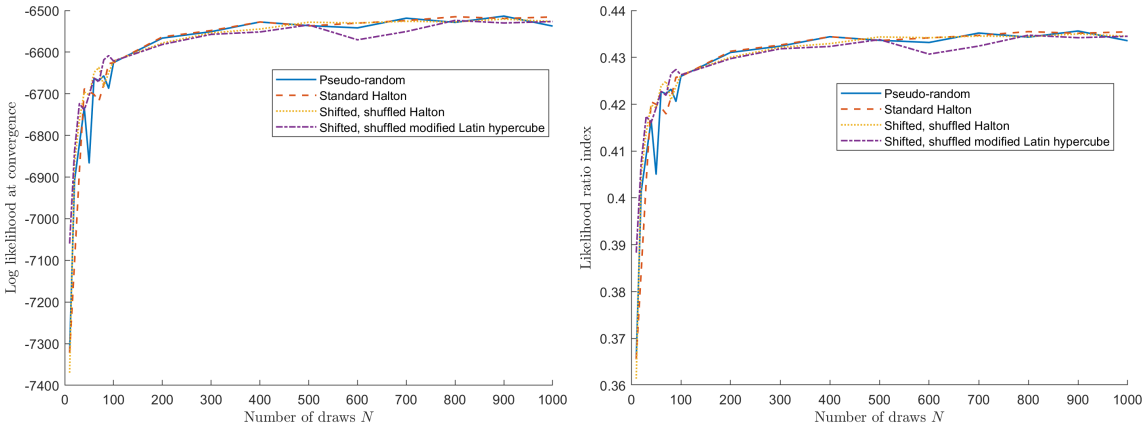

Figure 4-1: Variation of converged log-likelihood and likelihood ratio index with the number of random draws used.

We find that the value of the log likelihood function at convergence and consequently the likelihood ratio index, both increase with a larger number of simulated draws as seen in fig. 4-1, thus resulting in a more accurate model which better explains or predicts the observed data. The increase is very steep in going from 10 to around 200 draws, after which the values tend to plateau and only rise marginally with further increases in draws. This indicates that the simulating more random draws offers diminishing returns after a certain point. Thus, in this case, $N = 200$ draws would provide an optimal balance of (i) producing a model with sufficiently high predictive power and (ii) lower runtimes and computational burden.

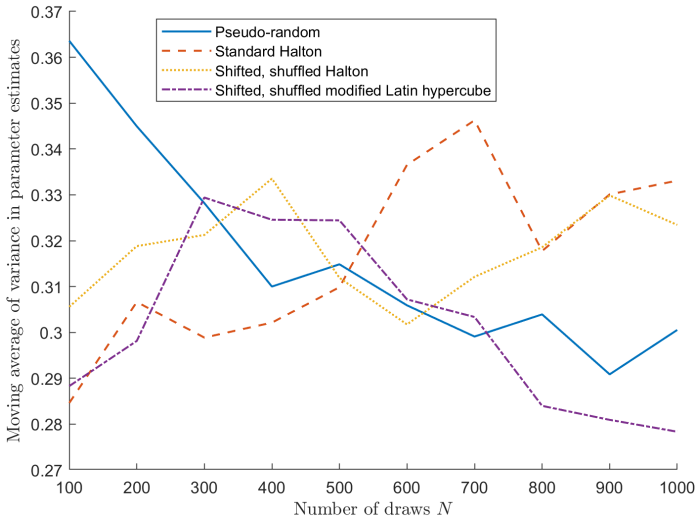

Figure 4-2: Average variance in the mode choice parameter versus number of simulations.

In general, we find that increasing the number of random draws reduces the simulation variance [54] as seen in fig. 4-2 where we performed a moving average on the data in order to show this trend more clearly. However, for certain draw types like standard Halton draws, the trend is not quite monotonic and the variance may even increase with the number of draws. This may be due to random variations from one run to the next, more definitive results could potentially be obtained by averaging over a large number of repeated trials. On the other hand, the average standard error of the estimated parameters increases slightly with more random draws, as seen infig. 4-3.

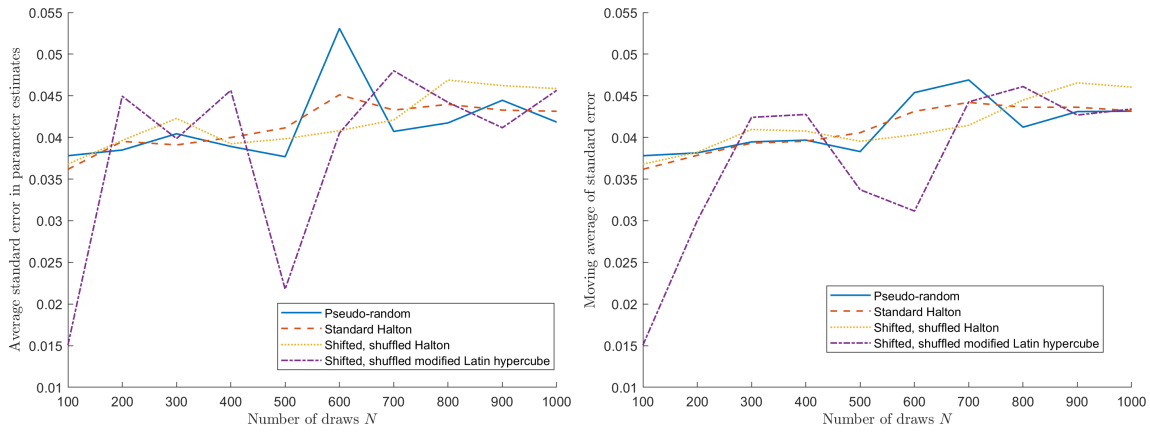

Figure 4-3: Variation in the average standard with the number of random draws used.

However, if the number of draws is too large relative to the sample size, we run into the issue of 'exploding parameters' which causes the standard errors to become very large and approach ∞, thus producing results that are no longer physically meaningful or accurate. This problem has been documented in past studies as well [60]. In our case, this is observed when the number of random draws N_{draws} gets larger than ≈ 1000, although the exact threshold also depends on other factors like the mixing distribution used, types of random draws taken and parameter initializations etc. On the other hand, if the number of draws is too small, we get very large standard errors and the Hessian of the log-likelihood may also become singular.

Type of draws used

I experimented with several different types of random draws such as:

1. Pseudo-random draws
2. Standard Halton draws
3. Shifted and shuffled Halton draws
4. Shifted and shuffled modified Latin hypercube sampling

It was found that the choice of draw only marginally affected the parameter estimates obtained. However, more sophisticated sampling techniques like (3) and (4)

do produce slightly better results (for the same number of draws) in terms of higher log-likelihoods at convergence, lower simulation variance and lower standard errors. This is because the sequences generated by such draws are closer to being truly random samples compared to standard draws. This implies that the simulation variance or error in the estimated parameters is lower when using such methods compared to pseudo-random draws, for the same number of runs. Thus, using Halton or Latin hypercube draws can significantly cut down the number of simulations needed and run times for MSL [54].

The difference is more apparent when using a small number of draws. For examples, in figures 4-1, 4-2 and 4-3, we see that for smaller values of N, Halton and Latin hypercube draws generally provide better results (i.e. higher likelihood ratio index ρ, lower variance σ^2 and lower standard error) than pseudo random draws using the same N. However, as N gets larger, the type of draw becomes less influential and all four sampling methods are comparable.

Type of assumed underlying distribution for parameters

There are several different possible choices for the mixing distributions of random parameters, such as:

- **Normal or Gaussian**: The parameter $\beta \sim \mathcal{N}(\mu, \sigma^2)$ with mean μ and standard deviation σ being estimated by MSL.
- **Log normal**: The model coefficients are e^{β} where $\beta \sim \mathcal{N}(\mu, \sigma^2)$, with μ and σ being estimated.
- **Truncated normal**: A normal distribution where the probability share below zero is massed at zero, i.e. the parameters are given by $max(0, \beta)$ where $\beta \sim \mathcal{N}(\mu, \sigma^2)$, with μ and σ being estimated.
- **Bounded Johnson's S_B distribution**: The parameters are given by $\frac{e^{\beta}}{1+e^{\beta}}$ where $\beta \sim \mathcal{N}(\mu, \sigma^2)$, with μ and σ being estimated. Note that this always produces values between 0 and 1.

- **Triangular**: Parameters are distributed according to $\beta \sim \mu + t\sigma$ where t is a piecewise linear, triangular function between -1 and 1 and the mean μ and spread σ are estimated.

From intuition, we always expect the coefficients on the price and travel time terms to be negative since they cause disutility to passengers. Thus, the truncated normal, log normal or S_B distributions should theoretically be good choice for these disutility coefficients since these always output parameter estimates of the same sign. However, the objective loss function (i.e. the negative log-likelihood) sometimes becomes unbounded or undefined (i.e. $\rightarrow \pm\infty$) at certain points for these mixing distributions, depending on the other hyperparameters used for the MSL estimation. This caused practical issues while minimizing the objective and frequently resulted in errors during the estimation process. Moreover, these distributions also sometimes resulted in spurious results that didn't agree well with intuition. For example, table 4.5 shows results obtained using a truncated normal mixing distribution for the travel time and tariff coefficients and a normal distribution over the alternative specific constants, using $N = 500$ modified Latin hypercube draws. As we can see, the coefficient on price is much smaller in magnitude than the travel time coefficients. This results in grossly overestimated VOT estimates across all modes, around 5-10 times higher than the average hourly wage. This leads us to have lower confidence in these estimates in spite of the likelihood ratio being quite high.

Thus, I avoided these distribution types and restricted myself to mainly using normal distributions during the estimation process and also while experimenting with different combinations of other hyperparameters. A normal distribution was also assumed for the constant bias terms (i.e., the ASCs) since the exclusive ride hailing and ride pooling modes could offer either positive or negative utilities relative to public transit.

Parameter	**Mean** μ_β	**SE**	**SD** σ_β	**SE**
Truncated normal $\sim -max(0, \beta \sim \mathcal{N}(\mu, \sigma^2))$				
a_{walk}	-0.3017	0.0354	0.5757	0.0618
a_{wait}	0.7712	0.0000	0.2547	0.0000
$a_{ride,\ transit}$	-0.4802	0.0739	1.0103	0.1440
$a_{ride,\ exclusive}$	-0.0552	0.0119	0.3639	0.0285
$a_{ride,\ pooled}$	0.0540	0.0057	0.2621	0.0395
b	-0.0100	0.0086	1.2641	0.1585
Normal $\beta \sim \mathcal{N}(\mu, \sigma^2)$				
$ASC_{exclusive}$	-0.0584	0.0831	0.1451	0.0000
ASC_{pooled}	-0.3485	0.0817	0.2181	0.0000
Log-likelihood value at convergence	-6.5808×10^3			
Likelihood ratio index ρ	0.4298			

Table 4.5: Mode choice utility functions estimated using truncated normal distributions for the price and travel time coefficients.

All these results seem to indicate that normal mixing distributions are most appropriate for all of the parameters in our mode choice model. This also agrees with the theory behind mixed logit and maximum simulated likelihood since MSL estimators are approximately asymptotically normal for large enough sample sizes [10].

Starting parameter values for the simulation

During MSL estimation, we have the freedom to set the initial values for the simulations or random draws. These correspond to the starting means and standard deviations for random coefficients, and starting magnitudes for fixed coefficients. I experimented with several different initializations but found that the model parameters estimated through MSL are largely invariant with respect to these, which agrees with intuition.

However, the estimation process overall is quite sensitive to the combinations of the different parameters described above, that are used. Under certain combinations, the objective function became unbounded or failed to converge, and it often wasn't clear which hyperparameter caused the issue. This made the tuning of these multiple hyperparameters more challenging.

4.2 CPT model estimation results

4.2.1 Validity of responses

During the pre-processing step as described in section 3.5.1, the responses that violated first order stochastic dominance were screened out. It was found that $\approx 66.53\%$ of the total responses (i.e. $n \approx 664$ respondents) were valid. All of the subsequent analysis and results presented below used only this valid subset of the data. Fortunately, the current survey still provides a large enough sample of respondents to obtain meaningful estimates and observe some heterogeneity in the results. However, future iterations of the survey could enforce more quality checks early on in the survey to prevent invalid respondents from proceeding further. For example, it may be possible to incorporate tests for validity similar to those described in section 3.5.2 within the survey itself to filter out such respondents in real time, albeit at the cost of increased complexities and effort needed in coding up the survey logic and design.

4.2.2 Detection of CPT effects

From table 4.6, we see that the valid responses clearly display CPT effects. The reflection or framing effect is shown by nearly all the valid respondents, indicating that our proposed value function is likely an accurate descriptor of how the passengers perceive their gains and losses. The probability weighting effect is not as dominant but it is still quite significant. We find that majority of them ($> 72\%$) show at least some overweighting of probabilities and it is also most common in the lower probability ranges (between $10 - 60\%$). This agrees with CPT theory since ti postulates that

people tend to overestimate the likelihood of rare events. The relatively large value of the mean gain/loss ratio (> 1) in the mixed lotteries indicates a significant degree of loss aversion among the surveyed passengers. However, the median value is quite close to 1 indicating that loss aversion may not be as prevalent for a sizeable portion of the passengers sampled in this study.

CPT effect tested	**% of valid responses**
Reflection effect	95.03
Probability overweighting between	
10% and 60% probability	62.56 %
60% and 90% probability	40.51 %
10% and 90% probability	51.05 %
Any probability weighting	72.44 %
Mean gain/loss ratio for mixed outcome lotteries	3.7254
Median gain/loss ratio for mixed outcome lotteries	1.0250

Table 4.6: Summary of key CPT effects observed.

4.2.3 Initial results for CPT model parameters

We first attempted to estimate the originally proposed CPT model having four parameters and the weighting function as given by eq. (3.17). We experimented with both the Levenberg-Marquardt and Trust-region-reflective algorithms, to approximately solve the nonlinear least squares problem in MATLAB. Both of these methods gave comparable results and are interchangeable for the purposes of this study.

Financial risk CPT parameters

As can be seen from the scatter plots in fig. 4-4 and the histograms in fig. 4-6, the financial risk parameters estimated from the lottery questions have a very poor distribution. For a large number of parameters, all four CPT parameters tend to go

towards either their upper or lower bounds. This is especially noticeable in the case of the loss aversion parameter λ which stays at the lower bound of 1 for almost all the respondents, which directly contradicts our observations in table 4.6 that indicate our respondents do display significant loss aversion. Although the results do look slightly better for the other three parameters, there is clearly a lot of room for improvement. This is also evident in fig. 4-5, which shows relatively large estimation errors for a large portion of respondents.

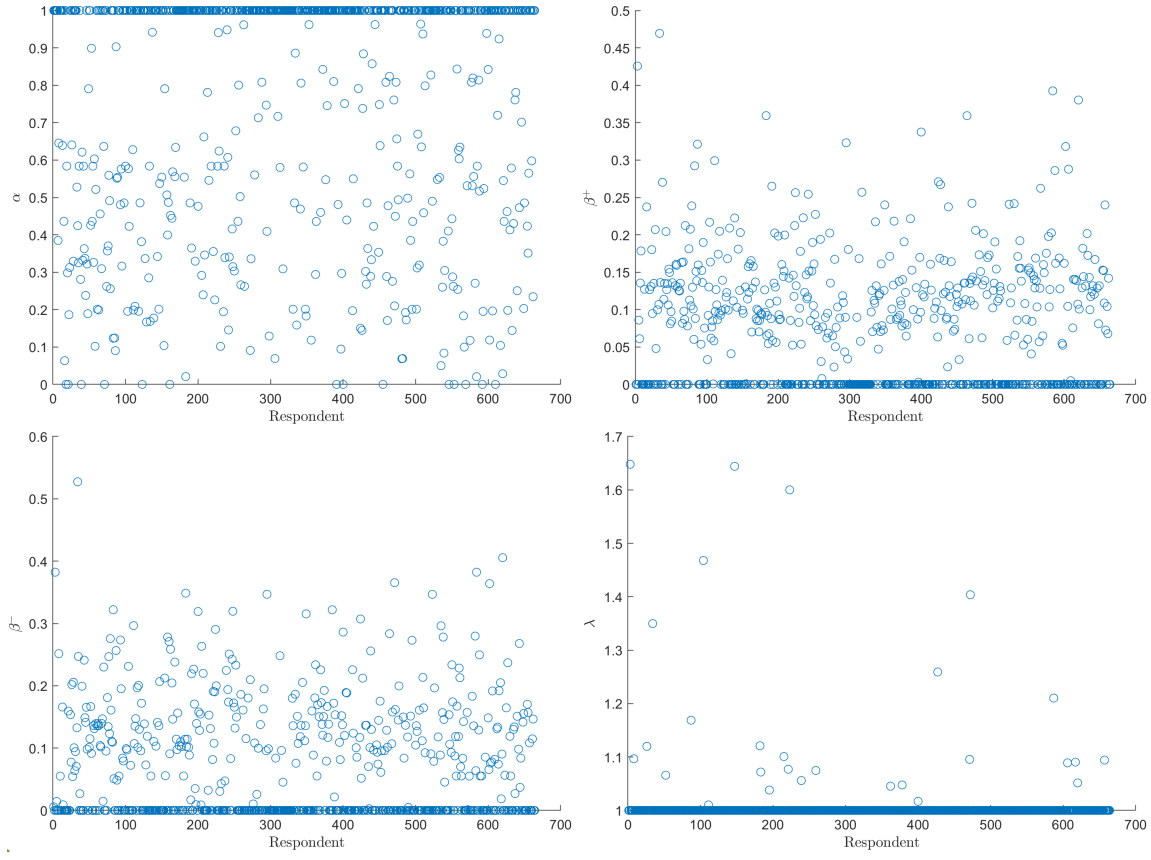

Figure 4-4: Scatter plots showing the poor distribution of financial CPT parameters across all the valid respondents.

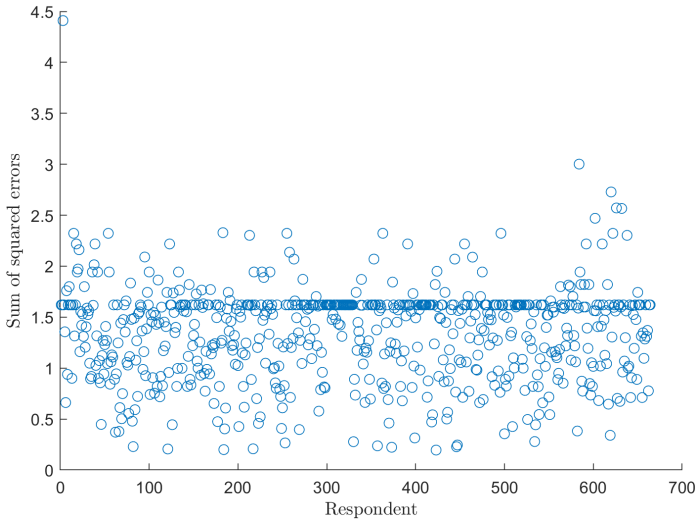

Figure 4-5: Scatter plot showing the squared norm of the CPT-lottery estimation errors for each of the valid respondents in the sample

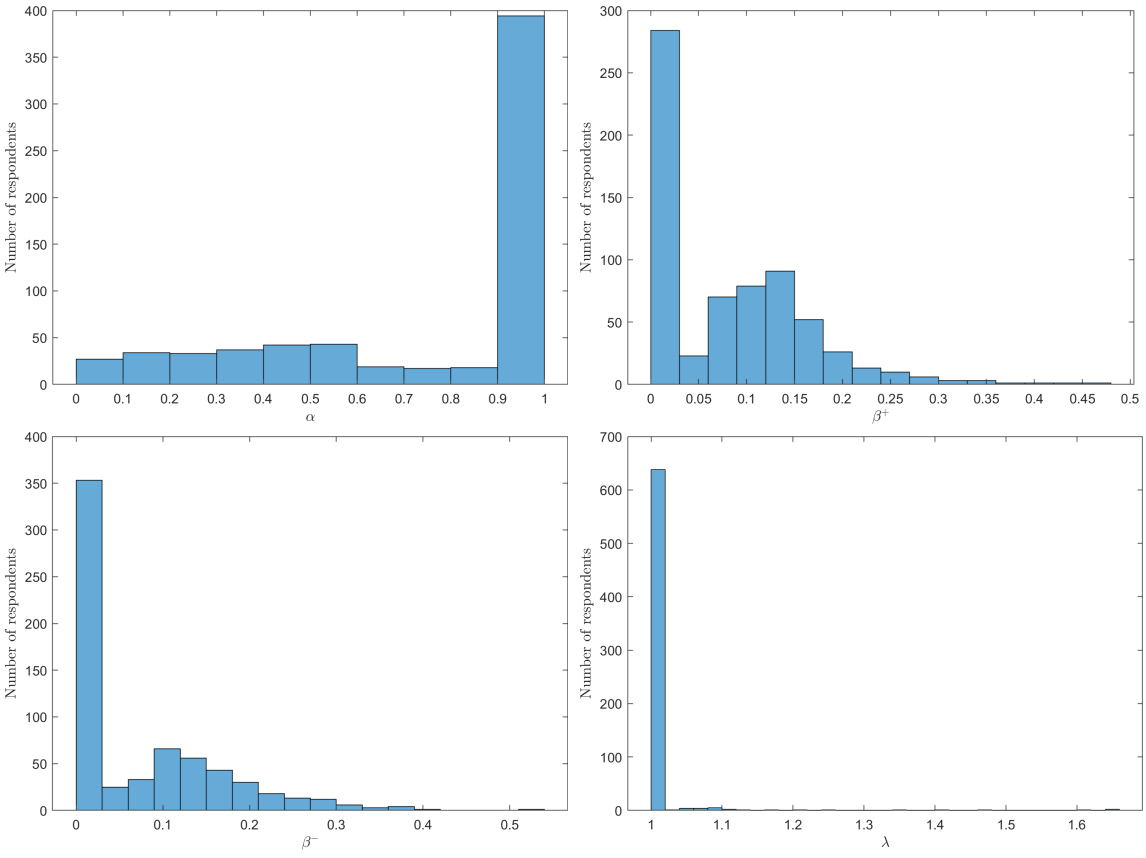

Figure 4-6: Histograms showing the skewed distribution of financial CPT parameters across all the valid respondents.

