Essays on Inventory Management and Conjoint Analysis

Yupeng Chen

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2015

ABSTRACT

Essays on Inventory Management and Conjoint Analysis

Yupeng Chen

With recent theoretic and algorithmic advancements, modern optimization methodologies have seen a substantial expansion of modeling power, being applied to solve challenging problems in impressively diverse areas. This dissertation aims to extend the modeling frontier of optimization methodologies in two exciting fields-inventory management and conjoint analysis. Although the three essays concern distinct applications using different optimization methodologies, they share a unifying theme, which is to develop intuitive models using advanced optimization techniques to solve problems of practical relevance.

The first essay (Chapter 2) applies robust optimization to solve a single installation inventory model with non-stationary uncertain demand. A classical problem in operations research, the inventory management model could become very challenging to analyze when lost-sales dynamics, non-zero fixed ordering cost, and positive lead time are introduced. In this essay, we propose a robust cycle-based control policy based on an innovative decomposition idea to solve a family of variants of this model. The policy is simple, flexible, easily implementable and numerical experiments suggest that the policy has very promising empirical performance. The policy can be used both when the excess demand is backlogged as well as when it is lost; with non-zero fixed ordering cost, and also when lead time is non-zero. The policy decisions are computed by solving a collection of linear programs even when there is a positive fixed ordering

cost. The policy also extends in a very simple manner to the joint pricing and inventory control problem.

The second essay (Chapter 3) applies sparse machine learning to model multimodal continuous heterogeneity in conjoint analysis. Consumers' heterogeneous preferences can often be represented using a multimodal continuous heterogeneity (MCH) distribution. One interpretation of MCH is that the consumer population consists of a few distinct segments, each of which contains a heterogeneous sub-population. Modeling of MCH raises considerable challenges as both across- and within-segment heterogeneity need to be accounted for. In this essay, we propose an innovative sparse learning approach for modeling MCH and apply it to conjoint analysis where adequate modeling of consumer heterogeneity is critical. The sparse learning approach models MCH via a two-stage divide-and-conquer framework, in which we first decompose the consumer population by recovering a set of candidate segmentations using structured sparsity modeling, and then use each candidate segmentation to develop a set of individual-level representations of MCH. We select the optimal individual-level representation of MCH and the corresponding optimal candidate segmentation using cross-validation. Two notable features of our approach are that it accommodates both across- and within-segment heterogeneity and endogenously imposes an adequate amount of shrinkage to recover the individual-level partworths. We empirically validate the performance of the sparse learning approach using extensive simulation experiments and two empirical conjoint data sets.

The third essay (Chapter 4) applies dynamic discrete choice models to investigate the impact of return policies on consumers' product purchase and return behavior. Return policies have been ubiquitous in the marketplace, allowing consumers to use and evaluate a product before

fully committing to purchase. Despite the clear practical relevance of return policies, however, few studies have provided empirical assessments of how consumers' purchase and return decisions respond to the return policies facing them. In this essay, we propose to model consumers' purchase and return decisions using a dynamic discrete choice model with forward-looking and Bayesian learning. More specifically, we postulate that consumers' purchase and return decisions are optimal solutions for some underlying dynamic expected-utility maximization problem in which consumers learn their true evaluations of products via usage in a Bayesian manner and make purchase and return decisions to maximize their expected present value of utility, and return policies impact consumers' purchase and return decisions by entering the dynamic expected-utility maximization problem as constraints. Our proposed model provides a behaviorally plausible approach to examine the impact of return policies on consumers' purchase and return behavior.

Contents

List of Figures

List of Tables

Acknowledgements

I would like to start by thanking my advisor, Professor Garud Iyengar, for his constant guidance and help. Professor Garud Iyengar has been an extremely supportive mentor, always encouraging me to explore research questions that I am passionate about and sharing his insights on solving challenging problems. I am extremely fortunate to have the opportunity to work with him on a series of exciting research projects, which has been an invaluable learning process for me. Professor Garud Iyengar also provides a lot of support for my personal and professional development, for which I always feel grateful.

I would like to thank Professor Raghu Iyengar, who, together with Professor Garud Iyengar, motivates and guides my research on the intersection between operations research and quantitative marketing. Working with Professor Raghu Iyengar has been an extremely rewarding experience, as he not only shares his thoughtful perspectives on how to ask and answer research questions, but also influences me with his constant passion and curiosity towards research.

I am deeply indebted to Professors Martin Haugh, Soulaymane Kachani, and Olivier Toubia, for serving on my dissertation committee and providing insightful feedback.

I am extremely grateful to have met many good friends during the past years and their warm friendships is one huge reason why I have enjoyed so much of my time at Columbia. My most sincere appreciation goes for them.

Finally, I owe my deepest gratitude to my parents and my wife, Zhe, without whose love and support this dissertation would not have been possible.

To My Family

Chapter 1

Overview

This dissertation discusses three applications of modern optimization methodologies in inventory management and conjoint analysis. Inventory management is a field of operations research concerning optimal dynamic control of stochastic inventory systems, and conjoint analysis is a field of quantitative marketing developing accurate measurement and modeling of consumers' preferences. Through this dissertation, we aim to show that both fields are constantly presenting exciting challenges despite their long and rich intellectual histories, and, more importantly, we aim to demonstrate the powerful modeling capacity of recently developed optimization methodologies in these fields. In particular, for each of the three applications we consider in this dissertation, we propose an innovative solution method based on optimization techniques which delivers promising empirical performance.

Single installation inventory management model has been extensively researched over the past 50 years. However, in the presence of lost-sales dynamics, non-zero fixed ordering cost, and positive lead time, finding a computationally tractable control policy with near-optimal

performance remains a challenging task. The standard approach to solve for the optimal control policy is dynamic programming (DP). DP is effective in characterizing the structure of optimal policies for simple inventory models, yet its computational complexity easily explodes as the state space expands, which severely limits the application of DP in practice. Recent research applies the robust optimization (RO) methodology to develop computationally tractable solutions for inventory models, in which control policies are parameterized in realized demands and good control policies are identified by solving mathematical optimization problems. One key limitation of the current RO-based approaches is that they rely critically on linear system dynamics, which prohibits their application to lost-sales inventory models. In addition, the RO-based approaches model the non-zero fixed ordering cost using integer decision variables, rendering the resulting optimization problems difficult to solve.

In Chapter 2, we propose a new robust cycle-based control policy for solving the single installation, finite horizon inventory management model with non-stationary uncertain demand. Inspired in part by the classic Economic-Order-Quantity model, our proposed control policy dynamically splits the planning horizon into ordering cycles in an online fashion and chooses an appropriate order quantity and cycle length at the beginning of each ordering cycle. We compute good order quantity and cycle length decisions using the RO-framework to account for demand uncertainty. We show that by focusing on each ordering cycle as well as exploiting the special structure of inventory management models, the decisions in our policy are efficiently computable under various model assumptions, including lost-sales dynamics, non-zero fixed ordering cost, and positive lead time. More specifically, we show that the decisions can be computed by solving a collection of linear programs of modest size, and therefore our policy is

easily implementable using any linear program solvers. Using extensive numerical experiments, we demonstrate that our proposed control policy has very promising performance compared to strong benchmarks across various scenarios despite its simplicity.

Consumers have heterogeneous preferences and an adequate modeling of consumer heterogeneity is a prerequisite for accurate conjoint estimation. In practice, consumer preferences can often be represented using a multimodal continuous heterogeneity (MCH) distribution, in which the consumer population can be interpreted as consisting of a few distinct segments, each of which contains a heterogeneous sub-population. Modeling of MCH raises considerable challenges as both across- and within-segment heterogeneity need to be accounted for. Standard approaches to modeling consumer heterogeneity, including the finite mixture models and the unimodal hierarchical Bayes models, make simplifying specifications regarding the underlying heterogeneity distribution. Consequently, they are not capable of fully capturing the variation in individual-level partworths and may lead to misleading inferences when MCH is present. Recent work has focused on enhancing the modeling power of hierarchical Bayes models by incorporating more sophisticated heterogeneity distributions, e.g., the normal component mixture models and the Dirichlet process models. These advanced models have demonstrated great capacity to recover individual-level partworths in the presence of MCH; however, their performance in modeling MCH suffers from the fact that the amount of Bayesian shrinkage imposed by them is influenced by exogenously chosen parameters of the second-stage priors and hence may be suboptimal, and that they face inferential difficulties when conducting a segment-level analysis.

In Chapter 3, we propose and test a new sparse learning approach to modeling MCH in

the context of both metric and choice-based conjoint analysis. Our sparse learning approach models MCH via a two-stage divide-and-conquer framework, in which we first decompose the consumer population by recovering a set of candidate segmentations using structured sparsity modeling, and then use each candidate segmentation to develop a set of individual-level representations of MCH. We select the optimal individual-level representation of MCH and the corresponding optimal candidate segmentation using cross-validation. Two notable features of our approach are that it accommodates both across- and within-segment heterogeneity and endogenously imposes an adequate amount of shrinkage to recover the individual-level partworths. We empirically validate the performance of the sparse learning approach using extensive simulation experiments and two empirical conjoint data sets.

Return policies have been ubiquitous in today's marketplace, yet research empirically investigating the impact of return policies on consumers' purchase and return decisions based on individual-level data has been limited. The extant research assumes that consumers' uncertainty on the product's quality or fit is resolved immediately following the receipt of the product, allowing the researchers to model return policies solely as options. This assumption is not appropriate in many product categories and ignores potentially important aspects of return policies including the length of a return policy.

In Chapter 4, we propose a dynamic discrete choice model with forward-looking and Bayesian learning to model consumers' purchase and return behavior and to investigate the impact of return policies. In our proposed model, consumers' purchase and return decisions are postulated as optimal solutions for some underlying dynamic expected-utility maximization problem in which consumers learn their true evaluations of products via usage in a Bayesian manner and

make purchase and return decisions to maximize their expected present value of utility; on the other hand, return policies enter the dynamic expected-utility maximization problem as constraints and hence impact consumers' purchase and return decisions. Our proposed model provides a behaviorally plausible approach to capture the impact of various aspects of return policies, including the restocking fee and the length of a return policy.

Chapter 2

A New Robust Cycle-Based Inventory Control Policy

2.1 Introduction

Single installation inventory management is a classical research topic which has received persistent attention from various research communities. The standard approach to compute the optimal inventory control policy uses dynamic programming (DP). For simple inventory control problems, DP-based methods are able to *characterize* the structure of the optimal policies, e.g., [72] showed that the optimal policy for an inventory control problem with fixed ordering cost is a so-called (s,S)-policy. Many other researchers have also used DP to derive structural results of the optimal policies under various model assumptions [53, 87]. Recently, [77] and [73] have extended the structural results to inventory control problems with a more general Markovian demand. In practice, however, computing the optimal parameters is often challenging. Since

DP-based methods compute a value function for each possible state in each period, the computational complexity of DP is prohibitive when state space is large. This *curse of dimensionality* has severely limited the application of DP in practice.

The DP methodology demands that the demand distribution be *perfectly* known. In practice, the available historical data on demand distribution is limited, and moreover, the demand distribution is very often non-stationary. Thus, the perfect information assumption almost never holds in practice. The performance of DP-based policies is often very sensitive to errors in the demand distribution [46, 64]. See [24] for an empirical investigation. Consequently, practitioners prefer simple policies that can be computed using limited information about the demand distribution [94]. [71] formulated a min-max single-period inventory problem where only the mean and the variance of the demand are known. Moon and Gallego extended this approach to single-period news-vendor problems [39] and to single-stage periodic review inventory models with a fixed reorder quantity [61]. [40] later extended the Scarf model to a finite horizon model under the assumption that the demand is a discrete random variable. However, the complexity of computing the optimal policy is exponential in the time horizon.

Recent research efforts have focused on developing computationally tractable robust solutions for the dynamic program. One such approach is to use the robust optimization (RO) methodology to solve inventory control problems. RO is a relatively new methodology for tractably computing good solutions for optimization problems with uncertain parameters in which the parameters are assumed to belong to a known bounded *uncertainty* set and decisions are chosen assuming worst case behavior of these uncertain parameters [13, 15–17, 22, 23]. [24] introduced a robust inventory control model where the uncertain parameters are the realized demand, and showed

that for a family of polyhedral uncertainty sets introduced in [23] the open-loop control policy can be computed by solving a linear program (LP) (resp. mixed integer program (MIP)) when fixed ordering cost is zero (resp. positive). [25] showed that the optimal base-stock policy for a closely related model can be computed efficiently using a Benders' decomposition approach. [14] showed that the optimal affine policy for a robust inventory model can be efficiently computed for a large class of uncertainty sets. Recently, [21] showed that affine policies are optimal for hypercube-like uncertainty sets. [20] also showed how to compute optimal polynomial policies.

The RO-based policies in the literature treat the inventory control problem as a special case of the more general problem of controlling a linear system. These policies typically do not exploit the special structure of inventory control problems. Since RO-based policies work only for linear dynamics, these policies are only applicable when the excess demand is fully backlogged, and not when the excess demand is lost. Moreover, the ordering decisions in the RO-based policies for inventory control problems with a positive fixed ordering cost have to be computed by solving an MIP. Since these policies are typically implemented in a rolling horizon fashion, one needs to solve several, typically large, MIPs. MIP solvers are neither as ubiquitous nor as stable as LP solvers; consequently, MIP-based control polices are less likely to be widely adopted.

In this paper, we propose a new RO-based control policy that is inspired in part by the Economic-Order-Quantity (EOQ) model [94] and the periodic review (R,T) or (S,T)-type policies [9, 66, 74, 94]. The EOQ model is an infinite horizon, continuous review inventory model with constant deterministic demand. In order to balance the inventory ordering cost,

which includes a fixed cost, and the inventory holding cost and backlogging cost, EOQ optimal policy splits the planning horizon into fixed length *cycles* and chooses the cycle length and the order quantity jointly to minimize the average cost over a cycle. The periodic review (R,T)-type policies generalize the EOQ model to the case where the demand is a stationary stochastic process, specifically Poisson process. Under an (R,T) policy, the inventory position is reviewed every T periods, i.e., the planning horizon is split into length T cycles, and, if required, the policy raises the inventory position up to R. As in the EOQ model, in the (R,T) policy the cycle length T is an explicit decision variable. Although the (R,T) policy is generally sub-optimal, it is very attractive in practice because it specifies the reorder epochs upfront and is, thus, easy to implement. In particular, the (R,T) policy can substantially simplify the coordination within a more complex inventory system [66, 94]. [66] showed that the long-run average cost function for the (R,T) policy is jointly convex in R and T, and provided performance guarantee for the (R,T) policy. See [9, 74] for a detailed study of the empirical performance of the (R,T) policy and further applications of the policy.

The EOQ model and the (R,T) policy are simple, yet powerful, control policies for a single installation inventory system. However, these policies have so far been restricted to models where the goal is to minimize the long-run average cost and the demand process is either stationary, deterministic, i.e., constant, or stationary, stochastic. Our RO-based approach extends the "cycle-based" perspective to a single installation, finite horizon inventory model with non-stationary uncertain demand. Specifically, the contributions of this paper are as follows.

1. *New robust cycle-based policy:* We propose a new robust cycle-based control policy for a single installation, finite horizon inventory model with non-stationary uncertain demand. Our policy dynamically breaks the planning horizon into *ordering cycles* and chooses an appropriate order quantity and cycle length at the beginning of each ordering cycle assuming the worst case behavior of the uncertain demand. Thus, our policy can be viewed as an *adaptive* version of the (R,T) policy for an inventory model with more general demand process. The policy has a simple intuitive economic interpretation and is very easy to implement. Consequently, we believe that our proposed policy has a higher likelihood of being adopted in practice.

2. *Computationally efficient solution for both backlogging and lost-sales models:* We show that our policy provides a unified and efficient approach for solving inventory control problems under various model assumptions. In particular, the decisions in our proposed robust cycle-based policy can be efficiently computed both when the excess demand is backlogged or lost. To the best of our knowledge, this is the first RO-based inventory control policy for the lost-sales model. The decisions are also efficiently computable when there is a non-zero fixed ordering cost and also when lead time is non-zero. More specifically, we show that for all these cases the decisions in our policy can be computed by solving a collection of LPs of modest size even when the fixed ordering cost is non-zero, i.e., we do not have to solve any MIPs. Consequently, the computation of our policy is extremely efficient (see Section 2.6). On the other hand, since LP solvers are significantly

more easily accessible and stable when compared to MIP solvers, our policy should be an attractive candidate for practical implementation.

Our policy also extends to the joint pricing and inventory control problem. In this case we need to solve a small collection of convex problems to compute the optimal decisions. These problems reduce to QPs in special cases.

3. *Near-optimal empirical performance:* The robust cycle-based policy has very promising empirical performance. We conducted an extensive numerical experiment to test the performance of our proposed simple policy. We compared the performance of our policy with the rolling-horizon version of the robust policy in [24] when the excess demand is backlogged. For the lost-sales model, we compared the performance of our policy with a DP-based policy, the best base-stock policy in hindsight, and the set of heuristic policies considered in [95]. In the joint pricing and inventory control model, we compared the performance of the simple extension of our policy with a DP-based policy. For the backlogging model, the lost-sales model without lead time and the joint pricing and inventory control model, the cost of the robust cycle-based policy is typically within $\pm 3\%$ of the cost of benchmark methods. For the lost-sales model with positive lead times, our policy incurs no more than 7% of the cost of the optimal policy in most scenarios and is very competitive with other heuristics. In most cases, the computational time required to compute the decisions in our policy is at least 5–10 times smaller than that required for the competing methods.

The rest of the paper is organized as follows. In Section 2.2, we introduce our basic

inventory model with zero lead time. In Section 2.3, we propose our robust cycle-based inventory control policy, and generalize it to the case with non-zero lead time in Section 2.4. We briefly discuss the extension to a joint pricing and inventory control model in Section 2.5. In Section 2.6, we report the preliminary numerical results. In Section 2.7 we include some concluding remarks and avenues for further extensions.

2.2 Inventory model with zero lead time

We consider a single product single installation stochastic inventory problem over a finite discrete horizon of T periods. Let x_k denote the inventory available at the beginning of period k, u_k denote the quantity ordered at the beginning of period k, and d_k denote the realized demand in period k, for $k = 1, 2, \ldots, T$. We assume that the lead time is zero, i.e., the quantity ordered at the beginning of any period is delivered immediately. In particular, we assume that the quantity u_k is delivered before the demand d_k is realized. We discuss constant positive lead time in Section 2.4.

Dynamics We consider the following two most prevalent system dynamics [19, 94]:

1. Backlogging: The demand in excess of the inventory in any period is backlogged and carried over into next period, i.e., $x_{k+1} = x_k + u_k - d_k$, for all $1 \leq k \leq T-1$. Note that x_k denotes the net inventory and can take negative values.

2. Lost-sales: The demand in excess of the inventory in any period is permanently lost, i.e.,

$x_{k+1} = \max\{x_k + u_k - d_k, 0\}$, for all $1 \leq k \leq T-1$. Here, x_k denotes the actual inventory at hand and must be nonnegative.

We assume that the installation has a positive capacity M (possibly $+\infty$). Since u_k arrives before the demand, d_k, is realized, we require that $x_k + u_k \leq M$, for all $1 \leq k \leq T$.

Cost function The ordering cost in period k is given by $K\mathbf{1}(u_k > 0) + c \cdot u_k$, where K denotes the fixed ordering cost, c denotes the per unit variable ordering cost, and $\mathbf{1}(\cdot)$ denotes an indicator function that takes value 1 when argument is true, and zero otherwise. The inventory holding cost at the end of period k is given by $h \cdot \max\{x_k + u_k - d_k, 0\}$, where h denotes the per unit holding cost, and the shortage cost at the end of period k is given by $b \cdot \max\{-(x_k + u_k - d_k), 0\}$, where b denotes the per unit shortage cost; i.e., the total inventory holding and shortage cost in period k is given by $\max\left\{h \cdot (x_k + u_k - d_k), -b \cdot (x_k + u_k - d_k)\right\}$. Thus, the total cost incurred in period k is given by $C_k = K\mathbf{1}(u_k > 0) + c \cdot u_k + \max\left\{h \cdot (x_k + u_k - d_k), -b \cdot (x_k + u_k - d_k)\right\}$.

Demand distribution We assume that the demand process $\mathfrak{D} = (D_1, D_2, \ldots, D_T)$ is a stochastic process that satisfies the following assumptions.

Assumption 2.2.1. *The random variables $\{D_k\}_{k=1}^T$ are mutually independent, but not necessarily identically distributed.*

Assumption 2.2.2. *Only the first two moments of the distribution are known, i.e., only $E(D_k) = \bar{d}_k$ and $var(D_k) = \sigma_k^2$ are known for all k; all other moments and the distribution are unknown.*

Assumption 2.2.1 is common in inventory literature. It allows demand sequence to be non-stationary, e.g., seasonally cyclic. Assumption 2.2.2 is a departure from the classical inventory literature where one assumes perfect information about the demand distributions. The perfect information assumption leads to a tractable inventory control problem; however, it rarely holds in practice. In most practical settings, the inventory manager may have sufficient data to accurately estimate the first two moments of the demand distribution; however, higher moments or the shape of the distribution cannot be estimated accurately [40].

We denote the set of all distributions for the demand process $\mathfrak{D}$ which are consistent with Assumption 2.2.1-2.2.2 by the set $\mathcal{F}$. We denote a realization of $\mathfrak{D}$ by the lower case vector as $\mathbf{d} = (d_1, d_2, \ldots, d_T)$. For $1 \leq k \leq T$, $\mathbf{d}[k]$ denotes the sub-sequence $(d_k, d_{k+1}, \ldots, d_T)$, and for $1 \leq j \leq T-k+1$, $\mathbf{d}[k, j]$ denotes the sub-sequence $(d_k, d_{k+1}, \ldots, d_{k+j-1})$.

Optimality criterion An inventory control policy $\pi = \{\pi_k : k \in \{1, 2, \ldots, T\}\}$ consists of ordering functions π_k that prescribe the order quantity u_k for period k. We will call a policy π admissible if π_k is non-anticipatory, i.e., the function π_k is measurable with the respect to the filtration generated by the demand up to (and including) period $k-1$. Let Π denote the set of all admissible policies. The cost $J(\pi)$ of an admissible policy π is defined as

$$J(\pi) = \sup_{\mathfrak{D} \in \mathcal{F}} \mathbf{E}_{\mathfrak{D}} \Big[\sum_{k=1}^{T} \Big(K\mathbf{1}(\pi_k > 0) + c \cdot \pi_k + \max \{ h \cdot (x_k + \pi_k - d_k), -b \cdot (x_k + \pi_k - d_k) \} \Big) \Big], \tag{2.2.1}$$

where the expectation is taken with respect to the distribution of $\mathbf{d}$. We want to characterize $J^* = \inf_{\pi \in \Pi} J(\pi)$, and identify an optimal policy π^* that achieves J^*.

The optimality criterion (2.2.1) and the associated optimization problem J^* is not new. [71] considers a single period version of J^*. [40] consider J^* for the setting with discrete demand distribution and backlogging of excess demand. They show that the optimal policy is of the (s,S)-type, and provide a dynamic programming (DP) recursion to compute the optimal parameters; the solution algorithm is, however, exponential in T.

Robust inventory control Robust inventory control was introduced in [14,24,25]. We briefly review the RO-based approach proposed in [24]. Instead of using (2.2.1), Bertsimas and Thiele define the *robust* cost $J_r(\pi)$ for policy π as

$$J_r(\pi) = \sup_{\mathbf{d}\in\mathcal{D}} \Big[\sum_{k=1}^{T} \Big(K\mathbf{1}(\pi_k > 0) + c\cdot\pi_k + \max\big\{ h\cdot(x_k+\pi_k-d_k), -b\cdot(x_k+\pi_k-d_k) \big\} \Big) \Big], \quad (2.2.2)$$

where $\mathcal{D}$ is an *uncertainty set* of demand realizations, i.e., in the robust approach the uncertainty set $\mathcal{F}$ of demand *distributions* is replaced by an uncertainty set $\mathcal{D}$ of demand *realizations*. The uncertainty set $\mathcal{D}$ is constructed using the available moment information and the optimal control policy is computed by solving min-max problem $J_r^* = \inf_{\pi\in\Pi} J_r(\pi)$. The min-max problem J_r^* is still a hard problem. By restricting attention to open-loop control policies, Bertsimas and Thiele are able to reformulate an appropriate relaxation of J_r^* as a mathematical program and solve it efficiently. They show that the empirical performance of the robust policy is close to optimal for the original min-max optimality criterion J^*.

2.3 A robust cycle-based inventory control policy

In this section, we propose a simple robust cycle-based inventory control policy that provides a good solution for J^*. Intuitively, we expect our policy to show good *ex post* average cost across different demand distributions. We first motivate the main conceptual insight underlying our policy; next, we show how to compute the policy decisions in detail. We also show that the policy presented here can be applied to both backlogging dynamics and lost-sales dynamics.

Cycle-based policies in inventory management Consider an infinite horizon inventory control problem where the goal is to minimize the long-run average cost for flat deterministic demand, i.e., $D_k = d$ for all k. A good inventory control policy has to "balance" the ordering cost with other costs. On one extreme is the policy that orders the quantity d in each period; thereby, incurring no inventory holding and shortage cost but a large fixed ordering cost. A policy that places larger but less frequent orders incurs a smaller fixed ordering cost; but the policy either incurs inventory holding cost, or shortage cost, or both. The optimal policy breaks the horizon into fixed length *cycles* where an order is placed in the first period of each cycle. For any given cycle, the cycle length and the order quantity is chosen to minimize the *average* cost over the cycle. This type of analysis underlies the well known Economic-Order-Quantity (EOQ) model [94] and extends to the (R,T)-type policies with stationary stochastic demand processes [9, 66, 74, 94].

Drawing insights from the EOQ model and the (R,T)-type policies, we propose a robust cycle-based policy that is tailored for the finite horizon inventory model with non-stationary uncertain demand. Without loss of generality, assume the initial inventory $x_1 = 0$. At the

beginning of period $\tau_1 = 1$, we make the following two decisions: order quantity u_{τ_1} and the *cycle length* ξ_1. The order quantity u_{τ_1} and the cycle length ξ_1 is chosen to minimize an appropriately defined worst-case average cost over the *ordering cycle* $(\tau_1, \ldots, \tau_1 + \xi_1 - 1)$. We do not order until period $\tau_2 = \tau_1 + \xi_1$. At the beginning of period τ_2, we choose order quantity u_{τ_2} and the next cycle length ξ_2 that minimizes the worst-case average cost over the next ordering cycle $(\tau_2, \ldots, \tau_2 + \xi_2 - 1)$. We proceed in this manner until the end of the planning horizon. In the rest of this section, we precisely define the worst-case average cost that defines our policy and show that the decisions (u_τ, ξ) can be computed efficiently in an *online* fashion.

Note that our proposed policy is different from the (r,q)-type policies [94]. In the (r,q) policy the decision maker chooses the reorder point r and order quantity q, and a new cycle is triggered when the inventory position hit the reorder point; thus, the cycle length is a function of r, q and the demand dynamics. Consequently, the cycle length of an (r,q) policy cannot be easily computed when the demand distribution is uncertain. In our cycle-based policy, the cycle length is an explicit decision variable, and once the cycle length is determined at the beginning of cycle, it is held constant. This alternative formulation leads to tractable approximation for both backlogging and lost-sales dynamics, even in the presence of positive lead time and uncertain demand. Moreover, the (r,q) policy was proposed for a continuous review inventory model with stationary demand, and the decision maker computes the optimal *stationary* pair (r^*, q^*) *offline*. In contrast, our policy computes (u_τ, ξ) in an *online* fashion and the (u_τ, ξ) are adapted to the sample path of demands, and are generally different across cycles. Thus, our cycle-based policy is capable of handling non-stationary demand.

Convexity of inventory holding and shortage costs Consider an ordering cycle $(\tau, \tau+1, \ldots, \tau+\xi-1)$ that starts from period τ and has cycle length ξ. Let x denote available inventory at the beginning of the ordering cycle, u denote the order placed at the beginning of the cycle, and let $H(\tau, x, \xi, u, \mathbf{d}[\tau, \xi])$ denote the inventory holding and shortage cost incurred over the ordering cycle as a function of τ, x, ξ, u and the realized demand sequence $\mathbf{d}[\tau, \xi]$.

Lemma 2.3.1. *For fixed* (τ, x, ξ), $H(\tau, x, \xi, u, \mathbf{d}[\tau, \xi])$ *is a piecewise linear convex function of* $(u, \mathbf{d}[\tau, \xi])$ *with* $\xi + 1$ *linear pieces on* $\mathcal{B} \triangleq \Big\{(u, \mathbf{d}[\tau, \xi]) \mid u \in \mathbb{R},\ \mathbf{d}[\tau, \xi] \in \mathbb{R}_+^{\xi}\Big\}$.

Proof. Proof. First consider backlogging dynamics. Since $u_t = 0$, for all $\tau + 1 \leq t \leq \tau + \xi - 1$, $x_t = x + u - \sum_{\ell=\tau}^{t-1} d_\ell$, for $\tau + 1 \leq t \leq \tau + \xi - 1$; thus, the inventory holding and shortage cost incurred over the cycle is given by

$$H\big(\tau, x, \xi, u, \mathbf{d}[\tau, \xi]\big) = \sum_{t=\tau}^{\tau+\xi-1} \max\Big\{h\big(x + u - \sum_{\ell=\tau}^{t} d_\ell\big), -b\big(x + u - \sum_{\ell=\tau}^{t} d_\ell\big)\Big\}. \tag{2.3.1}$$

Partition $\mathcal{B} = \cup_{r=0}^{\xi} \mathcal{B}^r$, where

$$\mathcal{B}^r = \begin{cases} \Big\{(u, \mathbf{d}[\tau, \xi]) \in \mathcal{B} \mid u \in (-\infty, d_\tau - x]\Big\}, & r = 0, \\ \Big\{(u, \mathbf{d}[\tau, \xi]) \in \mathcal{B} \mid u \in \big[\sum_{\ell=\tau}^{\tau+r-1} d_\ell - x, \sum_{\ell=\tau}^{\tau+r} d_\ell - x\big]\Big\}, & 1 \leq r \leq \xi - 1, \\ \Big\{(u, \mathbf{d}[\tau, \xi]) \in \mathcal{B} \mid u \in \big[\sum_{\ell=\tau}^{\tau+\xi-1} d_\ell - x, +\infty\big)\Big\}, & r = \xi. \end{cases}$$

From (2.3.1) it follows that

- $(u, \mathbf{d}[\tau, \xi]) \in \mathcal{B}^0$: $H = f^0\big(u, \mathbf{d}[\tau, \xi]\big) \triangleq -\sum_{t=\tau}^{\tau+\xi-1} b(x + u - \sum_{\ell=\tau}^{t} d_\ell)$.

- $(u,\mathbf{d}[\tau,\xi]) \in \mathcal{B}^r$, for $1 \le r \le \xi - 1$:

  $H = f^r\big(u,\mathbf{d}[\tau,\xi]\big) \triangleq \sum_{t=\tau}^{\tau+r-1} h(x+u-\sum_{\ell=\tau}^{t} d_\ell) - \sum_{t=\tau+r}^{\tau+\xi-1} b(x+u-\sum_{\ell=\tau}^{t} d_\ell)$.

- $(u,\mathbf{d}[\tau,\xi]) \in \mathcal{B}^\xi$: $H = f^\xi\big(u,\mathbf{d}[\tau,\xi]\big) \triangleq \sum_{t=\tau}^{\tau+\xi-1} h(x+u-\sum_{\ell=\tau}^{t} d_\ell)$.

It's straightforward to verify that $H = \max\Big\{f^0\big(u,\mathbf{d}[\tau,\xi]\big), f^1\big(u,\mathbf{d}[\tau,\xi]\big), \ldots, f^\xi\big(u,\mathbf{d}[\tau,\xi]\big)\Big\}$ on $\mathcal{B}$. Note that $\{f^r\}_{r=0}^{\xi}$ are linear functions of $\big(u,\mathbf{d}[\tau,\xi]\big)$, the conclusion follows immediately.

Next, we consider the case with lost-sales dynamics. Since $u_t = 0$, for all $\tau+1 \le t \le \tau+\xi-1$, $x_t = \max\{x+u-\sum_{\ell=\tau}^{t-1} d_\ell, 0\}$, for $\tau+1 \le t \le \tau+\xi-1$. Therefore,

$$\begin{aligned} H(\tau,x,\xi,u,\mathbf{d}[\tau,\xi]) &= \max\Big\{h(x+u-d_\tau), -b(x+u-d_\tau)\Big\} \\ &+ \sum_{t=\tau+1}^{\tau+\xi-1} \max\left\{h\Big(\max\{x+u-\sum_{\ell=\tau}^{t-1} d_\ell, 0\} - d_t\Big), -b\Big(\max\{x+u-\sum_{\ell=\tau}^{t-1} d_\ell, 0\} - d_t\Big)\right\} \end{aligned} \quad (2.3.2)$$

From (2.3.2), it follows that

- $(u,\mathbf{d}[\tau,\xi]) \in \mathcal{B}^0$: $H = g^0\big(u,\mathbf{d}[\tau,\xi]\big) \triangleq -b(x+u-\sum_{t=\tau}^{\tau+\xi-1} d_t)$.

- $(u,\mathbf{d}[\tau,\xi]) \in \mathcal{B}^r$, for $1 \le r \le \xi - 1$:

  $H = g^r\big(u,\mathbf{d}[\tau,\xi]\big) \triangleq \sum_{t=\tau}^{\tau+r-1} h\big(x+u-\sum_{\ell=\tau}^{t} d_\ell\big) - b\big(x+u-\sum_{t=\tau}^{\tau+\xi-1} d_t\big)$.

- $(u,\mathbf{d}[\tau,\xi]) \in \mathcal{B}^\xi$: $H = g^\xi\big(u,\mathbf{d}[\tau,\xi]\big) \triangleq \sum_{t=\tau}^{\tau+\xi-1} h(x+u-\sum_{\ell=\tau}^{t} d_\ell)$.

It's straightforward to verify that $H = \max\Big\{g^0\big(u,\mathbf{d}[\tau,\xi]\big), g^1\big(u,\mathbf{d}[\tau,\xi]\big), \ldots, g^\xi\big(u,\mathbf{d}[\tau,\xi]\big)\Big\}$ on $\mathcal{B}$. Note $\{g^r\}_{r=0}^{\xi}$ are linear functions of $\big(u,\mathbf{d}[\tau,\xi]\big)$, the conclusion follows immediately. □

Lemma 2.3.1 reveals a surprising fact that when restricted to a given ordering cycle the inventory holding and shortage cost function in *both* backlogging model and lost-sales model is piecewise

linear convex! In the following, we will exploit this simple structure of the inventory holding and shortage cost function to devise a unified and efficient approach to compute policy decisions.

