Assortment Optimization under the Decision Forest Model

Yi-Chun Akchen
School of Management, University College London, London E14 5AB, United Kingdom.
yi-chun.akchen@ucl.ac.uk

Velibor V. Mišić
UCLA Anderson School of Management, University of California, Los Angeles, California 90095, United States,
velibor.misic@anderson.ucla.edu

Problem definition: We study the problem of finding the optimal assortment that maximizes expected revenue under the decision forest model, a recently proposed nonparametric choice model that is capable of representing any discrete choice model and in particular, can be used to represent non-rational customer behavior. This problem is of practical importance because it allows a firm to tailor its product offerings to profitably exploit deviations from rational customer behavior, but at the same time is challenging due to the extremely general nature of the decision forest model. *Methodology/Results*: We approach this problem from a mixed-integer optimization perspective and propose two different formulations. We theoretically compare the two formulations in strength, and analyze when they are integral in the special case of a single tree. We further propose a methodology for solving the two formulations at a large-scale based on Benders decomposition, and show that the Benders subproblem can be solved efficiently by primal-dual greedy algorithms when the master solution is fractional for one of the formulations, and in closed form when the master solution is binary for both formulations. Using synthetically generated instances, we demonstrate the practical tractability of our formulations and our Benders decomposition approach, and their edge over heuristic approaches. *Managerial implications:* In a case study based on a real-world transaction data, we demonstrate that our proposed approach can factor the behavioral anomalies observed in consumer choice into assortment decision and create higher revenue.

Key words: decision trees; choice modeling; integer optimization; Benders decomposition; choice overload

1. Introduction

Assortment optimization is a basic operational problem faced by many firms. In its simplest form, the problem can be posed as follows. A firm has a set of products that it can offer, and a set of customers who have preferences over those products; what is the set of products the firm should offer so as to maximize the revenue that results when the customers choose from these products?

While assortment optimization and the related problem of product line design have been studied extensively under a wide range of choice models, the majority of research in this area focuses on rational choice models, specifically those that follow the random utility maximization (RUM)

assumption. A significant body of research in the behavioral sciences shows that customers behave in ways that deviate significantly from predictions made by RUM models. In addition, there is a substantial number of empirical examples of firms that make assortment decisions in ways that directly exploit customer irrationality. For example, the paper of Kivetz et al. (2004) provides an example of an assortment of document preparation systems from Xerox that are structured around the decoy effect, and an example of an assortment of servers from IBM that are structured around the compromise effect.

A recent paper by Chen and Mišić (2022) proposed a new choice model called the *decision forest model* for capturing customer irrationalities. This model involves representing the customer population as a probability distribution over binary trees, with each tree representing the decision process of one customer type. In a key result of the paper, the authors showed that this model is universal: *every discrete choice model is representable as a decision forest model.* While the paper of Chen and Mišić (2022) alludes to the downstream assortment optimization problem, it is entirely focused on model representation and prediction: it does not provide any answer to how one can select an optimal assortment with respect to a decision forest model.

In the present paper, we present a methodology for assortment optimization under the decision forest model, based on mixed-integer optimization. Our approach allows the firm to obtain assortments that are either provably optimal or within a desired optimality gap for a given decision forest model. At the same time, the approach easily allows the firm to incorporate business rules as linear constraints in the MIO formulation. Most importantly, given the universality property of the decision forest model, our optimization approach allows a firm to optimally tailor its assortment to any kind of predictable irrationality in the firm's customer population.

We make the following specific contributions:

1. We propose two different integer optimization models – SplitMIO and ProductMIO– for the problem of assortment optimization under the decision forest model. In terms of formulation strength, ProductMIO is stronger than SplitMIO. In the special case of a single purchase decision tree, we show that SplitMIO is integral in the special case that a product appears at most once in the splits of a purchase decision tree, and ProductMIO is always integral regardless of the structure of the tree.

2. We propose a Benders decomposition approach for solving the two formulations at a large scale. We show that Benders cuts for the linear optimization relaxations of SplitMIO can be obtained via a greedy algorithm that solves both the primal and dual of the subproblem. We also provide a simple example to show that the same type of greedy algorithm fails to solve the primal subproblem of ProductMIO. We further show how to obtain Benders cuts for the integer solutions of both SplitMIO and ProductMIO in closed form.

3. We present numerical experiments using both synthetic and real-world data. We first use the synthetic data to examine the formulation strength of SplitMIO and ProductMIO models, and to demonstrate the scalability of the Benders decomposition approach for the SplitMIO formulation in problem instances involving up to 3000 products, 500 trees, and 512 leaves per tree. We then use a real-world dataset to demonstrate how the decision forest model can factor a behavioral anomaly (*choice overload*) into its assortment decision and create higher revenue.

We organize the paper as follows. Section 2 reviews the related literature in choice modeling and assortment optimization. Section 3 defines the assortment problem and presents the two MIO formulations. Section 4 proposes a Benders decomposition approach to our formulations, and analyzes the subproblem for each of the formulations for fractional and binary solutions of the master problem. Sections 5 and 6 present the numerical results involving both the synthetic and real-world data. All proofs are relegated to the appendix.

2. Literature review

The problem of assortment optimization has been extensively studied in the operations management community; we refer readers to Gallego and Topaloglu (2019) for a recent review of the literature. The literature on assortment optimization has focused on developing approaches for finding the optimal assortment under many different rational choice models, such as the MNL model (Talluri and Van Ryzin 2004, Sumida et al. 2020), the latent class MNL model (Bront et al. 2009, Méndez-Díaz et al. 2014), the Markov chain choice model (Feldman and Topaloglu 2017, Désir et al. 2020) and the ranking-based model (Aouad et al. 2020, 2018, Feldman et al. 2019).

In addition to the assortment optimization literature, our paper is also related to the literature on product line design found in the marketing community. While assortment optimization is more often focused on the tactical decision of selecting which existing products to offer, where the products are ones that have been sold in the past and the choice model comes from transaction data involving those products, the product line design problem involves selecting which new products to offer, where the products are candidate products (i.e., they have not been offered before) and the choice model comes from conjoint survey data, where customers are asked to rate or choose between hypothetical products. Research in this area has considered different approaches to solve the problem under the ranking-based/first-choice model (McBride and Zufryden 1988, Belloni et al. 2008, Bertsimas and Mišić 2019) and the multinomial logit model (Chen and Hausman 2000, Schön 2010); for more details, we refer the reader to the literature review of Bertsimas and Mišić (2019).

Our paper is related to Bertsimas and Mišić (2019), which presents integer optimization formulations of the product line design problem when the choice model is a ranking-based model. As we will see later, our formulation SplitMIO can be viewed a generalization of the formulation

Bertsimas and Mišić (2019), to the decision forest model. In addition, the paper of Bertsimas and Mišić (2019) develops a specialized Benders decomposition approach for its formulation, which uses the fact that one can solve the subproblem associated with each customer type by applying a greedy algorithm. We will show in Section 4 that this same property generalizes to the SplitMIO formulation, leading to a tailored Benders decomposition algorithm for solving SplitMIO at scale.