Travel risk CPT parameters for SMoDS

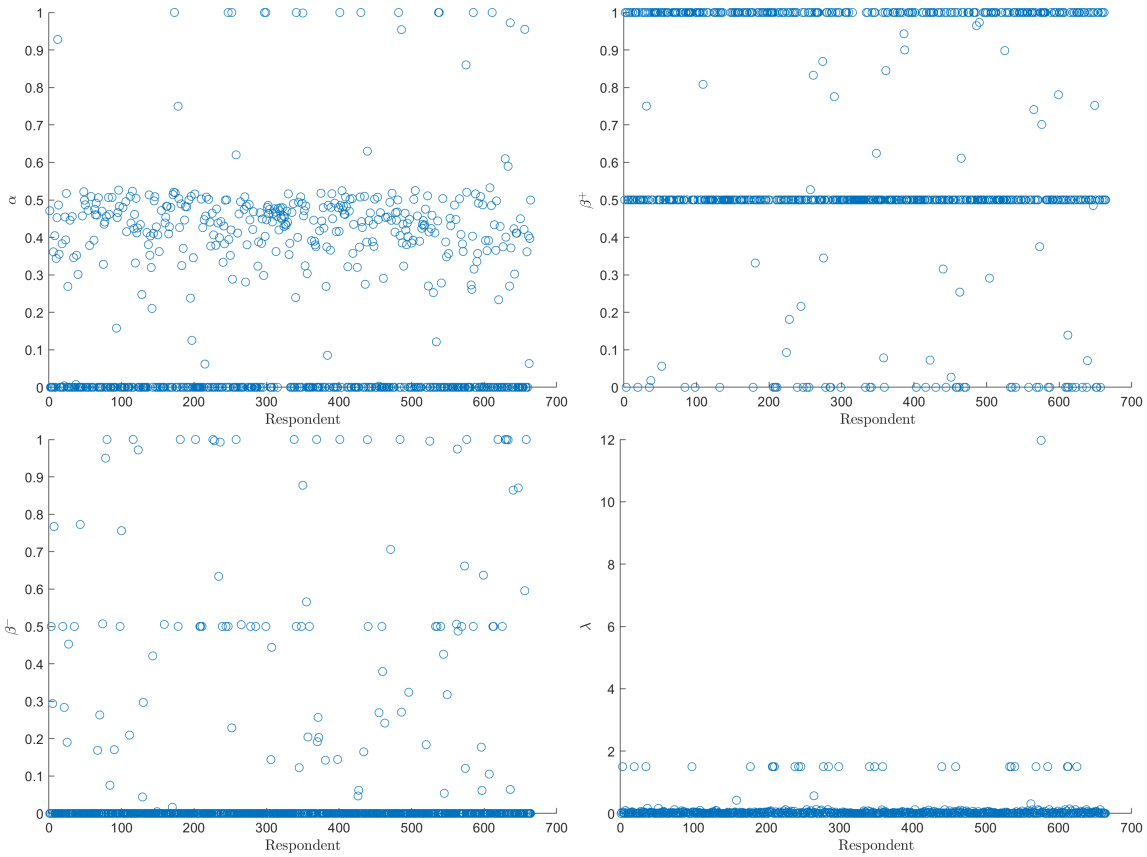

Figure 4-7: Scatter plots showing the poor distribution of travel CPT parameters across all the valid respondents.

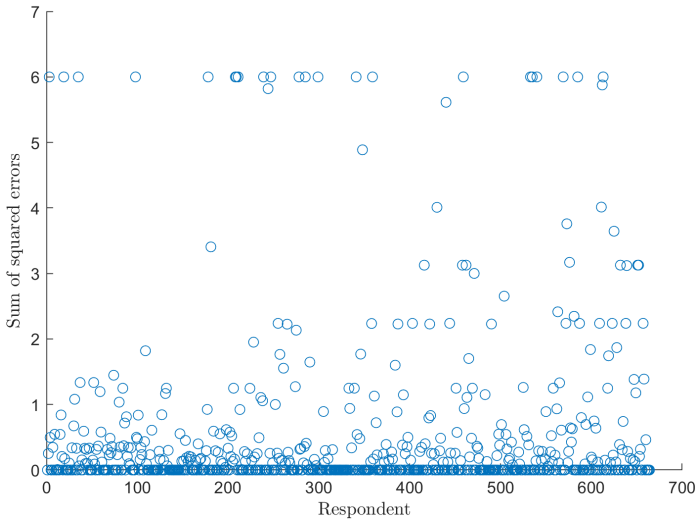

Figure 4-8: Scatter plot showing the squared norm of the CPT-SMoDS estimation errors for each of the valid respondents in the sample.

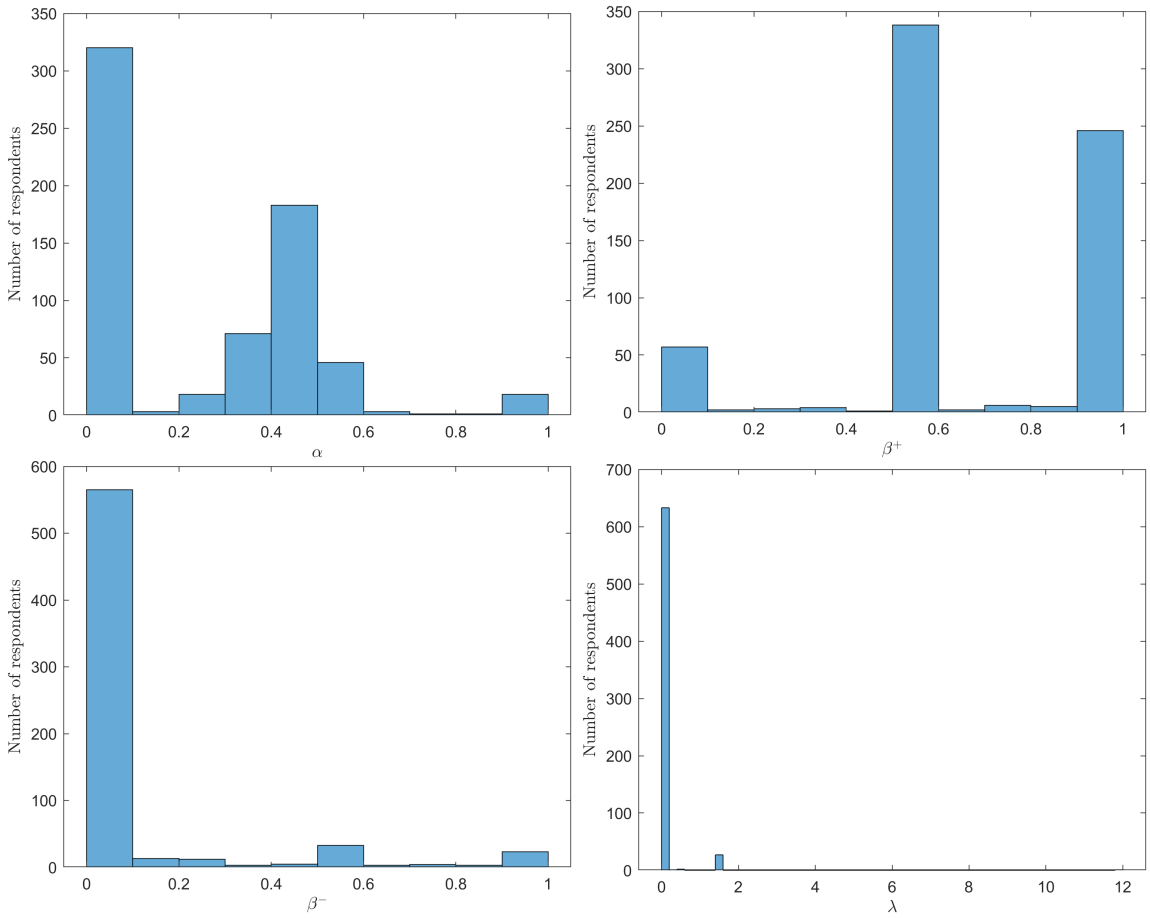

Figure 4-9: Histograms showing the skewed distribution of travel CPT parameters across all the valid respondents.

Similarly, the distribution of CPT parameter estimates for the SMoDS scenarios were also quite skewed and not very meaningful since most of the values either stay close to their initializations or go towards the bounds. This inferior distribution of both the travel and financial CPT risk parameters can also be inferred from table 4.7.

Type	**Lotteries**				**SMoDS**			
CPT parameter	α	β^+	β^-	λ	α	β^+	β^-	λ
Mean	0.7584	0.0774	0.0715	1.0077	0.2392	0.6448	0.0841	1.0989
Median	1.000	0.0745	10^{-6}	1.0000	0.2696	0.5000	10^{-7}	1.0000
Standard deviation	0.3257	0.0827	0.0917	0.0537	0.2533	0.3077	0.2290	0.5477

Table 4.7: Summary statistics for initial CPT parameter estimates.

4.2.4 Troubleshooting efforts

There may be several reasons behind the poor results obtained initially for the CPT estimates. Thus, I ran numerous computational experiments to diagnose what these issues might be. This in turn guided my attempts to resolve some of these problems and obtain better results for the CPT parameters.

Solvers, numerical settings and hyperparameters

Firstly, the nonlinear least squares solver as well as most other constrained optimization algorithms in MATLAB are often prone to local minima. This may cause it to prematurely converge to solutions that aren't globally optimal. In fact, many of the global optimization toolbox solvers in MATLAB actually use multiple local solvers to try and reach global optima. In order to prevent the algorithm from getting stuck at local extrema, I experimented with different solution techniques and also tweaked numerous settings within these solvers away from their default values. These include various tolerances like the optimality and function tolerances used to determine convergence and stopping criteria. Since the objective function was sometimes found to be insensitive to small parameter changes, I increased the finite difference step size δ from 10^{-8} to 10^{-3}. Too small values of δ may incorrectly result in the estimated gradients being zero and cause convergence to local minima. Thus, I also increased the minimum change in variables needed to compute finite difference derivatives from 0 to 10^{-3}. In order to improve accuracy from 1^{st} to 2^{nd} order, the finite difference method was changed from forward to central differences. Finally, I significantly increased the upper limits on the allowed number of iterations and function evaluations, to prevent the solver from stopping prematurely. Tweaking all of these settings resulted in some *marginal* improvements in the quality of results, giving a slightly better parameter distribution over the respondents.

Modifying model assumptions

Another avenue we explored was altering some of our modelling assumptions and trying out other alternative parametric forms for the CPT model that have been proposed in the literature. In addition to Prelec's probability weighting function eq. (3.17) that we proposed in our original model [51], we examined another form for the weighting function that has been widely used in past studies [61, 62] including some applicationsin the transportation space [39, 40]:

$$\tilde{\pi}(p) = \frac{p^{\alpha}}{(p^{\alpha} + (1-p)^{\alpha})^{\frac{1}{\alpha}}} \tag{4.1}$$

We found that both the original and modified weighting functions produce similar results, although the above form in eq. (4.1) does prove to better suited for analytical tractability due to its simpler parametrization when compared to to eq. (3.17).

In order to better fit the model to the data, we also allowed for different probability distortion parameters in the loss and gain regimes, i.e. α^{+} and α^{-} rather than simply α. Intuitively, this approach is reasonable since once would anticipate people to weight the likelihood of outcomes differently depending on whether they're losses or gains. Thus, both our original and modified weighting functions now become:

$$\pi_{\pm}(p) = e^{-(-ln(p))^{\alpha_{\pm}}} \tag{4.2}$$

$$\tilde{\pi}_{\pm}(p) = \frac{p^{\alpha_{\pm}}}{(p^{\alpha_{\pm}} + (1-p)^{\alpha_{\pm}})^{\frac{1}{\alpha_{\pm}}}} \tag{4.3}$$

As expected, increasing the number of model parameters from four to five does allow for better fit of the nonlinear model to the survey data and results in lower estimation errors overall, due to the higher dimensionality.

Rescaling and normalization

Optimization routines can sometimes output sub-optimal or even incorrect results if the variables involved (i.e. the inputs to the objective function) are on very different scales, resulting in a poorly posed or ill-conditioned problem. For example, in our

case with the extended 5-parameter CPT model, four of the risk parameters always lie between zero and one ($0 < \alpha^+, \alpha^-, \beta^+, \beta^- < 1$) while the loss aversion parameter is always greater than 1 by definition ($\lambda > 1$). Thus, I rescaled λ to also lie in a similar range i.e. $\in (0.01, 1]$ and accordingly also made changes to the objective function computations.

Furthermore, taking inspiration from [2], I decided to normalize the estimation errors or residuals in order to comparable across different scenarios that may differ greatly in their outcomes. Without such normalization, it may not be possible to meaningfully compare two risk scenarios whose respective outcomes are very different in their scales or magnitudes. Before normalization, the error for both the SMoDS and lottery scenarios is defined as the difference between the certainty equivalent predicted by our CPT model and the actual certainty equivalent value elicited from the survey response. For the financial CPT parameters using the lottery questions, two different normalization factors were considered:

1. Normalize the error for each question as the proportion of the above mentioned difference and the maximum monetary outcome possible in that lottery. Thus the normalized error is given by:

$$e_n = \frac{CE_{true} - CE_{pred}}{max\{|u_1|, |u_2|\}} \tag{4.4}$$

$$= \frac{U_R^s - \widehat{U}_R^s}{max\{|u_1|, |u_2|\}} \tag{4.5}$$

   where u_1 and u_2 are the utilities (i.e. monetary gains or losses) corresponding to the possible outcomes of that lottery.

2. Normalize the error in the certainty equivalent by the true certainty equivalent derived from the survey. The normalized error here is:

$$e_n = \frac{CE_{true} - CE_{pred}}{CE_{true}} \tag{4.6}$$

$$= \frac{U_R^s - \widehat{U}_R^s}{U_R^s} \tag{4.7}$$

Both of these normalizations are reasonable approaches, although option 2 is more intuitive. However, option 1 has also been shown to produce good results in past studies. For the SMoDS travel scenarios, only one normalization procedure was explored. The error for each scenario was divided by the true certainty equivalent from the survey, i.e. the same as option 2 from above.

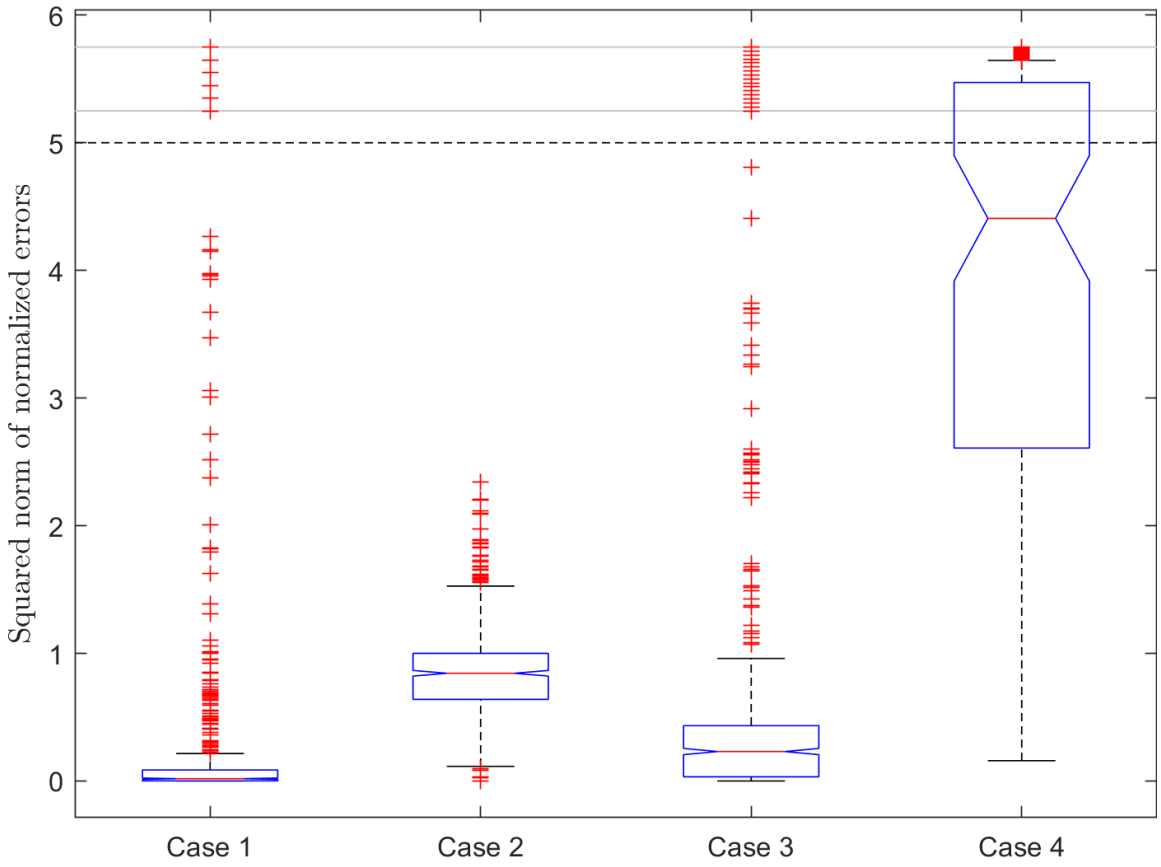

Figure 4-10: Box plots showing the distribution of estimation errors across respondents for the lottery questions, under four distinct approaches.

In fig. 4-10, we compare the estimation errors for determining financial CPT parameters from the lottery questions, under four different combinations of normalization factors and probability weighting functions. The four cases considered are:

1. Using the modified weighting function from eq. (4.3) and normalizing estimation errors by the maximum outcome of each lottery as in eq. (4.5).

2. Using the modified weighting function from eq. (4.3) and normalizing estimation

errors by the true certainty equivalent of each lottery as in eq. (4.7).

3. Using the original Prelec weighting function from eq. (4.3) and normalizing estimation errors by the maximum outcome of each lottery as in eq. (4.5).

4. Using the original Prelec weighting function from eq. (4.3) and normalizing estimation errors by the true certainty equivalent of each lottery as in eq. (4.7).

On each box, the central notched mark indicates the median value while the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme data points that are not considered outliers and the outliers are plotted individually using the '+' symbols. As we can see clearly from fig. 4-10, case 4 has a much higher median error than the other 3 cases. Case 1 has the lowest median error of all and the smallest spread but with a large number of outlier respondents for whom the error is very large. Similarly, case 3 has a lower median error than case 2 but also has many more extreme outliers. Thus, case 2 is likely the best approach to minimize errors relatively uniformly across all respondents in our sample.

Random draws for mode choice coefficients

Initially, for simplicity, the CPT estimation was performed assuming all coefficients in the mode choice model to be fixed. Thus, we just used the mean values for each parameter and did not allow variation among respondents. Later on, this assumption was relaxed and we used random parameter mode choice models to determine the CPT parameters. For each respondent in the sample, their mode choice parameters were randomly drawn from mixing distributions (in our case, normal) specified by the mean and variance estimated previously. Allowing for such heterogeneity among passengers is much more realistic and also produces better results.

Transfer learning

One of the main challenges with estimating SMoDS CPT parameters lies in determining the true reference point R for the travel scenarios. However, this is not an issue

for the lottery scenarios since $R = \$0$ (i.e. no loss or gain of money) is unambiguously the correct reference point, regardless of the respondent or outcomes being considered. This makes the CPT estimation process for the lotteries more straightforward. Thus, we tried to use the idea of 'transfer learning' from the lottery data, in order to leverage the relative simplicity of the lottery questions relative to the SMoDS mode choice scenarios. The process used was as follows:

1. Use the financial CPT parameters estimated from the lottery questions to first determine the reference for each of the travel scenarios, where R is taken to be the only unknown variable. Either a static or dynamic reference can be used.

2. Then taking these values of R as given, rerun the estimation process on the travel scenarios and determine the CPT parameters for the SMoDS, taking $\alpha^+, \alpha^-, \beta^+, \beta^-$ and λ as unknowns in this step.

Although this approach seems promising, implementing this procedure actually produced no visible improvement in the results for the SMoDS CPT parameters. Furthermore, it implicitly assumes that CPT risk attitudes and parameters are similar for both financial and travel-related outcomes, which may not always be the case.

Regularization

Finally, we considered regularizing the nonlinear least squares minimization problem (eq. (3.62)) in order to tackle the issue of parameter estimates tending to move towards either the specified upper or lower bound constraints. We implemented explicit regularization by adding L_2-norm penalty terms to the original loss function, similar to the Lagrangian while performing ridge regression or Tikhonov regularization for linear least squares. Thus, solutions that were too close to either limit were penalized equally:

$$\min_{\boldsymbol{\Theta}} \|\mathbf{e}_n(\boldsymbol{\Theta})\|_2^2 + \nu_1 \|\boldsymbol{\Theta} - \overline{\boldsymbol{\Theta}}\|_2^2 + \nu_2 \|\boldsymbol{\Theta} - \underline{\boldsymbol{\Theta}}\|_2^2 \tag{4.8}$$

$$\text{s.t. } 0 < \alpha^+, \alpha^-, \beta^+, \beta^- < 1,\ \lambda > 1 \tag{4.9}$$

where $\boldsymbol{\Theta} = [\alpha^+, \alpha^-, \beta^+, \beta^-, \lambda]^\intercal$, while $\overline{\boldsymbol{\Theta}}$ and $\underline{\boldsymbol{\Theta}}$ represent the enforced upper and lower limits. $\mathbf{e}_n(\boldsymbol{\Theta})$ represents the *normalized* error vector across all scenarios, for each respondent. Note that in the current implementation, we set the regularization hyperparameters on both terms as equal due to the symmetry in the problem statement i.e., $\nu_1 = \nu_2 = \nu$.

We found that adding such regularization to the objective does in fact improve the distribution of CPT parameter estimates for both the lotteries and SMoDS. In general, stronger regularization i.e. higher ν produce better distributions of CPT parameters. However, adding these penalty terms is a form of artificial bias. Thus, too much bias can result in models that are no longer as accurate in terms of their ability to predict the user's choices. This is similar to the bias-variance tradeoff encountered in supervised machine learning [63]. Thus, we need to carefully tune this hyperparameter through leave-one-out *cross-validation* to select the optimal ν. This can be done by repeatedly training the model parameters on a subset of the responses (for each respondent) and then evaluating the resulting model's predictions (of the certainty equivalent) using the remaining responses. The regularization parameter can then be determined using a uniform grid search to minimize this test error, while also taking into consideration the distribution of CPT parameters.

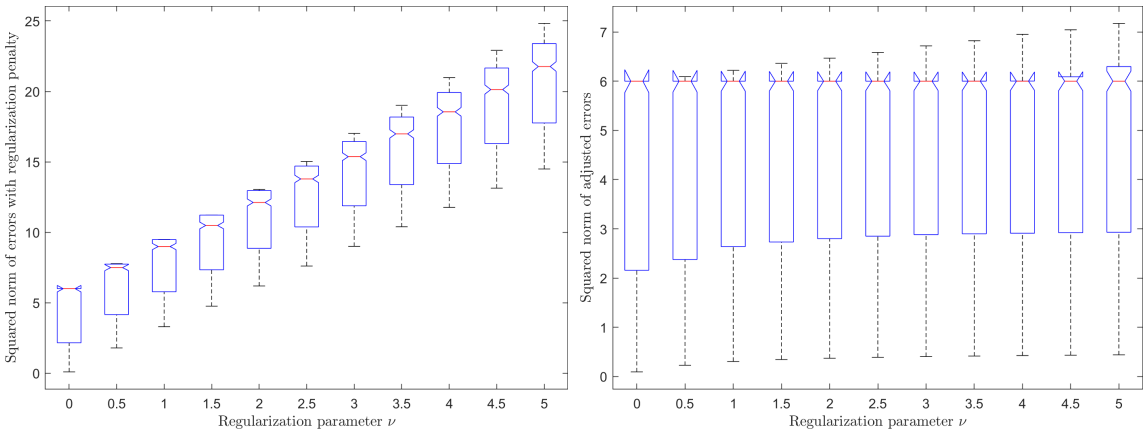

(a) Errors including regularization penalties and with $R =$ dynamic u_0.

(b) ν-*Adjusted* estimation errors and using $R =$ dynamic u_0.