Robust optimization approach for computing (u_τ, ξ) Let τ denote the first period of a new ordering cycle. Let x_τ denote the available inventory and $\mathbf{d}[1, \tau-1]$ denote the past demand realization. By Assumption 2.2.1, $\mathbf{d}[1, \tau-1]$ does not provide any new information about the future demand $(D_\tau, D_{\tau+1}, \ldots, D_T)$. Following [24], we introduce the uncertainty set

$$\mathcal{D}[\tau] = \left\{ \mathbf{d}[\tau] \middle| \bar{\omega}_t - \hat{\omega}_t \le d_t \le \bar{\omega}_t + \hat{\omega}_t, \quad \sum_{\ell=\tau}^{t} \frac{|d_\ell - \bar{\omega}_\ell|}{\hat{\omega}_\ell} \le \Gamma^\tau_{t-\tau+1}, \quad \forall\, \tau \le t \le T \right\}, \tag{2.3.3}$$

where $\bar{\omega}_t$ and $\hat{\omega}_t$ are, respectively, the *nominal value* and *deviation* of d_t, and are computed from $\{(\bar{d}_t, \sigma_t^2) : \tau \le t \le T\}$. In [24], $\bar{\omega}_t = \bar{d}_t$, $\hat{\omega}_t = \min\{2\sigma_t, \bar{\omega}_t\}$. The *budget of uncertainty* constraint $\sum_{\ell=\tau}^{t} \frac{|d_\ell - \bar{\omega}_\ell|}{\hat{\omega}_\ell} \le \Gamma^\tau_{t-\tau+1}$ was introduced in [22, 23]. It rules out large deviations in the cumulative demand; and, therefore, controls conservativeness of the robust constraints. Note that the uncertainty set $\mathcal{D}[\tau]$ defined in (2.3.3) is symmetric about the nominal demand $\bar{\omega}_t$. When the demand distribution is asymmetric about the mean, an asymmetric uncertainty set may be more appropriate. We consider such extensions in Section 2.6.2 and show that asymmetric uncertainty sets lead to significantly improved performance when the demand distribution is heavy tailed.

The worst-case average cost $F(u_\tau, \xi)$ over the cycle $(\tau, \tau+1, \ldots, \tau+\xi-1)$ is defined as follows:

$$F(u_\tau, \xi) = \max_{\mathbf{d}[\tau] \in \mathcal{D}[\tau]} \left\{ \frac{K\mathbf{1}(u_\tau > 0) + cu_\tau + H\Big(\tau, x_\tau, \xi, u_\tau, \mathbf{d}[\tau, \xi]\Big)}{\xi} \right\}, \tag{2.3.4}$$

where H denotes the total inventory holding and shortage costs over the cycle. The optimal decision (u_τ^*, ξ^*) is chosen by solving the optimization problem

$$\min_{\xi \in \left\{1,2,\ldots,\min\{T-\tau+1,U\}\right\}} \min_{0 \le u_\tau \le M - x_\tau} \Big\{ F(u_\tau, \xi) \Big\}, \tag{2.3.5}$$

where U is a user-defined parameter. Setting $U = \infty$ allows the cycle length ξ to be set to any positive integer less than or equal to the number of remaining periods in the planning horizon. In our numerical experiments, we found that $U = 12$ was sufficient for very good performance. Note that in (2.3.4) and (2.3.5), τ and x_τ are known constants, i.e., part of the inputs.

It is clear that the traditional myopic policy is a special case of our policy with $\xi = 1$. The following example highlights the benefits of the flexibility of allowing longer cycle lengths. Consider a lost-sales inventory model with deterministic demand $d_t = \bar{d} = 100$, $T = 10$, $K = 1000$, $c = h = 0$, $b = 5$ and $x_1 = 0$. Since $K > b\bar{d}$, the myopic policy never places an order, and incurs a total cost of 5000. In contrast, our policy selects $\xi = 10$, $u_1 = 1000$, and incurs a total cost of 1000.

Solution algorithm To solve (2.3.5), we fix ξ, solve the inner minimization problem $\phi(\xi) = \min_{0 \le u_\tau \le M - x_\tau} F(u_\tau, \xi)$, and then enumerate all feasible value of ξ. Since the number of feasible ξ is at most U, and we show below that $\phi(\xi)$ can be solved efficiently, an optimal solution to (2.3.5) can be computed efficiently. The optimization problem $\phi(\xi)$ is equivalent to

$$\min_{0 \le u_\tau \le M - x_\tau} \max_{\mathbf{d}[\tau] \in \mathcal{D}[\tau]} \Big\{ K\mathbf{1}(u_\tau > 0) + c u_\tau + H\big(\tau, x_\tau, \xi, u_\tau, \mathbf{d}[\tau, \xi]\big) \Big\}. \tag{2.3.6}$$

Since $\left(\tau, x_\tau, \xi\right)$ are known constants in (2.3.6), we abbreviate $H\left(\tau, x_\tau, \xi, u_\tau, \mathbf{d}[\tau,\xi]\right)$ as $H\left(u_\tau, \mathbf{d}[\tau,\xi]\right)$. Lemma 2 reformulates (2.3.6) into a more tractable form.

Lemma 2.3.2. *The optimization problem (2.3.6) is equivalent to the following optimization problem:*

$$\min \left\{ \begin{array}{c} \min_{0\le u_\tau \le M - x_\tau} \max_{\mathbf{d}[\tau,\xi]\in\mathcal{D}[\tau,\xi]} \Big\{ K + cu_\tau + H\left(u_\tau, \mathbf{d}[\tau,\xi]\right) \Big\}, \\ \max_{\mathbf{d}[\tau,\xi]\in\mathcal{D}[\tau,\xi]} H\left(0, \mathbf{d}[\tau,\xi]\right) \end{array} \right\}, \tag{2.3.7}$$

where $\mathcal{D}[\tau,\xi]$ denotes the projection of the set $\mathcal{D}[\tau]$ to the space spanned by $\mathbf{d}[\tau,\xi]$ and is given by

$$\mathcal{D}[\tau,\xi] = \Big\{ \mathbf{d}[\tau,\xi] \Big| \bar{\omega}_t - \hat{\omega}_t \le d_t \le \bar{\omega}_t + \hat{\omega}_t, \quad \sum_{\ell=\tau}^{t} \frac{|d_\ell - \bar{\omega}_\ell|}{\hat{\omega}_\ell} \le \Gamma^\tau_{t-\tau+1}, \quad \forall\, \tau \le t \le \tau+\xi-1 \Big\}. \tag{2.3.8}$$

Proof. Proof. Since $\mathbf{d}[\tau,\xi]$ is a sub-vector of $\mathbf{d}[\tau]$, it follows that

$$\max_{\mathbf{d}[\tau]\in\mathcal{D}[\tau]} \Big\{ K\mathbf{1}(u_\tau > 0) + cu_\tau + H\left(u_\tau, \mathbf{d}[\tau,\xi]\right) \Big\} = \max_{\mathbf{d}[\tau,\xi]\in\mathcal{D}[\tau,\xi]} \Big\{ K\mathbf{1}(u_\tau > 0) + cu_\tau + H\left(u_\tau, \mathbf{d}[\tau,\xi]\right) \Big\}.$$

Thus, (2.3.6) is equivalent to

$$\min_{0\le u_\tau \le M - x_\tau} \max_{\mathbf{d}[\tau,\xi]\in\mathcal{D}[\tau,\xi]} \Big\{ K\mathbf{1}(u_\tau > 0) + cu_\tau + H\left(u_\tau, \mathbf{d}[\tau,\xi]\right) \Big\}. \tag{2.3.9}$$

To account for the discontinuity of the cost function at $u_\tau = 0$, we split (2.3.9) into two terms and take the minimum:

$$\min \left\{ \begin{array}{c} \min_{0<u_\tau\le M-x_\tau} \max_{\mathbf{d}[\tau,\xi]\in\mathcal{D}[\tau,\xi]} \Big\{ K + cu_\tau + H\big(u_\tau, \mathbf{d}[\tau,\xi]\big) \Big\}, \\ \\ \max_{\mathbf{d}[\tau,\xi]\in\mathcal{D}[\tau,\xi]} H\big(0, \mathbf{d}[\tau,\xi]\big) \end{array} \right\}.$$

(2.3.7) follows from observing that $\max_{\mathbf{d}[\tau,\xi]\in\mathcal{D}[\tau,\xi]} K + H\big(0,\mathbf{d}[\tau,\xi]\big) \ge \max_{\mathbf{d}[\tau,\xi]\in\mathcal{D}[\tau,\xi]} H\big(0,\mathbf{d}[\tau,\xi]\big)$.

□

Next, we show how to efficiently solve (2.3.7). We focus on the first program and show that the solution to the second program is a byproduct of solving the first. For notational simplicity, we drop all the subscripts appear in (2.3.7), and let $u = u_\tau$, $x = x_\tau$, $\mathbf{d} = \mathbf{d}[\tau,\xi]$, and $\mathcal{D} = \mathcal{D}[\tau,\xi]$. Then the first program in (2.3.7) can be written as

$$\min_{0\le u\le M-x} \max_{\mathbf{d}\in\mathcal{D}} \Big\{ K + cu + H(u,\mathbf{d}) \Big\}. \tag{2.3.10}$$

Since Lemma 2.3.1 shows that $cu + H(u,\mathbf{d})$ is convex in u, (2.3.10) is a convex optimization problem. We solve (2.3.10) using the Benders' decomposition approach [18, 25] displayed in Figure 2.1. In this procedure, one iteratively solves the inventory manager's problem that selects a candidate order quantity $\tilde{u}$ and the adversary's problem that computes the worst-case demand sequence $\tilde{\mathbf{d}} \in \mathcal{D}$ corresponding to $\tilde{u}$ that is added to the working set $\tilde{D}$ of candidate demand sequences for the inventory manager's problem. For this procedure to converge quickly, the inventory manager's problem and the adversary's problem should be efficiently solvable, and,

1. **Initialize:** $\tilde{D} = \left\{\tilde{\mathbf{d}} = (\bar{\omega}_\tau, \bar{\omega}_{\tau+1}, \ldots, \bar{\omega}_{\tau+\xi-1})\right\}$, $LB = 0$ and $UB = +\infty$.

2. **Inventory manager's problem:** Solve the inventory manager's problem with $\mathbf{d} \in \tilde{D}$, i.e., the adversary is restricted to the set of demands in the working list $\tilde{D}$. Let

$$\tilde{u} \in \underset{0 \le u \le M-x}{\operatorname{argmin}} \max_{\mathbf{d} \in \tilde{D}} K + cu + H(u, \mathbf{d}). \tag{2.3.11}$$

   Set $LB \leftarrow \max_{\mathbf{d} \in \tilde{D}} K + c\tilde{u} + H(\tilde{u}, \mathbf{d})$.

3. **Adversary's problem:** Solve the worst case demand sequence $\tilde{\mathbf{d}}$ corresponding to the policy $\tilde{u}$ computed in Step 2, i.e., compute

$$\tilde{\mathbf{d}} \in \underset{\mathbf{d} \in \mathcal{D}}{\operatorname{argmax}} K + c\tilde{u} + H(\tilde{u}, \mathbf{d}). \tag{2.3.12}$$

   Note that $\tilde{\mathbf{d}}$ is selected from the full set $\mathcal{D}$ and not the current working list $\tilde{D}$. Set $UB \leftarrow \min\left(UB, K + c\tilde{u} + H(\tilde{u}, \tilde{\mathbf{d}})\right)$.

4. **Termination test and update:** If $UB - LB$ is small enough, then **EXIT**. Otherwise, set $\tilde{D} \leftarrow \tilde{D} \cup \{\tilde{\mathbf{d}}\}$ and return to Step 2.

Figure 2.1: Benders' decomposition

in addition, the total number of iterations should be small. Since (2.3.10) is a one-dimensional convex optimization problem in u, and Benders' decomposition is a sub-gradient algorithm for solving (2.3.10), we expect the algorithm to converge quickly, and, indeed in our numerical experiments, we find that in nearly all cases the Benders' decomposition requires at most 5 iterations to guarantee that $\frac{UB-LB}{LB} < 10^{-5}$. In the rest of this section we show that both the inventory manager's problem and the adversary's problem can be solved efficiently.

Lemma 2.3.3. *The inventory manager's problem (2.3.11) in Figure 2.1 can be reformulated as an LP with* 2 *decision variables and* $(\xi+1)|\tilde{D}|+2$ *constraints, where* $|\tilde{D}|$ *denotes the cardinality of the set* $\tilde{D}$.

Proof. Proof. Let $\tilde{D} = \{\mathbf{d}^1, \mathbf{d}^2, \ldots, \mathbf{d}^{|\tilde{D}|}\}$, where $\mathbf{d}^j = (d^j_\tau, d^j_{\tau+1}, \ldots, d^j_{\tau+\xi-1})$. It is clear that (2.3.11) is equivalent to $\min\left\{ t \,\middle|\, u \in [0, M-x], \quad t \geq K + cu + H(u, \mathbf{d}^j), \quad \forall\, j = 1, 2, \ldots, |\tilde{D}| \right\}$, where (u,t) are the decision variables. Recall that Lemma 2.3.1 implies that $H(u, \mathbf{d}^j)$ is the maximum of $\xi + 1$ linear functions of u; therefore, it follows that for each $1 \leq j \leq |\tilde{D}|$, constraint $t \geq K + cu + H(u, \mathbf{d}^j)$ can be reformulated as $\xi + 1$ linear constraints in (u,t). □

The inventory manager's problem (2.3.11) can be reduced to an LP with no assumption on the uncertainty set $\mathcal{D}$. Next, we show that the adversary's problem (2.3.12) can be solved efficiently. The key insight in establishing the result is that, as indicated by Lemma 2.3.1, the cost function is a piecewise linear function with at most $\xi + 1$ linear pieces.

Lemma 2.3.4. *The adversary's problem (2.3.12) can be reduced to solving at most* $\xi + 1$ *LPs.*

Proof. Proof. See online supplement. □

Lemma 2.3.3 and 2.3.4 imply that the Benders' decomposition approach should be efficient for solving (2.3.10), i.e., the first program in (2.3.7). Since the second program in (2.3.7) is simply the adversary's problem with $\tilde{u} = 0$, we now have an efficient algorithm for solving (2.3.7).

2.4 Inventory model with non-zero lead time

In this section, we consider an inventory model which is identical to our basic model in all other respects except that the lead time is a positive constant L, i.e., the quantity ordered at the beginning of period k arrives at the beginning of period $k + L$. Thus, the inventory dynamics are as follows:

1. When the excess demand is backlogged, $x_{k+1} = x_k + u_{k-L} - d_k$ for $k \in \{1,2,\ldots,T-1\}$.

2. When the excess demand is lost, $x_{k+1} = \max\{x_k + u_{k-L} - d_k, 0\}$ for $k \in \{1,2,\ldots,T-1\}$.

Note that u_t, for $t \leq 0$, denotes orders placed before period 1, and are known constants at the beginning of the planning horizon. The cost C_k incurred during period k is given by $C_k = K\mathbf{1}(u_k > 0) + c \cdot u_k + \max\left\{h \cdot (x_k + u_{k-L} - d_k), -b \cdot (x_k + u_{k-L} - d_k)\right\}$[1]. For simplicity, we assume that the installation capacity $M = +\infty$.

Positive lead times lead to a larger state space which makes computing the optimal policy for the inventory model much harder, especially when the excess demand is lost. Our understanding of the lost-sales inventory model with positive lead time is still very limited [95], and this model is still a very active area of research. See [54], [62, 63], [49], [95], [45] and references therein for the progression of our understanding of this model.

In order to generalize the basic policy in Section 2.3 to the inventory model with positive lead time L, we associate the *inventory cycle* $(\tau+L, \tau+L+1, \ldots, \tau+L+\xi-1)$ with the ordering cycle $(\tau, \tau+1, \ldots, \tau+\xi-1)$, i.e., the inventory cycle is obtained by shifting the corresponding ordering cycle forward by L periods. Our modified policy is defined as follows. As before orders are only placed at the beginning of each *ordering cycle*. However, unlike in the basic policy, (u_τ, ξ) are chosen to minimize the cost associated with placing the order u_τ and the inventory holding cost and shortage cost incurred over the corresponding *inventory cycle* averaged over cycle length ξ. The cost function is formally defined in (2.4.1). The motivation here is that the order decision u_τ, and the implicit decision $u_t = 0$ for $\tau+1 \leq t \leq \tau+\xi-1$, are matched with

[1] We do not include the pipeline inventory costs; however, since the lead time is a constant L, this cost can be readily incorporated into the model by redefining the variable ordering cost as $c' = c + Lh'$, where h' denotes the per unit per period pipeline inventory cost.

the inventory holding and shortage cost incurred over the corresponding inventory cycle. Note that when $L = 0$, the modified policy reduces to the basic policy.

Robust optimization approach for computing (u_τ, ξ) Let x_τ denote the available inventory and $\bar{\mathbf{u}} = (\bar{u}_{\tau-L}, \bar{u}_{\tau-L+1}, \ldots, \bar{u}_{\tau-1})$ the vector of outstanding orders at the beginning of the first period τ of a new ordering cycle. The inventory cycle corresponding to this ordering cycle is $(\tau+L, \tau+L+1, \ldots, \tau+L+\xi-1)$. Since $u_t = 0$ for all $t \in \{\tau+1, \ldots, \tau+\xi-1\}$, the order u_τ arrives at the beginning of the inventory cycle and the realized demand over the inventory cycle is $\mathbf{d}[\tau+L, \xi]$, the inventory holding cost and shortage cost incurred over the inventory cycle is given by $H\big(\tau+L, x_{\tau+L}(\mathbf{d}[\tau, L]), \xi, u_\tau, \mathbf{d}[\tau+L, \xi]\big)$, where the functional form of H is given by (2.3.1) (resp. (2.3.2)) when the unmet demand is backlogged (resp. lost), and we emphasize the fact that the inventory level $x_{\tau+L}$ is a function of $\mathbf{d}[\tau, L]$. The cost $G(u_\tau, \xi)$ associated with the decision (u_τ, ξ) is given by

$$G(u_\tau, \xi) = \max_{\mathbf{d}[\tau] \in \mathcal{D}[\tau]} \left\{ \frac{K\mathbf{1}(u_\tau > 0) + cu_\tau + H\Big(\tau+L, x_{\tau+L}\big(\mathbf{d}[\tau, L]\big), \xi, u_\tau, \mathbf{d}[\tau+L, \xi]\Big)}{\xi} \right\} \quad (2.4.1)$$

The order quantity u_τ and cycle length ξ are computed by solving

$$\min_{\xi \in \left\{1, 2, \ldots, \min\{T-\tau-L+1, U\}\right\}} \min_{u_\tau \geq 0} \Big\{ G(u_\tau, \xi) \Big\}. \quad (2.4.2)$$

Solution algorithm We use the solution algorithm proposed in Section 2.3 to solve program (2.4.2). Specifically, we first fix ξ, solve the inner minimization problem, and then enumerate all feasible value of ξ. In (2.4.1) the adversary chooses a demand sequence $\mathbf{d}[\tau]$ that belongs to the uncertainty

set $\mathcal{D}[\tau]$ which is defined in (2.3.3). A simple modification of the argument in Lemma 2.3.2 establishes that the adversary can be restricted to the set $\mathcal{D}[\tau, L+\xi]$ which is defined analogous to the set in (2.3.8). To simplify notation, let $u = u_\tau$, $\mathbf{d}^O = \mathbf{d}[\tau, L]$, $x(\mathbf{d}^O) = x_{\tau+L}(\mathbf{d}[\tau, L])$, $\mathbf{d}^I = \mathbf{d}[\tau+L, \xi]$, $\mathbf{d} = (\mathbf{d}^O, \mathbf{d}^I) = \mathbf{d}[\tau, L+\xi]$, and $\mathcal{D} = \mathcal{D}[\tau, L+\xi]$.

Since $(\tau+L, \xi)$ are constants in the inner minimization problem, we write $H\left(\tau+L, x(\mathbf{d}^O), \xi, u, \mathbf{d}^I\right)$ as $H\left(x(\mathbf{d}^O), u, \mathbf{d}^I\right)$. The inner minimization problem in u is equivalent to

$$\min_{u\geq 0} \max_{\mathbf{d}\in\mathcal{D}} \left\{ K\mathbf{1}(u>0) + cu + H\left(x(\mathbf{d}^O), u, \mathbf{d}^I\right) \right\}. \tag{2.4.3}$$

Using results in Lemma 2.3.2, it is easy to show that (2.4.3) can be reformulated as

$$\min \left\{ \begin{array}{c} \min_{u\geq 0} \max_{\mathbf{d}\in\mathcal{D}} \left\{ K + cu + H\left(x(\mathbf{d}^O), u, \mathbf{d}^I\right) \right\}, \\ \max_{\mathbf{d}\in\mathcal{D}} \left\{ H\left(x(\mathbf{d}^O), 0, \mathbf{d}^I\right) \right\} \end{array} \right\}. \tag{2.4.4}$$

We use Benders' decomposition to solve the first program in (2.4.4); the solution to the second program will be obtained as a special case.

Lemma 2.4.1. *The inventory manager's problem in the Benders' decomposition approach to solve the first problem in (2.4.4) is given by* $\min_{u\geq 0} \max_{\mathbf{d}\in\tilde{D}} K + cu + H\left(x\left(\mathbf{d}^O\right),\ u,\ \mathbf{d}^I\right)$,*where* $\tilde{D}$ *is a given finite working list of demand sequences. This problem can be reformulated as an LP with* 2 *decision variables and* $(\xi+1)|\tilde{D}|+1$ *constraints, where* $|\tilde{D}|$ *denotes the cardinality of the set* $\tilde{D}$.

The proof of this result is identical to that of Lemma 2.3.3. Next, we address the adversary's problem.

Lemma 2.4.2. *The adversary's problem in the Benders' decomposition applied to the first problem in (2.4.4) is given by*

$$\max_{\mathbf{d}=(\mathbf{d}^O,\mathbf{d}^I)\in\mathcal{D}} K + c\tilde{u} + H\Big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\Big), \tag{2.4.5}$$

where $\tilde{u}$ denotes the current working decision of the inventory manager. For backlogging dynamics, (2.4.5) is equivalent to solving at most $\xi+1$ LPs, and for lost-sales dynamics, (2.4.5) is equivalent to solving at most $(\xi+1)(L+1)$ LPs.

Proof. Proof. See online supplement. □

The proof of Lemma 2.4.2 is similar to that of Lemma 2.3.4. The complicating factor is that here the inventory position at the beginning of the inventory cycle is a (possibly nonlinear) function of the demand sequence. However, a careful case analysis establishes that the number of LPs that need to be solved still remains quite modest. In particular, in the backlogging model, the number of LPs one needs to solve remains the same as that in the zero lead time case; in the lost-sales model, the number of LPs is linear in the lead time L. Thus, the complexity of our solution algorithm scales gracefully with lead time.

2.5 Extension to joint pricing and inventory control

In this section, we consider a single installation joint pricing and inventory control model which is similar to that proposed in [38] and [30]. See, also, [31] and [78], for similar models. We extend our robust cycle-based control policy to efficiently compute a good solution for this model.

Dynamics We consider a single product single installation joint pricing and inventory control model over a finite discrete horizon of T periods. As in Section 2.2, we let x_k, u_k and d_k denote the inventory at the beginning of period k, the order quantity in period k, and the demand in period k, respectively. We denote by p_k the unit price of the product in period k. At the beginning of period k, the inventory manager chooses the order quantity $u_k \in [0,\ M - x_k]$ and the price $p_k \in \left[p^k_{\min},\ p^k_{\max}\right]$, where M denotes the capacity of the installation, and $\left[p^k_{\min},\ p^k_{\max}\right]$ denotes the feasible pricing interval. We assume that the excess demand in each period is backlogged, and the lead time is zero. Therefore, the system dynamics is given by $x_{k+1} = x_k + u_k - d_k$, for all $1 \leq k \leq T-1$.

Cost function The ordering cost and the inventory holding and shortage cost are assumed to be the same as in Section 2.2. In addition, the inventory manager receives revenue $p_k d_k$ in period k. Thus, the total cost incurred during period k is given by $C_k = K\mathbf{1}(u_k > 0) + c \cdot u_k + \max\left\{h \cdot (x_k + u_k - d_k), -b \cdot (x_k + u_k - d_k)\right\} - p_k d_k$. Note that we assume that we receive full revenue even for the units that are backlogged. This assumption is common in the joint pricing and inventory control literature [30, 38].

Demand distribution In the joint pricing and inventory control model, the demand process $\mathfrak{D} = (D_1, D_2, \ldots, D_T)$ is no longer *exogenous*; the inventory manager can control the stochastic demand D_k by choosing the pricing decision p_k. We assume that the demand D_k is given by a linear stochastic demand function $Q_k^{(l)}(p_k, z_k)$ defined as follows.

Assumption 2.5.1. *The linear stochastic demand functions $\{Q_k^{(l)}(p_k, z_k)\}$ satisfy the following three conditions.*

1. *$Q_k^{(l)}(p_k, z_k)$ is separately linear in p_k and z_k, i.e., linear in p_k for fixed z_k, and vice versa.*
2. *$Q_k^{(l)}(p_k, z_k)$ is non-increasing in p_k for fixed z_k.*
3. *$\{z_k\}_{k=1}^T$ are mutually independent, and only the mean and variance of z_k are known.*

Assumption 2.5.1(3) is similar to Assumption 2.2.1 and 2.2.2. This class of demand functions includes the widely-used additive demand function with linear demand curve $Q(p,z) = \alpha - \beta p + z$ ($\alpha > 0, \beta \geq 0$) as a special case. Although the demand functions in [38] are allowed to be more general, the convexity related regularity conditions essentially restrict the demand functions to the class of linear demand functions defined here. We will briefly discuss a broad family of multiplicative stochastic demand functions in online supplement.

Extension of the robust cycle-based policy Our robust cycle-based control policy extends to this setting as follows. As in Section 2.3, we split the planning horizon into ordering cycles. At the first period τ of a new ordering cycle, after observing the inventory level x_τ, the inventory manager chooses order quantity u_τ, cycle length ξ and price sequence $p[\tau, \xi]$ to minimize the worst-case average cost over the cycle $(\tau, \tau+1, \ldots, \tau+\xi-1)$, $F(u_\tau, \xi, p[\tau, \xi])$, which is defined

as

$$F(u_\tau, \xi, p[\tau,\xi]) = \max_{z[\tau]\in\mathcal{Z}[\tau]} \left\{ \frac{K\mathbf{1}(u_\tau > 0) + cu_\tau + W(u_\tau, p[\tau,\xi], z[\tau,\xi])}{\xi} \right\},$$

where $z[\tau]$ denotes $(z_\tau, z_{\tau+1}, \ldots, z_T)$, $z[\tau,\xi]$ denotes $(z_\tau, z_{\tau+1}, \ldots, z_{\tau+\xi-1})$, $\mathcal{Z}[\tau]$ is defined similarly as $\mathcal{D}[\tau]$ in (2.3.3), i.e.,

$$\mathcal{Z}[\tau] = \left\{ z[\tau] \middle| \bar{\delta}_t - \hat{\delta}_t \le z_t \le \bar{\delta}_t + \hat{\delta}_t, \quad \sum_{\ell=\tau}^{t} \frac{|z_\ell - \bar{\delta}_\ell|}{\hat{\delta}_\ell} \le \Gamma^\tau_{t-\tau+1}, \quad \forall\, \tau \le t \le T \right\}, \tag{2.5.1}$$

where $\bar{\delta}_t$ and $\hat{\delta}_t$ are, respectively, the *nominal value* and *deviation* of z_t, and $W(u_\tau, p[\tau,\xi], z[\tau,\xi]) = \sum_{t=\tau}^{\tau+\xi-1} \max\left\{h(x_\tau + u_\tau - \sum_{\ell=\tau}^{t} Q_\ell^{(l)}(p_\ell, z_\ell)), -b(x_\tau + u_\tau - \sum_{\ell=\tau}^{t} Q_\ell^{(l)}(p_\ell, z_\ell))\right\} - \sum_{t=\tau}^{\tau+\xi-1} p_t Q_t^{(l)}(p_t, z_t)$. Note that for fixed z-sequence $z[\tau,\xi]$, $W(u_\tau, p[\tau,\xi], z[\tau,\xi])$ is a convex piecewise quadratic function in $(u_\tau, p[\tau,\xi])$. The optimal decision $(u^*_\tau, \xi^*, p^*[\tau,\xi])$ is chosen by solving the optimization problem

$$\min_{\xi\in\left\{1,2,\ldots,\min\{T-\tau+1,U\}\right\}} \quad \min_{\left\{0\le u_\tau\le M-x_\tau, p[\tau,\xi]\in\mathcal{P}[\tau,\xi]\right\}} F(u_\tau, \xi, p[\tau,\xi]), \tag{2.5.2}$$

where $\mathcal{P}[\tau,\xi] = \left\{ p[\tau,\xi] \;\middle|\; p^t_{\min} \le p_t \le p^t_{\max}, \text{ for } \tau \le t \le \tau+\xi-1 \right\}$.

Solution algorithm To solve (2.5.2), we fix ξ, solve the inner minimization problem, and then enumerate all feasible value of ξ. A simple modification of Lemma 2.3.2 establishes that the adversary can be restricted to the set $\mathcal{Z}[\tau,\xi]$ which is defined analogously to the set in (2.3.8). To simplify notation, let $u = u_\tau$, $x = x_\tau$, $p = p[\tau,\xi]$, $z = z[\tau,\xi]$, $\mathcal{P} = \mathcal{P}[\tau,\xi]$ and $\mathcal{Z} = \mathcal{Z}[\tau,\xi]$. Following the arguments in Lemma 2.3.2, it is clear that solving the inner minimization problem

of (2.5.2) is equivalent to solving the following optimization problem:

$$\min \left\{ \begin{array}{c} \min_{\left\{0\le u\le M-x, p\in\mathcal{P}\right\}} \max_{z\in\mathcal{Z}} \Big\{K+cu+W(u,p,z)\Big\}, \\ \min_{p\in\mathcal{P}} \max_{z\in\mathcal{Z}} \Big\{W(0,p,z)\Big\} \end{array} \right\}. \tag{2.5.3}$$

We use Benders' decomposition to solve each of the optimization problems in (2.5.3). We focus on the first program without loss of generality, i.e., we show how to apply Benders' decomposition to efficiently solve

$$\min_{\left\{0\le u\le M-x, p\in\mathcal{P}\right\}} \max_{z\in\mathcal{Z}} \Big\{K+cu+W(u,p,z)\Big\}. \tag{2.5.4}$$

It is easy to see that (2.5.4) is a convex optimization problem. Next, we show that the inventory manager's problem and the adversary's problem of (2.5.4) can be solved efficiently.

Lemma 2.5.2. *The inventory manager's problem of (2.5.4) is* $\min_{\left\{0\le u\le M-x, p\in\mathcal{P}\right\}} \max_{z\in\tilde{Z}} \Big\{K+cu+W(u,p,z)\Big\}$*, where* $\tilde{Z}$ *is a given finite working list of random terms. The inventory manager's problem is a convex quadratically constrained quadratic program (QCQP). When the linear demand function* $Q_k^{(l)}(p,z)$ *is additive, i.e.,* $Q_k^{(l)}(p,z) = \alpha_k - \beta_k p + z$ $(\alpha_k > 0, \beta_k \ge 0)$*, the inventory manager's problem reduces to a convex quadratic program (QP).*

Proof. Proof. The inventory manager's problem can be reformulated as $\min\Big\{ K+cu+t \;\Big|\; u\in [0,M-x], \quad p\in\mathcal{P}, \quad W(u,p,z^j)\le t, \quad \forall z^j\in\tilde{Z}\Big\}$. Recall that $W(u,p,z)$ is a convex piecewise quadratic function of (u,p) for a fixed z. Therefore, the constraint $W(u,p,z^j)\le t$ can be reformulated as a collection of convex quadratic constraints.

When the linear demand function is additive, the quadratic term in p in

$$\sum_{t=\tau}^{\tau+\xi-1} p_t Q_t^{(l)}(p_t, z_t) = \sum_{t=\tau}^{\tau+\xi-1} p_t\left(\alpha_t - \beta_t p_t + z_t\right) = \sum_{t=\tau}^{\tau+\xi-1} p_t z_t + \sum_{t=\tau}^{\tau+\xi-1} p_t\left(\alpha_t - \beta_t p_t\right)$$

is independent of z. Consequently, the inventory manager's problem can be reformulated as a convex quadratic program, i.e., with a convex quadratic objective but only linear constraints.

□

Lemma A.3.2 implies that the inventory manager's problem is efficiently solvable. Next, we consider the adversary's problem.

Lemma 2.5.3. *The adversary's problem of (2.5.4) is* $\max_{z\in\mathcal{Z}}\left\{K + c\tilde{u} + W(\tilde{u},\tilde{p},z)\right\}$, *where* $(\tilde{u},\tilde{p})$ *denotes the current working decision of the inventory manager. The adversary's problem can be reduced to solving at most* $\xi+1$ *LPs.*

The proof of this lemma is similar to that of Lemma 2.3.4. Thus we have completed our algorithm.

2.6 Numerical experiments

In this section, we report the results of a set of numerical experiments that investigate the empirical performance of our cycle-based inventory control policy (CI).

2.6.1 Backlogging dynamics

We compared CI with the rolling horizon version of the open-loop policy (RHBT) proposed in [24]. RHBT efficiently computes the ordering decisions for a backlogging inventory control problem by solving LPs or MIPs. In [24] it was shown that the performance of RHBT is comparable or superior to that of the DP-based policies, and in [25] it was shown that the performance of RHBT is superior to that of the robust base-stock policy.

We set the parameters $\bar{\omega}_t$, $\hat{\omega}_t$ and Γ_j^t defining the uncertainty sets (2.3.3) to $\bar{\omega}_t = \bar{d}_t$, $\hat{\omega}_t = \min\{2\sigma_t, \bar{\omega}_t\}$, $\Gamma_j^t = \sqrt{j}$, for $1 \leq t \leq T$ and $1 \leq j \leq T-t+1$ (see Section 5 in [24]). To ensure a fair comparison, we use the same uncertainty sets for both CI and RHBT. We set $U = 12$ for CI. All the LPs and MIPs in the implementation of CI and RHBT are solved using Gurobi Optimizer v4.5. We call Gurobi Optimizer from MATLAB using Gurobi Mex [92].