Beyond these specific connections, the majority of the literature on assortment optimization and product line design considers rational choice models, whereas our paper contributes a methodology for non-rational assortment optimization. Fewer papers have focused on choice modeling for irrational customer behavior; besides the decision forest model, other models include the generalized attraction model (GAM; Gallego et al. 2015), the generalized stochastic preference model (Berbeglia 2018) and the generalized Luce model (Echenique and Saito 2019). An even smaller set of papers has considered assortment optimization under non-rational choice models, which we now review. The paper of Flores et al. (2017) considers assortment optimization under the two-stage Luce model, and develops a polynomial time algorithm for solving the unconstrained assortment optimization problem. The paper of Rooderkerk et al. (2011) considers a context-dependent utility model where the utility of a product can depend on other products that are offered and that can capture compromise, attraction and similarity effects; the paper empirically demonstrates how incorporating context effects leads to a predicted increase of 5.4% in expected profit.

Relative to these papers, our paper differs in that it considers the decision forest model. As noted earlier, the decision forest model can represent any type of choice behavior, and as such, an assortment optimization methodology based on such a model is attractive in terms of allowing a firm to take the next step from a high-fidelity model to a decision. In addition, our methodology is built on mixed-integer optimization. This is advantageous because it allows a firm to leverage continuing improvements in solution software for integer optimization (examples include commercial solvers like Gurobi and CPLEX), as well as continuing improvements in computer hardware. At the same time, integer optimization allows firms to accommodate business requirements using linear constraints, enhancing the practical applicability of the approach. Lastly, integer optimization also allows one to take advantage of well-studied large-scale solution methods for integer optimization problems. One such method that we focus on in this paper is Benders decomposition, which has seen an impressive resurgence in recent years for delivering state-of-the-art performance on large-scale problems such as hub location (Contreras et al. 2011), facility location (Fischetti et al. 2017) and set covering (Cordeau et al. 2019); see also Rahmaniani et al. (2017) for a review of the recent literature. Stated more concisely, the main contribution of our paper is *a general-purpose methodology for assortment optimization under a general-purpose choice model.*

In addition to the assortment optimization and product line design literatures, our formulations also have connections with others that have been proposed in the broader optimization literature. The formulation SplitMIO that we will present later can be viewed as a special case of the mixed-integer optimization formulation of Mišić (2020) for optimizing the predicted value of a tree ensemble model, such as a random forest or a boosted tree model. The formulation ProductMIO, which is our strongest formulation, also has a connection to the literature in the integer optimization community on formulating disjunctive constraints through independent branching schemes (Vielma et al. 2010, Vielma and Nemhauser 2011, Huchette and Vielma 2019); we also discuss this connection in more detail in Section 3.3.

3. Optimization model

In this section, we define the decision forest assortment optimization problem (Section 3.1) and subsequently develop our formulations, SplitMIO (Section 3.2) and ProductMIO (Section 3.3).

3.1. Problem definition

In this section, we briefly review the decision forest model of Chen and Mišić (2022), and then formally state the assortment optimization problem. We assume that there are n products, indexed from 1 to n, and let $\mathcal{N} = \{1, \ldots, n\}$ denote the set of all products. An assortment S corresponds to a subset of $\mathcal{N}$. When offered S, a customer may choose to purchase one of the products in S, or to not purchase anything from S at all; we use the index 0 to denote the latter possibility, which we will also refer to as the no-purchase option or the outside option.

The basic building block of the decision forest model is a purchase decision tree. A purchase decision tree is a directed binary tree, with each leaf node corresponding to an option in $\mathcal{N} \cup \{0\}$, and each non-leaf (or *split*) node corresponding to a product in $\mathcal{N}$. We use $\mathbf{splits}(t)$ to denote the set of split nodes of tree t, and $\mathbf{leaves}(t)$ to denote the set of leaf nodes. We use $c(t, \ell)$ to denote the purchase decision of leaf ℓ of tree t, i.e., the option chosen by tree t if the assortment is mapped to leaf ℓ. We use $v(t, s)$ to denote the product that is checked at split node s in tree t.

Each tree represents the purchasing behavior of one type of customer. Specifically, for an assortment S, the customer behaves as follows: the customer starts at the root of the tree. The customer checks whether the product corresponding to the root node is contained in S; if it is, she proceeds to the left child, and if not, she proceeds to the right child. She then checks again with the product at the new node, and the process repeats, until the customer reaches a leaf; the option that is at the leaf represents the choice of that customer. Figure 1 shows an example of a purchase decision tree being used to map an assortment to a purchase decision.

We make the following assumption about our purchase decision trees.

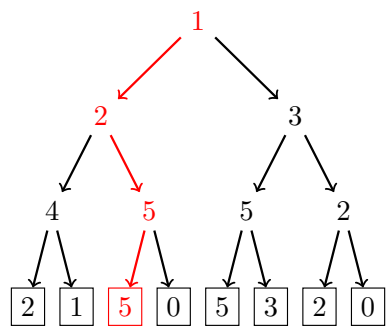

Figure 1 **Example of a purchase decision tree for $n=5$ products. Leaf nodes are enclosed in squares, while split nodes are not enclosed. The number on each node corresponds either to $v(t,s)$ for splits, or $c(t,\ell)$ for leaves. The path highlighted in red indicates how a customer following this tree maps the assortment $S=\{1,3,4,5\}$ to a leaf. For this assortment, the customer's decision is to purchase product 5.**

Assumption 1. *Let t be a purchase decision tree. For any two split nodes s and s' of t such that s' is a descendant of s, $v(t,s) \neq v(t,s')$.*

This assumption states that once a product appears on a split s, it cannot appear on any subsequent split s' that is reached by proceeding to the left or right child of s; in other words, each product in $\mathcal{N}$ appears at most once along the path from the root node to a leaf node, for every leaf node. As discussed in Chen and Mišić (2022), this assumption is not restrictive, as any tree for which this assumption is violated has a set of splits and leaves that are redundant and unreachable, and the tree can be modified to obtain an equivalent tree that satisfies the assumption.

The decision forest model assumes that the customer population is represented by a collection of trees or a *forest* F. Each tree $t \in F$ corresponds to a different customer type. We use λ_t to denote the probability associated with customer type/tree t, and $\boldsymbol{\lambda} = (\lambda_t)_{t \in F}$ to denote the probability distribution over the forest F. For each tree t, we use $\hat{A}(t,S)$ to denote the choice that a customer type following tree t will make when given the assortment S. For a given assortment $S \subseteq \mathcal{N}$ and a given choice $j \in S \cup \{0\}$, we use $\mathbf{P}^{(F,\boldsymbol{\lambda})}(j \mid S)$ to denote the choice probability, i.e., the probability of a random customer customer choosing j when offered the assortment S. It is defined as

$$\mathbf{P}^{(F,\boldsymbol{\lambda})}(j \mid S) = \sum_{t \in F} \lambda_t \cdot \mathbb{I}\{\hat{A}(t,S) = j\}. \tag{1}$$

We now define the assortment optimization problem. We use $\bar{r}_i$ to denote the marginal revenue of product i; for convenience, we use $\bar{r}_0 = 0$ to denote the revenue of the no-purchase option. The assortment optimization problem that we wish to solve is

$$\underset{S \subseteq \mathcal{N}}{\text{maximize}} \ \sum_{i \in S} \bar{r}_i \cdot \mathbf{P}^{(F,\boldsymbol{\lambda})}(i \mid S). \tag{2}$$

This is a challenging problem because of the general nature of the choice model $\mathbf{P}^{(F,\boldsymbol{\lambda})}(\cdot \mid \cdot)$. It turns out that problem (2) is theoretically intractable.