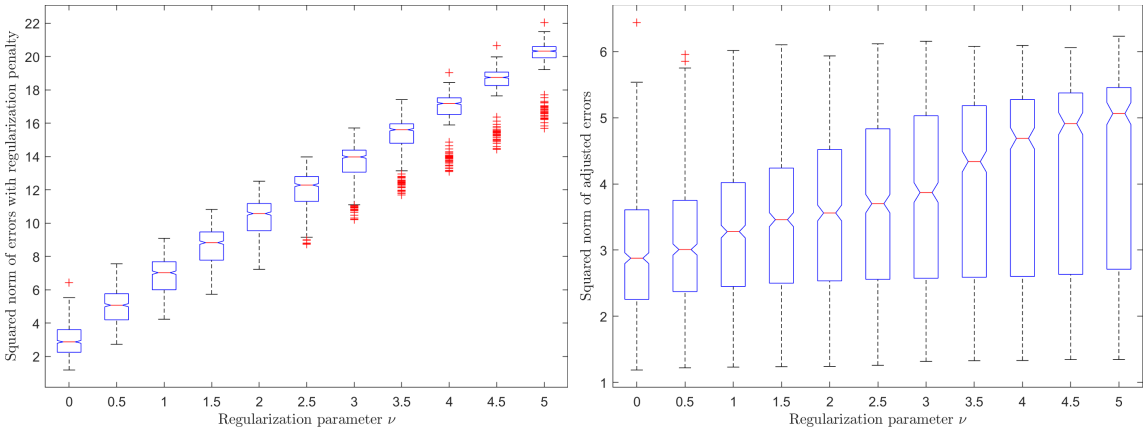

(c) Errors including regularization penalties and with $R =$ dynamic $SMoDS$.

(d) ν-*Adjusted* estimation errors and using $R =$ dynamic $SMoDS$.

Figure 4-11: Box plots showing the distribution of squared estimation error norms in the lottery CPT parameters across all valid respondents, for various regularization parameter values.

The plots in fig. 4-11 show both the raw and ν-*adjusted* estimation error distributions for various values of ν. The raw errors also include the contributions of the two squared L_2 norm penalty terms from eq. (4.9), while these were removed in the ν-*adjusted* values in order to isolate the errors arising from only the certainty equivalent prediction. As expected, we see that the raw error increases monotonically with increasing ν. However, the increases in the adjusted estimation errors are much smaller and more gradual. This implies that we can regularize the problem at least moderately in order to improve the distribution of parameter estimates, without significantly affecting the accuracy of the resulting model. It seems that a value of ν

between $\approx 0.5 - 2$ causes relatively small increases in the estimation error. Furthermore, these plots also show that $R =$ dynamic $SMoDS$ (expected utility of the two SMoDS outcomes) is a more accurate reference point since the median errors with $R =$ dynamic u_0 (objective utility of current baseline travel alternative) are much larger. The estimation errors when using *static* reference types were even larger and thus are not included here for conciseness. In contrast to the *dynamic* reference points which depend on the specific passenger as well as the current travel scenario under consideration, static references are objective utilities calculated using average-case prices and travel times per-mile. These vary for different passengers in the sample but do not depend on the specific trip attributes of the current travel scenario.

4.2.5 Improved results

Through the incremental improvements obtained by using the troubleshooting steps described above, the distributions of the financial CPT parameter estimates are greatly improved. In fact, we obtain satisfactory results for the lottery games, even without applying any regularization i.e. with $\nu = 0$. The scatter plots in fig. 4-12 and histograms in fig. 4-13 clearly show a much more heterogeneous and realistic distribution of parameters in the sample, while still keep the estimation error relatively low. These results were obtained using the modified weighting function, but estimates using the original weighting function were also decently good.

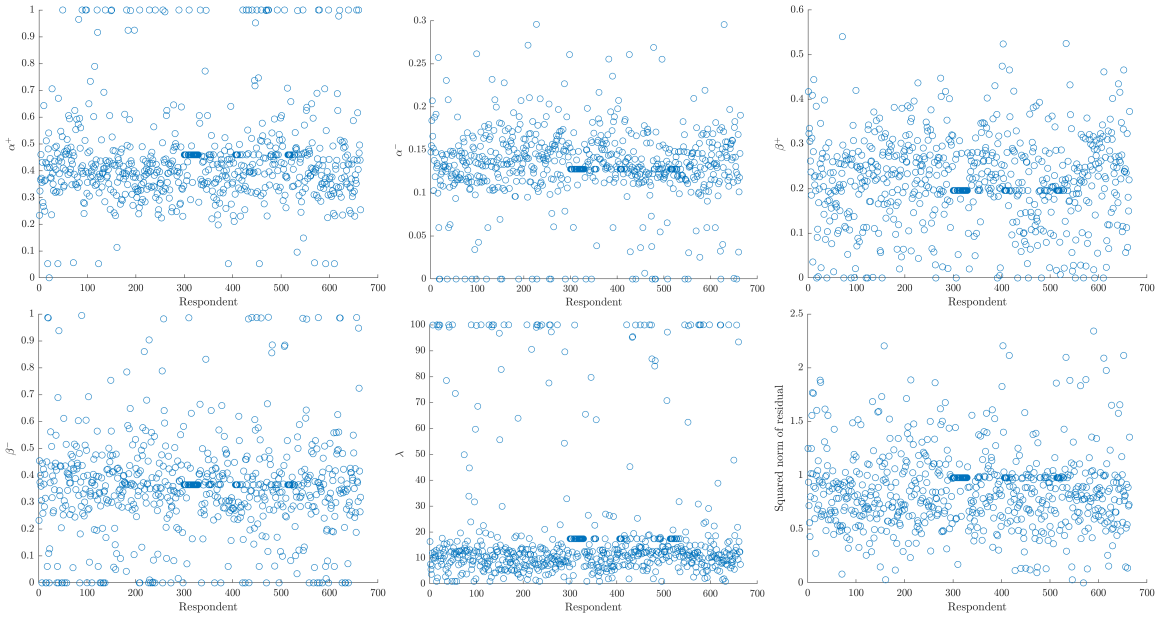

Figure 4-12: Scatter plots showing the distribution of financial CPT parameters and squared norm of estimation errors across all valid respondents, without any regularization.

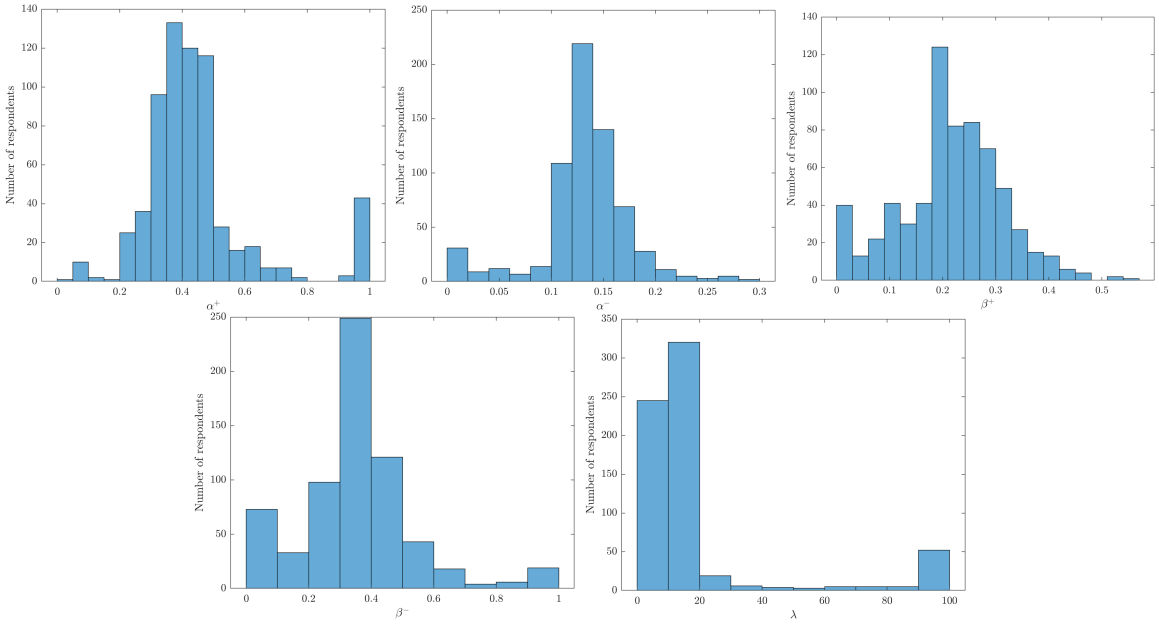

Figure 4-13: Histograms showing the distribution of financial CPT parameters across all the valid respondents, without any regularization.

The distributions of the SMoDS CPT parameters are also better now, but the improvement isn't quite as significant. This indicates that in spite of all our efforts above, there still remain some underlying issues with the travel risk scenarios and

the corresponding CPT model estimation for the SMoDS. We obtain decently good distributions for α^+ and β^-. However, even with a relatively high regularization penalty of $\nu = 2$, the estimates for α^-, β^+ and λ still noticeably tend to be skewed towards either the bounds or remain close to their initial values (0.5 for α^+, α^-, β^+, β^- and 50 for λ).

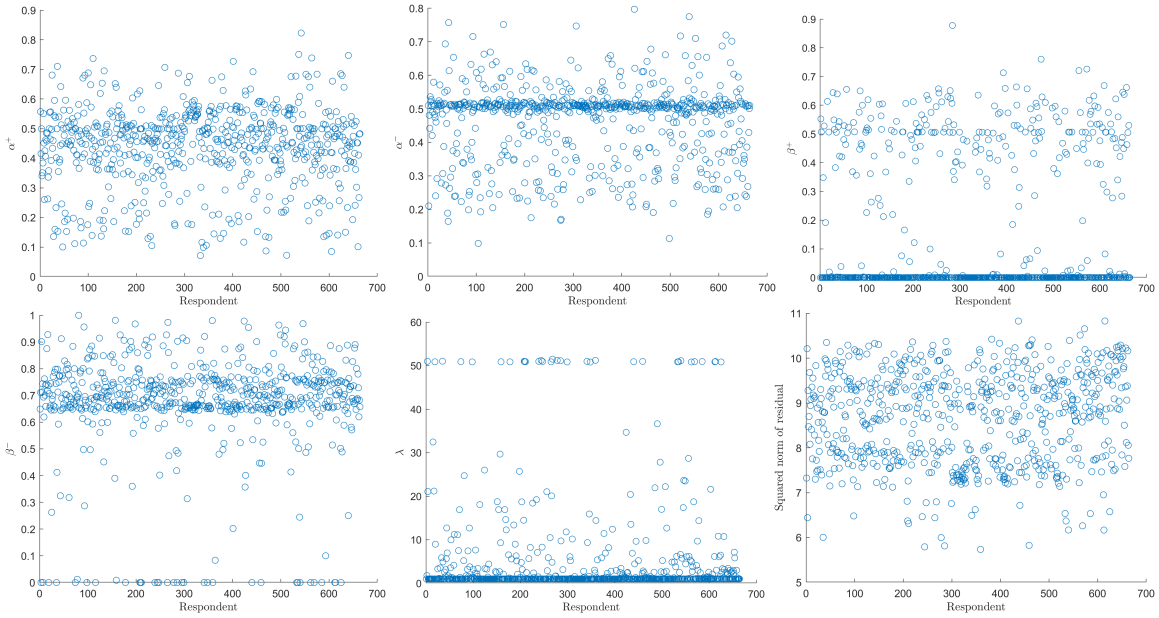

Figure 4-14: Scatter plots showing the distribution of travel CPT parameters and squared norm of (unadjusted) estimation errors across all valid respondents, using regularization with $\nu = 1.5$, $R =$ dynamic $SMoDS$ and the original weighting function.

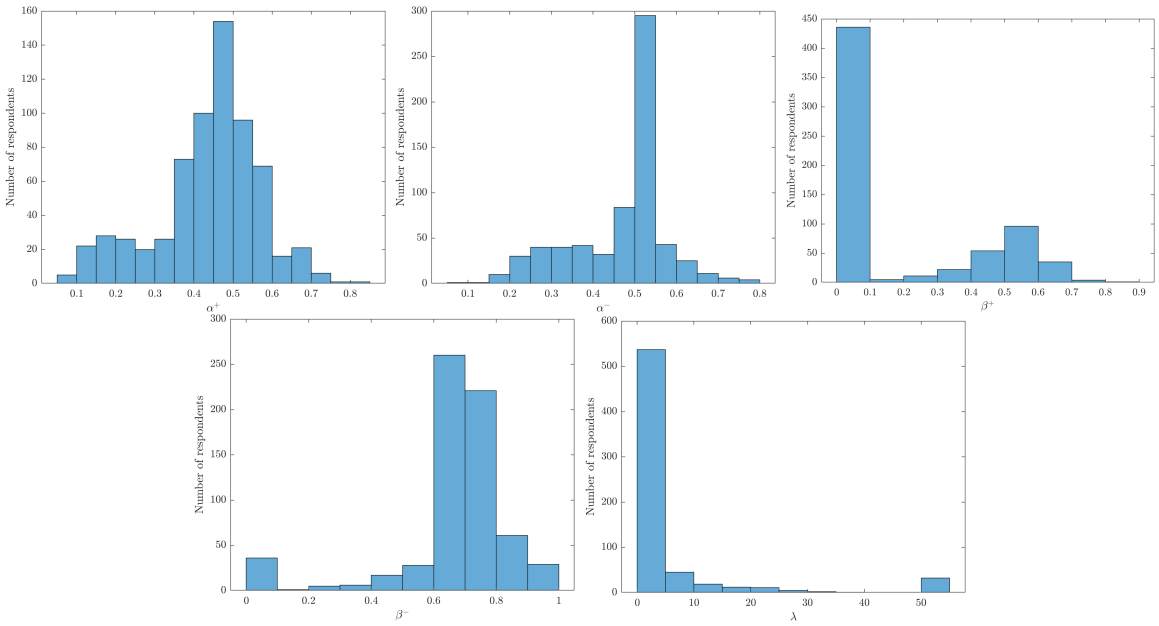

Figure 4-15: Histograms showing the distributions of SMoDS CPT parameters across all the valid respondents, without any regularization.

	α^+	α^-	β^+	β^-	λ	Squared error norm
Mean	0.4456	0.1315	0.2166	0.3550	20.0494	0.8625
Median	0.4124	0.1320	0.2188	0.3649	11.8715	0.8439
SD	0.1828	0.0448	0.0985	0.1906	25.8554	0.3605

Table 4.8: Summary statistics for final lottery CPT parameter estimates.

	α^+	α^-	β^+	β^-	λ	Squared error norm (adjusted)
Mean	0.4390	0.4653	0.1721	0.6664	5.3998	3.4915
Median	0.4607	0.5054	10^{-5}	0.6930	1.0000	3.4599
SD	0.1333	0.1130	0.2438	0.1922	11.4421	1.1016

Table 4.9: Summary statistics for final travel CPT parameter estimates.

4.2.6 Pending issues with SMoDS CPT parameters

Effect of reference point

One of the key challenges that arises with the estimation of travel CPT parameters is in the accurate determination of the reference utility R for each passenger and trip. This appears to be the dominant factor influencing both estimation residuals and the quality of parameter distributions. While the choice of reference is quite obvious and intuitive in the case of the lottery questions ($R = \$0$), this is not the case for the SMoDS travel scenarios. Here, the reference function is essentially unknown and likely not unique. The reference points could vary for each passenger in the population and also depending on the specific scenario and trip attributes under consideration. Our initial hypothesis was that we could reasonably approximate this quantity using a heuristics-based approach. In our experiments, we tried several different heuristics for both static and dynamic references, based on the objective utilities of either the SMoDS or the user's baseline travel alternative (i.e. the certain prospect). As we saw in fig. 4-11, the estimation errors can vary quite significantly depending on the value of R that's assumed. Thus, if our choice of R differs from the true value by a large margin, then the estimation errors remain large regardless of the values assumed by the CPT parameters. This would make the estimation process largely insensitive to the other parameters, causing the observed behavior where estimates either stay at their initializations or stick to the bounds of the interval.

In general, we found that dynamic references are more realistic and intuitive. These produce better parameter estimates and have lower residuals compared to static reference points. However, how each passenger dynamically computes their reference is still unclear. We obtained the lowest errors while using the expected utility of the SMoDS as our reference. However, even the residuals in this case are relatively large with a median value of ≈ 3 for $\nu = 0$ as seen in fig. 4-11. This is in contrast to the financial parameters where the estimation error norms are close to 1 as seen from fig. 4-12.

Thus, we also experimented with an alternative approach for calculating R. Rather

than assuming a heuristic, we also included R itself as a variable (along with the CPT parameters) in the nonlinear least squares optimization problem. Thus, we now have 6 unknowns that are estimated using 6 nonlinear equations, per respondent. We found that optimizing R separately for each passenger did help reduce the estimation errors to nearly 0 for all respondents. However, it *did not* improve the distributions of the CPT parameter estimates. In fact, the results actually got worse for a large portion of respondents. This is likely because optimizing R in this manner results in severe overfitting of the model to the training data. Thus, driving the estimation error to 0 in this manner does not necessarily lead to better, more meaningful parameter estimates. Another caveat with this approach is that it can only be used to model a static reference i.e. it implicitly assumes a single value of R for each passenger, that's shared across all 6 of the travel scenarios they face. Thus, for many respondents, the minimization routine may set R in such a way that the user perceives *all* of their travel scenarios as either pure gains (or losses). This would make their model insensitive to parameters in the other loss (or gain) regime. We cannot implement dynamic references (where R varies for both respondents and for each scenario) since this would result in 11 unknowns per respondent, which cannot be reliably estimated using only 6 equations

4.2.7 Data quality challenges

I ran some computational experiments to diagnose root causes for why the distribution of CPT travel parameters remains quite skewed. After making several adjustments as outlined in section 4.2.4, we are able to obtain good results for financial CPT parameter estimates. This confirms that our data processing, analysis and estimation procedures are technically sound. These findings point towards potential issues with the data used to estimate the SMoDS CPT parameters.

I tested two different hypotheses computationally using the responses to the lottery questions, to see if the issue was with either the *quantity* or *quality* of data collected, or both. Both of these tests used the lottery data and its associated financial CPT parameters as the 'ground truth':

1. **Quantity issue**: The inferior results for the travel parameters are due to insufficient amount of data (i.e. the number of equations) available for estimation. The SMoDS parameters are obtained using only 6 equations or SMoDS choice scenarios, whereas the lottery parameters use 10 equations. In order to test this hypothesis, I reran the CPT lottery estimation using fewer equations i.e. 6 equations instead of the 10 originally used, to see if this would produce worse results. I randomly selected a subset of the scenarios to be used for the non-linear least squares estimation process and these same 6 survey questions were used for all the respondents.

2. **Quality issue**: The inferior results are due to poor, sub-optimal design of the choice scenarios that caused incorrect or inaccurate responses, in particular as they relate to the reference point R. I tested this hypothesis by rerunning the CPT lottery estimation by optimizing R specifically for each respondent, rather than using $R = \$0$ as a heuristic.

From test 1, I found that what mattered was not necessarily the *number* of scenarios used, but rather the *types* of scenarios considered. The estimation errors swung considerably depending on which 6 scenarios were used. From the plots in fig. 4-16, we see that in general, the errors are still close to zero for nearly all the respondents even though only six equations were used. However, the exact distributions and range of error magnitudes obtained depend quite heavily on which subset of 6 lottery scenarios are chosen. Thus, we may deduce from this that it is the *quality* of data that matters more than the *quantity*, i.e. including more questions or choice games in the CPT-SMoDS risk section of the survey would not have necessarily led to better results.

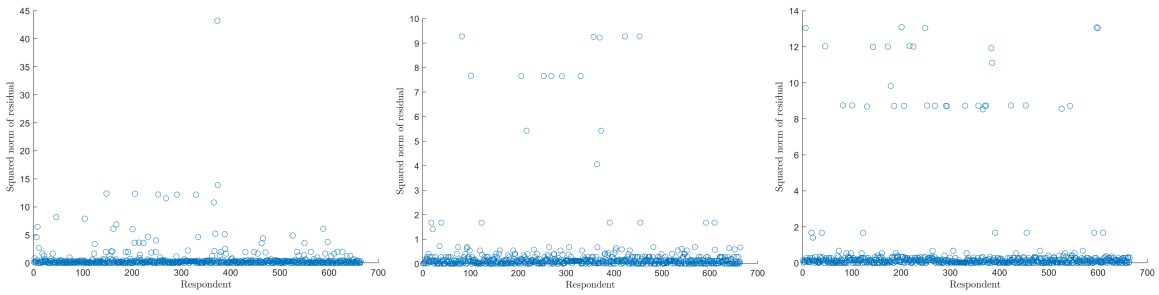

(a) Using lotteries $\mathcal{S} = \{1, 2, 5, 8, 9, 10\}$ (b) Using lotteries $\mathcal{S} = \{1, 4, 5, 6, 7, 9\}$ (c) Using lotteries $\mathcal{S} = \{1, 4, 5, 6, 8, 10\}$

Figure 4-16: Testing hypothesis 1: Effects of lottery scenarios selected for estimation, on the resulting errors in parameters.

From test 2, we found that empirically optimizing the reference point R separately for each individual respondent in the sample instead of directly setting $R = \$0$, still produces reasonably good results in terms of both low estimation errors and decent distributions of parameter estimates, as seen in fig. 4-17. Furthermore, it turns out that even when the reference is included as an optimization variable, the estimated R values are still quite close to \$ 0. This indicates that the lottery survey scenarios are still able to infer the respondents' true reference reasonably accurately and output meaningful parameter estimates, even without the valuable prior information that R should intuitively equal 0.

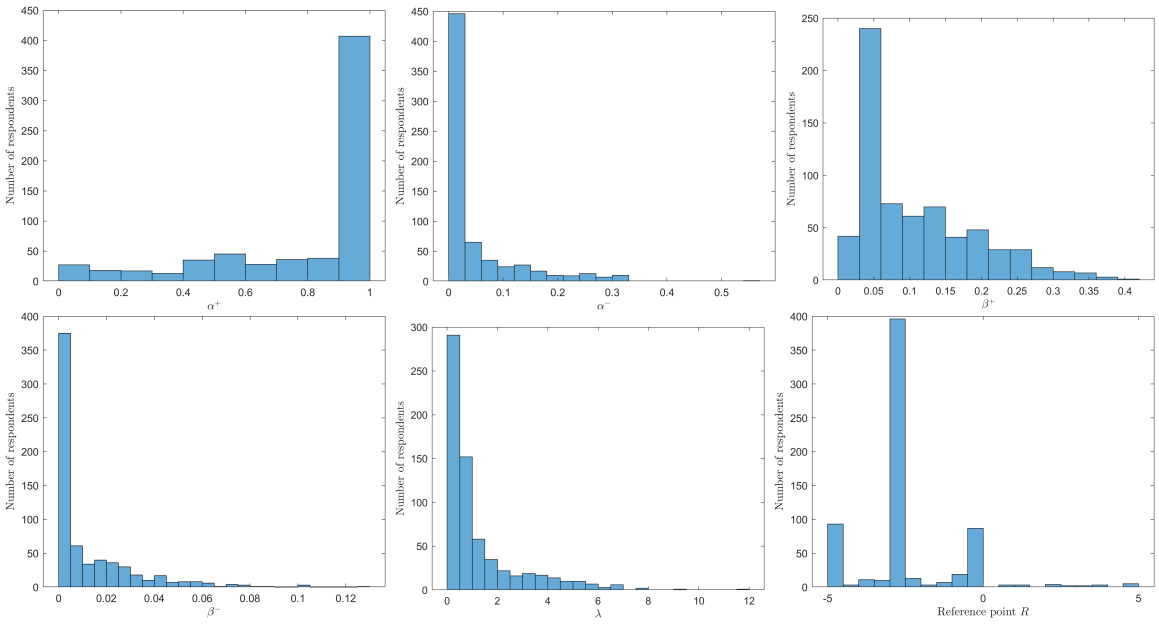

Figure 4-17: Histograms showing the distribution of financial CPT parameters and of the self-optimized reference points R for each passenger, when testing hypothesis 2.