Our goal in the numerical experiments was to understand the impact of problem parameters on the performance of CI and RHBT. In particular, we investigated the impact of different choices of the realized demand distribution, the mean demand trajectory, the standard deviation of demand, the cost parameters (K, c, h, b) and the lead time L. The parameters for the base model are displayed in Table 2.1, where the choice of c, h and b is suggested in [24]. To study the impact of the parameters, we change one parameter at a time while keeping all other parameters at their respective base values. For each configuration, we randomly generated $N = 100$ demand sequences and tested the performance of CI and RHBT on the same demand realizations. In these numerical experiments, we set $c = 0$ when *computing* the CI policy decisions. This is motivated by the fact that variable ordering costs incurred over a fixed finite horizon can be

viewed as a sunk cost when c is smaller than backlogging cost b. We found this modification lead to better performance for the CI policy. Note we still used the true value of c when *evaluating* the performance of the CI policy. We report the mean and standard deviation of the total cost incurred by CI and RHBT on the $N = 100$ sample paths. At the end of this section, we also report the running time of CI and RHBT.

Table 2.1: Parameters for base backlogging model.

Parameter	Value	Parameter	Value
Time horizon T	48	Fixed ordering cost K	500
Initial inventory x_1	0	Variable ordering cost/unit c	1
Demand distribution D_t	$N(\bar{d}_t, \sigma_t^2)$	Inventory holding cost/unit h	4
Mean of demand $\bar{d}_t$	$100 + 40\sin(\frac{2\pi}{12}t)$	Shortage cost/unit b	6
Std. dev. of demand σ_t	$0.25\bar{d}_t$	Lead time L	0

Impact of the realized demand distribution. We considered the following 5 different families of realized demand distributions: normal, Student's-$t(4)$, gamma, uniform and lognormal. In each configuration, $\{D_t\}_{t=1}^T$ were drawn from the same family, and with pre-specified mean $\{\bar{d}_t\}_{t=1}^T$ and variance $\{\sigma_t^2\}_{t=1}^T$ in the base model. Table 2.2 summaries the performance measures. Since CI and RHBT both only use the first two moments of demand distributions, their performance is fairly robust across various realized demand distributions. The performance of the two policies is comparable – the cost of the CI policy is 1.5% to 2.2% less on average, but volatility of the cost is higher.

Impact of the mean demand trajectory. We considered mean demand trajectory $\bar{d}_t = 100 + \beta\sin(\frac{2\pi}{12}t)$, with $\beta = 0, 20, 40, 60, 80$. In the base model, $\beta = 40$. Note that a larger β implies

Table 2.2: Impact of realized demand distribution.

Distribution	RHBT Mean: std dev	CI Mean(%): std dev
Normal	$29,330.6: 646.1$	$28,894.3(-1.49\%): 754.3$
T(4)	$28,990.6: 654.4$	$28,562.1(-1.48\%): 693.1$
Gamma	$29,422.5: 561.1$	$28,782.3(-2.17\%): 727.8$
Uniform	$29,649.1: 573.8$	$28,996.7(-2.20\%): 674.9$
Lognormal	$29,446.3: 704.1$	$28,823.5(-2.11\%): 723.1$

greater variability in $\{\bar{d}_t\}_{t=1}^T$. Table 2.3 summarizes the performance measures. When $\beta =$

Table 2.3: Impact of β.

β	RHBT Mean: std dev	CI Mean(%): std dev
0	$29,919.9: 516.7$	$28,684.7(-4.13\%): 648.7$
20	$29,546.2: 574.4$	$28,867.0(-2.30\%): 654.2$
40	$29,330.6: 646.1$	$28,894.3(-1.49\%): 754.3$
60	$28,596.9: 632.2$	$28,984.6(+1.36\%): 788.2$
80	$27,844.3: 654.6$	$29,609.1(+6.34\%): 1,087.9$

0, i.e., $\{D_t\}_{t=1}^T$ is an IID demand, the CI cost is approximately 4% lower than RHBT. As β increases, the cost of RHBT decreases while the cost of CI increases. When $\beta = 80$, i.e., the demand trajectory is highly variable, CI costs 6% more on average and incurs a higher standard deviation. Thus, CI outperforms RHBT when the demand is changing smoothly, and the reverse is true when the demand is changing sharply.

Impact of the standard deviation of demand. We fixed $\gamma = \frac{\sigma_t}{\bar{d}_t}$ for $1 \leq t \leq T$, therefore the ratio γ determines the relative magnitude of standard deviation. In the base model, $\gamma =$ 0.25. Table 2.4 summarizes the performance measures. As γ increases, the performance of CI improves relative to RHBT. When $\gamma = 0.35$, i.e., the standard deviation of demand is high, *both*

Table 2.4: Impact of γ.

γ	RHBT Mean: std dev	CI Mean(%): std dev
0.15	$28,303.1: 351.8$	$28,236.7(-0.23\%): 439.8$
0.20	$28,830.5: 489.7$	$28,714.6(-0.40\%): 596.9$
0.25	$29,330.6: 646.1$	$28,894.3(-1.49\%): 754.3$
0.30	$29,709.3: 728.6$	$29,203.1(-1.70\%): 790.1$
0.35	$30,195.1: 937.6$	$29,629.2(-1.87\%): 888.5$

the average cost and volatility of the cost of the CI policy are lower than those of the RHBT policy.

Impact of the fixed ordering cost. Table 2.5 summarizes the performance measures. CI has

Table 2.5: Impact of K.

K	RHBT Mean: std dev	CI Mean(%): std dev
0	$9,492.5: 536.6$	$9,492.5(-0.00\%): 536.6$
250	$21,277.1: 479.6$	$21,378.1(+0.47\%): 575.7$
500	$29,330.6: 646.1$	$28,894.3(-1.49\%): 754.3$
750	$34,101.2: 813.3$	$34,301.4(+0.59\%): 876.1$
1000	$38,070.8: 847.4$	$38,619.8(+1.44\%): 871.1$

a slightly higher average cost in 3 cases ($K = 250, 750, 1000$). When $K = 0$, we found that CI and RHBT incurred the same cost on every sample path.

Impact of the variable ordering cost. Table 2.6 summarizes the performance measures. CI has a lower average cost across all the 5 configurations, but the magnitudes of savings are decreasing as c increases. When c increases to 3, there are jumps in the volatility of the cost of both CI and RHBT.

Table 2.6: Impact of c.

	RHBT	CI
c	Mean: std dev	Mean(%): std dev
0	$24,545.7: 519.1$	$23,994.9(-2.24\%): 615.4$
1	$29,330.6: 646.1$	$28,894.3(-1.49\%): 754.3$
2	$33,923.3: 658.5$	$33,537.2(-1.14\%): 678.5$
3	$38,702.7: 947.1$	$38,323.1(-0.98\%): 1,085.1$
4	$43,556.1: 849.5$	$43,339.9(-0.50\%): 1,075.6$

Impact of the inventory holding cost. Table 2.7 summarizes the performance measures. CI

Table 2.7: Impact of h.

	RHBT	CI
h	Mean: std dev	Mean(%): std dev
2	$24,254.5: 489.5$	$23,883.2(-1.53\%): 555.3$
4	$29,330.6: 646.1$	$28,894.3(-1.49\%): 754.3$
6	$31,637.7: 755.9$	$31,530.0(-0.34\%): 900.9$
8	$33,106.6: 875.7$	$33,100.8(-0.02\%): 838.7$
10	$34,162.2: 952.8$	$34,299.6(+0.40\%): 1,024.7$

has a lower average cost in 4 out of the 5 cases, but the relative performance of CI over RHBT deteriorates as h increases.

Impact of the shortage cost. Table 2.8 summarizes the performance measures. CI has a

Table 2.8: Impact of b.

	RHBT	CI
b	Mean: std dev	Mean(%): std dev
2	$22,430.9: 664.4$	$21,917.1(-2.29\%): 539.8$
4	$26,639.4: 625.1$	$26,273.5(-1.37\%): 623.4$
6	$29,330.6: 646.1$	$28,894.3(-1.49\%): 754.3$
8	$30,870.1: 612.9$	$30,813.5(-0.18\%): 799.6$
10	$31,881.5: 685.4$	$31,492.1(-1.22\%): 664.6$

lower average cost across all the 5 cases, but also incurs significantly higher volatility of cost when $b = 6, 8$.

Impact of the lead time. When lead time is L, we assumed that the outstanding orders at the beginning of the planning horizon, $(u_{1-L}, \ldots, u_0) = (\bar{d}_1, \bar{d}_2, \ldots, \bar{d}_L)$. Table 2.9 summarizes the performance measures. It is clear from the table, that the average CI cost is significantly lower

Table 2.9: Impact of L.

	RHBT	CI
L	Mean: std dev	Mean(%): std dev
0	$29,330.6 : \quad 646.1$	$28,894.3(-1.49\%) : \quad 754.3$
1	$29,951.6 : \quad 638.5$	$28,899.3(-3.51\%) : \quad 773.4$
2	$30,519.4 : 1,154.7$	$29,214.4(-4.28\%) : 1,231.2$
4	$31,255.0 : 1,788.9$	$29,957.6(-4.15\%) : 1,747.9$
6	$31,812.8 : 2,227.6$	$31,057.7(-2.37\%) : 2,142.3$

when compared to the RHBT cost across all 5 configurations. The volatility of the cost of the two policies is comparable; it is also clear that the volatility of the cost of both policies increases relatively quickly as lead time L increases.

Next, we compared the running time of CI and RHBT. In particular, we investigated how the time horizon, T, and the fixed ordering cost, K, impact the running time of the two policies. For each configuration below, we randomly generated $N = 10$ demand sequences, and implemented CI and RHBT on the same demand realizations. We report the mean and standard deviation of the running time of CI and RHBT on the $N = 10$ sample paths.

Impact of the time horizon on running time. We considered $T = 24, 36, 48, 60, 72$, with all other parameters were chosen from the base model in Table 2.1. Table 2.10 reports the running time. We found that the running time of CI grows linearly with T, whereas the running time of

Table 2.10: Running time: Impact of T.

	RHBT	CI
T	Mean: std dev (in seconds)	Mean: std dev (in seconds)
24	0.72 : 0.02	4.74 : 0.21
36	5.16 : 0.24	7.77 : 0.38
48	28.25 : 0.88	10.38 : 0.25
60	72.99 : 1.04	13.16 : 0.34
72	141.77 : 1.87	16.24 : 0.26

RHBT increases exponentially as T increases. This result is not surprising – CI only relies on solving LPs, whereas RHBT solves MIPs with increasing size as T increases.

Impact of the fixed ordering cost on running time. We fixed $T = 48$ and let $K = 0, 250, 500, 750, 1000$ in each of the 5 scenarios. Table 2.11 reports the running time. Table 2.11 shows some

Table 2.11: Running time: Impact of K.

	RHBT	CI
K	Mean: std dev (in seconds)	Mean: std dev (in seconds)
0	1.24 : 0.01	26.73 : 1.76
250	1.65 : 0.02	21.49 : 0.50
500	28.25 : 0.88	10.38 : 0.25
750	48.68 : 0.89	10.26 : 0.41
1000	46.05 : 1.08	9.12 : 0.55

interesting results. When $K = 0$ or 250, RHBT costs less than 2 seconds – approximately 5% of that of CI. However, as K increases, the running time of RHBT increases significantly, indicating that the MIPs in RHBT are becoming harder to solve as K increases. When $K = 750$

or 1000, it takes RHBT almost 50 seconds to solve the problem. On the other hand, the running time of CI in fact decreases as K increases, since a larger K implies a lower ordering frequency, and therefore fewer decisions to make. It is clear that, when $K = 1000$, the running time of RHBT is more than 5 times of that of CI.

2.6.2 Lost-sales dynamics

Inventory control when lead time is zero

In the lost-sales model with zero lead time, we compared CI with the DP solution and the optimal base-stock policy [94] in *hindsight* (BH). Following [24], we compute the DP policy *assuming* $D_t \sim N(\bar{d}_t, \sigma_t^2)$ for $1 \leq t \leq T$, and approximate $N(\bar{d}_t, \sigma_t^2)$ by a probability mass function supported on the five points $\{\bar{d}_t - 2\sigma_t, \bar{d}_t - \sigma_t, \bar{d}_t, \bar{d}_t + \sigma_t, \bar{d}_t + 2\sigma_t\}$. The BH policy is base-stock policy but the order-up-to level S is computed in hindsight, i.e., *after* observing the entire demand realization $\mathbf{d}$. We compute the optimal base-stock level for a particular $\mathbf{d}$ by an exhaustive search. It is clear that the cost of BH is a lower bound for the cost of any non-anticipatory base-stock policy. The choice of BH as a comparison policy is motivated by the fact that base-stock policies with constant order-up-to level, are commonly used in practice in lost-sales setting [49].

We set the parameters $\bar{\omega}_t$, $\hat{\omega}_t$, Γ_j^t and U in the CI policy to those used in Section 2.6.1. In the DP policy, the state at each time instance is the inventory at hand. We restricted the state to the interval $[0, 500]$, i.e., set $M = 500$, and further restricted both the state and action values to multiples of 0.1.

We investigated the impact of the same set of parameters as in Section 2.6.1 on the performance of CI, DP and BH, except that we did not consider positive lead time since DP does not scale with lead time L. The parameters for the base model are displayed in Table 2.12. To study the impact of the parameters, we vary one parameter at a time while keeping all other parameters at their respective base values. For each configuration, we randomly generated $N = 100$ demand sequences and tested the performance of CI, DP and BH on the same demand realizations. In this section, we still set $c = 0$ when *computing* the CI policy decisions and used the true value of c when *evaluating* the performance of the CI policy. We report the mean and standard deviation of the total cost incurred by CI, DP and BH on the $N = 100$ sample paths. We report the running time for CI and DP at the end of this section.

Table 2.12: Parameters for base lost-sales model.

Parameter	Value	Parameter	Value
Time horizon T	48	Fixed ordering cost K	0
Initial inventory x_1	0	Variable ordering cost/unit c	1
Demand distribution D_t	$N(\bar{d}_t, \sigma_t^2)$	Inventory holding cost/unit h	4
Mean of demand $\bar{d}_t$	100	Shortage cost/unit b	12
Std. dev. of demand σ_t	$0.25\bar{d}_t$	Lead time L	0

Impact of the realized demand distribution. As in Section 2.6.1, we considered the following 5 different families of realized demand distributions: normal, Student's-$t(4)$, gamma, uniform and lognormal. Table 2.13 summaries the performance measures. DP and CI have comparable performance across the 5 scenarios; they are both robust with respect to the changes in realized

Table 2.13: Impact of realized demand distribution.

Distribution	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
Normal	$11,090.3: 545.5$	$11,105.8(+0.14\%): 520.1$	$10,611.4(-4.32\%): 619.2$
T(4)	$10,845.0: 774.7$	$10,871.7(+0.25\%): 750.2$	$9,954.9(-8.21\%): 925.9$
Gamma	$11,414.9: 717.9$	$11,408.8(-0.05\%): 668.3$	$10,812.9(-5.27\%): 834.0$
Uniform	$11,009.3: 421.6$	$10,998.1(-0.10\%): 407.8$	$10,780.0(-2.08\%): 452.3$
Lognormal	$11,418.7: 670.9$	$11,405.9(-0.11\%): 673.6$	$10,691.6(-6.37\%): 953.2$

demand distributions. BH incurs significantly less cost, since the mean demand trajectory is flat and there is no fixed ordering cost in the base model.

Impact of the mean demand trajectory. We considered mean demand trajectory $\bar{d}_t = 100 + \beta \sin(\frac{2\pi}{12}t)$, with $\beta = 0, 20, 40, 60, 80$. In the base model, $\beta = 0$, i.e., demands are IID random variables. Table 2.14 summarizes the performance measures. The costs of DP and CI are almost

Table 2.14: Impact of β.

β	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
0	$11,090.3: 545.5$	$11,105.8(+0.14\%): 520.1$	$10,611.4(-4.32\%): 619.2$
20	$11,229.2: 532.5$	$11,215.6(-0.12\%): 517.7$	$11,763.6(+4.76\%): 799.7$
40	$11,247.0: 555.7$	$11,209.8(-0.33\%): 529.9$	$14,641.3(+30.18\%): 989.3$
60	$11,144.8: 603.9$	$11,145.8(+0.01\%): 587.7$	$17,674.9(+58.59\%): 1,345.3$
80	$11,206.2: 609.4$	$11,218.7(+0.11\%): 643.5$	$21,180.5(+89.01\%): 1,415.5$

equal; however, the cost of BH increases significantly as β increases. This result is not surprising – the performance of base-stock policies with constant base-stock level is very sensitive to the variability of the mean demand trajectory. In cases where the mean demands vary relatively sharply, base-stock policies deteriorate quickly.

Impact of the standard deviation of demand. We fixed $\gamma = \frac{\sigma_t}{d_t}$ for $1 \leq t \leq T$. In the base model, $\gamma = 0.25$. Table 2.15 summarizes the performance measures. It is clear from the table,

Table 2.15: Impact of γ.

γ	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
0.15	8,568.6 : 307.9	8,576.3(+0.09%) : 297.0	8,293.9(−3.21%) : 338.2
0.20	9,800.2 : 493.2	9,790.2(−0.10%) : 470.9	9,342.9(−4.67%) : 601.1
0.25	11,090.3 : 545.5	11,105.8(+0.14%) : 520.1	10,611.4(−4.32%) : 619.2
0.30	12,326.4 : 752.1	12,337.4(+0.09%) : 732.2	11,728.2(−4.85%) : 881.5
0.35	13,699.2 : 854.7	13,707.5(+0.06%) : 836.4	13,064.4(−4.63%) : 1,036.1

that the costs of DP and CI are still very close.

Impact of the fixed ordering cost. Table 2.16 summarizes the performance measures. CI

Table 2.16: Impact of K.

K	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
0	11,090.3 : 545.5	11,105.8(+0.14%) : 520.1	10,611.4(− 4.32%) : 619.2
250	22,935.7 : 574.0	23,138.0(+0.88%) : 522.6	22,613.9(− 1.40%) : 668.4
500	31,234.3 : 913.5	32,170.9(+3.00%) : 792.4	34,630.4(+10.87%) : 733.8
750	37,061.8 : 951.0	37,754.4(+1.87%) : 711.0	46,562.9(+25.64%) : 737.5
1000	42,124.8 : 1,043.6	43,339.0(+2.88%) : 953.7	58,557.4(+39.01%) : 735.6

incurs 1% to 3% higher average cost compared to DP when $K \neq 0$, but saves approximately 10% to 20% in standard deviation. On the other hand, the performance of BH deteriorates significantly as K grows, since base-stock policies place positive orders in every period as long as demand is positive, regardless of the value of K. This suggests that base-stock policies are in general not good candidates when there are non-zero fixed ordering costs.

Impact of the variable ordering cost. Table 2.17 summarizes the performance measures. The performance of DP and CI is comparable when $c = 0, 1, 2, 3$ – DP incurs a slightly lower

Table 2.17: Impact of c.

c	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
0	$6,411.9$: 681.1	$6,408.7(-0.05\%)$: 681.3	$5,950.7(-7.19\%)$: 732.4
1	$11,090.3$: 545.5	$11,105.8(+0.14\%)$: 520.1	$10,611.4(-4.32\%)$: 619.2
2	$15,785.7$: 549.7	$15,829.4(+0.28\%)$: 539.0	$15,238.5(-3.47\%)$: 715.9
3	$20,421.0$: 470.9	$20,468.4(+0.23\%)$: 445.1	$19,710.3(-3.48\%)$: 840.0
4	$24,886.2$: 1,171.6	$25,268.1(+1.53\%)$: 553.0	$24,330.1(-2.23\%)$: 927.8

average cost while CI incurs a slightly lower standard deviation. When $c = 4$, CI costs 1.5% more on average but is more stable in that the standard deviation of CI is only half of that of DP.

Impact of the holding cost. Table 2.18 summarizes the performance measures. CI incurs 3%

Table 2.18: Impact of h.

h	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
2	$8,508.8$: 507.5	$8,778.9(+3.17\%)$: 325.6	$8,411.0(-1.15\%)$: 501.2
4	$11,090.3$: 545.5	$11,105.8(+0.14\%)$: 520.1	$10,611.4(-4.32\%)$: 619.2
6	$13,059.5$: 1,110.3	$12,732.3(-2.51\%)$: 665.2	$12,329.4(-5.59\%)$: 763.5
8	$14,154.8$: 1,023.1	$14,095.3(-0.42\%)$: 854.7	$13,780.5(-2.64\%)$: 877.8
10	$14,918.8$: 1,079.7	$14,920.1(+0.01\%)$: 1,007.9	$14,681.5(-1.59\%)$: 1,004.0

more in average when $h = 2$ and 2.5% less when $h = 6$ compared to DP; in other scenarios the differences are small. On the other hand, CI has a smaller standard deviation across all 5 cases. In particular, when $h = 6$, the standard deviation of CI cost is approximately 40% less than that of DP cost.

Impact of the shortage cost. Table 2.19 summarizes the performance measures. In all cases

Table 2.19: Impact of b.

	DP	CI	BH
b	Mean: std dev	Mean(%): std dev	Mean(%): std dev
8	$9,943.5:\ 682.4$	$9,987.1(+0.44\%):\ 483.2$	$9,610.7(-3.35\%):\ 575.9$
10	$10,951.5:\ 485.9$	$10,652.9(-2.73\%):\ 519.1$	$10,234.4(-6.55\%):\ 650.7$
12	$11,090.3:\ 545.5$	$11,105.8(+0.14\%):\ 520.1$	$10,611.4(-4.32\%):\ 619.2$
14	$11,236.6:\ 610.4$	$11,428.5(+1.71\%):\ 581.3$	$10,917.3(-2.84\%):\ 701.4$
16	$11,520.0:\ 684.2$	$11,863.4(+2.98\%):\ 591.4$	$11,233.5(-2.49\%):\ 818.2$

except when $b = 10$, CI incurs a higher average cost and a lower standard deviation compared to DP. When $b = 10$, CI costs 2.7% less but incurs a higher standard deviation of cost.

Impact of the fixed ordering cost on running time. Next, we report the running time for CI and DP. In this set of experiments we were interested in investigating the impact of the fixed ordering cost K on the running time of these two policies. We set $K = 0, 250, 500, 750, 1000$, and set all other parameters to the values in the base model displayed in Table 2.12. Table 2.20 reports the running time. As in the backlogging case, the running time for CI typically decreases as K increases, and CI goes from being approximately 10 times faster to approximately 20 times faster as we increase K from 0 to 1000.

Table 2.20: Running time: Impact of K.

	DP	CI
K	(in seconds)	Mean: std dev (in seconds)
0	154.32	16.88 : 1.03
250	155.94	15.76 : 0.28
500	156.39	7.85 : 0.13
750	156.09	7.96 : 0.21
1000	157.13	8.34 : 0.20

Inventory control with positive lead time L

Due to the "curse of dimensionality", comparing CI with DP for this model is impractical. Therefore, we compared CI with the optimal stationary policy and eight plausible heuristics in [95] – two myopic policies, the dual-balancing policy [58], four base-stock policies and the constant-order policy [67]. For most scenarios with Poisson demand, the CI cost exceeds the optimal long-run average cost by no more than 5%, and the CI performance typically improves as lead time L increases. For the heavier tailed geometric demand, the performance of CI with *asymmetric* uncertainty sets is very close to optimal. See online supplement for details of this set of numerical experiments.

2.6.3 Joint pricing and inventory control

In this section, we report the empirical performance of the cycle-based joint pricing and inventory control policy in Section 2.5. We continue to refer to our proposed cycle-based policy as CI.

We considered two scenarios which are adapted from the two scenarios considered in [38]. In both scenarios, we assumed that the stochastic demand functions are of the form $Q_k(p_k, z_k) = \alpha_k - \beta_k p_k + z_k$, i.e., additive demand functions with linear demand curve. Thus, Lemma A.3.2 implies that the inventory manager's problem can be reformulated as a convex quadratic program. We use Gurobi Optimizer v4.5 to solve the QPs in the implementation of CI. We set the parameters $\bar{\delta}_t$, $\hat{\delta}_t$ and Γ^t_j defining the uncertainty sets (2.5.1) to $\bar{\delta}_t = \bar{z}_t$, $\hat{\delta}_t = 2\sigma_t$, $\Gamma^t_j = \frac{1}{4}\sqrt{j}$, for $1 \le t \le T$ and $1 \le j \le T - t + 1$. We also set $U = 15$.

We compared the performance of CI with the stochastic DP solution computed by assuming

that the random term z_t was distributed according to a 5-point approximation to the $N(\bar{z}_t, \sigma_t^2)$ distribution supported on the set $\{\bar{z}_t - 2\sigma_t, \bar{z}_t - \sigma_t, \bar{z}_t, \bar{z}_t + \sigma_t, \bar{z}_t + 2\sigma_t\}$. The state at each time instance is, as before, the inventory at hand. We restricted the state $x \in [-200, 500]$ and discretized the state to take values in the set of integers. The action space in the DP was two dimensional – u and p – and we restricted each of the two dimensions to take values in the set of integers.

Table 2.21 displays the parameters of the first scenario – the "Dress Scenario" in [38]. We considered the following 5 different families of realized random term distributions: normal, Student's-$t(4)$, gamma, uniform and lognormal. In each configuration, $\{z_t\}_{t=1}^T$ were drawn from the same family, and with pre-specified mean $\{\bar{z}_t\}_{t=1}^T$ and variance $\{\sigma_t^2\}_{t=1}^T$. For each configuration, we randomly generated $N = 100$ random term sequences and tested the performance of CI and DP on the same random term realizations. Since in the two scenarios of this section the variable ordering cost c is close to or larger than the backlogging cost b and the CI policy is myopic, the optimal decision of the policy tends to order less over a given cycle and "transfers" the variable ordering cost to the subsequent cycle. In order to correctly account for the externality of the ordering decision over a given cycle, in the numerical experiments we added another penalty term $c \cdot \max\{-(x_\tau + u_\tau - \sum_{\ell=\tau}^{\tau+\xi-1} Q_\ell^{(l)}(p_\ell, z_\ell)), 0\}$ to $W(u_\tau, p[\tau, \xi], z[\tau, \xi])$, i.e., we penalized the end of the cycle backlogged demand by $b + c$ per unit. It is easy to see that after the modification our algorithm can be implemented in an identical fashion. We report the mean and standard deviation of the total *profit* generated by CI and DP on the $N = 100$ sample paths. Table 2.22 summarizes the performance measures. Table 2.23 displays the parameters of the second scenario – the "Skirt Scenario" in [38]. We repeated the procedures as in the

first scenario. Table 2.24 summarizes the performance measures. In all cases considered, the average profit of CI policy is approximately 2-2.5% below that of the DP-based policy, and the difference in the average profit of CI policy and the DP-based policy is around half of the confidence interval for the average profit of the DP-based policy.

Table 2.21: Parameters of Scenario 1 – "Dress Scenario".

Parameter	Value	Parameter	Value
Time horizon T	48	Std. dev. of random term σ_t	15
Fixed ordering cost K	500	Intercept of demand curve α_t	174
Variable ordering cost/unit c	22.15	Slope of demand curve β_t	3
Inventory holding cost/unit h	0.22	$p^t_{\max}$	44
Shortage cost/unit b	21.78	$p^t_{\min}$	25
Mean of random term $\bar{z}_t$	0	Capacity M	500

Table 2.22: Performance summary of Scenario 1.

	DP	CI
Distribution	Mean: std dev	Mean(%): std dev
Normal	$40,668.4:\ 2,100.3$	$39,652.3(-2.50\%):\ 2,236.9$
T(4)	$41,123.2:\ 1,646.4$	$40,224.8(-2.18\%):\ 1,746.6$
Gamma	$41,003.5:\ 1,615.3$	$40,028.7(-2.38\%):\ 1,681.2$
Uniform	$40,782.6:\ 1,878.0$	$39,792.9(-2.43\%):\ 1,930.2$
Lognormal	$40,937.1:\ 1,804.9$	$40,013.7(-2.26\%):\ 1,838.7$

The main difference between the two policies is the running time. We display the running time of both policies in Table 2.25. From the table, it is clear that for the two scenarios we considered, CI is approximately 10 times faster than DP. Moreover, DP costs around 50% more time in Scenario 2 than in Scenario 1 because the price range $[15,44]$ contains 50% more grid points than the price range $[25,44]$, whereas the running times of CI in both scenarios are

Table 2.23: Parameters of Scenario 2 – "Skirt Scenario".

Parameter	Value	Parameter	Value
Time horizon T	48	Std. dev. of random term σ_t	6.5
Fixed ordering cost K	500	Intercept of demand curve α_t	57
Variable ordering cost/unit c	14.05	Slope of demand curve β_t	1
Inventory holding cost/unit h	0.17	$p^t_{\max}$	44
Shortage cost/unit b	16.83	$p^t_{\min}$	15
Mean of random term $\bar{z}_t$	0	Capacity M	500

Table 2.24: Performance summary of Scenario 2.

	DP	CI
Distribution	Mean: std dev	Mean(%): std dev
Normal	$19,307.4:\ 874.1$	$18,852.5(-2.36\%):\ 834.2$
T(4)	$19,382.3:\ 1,038.2$	$18,930.1(-2.33\%):\ 1,041.8$
Gamma	$19,200.2:\ 941.2$	$18,765.2(-2.27\%):\ 936.8$
Uniform	$19,086.8:\ 997.7$	$18,652.2(-2.28\%):\ 966.1$
Lognormal	$19,179.8:\ 847.1$	$18,746.4(-2.26\%):\ 850.3$

almost identical. Thus, the running time of the DP-based policy is very sensitive to the size of the parameter ranges, but that the running time of CI is completely insensitive to the size of the parameter sets.

Table 2.25: Running time comparison.

	DP	CI
Scenario	(in seconds)	Mean: std dev (in seconds)
1	72.21	8.15 : 1.20
2	106.45	8.15 : 1.29

2.7 Conclusions

In this paper we have proposed a new robust cycle-based control policy for a single installation, finite horizon inventory model with non-stationary uncertain demand. Our policy is an extension of the EOQ and (R,T)-type policies, and, like these policies, it is structurally simple and the decisions in the policy have an intuitive appeal. Our policy provides a unified framework for efficiently solving inventory control problems under various model assumptions. In particular, our policy can be efficiently implemented both when the excess demand is backlogged or lost, when the fixed ordering cost is non-zero, and when the lead time is non-zero. The policy can also be extended to the joint pricing and inventory control problem with a simple generalization.

The optimal decisions in our policy are computed by a Benders decomposition approach where all the sub-problems are linear programs of modest size, even when the problem contains positive fixed ordering costs. Since our policy only requires LP solvers, it is easily implementable and has a higher likelihood of being adopted in practice. Our numerical experiments showed that the performance of our policy is very close to that of the competing polices while requiring significantly lower computing resources.

In this paper, we have restricted our attention to the single product, single installation inventory models. Just as in the (R,T) policy, the cycle length, or equivalently, the next reorder epoch, is an explicit decision variable in our policy. Consequently, coordinating orders in more complex inventory networks is relatively easy in the context of our policy [66]. We are currently investigating extensions of our robust cycle-based policy to more general supply chain settings, e.g., serial system [33] and one-warehouse-multi-retailer system (OWMR) [74].

Chapter 3

Modeling Multimodal Continuous Heterogeneity in Conjoint Analysis – A Sparse Learning Approach

3.1 Introduction

Marketing researchers and practitioners frequently use conjoint analysis to recover consumers' heterogeneous preferences [42,90], which serve as a critical input for many important marketing decisions, such as market segmentation [88] and differentiated product offerings and pricing [2]. In practice, consumer preferences can often be represented using a multimodal continuous heterogeneity (MCH) distribution, in which the consumer population can be interpreted as consisting of a few distinct segments, each of which contains a heterogeneous sub-population. Since in most conjoint applications the amount of information elicited from each respondent

through conjoint questionnaires is limited, adequate modeling of MCH is critical for accurate conjoint estimation in the presence of MCH.

In marketing, the most widely used approach for modeling MCH is the finite mixture (FM) model [32, 52]. The FM model approximates MCH using discrete mass points, with each mass point representing a segment of homogeneous consumers. Such a discrete representation of the heterogeneity distribution provides a flexible approach for characterizing multimodality and is computationally tractable. However, as within-segment heterogeneity is not allowed for and hence the variations in individual-level partworths are not fully captured, the FM model may not provide an adequate representation of MCH [1].

To better characterize MCH, marketing researchers proposed hierarchical Bayes (HB) models with flexible parametric specifications for the heterogeneity distribution. [1] generalized the FM model by developing a normal component mixture (NCM) model in which a mixture of multivariate normal distributions is utilized to represent the heterogeneity distribution. On the one hand, the NCM model is flexible and is capable of modeling a wide variety of heterogeneity distributions. Using multiple data sets, [1] showed that the NCM model outperforms the FM model. On the other hand, the NCM model may suffer from two limitations. First, the amount of Bayesian shrinkage imposed by the NCM model is influenced by exogenously chosen parameters of the second-stage priors. [37] showed that for unimodal continuous heterogeneity (UCH), the amount of shrinkage imposed by a unimodal HB model is often suboptimal under the standard practice of choosing parameters to induce diffuse second-stage piors, and that endogenously selecting parameters of the second-stage priors to optimize the amount of shrinkage is challenging within the HB framework. We expect the issue of a suboptimal amount of shrinkage to impact

the performance of the NCM model even when modeling MCH. Second, the NCM model faces inferential difficulties when conducting a segment-level analysis. In particular, the NCM model suffers from the overlapping mixtures problem [55] and the label switching problem [28, 80], neither of which has received definitive answers.[1]

In this paper, we aim to provide a new perspective on MCH by proposing and testing a sparse learning (SL) approach to modeling MCH in the context of both metric and choice-based conjoint analysis. Sparse learning [8, 81, 93] is a field of machine learning that aims at learning a *sparse* model that best explains the observed data, and has found successful applications in areas such as econometrics [12], image processing [51], and bioinformatics [82]. In the context of linear models, a sparse model is defined as a vector of parameters of which a large proportion of parameters takes zero value. Since only those explanatory variables associated with nonzero parameters are used to model the dependent variables, learning a sparse model is essentially selecting the most important explanatory variables [10]. To expand the modeling power of the sparse learning methodology beyond unstructured variable selection, a more refined sparsity concept called *structured sparsity* has been recently proposed in the machine learning literature [50, 56]. A structured sparse model could be either a sparse model of which the pattern of the zero parameters satisfies certain structure, or a nonsparse model of which the parameters display sparsity after some linear transformation. We model MCH using structured sparsity in the latter

[1]More recently, nonparametric Bayesian methods have been introduced to marketing. Examples include the Dirichlet process mixture (DPM) model [7, 55] and the centered Dirichlet process mixture (CDPM) model [59]. While nonparametric Bayesian methods provide more flexibility, they may still suffer from the same two limitations faced by the NCM model. With ongoing research in this area, we expect to see systematic comparisons between the benefits of using parametric and nonparametric Bayesian methods. In this paper, we compare our approach with the FM and NCM models, which are more established modeling frameworks. Future research may compare our approach with nonparametric Bayesian models.

sense. The underlying intuition of our SL approach is that any two respondents from the same segment have identical segment-level partworths, i.e., the difference between their respective segment-level partworths is the zero vector, which implies that respondents' segment-level partworths are structured sparse (as the pairwise differences of respondents' segment-level partworths are sparse). As a key step for modeling MCH, our SL approach leverages this intuition to recover candidate segmentations of the consumer population by imposing sparsity on the pairwise differences of respondents' segment-level partworths when learning individual-level and segment-level partworths from the conjoint data.