PROPOSITION 1. *The decision forest assortment optimization problem* (2) *is NP-Hard.*

The proof of this result (see Appendix B.1) follows by a reduction from the MAX 3SAT problem. In the next two sections, we present different mixed-integer optimization (MIO) formulations of this problem.

3.2. Formulation 1: SplitMIO

We now present our first formulation of the assortment optimization problem (2) as a mixed-integer optimization (MIO) problem. To formulate the problem, we introduce some additional notation. For notational convenience we let $r_{t,\ell} = \bar{r}_{c(t,\ell)}$ be the revenue of the purchase option of leaf ℓ of tree t. We let $\mathbf{left}(s)$ denote the set of leaf nodes that are to the left of split s (i.e., can only be reached by taking the left branch at split s), and similarly, we let $\mathbf{right}(s)$ denote the leaf nodes that are to the right of s. We introduce two sets of decision variables. For each $i \in \mathcal{N}$, we let x_i be a binary decision variable that is 1 if product i is included in the assortment, and 0 otherwise. For each tree $t \in F$ and leaf $\ell \in \mathbf{leaves}(t)$, we let $y_{t,\ell}$ be a binary decision variable that is 1 if the assortment encoded by $\mathbf{x}$ is mapped to leaf ℓ of tree t, and 0 otherwise.

With these definitions, our first formulation, SPLITMIO, is given below.

$$
\begin{aligned}
\text{SPLITMIO}: \quad & \underset{\mathbf{x},\mathbf{y}}{\text{maximize}} && \sum_{t \in F} \lambda_t \cdot \left[\sum_{\ell \in \mathbf{leaves}(t)} r_{t,\ell} y_{t,\ell} \right] && \text{(3a)} \\
& \text{subject to} && \sum_{\ell \in \mathbf{leaves}(t)} y_{t,\ell} = 1, \quad \forall\, t \in F, && \text{(3b)} \\
& && \sum_{\ell \in \mathbf{left}(s)} y_{t,\ell} \leq x_{v(t,s)}, \quad \forall\, t \in F,\ s \in \mathbf{splits}(t), && \text{(3c)} \\
& && \sum_{\ell \in \mathbf{right}(s)} y_{t,\ell} \leq 1 - x_{v(t,s)}, \quad \forall\, t \in F,\ s \in \mathbf{splits}(t), && \text{(3d)} \\
& && x_i \in \{0,1\}, \quad \forall\, i \in \mathcal{N}, && \text{(3e)} \\
& && y_{t,\ell} \geq 0, \quad \forall\, t \in F,\ \ell \in \mathbf{leaves}(t). && \text{(3f)}
\end{aligned}
$$

In order of appearance, the constraints in this formulation have the following meaning. Constraint (3b) requires that for each customer type t, the assortment encoded by $\mathbf{x}$ is mapped to exactly one leaf. Constraint (3c) imposes that for a split s in tree t, if product $v(t,s)$ is not in the assortment, then the assortment cannot be mapped to any of the leaves that are to the left of split s in tree t. Constraint (3d) is the symmetric case of constraint (3c) for the right-hand subtree of split s of tree t. The last two constraints require that $\mathbf{x}$ is binary and $\mathbf{y}$ is nonnegative. Note that it is not necessary to require $\mathbf{y}$ to be binary, as the constraints ensure that each $y_{t,\ell}$ automatically takes the correct value whenever $\mathbf{x}$ is binary. Finally, the objective function corresponds to the expected per-customer revenue of the assortment.

The SplitMIO formulation is related to two existing MIO formulations in the literature. First, it can be viewed as a specialized case of the MIO formulation in Mišić (2020). In that paper, the author develops a formulation for tree ensemble optimization, i.e., the problem of setting the independent variables in a tree ensemble model (e.g., a random forest or a gradient boosted tree model) to maximize the value predicted by that ensemble. Second, the SplitMIO formulation also relates to the MIO formulation for the product line design problem under the ranking-based model (Bertsimas and Mišić 2019). Note that Chen and Mišić (2022) showed that the ranking-based model can be regarded as a special case of the decision forest model. In the special case that each tree in the forest corresponds to a ranking and the decision forest corresponds to a ranking-based choice model, it can be verified that the formulation (3) can be reduced to the MIO formulation for product line design under ranking-based models presented in Bertsimas and Mišić (2019).

Before continuing to our second formulation, we establish an important property of problem (3) when $|F| = 1$. When $|F| = 1$, we can show that $\mathcal{F}_{\text{SplitMIO}}$ is integral in a particular special case. (Note that in the statement of the proposition below, we drop the index t to simplify notation.)

Proposition 2. *Let $(F, \boldsymbol{\lambda})$ be a decision forest model consisting of a single tree, i.e., $|F| = 1$. In addition, assume that for every $i \in \mathcal{N}$, $v(s) = i$ for at most one $s \in \textbf{splits}$. Then $\mathcal{F}_{\text{SplitMIO}}$ is integral, i.e., every extreme point $(\mathbf{x}, \mathbf{y})$ of the polyhedron $\mathcal{F}_{\text{SplitMIO}}$ satisfies $\mathbf{x} \in \{0,1\}^N$.*

The proof of this result (see Appendix B.2) follows by showing that the constraint matrix defining $\mathcal{F}_{\text{SplitMIO}}$ is totally unimodular. Proposition 2 is significant because it implies that for $|F| = 1$, the distinction between trees where each product appears at most once in any split and trees where a product may appear two or more times as a split is sharp. This insight provides the motivation for our second formulation, ProductMIO, which we present next.