Another challenge I faced in this study was the lack of any 'ground truth' data that I could compare my results and parameter estimates against, especially for the SMoDS choice scenarios. It would have been very helpful to have some actual empirical data for customer decision making for pooled ridesharing, that I could use to explicitly test the accuracy of my CPT model's predictions and compare their choice probabilities. However, ridesharing firms like Uber and Lyft are notoriously secretive about their pricing algorithms and it was thus impossible to get access to such data from the field. Thus, our only viable option for the estimation steps was to collect data ourselves via this survey, and generate synthetic data via simulations for the sensitivity-robustness analyses (section 3.7) and feedback control steps (chapter 5). Unlike with the mode choice coefficient estimates, we also can't check the statistical significance of these CPT parameter estimates via hypothesis testing due to the small sample size ($n = 6$) of survey responses used per respondent. Neither is it possible to get their standard errors via bootstrapping since each of these are individual specific and not population-wide parameters.

4.2.8 Guidelines for future survey design

Put together, our experimental results testing both the hypotheses indicate that there are potentially some underlying flaws in the design of the risky SMoDS choice scenarios, i.e. the survey section dealing with travel CPT parameters. As a result, these questions weren't able to elicit the passengers' risk preferences to a high degree of accuracy. There remain some technical barriers in designing such mode choice scenarios to aid prospect-theoretic modelling of travel risk attitudes. Constructing surveys involving questions like the SMoDS travel risk scenarios is much more involved and complex, compared to the much simpler lottery questions used for financial risk.

For instance, in the case of lotteries, it is inherently obvious whether a particular monetary outcome is a gain or loss. However, whether certain trip characteristics (like travel times, tariffs etc.) corresponds to gains or losses depends on a host of external factors such as the mode-specific utility functions, environmental considerations and the specific passenger's true reference point, which remains unknown. Accounting for all of these features and enumerating all the possibilities while designing such choice experiments proves to be very cumbersome. This requires the travel scenarios presented to be highly customized in real time for each user, while also being dependent on their responses to preceding questions. Nevertheless, implementing such surveys that more systematically cover all possible combinations of gains/losses/mixed outcomes, could produce better quality data which in turn would give better results for the travel CPT parameters.

Future iterations of the survey could include a few additional survey sections. One section would construct augmented discrete choice scenarios similar to our mode choice questions (section 3.3.3) but also incorporating uncertainty in the outcomes of the SMoDS (similar to section 3.3.4). This would allow us to directly compare the predicted and true choice probabilities, thus giving us a reliable metric for the CPT model's goodness of fit. This would increase our confidence in our model's predictive power, rather than relying only on the estimation errors. It would also be useful to create a section with new types of questions aimed solely at precisely determining

the reference for each passenger and trip, instead of having to resort to an arbitrary heuristic. However, even with highly sophisticated and expertly designed surveys, one must accept that there is always going to be some finite amount of human error that can't be avoided. This could be due to several factors like the respondent not understanding the question well enough, accidentally mixing up numbers while filling out their responses etc. [2].

4.2.9 Demographic heterogeneity in risk preferences

We were also able to leverage the anonymized demographic data that we collected as part of the survey, to examine how these correlated with the respondents' CPT parameters and general attitudes towards uncertainty and risk. Since the distributions were superior for the financial CPT parameter estimates, these results were generated using the lottery data. The box plots are helpful to concisely display the various summary statistics for each sub-group. Simultaneously, swarm scatter charts can provide deeper insights into demographic trends by visualizing the underlying parameter distributions for each sub-group within the population. Thus, both types of graphs have been provided here.

4.2.10 Variation of risk attitudes with gender

We observe that there are noticeable differences in risk preferences between males and females. Note that we can't reliably comment on the trends for non-binary respondents since a very small number of people identified as belonging to that group. Although the median values tend to be quite similar for both men and women across all five parameters in fig. 4-18, there are other discrepancies that are quite significant. I general, we see that men to exhibit more spread in their risk parameters as seen in fig. 4-19, with a larger inter quartile range and more extreme points. Women on average have higher values of α^-, β^+ and λ, indicating that they tend to over or under-weight losses more, are relatively more sensitive to incremental gains and are more risk averse than men. On the other hand, men tend to over or underestimate

gains more and are more sensitive to incremental losses below their point of reference.

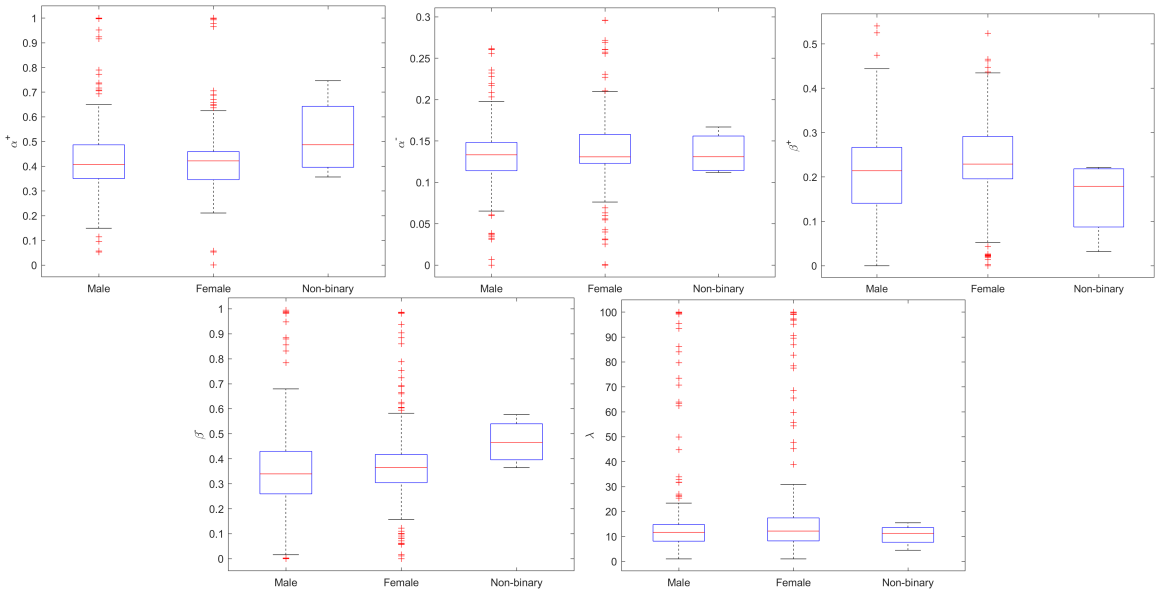

Figure 4-18: Box plots showing the variation in CPT parameters by gender.

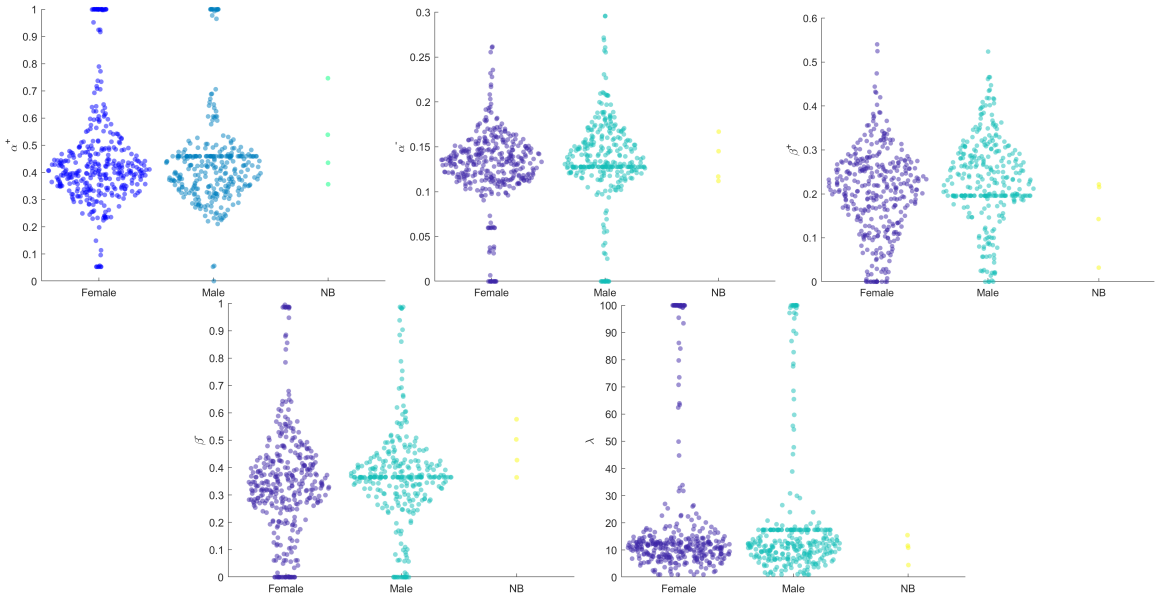

Figure 4-19: Swarm scatter charts showing the variation in CPT parameters by gender.

4.2.11 Variation of risk attitudes with income

Unlike with gender, there aren't as many conclusive or clear trends that we can deduce immediately, in risk preferences by gross annual income. One observation

from fig. 4-22 is that in general, medium-to-high income brackets tend to display more variation in their risk parameters across individuals when compared to lower income respondents. Looking at fig. 4-21, we notice that the all of the risk parameters seem to have a somewhat bimodal distribution with income, with their values reaching peaks or troughs either near the middle or towards the extreme ends of the income spectrum. Thus, there do not appear to be many decisive patterns that we can infer.

However, it is interesting to note that richer respondents are likely to be more loss averse. Another important insight we obtain from the swarm plots is that our data set seems to be skewed towards higher incomes, with a disproportionately large number of respondents belong to the $100-150,000 bracket in particular. We only enforced quotas for gender and age in our survey but it may be worthwhile to impose income quotas as well in future iterations to obtain a more representative sample.

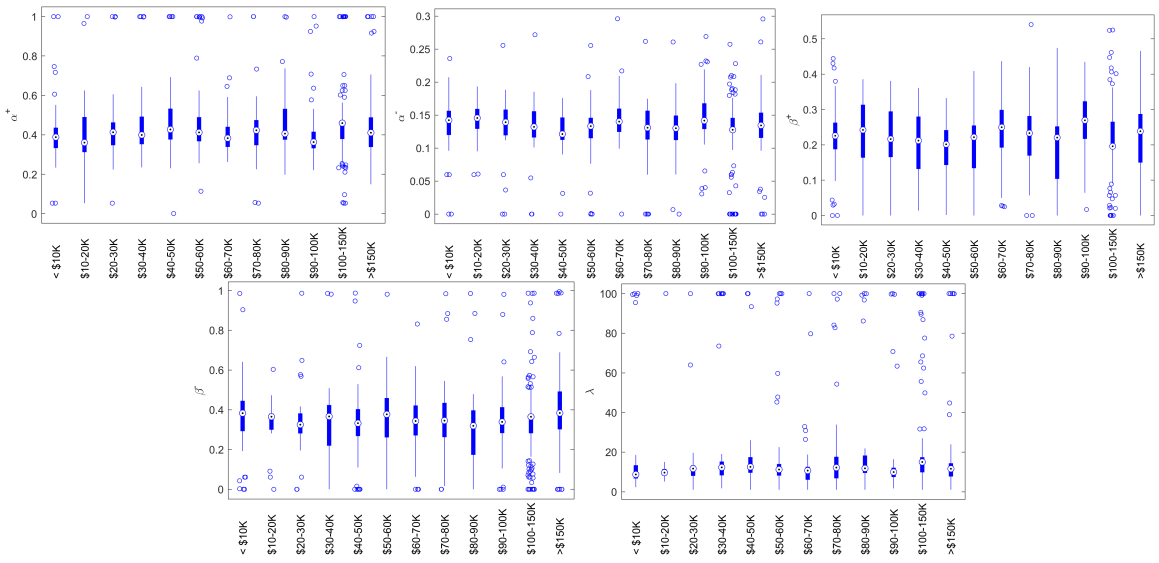

Figure 4-20: Variations in CPT parameters by income.

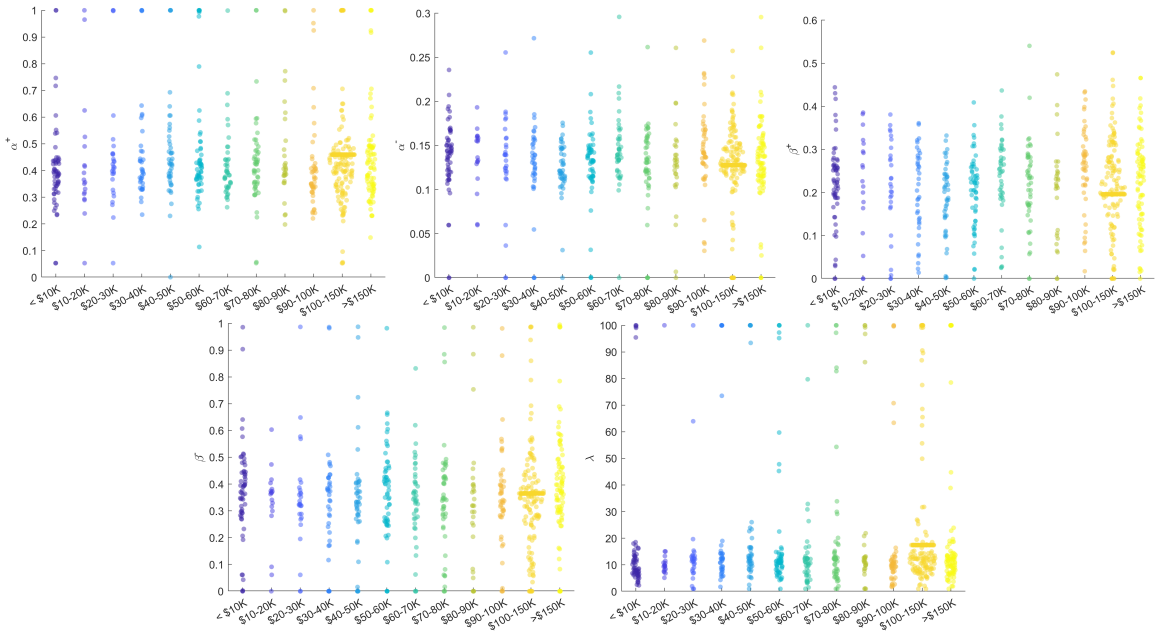

Figure 4-21: Swarm scatter charts showing the variation in CPT parameters by income.

4.2.12 Variation of risk attitudes with age

Similar to what we see with gender, there are some trends in CPT parameters with respect to age that do stand out in fig. 4-22. For example, older age groups show larger spread in their risk attitudes and more extreme variations among different respondents. Younger age groups are less loss averse (with lower λ) and are willing to tolerate higher levels of risk, which agrees with common beliefs as well. The trends in probability distortion parameters is a bit uneven across age groups, but older groups in general do seem to either over or under-weight the likelihood of gains and losses more (i.e. they tend to have higher values for both α^+ and α^-). Also, older age groups on average seem to have lower values for both β^+ and β^-, implying that they are less sensitive to either incremental gains or losses away from their reference baseline.

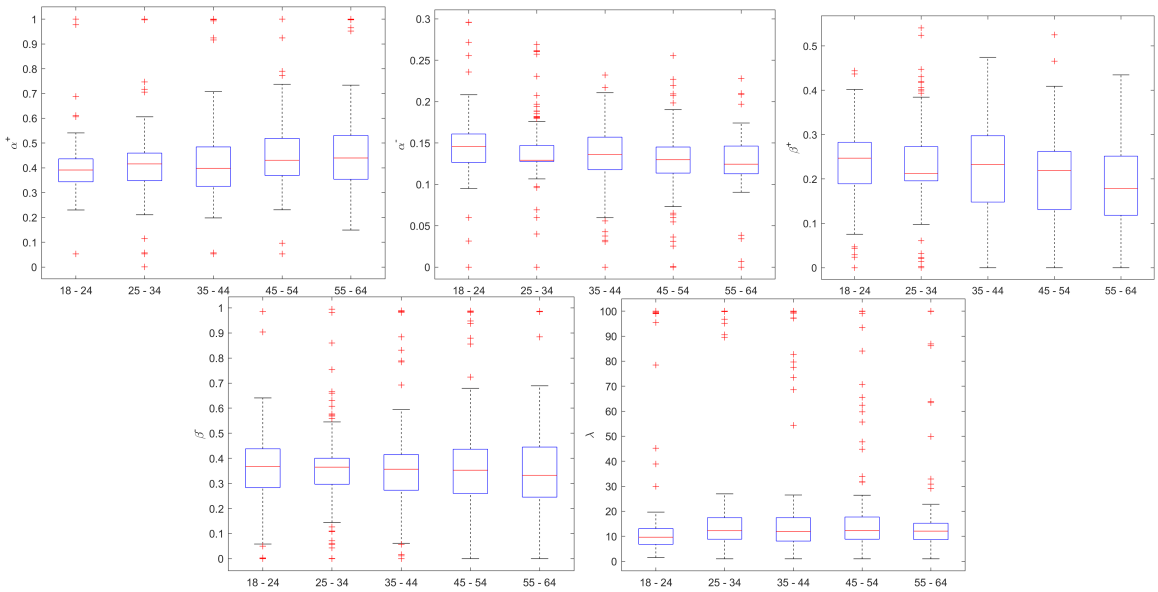

Figure 4-22: Variations in CPT parameters by age.

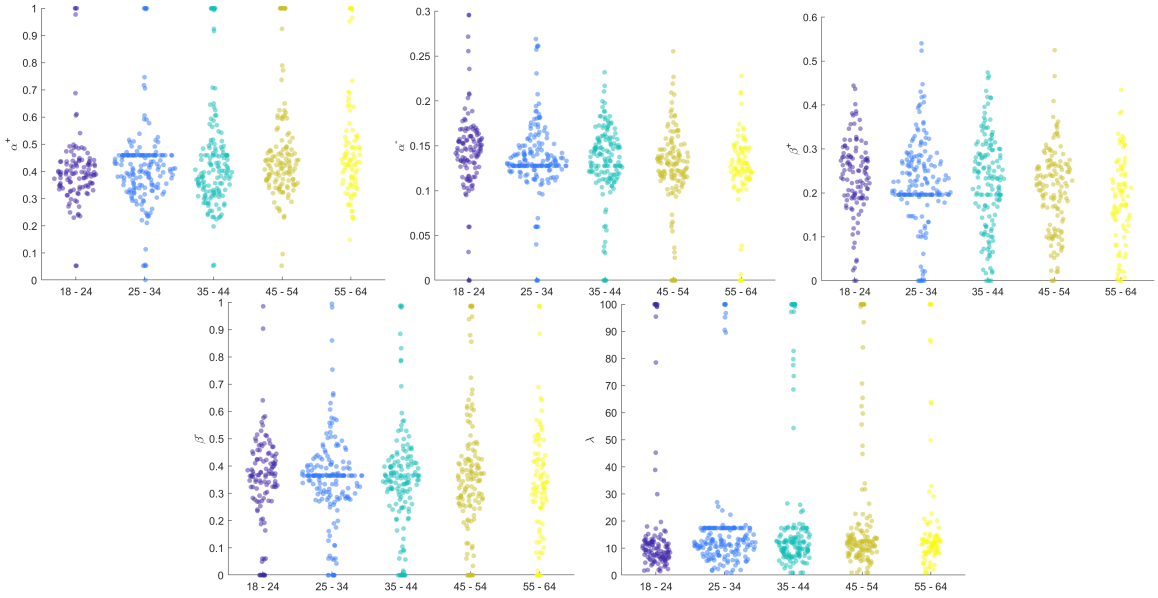

Figure 4-23: Swarm scatter charts showing the variation in CPT parameters by age.

4.2.13 Correlations between different CPT parameters

We examined pairwise correlations between all the CPT parameters, among all the valid respondents. The matrix in fig. 4-24 show the 2D scatter plots between each pair or parameters as well as their linear least squares regression lines. The slope of

these lines are also shown and correspond to the value of Pearson's linear correlation coefficient between the pair of parameters.

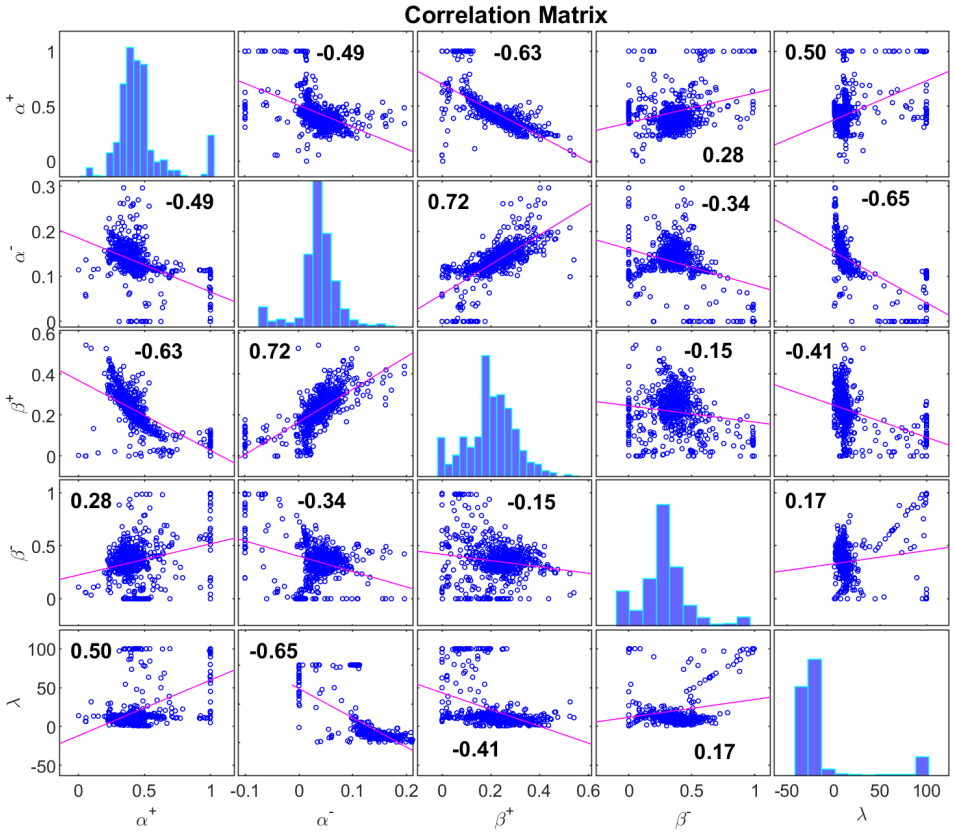

Figure 4-24: Correlation matrix between different pairs of CPT risk parameters.

Recall that higher values of these CPT parameters generally indicate a greater degree of irrationality since these deviate from traditional EUT behavior. Thus intuitively, one would accept that the parameters in the gain regime (α^+, β^+) would be positively correlated with one another and the same should be true for loss regime parameters $(\alpha^-, \beta^-, \lambda)$. However, what we find from fig. 4-24 is that this holds true in our sample only for the pair λ and β^-. This makes sense since people who are more sensitive to additional losses also tend to be more risk averse. But the opposite trend is true for all the other pairs i.e. (i) α^+ is negatively correlated with β^+, (ii) α^- is negatively correlated to both β^- and λ. Trend (i) implies that people who have a more distorted view of the likelihood of occurrence of gains tend to be less sensitive

to such gradational gains. Similarly, trend (ii) implies that people who have a more distorted view of the likelihood of occurrence of losses tend to be less sensitive to gradational losses and also less risk averse. Although these trends seem surprising at first, they actually do make intuitive sense upon closer inspection. Finally α^+ and α^- are negatively correlated, thus meaning that probability distortion in gains tends to be more or less balanced by that in losses.