Our SL approach applies machine learning and optimization techniques to model MCH via a two-stage divide-and-conquer framework. In the first stage, we "divide" MCH by using structured sparsity modeling to recover a set of candidate segmentations of the consumer population, each of which is interpreted as a potential decomposition of the MCH distribution into a small collection of within-segment unimodal continuous heterogeneity (UCH) distributions. In the second stage, we use each candidate segmentation to develop a set of individual-level representations of MCH by "conquering" UCH's, i.e., separately modeling the UCH distribution for each segment of the candidate segmentation. We select the optimal individual-level representation of MCH and the corresponding optimal candidate segmentation using cross-validation [43, 75, 86, 89]. We note that our approach explicitly models both across- and within-segment heterogeneity and is capable of endogenously selecting an adequate amount of shrinkage to recover the individual-level partworths.

Our SL approach adds to the growing literature of machine learning and optimization-based methods for conjoint estimation [34, 36, 37, 83, 84]. This stream of research has largely ignored

the modeling of consumer heterogeneity, with the exception of [37], who proposed a convex optimization (CO) model based on the multitask learning framework to modeling UCH. The authors showed that, by endogenously selecting the amount of shrinkage using cross-validation, their CO model outperforms a standard unimodal HB model. Our work contributes to this literature by developing the first machine learning-based approach to modeling the more general MCH using novel sparse learning techniques. In addition, as will be made clear later, our approach nests the CO model as a special case.

We compare the SL approach to the FM model, the NCM model, and the CO model using extensive simulation experiments and two empirical conjoint data sets. In simulation experiments where the conjoint data is generated according to MCH, the SL approach demonstrates strong performance in terms of both parameter recovery and predictive accuracy on holdout samples. In the empirical conjoint data sets, the SL approach also shows competitive performance in terms of predictive accuracy and the individual-level partworths estimates of the SL approach display shapes consistent with MCH. Therefore, we empirically validate the performance of our SL approach in modeling MCH.

The remainder of the paper is organized as follows. In Section 3.2 we present our SL approach to modeling MCH in both metric and choice-based conjoint analysis. We empirically compare the SL approach and the benchmark methods using simulation experiments in Section 3.3 and two empirical conjoint data sets in Section 3.4. Section 3.5 concludes with our contributions, key results, and a discussion of future work.

3.2 Model

In this section, we present our sparse learning (SL) approach to model multimodal continuous heterogeneity (MCH) in conjoint analysis. We first give a detailed description of our approach in the context of metric conjoint analysis; we then discuss a few modifications needed for choice-based conjoint analysis.

3.2.1 Metric Conjoint Analysis

Conjoint Setup

We assume a total of I consumers (or respondents), each rating J profiles with p attributes. Let the $1 \times p$ row vector x_{ij} represent the j-th profile rated by the i-th respondent, for $i = 1, 2, \ldots, I$ and $j = 1, 2, \ldots, J$, and denote $X_i \triangleq \left[x_{i1}^T, x_{i2}^T, \ldots, x_{iJ}^T\right]^T$ as the $J \times p$ design matrix for the i-th respondent. For respondent i, the $p \times 1$ column vector β_i is used to denote her partworths, and her ratings are contained in the $J \times 1$ column vector $Y_i \triangleq \left(y_{i1}, y_{i2}, \ldots, y_{iJ}\right)^T$. We assume additive utility functions, i.e., $Y_i = X_i\beta_i + \varepsilon_i$, for $i = 1, 2, \ldots, I$, where ε_i denotes the random error. The additive specification of the utility functions is a standard assumption in the conjoint analysis literature [42].

Model Overview

Under MCH, the consumer population can be interpreted as consisting of a few distinct segments of heterogeneous consumers. This interpretation motivates our divide-and-conquer framework for modeling MCH, in which the distinct segments are first recovered, followed by separate

modeling of within-segment heterogeneity for each segment. We propose the following two-stage approach to implement this modeling strategy.

In the first stage, we use structured sparsity modeling to recover a set of candidate segmentations of the consumer population, each of which provides a potential decomposition of the MCH distribution into a small collection of within-segment heterogeneity distributions. As will be made clear in Section 3.2.1, the key insight behind the proposed structured sparsity modeling is that any two respondents from the same segment share identical segment-level partworths, i.e., the difference between their respective segment-level partworths is the zero vector, and therefore candidate segmentations can be recovered by imposing sparsity on the pairwise differences of respondents' segment-level partworths when learning individual-level and segment-level partworths from the conjoint data. One limitation of the structured sparsity modeling is that it produces biased individual-level partworths estimates despite generating informative candidate segmentations; consequently, we retain only the set of candidate segmentations and use the second stage to recover an accurate individual-level representation of MCH.[2]

In the second stage, we leverage each candidate segmentation to develop a set of individual-level representations of MCH. Given a candidate segmentation, we separately model the within-segment heterogeneity distribution of each segment assuming an unimodal continuous heterogeneity (UCH) distribution. UCH is considerably easier to model compared to MCH, and we choose RR-Het, the metric version of the convex optimization (CO) model of [37], to model the within-segment UCH distributions. We then select the optimal individual-level representation

[2]We provide a detailed discussion on this limitation at the end of Section 3.2.1.

of MCH and the corresponding optimal candidate segmentation using cross-validation [43, 75, 86, 89].

The First Stage: Recovering Candidate Segmentations

The first stage of our SL approach aims at learning a set of candidate segmentations to provide decompositions of the MCH distribution. To motivate, we introduce a standard characterization of the data-generating process underlying MCH [4, 5]. The data-generating process selects the number of segments L, the segment-level partworths $\{\widehat{\beta}_l^S\}_{l=1}^L$, and the segment-membership matrix $Q \in \mathbb{R}^{I\times L}$, where $Q_{il} = 1$ if respondent i is assigned to segment l and $Q_{il} = 0$ otherwise. If respondent i belongs to segment l, she receives a copy of segment-level partworths $\beta_i^S = \widehat{\beta}_l^S$ and her individual-level partworths are determined by $\beta_i = \beta_i^S + \xi_i$, where ξ_i denotes the difference between respondent i's segment-level and individual-level partworths, i.e., the within-segment heterogeneity. Let $\widehat{B^S} \triangleq \{\widehat{\beta}_l^S\}_{l=1}^L$, $B^S \triangleq \{\beta_i^S\}_{i=1}^I$, and $B \triangleq \{\beta_i\}_{i=1}^I$.

Assuming the above data-generating process, recovering candidate segmentations can be achieved by learning the set of model parameters $\{L, \widehat{B^S}, Q, B^S, B\}$ from the conjoint data. A closer examination reveals that learning $\{B^S, B\}$ is sufficient, as other model parameters $\{L, \widehat{B^S}, Q\}$ can be easily derived from $\{B^S, B\}$. We note that the following assumptions on the structure of the data-generating process are relevant to learning $\{B^S, B\}$:

A1. The ratings vector Y_i is generated based on β_i, i.e., $Y_i = X_i\beta_i + \varepsilon_i$.

A2. The individual-level partworths β_i is generated based on the segment-level partworths β_i^S, i.e., $\beta_i = \beta_i^S + \xi_i$.

A3. Respondents i and k belong to the same segment if and only if $\beta_i^S - \beta_k^S = 0$.

Within an optimization framework in which $\left\{B^S, B\right\}$ are decision variables, A1 and A2 suggest to learn $\left\{B^S, B\right\}$ by penalizing the discrepancy between Y_i and $X_i\beta_i$ and that between β_i and β_i^S. On the other hand, A3 suggests that for a substantial proportion of $i-k$ pairs, i.e., the $i-k$ pairs from the same segment, the discrepancy between β_i^S and β_k^S is the zero vector, implying that the set of pairwise discrepancies of the true B^S is sparse; consequently, a sparse structure should be imposed on the pairwise discrepancies of B^S when learning $\left\{B^S, B\right\}$. In particular, A3 suggests to learn whether respondents i and k are drawn from the same segment by penalizing the discrepancy between β_i^S and β_k^S using some *sparsity-inducing penalty function*, which we discuss soon.

Motivated by these considerations, we propose the following sparse learning problem, which we term as **Metric-SEG**, to recover candidate segmentations.

$$\begin{aligned} \min \quad & \sum_{i=1}^{I} ||\, Y_i - X_i\beta_i \,||_2^2 + \gamma \sum_{i=1}^{I} (\beta_i - \beta_i^S)^T D^{-1} (\beta_i - \beta_i^S) + \lambda \sum_{1 \le i < k \le I} \theta_{ik} ||\, \beta_i^S - \beta_k^S \,||_2, \\ \text{s.t.} \quad & D \text{ is a positive semidefinite matrix scaled to have trace 1,} \\ & \beta_i, \beta_i^S \in \mathbb{R}^p, \text{ for } i = 1, 2, \ldots, I. \end{aligned} \tag{3.2.1}$$

We note that γ, λ, and $\left\{\theta_{ik}\right\}$ are the regularization parameters for Metric-SEG. The regularization parameters control the relative strength of each penalty term in Metric-SEG and we will discuss their specification later in the section. In Metric-SEG, the first two penalty terms are standard quadratic functions measuring the discrepancy between Y_i and $X_i\beta_i$ and that between β_i and β_i^S, respectively. We note that the matrix D is a decision variable and is related to the covariance

matrix of the partworths within each segment [37]. The third penalty term aims to impose the penalty suggested by A3 and is the key to the formulation of Metric-SEG. In particular, it aims to learn whether respondents i and k belong to the same segment by penalizing the ℓ_2-norm of $\beta_i^S - \beta_k^S$, i.e., $||\beta_i^S - \beta_k^S||_2$, for all $i-k$ pairs. We choose the ℓ_2-norm to measure the discrepancy between β_i^S and β_k^S since, unlike most standard measures of magnitude of vectors, e.g., the sum-of-squares measure, the ℓ_2-norm is a sparsity-inducing penalty function in that it is capable of enforcing *exact* zero value in optimal solutions under a suitable level of penalty.[3] Sparsity-inducing penalty functions play a fundamental role in sparse learning [10, 81, 93]. Our use of the ℓ_2-norm to penalize the pairwise differences of B^S can be viewed as a generalization of the overlapping ℓ_1/ℓ_2-norm [51, 56] and the Fused Lasso penalty [82], and was recently introduced in the context of unsupervised learning [44].

One simple intuition underlies our modeling choice of assessing whether respondents i and k are from the same segment by penalizing the ℓ_2-norm of $\beta_i^S - \beta_k^S$. To illustrate the intuition, we set the regularization parameter $\theta_{ik} = 1$ for all $i-k$ pairs in Metric-SEG and thus homogenize the penalty imposed on $||\beta_i^S - \beta_k^S||_2$'s. For any two respondents i and k, we consider the following

[3]The sparsity-inducing penalty functions are capable of enforcing exact zero value since, compared to non-sparsity-inducing penalty functions, they impose a sharp penalty on parameters in the neighborhoods of zeros. Such a sparsity-inducing characteristic has been demonstrated in the statistics and machine learning literature [81, 93, 96]. We offer a simple intuition by comparing the sparsity-inducing ℓ_2-norm, i.e., $\Omega_1(x) \triangleq ||x||_2$, and the non-sparsity-inducing sum-of-squares, i.e., $\Omega_2(x) \triangleq ||x||_2^2 = x_1^2 + x_2^2 + \ldots + x_p^2$. We focus on a simple one-dimensional case with $p = 1$, in which $\Omega_1(x) = |x|$ and $\Omega_2(x) = x^2$, and consider the following two optimization problems: (P1) $\min_{x\in\mathbb{R}} (x-a)^2 + \lambda\Omega_1(x)$, and (P2) $\min_{x\in\mathbb{R}} (x-a)^2 + \lambda\Omega_2(x)$, where $a > 0$. Simple algebra shows that the solution of (P1) is $x^{1*} = \max(a - \frac{\lambda}{2}, 0)$ and the solution of (P2) is $x^{2*} = \frac{a}{1+\lambda}$; that is, when $\lambda \geq 2a$, $\Omega_1(x)$ induces an exact zero solution, whereas $\Omega_2(x)$ always leads to a positive solution. This contrast is not surprising: the marginal penalty of $\Omega_1(x)$ is constant 1, i.e., $|\frac{d}{dx}\Omega_1(x)| = 1$ for $x \neq 0$, implying that any solution gets rewarded (or less penalized) by one unit as it moves toward zero by one unit; on the contrary, the marginal penalty of $\Omega_2(x)$ diminishes quickly toward the origin, i.e., $\lim_{x\to 0} |\frac{d}{dx}\Omega_2(x)| = 0$. Therefore, in (P1) there is always incentive to move toward zero as long as the marginal penalty of $(x-a)^2$ remains moderate, whereas in (P2) the solution will always get stuck somewhere on the way to zero.

components of the objective function of Metric-SEG:

$$G_{i,k} \triangleq \sum_{r=i,k} ||Y_r - X_r\beta_r||_2^2 + \gamma \sum_{r=i,k} (\beta_r - \beta_r^S)^T D^{-1} (\beta_r - \beta_r^S) + \lambda ||\beta_i^S - \beta_k^S||_2.$$

Within an optimization framework, the three penalty terms in $G_{i,k}$ induce competing shrinkage over the decision variables $\{\beta_r, \beta_r^S\}_{r=i,k}$. In particular, the first penalty term approximately shrinks β_r toward the true individual-level partworths $\beta_r(T)$, and the second penalty term shrinks β_r and β_r^S toward each other, for $r = i,k$, whereas the third penalty term shrinks β_i^S and β_k^S toward each other. We note that whether $\beta_i^S - \beta_k^S = 0$ holds in the optimal solution of Metric-SEG is to a large extent determined by the tradeoff among the three competing shrinkage, which is in turn determined by the distance between the true individual-level partworths $\beta_i(T)$ and $\beta_k(T)$ as well as the regularization parameters γ and λ. Intuitively, if respondents i and k are from the same segment the distance between $\beta_i(T)$ and $\beta_k(T)$ should be relatively small, implying that a moderate penalty imposed on $||\beta_i^S - \beta_k^S||_2$, i.e., a small λ, should be sufficient to enforce $\beta_i^S - \beta_k^S = 0$ due to the sparsity-inducing property of the ℓ_2-norm. On the other hand, if respondents i and k are from distinct segments the distance between $\beta_i(T)$ and $\beta_k(T)$ should be relatively large and therefore enforcing $\beta_i^S - \beta_k^S = 0$ cannot be achieved unless a strong penalty is imposed on $||\beta_i^S - \beta_k^S||_2$, i.e., a large λ is specified.[4] This intuition suggests that if γ and particularly λ are appropriately specified it is possible to recover the underlying segmentation of the consumer population by solving Metric-SEG and identifying $i-k$ pairs with $\beta_i^S - \beta_k^S = 0$ in the optimal solution.

[4]We note that specifying $\lambda = +\infty$ enforces $\beta_i^S - \beta_k^S = 0$ for all $i-k$ pairs, in which case Metric-SEG reduces to the key convex optimization problem (Equation (3)) in [37]. We show this link in the appendix.

In the above discussion, we set $\theta_{ik} = 1$ for all $i-k$ pairs only for the purpose of illustration. When implementing Metric-SEG, a heterogeneous specification for $\{\theta_{ik}\}$ is useful as it allows us to incorporate information that could potentially facilitate the recovery of the underlying segmentation. For example, we can incorporate information suggesting that respondents i and k are more likely to be drawn from the same segment compared to respondents i' and k' in Metric-SEG via a specification of $\{\theta_{ik}\}$ in which $\theta_{ik} > \theta_{i'k'}$, such that a relatively stronger sparsity-inducing penalty is imposed on $\beta_i^S - \beta_k^S$ to enforce $\beta_i^S - \beta_k^S = 0$ in the optimal solution. One way to assess the relative "likelihood" of respondents i and k belonging to the same segment and specify θ_{ik} is as follows:

$$\theta_{ik} = R\big(||\bar{\beta}_i - \bar{\beta}_k||_2\big), \tag{3.2.2}$$

where $\{\bar{\beta}_i\}_{i=1}^I$ are initial estimates of the individual-level partworths, and $R(\cdot)$ is a positive, non-increasing function. The intuition behind this specification for $\{\theta_{ik}\}$ is that, if the distance between the initial individual-level partworths estimates $\bar{\beta}_i$ and $\bar{\beta}_k$ is relatively small, then it is likely that respondents i and k are from the same segment and hence a relatively large regularization parameter θ_{ik} is specified. In Equation (3.2.2), the choices of $\{\bar{\beta}_i\}_{i=1}^I$ and $R(\cdot)$ are quite flexible. We choose to compute $\{\bar{\beta}_i\}_{i=1}^I$ using RR-Het, the metric version of the CO model of [37], which provides an effective approach to recovering consumers' heterogeneous preferences. We also specify $R(x) = e^{-\omega x}$, a simple positive, non-increasing function parameterized by a regularization parameter $\omega \geq 0$. Consequently, we adopt the following specification for $\{\theta_{ik}\}$:

$$\theta_{ik} = e^{-\omega||\bar{\beta}_i - \bar{\beta}_k||_2}, \tag{3.2.3}$$

where $\{\bar{\beta}_i\}_{i=1}^I$ are the initial estimates produced by RR-Het using the conjoint data.[5][6] The regularization parameter ω controls the extent to which $\{\bar{\beta}_i\}_{i=1}^I$ are used to facilitate recovering candidate segmentations. When $\omega = 0$, $\{\bar{\beta}_i\}_{i=1}^I$ do not enter the specification of $\{\theta_{ik}\}$ and a homogeneous penalty is imposed on the pairwise discrepancies of B^S; as ω increases, $\{\theta_{ik}\}$ become more heterogeneous and pairs of respondents with closer initial estimates, i.e., those deemed as more likely to be drawn from the same segment, are penalized more heavily than those with farther initial estimates.

As the specification of $\{\theta_{ik}\}$ reduces to the choice of ω, the regularization parameters of Metric-SEG become the triplet $(\gamma, \lambda, \omega)$, which we denote as $\Gamma \triangleq (\gamma, \lambda, \omega)$. Since an appropriate value for Γ cannot be determined a priori, in the first stage of the SL approach we specify a finite grid $\Theta \subset \mathbb{R}^3$ and for each $\Gamma \in \Theta$ we solve Metric-SEG and recover the candidate segmentation implied by the optimal B^S.[7] Consequently, we obtain a set of candidate segmentations which is the output of the first stage. Given Γ, Metric-SEG is a convex optimization problem, which implies that any of its local optimum is automatically a global optimum and it is efficiently solvable to global optimum in theory [27]. In practice, however, solving Metric-SEG poses a considerable algorithmic challenge since its third penalty term, $\lambda \sum_{1 \le i < k \le I} \theta_{ik} || \beta_i^S - \beta_k^S ||_2$, is both non-differentiable and non-separable, i.e., β_i^S is present in multiple ℓ_2-norms for each i. Non-differentiability implies that standard convex optimization methods requiring a differentiable

[5]The specification for $\{\theta_{ik}\}$ in Equation (3.2.3) uses only information contained in the conjoint data. Other information sources, e.g., consumers' demographic variables, can be readily incorporated in the specification for $\{\theta_{ik}\}$ and hence our SL approach via a simple extension of Equation (3.2.3). We discuss the extension in the appendix.

[6]We note that in Metric-SEG the amount of penalty imposed on $|| \beta_i^S - \beta_k^S ||_2$ is controlled by $\lambda \theta_{ik}$. In the empirical implementation of our SL approach, we normalize $\theta = (\theta_{ik})$ such that $||\theta||_2 = 1$ and interpret the regularization parameter λ as controlling the "total" amount of penalty imposed on $|| \beta_i^S - \beta_k^S ||_2$'s.

[7]The specifications of the finite grid Θ used in the simulation experiments and two empirical applications are summarized in the appendix.

objective function, e.g., the Newton's method, cannot be applied to solve Metric-SEG; non-separability also adds to the complexity [29]. In order to address this challenge, we apply an algorithm based on variable splitting and the Alternating Direction Augmented Lagrangian (ADAL) method that was proposed in [65] to solve Metric-SEG. This algorithm is specifically designed for handling complex sparsity-inducing penalty functions and is capable of solving for the global optimum of Metric-SEG. The details of the algorithm are available from the authors upon request.

After solving Metric-SEG for any given Γ, we aim to recover a candidate segmentation, represented by a binary segment-membership matrix Q, based on the optimal B^S. In the empirical implementation of Metric-SEG, we observe that a typical candidate segmentation Q contains both a few substantive segments which constitute the majority of the consumer population, and a few very small segments each consisting of few respondents, mostly one or two. Since these small segments bear little practical meaning, an appropriate interpretation of Q is needed. We propose a simple procedure to address this issue by combining each of the small segments with its closest substantive segment. Formally, we introduce a pre-specified threshold M, and define a segment in Q as a valid segment if it contains at least M respondents, and as an outlier segment otherwise. Without loss of generality, we assume the first $\bar{L}$ segments of Q are valid. We retain all valid segments, and for each outlier segment, i.e., the l-th segment with $l > \bar{L}$, we determine its closest valid segment by computing $c(l) \triangleq \left\{ v \in \left\{1,2,\ldots,\bar{L}\right\} \mid ||\widehat{\beta}_v^S - \widehat{\beta}_l^S||_2 < ||\widehat{\beta_{v'}^S} - \widehat{\beta}_l^S||_2, \text{ for } v' \in \left\{1,2,\ldots,\bar{L}\right\}, v' \neq v \right\}$, and combine the l-th segment (an outlier segment) and the $c(l)$-th segment (a valid segment). Mathematically, we introduce a new binary segment-membership matrix $\bar{Q} \in \mathbb{R}^{I \times \bar{L}}$, such that $\bar{Q}_{\cdot,v} \triangleq Q_{\cdot,v} + \sum_{l:l>\bar{L},c(l)=v} Q_{\cdot,l}$, where $\bar{Q}_{\cdot,v}$ and $Q_{\cdot,l}$ denote the v-th column of $\bar{Q}$ and the l-th column of Q, respectively. We define $\bar{Q}$, the

segmentation obtained after this processing as the candidate segmentation, but still refer to it using Q for simplicity hereafter.[8] We note it is possible that no valid segment exists in a segmentation, i.e., $\bar{L} = 0$. In such a case, we simply claim that no candidate segmentation is identified for this particular Γ.

In summary, the first stage of our SL approach recovers a set of candidate segmentations in the following manner. We first specify a finite grid $\Theta \subset \mathbb{R}^3$ from which the regularization parameters for Metric-SEG, $\Gamma = \left(\gamma, \lambda, \omega\right)$, are chosen. For each $\Gamma \in \Theta$, we solve Metric-SEG and obtain the candidate segmentation $Q(\Gamma)$. $Q(\Gamma)$ could be an empty matrix in cases where no candidate segmentation is identified. We also include the trivial segmentation in which all respondents are included in a single segment as a candidate segmentation, i.e., $Q(\text{Trivial}) \triangleq 1_{I\times 1}$. We denote the set of candidate segmentations as Φ, that is, $\Phi \triangleq \left\{Q(\Gamma)\right\}_{\Gamma\in\Theta: Q(\Gamma)\neq\emptyset} \bigcup \left\{Q(\text{Trivial})\right\}$. Φ is the output of the first stage of the SL approach.

Before concluding our presentation of the first stage of the SL approach, it is worthwhile to discuss the rationale behind our modeling choice of retaining only the set of candidate segmentations Φ and excluding the set of individual-level partworths estimates $\left\{B(\Gamma)\right\}$ obtained by solving Metric-SEG as the output of the first stage. Recall that in Metric-SEG we use the penalty term $\lambda \sum_{1\leq i<k\leq I} \theta_{ik} || \beta_i^S - \beta_k^S ||_2$ to impose sparsity on the pairwise discrepancies of B^S so as to identify the candidate segmentations. Given that a priori any pair of respondents could

[8] In the empirical implementation of our SL approach, we set $M = 10\%I$, such that any valid segment contains a non-negligible portion of the population. Note that other choices of M are also possible. Some readers may find pre-specifying a value for M ad-hoc; in such case, we may vary M in $\left\{1, 2, \ldots, I\right\}$ for any given Q and obtain a set of different $\bar{Q}$. This procedure leads to a more extensive set of candidate segmentations, which, as we will see in Section 3.2.1, does not negatively impact the performance of our SL approach, but significantly increases the computational demand. Consequently, we adopt the more efficient approach of pre-specifying M. The simulation experiments and empirical applications confirm the effectiveness of our choice of M.

be drawn from the same segment, we need to penalize $||\beta_i^S - \beta_k^S||_2$ for all $i-k$ pairs, including those which truly belong to distinct segments. Such across-segment shrinkage, i.e., penalty imposed on $||\beta_i^S - \beta_k^S||_2$ where respondents i and k are from distinct segments, is unlikely to enforce $\beta_i^S - \beta_k^S = 0$ in the optimal solution as we previously discussed; however, it has the undesirable consequence of biasing the segment-level partworths estimates B^S as well as the individual-level partworths estimates B.[9] Such bias suggests that $\{B(\Gamma)\}$ fail to accurately characterize MCH, but the recovered candidate segmentations $\{Q(\Gamma)\}$ remain informative as $Q(\Gamma)$ depends only on the pairwise discrepancies of $B^S(\Gamma)$ and thus is largely immune to the bias in the absolute location of $B^S(\Gamma)$. Consequently, our strategy is to use the set of candidate segmentations Φ to develop an accurate individual-level representation of MCH, which is the focus of the second stage of the SL approach.

The Second Stage: Recovering Individual-level Partworths

As we discussed in Section 3.2.1, the first stage of the SL approach generates a set of good candidate segmentations but not accurate individual-level partworths estimates. The second stage of the SL approach aims at leveraging the set of candidate segmentations to accurately recover the individual-level partworths. Toward this end, we develop a set of individual-level representations of MCH based on each candidate segmentation, and select the optimal individual-level representation of MCH and the corresponding optimal candidate segmentation using cross-validation.

Given any candidate segmentation $Q \in \Phi$, it is intuitive to represent MCH by separately modeling the heterogeneity distribution for each segment assuming an unimodal continuous

[9] We provide empirical evidence for the bias in $\{B(\Gamma)\}$ using simulation studies and report the results in the appendix.

heterogeneity (UCH) distribution. That is, the candidate segmentation Q is interpreted as a decomposition of the MCH distribution into a small collection of UCH distributions that are considerably easier to model. In the marketing literature, numerous effective approaches for modeling UCH have been proposed, including the unimodal hierarchical Bayes (HB) models [57, 68] and RR-Het, the metric version of the CO model of [37]. We choose RR-Het as a building block for the second stage of the SL approach to model within-segment heterogeneity distributions since (1) RR-Het overall outperforms standard unimodal HB models [37] and (2) RR-Het can be readily incorporated in the cross-validation framework which will be introduced soon.

Formally, for any candidate segmentation Q with L segments, we denote $\Upsilon(Q;l) \triangleq \{i : Q_{il} = 1\}$ as the l-th segment of Q, and define a set of *modeling strategies* $S \triangleq \Big(Q, \psi, \text{COV}\Big)$, which are parameterized by $\psi = \big(\psi^1, \psi^2, \ldots, \psi^L\big)$ and $\text{COV} \in \big\{\text{General } (G), \text{Restricted } (R)\big\}$. When $\text{COV} = G$, the modeling strategy S obtains the individual-level partworths estimates $\big\{\widetilde{\beta}_i\big\}_{i=1}^{I}$ by solving L convex optimization problems $\Big\{\textbf{Metric-HET-General}(Q;l;\psi^l)\Big\}_{l=1}^{L}$, with the l-th optimization problem Metric-HET-General$(Q;l;\psi^l)$ presented as follows:

$$\begin{aligned} \min \quad & \sum_{i\in\Upsilon(Q;l)} ||\, Y_i - X_i\widetilde{\beta}_i \,||_2^2 + \psi^l \sum_{i\in\Upsilon(Q;l)} \big(\widetilde{\beta}_i - \widetilde{\beta}_0^l\big)^T (D^l)^{-1} \big(\widetilde{\beta}_i - \widetilde{\beta}_0^l\big), \\ \text{s.t.} \quad & D^l \text{ is a positive semidefinite matrix scaled to have trace 1,} \\ & \widetilde{\beta}_i \in \mathbb{R}^p, \text{ for } i \in \Upsilon(Q;l); \quad \widetilde{\beta}_0^l \in \mathbb{R}^p. \end{aligned} \tag{3.2.4}$$

In Metric-HET-General, the regularization parameter ψ^l controls the tradeoff between fit and shrinkage, and the matrix D^l is related to the covariance matrix of the partworths within the

l-th segment [37]. Explicitly modeling D^l and allowing for a general covariance structure gives rise to much flexibility in modeling within-segment heterogeneity; however, it may also lead to overfitting especially when the number of respondents within each segment is moderate. Therefore, it could be beneficial to consider modeling strategies that model within-segment heterogeneity distributions using a restricted version of Metric-HET-General. Formally, when COV $= R$, the modeling strategy S obtains the individual-level partworths estimates $\left\{\widetilde{\beta}_i\right\}_{i=1}^{I}$ by solving L convex optimization problems $\left\{\textbf{Metric-HET-Restricted}(Q;l;\psi^l)\right\}_{l=1}^{L}$, with the l-th optimization problem Metric-HET-Restricted$(Q;l;\psi^l)$ presented as follows:

$$\begin{aligned} \min \quad & \sum_{i\in\Upsilon(Q;l)} ||\, Y_i - X_i\widetilde{\beta}_i\,||_2^2 + \psi^l \sum_{i\in\Upsilon(Q;l)} \left(\widetilde{\beta}_i - \widetilde{\beta}_0^l\right)^T (I/p)^{-1} \left(\widetilde{\beta}_i - \widetilde{\beta}_0^l\right), \\ \text{s.t.} \quad & \widetilde{\beta}_i \in \mathbb{R}^p, \text{ for } i \in \Upsilon(Q;l); \quad \widetilde{\beta}_0^l \in \mathbb{R}^p. \end{aligned} \tag{3.2.5}$$

Metric-HET-Restricted is obtained from Metric-HET-General by restricting $D^l = I/p$. Compared to Metric-HET-General, Metric-HET-Restricted has a more restrictive covariance structure and is less flexible in modeling within-segment heterogeneity distributions; on the other hand, it is more parsimonious and hence more robust with respect to overfitting. The tradeoff between Metric-HET-General and Metric-HET-Restricted will be made based on assessments of their respective capabilities of accurately modeling within-segment heterogeneity distributions using cross-validation.

We note that each modeling strategy S gives rise to a distinct individual-level representation of MCH. In order to select the optimal modeling strategy S (i.e., the optimal individual-level representation of MCH) and its corresponding Q, we evaluate the *cross-validation error* of each

modeling strategy S. Cross-validation is a standard technique used in the statistics and machine learning literature for model selection [43, 75, 86, 89], and has been adopted in the recent literature of machine learning and optimization-based methods for conjoint estimation [36, 37]. The cross-validation error of a modeling strategy $S = \left(Q, \psi, \text{COV}\right)$, $CV(S)$, is measured as follows (the presentation here closely follows [37]):

(1) Set $CV(S) = 0$.

(2) For $j = 1$ to J:

(a) Divide the conjoint data $\left\{X_i, Y_i\right\}_{i=1}^{I}$ into two disjoint subsets, $Z^{(-j)} \triangleq \left\{X_i^{(-j)}, Y_i^{(-j)}\right\}_{i=1}^{I}$, and $Z^{(j)} \triangleq \left\{x_{ij}, y_{ij}\right\}_{i=1}^{I}$, where $X_i^{(-j)} = \left[x_{i1}^T, x_{i2}^T, \ldots, x_{i(j-1)}^T, x_{i(j+1)}^T, \ldots, x_{iJ}^T\right]^T$ and $Y_i^{(-j)} = \left(y_{i1}, y_{i2}, \ldots, y_{i(j-1)}, y_{i(j+1)}, \ldots, y_{iJ}\right)^T$. $Z^{(-j)}$ is deemed as the "calibration" set containing all conjoint data except the j-th profile for each respondent, and the "holdout" set $Z^{(j)}$ consists of the j-th profile for each respondent.

(b) Apply the modeling strategy S to the calibration set $Z^{(-j)}$. That is, estimate the individual-level partworths $\left\{\beta_i^{(-j)}\right\}_{i=1}^{I}$ by solving $\left\{\text{Metric-HET-General}(Q; l; \psi^l)\right\}_{l=1}^{L}$ defined in Equation (3.2.4) or $\left\{\text{Metric-HET-Restricted}(Q; l; \psi^l)\right\}_{l=1}^{L}$ defined in Equation (3.2.5), depending on the value of COV, with X_i and Y_i substituted by $X_i^{(-j)}$ and $Y_i^{(-j)}$, respectively.

(c) Use the estimated individual-level partworths $\left\{\beta_i^{(-j)}\right\}_{i=1}^{I}$ to compute the ratings on $\left\{x_{ij}\right\}_{i=1}^{I}$, and let $\Delta(j)$ be the sum of squared differences between the estimated and observed ratings for the I holdout profiles in $Z^{(j)}$.

(d) Set $CV(S) = CV(S) + \Delta(j)$.

The cross-validation error $CV(S)$ provides an effective estimate of the predictive accuracy of the modeling strategy S on *out-of-sample* data using only *in-sample* data, i.e., the data available to the researcher for model calibration. To implement cross-validation, we pre-specify a finite grid $\Xi \subset \mathbb{R}$, and for each candidate segmentation Q, we consider modeling strategies $S = \left(Q, \psi, \text{COV}\right)$ such that $\psi^l \in \Xi$, for $l = 1, 2, \ldots, L$, and $\text{COV} \in \{G, R\}$.[10] We select S that minimizes the cross-validation error $CV(S)$ as the optimal modeling strategy and its corresponding Q as the optimal candidate segmentation, and denote them as S^* and Q^*, respectively. We recover the optimal individual-level partworths estimates $\left\{\widetilde{\beta}_i^*\right\}_{i=1}^I$ by applying S^* to the complete data set $\left\{X_i, Y_i\right\}_{i=1}^I$. We note that both across- and within-segment heterogeneity are explicitly modeled in our approach and that the use of cross-validation allows us to endogenously choose the amount of shrinkage imposed on each segment to recover the individual-level partworths.