3.3. Formulation 2: ProductMIO

The second formulation of problem (2) that we will present is motivated by the behavior of SplitMIO when a product participates in two or more splits. In particular, observe that in a given purchase decision tree, a product i may participate in two different splits s_1 and s_2 in the same tree. In this case, constraint (3c) in the SplitMIO will result in two constraints:

$$\sum_{\ell \in \textbf{left}(s_1)} y_{t,\ell} \leq x_i, \tag{4}$$

$$\sum_{\ell \in \textbf{left}(s_2)} y_{t,\ell} \leq x_i. \tag{5}$$

In the above two constraints, observe that $\textbf{left}(s_1)$ and $\textbf{left}(s_2)$ are disjoint (this is a direct consequence of Assumption 1). Given this and constraint (3b) that requires the $y_{t,\ell}$ variables to sum to 1, we can come up with a constraint that strengthens constraints (4) and (5) by combining them:

$$\sum_{\ell \in \textbf{left}(s_1)} y_{t,\ell} + \sum_{\ell \in \textbf{left}(s_2)} y_{t,\ell} \leq x_i. \tag{6}$$

In general, one can aggregate all the $y_{t,\ell}$ variables that are to the left of all splits involving a product i to produce a single left split constraint for product i. The same can also be done for the right split constraints. Generalizing this principle leads to the following alternate formulation, which we refer to as ProductMIO:

$$\text{ProductMIO}: \quad \underset{\mathbf{x},\mathbf{y}}{\text{maximize}} \quad \sum_{t\in F} \lambda_t \cdot \left[\sum_{\ell \in \mathbf{leaves}(t)} r_{t,\ell} y_{t,\ell} \right] \tag{7a}$$

$$\text{subject to} \quad \sum_{\ell \in \mathbf{leaves}(t)} y_{t,\ell} = 1, \quad \forall\, t \in F, \tag{7b}$$

$$\sum_{\substack{s \in \mathbf{splits}(t):\\ v(t,s)=i}} \sum_{\ell \in \mathbf{left}(s)} y_{t,\ell} \le x_i, \quad \forall\, t \in F,\ i \in \mathcal{N}, \tag{7c}$$

$$\sum_{\substack{s \in \mathbf{splits}(t):\\ v(t,s)=i}} \sum_{\ell \in \mathbf{right}(s)} y_{t,\ell} \le 1 - x_i, \quad \forall\, t \in F,\ i \in \mathcal{N}, \tag{7d}$$

$$x_i \in \{0,1\}, \quad \forall\, i \in \mathcal{N}, \tag{7e}$$

$$y_{t,\ell} \ge 0, \quad \forall\, t \in F,\ \ell \in \mathbf{leaves}(t). \tag{7f}$$

Relative to SplitMIO, ProductMIO differs in several ways. First, note that while both formulations have the same number of variables, formulation ProductMIO has a smaller number of constraints. In particular, SplitMIO has one left and one right split constraints for each split in each tree, whereas ProductMIO has one left and one right split constraint for each product. When the trees involve a large number of splits, this can lead to a sizable reduction in the number of constraints. Note also that when a product does not appear in any splits of a tree, we can also safely omit constraints (7c) and (7d) for that product.

The second difference with formulation SplitMIO, as we have already mentioned, is in formulation strength. Let $\mathcal{F}_{\text{ProductMIO}}$ be the feasible region of the linear optimization (LO) relaxation of ProductMIO. The following proposition formalizes the fact that formulation ProductMIO is at least as strong as formulation SplitMIO.

Proposition 3. *For any decision forest model* $(F, \boldsymbol{\lambda})$*,* $\mathcal{F}_{\text{ProductMIO}} \subseteq \mathcal{F}_{\text{SplitMIO}}$.

The proof follows straightforwardly using the logic given above; we thus omit the proof.

The last major difference is in how ProductMIO behaves when $|F| = 1$. We saw that a sufficient condition for $\mathcal{F}_{\text{SplitMIO}}$ to be integral when $|F| = 1$ is that each product appears in at most one split in the tree. In contrast, formulation ProductMIO is *always* integral when $|F| = 1$.

Proposition 4. *For any decision forest model* $(F, \boldsymbol{\lambda})$ *with* $|F| = 1$*,* $\mathcal{F}_{\text{ProductMIO}}$ *is integral.*

The proof of this proposition, given in Appendix B.3, follows by recognizing the connection between ProductMIO and another type of formulation in the literature. In particular, a stream of papers

in the mixed-integer optimization community (Vielma et al. 2010, Vielma and Nemhauser 2011, Huchette and Vielma 2019) has considered a general approach for deriving small and strong formulations of disjunctive constraints using independent branching schemes; we briefly review the most general such approach from Huchette and Vielma (2019) and showcase its connection to ProductMIO. In this approach, one has a finite ground set J, and is interested in optimizing over a particular subset of the $(|J|-1)-$dimensional unit simplex over J, $\Delta^J = \{\boldsymbol{\lambda} \in \mathbb{R}^J \mid \sum_{j \in J} \lambda_j = 1; \boldsymbol{\lambda} \geq \mathbf{0}\}$. The specific subset that we are interested in is called a *combinatorial disjunctive constraint* (CDC), and is given by

$$\mathrm{CDC}(\mathcal{S}) = \bigcup_{S \in \mathcal{S}} Q(S), \tag{8}$$

where $\mathcal{S}$ is a finite collection of subsets of J and $Q(S) = \{\boldsymbol{\lambda} \in \Delta \mid \lambda_j \leq 0 \text{ for } j \in J \setminus S\}$ for any $S \subseteq J$. This approach is very general: for example, by associating each j with a point $\mathbf{x}^j$ in $\mathbb{R}^n$, one can use $\mathrm{CDC}(\mathcal{S})$ to model an optimization problem over a union of polyhedra, where each polyhedron is the convex hull of a collection of vertices in $S \in \mathcal{S}$.

A *k-way independent branching scheme of depth t* is a representation of $\mathrm{CDC}(\mathcal{S})$ as a sequence of t choices between k alternatives:

$$\mathrm{CDC}(\mathcal{S}) = \bigcap_{m=1}^{t} \bigcup_{i=1}^{k} Q(L_i^m), \tag{9}$$

where $L_i^m \subseteq J$. In the special case that $k = 2$, we can write $\mathrm{CDC}(\mathcal{S}) = \cap_{m=1}^{t}(L_m \cup R_m)$ where $L_m, R_m \subseteq J$. This representation is known as a *pairwise independent branching scheme* and the constraints of the corresponding MIO can be written simply as

$$\sum_{j \in L_m} \lambda_j \leq z_m, \quad \forall\, m \in \{1, \ldots, k\}, \tag{10a}$$

$$\sum_{j \in R_m} \lambda_j \leq 1 - z_m, \quad \forall\, m \in \{1, \ldots, k\}, \tag{10b}$$

$$z_m \in \{0, 1\}, \quad \forall\, m \in \{1, \ldots, k\}, \tag{10c}$$

$$\sum_{j \in J} \lambda_j = 1, \tag{10d}$$

$$\lambda_j \geq 0, \quad \forall\, j \in J. \tag{10e}$$

This particular special case is important because it is always integral (see Theorem 1 of Vielma et al. 2010). Moreover, we can see that ProductMIO bears a strong resemblance to formulation (10). Constraints (10a) and (10a) correspond to constraints (7c) and (7d), respectively. In terms of variables, the λ_j and z_m variables in formulation (10) correspond to the $y_{t,\ell}$ and x_i variables in ProductMIO, respectively.