4.3 Results from sensitivity and robustness analyses

4.3.1 Key insights from representative scenarios

Both analytical and numerical results were obtained for more than 100 randomly generated travel scenarios. These were created by varying u_0, $\underline{x}$, $\overline{x}$, b_{sm}, $\underline{\gamma}$, $\overline{\gamma}$ while ensuring that the resulting choice set was valid and involved fair comparisons between the uncertain SMoDS and the certain alternative travel option. The objective utilities of the SMoDS and alternative can take positive or negative values according to eq. (3.65). The sensitivity analysis was performed for each of these scenarios using nominal parameter values $\alpha_0 = 0.82$, $\beta_0 = 0.8$ and $\lambda_0 = 2.25$ estimated in [1]. The probability of the worst-case SMoDS outcomes was set as $p_0 = 0.75$. In order to compute the mismatch loss, the passenger's true behavioral parameters were set equal to their initial values at the nominal optimum, i.e., $\theta_{true} = \theta_0$ in eq. (3.93). All five of these model parameters are then varied by as much as $\pm 20\%$. Only five select scenarios are presented here to show some distinct behaviors and trends, as summarized in table 4.10. For all scenarios, the best-case SMoDS outcome was used as reference, i.e., $R = \overline{u}$. In this case, both the alternative A and the worst case SMoDS outcome are perceived as losses by the passenger. The acceptance probability can then be derived using eq. (1.1) and equations (3.64)-(3.73):

$$A_R^s = V(u_0) = -\lambda(\overline{u} - u_0)^\beta \tag{4.10}$$

$$U_R^s = \pi(p)V(\underline{u}) = -\lambda e^{-(-ln(p))^\alpha}(\overline{u} - \underline{u})^\beta \tag{4.11}$$

$$f = \gamma p_R^s = \frac{\gamma}{1 + e^{\lambda\left(e^{-(-ln(p))^\alpha}(\overline{x}-\underline{x})^\beta - (\overline{x}+b\gamma-u_0)^\beta\right)}} \tag{4.12}$$

Table 4.10: Selected representative scenarios.

Scenario	$u_0[-]$	$b_{sm}[\$^{-1}]$	$\underline{\gamma}[\$]$	$\overline{\gamma}[\$]$	$\underline{x}[-]$	$\overline{x}[-]$
$S1$	8.17	-0.14	4.66	8.41	2.46	15.45
$S2$	-7.62	-0.08	2.04	19.26	-8.67	-5.15
$S3$	-2.54	-0.72	4.12	12.99	0.32	10.98
$S4$	9.51	-0.40	1.11	13.49	1.06	24.36
$S5$	9.55	-0.04	4.24	7.92	4.41	12.92

Table 4.11: Sensitivity differentials of the optimal solution and value function near the nominal operating point, all evaluated at their respective nominal parameter values.

Scenario	$\frac{d\gamma^*}{d\alpha}$	$\frac{d\gamma^*}{d\beta}$	$\frac{d\gamma^*}{d\lambda}$	$\frac{d\gamma^*}{dp}$	$\frac{df^*}{d\alpha}$	$\frac{df^*}{d\beta}$	$\frac{df^*}{d\lambda}$	$\frac{df^*}{dp}$
$S1$	-2.8	-24.1	-2.9	-8.7	2.9	6.9	0.5	8.9
$S2$	-3.9	-20.8	-4.2	-12.1	4.3	8.4	0.9	12.9
$S3$	-4.5	-3.1	0.2	-13.8	4.6	-0.8	-0.7	14.1
$S4$	-4.3	-15.7	-1.3	-13.3	4.4	6.9	0.3	13.5
$S5$	-27.7	-99.7	-12.2	-84.5	1.2	3.8	0.4	3.6

Table 4.12: Valid domain for local sensitivity analysis (in %).

Scenario	$\lvert\Delta\alpha_{max}\rvert$	$\lvert\Delta\beta_{max}\rvert$	$\lvert\Delta\lambda_{max}\rvert$	$\lvert\Delta p_{max}\rvert$
$S1$	71.7	8.7	24.6	25.7
$S2$	250.8	48.9	85.6	89.9
$S3$	88.8	131.9	706.9	31.8
$S4$	51.9	14.8	63.5	18.6
$S5$	3×10^{-6}	10^{-6}	3×10^{-6}	10^{-6}

The plots in fig. 4-25 display variations in the optimal tariff, objective function value (maximum expected revenue) and mismatch loss for scenario $S1$. The variations in optimal solution and mismatch loss sometimes become flat w.r.t. to certain parameters as seen with β and p in fig. 4-25. This happens because γ^* hits either the lower or upper bound and doesn't change further unless the active set changes yet again due to even larger perturbations. We now present three major implications from analyzing these scenarios.

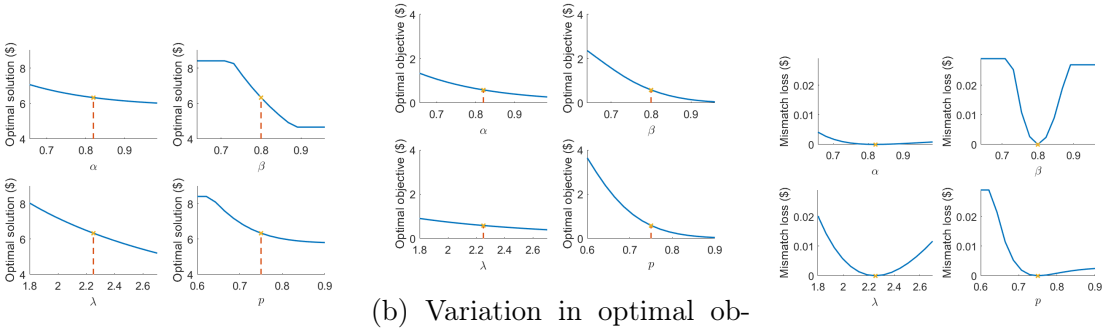

(a) Variation in optimal solution γ^* (SMoDS tariff).

(b) Variation in optimal objective f^* (maximum expected revenue).

(c) Variation in mismatch loss Δf.

Figure 4-25: Sensitivity results for $S1$.

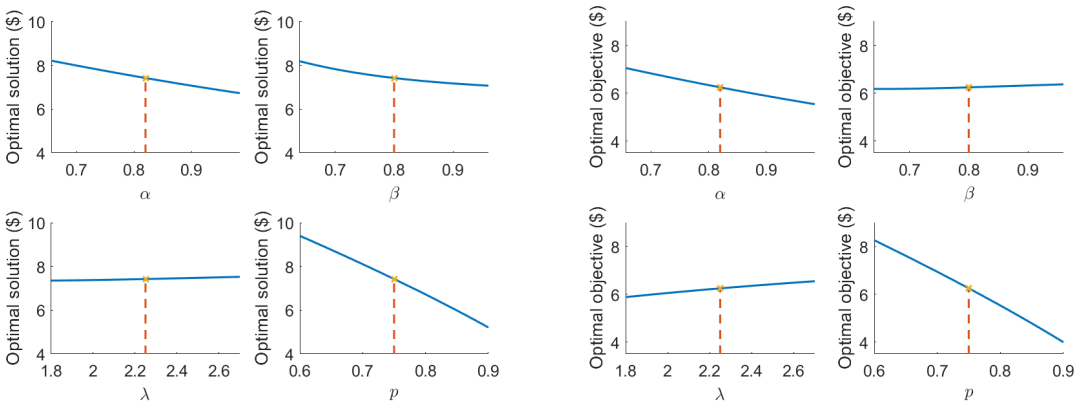

(a) Variation in optimal solution. (b) Variation in optimal objective.

Figure 4-26: Sensitivity results for $S3$.

General sensitivity trends for different parameters

Across all the scenarios studied, the optimal solution was found to be monotonically decreasing with respect to both α and p. This makes intuitive sense since as p increases, the likelihood of the best-case outcome decreases while the worst-case outcome becomes more likely. This makes the SMoDS relatively less attractive compared to the certain alternative A and thus the optimal SMoDS tariff γ^* must be reduced accordingly for it to remain competitive. Similarly, an increase in α indicates weaker probability distortion (i.e. under weighting high probability events and over weighting low probability events). Thus, as α increases, the user underestimates the likelihood of the more likely, worst-case outcome occurring to a lesser extent or equivalently overestimates the rarer, best-case outcome less. This reduces the SMoDS' relative attractiveness as well as γ^*. The opposite (monotonically increasing) trend would be observed with respect to α if best-case SMoDS outcome occurred with higher probability instead of the worst-case outcome, i.e., if $p < 0.5$. This was verified by also testing a few scenarios with $p_0 = 0.25$ instead of $p_0 = 0.75$, one of which is shown in fig. 4-27.

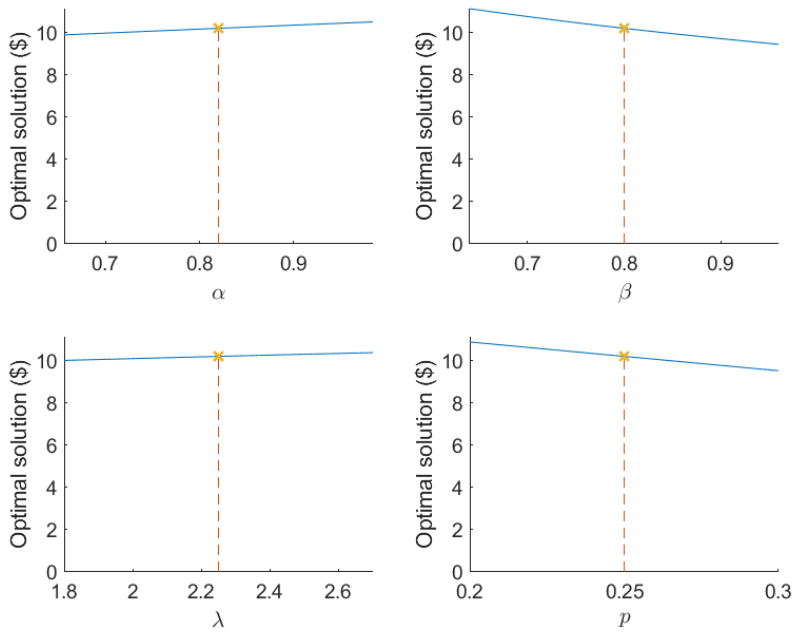

(a) Variation in optimal solution (SMoDS tariff).

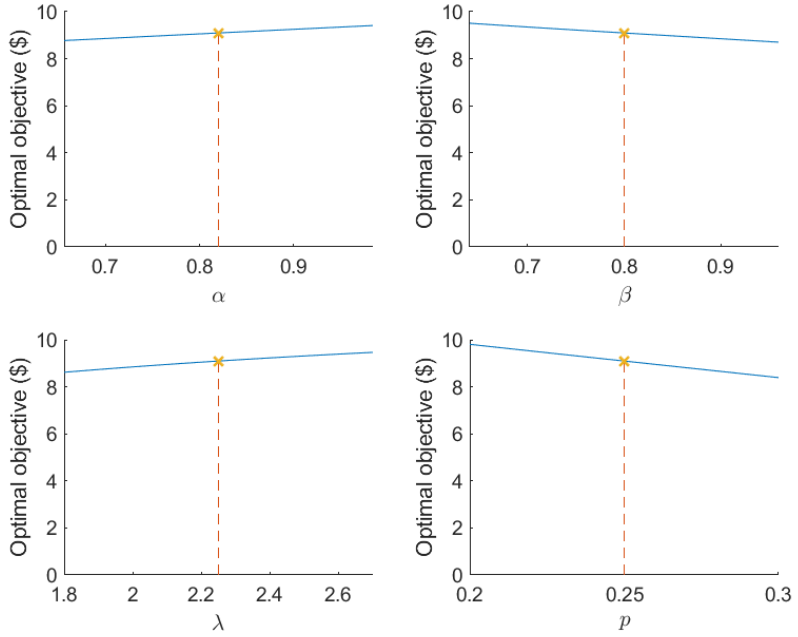

(b) Variation in optimal objective (expected revenue).

Figure 4-27: Sensitivity results for $S6$ with $p = 0.25$.

The variations in optimal objective are less straightforward to predict. Since $f^*(\gamma^*; \vec{\theta}) = \gamma^* \cdot p_R^s(\gamma^*; \vec{\theta})$ and p_R^s is strictly monotonically decreasing in γ [1], the

variation in expected revenue due to parameter changes will depend on how much the acceptance probability rises in response to a fall in γ^* and vice versa.

High degree of scenario specificity

Although we can make some broad statements about trends w.r.t. α and p, the same cannot be said for other parameters. Changes in α and p only affect the subjective utility of the SMoDS but variations in β and λ affect both the SMoDS as well as the alternative (eq. (4.12)). Thus, either a monotonically increasing or decreasing trend could be obtained depending on the relative magnitudes of $e^{-(-ln(p))^{\alpha}}(\overline{x}-\underline{x})^{\beta}$ and $(\overline{x}+b\gamma^*-u_0)^{\beta}$ in eq. (4.12). In other words, the behavior of the optimal solution in response to parameter perturbations depends on (1) the spread between the two possible SMoDS outcomes and (2) how this compares with the alternative. For instance, if $e^{-(-ln(p))^{\alpha}}(\overline{x}-\underline{x})^{\beta} > (\overline{x}+b\gamma^*-u_0)^{\beta} \implies U_R^s < A_R^s < 0$, i.e., the passenger perceives the worst-case SMoDS outcome as a greater loss (lower subjective utility) than the alternative. An increase in λ would make the user more averse to losses and thus persuade them away from the SMoDS in favor of the alternative. This reduces the relative attractiveness of SMoDS, causing γ^* to fall. The opposite effect occurs if $e^{-(-ln(p))^{\alpha}}(\overline{x}-\underline{x})^{\beta} < (\overline{x}+b\gamma^*-u_0)^{\beta}$ and the optimal tariff γ^* increases monotonically with λ, as seen for scenario $S3$ in fig. 4-26.

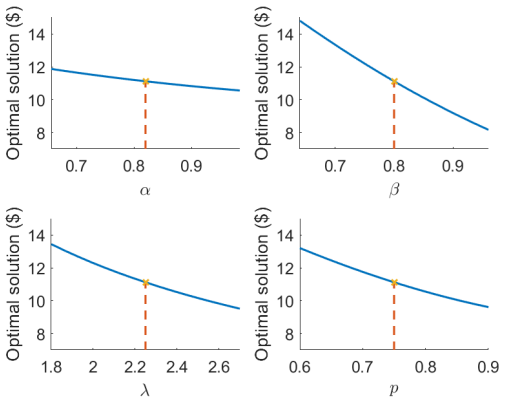

(a) Variation in optimal solution.

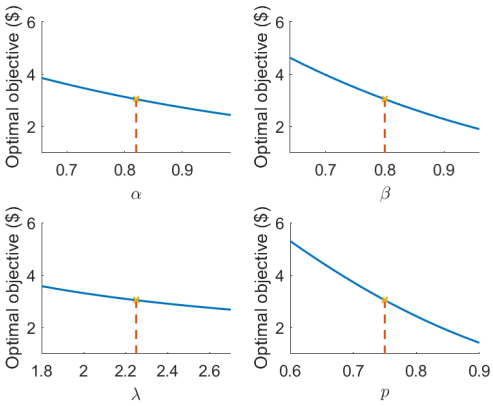

(b) Variation in optimal objective.

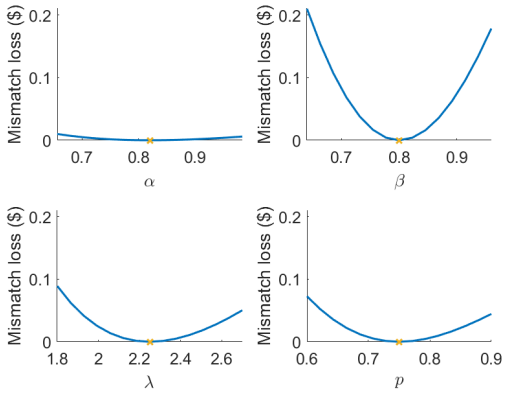

(c) Variation in mismatch loss.

Figure 4-28: Sensitivity results for $S2$.

Similarly, across a vast majority of the scenarios studied, γ^* was found to be decreasing with β. This agrees with intuition since $0 < \beta < 1$ describes the rate at which sensitivity (to both losses and gains) decreases as the outcome moves farther away from the reference. Thus, as β increases, the passenger becomes relatively more sensitive to losses even away from $R = \overline{u}$. Since $\underline{u} \leq u_0 \leq \overline{u}$, the passenger is now hurt more by the worst-case SMoDS outcome than before because the diminishing sensitivity effect is less influential. Thus, the optimal tariff γ^* decreases to compensate for this.

Furthermore, although we can sometimes predict the general direction of the variation as in section 4.3.1, the exact nature of the trend observed, i.e., the magnitude and rate of resulting changes, depends significantly upon the specific choice set and

travel scenario under consideration. These are determined by the objective utility of the certain alternative (u_0), travel time outcomes possible with the SMoDS ($\underline{x}$, $\overline{x}$), disutility associated with the SMoDS tariff (b_{sm}) and the bounds placed on the dynamic price ($[\underline{\gamma}, \overline{\gamma}]$). Even after experimenting with a large number of randomized scenarios, we were unable to obtain very specific scenario-agnostic insights that could be generalized. This is likely a consequence of the strongly nonlinear nature of the objective function. Thus, we concluded that the exact nature of variations in the optimal solution, value function and mismatch loss are determined to a large extent by the particular characteristics of each family of scenarios. In addition, the trends observed are likely to be affected by the reference level (R) used as well as the nominal values of the CPT parameters. For instance, we obtain almost linear variations for travel scenario $S2$ in fig. 4-28, which is quite different from the behavior seen for $S1$ in fig. 4-25. The relative sensitivities w.r.t. to each parameter can also vary quite significantly depending on the specific travel scenario as well as the nominal parameters, as indicated by the sensitivity differentials in table 4.11.

Robustness under certain scenarios

The local sensitivity domain was found to be quite large for most combinations of nominal parameters and travel scenarios, as shown by $S1$-$S4$ in table 4.12. This implies that a relatively large perturbation in parameters is needed to alter the nominal active set. Such a property can be exploited to set more robust dynamic tariffs. If the ride offer and travel scenarios can be designed such that the optimal tariff is at either the upper ($\overline{\gamma}$) or lower ($\underline{\gamma}$) bound, then γ^* would remain unchanged even with large errors or misconceptions in the parameter values assumed for the passenger. Although we assumed the tariff bounds to be exogenous to the travel scenarios in this study, the choice of this set $[\underline{\gamma}, \overline{\gamma}]$ will likely involve another optimization problem in itself. In practical applications, the bounds specified on the tariff would be determined depending on the given scenario and would need to be updated for each new passenger and trip. This aspect will be looked into further as part of future work.

4.3.2 Comparison of numerical results and analytical solutions

In general, for a majority of the scenarios considered, linearized approximations using 1^{st} order Taylor expansions provide a reasonable estimate of the optimal solution for small *local* perturbations (e.g. $\pm 20\%$) in the neighbourhood of the nominal parameter value. The analytical solutions obtained using a local sensitivity analysis, for both the optimal tariff and value function are very close to the exact values obtained numerically through repeated optimization, with $< 0.5\%$ error for most scenarios. However, larger discrepancies were obtained for a small subset of scenarios. For each of these scenarios, at least one of the following conditions was found to be true:

1. The variations in optimal tariff and optimal objective are clearly nonlinear and display significant curvature even for small perturbations from the nominal parameter values. In such cases, the error in the 1^{st} order linear approximation of the optimal solution is no longer negligible, as can be seen for parameter β and p in fig. 4-29a. However, for most of the scenarios, we observe close to linear variations in optimal objective for small to moderate perturbations. This makes intuitive sense since all four parameters (α, β, λ and p) are of the order of 1. Thus, the magnitude of deviations $\Delta\theta = \theta - \theta_0 < 1$ for small to medium perturbations, implying that the 2^{nd} and higher order terms in eq. (3.88) can usually be ignored for scenarios where curvature is not important i.e. $\frac{d^2 f^*}{d\theta^2}(\theta_0) << \frac{df^*}{d\theta}(\theta_0)$.

2. The local domain ($\Delta\theta_{max}$) is very small as for scenario $S5$ in table 4.12, implying that even a slight perturbation in the parameter would change the active constraint set. This makes the optimal solution more sensitive to parameter uncertainty since γ^* either (1) switches between the two bounds, (2) hits one of the bounds from the interior of $[\underline{\gamma}, \overline{\gamma}]$ or (3) leaves one of the bounds for the interior of the set. This causes stronger nonlinear behavior in the optimal value function too and makes a local analysis insufficient, as seen for β and p in fig. 4-29b.

For such cases, a global analytical method like that described in section 3.7.8 could be used instead to obtain more accurate analytical solutions through better, piecewise linear approximations for the optimal tariff.

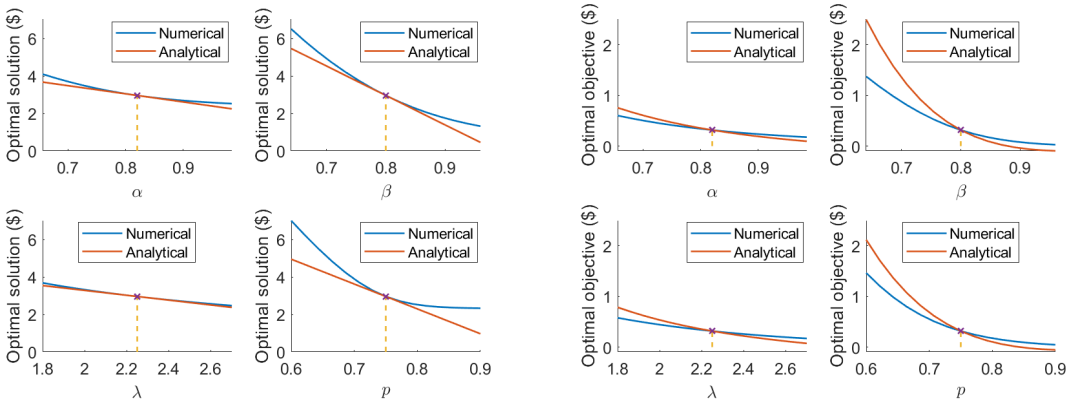

(a) Optimal price variation in $S4$. (b) Optimal objective variation in $S5$.

Figure 4-29: Analytical (local) vs numerical results.

Chapter 5

Towards transactive control

5.1 Background and motivation

After estimating and calibrating the CPT model, we then proceeded to assess the sensitivity and robustness of the SMoDS system to the uncertainty involved in the various risk parameters. We observed that such parametric uncertainty can significantly impact the key performance indicators (KPIs) of the SMoDS. This chapter builds upon these insights to explore how we can improve the robustness of the system using feedback control schemes. In particular, we describe our preliminary efforts at using a few different control methods to respond to, and rectify CPT model parametrization errors.