One interesting observation is that, when the set of candidate segmentations Φ contains only the trivial segmentation $Q(\text{Trivial})$ and we set $\text{COV} = G$, the complete SL approach consisting of both stages reduces to the complete RR-Het approach of [37]. In fact, given that the set of modeling strategies considered for cross-validation in our approach subsumes that in RR-Het, we expect the predicative performance of the former on out-of-sample data to be at least as good as that of the latter.

Summary

We briefly summarize our SL approach **Metric-SL** in the following.

[10] The specifications of the finite grid Ξ used in the simulation experiments and two empirical applications are summarized in the appendix.

The First Stage.

Step 1a. Obtain initial estimates $\{\bar{\beta}_i\}_{i=1}^{I}$ using RR-Het [37].

Step 1b. For each $\Gamma \in \Theta$, set $\theta_{ik} = e^{-\omega||\bar{\beta}_i - \bar{\beta}_k||_2}$, and solve Metric-SEG (Equation (3.2.1)). Recover the candidate segmentation $Q(\Gamma)$ from the optimal B^S.

Step 1c. Repeat Step 1b for each $\Gamma \in \Theta$, and obtain the set of candidate segmentations:

$$\Phi = \{Q(\Gamma)\}_{\Gamma \in \Theta: Q(\Gamma) \neq \emptyset} \bigcup \{Q(\text{Trivial})\}. \tag{3.2.6}$$

The Second Stage.

Step 2a. For each $Q \in \Phi$, define a set of modeling strategies $\Big\{S \mid S = \Big(Q, \psi, \text{COV}\Big)$ s.t. $\psi^l \in \Xi$, for $l = 1, 2, \ldots, L$, $\text{COV} \in \{G, R\}\Big\}$. A modeling strategy S recovers the individual-level partworths by solving a set of L optimization problems $\Big\{\text{Metric-HET-General}(Q; l; \psi^l)\Big\}_{l=1}^{L}$ defined in Equation (3.2.4) when $\text{COV} = G$, or $\Big\{\text{Metric-HET-Restricted}(Q; l; \psi^l)\Big\}_{l=1}^{L}$ defined in Equation (3.2.5) otherwise.

Step 2b. Select the modeling strategy $S^* = \Big(Q^*, \psi^*, \text{COV}^*\Big)$ with the minimum cross-validation error, i.e., $S^* = \underset{S}{\text{argmin}}\, CV(S)$. Q^* is selected as the optimal candidate segmentation.

Step 2c. Recover the optimal individual-level partworths $\{\widetilde{\beta}_i^*\}_{i=1}^{I}$ by applying S^* to $\{X_i, Y_i\}_{i=1}^{I}$, i.e., by solving the set of L^* optimization problems $\Big\{\text{Metric-HET-General}(Q^*; l; \psi^{l*})\Big\}_{l=1}^{L^*}$ when $\text{COV}^* = G$, or $\Big\{\text{Metric-HET-Restricted}(Q^*; l; \psi^{l*})\Big\}_{l=1}^{L^*}$ when $\text{COV}^* = R$. The

outputs of the second stage are $\left(\{\widetilde{\beta}_i^*\}_{i=1}^I, Q^*\right)$, which are also the final outputs of the complete Metric-SL approach.

3.2.2 Choice-based Conjoint Analysis

In this section, we show that our SL approach applies to the choice-based conjoint analysis by simply replacing the squared-error loss functions in all optimization problems in Metric-SL with the logistic loss functions. Choice-based conjoint analysis (CBC) has been the dominant conjoint approach recently [48]. In CBC, we assume a total of I respondents, each making choice decisions over J choice sets. In each choice set, there are H conjoint profiles with p attributes. We use the $1 \times p$ row vector x_{ijh} to represent the h-th profile in respondent i's j-th choice set, i.e., the j-th task of respondent i is to choose her most preferred profile from the set $\{x_{ijh}\}_{h=1}^H$. We use the $p \times 1$ column vector β_i to denote the partworths of respondent i, and assume that she chooses the profile x_{ijh*} from her j-th choice set $\{x_{ijh}\}_{h=1}^H$ using a standard logit model, i.e., $U_{ijh*} = \max\{U_{ijh}\}_{h=1}^H$, where $U_{ijh} = x_{ijh}\beta_i + \varepsilon_{ijh}$ and $\{\varepsilon_{ijh}\}$ is a set of independently and identically distributed type-1 extreme value random variables [85]. We summarize our approach in the context of CBC, which we term as **Choice-SL**, in the following.

The First Stage.

Step 1a. Obtain initial estimates $\{\bar{\beta}_i\}_{i=1}^I$ using LOG-Het [37].

Step 1b. For each $\Gamma \in \Theta$, set $\theta_{ik} = e^{-\omega||\bar{\beta}_i - \bar{\beta}_k||_2}$, and solve **Choice-SEG** presented as follows:

$$\begin{aligned} \min \quad & -\sum_{i=1}^{I}\sum_{j=1}^{J} \log \frac{e^{x_{ijh*}\beta_i}}{\sum_{h=1}^{H} e^{x_{ijh}\beta_i}} + \gamma \sum_{i=1}^{I} (\beta_i - \beta_i^S)^T D^{-1} (\beta_i - \beta_i^S) + \lambda \sum_{1 \le i < k \le I} \theta_{ik} || \beta_i^S - \beta_k^S ||_2, \\ \text{s.t.} \quad & D \text{ is a positive semidefinite matrix scaled to have trace 1,} \\ & \beta_i, \beta_i^S \in \mathbb{R}^p, \text{ for } i = 1, 2, \ldots, I. \end{aligned} \tag{3.2.7}$$

Note that Choice-SEG is obtained from Metric-SEG by replacing the squared-error loss with the logistic loss. We recover the candidate segmentation $Q(\Gamma)$ from the optimal B^S.

Step 1c. Repeat Step 1b for each $\Gamma \in \Theta$, and obtain the set of candidate segmentations Φ defined in Equation (3.2.6).

The Second Stage.

Step 2a. For each $Q \in \Phi$, define a set of modeling strategies $\left\{ S \mid S = \left(Q, \psi, \text{COV} \right) \text{ s.t. } \psi^l \in \Xi, \text{ for } l = 1, 2, \ldots, L, \text{ COV} \in \{G, R\} \right\}$. A modeling strategy S recovers the individual-level partworths by solving a set of L optimization problems $\left\{ \text{Choice-HET-General}(Q; l; \psi^l) \right\}_{l=1}^{L}$ when COV $= G$, where Choice-HET-General$(Q; l; \psi^l)$ is defined as follows:

$$\begin{aligned} \min \quad & -\sum_{i \in \Upsilon(Q;l)} \sum_{j=1}^{J} \log \frac{e^{x_{ijh*}\widetilde{\beta}_i}}{\sum_{h=1}^{H} e^{x_{ijh}\widetilde{\beta}_i}} + \psi^l \sum_{i \in \Upsilon(Q;l)} (\widetilde{\beta}_i - \widetilde{\beta}_0^l)^T (D^l)^{-1} (\widetilde{\beta}_i - \widetilde{\beta}_0^l), \\ \text{s.t.} \quad & D^l \text{ is a positive semidefinite matrix scaled to have trace 1,} \\ & \widetilde{\beta}_i \in \mathbb{R}^p, \text{ for } i \in \Upsilon(Q;l); \quad \widetilde{\beta}_0^l \in \mathbb{R}^p, \end{aligned} \tag{3.2.8}$$

or a set of L optimization problems $\left\{ \text{Choice-HET-Restricted}(Q; l; \psi^l) \right\}_{l=1}^{L}$ otherwise,

where Choice-HET-Restricted$(Q;l;\psi^l)$ is defined as follows:

$$\begin{aligned} \min \quad & -\sum_{i\in\Upsilon(Q;l)}\sum_{j=1}^{J}\log\frac{e^{x_{ijh*}\widetilde{\beta}_i}}{\sum_{h=1}^{H}e^{x_{ijh}\widetilde{\beta}_i}}+\psi^l\sum_{i\in\Upsilon(Q;l)}\left(\widetilde{\beta}_i-\widetilde{\beta}_0^l\right)^T(I/p)^{-1}\left(\widetilde{\beta}_i-\widetilde{\beta}_0^l\right), \\ \text{s.t.} \quad & \widetilde{\beta}_i\in\mathbb{R}^p, \text{ for } i\in\Upsilon(Q;l); \quad \widetilde{\beta}_0^l\in\mathbb{R}^p. \end{aligned} \tag{3.2.9}$$

Step 2b. Select the modeling strategy $S^*=\left(Q^*,\psi^*,\text{COV}^*\right)$ with the minimum cross-validation error, i.e., $S^*=\underset{S}{\text{argmin}}\,CV(S)$. Here the cross-validation error is measured by the logistic loss identically defined as that in [37]. Q^* is selected as the optimal candidate segmentation.

Step 2c. Recover the optimal individual-level partworths $\left\{\widetilde{\beta}_i^*\right\}_{i=1}^I$ by applying S^* to $\left\{\left\{x_{ijh}\right\}_{h=1}^H,x_{ijh*}\right\}_{i,j}$, i.e., by solving the set of L^* optimization problems $\left\{\text{Choice-HET-General}(Q^*;l;\psi^{l*})\right\}_{l=1}^{L^*}$ when $\text{COV}^*=G$, or $\left\{\text{Choice-HET-Restricted}(Q^*;l;\psi^{l*})\right\}_{l=1}^{L^*}$ when $\text{COV}^*=R$. The outputs of the second stage are $\left(\left\{\widetilde{\beta}_i^*\right\}_{i=1}^I,Q^*\right)$, which are also the final outputs of the complete Choice-SL approach.

3.3 Simulation Experiments

In this section, we test the empirical performance of our sparse learning (SL) approach via an extensive simulation study. Simulation experiments have been widely adopted in the marketing literature to evaluate conjoint estimation methods [5, 88]. One major advantage offered by simulation experiments over real-world conjoint applications is that in simulation experiments researchers have complete control of the data-generating process and therefore are able to explore various experimental domains of interest. In addition, the true individual-level partworths

are known so that the difference between actual and estimated partworths can be computed. In the following, we consider both metric and choice-based conjoint simulation experiments.

3.3.1 Metric Conjoint Simulation Experiments

We compared the metric version of our SL approach, Metric-SL, to three benchmark methods: (1) the finite mixture (FM) model [32, 52], (2) the Bayesian normal component mixture (NCM) model [1], and (3) RR-Het, the metric version of the convex optimization (CO) model of [37]. The FM model represents multimodal continuous heterogeneity (MCH) using discrete mass points and its effectiveness in modeling MCH has been well documented in the marketing literature [4, 5]. The NCM model specifies a mixture of multivariate normal distributions to characterize the population distribution and is capable of representing a wide variety of heterogeneity distributions. It is shown that the NCM model is more flexible than and empirically ourperforms the FM model [1]. On the other hand, RR-Het is not specifically designed to model MCH; however, we included it as a benchmark method since it has been shown to outperform a standard unimodal hierarchical Bayes (HB) model [37] and is nested within Metric-SL and hence it might be of interest to assess the improvement made by adopting a more general and sophisticated method.

The implementation of the three benchmark methods closely followed the extant literature. In particular, the FM model was calibrated using the Bayesian information criterion (BIC) [5], and for the NCM model the number of components was selected using the deviance information criterion (DIC) [60, 79]. We provide the setup of the NCM model including the specification of parameters for the second-stage priors in the appendix.

Data

In our metric conjoint simulation experiments, both the experimental design and the data-generating process closely followed those in [5].

Experimental Design. We experimentally manipulated the following four data characteristics:

Factor 1. The number of segments: 2 or 3;

Factor 2. The number of profiles per respondent (for calibration): 18 or 27;

Factor 3. The within-segment variances of distributions: 0.05, 0.10, 0.20, 0.40, 0.60, 0.80 or 1.00;[11]

Factor 4. The error variance: 0.5 or 1.5.

Hence, we used a $2^3 \times 7$ design, resulting in a total of 56 experimental conditions. We randomly generated 5 data sets for each experimental condition and estimated all conjoint models separately on each data set.

The Data-generating Process. The conjoint designs were taken from [5], which vary six product attributes at three levels each.[12] The responses of 100 synthetic respondents were generated in each data set according to the following three-step process: (1) we generated the true segment-level partworths, (2) assigned each respondent to a segment and generated her true individual-level partworths, and (3) generated her response vector. More specifically, the

[11] The range of within-segment variances is wider than that used in [5] and ensures a thorough understanding of the impact of within-segment heterogeneity.

[12] We thank Rick Andrews for kindly sharing the conjoint designs with us.

true segment-level partworths for each segment were uniformly randomly generated to be in the range of -1.7 to 1.7; that is, if we denote $\phi_{\ell t}$ as the true partworth of the t-th attribute/level for the ℓ-th segment, then $\phi_{\ell t}$'s are i.i.d. random variables with the distribution $U[-1.7, 1.7]$. Each respondent was randomly assigned to all segments with equal probabilities, and her true individual-level partworths β_i was generated by adding a normal random vector with mean 0 and a pre-specified variance (Factor 3) to her true segment-level partworths. Given β_i, the response vector Y_i was computed as $Y_i = X\beta_i + \varepsilon_i$, where ε_i is a normal random vector with mean 0 and a pre-specified variance (Factor 4). In order to evaluate the predictive accuracy of the conjoint estimation methods, we generated 8 holdout profiles for each respondent regardless of whether 18 or 27 profiles (Factor 2) were used for calibration.

Results

We compared the four conjoint estimation methods in terms of parameter recovery and predictive accuracy. Parameter recovery was assessed using the root mean square error between the true individual-level partworths, $\beta_i(T)$, and the estimated individual-level partworths, $\beta_i(E)$, which we denote as RMSE(β). For respondent i, the RMSE(β) was computed as follows:

$$\text{RMSE}(\beta) = \sqrt{\frac{1}{p}||\beta_i(E) - \beta_i(T)||_2^2}. \tag{3.3.1}$$

Following [37], we computed the RMSE(β) for each respondent and summarized the average RMSE(β) across respondents for each data set. We report the mean of the average RMSE(β)'s across 5 data sets in each experimental condition. Predictive accuracy was measured using the

root mean square error between the observed ratings, $Y_i(O)$, and the predicted ratings, $Y_i(P)$, for the holdout sample, which we denote as RMSE(Y). For respondent i, the RMSE(Y) was computed as follows:

$$\text{RMSE}(Y) = \sqrt{\frac{1}{\tilde{J}}||Y_i(P) - Y_i(O)||_2^2}, \tag{3.3.2}$$

where $\tilde{J}$ denotes the number of holdout profiles for each respondent, i.e., $\tilde{J} = 8$ in our study. Again, we computed the RMSE(Y) for each respondent and summarized the average RMSE(Y) across respondents for each data set; we report the mean of the average RMSE(Y)'s across 5 data sets in each experimental condition. In addition to parameter recovery and predictive accuracy, we also compared the computation time of Metric-SL and the NCM model and report the results in the appendix.

Tables 3.1 and 3.2 summarize the results, where Num-S denotes the number of segments in the population distribution (Factor 1), Num-P denotes the number of profiles per respondent for calibration (Factor 2), EV denotes the error variance (Factor 4), and WSV denotes the within-segment variances of distributions (Factor 3).

Insert Tables 3.1 and 3.2 here.

Table 3.1 reports the mean of the average RMSE(β)'s for each of the four conjoint estimation methods over 5 randomly simulated data sets for each experimental condition. Metric-SL demonstrates an advantage in terms of parameter recovery, performing best or not significantly different from best (at $p < 0.05$) in 53 conditions. The comparisons are based on paired t-tests over the same 500 respondents, i.e., 100 respondents per data set $\times$ 5 data sets, in each experimental condition. The results of metric conjoint simulation experiments highlight the

importance of (1) explicitly modeling both across- and within-segment heterogeneity and (2) endogenously selecting the amount of shrinkage to recover the individual-level partworths in modeling MCH. Recall that in our SL approach, we develop a divide-and-conquer framework to decompose MCH so as to model both across- and within-segment heterogeneity, and use cross-validation to select the modeling strategy, i.e., the candidate segmentation, and the amount of shrinkage and the covariance structure (General versus Restricted) imposed on each segment, to estimate the individual-level partworths. On the other hand, none of the three benchmark methods takes both factors into account jointly: the FM model assumes a discrete heterogeneity distribution and does not allow for within-segment heterogeneity; RR-Het assumes an unimodal continuous heterogeneity (UCH) distribution and does not model across-segment heterogeneity; the NCM model, despite modeling both across- and within-segment heterogeneity, is not capable of endogenously selecting the amount of shrinkage as it is a function of the exogenously chosen parameters for the second-stage priors.

A closer examination of Table 3.1 reveals a systematic performance pattern of the four conjoint estimation methods with respect to WSV. In experimental conditions where WSV is small, e.g., WSV $= 0.05$, the NCM model and RR-Het perform substantially worse than Metric-SL whereas the FM model shows a strong performance; as WSV increases the performance gaps between Metric-SL, the NCM model, and RR-Het gradually shrink and the FM model quickly deteriorates. When WSV is large, e.g., WSV $= 0.80$ or 1.00, the RMSE(β)'s of Metric-SL, the NCM model, and RR-Het are very close in magnitude. This performance pattern is consistent with our intuition. In particular, the assumption of UCH made by RR-Het is more restrictive when the underlying heterogeneity distribution is of a more discrete shape. On the other

hand, the discrete distribution assumed by the FM model is not capable of fully capturing the variations in consumer preferences when within-segment heterogeneity is substantial. It also suggests that the amount of shrinkage imposed by the NCM model could be suboptimal when within-segment heterogeneity is moderate. In addition, it is evident from Table 3.1 that a larger number of profiles for calibration (Num-P) improves the performance of all models and a larger error variance (EV) worsens the performance of all models. In particular, the NCM model and RR-Het benefit (suffer) most substantially from more profiles (larger error variance), whereas the FM model and Metric-SL appear to be less sensitive with respect to these two factors. Moreover, the impact of the number of true segments (Num-S) is not substantial.

Qualitatively similar results are found in Table 3.2, which shows the comparison of the four conjoint estimation methods on the measure of predictive accuracy, RMSE(Y). Metric-SL performs best or not significantly different from best (at $p < 0.05$) in 53 conditions.

3.3.2 Choice-based Conjoint Simulation Experiments

We compared the choice version of our SL approach, Choice-SL, to three benchmark methods: (1) the FM model [32, 52], (2) the NCM model [1], and (3) LOG-Het, the choice version of the CO model of [37]. All benchmark methods were choice versions of those in Section 3.3.1 and the implementations were similar to their metric version counterparts.

Data

In our choice-based conjoint simulation experiments, both the experimental design and the data-generating process closely followed those in [4] and [6].

Experimental Design. We experimentally manipulated the following four data characteristics:

Factor 1. The number of segments: 2 or 3;

Factor 2. The number of choice sets per respondent (for calibration): 16 or 24;[13]

Factor 3. The within-segment variances of distributions: 0.05, 0.10, 0.20, 0.40, 0.60, 0.80 or 1.00;[14]

Factor 4. The error variance: standard (1.645) or high (3.290).

Hence, we used a $2^3 \times 7$ design, resulting in a total of 56 experimental conditions. We randomly generated 5 data sets for each experimental condition and estimated all conjoint models separately on each data set.

The Data-generating Process. In all data sets, each choice set consists of four conjoint profiles, each of which associates with a distinct brand. In addition to the three (i.e., $4-1=3$) brand dummies, the attributes also include one continuous variable and two binary variables. We created four levels for the continuous variable, each of them being a range: "low" $\big[-1.3,-0.65\big]$, "medium-low" $\big[-0.65,0\big]$, "medium-high" $\big[0,0.65\big]$, and "high" $\big[0.65,1.3\big]$. For each choice set we randomly selected a value from each range and assigned the four values to the profiles such that each profile had an equal chance to be assigned with the lowest value.[15] To

[13]The numbers of choice sets per respondent, i.e., 16 and 24, are larger than those in [4] and [6], which considered scanner panel applications, but are consistent with recent choice-based conjoint studies [48].

[14]The within-segment variances are larger than those in [4], and was motivated by [1] who observed that the magnitude of within-segment heterogeneity is typically substantial. It also ensures a thorough understanding of the impact of within-segment heterogeneity.

[15]The purpose of this design was to generate sufficient variations in the data. A similar design was adopted in [48].

generate the two binary variables, for each choice set we randomly selected two profiles, h_1 and h_2, and for the k-th binary variable, only h_k was set to have the value of 1, for $k = 1,2$. We note that the design of the continuous variable and the two binary variables was aimed at inducing sufficient variations in the data and is different from those in [4] and [6], which considered scanner panel applications rather than conjoint applications.

The choices of 100 synthetic respondents were generated in each data set using a three-step process similar to that in Section 3.3.1: (1) we generated the true segment-level partworths, (2) assigned each respondent to a segment and generated her true individual-level partworths, and (3) generated her choices. In order to produce well-separated segment-level partworths, we closely followed [4] and [6] and generated three levels of coefficients (low, medium, and high) for each of the six attributes (i.e., three brand dummies, one continuous variable, and two binary variables). As will soon become clear, the rationale behind the design of three levels of coefficients was to have different segments assigned with different levels of coefficients for each attribute and therefore create clear separations between segments. The medium-level coefficients were generated as follows: the brand-specific constants were generated uniformly from $\left[-1,1\right]$, the coefficient of the continuous variable was generated uniformly from $\left[-2.5,-2\right]$, and the coefficients of the binary variables were generated uniformly from $\left[2,2.5\right]$. The high-level (low-level) coefficients were generated by adding to (subtracting from) the corresponding medium-level coefficients a normal random variable drawn from $N(1.5,0.15^2)$, where 1.5 is the mean separation between segments [4,6]. In experimental conditions with three segments (Factor 1), we generated the true segment-level partworths by randomly assigning the three levels of coefficients of each attribute to all three segments; on the other hand, in

experimental conditions with two segments, we simply retained the true segment-level partworths of the first two segments generated in the three-segment conditions.

After obtaining the true segment-level partworths, we randomly assigned each respondent to all segments with equal probabilities, and generated her true individual-level partworths β_i by adding a normal random vector with mean 0 and a pre-specified variance (Factor 3) to her true segment-level partworths. Given β_i, respondent i's choices were stochastically generated according to the logit probabilities where the variance of the type-I extreme value random variables was a pre-specified value (Factor 4). In order to evaluate the predictive accuracy of the conjoint estimation methods, we generated 8 holdout choice sets for each respondent regardless of whether 16 or 24 choice sets (Factor 2) were used for calibration.

Results

We compared the four conjoint estimation methods in terms of parameter recovery and predictive accuracy. Parameter recovery was assessed using RMSE(β) defined in Equation (3.3.1), with a modification needed for experimental conditions with Factor 4 taking the "high" level, i.e., when the choice data was generated according to the logit model with the variance of the type-I extreme value random variables being 3.290. In these experimental conditions, all models were still estimated assuming the standard logit model with variance 1.645 for the type-I extreme value random variables, implying that the estimated individual-level partworths $\beta_i(E)$ were implicitly scaled by a factor of $\frac{1}{\sqrt{2}}$ compared to the true individual-level partworths $\beta_i(T)$. Consequently, we multiplied the obtained $\beta_i(E)$ by $\sqrt{2}$ before computing RMSE(β) to put $\beta_i(T)$ and $\beta_i(E)$ on the same scale. Similar to Section 3.3.1, we computed the RMSE(β) for each

respondent and summarized the average RMSE(β) across respondents for each data set. We report the mean of the average RMSE(β)'s across 5 data sets in each experimental condition. Predictive accuracy was measured using the holdout sample log-likelihood [4], which we denote as Holdout-LL. For respondent i, the Holdout-LL was computed as follows:

$$\text{Holdout-LL} = \frac{1}{\tilde{J}} \sum_{j=1}^{\tilde{J}} \log \frac{e^{\tilde{x}_{ijh*}\beta_i(E)}}{\sum_{h=1}^{H} e^{\tilde{x}_{ijh}\beta_i(E)}}, \tag{3.3.3}$$

where $\tilde{J}$ denotes the number of holdout choice sets for each respondent, i.e., $\tilde{J} = 8$ in our study, and $\left\{\tilde{x}_{ijh}\right\}_{h=1}^{H}$ denotes the j-th holdout choice set for respondent i. Again, we computed the Holdout-LL for each respondent and summarized the average Holdout-LL across respondents for each data set. We report the mean of the average Holdout-LL's across 5 data sets in each experimental condition.

Tables 3.3 and 3.4 summarize the results, where Num-S denotes the number of segments in the population distribution (Factor 1), Num-CS denotes the number of choice sets per respondent for calibration (Factor 2), EV denotes the error variance (Factor 4), and WSV denotes the within-segment variances of distributions (Factor 3).

Insert Tables 3.3 and 3.4 here.

In the choice-based conjoint simulations, we find qualitatively similar results as those in the metric conjoint simulations except that the FM model becomes the best performing model when WSV is small. In particular, for the measure of parameter recovery RMSE(β) (Table 3.3), Choice-SL performs best or nonsignificantly different from best in 32 out of 56 experimental conditions at the $p < 0.05$ level; for the measure of predictive accuracy Holdout-LL (Table 3.4),

Choice-SL performs best or nonsignificantly different from best in 46 out of 56 conditions at the $p < 0.05$ level. Similar to Section 3.3.1, the comparisons are based on paired t-tests over the same 500 respondents, i.e., 100 respondents per data set $\times$ 5 data sets, in each experimental condition. Moreover, we find that as WSV increases the relative performance of the NCM model (the FM model) improves (deteriorates) and the performance gap between Choice-SL and LOG-Het narrows.

3.4 Applications

In this section, we demonstrate the empirical performance of our sparse learning (SL) approach using data from one metric conjoint study and one choice-based conjoint study.

3.4.1 Metric Conjoint Application

We compared Metric-SL, the FM model, the NCM model, and RR-Het using a metric conjoint data set of personal computers that was first introduced in [57] and later also used in [37].[16] In this study, 180 consumers each rated 20 hypothetical personal computers on an 11-point scale (0 to 10). Each hypothetical profile is represented using 13 binary attributes and an intercept. The first 16 profiles form an orthogonal and balanced design and were used for calibration, and the last 4 were used for holdout validation. We refer the reader to the two studies for details of this conjoint data set.

The predictive accuracy of the four conjoint estimation methods was assessed using two

[16] We thank Peter Lenk for kindly sharing this data set with us.

performance measures: (1) the root mean square error between the observed ratings, $Y_i(O)$, and the predicted ratings, $Y_i(P)$, for the holdout sample, i.e., the RMSE(Y) defined in Equation (3.3.2); (2) the first choice hits in the holdout sample [5], which we denote as 1stCH. For respondent i, 1stCH was set to 1 if her most preferred profile in the holdout sample, i.e., the holdout profile with the highest observed rating, was correctly predicted, and 0 otherwise. Following [37] and consistent with Section 3.3, we computed the RMSE(Y) and the 1stCH for each consumer and summarized the averages across 180 consumers for each method. Table 3.5 shows the results.

Insert Table 3.5 here.

Using paired t-tests over the 180 consumers, we find that Metric-SL and RR-Het perform best or not significantly different from best in terms of both RMSE(Y) and 1stCH (at $p < 0.05$). The performance comparison empirically validates the predictive accuracy of Metric-SL; it also suggests that the assumption of an unimodal continuous heterogeneity (UCH) distribution made by RR-Het is not restrictive on this computer conjoint data set if the researcher is primarily interested in prediction.

For Metric-SL, we find that the optimal modeling strategy S^* selected using cross-validation specifies a general covariance structure when modeling within-segment heterogeneity, i.e., $\text{COV}^* = G$, suggesting that the benefit of flexibility obtained by explicitly modeling the D^l matrix in Metric-HET-General outweighs the potential risk of overfitting in accurately modeling within-segment heterogeneity distributions on the computer conjoint data set, which has a relatively large sample size (i.e., 180). We also note that RR-Het specifies a general covariance structure for the

individual-level partworths via the explicit modeling of the D matrix, i.e., the decision matrix related to the covariance matrix of the partworths [37], which gives it some flexibility even within its UCH framework to approximate a complex and potentially multimodal heterogeneity distribution. To demonstrate the effect of a general covariance structure in helping RR-Het capture the variations in individual-level partworths for the computer conjoint data set, we estimated a restricted version of RR-Het in which the D matrix was constrained to the identity matrix. We found that the restricted version of RR-Het performed significantly worse than RR-Het in terms of RMSE(Y) (at $p < 0.05$) with the RMSE(Y)=1.6774.

In addition to assessing the predictive accuracy of the four conjoint estimation methods, we provide a graphical illustration of the individual-level heterogeneity representations recovered by these methods. Figure 3.1 shows the density plot of individual-level intercept estimates for the 180 consumers generated by each method.[17]

Insert Figure 3.1 here.

Figure 3.1 shows different patterns for the intercept estimates of Metric-SL and RR-Het, the two methods that performed best in terms of RMSE(Y) and 1stCH - Metric-SL recovers a heterogeneity distribution that suggests the presence of MCH whereas RR-Het recovers one that is more consistent with UCH. This contrast highlights the different approaches that Metric-SL and RR-Het adopt to model consumer heterogeneity despite their comparable predictive performance on the computer conjoint data set. Metric-SL is designed to handle MCH and models consumer heterogeneity by first decomposing the consumer population into multiple segments and then

[17] We estimate each density curve by applying a kernel smoothing density estimator to the individual-level point estimates of the intercept for all consumers produced by each conjoint estimation method. We implement the kernel smoothing density estimation procedure using the MATLAB function *ksdensity*.

imposing shrinkage within each segment. On the other hand, RR-Het assumes UCH and models consumer heterogeneity by directly imposing shrinkage on the full consumer population; in particular, RR-Het is not specifically designed to capture the potential multimodality in the heterogeneity distribution. Between these two models, we expect Metric-SL to provide a more viable approach for researchers interested in recovering and understanding MCH from conjoint data.

3.4.2 Choice-based Conjoint Application

We compared Choice-SL, the FM model, the NCM model, and LOG-Het on a choice-based conjoint data set of cell phone plans in [48].[18] A total of 72 consumers participated in this study, and each of them was shown 18 choice sets that consisted of three profiles and a no-choice option. Six attributes were used for constructing the conjoint profiles: (1) access fee, (2) per-minute rate, (3) plan minutes, (4) service provider, (5) Internet access, and (6) rollover of unused minutes. We tested these conjoint estimation methods using an additive conjoint specification, which [48] found to be the best fitting model – the "nonlinear-effects model" that adds to the standard conjoint model logarithmic terms in the access fee, per-minute rate, and plan minutes. We refer the reader to [48] for details on this data set and the conjoint specification. We differed from [48] in that we standardized all continuous attributes, i.e., access fee, per-minute rate, plan minutes, and their respective logarithmic terms, before model estimation; that is, each continuous attribute was demeaned and divided by its standard deviation. The standardization aimed to put all continuous attributes on similar scales and is widely

[18] We thank Raghuram Iyengar for kindly sharing this data set with us.

adopted in the statistics and machine learning literature [81]. We randomly selected 15 out of the 18 choice sets for each consumer for calibration and the remaining 3 for holdout validation.

We measured the predictive performance of the four conjoint estimation methods using the holdout sample log-likelihood, i.e., Holdout-LL defined in Equation (3.3.3), and the holdout sample hit rate, which we denote as Holdout-HIT. For respondent i, the Holdout-HIT was computed as follows:

$$\text{Holdout-HIT} = \frac{1}{\tilde{J}} \sum_{j=1}^{\tilde{J}} \mathbf{1}\big(\tilde{x}_{ijh*}\beta_i(E) = \max_h \{\tilde{x}_{ijh}\beta_i(E)\}\big), \tag{3.4.1}$$

where $\beta_i(E)$ denotes the estimated individual-level partworths, $\mathbf{1}(\cdot)$ takes value 1 when the argument is true, and 0 otherwise. While both Holdout-LL and Holdout-HIT assess the predictive accuracy of a conjoint estimation method, the continuous Holdout-LL provides a more sensitive measure [47]. The two performance measures also differ on how they penalize a model for assigning a low choice probability for the chosen profile of a choice set - the logarithmic functional form of Holdout-LL imposes a heavy penalty on such a scenario whereas the stepwise functional form of Holdout-HIT imposes a constant penalty as long as the chosen profile is not assigned with the highest choice probability.

Due to the relatively small sample size of the data set, i.e., 72 respondents, most performance comparisons among the four conjoint estimation methods in terms of Holdout-LL and Holdout-HIT were statistically insignificant using a test procedure similar to that used in Section 3.4.1, i.e., computing the performance measure for each consumer and conducting paired t-tests over the 72 consumers. In order to tackle the issue of small sample size, we adopted the following

alternative statistical test procedure. We generated 5 random replications of the data set where for each replication we retained all 72 consumers and randomly selected 15 out of the 18 choice sets for each consumer for calibration and the remaining 3 for holdout validation. Each conjoint estimation method was separately applied to each of the 5 replications, and its performance measures (i.e., Holdout-LL and Holdout-HIT) for each consumer were computed in each replication. We summarized the average Holdout-LL and Holdout-HIT across consumers for each replication and report the mean of the average Holdout-LL's and Holdout-HIT's. We compared the four conjoint estimation methods using paired t-tests over the 360 consumer-replication pairs, i.e., 72 consumers $\times$ 5 replications. One potential caveat of this alternative test procedure is that the performance measures for consumer-replication pairs associated with the same consumer may be correlated and consequently we might have inflated the statistical power by using such a test.

Table 3.6 summarizes the results.

Insert Table 3.6 here.