One notable difference is that in practice, one would use formulation (10) in a modular way; specifically, one would be faced with a problem where the feasible region can be written as

$\text{CDC}(\mathcal{S}_1) \cap \text{CDC}(\mathcal{S}_2) \cap \cdots \cap \text{CDC}(\mathcal{S}_G)$, where each $\mathcal{S}_g$ is a collection of subsets of J. To model this feasible region, one would introduce a set of $z_{g,m}$ variables for the gth CDC, enforce constraints (10a) - (10c) for the gth CDC, and use only one set of λ_j variables for the whole formulation. Thus, the λ_j variables are the "global" variables, while the $z_{g,m}$ variables would be "local" and specific to each CDC. In contrast, in ProductMIO, the x_i variables (the analogues of z_m) are the "global" variables, while the $y_{t,\ell}$ variables (the analogues of λ_j) are the "local" variables.

4. Solution methodology based on Benders decomposition

While the formulations in Section 3 bring the assortment optimization problem under the decision forest choice model closer to being solvable in practice, the effectiveness of these formulations can be limited in large-scale problems. In particular, consider the case where there is a large number of trees in the decision forest model and each tree consists of a large number of splits and leaves. In this setting, both formulations – SplitMIO and ProductMIO– will have a large number of $y_{t,\ell}$ variables and a large number of constraints to link those variables with the x_i variables, and may require significant computation time.

At the same time, SplitMIO and ProductMIO share a common problem structure. In particular, they have two sets of variables: the $\mathbf{x}$ variables, which determine the products that are to be included, and the $(\mathbf{y}_t)_{t \in F}$ variables, which model the choice of each customer type. In addition, for any two trees t, t' such that $t \neq t'$, the $\mathbf{y}_t$ variables and $\mathbf{y}_{t'}$ variables do not appear together in any constraints. Thus, one can view each of our two formulations as *a two-stage stochastic program*, where each tree t corresponds to a scenario; the variable $\mathbf{x}$ corresponds to the first-stage decision; and the variable $\mathbf{y}_t$ corresponds to the second-stage decision under scenario t, which is appropriately constrained by the first-stage decision $\mathbf{x}$.

Thus, one can apply Benders decomposition to solve the problem. At a high level, Benders decomposition involves using linear optimization duality to represent the optimal value of the second-stage problem for each tree t as a piecewise-linear concave function of $\mathbf{x}$, and to eliminate the $(\mathbf{y}_t)_{t \in F}$ variables. One can then re-write the optimization problem in epigraph form, resulting in an optimization problem in terms of the $\mathbf{x}$ variable and an auxiliary epigraph variable θ_t for each tree t, and a large family of constraints linking $\mathbf{x}$ and θ_t for each tree t. Although the family of constraints for each tree t is too large to be enumerated, one can solve the problem through constraint generation.

The main message of this section of the paper is that, in most cases, the primal and the dual forms of the second-stage problem can be solved either in closed form (when $\mathbf{x}$ is binary) or via a greedy algorithm (when $\mathbf{x}$ is fractional), thus allowing one to identify violated constraints for either the relaxation or the integer problem in a computationally efficient manner. In the remaining sections, we carefully analyze the second-stage problem for both formulations.

For SplitMIO, we show that the second-stage problem can be solved by a greedy algorithm when $\mathbf{x}$ is fractional (Section 4.1). For ProductMIO, we show that the same greedy approach does not solve the second-stage problem in the fractional case (Section 4.2). For both formulations, when $\mathbf{x}$ is binary, we characterize the primal and dual solutions in closed form (Section 4.3). Lastly, in Section 4.4, we briefly describe our overall algorithmic approach to solving the assortment optimization problem, which involves solving the Benders reformulation of the relaxed problem, followed by the Benders reformulation of the integer problem.

4.1. Benders reformulation of the SplitMIO relaxation

The Benders reformulation of the LO relaxation of SplitMIO can be written as

$$\begin{aligned} \underset{\mathbf{x},\boldsymbol{\theta}}{\text{maximize}} \quad & \sum_{t\in F} \lambda_t \theta_t && \text{(11a)}\\ \text{subject to} \quad & \theta_t \le G_t(\mathbf{x}), \quad \forall\, t\in F, && \text{(11b)}\\ & \mathbf{x}\in[0,1]^n, && \text{(11c)} \end{aligned}$$

where function $G_t(\mathbf{x})$ is the optimal value of the following subproblem corresponding to tree t:

$$\begin{aligned} G_t(\mathbf{x}) = \quad \underset{\mathbf{y}_t}{\text{maximize}} \quad & \sum_{\ell\in\mathbf{leaves}(t)} r_{t,\ell}\cdot y_{t,\ell} && \text{(12a)}\\ \text{subject to} \quad & \sum_{\ell\in\mathbf{leaves}(t)} y_{t,\ell} = 1, && \text{(12b)}\\ & \sum_{\ell\in\mathbf{left}(s)} y_{t,\ell} \le x_{v(t,s)}, \quad \forall\, s\in\mathbf{splits}(t), && \text{(12c)}\\ & \sum_{\ell\in\mathbf{right}(s)} y_{t,\ell} \le 1 - x_{v(t,s)}, \quad \forall\, s\in\mathbf{splits}(t), && \text{(12d)}\\ & y_{t,\ell}\ge 0, \quad \forall\, \ell\in\mathbf{leaves}(t). && \text{(12e)} \end{aligned}$$

We now present a greedy algorithm for solving problem (12) in Algorithm 1. This algorithm requires as input an ordering τ of the leaves in nondecreasing revenue. In particular, we require a bijection $\tau:\{1,\dots,|\mathbf{leaves}(t)|\}\to\mathbf{leaves}(t)$ such that $r_{t,\tau(1)}\ge r_{t,\tau(2)}\ge\cdots\ge r_{t,\tau(|\mathbf{leaves}(t)|}$, i.e., an ordering of leaves in nondecreasing revenue. In addition, in the definition of Algorithm 1, we use $\mathbf{LS}(\ell)$ and $\mathbf{RS}(\ell)$ to denote the set of left and right splits, respectively, of ℓ, which are defined as

$$\mathbf{LS}(\ell) = \{s\in\mathbf{splits}(t)\mid \ell\in\mathbf{left}(s)\},$$
$$\mathbf{RS}(\ell) = \{s\in\mathbf{splits}(t)\mid \ell\in\mathbf{right}(s)\},$$

In words, $\mathbf{LS}(\ell)$ is the set of splits for which we must proceed to the left in order to be able to reach ℓ, and $\mathbf{RS}(\ell)$ is the set of splits for which we must proceed to the right to reach ℓ. A split $s\in\mathbf{LS}(\ell)$ if and only if $\ell\in\mathbf{left}(s)$, and similarly, $s\in\mathbf{RS}(\ell)$ if and only if $\ell\in\mathbf{right}(s)$.