Using eq. (3.13)-eq. (3.19), we can derive that the probability of acceptance of the SMoDS is a strictly monotonically decreasing function $f(\cdot)$ of the dynamic tariff γ [1]:

$$p_R^s = f(\gamma; \Theta, R) = f(\gamma; \alpha^+, \alpha^-, \beta+, \beta^-, \lambda, R)$$

where $\Theta = [\alpha^+, \alpha^-, \beta+, \beta^-, \lambda]$ represents all five CPT parameters with some level of associated uncertainty. For the purposes of this study, we assume that the true reference point R is known with zero uncertainty. The goal of the SMoDS operator

or server is to optimize the collective performance across the entire SMoDS platform serving a particular region. The main performance metric we consider here is the average estimated waiting time (EWT) which needs to be regulated around a specific desired value EWT^* in order to balance real time demand and supply for ride sharing services. Thus, the optimal $EWT^*(t)$ varies with time depending on travel demand, traffic conditions, weather etc. Previous work in our lab has developed an approximate dynamic programming or reinforcement learning based algorithm to derive the desired acceptance probability $p^*(t)$ needed for $EWT(t) \rightarrow EWT^*(t)$ [64]. The goal of this work is as follows: Given the value of $EWT^*(t)$ and consequently the desired $p^*(t)$ that needs to maintained at each time period t, we aim to regulate the actual probability of acceptance $p_R^s(t)$ such that it approaches and stays near this setpoint.

For this purpose, we can treat the SMoDS as a closed loop system as shown in fig. 5-1, by integrating all the individual modules such as (i) dynamic routing via alternating minimization, (ii) computing p^* from EWT^* via reinforcement learning and (iii) dynamic pricing via CPT. Analogous to traditional feedback control theory, we can treat the plant to be the SMoDS passenger requesting a ride while the dynamic tariff γ is the primary control input into the system. The plant is described by our nonlinear CPT passenger behavioral model parametrized by the parameters Θ. The true plant model is not known, but we assume that both the true and estimated models share the same functional form $f(\cdot)$. They differ only in terms of their parametrizations, with the estimated model parameters being $\widehat{\Theta}$ and the true parameters represented by Θ^*. If the true behavioral parameters were known for each passenger being served, we could simply invert the CPT model to readily determine the optimal dynamic price $\gamma^* = f^{-1}(p^*)$ needed to achieve the equilibrium probability of acceptance p^*, which can be viewed as the reference or setpoint. However, since Θ^* is not known *apriori*, we need to consider alternative approaches to update $\hat{\gamma}$ in real time such that it approaches γ^* and $p_R^s \rightarrow p^*$. Since the SMoDS server has access to the accept/reject decisions for all the passengers and trips at each time instant, these can be used to calculate the actual acceptance probabilities at each time step, as well

as the reference errors $p^* - p_R^s$. These output errors serve as feedback signals and an appropriately designed controller can then compute the new dynamic tariffs for each ride-offer in the next time step.

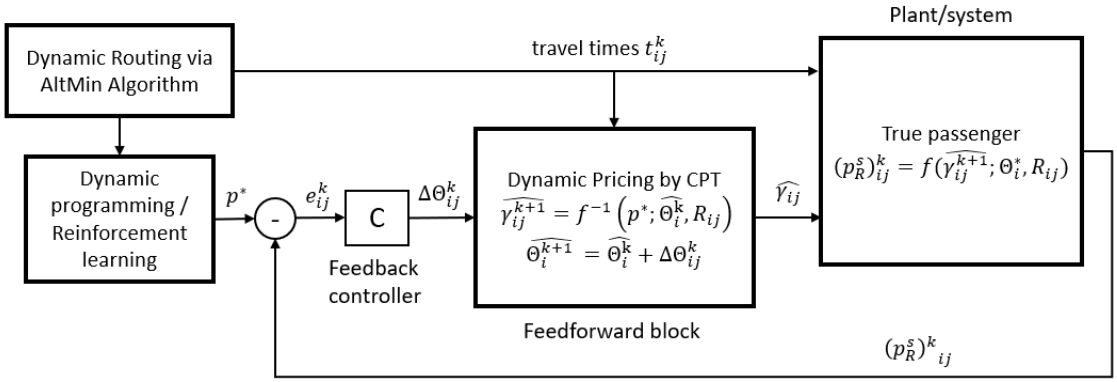

Figure 5-1: The closed loop SMoDS system, using CPT *parameter* updates as the feedback signal. The indices i, j and k correspond to the passengers, trip requests (or equivalently ride offers) and negotiation (or feedback) iterations, respectively.

Ideally, our control scheme would achieve all of the following objectives:

1. Achieve good tracking of the reference over time, i.e., $(p_R^s)_{ij} \to p^*$ for all relevant passengers i and trip requests (or ride offers) j.

2. Convergence to the true model parameters for each passenger with time i.e., $\widehat{\Theta}_i(t) \to \Theta_i^*(t)$ for all passengers i. Note that in addition to modelling uncertainty, another reason that these parameters may differ from their assumed values is because people's risk preferences often vary with time and depending on the specific circumstances at hand.

subject to several possible constraints:

- Achieve convergence (in both 1. and 2. above) in as few iterations or time steps k as possible. This constraint is needed to (i) prevent passengers from experiencing decision fatigue and (ii) ensure that the feedback mechanism and negotiation process itself don't markedly increase waiting times.

- Budgetary and fiduciary constraints of the SMoDS provider, which may be a profit-maximizing entity aiming to increase the expected revenue from the service.

- Other constraints such as serving a certain minimum ridership level, limits of trip duration, etc.

For this problem, we attempt to achieve both objectives but for simplicity, we largely focus on the 1st constraint of minimizing the number of iterations to attain fast convergence.

5.2 Brief literature survey

Feedback control has been applied to several problems similar to the SMoDS. These include transactive control using dynamic toll pricing [65, 66], adaptive pricing control for EVs [67], congestion pricing via optimal control [68] and threshold-based distributed control for road pricing [69]. Outside the transportation domain, it has been used to compute and update consumer preferences, trust and reputation scores for online shopping and e-commerce settings [70, 71]. In this study, we seek to use the notion of transactive control which utilizes variable price signals to influence the behavior of decision makers, as opposed to directly controlling the 'actuator'.

In terms of theoretical foundations, several different approaches have been suggested for feedback control of nonlinear systems, which are generally represented as:

$$\dot{x} = f(x, u) \tag{5.1}$$

$$y = g(x) \tag{5.2}$$

where x is the system's state, u is the control input and y is the measured output. In particular, the sub-fields of adaptive and robust control focus on controlling systems where the plant model f is uncertain. Adaptive control focuses on dual goals of both reference tracking and parameter convergence, as is the case with our problem

[72]. These are generally based on the parameter perturbation method [73] and its variants include direct adaptive fuzzy control [74, 75], model reference adaptive control (MRAC) [76] and extremum seeking control (ESC) [77]. In particular, ESC is a perturbation-based approach used for adaptive control of hard to model continuous or discrete time systems [78]. It is of special interest here since it seeks to minimize (i.e. reach an extremum of) some objective or loss function, using input perturbation signals. Thus, it is a popular tool used for real-time, model-free, online optimization and learning the true model parameters of dynamic systems. The scheme approximates the gradient of the objective function by passing its output through a high-pass filter and then multiplying it by the perturbation signal used. Sometimes, this is also followed by a low-pass filter in order to attenuate effects of any measurement noise in the gradient signal, as seen in section 5.2 below where a sinusoidal input perturbation is used.

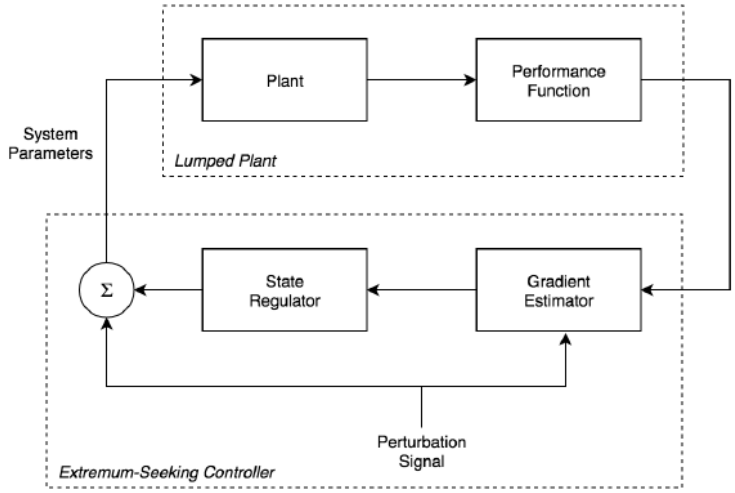

(a) Components of a general extremum seeking scheme [79].

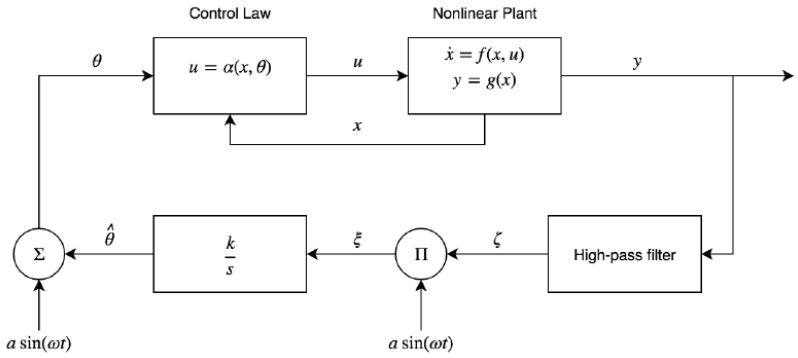

(b) Illustrative example of an extremum seeking controller [79].

Some other common approaches for nonlinear control include [80]:

1. Jacobian linearization of the dynamic system about one or more operating points.

2. Feedback linearization wherein the nonlinear system is transformed to an equivalent linear system using a change of variables.

3. Gain scheduling where the system's operational domain into several distinct sub-regions and the system can be approximated as being linear (with different properties) within each of these.

4. Lyapunov and K-function based methods such as (i) Lyapunov redesign wherein

the system is stabilized using a state feedback controller designed using its Lyapunov function and (ii) sliding mode control which applies discontinuous control signals.

However, our closed loop, SMoDS system differs quite significantly from the systems usually modelled using the above approaches. Firstly, our system or plant (i.e. the passenger or decision maker) is not modelled as having any internal state or dynamics and there is thus no explicit governing differential equation. Instead, we are only interested concerned with regulating the system's output (p_R^s) to track p^*. Furthermore, many of these methods have largely been applied to linear or linearized systems, usually under some idealized assumptions. In contrast, our CPT plant model $f(\gamma;\Theta)$ is highly nonlinear and complex. This makes it challenging to linearize and intractable to apply other involved analytical techniques. Finally, since the reference trajectory is known in our case, the exact form of the objective function can be computed here, which is not the case with methods like ESC. This leads us to gradient descent based control methods, which is the approach we ended up applying for this work. Such methods couple optimization problems with feedback, deriving control laws using gradient descent updates [81]. They have been long applied to optimize performance functions and control nonlinear systems for end goals such as output tracking [82, 83]. Feedback-based methods using gradient sampling [84], primal-dual algorithms [85] and discretized gradient flows [86] have achieved good convergence properties even for nonconvex and non-smooth problems.

5.3 Design of feedback controller

As explained earlier, we close the loop in the SMoDS system shown in fig. 3-1 by feeding back the error between the actual and desired acceptance probabilities, as illustrated in fig. 5-1. The goal of the feedback control loop can be formulated as an

equivalent optimization problem:

$$\min_{\hat{\gamma},\widehat{\Theta}} \quad E\left(\hat{\gamma},\widehat{\Theta}\right) = \mid p_R^s(\hat{\gamma};\widehat{\Theta}) - p^* \mid \tag{5.3}$$

$$\text{s.t.} \quad \underline{\gamma} \leq \hat{\gamma} \leq \overline{\gamma} \tag{5.4}$$

$$\underline{\widehat{\Theta}} \leq \widehat{\Theta} \leq \overline{\widehat{\Theta}} \tag{5.5}$$

For our purposes, minimizing the absolute error above is equivalent to instead minimizing the squared L_2 norm objective function instead:

$$L\left(\hat{\gamma},\widehat{\Theta}\right) = \frac{1}{2}\left\|\mathbf{p}_R^s(\hat{\gamma};\widehat{\Theta}) - \mathbf{p}^*\right\|_2^2 \tag{5.6}$$

While considering a single passenger trip pair, this reduces to:

$$L\left(\hat{\gamma}_{ij}^k,\widehat{\Theta}_i\right) = \frac{1}{2}\left(p_R^s\left(\hat{\gamma}_{ij}^k;\widehat{\Theta}_i\right) - p^*(t)\right)^2 \tag{5.7}$$

which is convex if the plant model function $p_R^s(\hat{\gamma};\widehat{\Theta}) = f$ is convex in both the input γ and parameters Θ. However, note that convexity of f is not guaranteed - it can be convex, convex and neither convex nor concave, depending on various factors like the reference R used, domain of possible values assumed by the CPT parameters and the range of allowed tariffs. We assumed that the desired acceptance probability p^* for the overall system changes on much slower timescales (if at all) compared to the time intervals between negotiation iterations since the server samples decisions very frequently (say every 30-60 s).

We propose the use of gradient descent based feedback control to repeatedly the dynamic tariff at each iteration or time step k. This is done for every passenger-trip pair (i, j). The idea then is that passengers and the SMoDS server will engage in several rounds of back-and-forth negotiation iterations where the server provides a ride offer and the user either accepts or rejects it. If the user accepts, they leave the current pool of passengers being considered by the controlller. If the ride offer is

rejected, the server uses the feedback controller to compute the updated price for the next ride offer at iteration $k+1$. This process continues for each passenger-trip request pair until the passenger either accepts the offers or leaves the SMoDS platform, with the goal being to achieve ride acceptance in as few rounds as possible. A centralized controller will be implemented for each SMoDS hub or server operating in a particular neighbourhood or region. This server samples all the ride offer decisions made by passengers within its purview at each iteration. These are then used to accordingly update the actual probability of acceptance. Thus, the goal of the controller sited at each each server is to have the average acceptance probability over its entire region or scope, approach the desired value p^* and also track time-varying references $p^*(t)$. Such a program can easily be implemented through platforms similar to those used by ride sharing apps today, keeping track of a customized CPT model unique to each rider. Broadly, there are two possible paradigms for such gradient-descent based control, when applied to our SMoDS problem.

5.3.1 Directly using tariff as control input

The first strategy varies the dynamic tariff $ga\hat{m}ma$ directly as the control input to influence the acceptance probability. This is a relatively easier approach and more intuitive since we can leverage the strictly monotonic relationship between p_R^s and γ, the acceptance probability monotonically decreases with increasing price. Thus, we can readily come up with naive update rule as follows for passenger i, trip j at iteration k:

$$\text{If } (p_R^s)_{ij}^k < p^*(t) \implies e_{ij}^k = p^* - ((p_R^s)_{ij}^k > 0 \implies \text{Lower } \hat{\gamma} \text{ to increase } (p_R^s)_{ij}^k \quad (5.8)$$

$$\text{If } (p_R^s)_{ij}^k > p^*(t) \implies e_{ij}^k = p^* - ((p_R^s)_{ij}^k < 0 \implies \text{Raise } \hat{\gamma} \text{ to decrease } (p_R^s)_{ij}^k \quad (5.9)$$

$$\hat{\gamma}^{k+1} = \hat{\gamma} - Ce^k \quad (5.10)$$

where C is the step size or proportionality gain that needs to be tuned. However, such a simple proportional control scheme is not expected to generalize well since

our SMoDS is highly nonlinear as shown earlier. For a specific example of what the acceptance probability function might look like, see eq. (4.12). Thus, we also formulate a direct gradient descent based control scheme.

When a new passenger-trip pair is added to the pool, the CPT model is inverted using the current p^* value to set the initial price at the 1^{st} iteration $k = 1$. This functional inverse is used to determine the tariff only when either the setpoint $p^*(t)$ changes or a new passenger-ride pair joins the sample. Otherwise, in subsequent iterations, the tariff is updated using gradient descent to minimize our squared error norm or loss function:

$$\hat{\boldsymbol{\gamma}}^{k+1} = \hat{\boldsymbol{\gamma}}^{k} - \eta_k \frac{\partial L}{\partial \boldsymbol{\gamma}}\bigg|_{\hat{\boldsymbol{\gamma}}^k} \tag{5.11}$$

$$\frac{\partial L}{\partial \boldsymbol{\gamma}} = \frac{1}{2}\frac{\partial \left\| \mathbf{p}_R^s(\hat{\gamma}; \widehat{\Theta}) - \mathbf{p}^* \right\|_2^2}{\partial \boldsymbol{\gamma}} = \frac{1}{2}\frac{\partial \left(\mathbf{p}_R^s - \mathbf{p}^*\right)^\intercal \left(\mathbf{p}_R^s - \mathbf{p}^*\right)}{\partial \boldsymbol{\gamma}} \tag{5.12}$$

$$= \frac{1}{2}\frac{\partial \left(\mathbf{p}_R^s - \mathbf{p}^*\right)^\intercal \left(\mathbf{p}_R^s - \mathbf{p}^*\right)}{\partial \mathbf{p}_R^s} \cdot \frac{\partial \mathbf{p}_R^s}{\partial \boldsymbol{\gamma}} = \left(\mathbf{p}_R^s - \mathbf{p}^*\right) \nabla_{\boldsymbol{\gamma}} \mathbf{p}_R^s \tag{5.13}$$

where we have applied the chain rule and used the fact that $\mathbf{p}^*$ is a constant vector between iterations (but may change over longer time periods). Note that $\boldsymbol{\gamma}$ and $\mathbf{p}_R^s$ are vectors of tariffs and acceptance probabilities spanning all the passenger-trip pairs at that iteration. The step size or learning rate η_k can either be fixed or vary with each iteration. Since the server has access to the passenger decisions at each iteration, the resulting iteration errors can be used to approximately compute the gradients of the acceptance probability with respect to the tariff, using a simple first order finite difference scheme:

$$\nabla_{\boldsymbol{\gamma}} \mathbf{p}_R^s\big|_k = \frac{\partial \mathbf{p}_R^s}{\hat{\boldsymbol{\gamma}}^k}\bigg|_{\hat{\boldsymbol{\gamma}}^k} \approx \frac{\Delta \mathbf{p}_R^s}{\Delta \boldsymbol{\gamma}}\bigg|_{\hat{\boldsymbol{\gamma}}^k} = \frac{\left(\mathbf{p}_R^s\right)^k - \left(\mathbf{p}_R^s\right)^{k-1}}{\hat{\boldsymbol{\gamma}}^k - \hat{\boldsymbol{\gamma}}^{k-1}} \tag{5.14}$$

While this method is relatively straightforward to implement, it does suffer from some limitations. Namely, such a scheme does not allow us to learn the true behavioral model of the passenger since there isn't a clear way to infer the CPT parameters Θ

based on how p_R^s changes in response to tariff updates.

5.3.2 Indirect control by updating the CPT parameters

As an alternative approach to updating the tariff directly, one can instead update the five CPT parameters contained in $\widehat{\Theta}$ at each iteration and for all the passengers currently in the pool. This is the closed loop workflow illustrated in fig. 5-1. These updated parameter values in turn also update the tariff at the next iteration, through the CPT model $f(\cdot)$. Similar to adaptive control, the key advantage of this method is that with a well-designed controller, both our goals of reference tracking (driving the racking error to zero) and parameter estimation can theoretically be achieved. In other words, with enough iterations, we can attain convergence of both the output $p_R^s \to p^*$ as well the true behavioral model parameters $\hat{\Theta} \to \Theta^*$. The caveat is that it is more challenging to construct, implement and analyze the properties of this controller due to the complex, nonlinear and non-monotonic relationships between p_R^s and each of the CPT parameters.

The algorithmic workflow is similar to that used in section 5.3.1 since we again use a gradient descent based update scheme:

$$\hat{\boldsymbol{\Theta}}^{k+1} = \hat{\boldsymbol{\Theta}}^k - \eta_k \frac{\partial L}{\partial \boldsymbol{\Theta}}\bigg|_{\hat{\boldsymbol{\Theta}}^k} \tag{5.15}$$

$$\frac{\partial L}{\partial \boldsymbol{\Theta}}\bigg|_{\hat{\boldsymbol{\Theta}}^k} = \left(\mathbf{p}_R^s - \mathbf{p}^*\right) \nabla_{\boldsymbol{\Theta}} \mathbf{p}_R^s \tag{5.16}$$

$$\nabla_{\boldsymbol{\Theta}} \mathbf{p}_R^s = \frac{\left(\mathbf{p}_R^s\right)^k - \left(\mathbf{p}_R^s\right)^{k-1}}{\hat{\boldsymbol{\Theta}}^k - \hat{\boldsymbol{\Theta}}^{k-1}} \tag{5.17}$$

One important difference from the previous scheme is that here, the CPT model is inverted using the updated parameters to update the tariff at *every* iteration, not just at $k = 1$:

$$\hat{\boldsymbol{\gamma}}^{k+1} = f^{-1}\left(\mathbf{p}^*; \hat{\boldsymbol{\Theta}}^{k+1}\right) \tag{5.18}$$

5.3.3 Learning rate considerations

If the CPT-based plant model f is convex in γ, then the 1^{st} method from section 5.3.1 is guaranteed to converge to unique globally optimal solutions that minimize the error in the acceptance probability. Similarly, if f is convex in each of the CPT parameters $\Theta = [\alpha^+, \alpha^-, \beta+, \beta^-, \lambda]$, the 2^{nd} method from section 5.3.2 has guaranteed global convergence. However, the plant model may sometimes be nonconvex for certain passengers or trips depending on the specific combinations of tariff values, CPT parameters and reference points used. In such cases, we cannot provide any robust theoretical guarantees for guarantees for global convergence. Although such batch gradient descent-based methods will converge to local minima, these may not be global minimizers of the objective [87].

In order to increase the speed of convergence, improve stability and prevent the controller from getting stuck at local minima, we experimented with several different learning rates or step sizes. The initial implementation was vanilla gradient descent using a fixed step size $\eta_k = \eta \; \forall \; k$, and tuning was done via trial and error. However, it was observed that learning often stagnates after a few iterations resulting in a relatively large steady state error for the output. Furthermore, stochastic gradient descent (SGD) is not well-suited to this application since it would cause large, random fluctuations in the objective function. This is because it updates parameters more frequently with higher variance. However, we would ideally like to avoid such erratic swings in tariffs and acceptance probabilities, from one iteration to the next. Thus, we implemented a few different learning rate schedules as well as adaptive learning rates [87]:

- **Inverse time decay**: These are schedules which modify the learning rate at each iteration k based on the previous iteration's learning rate.

$$\eta_k = \frac{\eta_0}{1 + dk} \tag{5.19}$$

  where η_0 is the initial learning rate at $k = 1$ and d is the decay factor.

- **Exponential decay**: This schedule changes the learning rate at each iteration to follow a decaying exponential curve.

$$\eta_k = \eta_0 e^{-dk} \tag{5.20}$$

where η_0 is the initial learning rate and d is the decay hyperparameter.

- **Step decay**: This schedule drops the learning rate by a fixed factor d after every n iterations or epochs.

$$\eta_k = \eta_0 d^{\lfloor (\frac{k}{n}) \rfloor} \tag{5.21}$$

- **Momentum**: This is a method that dampens the oscillations associated with SGD and accelerates its descent in the desired direction. It computes the next update step as a linear combination of the current gradient and the previous iteration's update.

$$\Delta_k = \alpha \Delta_{k-1} + \eta \nabla_\theta L(\theta_{k-1}) \tag{5.22}$$

$$\theta_k = \theta_{k-1} - \Delta_k \tag{5.23}$$

where θ is the variable of interest that we're optimizing with respect to, i.e., either Θ or γ. The momentum term α is usually set equal to 0.9 or a value close to that.