Table 3.6 shows that Choice-SL performs best (at $p < 0.05$) in terms of Holdout-LL whereas the NCM model performs best (at $p < 0.05$) in terms of Holdout-HIT. For Choice-SL, we find that the optimal modeling strategy S^* selected using cross-validation (in all 5 replications) specifies a restrictive covariance structure when modeling within-segment heterogeneity, i.e., $\text{COV}^* = R$, suggesting that the parsimony of restricting $D^l = I/p$ in Choice-HET-Restricted helps accurately modeling within-segment heterogeneity distributions given the relatively small sample size (i.e., 72) of the cell phone plans data set. As discussed earlier, Holdout-LL differs from Holdout-HIT in imposing a heavy penalty when a model assigns a low choice probability

for the chosen profile of a choice set. This property of Holdout-LL leads to a relatively weak performance of the NCM model on this measure, which, despite correctly predicting the chosen profiles with the highest frequency, also assigns low choice probabilities for many chosen profiles.[19]

Similar to Section 3.4.1, we provide a graphical illustration of the individual-level heterogeneity representations recovered by these methods. In particular, we focus on one replication and show the density plot of individual-level parameter estimates for the plan minutes for the 72 consumers generated by each method in Figure 3.2.[20]

Insert Figure 3.2 here.

Figure 3.2 shows that the heterogeneity distributions recovered by the two best performing models - Choice-SL for Holdout-LL and the NCM model for Holdout-HIT - are qualitatively different. In particular, Choice-SL recovers a heterogeneity distribution that is consistent with MCH whereas the NCM model recovers one that is unimodal.

3.5 Conclusions

In practice, consumer preferences can often be represented by a multimodal continuous heterogeneity (MCH) distribution and adequate modeling of MCH is critical for accurate conjoint estimation.

[19] It is not uncommon to find that two different measures of predictive accuracy favor different models. For example, [4] compared multiple specifications to modeling heterogeneity using two measures of predictive accuracy, Holdout-LL and $\bar{P}(V)$, the average predicted choice probability for the chosen profile in the holdout sample. They found that while Holdout-LL favored the FM-L model $\bar{P}(V)$ favored the HB model.

[20] Similar to Figure 3.1, we estimate each density curve by applying a kernel smoothing density estimator to the individual-level point parameter estimates for the plan minutes for all consumers produced by each conjoint estimation method. We implement the kernel smoothing density estimation procedure using the MATLAB function *ksdensity*.

In this paper, we propose an innovative sparse learning (SL) approach for modeling MCH and apply it to conjoint analysis. The SL approach models MCH via a divide-and-conquer framework, in which MCH is decomposed into a small collection of within-segment unimodal continuous heterogeneity (UCH) distributions using structured sparsity modeling and each UCH is then separately modeled. Consequently, both across- and within-segment heterogeneity are explicitly accounted for in the SL approach. In addition, the amount of shrinkage imposed to model each UCH is endogenously selected using cross-validation.

We test the empirical performance of our SL approach and compare it with the finite mixture (FM) model [32,52], the Bayesian normal component mixture (NCM) model [1], and the convex optimization (CO) model of [37] using extensive simulation experiments and two empirical conjoint data sets. In simulation experiments where the data-generating process is consistent with MCH, the SL approach demonstrates strong performance in terms of both parameter recovery and predictive accuracy on holdout samples. In the empirical conjoint data sets, the SL approach also shows competitive performance in terms of predictive accuracy and the individual-level partworths estimates of the SL approach display shapes consistent with MCH. Therefore, we empirically validate the performance of our SL approach in modeling MCH.

There are several promising avenues for future research. First, we can consider an extension of our SL approach by incorporating kernel methods [86] which was introduced to marketing by [34] and [36]. Second, researchers can also consider other population based complexity controls to improve the capability for modeling MCH. Third, it may be fruitful to compare our SL approach with nonparametric Bayesian approaches for modeling heterogeneity. Finally, as

noted by [37], an interesting research direction is to explore the potential of optimization / machine learning methods in modeling other phenomena in conjoint analysis beyond consumer heterogeneity.

Table 3.1: RMSE(β): Metric Conjoint Simulations

Num-S	Num-P	EV	WSV	Metric-SL	NCM	FM	RR-Het	Num-S	Num-P	EV	WSV	Metric-SL	NCM	FM	RR-Het
2	18	0.5	0.05	**0.1921**	0.2938	0.2264	0.2536	3	18	0.5	0.05	**0.1945**	0.2960	0.2281	0.2599
			0.10	**0.2401**	0.2980	0.3087	0.2810				0.10	**0.2444**	0.3182	0.3131	0.2899
			0.20	**0.2879**	0.3219	0.4333	0.3131				0.20	**0.2843**	0.3179	0.4279	0.3162
			0.40	**0.3316**	0.3453	0.5910	0.3442				0.40	**0.3322**	0.3455	0.6062	0.3480
			0.60	**0.3462**	0.3508	0.7282	0.3534				0.60	**0.3554**	0.3599	0.7686	0.3648
			0.80	**0.3705**	**0.3696**	0.8665	**0.3694**				0.80	**0.3676**	0.3707	0.8712	0.3735
			1.00	**0.3716**	0.3759	0.9538	0.3767				1.00	**0.3744**	**0.3764**	1.0342	**0.3762**
2	18	1.5	0.05	**0.2322**	0.4151	0.2453	0.3605	3	18	1.5	0.05	**0.2357**	0.4510	0.2446	0.3935
			0.10	**0.2870**	0.4436	0.3240	0.4179				0.10	**0.2921**	0.4469	0.3246	0.4232
			0.20	**0.3612**	0.4465	0.4493	0.4309				0.20	**0.3743**	0.4703	0.4552	0.4575
			0.40	**0.4483**	0.4875	0.6015	0.4781				0.40	**0.4537**	0.4924	0.6395	0.5041
			0.60	**0.5016**	0.5157	0.7553	0.5318				0.60	**0.5018**	0.5192	0.7673	0.5313
			0.80	**0.5402**	**0.5437**	0.8822	0.5549				0.80	**0.5510**	**0.5508**	0.9113	0.5614
			1.00	0.5624	**0.5577**	0.9541	0.5664				1.00	0.5776	**0.5687**	1.0463	0.5781
2	27	0.5	0.05	**0.1775**	0.2514	0.2200	0.2109	3	27	0.5	0.05	**0.1854**	0.2553	0.2276	0.2266
			0.10	**0.2201**	0.2596	0.3106	0.2401				0.10	**0.2188**	0.2632	0.3039	0.2482
			0.20	**0.2548**	0.2750	0.4254	0.2714				0.20	**0.2541**	0.2746	0.4259	0.2707
			0.40	**0.2835**	0.2896	0.5949	0.2914				0.40	**0.2848**	0.2913	0.5865	0.2948
			0.60	**0.2957**	0.2990	0.7168	0.2995				0.60	**0.2982**	**0.2993**	0.7359	0.3010
			0.80	**0.3094**	**0.3105**	0.8168	**0.3105**				0.80	**0.3138**	**0.3136**	0.8274	**0.3143**
			1.00	**0.3078**	0.3092	0.9236	**0.3084**				1.00	**0.3122**	**0.3120**	0.9318	**0.3126**
2	27	1.5	0.05	**0.2145**	0.3476	0.2336	0.2797	3	27	1.5	0.05	**0.2392**	0.3725	**0.2401**	0.3340
			0.10	**0.2705**	0.3672	0.3198	0.3362				0.10	**0.2772**	0.3843	0.3206	0.3512
			0.20	**0.3404**	0.3897	0.4366	0.3783				0.20	**0.3446**	0.3973	0.4462	0.3884
			0.40	**0.4061**	0.4298	0.6042	0.4312				0.40	**0.4040**	0.4265	0.6089	0.4278
			0.60	**0.4375**	0.4510	0.7374	0.4629				0.60	**0.4540**	0.4616	0.7679	0.4692
			0.80	**0.4755**	0.4829	0.8351	0.4856				0.80	**0.4716**	0.4755	0.8683	0.4802
			1.00	**0.4928**	**0.4947**	0.9607	0.4954				1.00	0.4978	**0.4925**	0.9996	**0.4932**

Notes. RMSE(β) reported above are averaged over 5 randomly simulated data sets. Bold numbers in each experimental condition indicate best or not significantly different from best at the $p < 0.05$ level based on paired t-tests. Metric-SL performs best or not significantly different from best in 53 out of 56 experimental conditions.

Table 3.2: RMSE(Y): Metric Conjoint Simulations

Num-S	Num-P	EV	WSV	Metric-SL	NCM	FM	RR-Het	Num-S	Num-P	EV	WSV	Metric-SL	NCM	FM	RR-Het
2	18	0.5	0.05	**0.7584**	0.8214	0.8296	0.8015	3	18	0.5	0.05	**0.7577**	0.8150	0.8361	0.7992
			0.10	**0.7812**	0.8178	0.9030	0.8086				0.10	**0.8089**	0.8560	0.9535	0.8382
			0.20	**0.8283**	0.8523	1.0784	0.8502				0.20	**0.8387**	0.8601	1.0929	0.8645
			0.40	**0.8528**	0.8604	1.2860	0.8623				0.40	**0.8729**	0.8789	1.3468	0.8822
			0.60	**0.8732**	0.8781	1.4961	0.8788				0.60	**0.8628**	**0.8657**	1.5836	0.8691
			0.80	**0.8814**	**0.8810**	1.7308	**0.8801**				0.80	**0.8668**	**0.8687**	1.7530	0.8706
			1.00	**0.8522**	0.8551	1.8484	**0.8539**				1.00	**0.8693**	**0.8683**	1.9922	0.8711
2	18	1.5	0.05	**1.2615**	1.3659	1.2931	1.3450	3	18	1.5	0.05	**1.2823**	1.4158	1.2949	1.3919
			0.10	**1.2906**	1.3707	1.3824	1.3661				0.10	**1.2845**	1.3681	1.3647	1.3640
			0.20	**1.3278**	1.3745	1.5008	1.3720				0.20	**1.3798**	1.4422	1.5805	1.4458
			0.40	**1.3706**	1.3930	1.6006	1.3931				0.40	**1.3900**	1.4115	1.7083	1.4185
			0.60	**1.4195**	1.4317	1.8283	1.4481				0.60	**1.4152**	1.4236	1.9150	1.4363
			0.80	**1.4665**	**1.4687**	2.0489	1.4740				0.80	**1.4809**	**1.4793**	2.1182	1.4847
			1.00	1.4675	**1.4614**	2.1148	1.4693				1.00	1.5005	**1.4946**	2.2167	1.5026
2	27	0.5	0.05	**0.7234**	0.7606	0.7811	0.7413	3	27	0.5	0.05	**0.7443**	0.7820	0.8017	0.7678
			0.10	**0.7708**	0.7941	0.8743	0.7814				0.10	**0.7578**	0.7895	0.8607	0.7794
			0.20	**0.7768**	0.7891	1.0006	0.7862				0.20	**0.7687**	0.7865	0.9880	0.7829
			0.40	**0.8097**	0.8164	1.2003	0.8148				0.40	**0.7993**	0.8021	1.2233	0.8054
			0.60	**0.8201**	0.8225	1.3804	0.8246				0.60	**0.8220**	**0.8236**	1.4466	**0.8235**
			0.80	**0.8303**	**0.8297**	1.5274	**0.8298**				0.80	**0.8163**	**0.8164**	1.5668	**0.8173**
			1.00	**0.8100**	**0.8105**	1.6819	**0.8101**				1.00	**0.8200**	**0.8201**	1.6736	**0.8199**
2	27	1.5	0.05	**1.2419**	1.3131	1.2662	1.2787	3	27	1.5	0.05	**1.2453**	1.3076	1.2707	1.2927
			0.10	**1.2745**	1.3294	1.3445	1.3140				0.10	**1.2546**	1.3001	1.3343	1.2841
			0.20	**1.2926**	1.3178	1.4035	1.3154				0.20	**1.3085**	1.3434	1.4472	1.3346
			0.40	**1.3443**	1.3607	1.5685	1.3656				0.40	**1.3573**	1.3634	1.6318	**1.3602**
			0.60	**1.3498**	1.3615	1.7219	1.3681				0.60	**1.3747**	**1.3743**	1.7689	**1.3774**
			0.80	**1.3630**	1.3673	1.8497	1.3678				0.80	**1.4063**	**1.4072**	1.9040	**1.4080**
			1.00	**1.3989**	**1.4004**	2.0023	**1.3988**				1.00	1.4106	**1.4071**	2.0384	**1.4053**

Notes. RMSE(Y) reported above are averaged over 5 randomly simulated data sets. Bold numbers in each experimental condition indicate best or not significantly different from best at the $p < 0.05$ level based on paired t-tests. Metric-SL performs best or not significantly different from best in 53 out of 56 experimental conditions.

Table 3.3: RMSE(β): Choice-based Conjoint Simulations

Num-S	Num-CS	EV	WSV	Choice-SL	NCM	FM	LOG-Het	Num-S	Num-CS	EV	WSV	Choice-SL	NCM	FM	LOG-Het
2	16	1.645	0.05	0.4174	0.7864	**0.2536**	0.6020	3	16	1.645	0.05	0.4106	0.7396	**0.2885**	0.6440
			0.10	0.5028	0.6493	**0.3436**	0.5790				0.10	0.5243	0.6609	**0.4650**	0.7053
			0.20	0.5112	0.6359	**0.4744**	0.6424				0.20	0.5690	0.6497	**0.5206**	0.7334
			0.40	**0.6154**	0.6723	0.6659	0.7014				0.40	**0.6157**	0.7270	0.7187	0.7579
			0.60	**0.6906**	**0.6853**	0.8595	0.7976				0.60	**0.6902**	0.7308	0.8943	0.8125
			0.80	**0.7727**	**0.7667**	0.9715	0.8193				0.80	**0.7390**	0.8009	0.9395	0.8541
			1.00	**0.7821**	**0.7689**	0.9946	0.8680				1.00	**0.8176**	**0.8404**	1.0757	0.8969
2	16	3.290	0.05	0.4749	0.7300	**0.3027**	0.6760	3	16	3.290	0.05	0.4862	0.8303	**0.3467**	0.7298
			0.10	0.5099	0.7161	**0.3534**	0.7136				0.10	0.5409	0.7898	**0.4182**	0.7596
			0.20	0.4918	0.7371	**0.4797**	0.6087				0.20	0.5997	0.7768	**0.5208**	0.7361
			0.40	**0.5853**	0.7592	0.6287	0.6983				0.40	**0.6370**	0.8217	0.6887	0.8235
			0.60	**0.7102**	0.8113	0.8132	0.7751				0.60	**0.7148**	0.8430	0.8254	0.8539
			0.80	**0.8055**	**0.7902**	0.9583	0.8696				0.80	**0.8196**	0.9028	0.9742	0.9248
			1.00	**0.8811**	**0.8756**	1.0337	0.9388				1.00	**0.8614**	0.9542	1.0484	0.9283
2	24	1.645	0.05	0.3892	0.5514	**0.2586**	0.4537	3	24	1.645	0.05	0.4501	0.5786	**0.2962**	0.5448
			0.10	0.3967	0.5491	**0.3622**	0.5742				0.10	0.4245	0.5680	**0.3494**	0.6294
			0.20	**0.4342**	0.5260	0.4493	0.5282				0.20	**0.4606**	0.5672	0.4914	0.6056
			0.40	**0.5421**	0.5681	0.6527	0.5831				0.40	**0.5279**	0.6288	0.6771	0.6274
			0.60	**0.6070**	0.6358	0.7796	0.6487				0.60	**0.6034**	0.6516	0.8240	0.6553
			0.80	0.6793	**0.6364**	0.8908	0.7244				0.80	**0.6960**	**0.6891**	0.9597	0.7580
			1.00	**0.6846**	**0.6825**	0.9772	0.7372				1.00	0.7402	**0.7133**	1.0142	0.7913
2	24	3.290	0.05	0.4007	0.5875	**0.2877**	0.5757	3	24	3.290	0.05	0.3397	0.6507	**0.2701**	0.5879
			0.10	0.4524	0.6268	**0.3595**	0.5993				0.10	0.4773	0.6361	**0.3820**	0.6773
			0.20	0.5065	0.6418	**0.4735**	0.6215				0.20	0.5107	0.6391	**0.4879**	0.6595
			0.40	**0.5640**	0.6708	0.6512	0.6199				0.40	**0.5809**	0.7040	0.6739	0.6877
			0.60	**0.6189**	0.6949	0.7843	0.6931				0.60	**0.6989**	0.7316	0.8276	0.7666
			0.80	**0.6973**	0.7276	0.9078	0.7206				0.80	**0.6900**	0.7360	0.8943	0.7775
			1.00	**0.7454**	0.7651	0.9632	0.7825				1.00	**0.7489**	**0.7624**	1.0338	0.8210

Notes. RMSE(β) reported above are averaged over 5 randomly simulated data sets. Bold numbers in each experimental condition indicate best or not significantly different from best at the $p < 0.05$ level based on paired t-tests. Choice-SL performs best or not significantly different from best in 32 out of 56 experimental conditions.

Table 3.4: Holdout-LL: Choice-based Conjoint Simulations

Num-S	Num-CS	EV	WSV	Choice-SL	NCM	FM	LOG-Het	Num-S	Num-CS	EV	WSV	Choice-SL	NCM	FM	LOG-Het
2	16	1.645	0.05	−0.7090	−0.7570	**-0.6908**	−0.7559	3	16	1.645	0.05	−0.7355	−0.7807	**-0.7210**	−0.7873
			0.10	**-0.7160**	−0.7323	**-0.7131**	−0.7439				0.10	**-0.7468**	−0.7773	**-0.7447**	−0.7846
			0.20	**-0.7296**	−0.7572	−0.7552	−0.7599				0.20	**-0.7400**	−0.7672	−0.7577	−0.7742
			0.40	**-0.7135**	−0.7364	−0.7658	−0.7276				0.40	**-0.7678**	−0.7995	−0.8285	−0.7858
			0.60	**-0.6750**	−0.6915	−0.7466	−0.6950				0.60	**-0.7168**	−0.7487	−0.8068	−0.7327
			0.80	**-0.7670**	−0.7932	−0.8603	−0.7763				0.80	**-0.6873**	−0.7209	−0.7880	**-0.6920**
			1.00	**-0.7258**	−0.7475	−0.8189	−0.7353				1.00	**-0.7236**	−0.7603	−0.8390	−0.7350
2	16	3.290	0.05	−0.8992	−0.9439	**-0.8773**	−0.9480	3	16	3.290	0.05	−0.8961	−0.9548	**-0.8803**	−0.9434
			0.10	−0.8911	−0.9268	**-0.8759**	−0.9191				0.10	−0.9395	−0.9946	**-0.9231**	−0.9815
			0.20	**-0.8667**	−0.9053	**-0.8715**	−0.8941				0.20	**-0.8846**	−0.9146	**-0.8797**	−0.9134
			0.40	**-0.8974**	−0.9281	−0.9312	−0.9190				0.40	**-0.8836**	−0.9227	−0.9129	−0.9140
			0.60	**-0.8818**	−0.9079	−0.9237	−0.8949				0.60	**-0.8931**	−0.9151	−0.9298	−0.9141
			0.80	**-0.8741**	**-0.8808**	−0.9254	−0.8808				0.80	**-0.9522**	−0.9626	−1.0037	**-0.9458**
			1.00	**-0.9032**	−0.9235	−0.9654	−0.9149				1.00	**-0.9073**	−0.9334	−0.9878	**-0.9046**
2	24	1.645	0.05	−0.7547	−0.7756	**-0.7394**	−0.7769	3	24	1.645	0.05	−0.6886	−0.7271	**-0.6749**	−0.7321
			0.10	**-0.7659**	−0.7929	**-0.7702**	−0.7987				0.10	**-0.7621**	−0.7898	**-0.7629**	−0.7996
			0.20	**-0.7061**	−0.7269	−0.7218	−0.7291				0.20	**-0.6946**	−0.7183	−0.7238	−0.7208
			0.40	**-0.7499**	−0.7636	−0.8018	−0.7593				0.40	**-0.6941**	−0.7124	−0.7440	−0.7104
			0.60	**-0.6639**	−0.6784	−0.7327	−0.6755				0.60	**-0.7517**	−0.7709	−0.8407	**-0.7557**
			0.80	**-0.7249**	**-0.7293**	−0.7947	**-0.7277**				0.80	**-0.7163**	−0.7345	−0.8030	−0.7225
			1.00	**-0.7064**	−0.7187	−0.8005	**-0.7073**				1.00	**-0.6885**	−0.6967	−0.8002	**-0.6914**
2	24	3.290	0.05	−0.9124	−0.9409	**-0.9047**	−0.9407	3	24	3.290	0.05	−0.9222	−0.9637	**-0.9141**	−0.9624
			0.10	**-0.8231**	−0.8483	**-0.8188**	−0.8480				0.10	**-0.8934**	−0.9185	**-0.8893**	−0.9190
			0.20	**-0.8446**	−0.8694	**-0.8461**	−0.8708				0.20	**-0.8606**	−0.8839	−0.8728	−0.8818
			0.40	**-0.9147**	−0.9271	−0.9530	**-0.9178**				0.40	**-0.8742**	−0.8951	−0.9068	−0.8932
			0.60	**-0.8722**	−0.8850	−0.9285	−0.8809				0.60	**-0.9283**	−0.9417	−0.9754	**−0.9329**
			0.80	**-0.8948**	−0.9065	−0.9575	**-0.8961**				0.80	**-0.8962**	−0.9090	−0.9496	−0.9073
			1.00	**-0.8555**	−0.8637	−0.9205	−0.8602				1.00	**-0.8670**	−0.8761	−0.9562	−0.8725

Notes. Holdout-LL reported above are averaged over 5 randomly simulated data sets. Bold numbers in each experimental condition indicate best or not significantly different from best at the $p < 0.05$ level based on paired t-tests. Choice-SL performs best or nonsignificantly different from best in 46 out of 56 experimental conditions.

Table 3.5: The Personal Computer Conjoint Data Set of [57]

	Metric-SL	NCM	FM	RR-Het
RMSE(Y)	**1.6216**	1.6561	1.8639	**1.6096**
1stCH	**0.7000**	0.6722	0.5889	**0.6944**

Notes. For RMSE(Y) lower numbers indicate higher performance whereas for 1stCH higher numbers indicate higher performance. Bold numbers in each row indicate best or not significantly different from best at the $p < 0.05$ level based on paired t-tests.

Table 3.6: The Cell Phone Plans Data Set of [48]

	Choice-SL	NCM	FM	LOG-Het
Holdout-LL	**-0.8827**	-0.9403	-0.9928	-0.9257
Holdout-HIT	0.6361	**0.6537**	0.5944	0.6287

Notes. For both Holdout-LL and Holdout-HIT higher numbers indicate higher performance. Bold numbers in each row indicate best or not significantly different from best at the $p < 0.05$ level based on paired t-tests.

Figure 3.1: The Density Plots of the Individual-level Intercept Estimates for the Personal Computer Conjoint Data Set of [57]

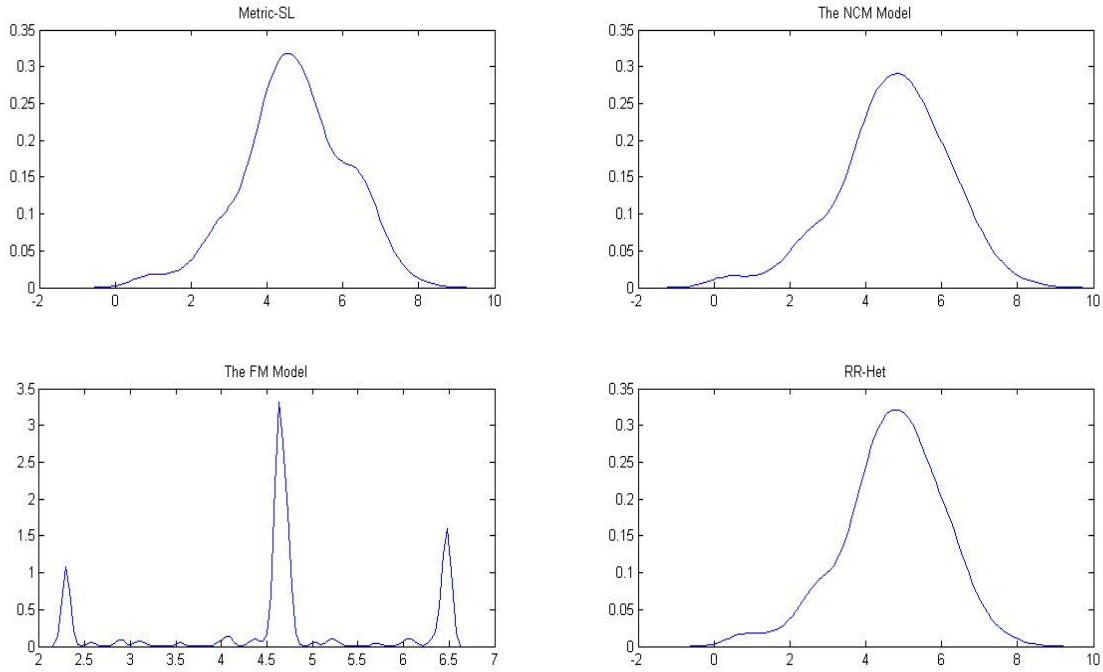

Figure 3.2: The Density Plots of the Individual-level Coefficient Estimates for Plan Minutes for the Cell Phone Plans Conjoint Data Set of [48]

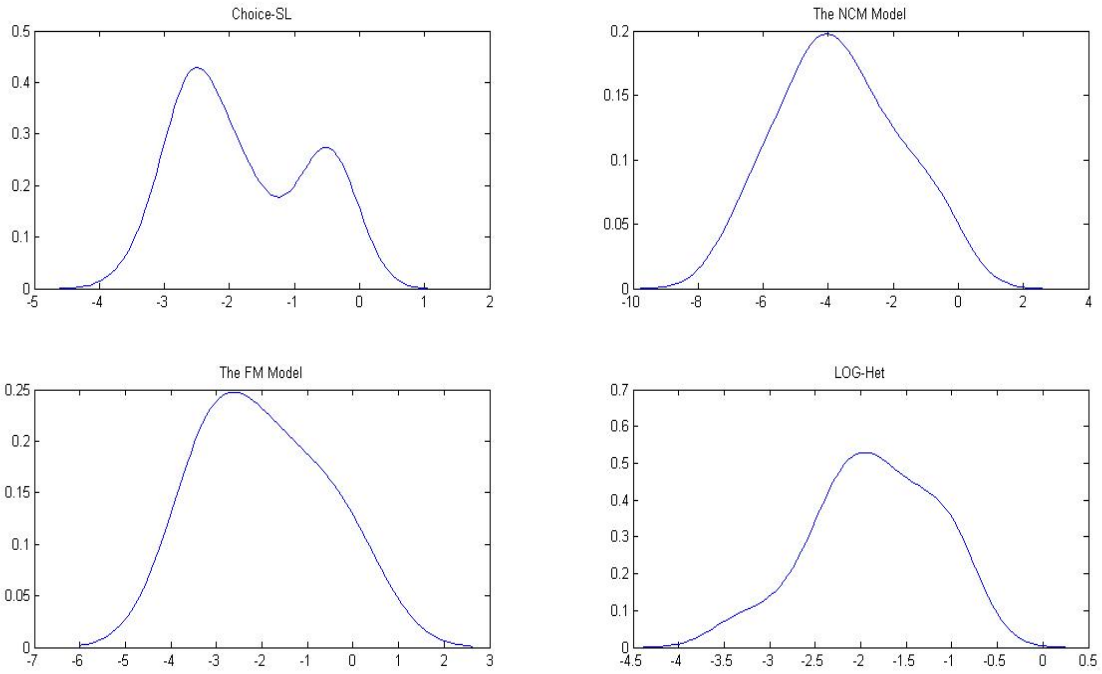

Chapter 4

Modeling Consumer Purchase and Return Decisions

4.1 Introduction

Uncertainty is always present when consumers make purchase decisions on any durable goods as they only have access to imperfect information about products' quality or fit before usage. Oftentimes, such uncertainty has a significant impact on consumers' choice behavior and willingness-to-pay. An ubiquitous approach for retailers to alleviate the impact of consumer uncertainty at the time of purchase is to offer a return policy, which allows consumers to use and experience a product for a certain amount of time before the end of which a return vs. retain decision must be made. This way, consumers are able to resolve their uncertainty and uncover their true evaluation regarding a product before a full commitment to purchase it.

Despite the widespread use of return policies in practice, a thorough understanding of the

impact of return policies on consumer purchase and return decisions remains elusive, especially in a competitive environment. Consequently, there seems to be little consensus over the composition of an optimal return policy, which has been reflected in debates over liberal vs. restrictive return policies in the popular press as well as the wide heterogeneity in return policies implemented by retailers. In addition, an estimated $100 billion loss annually through product depreciation and management of the return process [26] also suggests that the current practice is far from optimal. Therefore, from the perspectives of both academics and practitioners, a thorough empirical investigation of the impact of return policies on consumer purchase and return decisions is not only interesting but also relevant.

Marketing researchers have contributed to this important discussion using an array of research methodologies. Through lab experiments, [91] found that return policy leniency decreases deliberation time, increases expectations of product quality, reduces continued product search, and that differences in perceived quality due to return policy signals can persist after product receipt. [11] hypothesized product returns to be the result of exposure to disconfirming information, and tested the impact of the cognitive responses generated before purchase on consumers' return decisions. [76] developed an economic model to demonstrate how retailers can strategically use return policies and information provision to maximize profit. Empirical modeling research on return policies, on the other hand, has been quite sparse. To the best of our knowledge, [3] is the only published paper that calibrates an econometric model for consumer purchase and return decisions under return policies using individual-level data. Using their model, the authors calculated the individual-level option value of return policies in the context of a mail-order

apparel retailer, and demonstrated how the apparel retailer could use these information to guide the optimal design of return policies for each consumer.

One modeling assumption made by [3] is that consumers resolve their uncertainty on a product's quality or fit instantaneously after they receive the product, which allows the researchers to ignore the temporal dimension of return policies - the length of the "trial period" before the return deadline - and model return policies simply as options to a one-shot gamble. This assumption is appropriate in some product categories, e.g., apparel, shoes, etc., but its validity is questionable in many other product categories, e.g., consumer electronics, where consumer learning is unlikely to happen instantaneously. Instead, it is more plausible that consumers will gradually resolve their uncertainty and learn their true evaluation of a product over time by using the product, in which case the length of the trial period before the return deadline will have an impact on the amount of learning that could take place and hence affect consumer purchase and return decisions. On the other hand, since consumers not only learn their evaluation but also derive utility from using the product during the trial period, retailers usually charge restocking fee, which will also impact consumer purchase and return decisions by imposing a cost for returning the product.

In order to obtain a holistic understanding of the impact of return policies, we propose to collect a comprehensive data set on consumer purchase and return behavior using a field experiment. In the field experiment, each participant will be given a set of real products, each of which comes with a price and a return policy, and she will decide to purchase, use, and/or return one or more products from the set of products given to her over a finite horizon. In addition to observing the complete sequence of purchase and return decisions made by each

participant, we will also collect each participant's evaluation of the uncertain quality or fit of the product that she uses in each period which provides a measure of her learning process. We model the observed consumers' purchase and return decisions using a dynamic discrete choice model with forward-looking and Bayesian learning. Our modeling framework is similar to [35] in that consumers' learning of their true evaluations of products is explicitly modeled using a Bayesian learning model and is embedded in a dynamic expected-utility maximization problem in which consumers are assumed to make purchase and return decisions to maximize their expected present value of utility. Our proposed model provides a behaviorally plausible approach to capture the joint impact of return deadline and restocking fee of a return policy.

This is an ongoing project. In the following we present the design of the field experiment in Section 4.2 and discuss the proposed dynamic discrete choice model in Section 4.3. We leave the implementation of the field experiment and the dynamic discrete choice model as future work.

4.2 Field Experiment

In this section, we discuss the design of the field experiment for collecting consumer purchase and return decisions. We plan to recruit I participants and implement the experiment over a finite horizon of T periods.[1] At the beginning of the experiment, participant i will be given a set of J real products $\left\{\Lambda_{ij}\right\}_{j=1}^{J}$. The j-th product, Λ_{ij}, can be mathematically represented using the following attributes:

[1] In the experiment, periods can be operationalized as any reasonably defined time intervals, such as days, weeks, or months.

- $X_{ij} \in \mathbb{R}^d$, the search attributes of Λ_{ij}; X_{ij} are fully known to participant i at the beginning of the experiment.

- $A_{ij} \in \mathbb{R}$, the experiential attribute of Λ_{ij} which can be operationalized as quality or fit; A_{ij} is uncertain from participant i's perspective; as part of the experiment, we will manipulate participant i's initial belief about A_{ij} by instructing her that a subset of past participants has used the same product and the mean and variance of their evaluations of its experiential attribute is μ_{ij1} and σ^2_{ij1}.

- P_{ij}, the price of Λ_{ij}.

- L_{ij}, the return deadline of Λ_{ij}.

- R_{ij}, the restocking fee of Λ_{ij}.

During the experiment, participant i will make a sequence of purchase and return decisions regarding the set of products $\left\{\Lambda_{ij}\right\}_{j=1}^{J}$. The dynamics of the experiment closely resembles that of a real retailing setting, with the exception that participant i can possess at most one product at any time of the experiment. In the following, we provide a detailed description of the experiment and the decisions participant i will make.

Period 1**.** At the beginning of period 1, participant i will decide to purchase one or none of the products $\left\{\Lambda_{ij}\right\}_{j=1}^{J}$. If she purchases Λ_{ij}, she will pay the price P_{ij} and at the end of period 1 she will report her perception of the experiential attribute A_{ij} in period 1, A_{ij1}.

Period t **(**$2 \leq t \leq T$**).** At the beginning of period t, the decision that participant i will make depends on her previous decision sequence. The following scenarios could arise:

(a) Participant i possesses Λ_{ij} and the return deadline of Λ_{ij}, L_{ij}, has not passed. In this scenario, she can either choose to retain Λ_{ij} for period t, in which case she will report at the end of period t her perception of the experiential attribute A_{ij} in period t, A_{ijt}, or choose to return Λ_{ij} and receive the refund of the purchase price P_{ij} less the restocking fee R_{ij}. Conditional on participant i choosing to return Λ_{ij}, she will have an opportunity to decide whether to purchase one or none of the products. If she purchases $\Lambda_{ij'}$, she will pay the price $P_{ij'}$ and report $A_{ij't}$ at the end of period t.

(b) Participant i possesses Λ_{ij} and the return deadline L_{ij} has passed. In this scenario, she can only choose to retain Λ_{ij} as it can no longer be returned. In addition, she will not be able to purchase other products because of the restriction that at most one product can be possessed at any given time of the experiment. She will report A_{ijt} at the end of period t.

(c) Participant i does not possess any product. In this scenario, she will decide to purchase one or none of the products. If she purchases Λ_{ij}, she will pay the price P_{ij} and at the end of period t she will report A_{ijt}.