Intuitively, this algorithm progresses through the leaves from highest to lowest revenue, and sets the $y_{t,\ell}$ variable of each leaf ℓ to the highest value it can be set to without violating the left and right split constraints (12c) and (12d) and without violating the constraint $\sum_{\ell \in \mathbf{leaves}(t)} y_{t,\ell} \leq 1$. At each iteration, the algorithm additionally keeps track of which constraint became tight through the event set $\mathcal{E}$. An A_s event indicates that the left split constraint (12c) for split s became tight; a B_s event indicates that the right split constraint (12d) for split s became tight; and a C event indicates that the constraint $\sum_{\ell \in \mathbf{leaves}(t)} y_{t,\ell} \leq 1$ became tight. When a C event is not triggered, Algorithm 1 looks for the split which has the least remaining capacity (line 15). In the case that the $\arg\min$ is not unique and there are two or more splits that are tied, we break ties by choosing the split s with the lowest depth $d(s)$ (i.e., the split closest to the root node of the tree).

The function f keeps track of which leaf ℓ was being checked when an A_s / B_s / C event occurred. We note that both $\mathcal{E}$ and f are not needed to find the primal solution, but they are essential to determining the dual solution in the dual procedure (Algorithm 2, which we will define shortly).

Algorithm 1 Primal greedy algorithm for SplitMIO.

Require: Bijection $\tau : \{1, \ldots, |\mathbf{leaves}(t)|\} \to \mathbf{leaves}(t)$ such that $r_{t,\tau(1)} \geq r_{t,\tau(2)} \geq \cdots \geq r_{t,\tau(|\mathbf{leaves}(t)|)}$

1: Initialize $y_{t,\ell} \leftarrow 0$ for each $\ell \in \mathbf{leaves}(t)$.
2: **for** $i = 1, \ldots, |\mathbf{leaves}(t)|$ **do**
3: Set $q_C \leftarrow 1 - \sum_{j=1}^{i-1} y_{t,\tau(j)}$.
4: **for** $s \in \mathbf{LS}(\tau(i))$ **do**
5: Set $q_s \leftarrow x_{v(t,s)} - \sum_{\substack{j=1: \\ \tau(j) \in \mathbf{left}(s)}}^{i-1} y_{t,\tau(j)}$
6: **for** $s \in \mathbf{RS}(\tau(i))$ **do**
7: Set $q_s \leftarrow 1 - x_{v(t,s)} - \sum_{\substack{j=1: \\ \tau(j) \in \mathbf{right}(s)}}^{i-1} y_{t,\tau(j)}$
8: Set $q_{A,B} \leftarrow \min_{s \in \mathbf{LS}(\tau(i)) \cup \mathbf{RS}(\tau(i))} q_s$
9: Set $q^* \leftarrow \min\{q_C, q_{A,B}\}$
10: Set $y_{t,\tau(i)} \leftarrow q^*$
11: **if** $q^* = q_C$ **then**
12: Set $\mathcal{E} \leftarrow \mathcal{E} \cup \{C\}$.
13: Set $f(C) = \tau(i)$.
14: **else**
15: Set $s^* \leftarrow \arg\min_{s \in \mathbf{LS}(\tau(i)) \cup \mathbf{RS}(\tau(i))} q_s$
16: **if** $s^* \in \mathbf{LS}(\tau(i))$ **then**
17: Set $e = A_s$
18: **else**
19: Set $e = B_s$
20: **if** $e \notin \mathcal{E}$ **then**
21: Set $\mathcal{E} \leftarrow \mathcal{E} \cup \{e\}$.
22: Set $f(e) = \tau(i)$.

It turns out that Algorithm 1 returns a feasible solution that is an extreme point of the polyhedron defined in problem (12). We state the result formally as the following theorem.

THEOREM 1. *Fix $t \in F$. Let $\mathbf{y}_t$ be a solution to problem* (12) *produced by Algorithm 1. Then: (a) $\mathbf{y}_t$ is a feasible solution to problem* (12)*; and (b) $\mathbf{y}_t$ is an extreme point of the feasible region of problem* (12)*.*

Theorem 1 implies that problem (12) is feasible, and since the problem is bounded, it has a finite optimal value. By strong duality, the optimal value of problem (13) is equal to the optimal value of its dual:

$$\begin{aligned}
\underset{\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t}{\text{minimize}} \quad & \sum_{s \in \mathbf{splits}(t)} x_{v(t,s)} \cdot \alpha_{t,s} + \sum_{s \in \mathbf{splits}(t)} (1 - x_{v(t,s)}) \beta_{t,s} + \gamma_t && \text{(13a)} \\
\text{subject to} \quad & \sum_{s : \ell \in \mathbf{left}(s)} \alpha_{t,s} + \sum_{s : \ell \in \mathbf{right}(s)} \beta_{t,s} + \gamma_t \geq r_{t,\ell}, \quad \forall\ \ell \in \mathbf{leaves}(t), && \text{(13b)} \\
& \alpha_{t,s} \geq 0, \quad \forall\ s \in \mathbf{splits}(t), && \text{(13c)} \\
& \beta_{t,s} \geq 0, \quad \forall\ s \in \mathbf{splits}(t). && \text{(13d)}
\end{aligned}$$

Letting $\mathcal{D}_{t,\text{SplitMIO}}$ denote the set of feasible solutions to the dual subproblem (13), we can formulate the master problem (11) as

$$\begin{aligned}
\underset{\mathbf{x}, \boldsymbol{\theta}}{\text{maximize}} \quad & \sum_{t \in F} \lambda_t \theta_t && \text{(14a)} \\
\text{subject to} \quad & \theta_t \leq \sum_{s \in \mathbf{splits}(t)} x_{v(t,s)} \cdot \alpha_{t,s} + \sum_{s \in \mathbf{splits}(t)} (1 - x_{v(t,s)}) \beta_{t,s} + \gamma_t, \\
& \quad \forall\ (\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t) \in \mathcal{D}_{t,\text{SplitMIO}}, && \text{(14b)} \\
& \mathbf{x} \in [0,1]^n. && \text{(14c)}
\end{aligned}$$

The value of this formulation, relative to the original formulation, is that we have replaced the $(\mathbf{y}_t)_{t \in F}$ variables and the constraints that link them to the $\mathbf{x}$ variables, with a large family of constraints in terms of $\mathbf{x}$. Although this new formulation is still challenging, the advantage of this formulation is that it is suited to constraint generation.

The constraint generation approach to solving problem (14) involves starting the problem with no constraints and then, for each $t \in F$, checking whether constraint (14b) is violated. If constraint (14b) is not violated for any $t \in F$, then we conclude that the current solution $\mathbf{x}$ is optimal. Otherwise, for any $t \in F$ such that constraint (14b) is violated, we add the constraint corresponding to the $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ solution at which the violation occurred, and solve the problem again to obtain a new $\mathbf{x}$. The procedure then repeats at the new $\mathbf{x}$ solution until no more violated constraints have been found.

The crucial step to solving this problem is being able to solve the dual subproblem (13); that is, for a fixed $t \in F$, either asserting that the current solution $\mathbf{x}$ satisfies constraint (14b) or identifying

a $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ at which constraint (14b) is violated. This amounts to solving the dual subproblem (13) and comparing its objective value to θ_t.