- **Nesterov acceleration**: This is an anticipatory method that looks ahead and computes the update using the gradient relative to the approximate parameter values in the next time step, rather than w.r.t the current parameters. This prevents the algorithm from moving too quickly, increases the responsiveness

and improved performance.

$$\Delta_k = \alpha \Delta_{k-1} + \eta \nabla_\theta L(\theta_{k-1} - \alpha \Delta_{k-1}) \tag{5.24}$$

$$\theta_k = \theta_{k-1} - \Delta_k \tag{5.25}$$

Finally, in order to improve convergence properties, norm clipping was performed by normalizing the gradient updates by their vector L_2 norms at each iteration.

5.4 Numerical simulation results

The numerical simulation to emulate the closed loop SMoDS system was set up in a manner similar to that described in section 3.7. However, several of the assumptions made earlier (for analytical tractability) have been relaxed here to model more general scenarios. In particular, we allow several different types of reference points R and also scale up the setup to multiple passengers and rides. We show some preliminary results applying our control to several randomly generated travel situations for a sample size or pool of $N = 500$ passenger-trip pairs. The acceptance probabilities for all the passengers were initialized by adding random Gaussian noise to the desired value p^*. A large number of trials were conducted using different types of learning rates and multiple combinations of hyperparameters such as the initial step size, decay factor, iterations per step decay, momentum factor etc. For conciseness, only a few selected results are included below to illustrate some of the key trends.

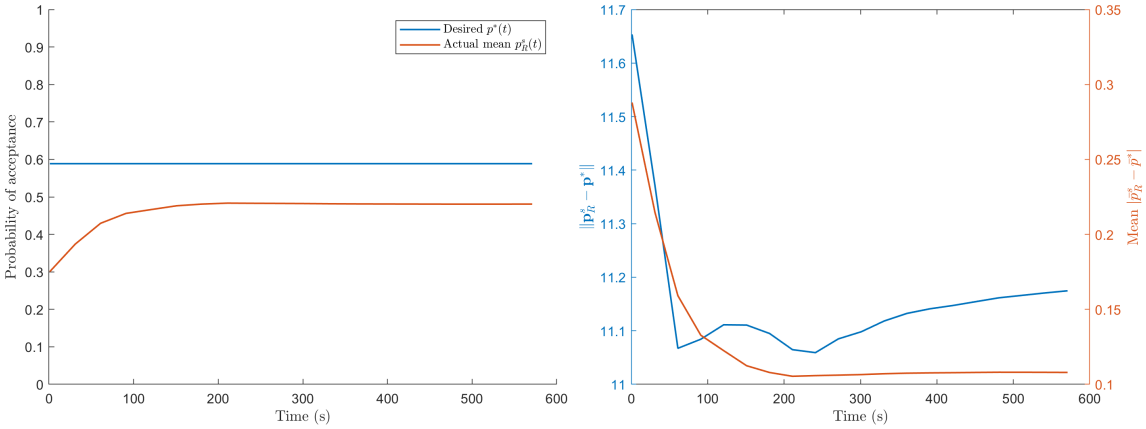

(a) Plot of desired vs actual probabilities of acceptance, averaged over the whole sample.

(b) Plot of the mean absolute error and the vector L_2 norm of error.

Figure 5-3: Direct tariff-based control strategy using adaptive gradient descent with momentum, $\alpha = 0.8, \eta = 0.1$.

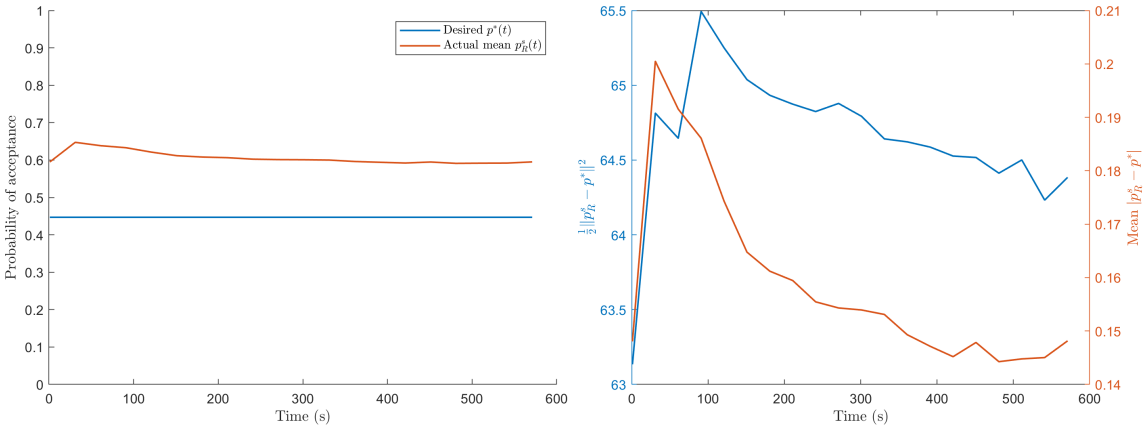

(a) Plot of desired vs actual probabilities of acceptance, averaged over the whole sample.

(b) Plot of the mean absolute error and the vector L_2 norm of error.

Figure 5-4: Direct tariff-based control strategy using adaptive gradient descent with momentum, $\alpha = 0.9, \eta = 0.4$.

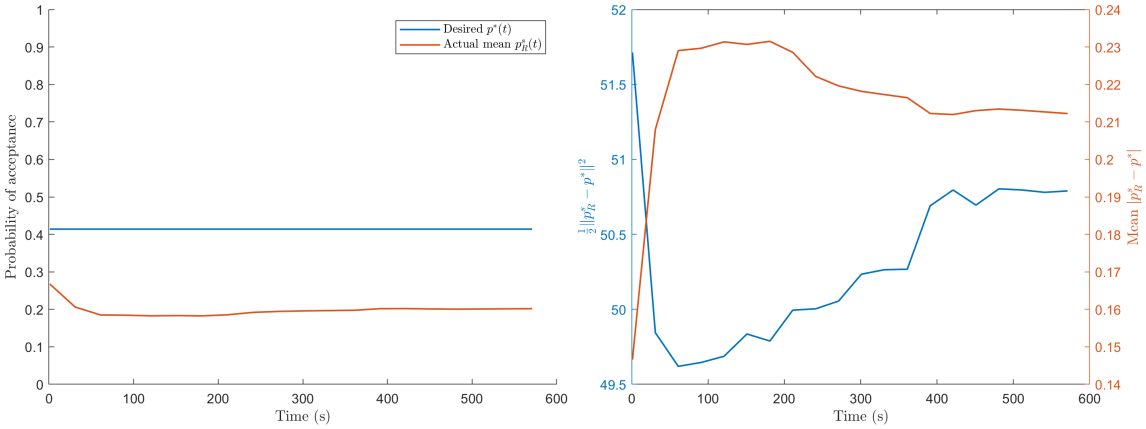

(a) Plot of desired vs actual probabilities of acceptance, averaged over the whole sample.

(b) Plot of the mean absolute error and the vector L_2 norm of error.

Figure 5-5: Direct tariff-based control strategy using adaptive gradient descent with momentum, $\alpha = 0.9, \eta = 0.1$.

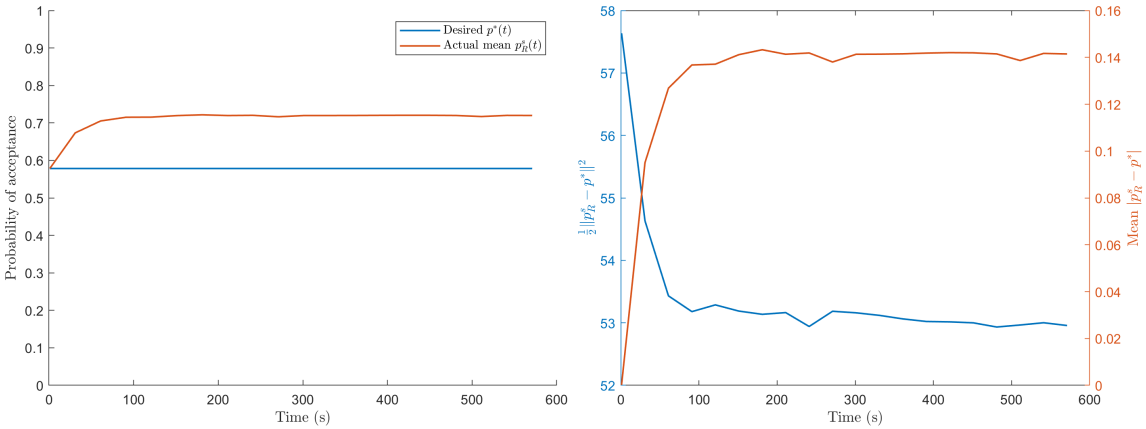

(a) Plot of desired vs actual probabilities of acceptance, averaged over the whole sample.

(b) Plot of the mean absolute error and the vector L_2 norm of error.

Figure 5-6: Direct tariff-based control strategy using gradient descent with fixed step size $\eta = 0.2$.

We find that the performance of the controller is quite sensitive to the hyperparameters used. Furthermore, the behavior was found to fluctuate quite drastically between runs depending on (i) attributes of the trips sampled as well as (ii) initializations of the parameters and acceptance probabilities. For several trials, the controller performed reasonably well in terms of tracking the reference and gradually reducing the error as p_R^s approaches the setpoint p^*, as seen in fig. 5-3 and fig. 5-4.

However, the learning does stagnate eventually, resulting in a non-zero steady state error. Even with adaptive methods using momentum and Nesterov acceleration, this decay in learning was observed. For certain hyperparameters, I was able to reduce the steady state error significantly, but never quite drive it to zero altogether.

However, for certain other scenarios, the control performed poorly in that the acceptance probabilities did not change much at all or they moved opposite to the desired direction, thus increasing the error even further. In some of these situations, this behavior was caused by local concavities in the plant model function $f(\cdot)$ and it could be rectified by simply negating the gradient updates. However, in other situations, this modification did not resolve the issue either as can be seen in fig. 5-5 and fig. 5-6.

In addition, for a few cases, the controller exhibits some undesirable and non-intuitive behaviors like oscillations or sudden step changes, like those seen in fig. 5-7 and fig. 5-8.

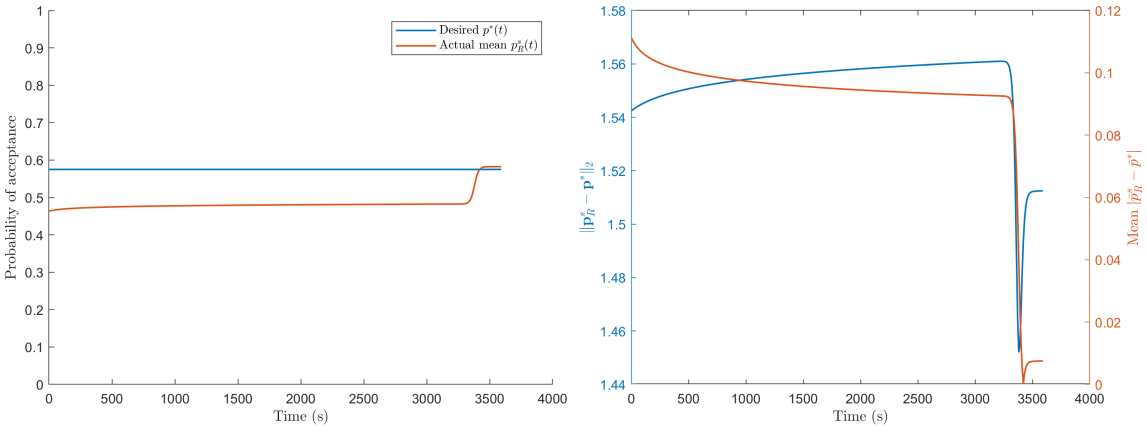

(a) Plot of desired vs actual probabilities of acceptance, averaged over the whole sample.

(b) Plot of the mean absolute error and the vector L_2 norm of error.

Figure 5-7: Direct tariff-based control strategy using gradient descent with fixed step size $\eta = 0.4$.

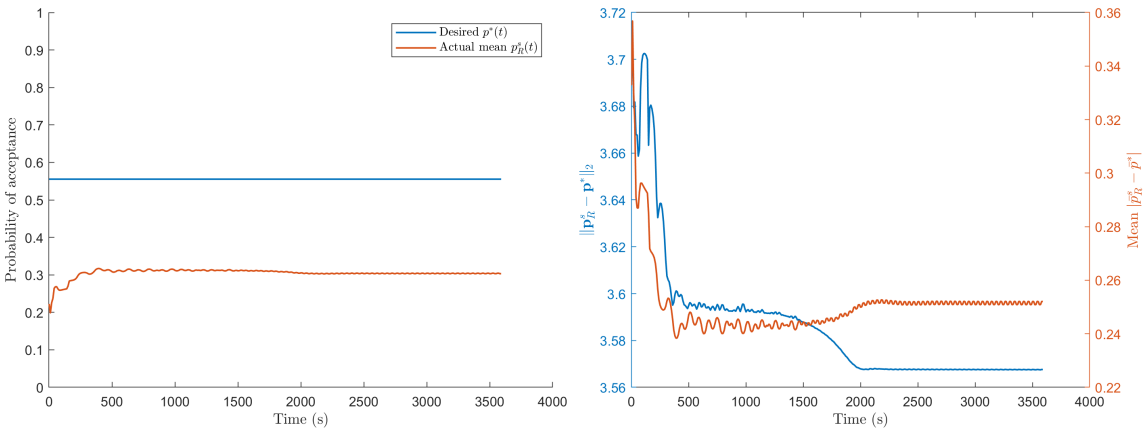

(a) Plot of desired vs actual probabilities of acceptance, averaged over the whole sample.

(b) Plot of the mean absolute error and the vector L_2 norm of error.

Figure 5-8: Direct tariff-based control strategy using Nesterov accelerated gradient descent with $\alpha = 0.9$ and $\eta = 0.1$.

Finally, fig. 5-9 compares fixed step sizes against various learning rate schedules. In this case, the decays cause the scheduled methods to actually converge slower while the fixed learning rate is still able to escape gradient plateaus. This highlights one of the main challenges with scheduling step sizes since since they need to be carefully tuned and customized for each data set and training task.

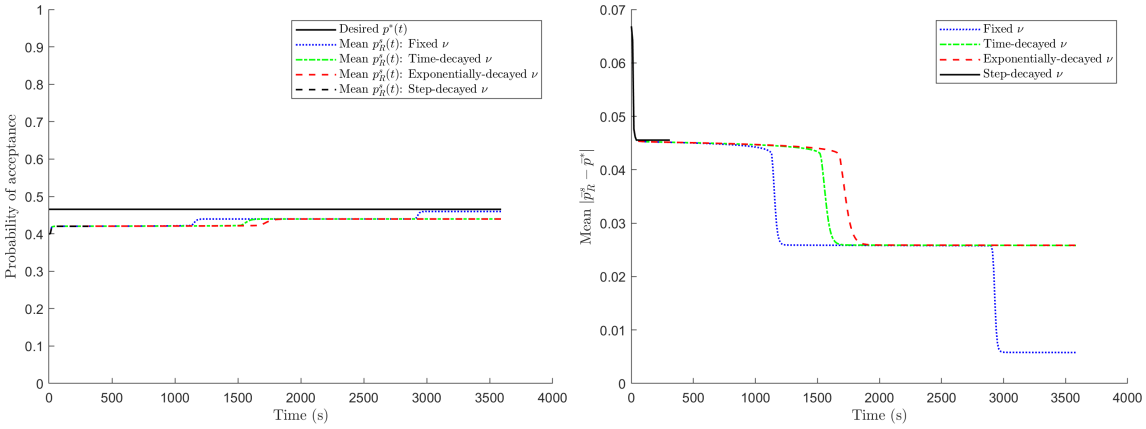

(a) Plot of desired vs actual probabilities of acceptance, averaged over the whole sample.

(b) Plot of the mean absolute error and the vector L_2 norm of error.

Figure 5-9: Comparison of fixed step size and various learning rate schedules for gradient descent with fixed step size $\eta = 0.3, d = 0.005$ and step decay after every 5 iterations.

5.5 Challenges and ideas for potential improvement

Even after implementing several different approaches to compute the gradient-based tariff and/or CPT parameter updates, and trying out different fixed, scheduled and adaptive learning rates, the performance of the controller is still quite far from optimal. In particular, the learning tends to stagnate after a few iterations even with methods using momentum and Nesterov acceleration, resulting in quite large steady state errors from the setpoint that cannot be fully eliminated. Furthermore, the behavior of the gradient descent based controller is quite erratic and varies significantly from one run to the next. Since the scenarios used in our trials were randomized, it is likely that certain trials had well-behaved or even convex plant models $f(\cdot)$ while other trials ended up having highly nonsmooth and nonconvex functions that prove challenging to minimize effectively.

Thus, the unresolved issues with our controller can largely be attributed to the possible non-convexity and non-smoothness of our CPT model for certain hyperparameter values, causing the objective function to have a large number of suboptimal local minima and possibly saddle points as well. This in turn causes the controller to often get trapped at these points and prevents it from reaching solutions that are truly globally optimal. Saddle points are especially hard to escape since the gradients around such points are nearly zero in all dimensions [88]. Another practical challenge while tuning the control loop parameters is that the simulations are quite computationally intensive with long runtimes. This is mainly because computing the functional inverse of the CPT model over all passenger-trip pairs is an expensive step.

A promising approach to take this forward and resolve some of these challenges is by further examining the learning rate used. In addition to the adaptive step sizes and schedules implemented here, there are several other even more sophisticated gradient descent optimization algorithms such as Adagrad, Adadelta, RMSProp and Adam. These are variants that adapt or customize learning rates to each parameter (or variable), rather than applying the same learning rate for all the parameter updates. In

our case, this could translate to customizing learning rates specific to each passenger-trip pair in algorithm I section 5.3.1 or having different learning rates for each for each of the five CPT parameters in algorithm II section 5.3.2. Adadelta and RMSProp are essentially extensions of Adagrad that seek to avoid its issue of rapidly diminishing learning rates [87].

There are also several other simplifications and assumptions that we made while constructing our CPT model and travel scenarios for the SMoDS. The difficulties we run into here may also be because some of these are unrealistic, infeasible or invalid in physical settings. For instance, our controller synthesis for both approaches I and II implicity assumes that the CPT model f is always an invertible bijective mapping from tariffs γ to subjective acceptance probabilities p_R^s. It is not hard to conceive of situations where this is not true - for certain combinations of R and Θ there could in fact be multiple dynamic prices that result in the same probability of acceptance.

Finally, after obtaining satisfactory performance with a fixed reference, it would also be interesting to test whether the controller can reasonably track a time-varying setpoint $p^*(t)$ over an extended time horizon, like that shown in fig. 5-10.

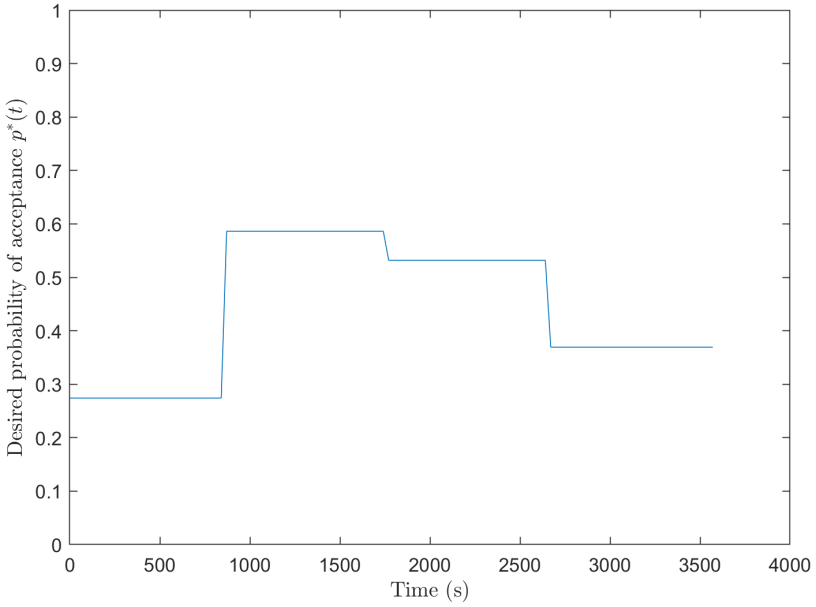

Figure 5-10: Desired probability of acceptance changing over time for the SMoDS.

Chapter 6

Conclusion and future work

In this study, we developed, implemented and analyzed methods to to fully characterize a novel, end-to-end solution for providing Shared Mobility on Demand Services (SMoDS), and subsequently used this to develop a transactive control framework to regulate ride sharing demand and traffic conditions in urban areas. In order to do so, a comprehensive survey was designed and launched at scale in the field. The survey design borrowed methods from several fields like discrete choice modelling, factorial experiment design, maximum likelihood estimation theory and cognitive psychology. The resulting data was used to first build the random utility functions for each mode using maximum simulated likelihood estimation. The resulting mode choice model was validated and shown to have high explanatory power. The parameters obtained here agreed with literature and also shed interesting insights on the value of time spent on different modes and legs, as perceived by passengers.

Secondly, the survey responses and the mode choice models were used to fit a new passenger behavioral model based on Cumulative Prospect theory, and used for real-time, dynamic pricing of such pooled ridesharing services. This was done using the method of certainty equivalents and non linear least squares. Numerous statistical and numerical techniques were applied in attempts to troubleshoot several estimation challenges and eventually obtain more accurate and meaningful distributions of parameter estimates across our sample of respondents. Some of the techniques ex-

perimented with include regularization, transfer learning, global optimization tools etc.

Thirdly, an extensive sensitivity and robustness analysis was conducted to assess the effects of modelling and parametrization errors on system performance and key objectives. This was done by applying tools from constrained optimization theory and nonlinear programming.

Fourthly, the SMoDS was simulated as a closed loop system and a feedback controller was synthesized to correct modelling errors and respond to changes in passengers' risk preferences, via price-based transactive control. This was achieved via gradient descent based control to update the dynamic SMoDS tariff, while also attempting to achieve convergence in the CPT behavioral model parameters towards their true values.

Finally, challenges arose at all stages throughout this study due to the highly nonconvex nonlinearities introduced by CPT models. We explored several different strategies to tackle and ultimately resolve some, but not all of these issues. To our knowledge, this is the first application of CPT to ridesharing and ride pooling applications. Accurate estimation and calibration of such models for the SMoDS, and the synthesis of feedback controller based on these, is a truly challenging task. We have made significant progress towards these goals in this thesis, in addition to developing standardized frameworks and methodologies for each of these steps. However, there are some unresolved issues and this remains an active, open area of research.

6.1 Future work

There are several areas that could be looked into as part of future work to take this research further. Some of the potential avenues to explore are discussed below.

6.1.1 Extensions to mode choice model

This could entail incorporating higher order terms (quadratic, cubic etc.) and interaction effects in the objective utility functions. For e.g., several models from the literature include interaction terms combining price and travel times, such as $\gamma_i t_{ride,i}$. In addition, we could also make the models more realistic by incorporating other modes like driving, mixed modes, accounting for externalities like weather etc.

6.1.2 Further Improving quality of CPT estimates

The lessons learned through this study could be used to launch revised and revamped iterations of the survey in the future, to improve the quality of survey responses and build richer data sets. Hierarchical parameter estimation approaches could also be explored which perform more systematic and sequential estimation of the CPT parameters. Some examples from the literature include hierarchical Bayesian [89] and hierarchical maximum likelihood based procedures [90]. Other studies also determine parameters in gain (β^+ and α^+) and loss (β^- and α^-) regimes separately from pure gain and pure loss scenario questions, respectively. The loss aversion parameter (λ) can then be inferred using the mixed outcome lotteries [2].