Compared to observational data sets on consumer purchase and return behavior which are potentially accessible from retailers, e.g., the data set in [3], our proposed field experiment provides several advantages for understanding the impact of return policies on consumer purchase and return decisions. First, in our experiment the researchers have full control over the design of return policies and are able to manipulate and generate variations in return policies across

products, whereas in practice a retailer is unlikely to vary return policies across products within the same category and hence limited variations in return policies could be observed in an observational data set. Second, in our experiment the researchers observe the choice set facing each participant and her complete sequence of purchase and return decisions; in contrast, based on an observational data set provided by a retailer the researchers can only observe consumers' purchase and return of products at the focal retailer and miss consumers' purchase and return of products at other retailers. Third, in our experiment the researchers observe participants' perceptions of the uncertain experiential attributes of the products they purchase over time, which provides a measure of the dynamics of their learning process.

4.3 Model

In this section, we consider a dynamic discrete choice model with forward-looking and Bayesian learning to model consumer purchase and return decisions collected using the field experiment discussed in Section 4.2. Our proposed model shares a similar modeling framework with the forward-looking dynamic structural model of [35] in that participants are assumed to make purchase and return decisions to maximize the expected present value of utility while learning uncertain experiential attributes of products in a Bayesian manner. In the following we discuss our model in terms of state variables and decisions, utility functions, dynamics of Bayesian learning, and finally the likelihood function.

4.3.1 State Variables and Decisions

We denote S_{it} as the state variables for participant i at the beginning of period t, i.e., S_{it} provide all relevant information for participant i to make decision at the beginning of period t. In our model, S_{it} consist of the following components: (1) t, the time index of the current period; (2) W_{it}, the product that participant i possesses at the beginning of period t; $W_{it} = j$ if Λ_{ij} is possessed, $j = 1, 2, \ldots, J$, or $W_{it} = 0$ if no product is possessed; (3) Q_{it}, the number of time periods remains until the return deadline of the product currently being possessed; we set $Q_{it} = +\infty$ if $W_{it} = 0$; (4) I_{it}, the information set available at the beginning of period t, i.e., $I_{it} \triangleq \left\{A_{ijs} \middle| \Lambda_{ij} \text{ is used in period } s, 1 \leq s \leq t-1\right\}$; and (5) ε_{it}, the unobserved state variable which is observable by participant i at the beginning of period t but unobservable by the researcher and which serves as the econometric error [70]; the dimension of ε_{it} equals the number of feasible actions, which is described in details below.

Let D_{it} denote the decision to be made by participant i at the beginning of period t. The set of feasible values for D_{it} is determined by the state variables S_{it}. We consider the following three scenarios:[2]

(a) $W_{it} = j$, $j = 1, 2, \ldots, J$, and $Q_{it} \geq 0$. In this scenario, participant i can choose to retain Λ_{ij}, in which case we denote $D_{it} = r$; or choose to return Λ_{ij} and purchase $\Lambda_{ij'}$, $j' = 1, 2, \ldots, J$, in which case we denote $D_{it} = j'$; or choose to return Λ_{ij} and purchase no product, in which case we denote $D_{it} = 0$. The state transition, depending on the value of D_{it}, is as follows: when $D_{it} = r$, we have $W_{i,t+1} = W_{i,t}$, $Q_{i,t+1} = Q_{i,t} - 1$, $I_{i,t+1} = I_{i,t} \cup \{A_{iW_{it}t}\}$;

[2] These three scenarios have already been introduced in Section 4.2.

when $D_{it} = j'$, we have $W_{i,t+1} = j'$, $Q_{i,t+1} = L_{ij'} - 1$, $I_{i,t+1} = I_{i,t} \cup \{A_{ij't}\}$; when $D_{it} = 0$, we have $W_{i,t+1} = 0$, $Q_{i,t+1} = +\infty$, $I_{i,t+1} = I_{i,t}$.

(b) $W_{it} = j$, $j = 1, 2, \ldots, J$, and $Q_{it} < 0$. In this scenario, participant i can only choose to retain Λ_{ij}, i.e., D_{it} can only take value r. The state transition is $W_{i,t+1} = W_{i,t}$, $Q_{i,t+1} = Q_{i,t} - 1$, $I_{i,t+1} = I_{i,t} \cup \{A_{iW_{it}t}\}$.

(c) $W_{it} = 0$. In this scenario, participant i can choose to purchase Λ_{ij}, $j = 1, 2, \ldots, J$, in which case we denote $D_{it} = j$; or choose to purchase no product, in which case we denote $D_{it} = 0$. The state transition, depending on the value of D_{it}, is as follows: when $D_{it} = j$, we have $W_{i,t+1} = j$, $Q_{i,t+1} = L_{ij} - 1$, $I_{i,t+1} = I_{i,t} \cup \{A_{ijt}\}$; when $D_{it} = 0$, we have $W_{i,t+1} = 0$, $Q_{i,t+1} = +\infty$, $I_{i,t+1} = I_{i,t}$.

We note that the state transition from $I_{i,t}$ to $I_{i,t+1}$ is stochastic from participant i's perspective at the beginning of period t as her perception of the experiential attributes in period t, i.e., A_{ijt}'s, are uncertain. As will become clear soon (Section 4.3.3), we will assume that participant i forms a subjective belief about the distribution of A_{ijt}'s and hence the state transition probability distribution according to a Bayesian learning model.

4.3.2 Utility Functions

In this section, we specify the per-period utility function $U_{it}\left(S_{it}, D_{it}, A_{it}, \varepsilon_{it}\right)$, where $A_{it} = \{A_{ijt}\}_{j=1}^{J}$. Following [35], we adopt additive compensatory specifications for U_{it}. Again, we discuss the specifications for U_{it} by considering the following three scenarios:

(a) $W_{it} = j$, $j = 1, 2, \ldots, J$, and $Q_{it} \geq 0$. When $D_{it} = r$,

$$U_{it}\left(S_{it}, D_{it}, A_{it}, \varepsilon_{it}\right) = X_{ij}^{T}\beta_i + \omega_i A_{ijt} - \omega_i r_i A_{ijt}^2 + \varepsilon_{irt}. \tag{4.3.1}$$

When $D_{it} = j'$, $j' = 1, 2, \ldots, J$,

$$U_{it}\left(S_{it}, D_{it}, A_{it}, \varepsilon_{it}\right) = \alpha_i(P_{ij} - R_{ij} - P_{ij'}) + X_{ij'}^{T}\beta_i + \omega_i A_{ij't} - \omega_i r_i A_{ij't}^2 + \varepsilon_{ij't}. \tag{4.3.2}$$

When $D_{it} = 0$,

$$U_{it}\left(S_{it}, D_{it}, A_{it}, \varepsilon_{it}\right) = \alpha_i(P_{ij} - R_{ij}) + \tau_i + \varepsilon_{i0t}. \tag{4.3.3}$$

In the above specifications, α_i is participant i's utility weight for price; β_i is her utility weight for the search attributes; ω_i is her utility weight for the perceived experiential attribute and r_i is her risk coefficient; τ_i is the baseline utility participant i derives from not possessing any product in a period; ε's are the unobserved state variables. The specifications for U_{it} in the other two scenarios are similar and are presented below.

(b) $W_{it} = j$, $j = 1, 2, \ldots, J$, and $Q_{it} < 0$. D_{it} can only take value r.

$$U_{it}\left(S_{it}, D_{it}, A_{it}, \varepsilon_{it}\right) = X_{ij}^{T}\beta_i + \omega_i A_{ijt} - \omega_i r_i A_{ijt}^2 + \varepsilon_{irt}. \tag{4.3.4}$$

(c) $W_{it} = 0$. When $D_{it} = j$, $j = 1, 2, \ldots, J$,

$$U_{it}\left(S_{it}, D_{it}, A_{it}, \varepsilon_{it}\right) = -\alpha_i P_{ij} + X_{ij}^{T}\beta_i + \omega_i A_{ijt} - \omega_i r_i A_{ijt}^2 + \varepsilon_{ijt}. \tag{4.3.5}$$

When $D_{it} = 0$,

$$U_{it}\left(S_{it}, D_{it}, A_{it}, \varepsilon_{it}\right) = \tau_i + \varepsilon_{i0t}. \tag{4.3.6}$$

We note that U_{it} is stochastic from participant i's perspective at the beginning of period t as her perception of the experiential attributes in period t, i.e., A_{it}, are uncertain. Consequently, when making the decision D_{it} participant i will use the expected utility $\mathbb{E}\left[U_{it}|S_{it}\right]$, where the expectation is taken with respect to A_{it}. As mentioned in Section 4.3.1, we will assume that participant i forms a subjective belief about the distribution of A_{it} according to a Bayesian learning model, which is introduced in Section 4.3.3.

4.3.3 Dynamics of Bayesian Learning

We assume participant i is uncertain about the experiential attributes A_{ij}'s and receives imperfect information about these uncertain attributes over time via product usage. In particular, if she possesses and uses Λ_{ij} in period t she will receive a noisy signal about A_{ij}, A_{ijt}, which will be later reported to the researcher. We assume the following information structure:

$$A_{ijt} = A_{ij} + \delta_{ijt}, \quad \delta_{ijt} \sim N(0, \sigma_\delta^2), \quad \forall i, j, t. \tag{4.3.7}$$

That is, participant i's perception of the experiential attribute A_{ij} in period t is A_{ij} plus a mean zero normal noise with constant variance. We assume that participant i knows the information structure including the variance σ_δ^2 but not the true value of A_{ij}'s, and she learns A_{ij}'s using Bayesian updating. In addition, we assume participant i's prior on A_{ij} at the beginning of

the experiment is $N(\mu_{ij1}, \sigma^2_{ij1})$, where μ_{ij1} and σ^2_{ij1} are given to participant i as the mean and variance of past participants' evaluations of A_{ij}. Under these assumptions, participant i's posterior distribution on A_{ij} at the beginning of period t and given the information set I_{it} is $N(\mu_{ijt}, \sigma^2_{ijt})$, where μ_{ijt} and σ^2_{ijt} can be computed using the following recursion [41]:

$$\mu_{ijt} = \begin{cases} \frac{\frac{A_{ij,t-1}}{\sigma^2_\delta} + \frac{\mu_{ij,t-1}}{\sigma^2_{ij,t-1}}}{\frac{1}{\sigma^2_\delta} + \frac{1}{\sigma^2_{ij,t-1}}}, & \text{if } A_{ij,t-1} \in I_{it}, \\ \mu_{ij,t-1}, & \text{if } A_{ij,t-1} \notin I_{it}, \end{cases} \tag{4.3.8}$$

and

$$\sigma^2_{ijt} = \begin{cases} \frac{1}{\frac{1}{\sigma^2_\delta} + \frac{1}{\sigma^2_{ij,t-1}}}, & \text{if } A_{ij,t-1} \in I_{it}, \\ \sigma^2_{ij,t-1}, & \text{if } A_{ij,t-1} \notin I_{it}. \end{cases} \tag{4.3.9}$$

That is, $\left[A_{ij}|S_{it}\right] \sim N(\mu_{ijt}, \sigma^2_{ijt})$ with μ_{ijt} and σ^2_{ijt} given in Equations (4.3.8) and (4.3.9). Moreover, the posterior predictive distribution of A_{ijt} is also normal, i.e., $\left[A_{ijt}|S_{it}\right] \sim N(\mu_{ijt}, \sigma^2_{ijt} + \sigma^2_\delta)$ [41].

4.3.4 Likelihood

The objective of participant i is finding the optimal sequence of decisions $\left\{D_{it}(S_{it})\right\}_{t=1}^{T}$ to maximize the expected present value of utility over the planning horizon:

$$\mathbb{E}\Big[\sum_{t=1}^{T} \rho_i^{t-1} \mathbb{E}\left[U_{it}\left(S_{it}, D_{it}, A_{it}, \varepsilon_{it}\right)\Big|S_{it}\right]\Big|S_{i1}\Big], \tag{4.3.10}$$

where ρ_i is the discount factor. We define the value function $V_{it}(S_{it})$ as the maximal expected value of the discounted expected utility for participant i at period t given the state variables S_{it},

i.e.,

$$V_{it}(S_{it}) = \max_{\{D_{i\tau}(S_{i\tau})\}_{\tau=t}^{T}} \mathbb{E}\Big[\sum_{\tau=t}^{T} \rho_i^{\tau-t} \mathbb{E}\big[U_{i\tau}\big(S_{i\tau}, D_{i\tau}, A_{i\tau}, \varepsilon_{i\tau}\big)\big|S_{i\tau}\big]\big|S_{it}\Big]. \tag{4.3.11}$$

The value function $V_{it}(S_{it})$ can be expressed as the maximum of the alternative specific value functions $V_{it}(S_{it}, D_{it})$, that is,

$$V_{it}(S_{it}) = \max_{D_{it}} V_{it}(S_{it}, D_{it}), \tag{4.3.12}$$

where the alternative specific value functions $V_{it}(S_{it}, D_{it})$ obey the Bellman equation:

$$V_{it}(S_{it}, D_{it}) = \mathbb{E}\big[U_{it}\big(S_{it}, D_{it}, A_{it}, \varepsilon_{it}\big)\big|S_{it}\big] + \rho_i \mathbb{E}\big[V_{i,t+1}(S_{i,t+1})\big|S_{it}, D_{it}\big], \quad t = 1, 2, \ldots, T-1, \tag{4.3.13}$$

and

$$V_{iT}(S_{iT}, D_{iT}) = \mathbb{E}\big[U_{iT}\big(S_{iT}, D_{iT}, A_{iT}, \varepsilon_{iT}\big)\big|S_{iT}\big]. \tag{4.3.14}$$

We note that the alternative specific value functions $V_{it}(S_{it}, D_{it})$ are deterministic from participant i's perspective at the beginning of period t, but are stochastic from the researcher's perspective as the researcher cannot observe the unobserved state variables ε_{it}. Assuming the individual components of ε_{it} are independently and identically distributed type-1 extreme value random variables, the researcher can form conditional choice probabilities for participant i's decisions [70]. More specifically, define $\bar{V}_{it}(S_{it}, D_{it}) = V_{it}(S_{it}, D_{it}) - \varepsilon_{i,D_{it},t}$, i.e., $\bar{V}_{it}(S_{it}, D_{it})$ is the part of the alternative specific value function that does not depend on the unobserved state variables and is therefore deterministic from the researcher's perspective. The likelihood function for

participant i's decision sequence is

$$L_i = \prod_{t=1}^{T} \frac{e^{\bar{V}_{it}(S_{it},D_{it})}}{\sum_{D'_{it}} e^{\bar{V}_{it}(S_{it},D'_{it})}}, \tag{4.3.15}$$

where the summation in the denominator is taken over all feasible decisions D'_{it} given the state variables S_{it}. Given the likelihood function, our model could be estimated using the method of simulated maximum likelihood following [35]. We leave the implementation of the estimation algorithm as future work.

Bibliography

[1] G.M. Allenby, N. Arora, and J.L. Ginter, *On the heterogeneity of demand*, Journal of Marketing Research (1998), 384–389.

[2] G.M. Allenby and P.E. Rossi, *Marketing models of consumer heterogeneity*, Journal of Econometrics **89** (1998), no. 1, 57–78.

[3] Eric T Anderson, Karsten Hansen, and Duncan Simester, *The option value of returns: Theory and empirical evidence*, Marketing Science **28** (2009), no. 3, 405–423.

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

[5] R.L. Andrews, A. Ansari, and I.S. Currim, *Hierarchical bayes versus finite mixture conjoint analysis models: A comparison of fit, prediction, and partworth recovery*, Journal of Marketing Research (2002), 87–98.

[6] R.L. Andrews and I.S. Currim, *A comparison of segment retention criteria for finite mixture logit models*, Journal of Marketing Research (2003), 235–243.

[7] A. Ansari and C.F. Mela, *E-customization*, Journal of Marketing Research (2003), 131–145.

[8] A. Argyriou, T. Evgeniou, and M. Pontil, *Convex multi-task feature learning*, Machine Learning **73** (2008), no. 3, 243–272.

[9] D.R. Atkins and P.O. Iyogun, *Periodic versus 'can-order' policies for coordinated multi-item inventory systems*, Management Sci. (1988), 791–796.

[10] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski, *Convex optimization with sparsity-inducing norms*, Optimization for Machine Learning (2011), 19–53.

[11] Nada Nasr Bechwati and Wendy Schneier Siegal, *The impact of the prechoice process on product returns*, Journal of marketing research (2005), 358–367.

[12] Alexandre Belloni, Daniel Chen, Victor Chernozhukov, and Christian Hansen, *Sparse models and methods for optimal instruments with an application to eminent domain*, Econometrica **80** (2012), no. 6, 2369–2429.

[13] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust optimization*, Princeton Series in Applied Mathematics, Princeton University Press, 2009.

[14] A. Ben-Tal, B. Golany, A. Nemirovski, and J.P. Vial, *Retailer-supplier flexible commitments contracts: A robust optimization approach*, Manufacturing Service Oper. Management **7** (2005), 248–271.

[15] A. Ben-Tal and A. Nemirovski, *Robust convex optimization*, Math. Oper. Res. **23** (1998), 769–805.

[16] _____, *Robust solutions of uncertain linear programs*, Oper. Res. Lett. **25** (1999), 1–13.

[17] _____, *Robust solutions of linear programming problems contaminated with uncertain data*, Math. Programming, Ser. A **88** (2000), 411–424.

[18] J. Benders, *Partitioning procedures for solving mixed variables programming problems*, Numerische Mathematik **4** (1962), 238–252.

[19] D. Bertsekas, *Dynamic programming and optimal control*, vol. 1, Athena Scientific, 1995.

[20] D. Bertsimas, D. Iancu, and P. Parrilo, *A hierarchy of near-optimal policies for multi-stage adaptive optimization*, Submitted for publication (2009).

[21] _____, *Optimality of affine policies in multistage robust optimization*, Math. Oper. Res. **35** (2010), no. 2, 363–394.

[22] D. Bertsimas and M. Sim, *Robust discrete optimization and network flows*, Math. Programming, Ser. B **98** (2003), 49–71.

[23] _____, *The price of robustness*, Oper. Res. **52** (2004), 35–53.

[24] D. Bertsimas and A. Thiele, *A robust optimization approach to inventory theory*, Oper. Res. **54** (2006), 150–168.

[25] D. Bienstock and N. Ozbay, *Computing robust basestock levels*, Discrete Optimization **5** (2008), 389–414.

[26] Dave Blanchard, *Moving forward in reverse*, Logistics Today **46** (2005), no. 7.

[27] S. Boyd and L. Vandenberghe, *Convex optimization*, Cambridge university press, 2004.

[28] G. Celeux, M. Hurn, and C.P. Robert, *Computational and inferential difficulties with mixture posterior distributions*, Journal of the American Statistical Association **95** (2000), no. 451, 957–970.

[29] X. Chen, Q. Lin, S. Kim, J.G. Carbonell, and E.P. Xing, *Smoothing proximal gradient method for general structured sparse regression*, The Annals of Applied Statistics **6** (2012), no. 2, 719–752.

[30] X. Chen and D. Simchi-Levi, *Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case*, Oper. Res. (2004), 887–896.

[31] _____, *Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The infinite horizon case*, Math. Oper. Res. (2004), 698–723.

[32] P.K. Chintagunta, D.C. Jain, and N.J. Vilcassim, *Investigating heterogeneity in brand preferences in logit models for panel data*, Journal of Marketing Research (1991), 417–428.

[33] A. Clark and H. Scarf, *Optimal policies for a multi-echelon inventory problem*, Management Sci. **6** (1960), no. 4, 475–490.

[34] D. Cui and D. Curry, *Prediction in marketing using the support vector machine*, Marketing Science **24** (2005), no. 4, 595–615.

[35] Tülin Erdem and Michael P Keane, *Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets*, Marketing science **15** (1996), no. 1, 1–20.

[36] T. Evgeniou, C. Boussios, and G. Zacharia, *Generalized robust conjoint estimation*, Marketing Science **24** (2005), no. 3, 415–429.

[37] T. Evgeniou, M. Pontil, and O. Toubia, *A convex optimization approach to modeling consumer heterogeneity in conjoint estimation*, Marketing Science **26** (2007), no. 6, 805–818.

[38] A. Federgruen and A. Heching, *Combined pricing and inventory control under uncertainty*, Oper. Res. **47** (1999), no. 3, 454–475.

[39] G. Gallego and I. Moon, *The distribution free newsboy problem: review and extensions*, Journal of the Operational Research Society (1993), 825–834.

[40] G. Gallego, J. Ryan, and D. Simchi-Levi, *Minimax analysis for finite horizon inventory models*, IIE Trans. **33** (2001), 861–874.

[41] Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin, *Bayesian data analysis*, CRC press, 2013.

[42] P.E. Green and V. Srinivasan, *Conjoint analysis in marketing: new developments with implications for research and practice*, The Journal of Marketing (1990), 3–19.

[43] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning*, Springer Series in Statistics, 2001.

[44] T.D. Hocking, A. Joulin, F. Bach, and J.P. Vert, *Clusterpath: An algorithm for clustering using convex fusion penalties*, Proceedings of the 28th International Conference on Machine Learning (2011).

[45] W.T. Huh, G. Janakiraman, J.A. Muckstadt, and P. Rusmevichientong, *An adaptive algorithm for finding the optimal base-stock policy in lost sales inventory systems with censored demand*, Math. Oper. Res. **34** (2009), no. 2, 397–416.

[46] G. N. Iyengar, *Robust dynamic programming*, Math. Oper. Res. **30** (2005), no. 2, 257–280.

[47] R. Iyengar and K. Jedidi, *A conjoint model of quantity discounts*, Marketing Science **31** (2012), no. 2, 334–350.

[48] R. Iyengar, K. Jedidi, and R. Kohli, *A conjoint approach to multipart pricing*, Journal of Marketing Research **45** (2008), no. 2, 195–210.

[49] G. Janakiraman and R. Roundy, *Lost-sales problems with stochastic lead times: Convexity results for base-stock policies*, Oper. Res. **52** (2004), no. 5, 795–803.

[50] R. Jenatton, J.Y. Audibert, and F. Bach, *Structured variable selection with sparsity-inducing norms*, Journal of Machine Learning Research **12** (2011), 2777–2824.

[51] R. Jenatton, A. Gramfort, V. Michel, G. Obozinski, E. Eger, F. Bach, and B. Thirion, *Multiscale mining of fmri data with hierarchical structured sparsity*, SIAM Journal on Imaging Sciences **5** (2012), no. 3, 835–856.

[52] W.A. Kamakura and G.J. Russell, *A probabilistic choice model for market segmentation and elasticity structure*, Journal of Marketing Research (1989), 379–390.

[53] S. Karlin, *Dynamic inventory policy with varying stochastic demands*, Management Sci. (1960), 231–258.

[54] S. Karlin and H. Scarf, *Inventory models of the Arrow-Harris-Marschak type with time lag*, Studies in the Mathematical Theory of Inventory and Production (1958), 155–178.

[55] J.G. Kim, U. Menzefricke, and F.M. Feinberg, *Assessing heterogeneity in discrete choice models using a dirichlet process prior*, Review of marketing Science **2** (2004), no. 1.

[56] S. Kim and E.P. Xing, *Tree-guided group lasso for multi-response regression with structured sparsity, with an application to eqtl mapping*, The Annals of Applied Statistics **6** (2012), no. 3, 1095–1117.

[57] P.J. Lenk, W.S. DeSarbo, P.E. Green, and M.R. Young, *Hierarchical bayes conjoint analysis: Recovery of partworth heterogeneity from reduced experimental designs*, Marketing Science **15** (1996), no. 2, 173–191.

[58] R. Levi, G. Janakiraman, and M. Nagarajan, *Provably near-optimal balancing policies for stochastic inventory control models with lost sales*, Math. Oper. Res. **33** (2008), no. 2, 351–374.

[59] Yang Li and Asim Ansari, *A bayesian semiparametric approach for endogeneity and heterogeneity in choice models*, Management Science **60** (2013), no. 5, 1161–1179.

[60] Lan Luo, *Product line design for consumer durables: an integrated marketing and engineering approach*, Journal of Marketing Research **48** (2011), no. 1, 128–139.

[61] I. Moon and G. Gallego, *Distribution free procedures for some inventory models*, Journal of the Operational research Society (1994), 651–658.

[62] T.E. Morton, *Bounds on the solution of the lagged optimal inventory equation with no demand backlogging and proportional costs*, SIAM Rev. **11** (1969), no. 4, 572–596.

[63] _____, *The near-myopic nature of the lagged-proportional-cost inventory problem with lost sales*, Oper. Res. (1971), 1708–1716.

[64] A. Nilim and L. El Ghaoui, *Robust control of Markov decision processes with uncertain transition matrices*, Oper. Res. **53** (2005), no. 5, 780–798.

[65] Z. Qin and D. Goldfarb, *Structured sparsity via alternating direction methods*, The Journal of Machine Learning Research **98888** (2012), 1435–1468.

[66] U.S. Rao, *Properties of the periodic review (R, T) inventory control policy for stationary, stochastic demand*, Manufacturing Service Oper. Management **5** (2003), no. 1, 37–53.

[67] M.I. Reiman, *A new and simple policy for the continuous review lost sales inventory model*, Unpublished manuscript (2004).

[68] P.E. Rossi, R.E. McCulloch, and G.M. Allenby, *The value of purchase history data in target marketing*, Marketing Science **15** (1996), no. 4, 321–340.

[69] Peter E. Rossi, Greg M. Allenby, and Robert McCulloch, *Bayesian statistics and marketing*, J. Wiley & Sons, 2005.

[70] John Rust, *Structural estimation of markov decision processes*, Handbook of econometrics **4** (1994), no. 4.

[71] H. Scarf, *A min-max solution of an inventory problem*, Studies in the Mathematical Theory of Inventory and Production (1958), 201–209.

[72] H. Scarf, *The optimality of (s, S) policies in the dynamic inventory problem*, Mathematical Methods in the Social Sciences (K. Arrow, S. Karlin, and P. Suppes, eds.), Stanford University Press, Stanford, 1960, pp. 196–202.

[73] S. Sethi and F. Cheng, *Optimality of (s, S) policies in inventory models with markovian demand*, Oper. Res. **45** (1997), no. 6, 931–939.

[74] K.H. Shang and S.X. Zhou, *Optimizing replenishment intervals for two-echelon distribution systems with stochastic demand*, Working paper (2011).

[75] J. Shao, *Linear model selection by cross-validation*, Journal of the American Statistical Association **88** (1993), no. 422, 486–494.

[76] Jeffrey D Shulman, Anne T Coughlan, and R Canan Savaskan, *Optimal restocking fees and information provision in an integrated demand-supply model of product returns*, Manufacturing & Service Operations Management **11** (2009), no. 4, 577–594.

[77] J.S. Song and P. Zipkin, *Inventory control in a fluctuating demand environment*, Oper. Res. (1993), 351–370.

[78] Y. Song, S. Ray, and T. Boyaci, *Optimal dynamic joint inventory-pricing control for multiplicative demand with fixed order costs and lost sales*, Oper. Res. **57** (2009), no. 1, 245–250.

[79] David J Spiegelhalter, Nicola G Best, Bradley P Carlin, and Angelika Van Der Linde, *Bayesian measures of model complexity and fit*, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **64** (2002), no. 4, 583–639.

[80] M. Stephens, *Dealing with label switching in mixture models*, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **62** (2000), no. 4, 795–809.

[81] R. Tibshirani, *Regression shrinkage and selection via the lasso*, Journal of the Royal Statistical Society. Series B (Methodological) (1996), 267–288.

[82] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight, *Sparsity and smoothness via the fused lasso*, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **67** (2004), no. 1, 91–108.

[83] O. Toubia, J.R. Hauser, and D.I. Simester, *Polyhedral methods for adaptive choice-based conjoint analysis*, Journal of Marketing Research (2004), 116–131.

[84] O. Toubia, D.I. Simester, J.R. Hauser, and E. Dahan, *Fast polyhedral adaptive conjoint estimation*, Marketing Science **22** (2003), no. 3, 273–303.

[85] K. Train, *Discrete choice methods with simulation*, Cambridge university press, 2003.

[86] V. Vapnik, *Statistical learning theory*, Wiley, 1998.

[87] A.F. Veinott Jr, *On the optimality of (s, S) inventory policies: New conditions and a new proof*, SIAM J. Appl. Math. (1966), 1067–1083.

[88] M. Vriens, M. Wedel, and T. Wilms, *Metric conjoint segmentation methods: A monte carlo comparison*, Journal of Marketing Research (1996), 73–85.

[89] G. Wahba, *Spline models for observational data*, Society for Industrial and Applied Mathematics, 1990.

[90] D.R. Wittink and P. Cattin, *Commercial use of conjoint analysis: an update*, The Journal of Marketing (1989), 91–96.

[91] Stacy L Wood, *Remote purchase environments: the influence of return policy leniency on two-stage decision processes*, Journal of Marketing Research (2001), 157–169.

[92] W. Yin, *Gurobi mex: A matlab interface for gurobi*, URL: http://convexoptimization.com/wikimization/index.php/gurobi_mex, 2009-2011.

[93] M. Yuan and Y. Lin, *Model selection and estimation in regression with grouped variables*, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **68** (2005), no. 1, 49–67.

[94] P. Zipkin, *Foundations of inventory management*, McGraw-Hill Higher Education, 2000.

[95] P. Zipkin, *Old and new methods for lost-sales inventory systems*, Oper. Res. **56** (2008), no. 5, 1256–1263.

[96] H. Zou and T. Hastie, *Regularization and variable selection via the elastic net*, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **67** (2005), no. 2, 301–320.

Appendix A

Appendix for Chapter 2

A.1 Proof of Lemma 4

Proof. Proof. Since the term $K + c\tilde{u}$ is fixed constant, the adversary's problem is equivalent to $\max_{\mathbf{d}\in\mathcal{D}} H(\tilde{u},\mathbf{d})$. We first consider the case with backlogging dynamics. Consider the following two cases:

1. $x+\tilde{u} \leq 0$: Since $\mathbf{d} \geq 0$, $x+\tilde{u} \leq 0$ implies $x+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell \leq 0$, for all $\tau \leq t \leq \tau+\xi-1$. Thus, $H(\tilde{u},\mathbf{d}) = \sum_{t=\tau}^{\tau+\xi-1} \max\left\{h(x+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell), -b(x+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell)\right\} = -\sum_{t=\tau}^{\tau+\xi-1} b(x+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell)$. Thus

$$\max_{\mathbf{d}\in\mathcal{D}} H(\tilde{u},\mathbf{d}) = \max_{\mathbf{d}\in\mathcal{D}} \Big\{ -\sum_{t=\tau}^{\tau+\xi-1} b\big(x+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell\big)\Big\}. \tag{A.1}$$

Since $\mathcal{D}$ is a polyhedron, (A.1) is LP. In this case, we only have to solve a single LP.

2. $x+\tilde{u} > 0$: It is clear that

$$\max_{\mathbf{d}\in\mathcal{D}} H(\tilde{u},\mathbf{d}) = \max_{r\in\{0,1,\ldots,\xi\}} \Big\{ \max_{\mathbf{d}\in\mathcal{D}^r} H(\tilde{u},\mathbf{d}) \Big\},$$

where

$$\mathcal{D}^r = \begin{cases} \Big\{ \mathbf{d}\in\mathcal{D} \;\Big|\; x+\tilde{u}-d_\tau \le 0 \Big\}, & r=0, \\ \Big\{ \mathbf{d}\in\mathcal{D} \;\Big|\; x+\tilde{u}-\sum_{\ell=\tau}^{\tau+r-1} d_\ell \ge 0,\; x+\tilde{u}-\sum_{\ell=\tau}^{\tau+r} d_\ell \le 0 \Big\}, & 1\le r\le \xi-1, \\ \Big\{ \mathbf{d}\in\mathcal{D} \;\Big|\; x+\tilde{u}-\sum_{\ell=\tau}^{\tau+\xi-1} d_\ell \ge 0 \Big\}, & r=\xi. \end{cases} \tag{A.2}$$

Fix r. Then

$$\begin{aligned} \max_{\mathbf{d}\in\mathcal{D}^r} H(\tilde{u},\mathbf{d}) &= \max_{\mathbf{d}\in\mathcal{D}^r} \left\{ \sum_{t=\tau}^{\tau+\xi-1} \max\left\{ h(x+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell), -b(x+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell) \right\} \right\} \\ &= \max_{\mathbf{d}\in\mathcal{D}^r} \left\{ \sum_{t=\tau}^{\tau+r-1} h(x+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell) - \sum_{t=\tau+r}^{\tau+\xi-1} b(x+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell) \right\}, \end{aligned} \tag{A.3}$$

where the last equality follows from the structure of $\mathcal{D}^r$. Since $\mathcal{D}^r$ is polyhedron for all r, it follows that (A.3) is an LP, and the adversary's problem reduces to solving $\xi+1$ LPs.

We now consider the case with lost-sales dynamics. In this case, $x+\tilde{u}$ is always non-negative. From Lemma 1, it follows that

$$\max_{\mathbf{d}\in\mathcal{D}^r} H(\tilde{u},\mathbf{d}) = \begin{cases} \max_{\mathbf{d}\in\mathcal{D}^r}\left\{\sum_{t=\tau}^{\tau+r-1} h\Big(x+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell\Big) - b\Big(x+\tilde{u}-\sum_{t=\tau}^{\tau+\xi-1} d_t\Big)\right\}, & 0\le r\le \xi-1, \\ \max_{\mathbf{d}\in\mathcal{D}^r}\left\{\sum_{t=\tau}^{\tau+\xi-1} h\Big(x+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell\Big)\right\}, & r=\xi, \end{cases}$$

where $\mathcal{D}^r$ is given by (A.2). Since each $\mathcal{D}^r$ is a polyhedron, it follows that $\max_{\mathbf{d}\in\mathcal{D}^r} H(\tilde{u},\mathbf{d})$ is an LP for all r; thus, the adversary's problem reduces to solving $\xi+1$ LPs. □

A.2 Proof of Lemma 6

Proof. Proof. Since $K+c\tilde{u}$ is a constant, we can simply omit it and rewrite the adversary's problem as

$$\max_{\mathbf{d}=(\mathbf{d}^O,\mathbf{d}^I)\in\mathcal{D}} H\Big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\Big). \tag{A.1}$$

In the backlogging case $x(\mathbf{d}^O) = x_\tau + \sum_{t=\tau}^{\tau+L-1}(\bar{u}_{t-L}-d_t)$. Therefore, (A.1) is equivalent to

$$\max_{\mathbf{d}\in\mathcal{D}} \sum_{t=\tau+L}^{\tau+L+\xi-1} \max\left\{h\Big(x_\tau+\sum_{\ell=\tau-L}^{\tau-1}\bar{u}_\ell+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell\Big), -b\Big(x_\tau+\sum_{\ell=\tau-L}^{\tau-1}\bar{u}_\ell+\tilde{u}-\sum_{\ell=\tau}^{t} d_\ell\Big)\right\}. \tag{A.2}$$

Let $q = x_\tau + \sum_{\ell=\tau-L}^{\tau-1}\bar{u}_\ell + \tilde{u}$. We consider the following two cases.