Fortunately, it turns out that we can solve the dual subproblem (13) using a specialized algorithm, in the same way that we can solve the primal subproblem (12) using Algorithm 1. Algorithm 2 uses auxiliary information obtained during the execution of Algorithm 1. In the definition of Algorithm 2, we use $d(s)$ to denote the depth of an arbitrary split, where the root split corresponds to a depth of 1, and $d_{\max} = \max_{s \in \mathbf{splits}(t)} d(s)$ is the depth of the deepest split in the tree. In addition, we use $\mathbf{splits}(t, d) = \{s \in \mathbf{splits}(t) \mid d(s) = d\}$ to denote the set of all splits at a particular depth d.

Algorithm 2 Dual greedy algorithm for SPLITMIO.

Initialize $\alpha_{t,s} \leftarrow 0, \beta_{t,s} \leftarrow 0$ for all $s \in \mathbf{splits}(t)$, $\gamma_t \leftarrow 0$
Set $\gamma \leftarrow r_{f(C)}$
for $d = 1, \ldots, d_{\max}$ **do**
  for $s \in \mathbf{splits}(t, d)$ **do**
    if $A_s \in \mathcal{E}$ **then**
      Set $\alpha_{t,s} \leftarrow r_{t,f(A_s)} - \gamma_t - \sum_{\substack{s' \in \mathbf{LS}(f(A_s)): \\ A_{s'} \in \mathcal{E}, \\ d(s') < d}} \alpha_{t,s'} - \sum_{\substack{s' \in \mathbf{RS}(f(A_s)): \\ B_{s'} \in \mathcal{E}, \\ d(s') < d}} \beta_{t,s'}$
    if $B_s \in \mathcal{E}$ **then**
      Set $\beta_{t,s} \leftarrow r_{t,f(B_s)} - \gamma_t - \sum_{\substack{s' \in \mathbf{LS}(f(A_s)): \\ A_{s'} \in \mathcal{E}, \\ d(s') < d}} \alpha_{t,s'} - \sum_{\substack{s' \in \mathbf{RS}(f(A_s)): \\ B_{s'} \in \mathcal{E}, \\ d(s') < d}} \beta_{t,s'}$

We provide a worked example of the execution of both Algorithms 1 and 2 in Appendix A.

We can show that the dual solution produced by Algorithm 2 is a feasible extreme point solution of the dual subproblem (13).

THEOREM 2. *Fix $t \in F$. Let $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ be a solution to problem* (13) *produced by Algorithm 2. Then: (a) $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ is a feasible solution to problem* (13)*; and (b) $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ is an extreme point of the feasible region of problem* (13)*.*

Lastly, and most importantly, we show that the solutions produced by Algorithms 1 and 2 are optimal for their respective problems. Thus, Algorithm 2 is a valid procedure for identifying values of $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ at which constraint (14b) is violated.

THEOREM 3. *Fix $t \in F$. Let $\mathbf{y}_t$ be a solution to problem* (12) *produced by Algorithm 1 and $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ be a solution to problem* (13) *produced by Algorithm 2. Then: (a) $\mathbf{y}_t$ is an optimal solution to problem* (12)*; and (b) $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ is an optimal solution to problem* (13)*.*

The proof of this result follows by verifying that the two solutions satisfy complementary slackness.

Before continuing, we note that Algorithms 1 and 2 can be viewed as the generalization of the algorithms arising in the Benders decomposition approach to the ranking-based assortment

optimization problem (Bertsimas and Mišić 2019). The results of that paper show that the primal subproblem of the MIO formulation in Bertsimas and Mišić (2019) can be solved via a greedy algorithm (analogous to Algorithm 1) and the dual subproblem can be solved via an algorithm that uses information from the primal algorithm (analogous to Algorithm 2). This generalization is not straightforward. The main challenge in this generalization is designing the sequence of updates in the greedy algorithm according to the tree topology. For example, in Algorithm 1, one considers all left/right splits and the y values of their left/right leaves when constructing the lowest upper bound of y_ℓ for each leaf node ℓ. Also, as shown in Algorithm 2, the dual variables $\alpha_{t,s}$ and $\beta_{t,s}$ have to be updated according to the tree topology and the events $A_{s'}$ and $B_{s'}$ of the split s' with smaller depth. In contrast, in the ranking-based assortment optimization problem, one only needs to calculate the "capacities" (the q_s values in Algorithm 1) by subtracting the y values of the preceding products in the ranking, which does not inherit any topological structure. For these reasons, the primal and dual Benders subproblems for the decision forest assortment problem are much more challenging than those of the ranking-based assortment problem.

4.2. Benders reformulation of the ProductMIO relaxation

We also consider a Benders reformulation of the relaxation of ProductMIO. The Benders master problem is given by formulation (11) where the function $G_t(\mathbf{x})$ is defined as the optimal value of the ProductMIO subproblem for tree t. To aid in the definition of the subproblem, let $P(t)$ denote the set of products that appear in the splits of tree t:

$$P(t) = \{i \in \mathcal{N} \mid i = v(t,s) \text{ for some } s \in \mathbf{splits}(t)\}.$$

With a slight abuse of notation, let $\mathbf{left}(i)$ denote the set of leaves for which product i must be included in the assortment for those leaves to be reached, and similarly, let $\mathbf{right}(i)$ denote the set of leaves for which product i must be excluded from the assortment for those leaves to be reached; formally,

$$\mathbf{left}(i) = \bigcup_{\substack{s \in \mathbf{splits}(t): \\ v(t,s)=i}} \mathbf{left}(s), \quad \text{and} \quad \mathbf{right}(i) = \bigcup_{\substack{s \in \mathbf{splits}(t): \\ v(t,s)=i}} \mathbf{right}(s).$$

With these definitions, we can write down the ProductMIO subproblem as follows:

$$G_t(\mathbf{x}) = \quad \underset{\mathbf{y}_t}{\text{maximize}} \quad \sum_{\ell \in \mathbf{leaves}(t)} r_{t,\ell} \cdot y_{t,\ell} \tag{15a}$$

$$\text{subject to} \quad \sum_{\ell \in \mathbf{leaves}(t)} y_{t,\ell} = 1, \tag{15b}$$

$$\sum_{\ell \in \mathbf{left}(i)} y_{t,\ell} \le x_i, \quad \forall\, i \in P(t), \tag{15c}$$

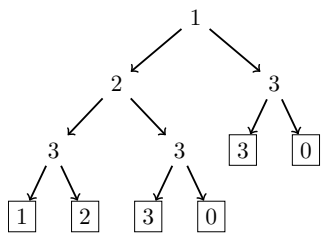

Figure 2 **Structure of tree for which the ProductMIO primal subproblem is not solvable via a greedy algorithm.**

$$\sum_{\ell \in \mathbf{right}(i)} y_{t,\ell} \leq 1 - x_i, \quad \forall\, i \in P(t), \tag{15d}$$

$$y_{t,\ell} \geq 0, \quad \forall\, \ell \in \mathbf{leaves}(t). \tag{15e}$$

In the same way as SplitMIO, one can consider solving problem (15) using a greedy approach, where one iterates through the leaves from highest to lowest revenue, and sets each leaf's $y_{t,\ell}$ variable to the highest possible value without violating any of the constraints. Unlike SplitMIO, it unfortunately turns out that this greedy approach is not always optimal, which is formalized in the following proposition.