6.1.3 Extensions to sensitivity and robustness analysis

Possible extensions here include implement the global sensitivity analysis method via both analytical approaches and numerical solutions, and verify that it does indeed improve the accuracy of our results in terms of the predicted mismatch loss as well as variations in the optimal solution and objective value. In addition, the analysis could be extended to multiple passengers and rides, to see how the losses due to parametrization errors scale with the size of the population being served. Global sensitivity analysis-based approaches can also be used to compute closed loop solution updates which can then be used for dynamic optimization of systems [59].

6.1.4 Other potential applications for CPT

In theory, CPT can be applied to any other field or use case that involves decision making in the presence of risk and uncertainty. Some other fields where such models and methods could be applied include, but are not limited to:

- Dynamic rates for electric vehicle charging that accurately account for range anxiety, variable waiting times at charging stations etc.
- Uncertainties introduced by the COVID-19 pandemic in numerous areas such as vaccine uptake, use of public transit and ridesharing services, risk mitigation while social distancing etc.
- Dynamic tolls and congestion pricing.
- Real-time pricing for electricity or ancillary services in power grids with an increasing penetration of intermittent, variable renewable sources and other distributed energy resources.
- Mitigation and preventative actions and decision making for various extreme events like pandemics, natural disasters induced by climate change etc.

Bibliography

[1] Y. Guan, A. M. Annaswamy, and H. E. Tseng, "Cumulative Prospect Theory Based Dynamic Pricing for Shared Mobility on Demand Services," *arXiv:1904.04824 [cs, math]*, Apr. 2019, arXiv: 1904.04824. [Online]. Available: http://arxiv.org/abs/1904.04824

[2] M. O. Rieger, M. Wang, and T. Hens, "Estimating cumulative prospect theory parameters from an international survey," *Theory Decis*, vol. 82, no. 4, pp. 567–596, Apr. 2017. [Online]. Available: https://doi.org/10.1007/s11238-016-9582-8

[3] G. Ambrosino, N. Di Volo, G. Ferilli, and B. Finn, "Mobility services accessibility: demand responsive transport service towards the flexible mobility agency," 2004.

[4] Z. Chong, B. Qin, T. Bandyopadhyay, T. Wongpiromsarn, B. Rebsamen, P. Dai, E. Rankin, and M. H. Ang, "Autonomy for mobility on demand," in *Intelligent Autonomous Systems 12*. Springer, 2013, pp. 671–682.

[5] P. Santi, G. Resta, M. Szell, S. Sobolevsky, S. H. Strogatz, and C. Ratti, "Quantifying the benefits of vehicle pooling with shareability networks," *Proceedings of the National Academy of Sciences*, vol. 111, no. 37, pp. 13 290–13 294, 2014.

[6] P. C. Fishburn, *Nonlinear preference and utility theory*. Johns Hopkins University Press Baltimore, 1988, no. 5.

[7] A. Tversky and D. Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *J Risk Uncertainty*, vol. 5, no. 4, pp. 297–323, Oct. 1992. [Online]. Available: https://doi.org/10.1007/BF00122574

[8] D. Kahneman and A. Tversky, "Prospect Theory: An Analysis of Decision Under Risk," in *Handbook of the Fundamentals of Financial Decision Making*, ser. World Scientific Handbook in Financial Economics Series. WORLD SCIENTIFIC, Jun. 2012, vol. Volume 4, no. Volume 4, pp. 99–127. [Online]. Available: https://www.worldscientific.com/doi/abs/10.1142/9789814417358_0006

[9] Y. Guan, A. M. Annaswamy, and H. E. Tseng, "A Dynamic Routing Framework for Shared Mobility Services," *arXiv:1903.11014 [math]*, Mar. 2019, arXiv: 1903.11014. [Online]. Available: http://arxiv.org/abs/1903.11014

[10] K. Train, "Discrete Choice Methods with Simulation," Jun. 2009.

[11] J. J. Louviere, D. A. Hensher, and J. D. Swait, "Stated Choice Methods: Analysis and Applications," Sep. 2009.

[12] P. Bansal, Y. Liu, R. Daziano, and S. Samaranayake, "Can Mobility-on-Demand services do better after discerning reliability preferences of riders?" *arXiv preprint arXiv:1904.07987*, Apr. 2019. [Online]. Available: https://arxiv.org/abs/1904.07987v1

[13] K. Train, "Mixed logit with a flexible mixing distribution," *Journal of Choice Modelling*, vol. 19, pp. 40–53, Jun. 2016. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1755534516300136

[14] K. Train and G. Sonnier, "Mixed Logit with Bounded Distributions of Part-worths," p. 21, 2003.

[15] K. E. Train, "Recreation Demand Models with Taste Differences over People," *Land Economics*, vol. 74, no. 2, p. 230, May 1998. [Online]. Available: http://www.jstor.org/stable/3147053?origin=crossref

[16] D. Revelt and K. Train, "Mixed Logit with Repeated Choices: Households' Choices of Appliance Efficiency Level," *Review of Economics and Statistics*, vol. 80, no. 4, pp. 647–657, Nov. 1998. [Online]. Available: http://www.mitpressjournals.org/doi/10.1162/003465398557735

[17] W. N. Evans and W. K. Viscusi, "Estimation of State-Dependent Utility Functions Using Survey Data," *The Review of Economics and Statistics*, vol. 73, no. 1, pp. 94–104, 1991. [Online]. Available: https://www.jstor.org/stable/2109691

[18] T. Andrejszki, A. Torok, and M. Csete, "Identifying the Utility Function of Transport Services From Stated Preferences," *Transport and Telecommunication Journal*, vol. 16, no. 2, pp. 138–144, Jun. 2015. [Online]. Available: https://content.sciendo.com/view/journals/ttj/16/2/article-p138.xml

[19] M. E. Ben-Akiva and S. R. Lerman, *Discrete choice analysis: theory and application to travel demand*, ser. MIT Press series in transportation studies. Cambridge, Mass: MIT Press, 1985, no. 9.

[20] K. B. Laskey and G. W. Fischer, "Estimating Utility Functions in the Presence of Response Error," *Management Science*, vol. 33, no. 8, pp. 965–980, 1987, publisher: INFORMS. [Online]. Available: http://www.jstor.org/stable/2632172

[21] T. C. Lam and K. A. Small, "The value of time and reliability: measurement from a value pricing experiment," *Transportation Research Part E: Logistics and Transportation Review*, vol. 37, no. 2-3, pp. 231–251, Apr. 2001. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1366554500000168

[22] C. Carrion and D. Levinson, "Value of travel time reliability: A review of current evidence," *Transportation Research Part A: Policy and Practice*, vol. 46, no. 4, pp. 720–741, May 2012. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0965856412000043

[23] Z. Li, D. A. Hensher, and J. M. Rose, "Willingness to pay for travel time reliability in passenger transport: A review and some new empirical evidence," *Transportation Research Part E: Logistics and Transportation Review*, vol. 46, no. 3, pp. 384–403, May 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1366554509001550

[24] K. A. Small, "Valuation of travel time," *Economics of Transportation*, vol. 1, no. 1, pp. 2–14, Dec. 2012. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S2212012212000093

[25] Y. Liu, P. Bansal, R. Daziano, and S. Samaranayake, "A framework to integrate mode choice in the design of mobility-on-demand systems," *Transportation Research Part C: Emerging Technologies*, vol. 105, pp. 648–665, Aug. 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0968090X18313718

[26] N. Buchholz, L. Doval, J. Kastl, F. Matějka, and T. Salz, "The Value of Time: Evidence From Auctioned Cab Rides," National Bureau of Economic Research, Cambridge, MA, Tech. Rep. w27087, May 2020. [Online]. Available: http://www.nber.org/papers/w27087.pdf

[27] R. Krueger, T. H. Rashidi, and J. M. Rose, "Preferences for shared autonomous vehicles," *Transportation Research Part C: Emerging Technologies*, vol. 69, pp. 343–355, Aug. 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0968090X16300870

[28] S. Wang and J. Zhao, "Risk preference and adoption of autonomous vehicles," *Transportation Research Part A: Policy and Practice*, vol. 126, pp. 215–229, Aug. 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0965856417310315

[29] O. Al Maghraoui, R. Vosooghi, A. Mourad, J. Kamel, J. Puchinger, F. Vallet, and B. Yannou, "Shared autonomous vehicle services and user taste variations: Survey and model applications," *Transportation Research Procedia*, vol. 47, pp. 3–10, Jan. 2020.

[30] J. Hawkins and K. N. Habib, "Heterogeneity in marginal value of urban mobility: evidence from a large-scale household travel survey in the Greater Toronto and Hamilton Area," *Transportation*, Jul. 2019. [Online]. Available: https://doi.org/10.1007/s11116-019-10041-7

[31] M. Montazery and N. Wilson, "A New Approach for Learning User Preferences for a Ridesharing Application," in *Transactions on Computational Collective Intelligence XXVIII*, ser. Lecture Notes in Computer Science, N. T. Nguyen, R. Kowalczyk, J. van den Herik, A. P. Rocha, and J. Filipe, Eds. Cham: Springer International Publishing, 2018, pp. 1–24. [Online]. Available: https://doi.org/10.1007/978-3-319-78301-7_1

[32] S. Naumov and D. R. Keith, "Hailing Rides Using On-Demand Mobility Platforms: What Motivates Consumers to Choose Pooling?" *Proceedings*, vol. 2019, no. 1, p. 19670, Aug. 2019. [Online]. Available: https://journals.aom.org/doi/abs/10.5465/AMBPP.2019.19670abstract

[33] G. Ewing and E. Sarigöllü, "Assessing Consumer Preferences for Clean-Fuel Vehicles: A Discrete Choice Experiment," *Journal of Public Policy & Marketing*, vol. 19, no. 1, pp. 106–118, Apr. 2000. [Online]. Available: https://doi.org/10.1509/jppm.19.1.106.16946

[34] El Khoury John and Srour F. Jordan, "Value of Dynamic Revenue-Maximizing Congestion Pricing in a Highly Congested Corridor," *Journal of Transportation Engineering*, vol. 141, no. 12, p. 04015029, Dec. 2015. [Online]. Available: https://ascelibrary.org/doi/full/10.1061/%28ASCE%29TE.1943-5436.0000798

[35] M. O. Rieger and T. Bui, "Too Risk-Averse for Prospect Theory?" *ME*, vol. 02, no. 04, pp. 691–700, 2011. [Online]. Available: http://www.scirp.org/journal/doi.aspx?DOI=10.4236/me.2011.24077

[36] M. O. Rieger, M. Wang, and T. Hens, "Risk Preferences Around the World," *Management Science*, vol. 61, no. 3, pp. 637–648, Feb. 2014. [Online]. Available: https://pubsonline.informs.org/doi/abs/10.1287/mnsc.2013.1869

[37] C. Bernard and M. Ghossoub, "Static portfolio choice under Cumulative Prospect Theory," *Math Finan Econ*, vol. 2, no. 4, pp. 277–306, Mar. 2010. [Online]. Available: https://doi.org/10.1007/s11579-009-0021-2

[38] T. Tanaka, C. F. Camerer, and Q. Nguyen, "Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam," *American Economic Review*, vol. 100, no. 1, pp. 557–571, Mar. 2010. [Online]. Available: https://www.aeaweb.org/articles?id=10.1257/aer.100.1.557

[39] H. Xu, J. Zhou, and W. Xu, "A decision-making rule for modeling travelers' route choice behavior based on cumulative prospect theory," *Transportation Research Part C: Emerging Technologies*, vol. 19, no. 2, pp. 218–228, Apr. 2011. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0968090X10000938

[40] R.-C. Jou and K.-H. Chen, "An application of cumulative prospect theory to freeway drivers' route choice behaviours," *Transportation Research Part A: Policy and Practice*, vol. 49, pp. 123–131, Mar. 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0965856413000189

[41] A. V. Fiacco, *Introduction to sensitivity and stability analysis in nonlinear programming.* Elsevier, 1983.

[42] A. V. Fiacco and Y. Ishizuka, "Sensitivity and stability analysis for nonlinear programming," *Ann Oper Res*, vol. 27, no. 1, pp. 215–235, Dec. 1990. [Online]. Available: https://doi.org/10.1007/BF0205no5196

[43] A. Shapiro, "Sensitivity Analysis of Nonlinear Programs and Differentiability Properties of Metric Projections," *SIAM J. Control Optim.*, vol. 26, no. 3, pp. 628–645, May 1988. [Online]. Available: https://epubs.siam.org/doi/abs/10.1137/0326037

[44] ——, "Second order sensitivity analysis and asymptotic theory of parametrized nonlinear programs," *Sensitivity analysis*, Dec. 1985.

[45] A. V. Fiacco, "Sensitivity analysis for nonlinear programming using penalty methods," *Mathematical Programming*, vol. 10, no. 1, pp. 287–311, Dec. 1976. [Online]. Available: https://doi.org/10.1007/BF01580677

[46] E. Castillo, A. J. Conejo, C. Castillo, R. Mínguez, and D. Ortigosa, "Perturbation Approach to Sensitivity Analysis in Mathematical Programming," *J Optim Theory Appl*, vol. 128, no. 1, pp. 49–74, Jan. 2006. [Online]. Available: http://link.springer.com/10.1007/s10957-005-7557-y

[47] D. Ralph and S. Dempe, "Directional derivatives of the solution of a parametric nonlinear program," *Mathematical Programming*, vol. 70, no. 1-3, pp. 159–172, Oct. 1995. [Online]. Available: http://link.springer.com/10.1007/BF01585934

[48] H. M. Wagner, "Global sensitivity analysis," *Operations Research*, vol. 43, no. 6, pp. 948–969, 1995.

[49] F. S. Koppelman and C. Bhat, "A self instructing course in mode choice modeling: multinomial and nested logit models," 2006.

[50] P. Belenky, "The value of travel time savings: departmental guidance for conducting economic evaluations, revision 2," *Washington DC: United States Department of Transportation*, 2011.

[51] D. Prelec, "The Probability Weighting Function," *Econometrica*, vol. 66, no. 3, pp. 497–527, 1998. [Online]. Available: https://www.jstor.org/stable/2998573

[52] W. F. Kuhfeld, "Experimental design, efficiency, coding, and choice designs," *Marketing research methods in sas: Experimental design, choice, conjoint, and graphical techniques*, Jan. 2005.

[53] I. J. Myung, "Tutorial on maximum likelihood estimation," *Journal of Mathematical Psychology*, vol. 47, no. 1, pp. 90–100, Feb. 2003. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0022249602000287

[54] K. Train, "Halton sequences for mixed logit," Aug. 2000.

[55] M. O. Rieger and M. Wang, "Cumulative prospect theory and the St. Petersburg paradox," *Economic Theory*, vol. 28, no. 3, pp. 665–679, Aug. 2006. [Online]. Available: https://doi.org/10.1007/s00199-005-0641-6

[56] A. Reuther, J. Kepner, C. Byun, S. Samsi, W. Arcand, D. Bestor, B. Bergeron, V. Gadepally, M. Houle, M. Hubbell *et al.*, "Interactive supercomputing on 40,000 cores for machine learning and data analysis," in *2018 IEEE High Performance extreme Computing Conference (HPEC)*. IEEE, Sep. 2018, pp. 1–6.

[57] D. P. Bertsekas, "Nonlinear programming," *Journal of the Operational Research Society*, vol. 48, no. 3, 1997.

[58] C. Büskens and H. Maurer, "Sensitivity Analysis and Real-Time Optimization of Parametric Nonlinear Programming Problems," in *Online Optimization of Large Scale Systems*, M. Grötschel, S. O. Krumke, and J. Rambau, Eds. Berlin, Heidelberg: Springer, 2001, pp. 3–16. [Online]. Available: https://doi.org/10.1007/978-3-662-04331-8_1

[59] J. V. Kadam and W. Marquardt, "Sensitivity-Based Solution Updates in Closed-Loop Dynamic Optimization," *IFAC Proceedings Volumes*, vol. 37, no. 9, pp. 947–952, Jul. 2004. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1474667017319304

[60] A. R. Hole, "A comment on 'Mixed logit models: Accuracy and software choice'," May 2011. [Online]. Available: http://www.shef.ac.uk/economics/research/serps/year.html

[61] R. Gonzalez and G. Wu, "On the shape of the probability weighting function," *Cognitive psychology*, vol. 38, no. 1, pp. 129–166, 1999.

[62] G. Wu and R. Gonzalez, "Curvature of the Probability Weighting Function," *Management Science*, vol. 42, no. 12, pp. 1676–1690, 1996, publisher: INFORMS. [Online]. Available: https://www.jstor.org/stable/2634546

[63] S. Geman, E. Bienenstock, and R. Doursat, "Neural networks and the bias/variance dilemma," *Neural computation*, vol. 4, no. 1, pp. 1–58, Jan. 1992.

[64] Y. Guan, A. M. Annaswamy, and H. E. Tseng, "Towards Dynamic Pricing for Shared Mobility on Demand using Markov Decision Processes and Dynamic Programming," *arXiv:1910.01993 [math]*, Jul. 2020, arXiv: 1910.01993. [Online]. Available: http://arxiv.org/abs/1910.01993

[65] A. M. Annaswamy, Y. Guan, H. E. Tseng, H. Zhou, T. Phan, and D. Yanakiev, "Transactive Control in Smart Cities," *Proceedings of the IEEE*, vol. 106, no. 4, pp. 518–537, Apr. 2018, conference Name: Proceedings of the IEEE.

[66] T. Phan, A. M. Annaswamy, D. Yanakiev, and E. Tseng, "A model-based dynamic toll pricing strategy for controlling highway traffic," in *2016 American Control Conference (ACC)*, Jul. 2016, pp. 6245–6252, iSSN: 2378-5861.

[67] Y. Nie, X. Wang, and K. E. Cheng, "Multi-Area Self-Adaptive Pricing Control in Smart City With EV User Participation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 7, pp. 2156–2164, Jul. 2018, conference Name: IEEE Transactions on Intelligent Transportation Systems.

[68] P. Kachroo, S. Gupta, S. Agarwal, and K. Ozbay, "Optimal Control for Congestion Pricing: Theory, Simulation, and Evaluation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 5, pp. 1234–1240, May 2017, conference Name: IEEE Transactions on Intelligent Transportation Systems.

[69] S. Bouchelaghem and M. Omar, "Reliable and Secure Distributed Smart Road Pricing System for Smart Cities," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 5, pp. 1592–1603, May 2019, conference Name: IEEE Transactions on Intelligent Transportation Systems.

[70] A. S. Krishen, R. L. Raschke, and P. Kachroo, "A feedback control approach to maintain consumer information load in online shopping environments," *Information & Management*, vol. 48, no. 8, pp. 344–352, Dec. 2011. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0378720611000796

[71] B. Tian, J. Han, and K. Liu, "Closed-Loop Feedback Computation Model of Dynamical Reputation Based on the Local Trust Evaluation in Business-to-Consumer E-Commerce," *Information*, vol. 7, no. 1, p. 4, Mar. 2016, number: 1 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/2078-2489/7/1/4

[72] K. S. Narendra and A. M. Annaswamy, *Stable adaptive systems*. Courier Corporation, 2012.

[73] E. Lavretsky, "Adaptive control: Introduction, overview, and applications," in *Lecture notes from IEEE Robust and Adaptive Control Workshop*, 2008.

[74] S. Labiod and T. M. Guerra, "Adaptive fuzzy control of a class of SISO nonaffine nonlinear systems," *Fuzzy Sets and Systems*, vol. 158, no. 10, pp. 1126–1137, May 2007. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0165011406005239

[75] ——, "Direct adaptive fuzzy control for a class of MIMO nonlinear systems," *International Journal of Systems Science*, vol. 38, no. 8, pp. 665–675, Aug. 2007, publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00207720701500583. [Online]. Available: https://doi.org/10.1080/00207720701500583

[76] B. B. Alagoz, G. Kavuran, A. Ates, and C. Yeroglu, "Reference-shaping adaptive control by using gradient descent optimizers," *PloS one*, vol. 12, no. 11, p. e0188527, 2017.

[77] K. B. Ariyur and M. Krstić, *Real time optimization by extremum seeking control*. Hoboken, NJ: Wiley Interscience, 2003.

[78] Joon-Young Choi, M. Krstic, K. B. Ariyur, and J. S. Lee, "Extremum seeking control for discrete-time systems," *IEEE Transactions on Automatic Control*, vol. 47, no. 2, pp. 318–323, Feb. 2002, conference Name: IEEE Transactions on Automatic Control.

[79] A. Skafte, "An introduction to extremum-seeking control," 2017.

[80] H. K. Khalil, *Nonlinear control*. Boston: Pearson, 2015.

[81] K. Shimizu, K. Otsuka, and J. Naiborhu, "Improved direct gradient descent control of general nonlinear systems," in *1999 European Control Conference (ECC)*, Aug. 1999, pp. 2006–2011.

[82] K. Jouili and N. Benhadj Braiek, "A gradient descent control for output tracking of a class of non-minimum phase nonlinear systems," *Journal of Applied Research and Technology*, vol. 14, no. 6, pp. 383–395, Dec. 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1665642316300943

[83] K. Shimizu, S. Ito, and S. Suzuki, "Tracking control of general nonlinear systems by direct gradient descent method," *IFAC Proceedings Volumes*, vol. 31, no. 17, pp. 183–188, 1998.

[84] J. V. Burke, A. S. Lewis, and M. L. Overton, "A Robust Gradient Sampling Algorithm for Nonsmooth, Nonconvex Optimization," *SIAM J. Optim.*, vol. 15, no. 3, pp. 751–779, Jan. 2005, publisher: Society for Industrial and Applied Mathematics. [Online]. Available: http://epubs.siam.org/doi/abs/10.1137/030601296

[85] Y. Tang, E. Dall'Anese, A. Bernstein, and S. H. Low, "A Feedback-Based Regularized Primal-Dual Gradient Method for Time-Varying Nonconvex Optimization," in *2018 IEEE Conference on Decision and Control (CDC)*, Dec. 2018, pp. 3244–3250, iSSN: 2576-2370.

[86] V. Häberle, A. Hauswirth, L. Ortmann, S. Bolognani, and F. Dörfler, "Non-convex Feedback Optimization with Input and Output Constraints," *IEEE Control Syst. Lett.*, pp. 1–1, 2020, arXiv: 2004.06407. [Online]. Available: http://arxiv.org/abs/2004.06407

[87] S. Ruder, "An overview of gradient descent optimization algorithms," *arXiv:1609.04747 [cs]*, Jun. 2017, arXiv: 1609.04747. [Online]. Available: http://arxiv.org/abs/1609.04747

[88] Y. Dauphin, R. Pascanu, C. Gulcehre, K. Cho, S. Ganguli, and Y. Bengio, "Identifying and attacking the saddle point problem in high-dimensional non-convex optimization," *arXiv preprint arXiv:1406.2572*, 2014.

[89] H. Nilsson, J. Rieskamp, and E.-J. Wagenmakers, "Hierarchical Bayesian parameter estimation for cumulative prospect theory," *Journal of Mathematical Psychology*, vol. 55, no. 1, pp. 84–93, Feb. 2011. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0022249610001070

[90] R. O. Murphy and R. H. W. ten Brincke, "Hierarchical Maximum Likelihood Parameter Estimation for Cumulative Prospect Theory: Improving the Reliability of Individual Risk Parameter Estimates," *Management Science*, vol. 64, no. 1, pp. 308–326, Jan. 2017, publisher: INFORMS. [Online]. Available: https://pubsonline.informs.org/doi/abs/10.1287/mnsc.2016.2591