1. $q\le 0$: Since $q-\sum_{\ell=\tau}^{t} d_\ell \le 0$, for all $\tau+L\le t\le \tau+L+\xi-1$, it follows that (A.2) reduces to the single LP, $\max_{\mathbf{d}\in\mathcal{D}}\Big\{-\sum_{t=\tau+L}^{\tau+L+\xi-1} b\big(q-\sum_{\ell=\tau}^{t} d_\ell\big)\Big\}$.

2. $q > 0$: As in Lemma 4, (A.1) is equivalent to $\max_{r\in\{0,\ldots,\xi\}} \max_{\mathbf{d}\in\mathcal{D}^r} \Big\{H\Big(x\big(\mathbf{d}^O\big),\, \tilde{u},\, \mathbf{d}^I\Big)\Big\}$, where

$$\mathcal{D}^r = \begin{cases} \Big\{\mathbf{d}\in\mathcal{D} \mid q - \sum_{\ell=\tau}^{\tau+L} d_\ell \le 0\Big\}, & r = 0, \\ \Big\{\mathbf{d}\in\mathcal{D} \mid q - \sum_{\ell=\tau}^{\tau+L+r-1} d_\ell \ge 0,\, q - \sum_{\ell=\tau}^{\tau+L+r} d_\ell \le 0\Big\}, & 1 \le r \le \xi-1, \\ \Big\{\mathbf{d}\in\mathcal{D} \mid q - \sum_{\ell=\tau}^{\tau+L+\xi-1} d_\ell \ge 0\Big\}, & r = \xi. \end{cases}$$

Since the definition of $\mathcal{D}^r$ implies that

$$\max_{\mathbf{d}\in\mathcal{D}^r} \Big\{H\Big(x\big(\mathbf{d}^O\big),\, \tilde{u},\, \mathbf{d}^I\Big)\Big\} = \max_{\mathbf{d}\in\mathcal{D}^r} \Big\{ \sum_{t=\tau+L}^{\tau+L+r-1} h\big(q - \sum_{\ell=\tau}^{t} d_\ell\big) - \sum_{t=\tau+L+r}^{\tau+L+\xi-1} b\big(q - \sum_{\ell=\tau}^{t} d_\ell\big)\Big\},$$

and each $\mathcal{D}^r$ is polyhedron, $\max_{\mathbf{d}\in\mathcal{D}^r} \Big\{H\Big(x\big(\mathbf{d}^O\big),\, \tilde{u},\, \mathbf{d}^I\Big)\Big\}$ is an LP, and the adversary's problem (A.2) reduces to solving $\xi+1$ LPs.

Next, we consider the case with lost-sales dynamics. Define

$$\theta^j(\mathbf{d}^O) = \begin{cases} x_\tau + \sum_{t=\tau}^{\tau+L-1} (\bar{u}_{t-L} - d_t), & j = 0, \\ \sum_{t=j+\tau}^{\tau+L-1} (\bar{u}_{t-L} - d_t), & 1 \le j \le L-1, \\ 0, & j = L. \end{cases}$$

Then the inventory level $x(\mathbf{d}^O) = \max\Big\{\theta^0(\mathbf{d}^O), \theta^1(\mathbf{d}^O), \ldots, \theta^L(\mathbf{d}^O)\Big\}$. Using this respresentation for $x(\mathbf{d}^O)$ we partition $\mathcal{D} = \cup_{j=0}^{L} \mathcal{E}^j$, where

$$\mathcal{E}^j \triangleq \Big\{\mathbf{d}\in\mathcal{D} \mid x(\mathbf{d}^O) = \theta^j(\mathbf{d}^O)\Big\} = \Big\{\mathbf{d}\in\mathcal{D} \mid \theta^j(\mathbf{d}^O) \ge \theta^\ell(\mathbf{d}^O), \forall\, \ell \ne j\Big\}.$$

We further partition $\mathcal{E}^j = \bigcup_{r=0}^{\xi} \mathcal{E}^{j,r}$, where

$$\mathcal{E}^{j,r} = \begin{cases} \Big\{ \mathbf{d} \in \mathcal{E}^j |\ \theta^j(\mathbf{d}^O) + \tilde{u} - d_{\tau+L} \le 0 \Big\}, & r = 0, \\ \Big\{ \mathbf{d} \in \mathcal{E}^j |\ \theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{\ell=\tau+L}^{\tau+L+r-1} d_\ell \ge 0, \theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{\ell=\tau+L}^{\tau+L+r} d_\ell \le 0 \Big\}, & 1 \le r \le \xi - 1, \\ \Big\{ \mathbf{d} \in \mathcal{E}^j |\ \theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{\ell=\tau+L}^{\tau+L+\xi-1} d_\ell \ge 0 \Big\}, & r = \xi. \end{cases}$$

Then $\max_{\mathbf{d}=(\mathbf{d}^O,\mathbf{d}^I)\in\mathcal{D}} H\big(x\big(\mathbf{d}^O\big),\, \tilde{u},\, \mathbf{d}^I\big) = \max_{j\in\{0,\ldots,L\}} \max_{r\in\{0,\ldots,\xi\}} \max_{\mathbf{d}\in\mathcal{E}^{j,r}} H\big(x\big(\mathbf{d}^O\big),\, \tilde{u},\, \mathbf{d}^I\big)$, and from the definition of $\mathcal{E}^{j,r}$ it follows that

$$\max_{\mathbf{d}\in\mathcal{E}^{j,r}} H\big(x\big(\mathbf{d}^O\big),\, \tilde{u},\, \mathbf{d}^I\big) = \begin{cases} \max_{\mathbf{d}\in\mathcal{E}^{j,r}} \Big\{ \sum_{t=\tau+L}^{\tau+L+r-1} h\Big(\theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{\ell=\tau+L}^{t} d_\ell\Big) \\ \qquad\qquad - b\Big(\theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{t=\tau+L}^{\tau+L+\xi-1} d_t\Big) \Big\}, & 0 \le r \le \xi - 1, \\ \max_{\mathbf{d}\in\mathcal{E}^{j,r}} \Big\{ \sum_{t=\tau+L}^{\tau+L+\xi-1} h\Big(\theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{\ell=\tau+L}^{t} d_\ell\Big) \Big\}, & r = \xi. \end{cases}$$

Since each $\theta^j(\mathbf{d}^O)$ is linear function of $\mathbf{d}^O$, and therefore of $\mathbf{d}$, $\mathcal{E}^{j,r}$ is polyhedron and $H\big(x\big(\mathbf{d}^O\big),\, \tilde{u},\, \mathbf{d}^I\big)$, when restricted to $\mathcal{E}^{j,r}$, is a linear function of $\mathbf{d}$, it follows that $\max_{\mathbf{d}\in\mathcal{E}^{j,r}} H\big(x\big(\mathbf{d}^O\big),\, \tilde{u},\, \mathbf{d}^I\big)$ is an LP. Thus, the adversary's problem reduces to solving $(\xi+1)(L+1)$ LPs. □

A.3 Joint Pricing and Inventory Control: Multiplicative Demand Case

In this section, we assume that in the joint pricing and inventory control model the demand D_k is given by a multiplicative stochastic demand function $Q_k^{(m)}\big(p_k, z_k\big)$ defined as follows.

Assumption A.3.1. *The multiplicative stochastic demand function* $Q_k^{(m)}(p_k, z_k) = \gamma_k(p_k) \cdot z_k$, *where the functions* $\{\gamma_k(\cdot)\}$ *and the random variables* $\{z_k\}$ *satisfy the following two conditions.*

1. *The function* $\gamma_k(\cdot)$ *is continuous and strictly decreasing on* $\left[p_{min}^k,\ p_{max}^k\right]$. *Let* $\rho_k(\cdot)$ *denote the inverse of* $\gamma_k(\cdot)$. *We assume* $\rho_k(\omega) \cdot \omega$ *is concave in* ω *on* $\left[\gamma_k(p_{max}^k),\ \gamma_k(p_{min}^k)\right]$.

2. *The random variables* $\{z_k\}_{k=1}^T$ *are nonnegative and mutually independent, and only the mean and variance of* z_k *are known.*

Assumption A.3.1(1) is a common assumption on stochastic demand function (see [30]). The multiplicative demand function here includes linear demand curve $\gamma(p) = \alpha - \beta p$ $(\alpha > 0, \beta > 0)$, log-linear demand curve $\gamma(p) = e^{\alpha - \beta p}$ $(\alpha > 0, \beta > 0)$ and iso-elastic demand curve $\gamma(p) = \alpha p^{-\beta}$ $(\alpha > 0, \beta > 1)$ as special cases.

Our policy can be extended to this case by making the following simple modification. Assumption A.3.1(1) implies that, instead of selecting the price p_k, the inventory manager can equivalently select the *nominal demand* $\omega_k = \gamma_k(p_k) \in \left[\gamma_k(p_{\max}^k),\ \gamma_k(p_{\min}^k)\right]$. In particular, at the first period τ of a new ordering cycle, the inventory manager chooses order quantity u_τ, cycle length ξ and nominal demand sequence $\omega[\tau, \xi]$ to minimize the worst-case average cost over the cycle $(\tau, \tau+1, \ldots, \tau+\xi-1)$, $F(u_\tau, \xi, \omega[\tau, \xi])$, which is defined as

$$F(u_\tau, \xi, \omega[\tau, \xi]) = \max_{z[\tau] \in \mathcal{Z}[\tau]} \left\{ \frac{K\mathbf{1}(u_\tau > 0) + cu_\tau + W(u_\tau, \omega[\tau, \xi], z[\tau, \xi])}{\xi} \right\},$$

where

$$W(u_\tau, \omega[\tau,\xi], z[\tau,\xi]) = \sum_{t=\tau}^{\tau+\xi-1} \max\left\{h(x_\tau + u_\tau - \sum_{\ell=\tau}^{t} \omega_\ell z_\ell), -b(x_\tau + u_\tau - \sum_{\ell=\tau}^{t} \omega_\ell z_\ell)\right\} - \sum_{t=\tau}^{\tau+\xi-1} \rho_t(\omega_t)\omega_t z_t.$$

The optimal decision $(u^*_\tau, \xi^*, \omega^*[\tau,\xi])$ is chosen by solving the optimization problem

$$\min_{\xi \in \left\{1,2,\ldots,\min\{T-\tau+1,U\}\right\}} \quad \min_{\left\{0 \le u_\tau \le M - x_\tau, \omega[\tau,\xi] \in \Omega[\tau,\xi]\right\}} F(u_\tau, \xi, \omega[\tau,\xi]), \tag{A.1}$$

where $\Omega[\tau,\xi] = \left\{\omega[\tau,\xi] \mid \omega_t \in [\gamma_t(p^t_{\max}),\ \gamma_t(p^t_{\min})],\ \text{for } \tau \le t \le \tau+\xi-1\right\}$.

We use the solution algorithm which is identical to that of (21) of the regular text to solve (A.1). In particular, we apply Benders' decomposition to solve the counterpart of program (23) of the regular text. We show that the inventory manager's problem and the adversary's problem can be solved efficiently.

Proposition A.3.2. *The inventory manager's problem is* $\min_{\left\{0 \le u \le M-x, \omega \in \Omega\right\}} \max_{z \in \tilde{Z}} \left\{K + cu + W(u,\omega,z)\right\}$*, where* $\tilde{Z}$ *is a given finite working list of random terms. The inventory manager's problem is a convex optimization problem.*

Proof. Proof. Simply note that $W(u,\omega,z)$ is convex in (u,ω) for any non-negative z. □

Proposition A.3.3. *The adversary's problem is* $\max_{z \in Z} \left\{K + c\tilde{u} + W(\tilde{u}, \tilde{\omega}, z)\right\}$*, where* $(\tilde{u}, \tilde{\omega})$ *denotes the current working decision of the inventory manager. The adversary's problem can be reduced to solving at most* $\xi + 1$ *LPs.*

Proof. Proof. Similar to that of Lemma 4. □

Thus we have completed our algorithm.

A.4 Numerical experiments: Lost-sales model with positive lead time

This section contains the numerical experiments of Section 6.2.2.. We adopt the simulation design in [95]. The parameters for this model were as follows: fixed cost $K = 0$, variable ordering cost $c = 0$, inventory holding cost $h = 1$, lost-sales penalty cost $b \in \{4, 9, 19, 39\}$, lead time $L \in \{1, 2, 3, 4\}$, and demand distribution D_t was either Poisson with mean 5 or geometric with mean 5. For details, see Section 6 of [95].

Since [95] considered the long-run average cost criterion, and our cycle-based policy CI is designed for finite horizon inventory problems, we used the batch means method to approximate the long-run average cost. The specific simulation strategy was as follows: for each combination of parameter values, we simulated *one* sample path under the CI policy for a total time horizon of $T = 50,000$ periods. We discarded the first $T_b = 10,000$ periods as the "burn-in" periods. We equally divided the rest $T_s = 40,000$ periods into $N = 10$ batches, each consists of $T_s/N = 4,000$ consecutive periods. We computed the average cost over each batch and then the mean and standard deviation across the $N = 10$ batches. We report the mean as an estimator for the long-run average cost of CI.

Since in all the cases we considered, the mean demand was 5, we set $\bar{\omega}_t = 5$. We set $\Gamma^t_j = \sqrt{j}$,

for $1 \leq t \leq T$ and $1 \leq j \leq T-t+1$, and set $U = 5$. The deviation $\hat{\omega}_t$ was selected as a function of the penalty cost b and the demand distribution. The specific choice for $\hat{\omega}_t$ are displayed in Table A.1. Note that the choice of $\hat{\omega}_t$ is independent of the lead time – this was done for

Table A.1: Choice of $\hat{\omega}_t$.

	b			
Demand	4	9	19	39
Poisson	2	2	3	4
Geometric	3	5	5	5

simplicity and one can further improve performance by choosing a lead time dependent $\hat{\omega}_t$. Our choice of $\hat{\omega}_t$ is non-decreasing in b, and for a given b, we chose a larger $\hat{\omega}_t$ for the geometric case. This choice is motivated by the fact that as the lost-sales penalty b increases the decision maker should prefer a "fatter" uncertainty set to protect against the very expensive stock-out events more aggressively. And since the geometric distribution has heavier tails, the uncertainty set is chosen to be larger in this case.

Poisson demand. We summarize the performance of CI when demand is Poisson in Table A.2. We report the optimal average cost taken from [95], the long-run average cost for CI, and the amount by which the CI cost exceeds the optimal cost as a percentage of the optimal cost. We also report the standard deviation of the batch averages in parenthesis. For the sake of brevity, we only report the rank of CI among the nine heuristics, i.e., the original eight heuristics in [95] and CI, for each set of parameter values. We refer the interested readers to [95] for the cost for each of the other eight heuristics.

Recall that the optimal policy and the eight heuristics in [95] utilize the full distribution

Table A.2: Poisson case.

	Lead time							
	1		2		3		4	
(a) $b = 4$								
Optimal	4.04		4.40		4.60		4.73	
CI	4.07	+0.74%	4.49	+2.00%	4.86	+5.66%	5.13	+8.42%
	(0.05)	2nd	(0.07)	4th	(0.06)	5th	(0.12)	5th
(b) $b = 9$								
Optimal	5.44		6.09		6.53		6.84	
CI	5.83	+7.12%	6.23	+2.25%	6.70	+2.62%	7.00	+2.37%
	(0.10)	8th	(0.11)	4th	(0.14)	3rd	(0.10)	4th
(c) $b = 19$								
Optimal	6.68		7.66		8.36		8.89	
CI	7.19	+7.62%	7.75	+1.22%	8.61	+2.93%	9.08	+2.14%
	(0.17)	7th	(0.25)	2nd	(0.25)	5th	(0.18)	3rd
(d) $b = 39$								
Optimal	7.84		9.11		10.04		10.79	
CI	8.21	+4.66%	9.42	+3.39%	10.37	+3.24%	10.93	+1.27%
	(0.16)	7th	(0.29)	7th	(0.30)	5th	(0.25)	2nd

information of demand and specifically optimize the long-run average cost. In contrast, CI only uses partial demand information and adopts a min-max approach to compute the order decisions. Consequently, the simulations favor of the polices in [95] over CI. The results in Table A.2 indicate that the performance of CI is still quite good. In 12 out of the 16 parameter combinations, the CI cost exceeds the optimal cost by no more than 5%. Note that when $b = 4$, i.e., the lost-sales penalty is small, the relative performance of CI deteriorates as lead time L increases; whereas for $b \in \{9, 19, 39\}$, the opposite is true, i.e., the relative performance improves with lead time. This behavior can be explained by considering the dynamics of the CI policy for positive lead times. Recall that the CI policy selects u_τ to balance the ordering

cost with the holding cost and lost-sales penalty over the *inventory cycle* $(\tau+L,\ldots,\tau+L+\xi-1)$. In order to counteract the possibility of stock-out due to the uncertain demand over the intervening interval $(\tau,\ldots,\tau+L-1)$, the CI policy tends to choose a larger value of u_τ when the lead time L is large. When the lost-sales penalty b is small, this higher inventory position becomes a "liability" because it forces the inventory manager to pay more inventory holding cost; however, when b is large, higher inventory position becomes an "asset" since it helps the inventory manager to avoid or alleviate the expensive stock out.

Geometric demand. Table A.3 summarizes the performance of CI when demand is geometric. CI performs quite well when $b \in \{4,9\}$, i.e., when the lost-sales penalty is modest. However, as b becomes large, the performance of CI deteriorates quickly. When $b = 39$, CI cost is significantly larger than the optimal cost and is beaten by most other heuristics.

The poor performance of CI for large b is because we use a symmetric uncertainty set while the true demand distribution is highly asymmetric and has a considerable probability mass outside of the uncertainty set which, in turn results, in a high lost-sales penalty. When $b = 39$, we found that the relatively infrequent stock-out events, which occur on average once every 7–9 periods, account for 70–80% of the total cost. In order to correct this we considered the following asymmetric uncertainty set:

$$\mathcal{D}[\tau] = \left\{ \mathbf{d}[\tau] \;\middle|\; \bar{\omega}_t - \hat{\omega}_t^l z_t \le d_t \le \bar{\omega}_t + \hat{\omega}_t^r z_t, \quad 0 \le z_t \le 1, \quad \sum_{\ell=\tau}^{t} z_\ell \le \Gamma^\tau_{t-\tau+1}, \quad \forall\, \tau \le t \le T \right\},$$

where $\hat{\omega}_t^l$ (resp. $\hat{\omega}_t^r$) denotes the *left* (resp. *right*) deviation. It is clear that we recover the

Table A.3: Geometric case.

	Lead time							
	1		2		3		4	
(a) $b = 4$								
Optimal	9.82		10.24		10.47		10.61	
CI	9.95	+1.28%	10.57	+3.26%	11.19	+6.84%	11.75	+10.74%
	(0.19)	3rd	(0.25)	4th	(0.23)	6th	(0.23)	6th
(b) $b = 9$								
Optimal	14.51		15.50		16.14		16.58	
CI	14.82	+2.13%	15.56	+0.37%	16.63	+3.06%	17.38	+4.85%
	(0.32)	5th	(0.38)	1st	(0.26)	3rd	(0.39)	3rd
(c) $b = 19$								
Optimal	19.22		20.89		22.06		22.95	
CI	22.70	+18.13%	22.36	+7.02%	23.25	+5.39%	23.95	+4.34%
	(1.16)	7th	(0.99)	6th	(0.68)	5th	(0.61)	4th
(d) $b = 39$								
Optimal	23.87		26.21		27.96		29.36	
CI	37.57	+57.37%	35.57	+35.69%	36.25	+29.64%	35.32	+20.31%
	(1.36)	9th	(1.00)	8th	(1.70)	7th	(1.49)	7th

symmetric uncertainty set (5) of the regular text by setting $\hat{\omega}_t^l = \hat{\omega}_t^r = \hat{\omega}_t$. We re-did the cases with $b = 19$ and 39, where CI with symmetric uncertainty sets performs poorly. To ensure a fair comparison, we tested CI with asymmetric uncertainty sets on the same sample paths that we tested CI with symmetric uncertainty sets. In our numerical experiments, we set $\hat{\omega}_t^l = 5$ and $\hat{\omega}_t^r = 10$. We report the results in Table A.4. The new results are quite encouraging. When $b = 39$, the cost of CI with asymmetric uncertainty sets exceeds the optimal cost by no more than 5%.

These experiments indicate that tailoring the uncertainty sets by using additional information of the demand distribution, e.g., whether the demand is asymmetric about the mean, the tail

Table A.4: Geometric case continued.

	Lead time							
	1		2		3		4	
(c) $b = 19$								
Optimal	19.22		20.89		22.06		22.95	
CI(Asy)	19.43	+1.08%	21.41	+2.47%	23.58	+6.87%	25.58	+11.44%
	(0.63)	5th	(0.55)	5th	(0.25)	5th	(0.48)	5th
(d) $b = 39$								
Optimal	23.87		26.21		27.96		29.36	
CI(Asy)	25.00	+4.75%	26.43	+0.82%	28.77	+2.90%	30.20	+2.87%
	(0.75)	7th	(0.67)	3rd	(1.03)	5th	(0.94)	4th

decays slowly, etc., and problem parameters, e.g., lost-sales penalty cost b, results in a close-to-optimal performance for a very large set of scenarios.

Appendix B

Appendix for Chapter 3

B.1 The Link Between Metric-SEG and the Key Convex Optimization Problem (Equation (3)) of [37]

In Section 3.2.1, we point out that specifying $\lambda = +\infty$ reduces Metric-SEG to the key convex optimization problem (Equation (3)) of [37]. To see this link, we note that with $\lambda = +\infty$ any optimal solution of Metric-SEG must satisfy $\beta_i^S = \beta_k^S$ for $1 \leq i < k \leq I$. We use $\beta_0 \in \mathbb{R}^p$ to denote this common vector, which is interpreted as the population mean, and Metric-SEG reduces to the following optimization problem:

$$\begin{aligned} \min \quad & \sum_{i=1}^{I} ||Y_i - X_i\beta_i||_2^2 + \gamma \sum_{i=1}^{I} (\beta_i - \beta_0)^T D^{-1} (\beta_i - \beta_0), \\ \text{s.t.} \quad & D \text{ is a positive semidefinite matrix scaled to have trace 1,} \\ & \beta_i \in \mathbb{R}^p, \text{ for } i = 1, 2, \ldots, I; \quad \beta_0 \in \mathbb{R}^p, \end{aligned} \tag{B.1}$$

which is Equation (3) of [37].

B.2 The Extension of the Specification for $\left\{\theta_{ik}\right\}$

In Section 3.2.1, we note that information sources other than the conjoint data, e.g., consumers' demographic variables, can be readily incorporated in the specification for $\left\{\theta_{ik}\right\}$. For the purpose of illustration, we use demographic information as an example and assume that consumers' heterogeneous demographic variables $\tau_i \in \mathbb{R}^{p'}$ are available. We introduce the following specification for $\left\{\theta_{ik}\right\}$:

$$\theta_{ik} = e^{-(\omega||\bar{\beta}_i - \bar{\beta}_k||_2 + \omega'||\tau_i - \tau_k||_2)}, \tag{B.1}$$

where $\omega' \geq 0$ is a new regularization parameter controlling the impact of $\left\{\tau_i\right\}_{i=1}^{I}$. We note that when $\omega' = 0$ we recover the specification in Equation (3.2.3). In Equation (B.1), $\left\{\tau_i\right\}_{i=1}^{I}$ enter $\left\{\theta_{ik}\right\}$ in a similar manner as $\left\{\bar{\beta}_i\right\}_{i=1}^{I}$, in that pairs of respondents with more similar demographic variables are deemed as more likely to be drawn from the same segment and hence are penalized more heavily than those with less similar demographic variables. This way, we take consumers' demographic variables into account when recovering candidate segmentations.

B.3 Empirical Evidence for Bias in $\left\{B(\Gamma)\right\}$

In Section 3.2.1, we argue that $\left\{B(\Gamma)\right\}$, the set of individual-level partworths estimates obtained by solving Metric-SEG for each $\Gamma \in \Theta$, does not accurately characterize MCH since $B(\Gamma)$'s are biased. In this section, we provide empirical evidence for bias in $\left\{B(\Gamma)\right\}$ using simulation

experiments. The basic setup of the simulation experiments is identical to that in Section 3.3.1 and we refer the reader to Section 3.3.1 for details. We consider 3 experimental conditions in Section 3.3.1, in which we set Factor 1 (the number of segments) at 2, Factor 2 (the number of profiles per respondents) at 18, Factor 4 (the error variance) at 0.5, and vary Factor 3 (the within-segment variances of distributions) at 0.05, 0.10, and 0.20. We randomly generate 5 data sets for each experimental condition.

For each data set, we implement the first stage of the sparse learning (SL) approach for metric conjoint analysis, termed as Metric-SL, and obtain $\{B(\Gamma)\}$. We measure parameter recovery for each $B(\Gamma)$ by computing the RMSE(β) defined in Section 3.3.1 for each respondent and summarize the average RMSE(β) across respondents. We select the lowest average RMSE(β) across $B(\Gamma)$'s for each data set, and report the mean of the lowest average RMSE(β)'s across 5 data sets for each experimental condition. Consequently, the RMSE(β) we report for the first stage of Metric-SL provides a lower bound for that of any $B(\Gamma)$. As a benchmark, we also report the mean of the average RMSE(β)'s incurred by the individual-level partworths estimates obtained by the complete Metric-SL approach across the same 5 data sets for each experimental condition. The results are presented in Table B.1. It is clear from Table B.1 that

Table B.1: RMSE(β): Evidence for bias in $\{B(\Gamma)\}$

Num-S	Num-P	EV	WSV	The First Stage of Metric-SL	The Complete Metric-SL Approach
2	18	0.5	0.05	0.2576	0.1921
			0.10	0.2806	0.2401
			0.20	0.3097	0.2879

the complete Metric-SL approach leads to better parameter recovery compared to the first stage

alone, suggesting that $\{B(\Gamma)\}$ is indeed biased and fails to provide accurate individual-level representations of MCH. It also suggests that the inclusion of the second stage in the SL approach is critical for accurately recovering individual-level partworths in the presence of MCH.

B.4 The Normal Component Mixture (NCM) Model

In Section 3.3.1 and Section 3.4.1, we use the following specification of the NCM model for metric conjoint analysis:

Likelihood:

$$y_{ij} = x_{ij}\beta_i + \varepsilon_{ij},$$
$$\varepsilon_{ij} \sim N(0, \sigma^2);$$

First-stage prior:

$$\beta_i \sim N\left(\mu_{\text{Ind}_i}, \Sigma_{\text{Ind}_i}\right),$$
$$\text{Ind}_i \sim \text{Multinomial}(\text{pvec});$$

Second-stage prior:

$$\text{pvec} \sim \text{Dirichlet}(\alpha),$$
$$\mu_k \sim N\left(\bar{\mu}, \Sigma_k \otimes a_\mu^{-1}\right), \quad k = 1, 2, \ldots, K,$$
$$\Sigma_k \sim \text{IW}(v, V),$$
$$\sigma^2 \sim \text{IG}(r_0/2, s_0/2).$$

Our choice of the parameters for the second-stage prior follows the literature and aims to induce diffuse second-stage priors. In particular, we set $\alpha_k = 2$, for $k = 1, 2, \ldots, K$, $\bar{\mu} = 0$, $a_\mu = 1/16$,

$v = p+3$, $V = vI$, $r_0 = 1$, and $s_0 = 1$. For any fixed number of components K, we use the Gibbs sampler to generate draws from the posterior distribution. We execute the Gibbs sampler for 30000 iterations, using the first 15000 iterations as the burn-in period and the last 15000 iterations to obtain parameter estimates. In addition, we estimate the NCM model for $K \in \{1,2,\ldots,10\}$ and select K using the deviance information criterion (DIC) [60, 79].

The NCM model for choice-based conjoint analysis is obtained by modifying the likelihood of the NCM model for metric conjoint analysis. More specifically, in Section 3.3.2 and Section 3.4.2, we use the following specification of the NCM model for choice-based conjoint analysis:

Likelihood:

$$\text{Prob}\big(x_{ijh*} \text{ is chosen}\big) = \frac{e^{x_{ijh*}\beta_i}}{\sum_{h=1}^{H} e^{x_{ijh}\beta_i}};$$

First-stage prior:

$$\beta_i \sim N\big(\mu_{\text{Ind}_i}, \Sigma_{\text{Ind}_i}\big),$$

$$\text{Ind}_i \sim \text{Multinomial}(\text{pvec});$$

Second-stage prior:

$$\text{pvec} \sim \text{Dirichlet}(\alpha),$$

$$\mu_k \sim N\big(\bar{\mu}, \Sigma_k \otimes a_\mu^{-1}\big), \quad k = 1,2,\ldots,K,$$

$$\Sigma_k \sim \text{IW}(v,V).$$

Similar to the metric conjoint case, we set $\alpha_k = 2$, for $k = 1,2,\ldots,K$, $\bar{\mu} = 0$, $v = p+3$, and $V = vI$. In the choice-based conjoint application (Section 3.4.2) we set $a_\mu = 1/16$. On the other hand, in the choice-based simulation experiments (Section 3.3.2) we find that setting $a_\mu = 1/16$ leads to extreme parameter estimates for the NCM model on many data sets; following the suggestion

of [69], we resolve this issue by setting a_μ to a larger value, i.e., $a_\mu = 1/3$, in the choice-based simulation experiments. For any fixed number of components K, we use the Gibbs sampler to generate draws from the posterior distribution. We execute the Gibbs sampler for 30000 iterations, using the first 15000 iterations as the burn-in period and the last 15000 iterations to obtain parameter estimates. Again, we estimate the NCM model for $K \in \{1,2,\ldots,10\}$ and select K using DIC.

B.5 The Specification of the Finite Grids for the Regularization Parameters of the Sparse Learning Approach

In our sparse learning approach, the regularization parameters, i.e., $\Gamma = (\gamma, \lambda, \omega)$ for Metric-SEG (and Choice-SEG) and ψ^l for Metric-HET-General/Restricted (and Choice-HET-General/Restricted), are selected from pre-specified finite grids. We present in the following the specification of the finite grids used in Sections 3.3 and 3.4.

Metric Simulation (Section 3.3.1) and Metric Application (Section 3.4.1).

- $\gamma \in \{100\}$, $\lambda \in \{10^{\frac{3}{2}+\frac{v}{8}}\}_{v=0}^{19}$, $\omega \in \{3,5,7\}$, $\psi^l \in \{10^{-2+\frac{v}{8}}\}_{v=0}^{29}$.

Choice-based Simulation (Section 3.3.2) and Choice-based Application (Section 3.4.2).

- $\gamma \in \{100\}$, $\lambda \in \{10^{\frac{3}{2}+\frac{v}{8}}\}_{v=0}^{19}$, $\omega \in \{3,5,7\}$, $\psi^l \in \{10^{-2+\frac{v}{4}}\}_{v=0}^{14}$.

B.6 The Computation Time Comparison Between the Sparse Learning Approach and the NCM Model

We report the computation time of the sparse learning (SL) approach and the NCM model on data sets considered in Sections 3.3 and 3.4. We note that the computation time reported in the following is conditional on (1) the specification of the finite grids for the regularization parameters of the SL approach presented in B.5 and (2) the estimation procedure for the NCM model outlined in B.4.

B.6.1 Metric Conjoint Simulation (Section 3.3.1)

We report the computation time of both models for a subset of experimental conditions considered in Section 3.3.1. In each experimental condition, we average the computation time of each model over 5 data sets and report the average computation time in hours. The results are summarized in Table B.2. It is evident from Table B.2 that the computation time of Metric-SL is much lower than that of the NCM model for each experimental condition. The comparisons for other experimental conditions are qualitatively similar and are available from the authors upon request. We note that one limitation of our computation time comparison is that the computation time may not be compared meaningfully across experimental conditions, since multiple servers with different computational speeds were used to conduct the metric conjoint simulation studies and different experimental conditions might have been assigned to different servers. Consequently, we are not able to conduct a systematic investigation on the impact of the experimental factors on the computation time of both models. On the other hand, comparisons

of the computation time of the two models within each experimental condition is valid since all model estimations for the same experimental condition was done using a single server.

Table B.2: The Computation Time Comparison for Metric Conjoint Simulation

Num-S	Num-P	EV	WSV	Metric-SL	NCM
2	18	0.5	0.05	0.5212	4.9376
			0.10	0.5806	4.9207
			0.20	0.5481	4.8652
2	27	0.5	0.05	0.6259	4.8927
			0.10	1.0768	7.0740
			0.20	1.4648	7.0039

Notes. The computation time reported above is averaged over 5 randomly simulated data sets and is reported in hours.

B.6.2 Choice-based Conjoint Simulation (Section 3.3.2)

We report the computation time of both models for a subset of experimental conditions considered in Section 3.3.2. In each experimental condition, we average the computation time of each model over 5 data sets and report the average computation time in hours. The results are summarized in Table B.3. Similar to the metric conjoint simulation case, it is evident from Table B.3 that the computation time of Choice-SL is much lower than that of the NCM model for each experimental condition. The comparisons for other experimental conditions are qualitatively similar and are available from the authors upon request. Again, we note that the computation time may not be compared meaningfully across experimental conditions.

Table B.3: The Computation Time Comparison for Choice-based Conjoint Simulation

Num-S	Num-CS	EV	WSV	Choice-SL	NCM
2	16	1.645	0.05	1.0352	6.3292
			0.10	1.7021	6.2121
			0.20	1.4311	6.0972
2	24	1.645	0.05	2.4757	7.4159
			0.10	3.3988	7.7346
			0.20	2.1034	7.8911

Notes. The computation time reported above is averaged over 5 randomly simulated data sets and is reported in hours.

B.6.3 Metric Conjoint Application (Section 3.4.1)

We report the computation time of both models for the computer conjoint data set used in Section 3.4.1. The results are summarized in Table B.4.

Table B.4: The Computation Time Comparison for Metric Conjoint Application

Metric-SL	NCM
2.8661	9.7850

Notes. The computation time is reported in hours.

B.6.4 Choice-based Conjoint Application (Section 3.4.2)

We report the computation time of both models for the cell phone plans data set used in Section 3.4.2. We average the computation time of each model over 5 replications and report the average computation time in hours. The results are summarized in Table B.5.

Table B.5: The Computation Time Comparison for Choice-based Conjoint Application

Choice-SL	NCM
3.0419	4.7709