Proposition 5. *There exists an $\mathbf{x} \in [0,1]^n$, a tree t and revenues $\bar{r}_1, \ldots, \bar{r}_n$ for which the greedy solution to problem* (15) *is not optimal.*

We provide a counterexample as follows. Consider an instance where $\mathcal{N} = \{1,2,3\}$ and $\mathbf{x} = (0.5, 0.5, 0.5)$. Assume the revenues of the products are $\bar{r}_1 = 20$, $\bar{r}_2 = 19$ and $\bar{r}_3 = 18$. Consider the tree shown in Figure 2. We label the leaves as 1, 2, 3, 4, 5 and 6 from left to right. When we apply the greedy algorithm to solve this LO problem, we can see that there are multiple orderings of the leaves in decreasing revenue: (i) 1,2,3,5,4,6; (ii) 1,2,5,3,4,6; (iii) 1,2,3,5,6,4; and (iv) 1,2,5,3,6,4. For any of these orderings, the greedy solution will turn out to be $(y_1, y_2, y_3, y_4, y_5, y_6) = (0.5, 0, 0, 0, 0, 0.5)$, resulting in an objective value of 10. However, the actual optimal solution $G_t(\mathbf{x})$ turns out to be $(y_1^*, y_2^*, y_3^*, y_4^*, y_5^*, y_6^*) = (0, 0.5, 0, 0, 0.5, 0)$, for which the objective value is 18.5. This counterexample shows that in general, the ProductMIO primal subproblem cannot be solved to optimality via the same type of greedy algorithm as for SplitMIO.

4.3. Bender Cuts for Integer Master Solutions

We further propose closed form expressions for the structure of the optimal primal and dual Benders subproblem solutions for integer solutions $\mathbf{x}$.

4.3.1. SplitMIO Our results in Section 4.1 for obtaining primal and dual solutions for the subproblem of SplitMIO apply for any $\mathbf{x} \in [0,1]^n$; in particular, they apply for fractional choices of $\mathbf{x}$, thus allowing us to solve the Benders reformulation of the relaxation of SplitMIO.

In the special case that $\mathbf{x}$ is a candidate integer solution of SplitMIO, we can find optimal solutions to the primal and dual subproblems of SplitMIO in closed form.

Theorem 4. *Fix $t \in F$, and let $\mathbf{x} \in \{0,1\}^n$. Define the primal subproblem solution $\mathbf{y}_t$ as*

$$y_{t,\ell} = \begin{cases} 1 & \text{if } \ell = \ell^*, \\ 0 & \text{if } \ell \neq \ell^*, \end{cases}$$

where ℓ^* denotes the leaf that the assortment encoded by $\mathbf{x}$ is mapped to. Define the dual subproblem solution $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ as

$$\begin{aligned} \alpha_{t,s} &= \begin{cases} \max\{0, \max_{\ell \in \mathbf{left}(s)} r_{t,\ell} - r_{t,\ell^*}\} & \text{if } s \in \mathbf{RS}(\ell^*), \\ 0 & \text{otherwise}, \end{cases} \\ \beta_{t,s} &= \begin{cases} \max\{0, \max_{\ell \in \mathbf{right}(s)} r_{t,\ell} - r_{t,\ell^*}\} & \text{if } s \in \mathbf{LS}(\ell^*), \\ 0 & \text{otherwise}, \end{cases} \\ \gamma_t &= r_{t,\ell^*}. \end{aligned}$$

Then: (a) $\mathbf{y}_t$ is a feasible solution to problem (12); (b) $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ is a feasible solution to problem (13); and (c) $\mathbf{y}_t$ and $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ are optimal for problems (12) and (13), respectively.

The significance of Theorem 4 is that it provides a simpler means to checking for violated constraints when $\mathbf{x}$ is binary than applying Algorithms 1 and 2. In particular, for the integer version of SplitMIO, a similar derivation as in Section 4.1 leads us to the following Benders reformulation of the integer problem for the SplitMIO formulation:

$$\begin{aligned} \underset{\mathbf{x},\boldsymbol{\theta}}{\text{maximize}} \quad & \sum_{t \in F} \lambda_t \theta_t && \text{(16a)} \\ \text{subject to} \quad & \theta_t \leq \sum_{s \in \mathbf{splits}(t)} x_{v(t,s)} \cdot \alpha_{t,s} + \sum_{s \in \mathbf{splits}(t)} (1 - x_{v(t,s)}) \beta_{t,s} + \gamma_t, \\ & \quad \forall\ (\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t) \in \mathcal{D}_{t,\text{SplitMIO}}, && \text{(16b)} \\ & \mathbf{x} \in \{0,1\}^n. && \text{(16c)} \end{aligned}$$

To check whether constraint (16b) is violated for a particular $\mathbf{x}$ and a tree t, we can simply use Theorem 4 to determine the optimal value of the subproblem, and compare it against θ_t; if the constraint corresponding to the dual solution of Theorem 4 is violated, we add that constraint to the problem. In our implementation of Benders decomposition, we embed the constraint generation process for the integer problem (16) within the branch-and-bound tree, using a technique referred to as *lazy constraint generation*; we discuss this more in Section 4.4.

4.3.2. ProductMIO We also consider ProductMIO. We begin by writing down the dual of the subproblem, for which we need to define several additional sets. We let $\mathbf{LP}(\ell)$ denote the set of "left products" of leaf ℓ (those products that must be included in the assortment for leaf

ℓ to be reached), and let $\mathbf{RP}(\ell)$ denote the set of "right products" of leaf ℓ (those products that must be excluded from the assortment for leaf ℓ to be reached). Note that $\ell \in \mathbf{left}(i)$ if and only if $i \in \mathbf{LP}(\ell)$, and similarly $\ell \in \mathbf{right}(i)$ if and only if $i \in \mathbf{RP}(\ell)$.

With these definitions, the dual of the primal subproblem (15) is

$$\underset{\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t}{\text{minimize}} \quad \sum_{i \in P(t)} \alpha_{t,i} x_i + \sum_{i \in P(t)} \beta_{t,i}(1 - x_i) + \gamma_t \tag{17a}$$

$$\text{subject to} \quad \sum_{i \in \mathbf{LP}(\ell)} \alpha_{t,i} + \sum_{i \in \mathbf{RP}(\ell)} \beta_{t,i} + \gamma_t \geq r_\ell, \quad \forall\, \ell \in \mathbf{leaves}(t), \tag{17b}$$

$$\alpha_{t,i} \geq 0, \quad \forall\, i \in P(t), \tag{17c}$$

$$\beta_{t,i} \geq 0, \quad \forall\, i \in P(t). \tag{17d}$$

In the case that $\mathbf{x}$ is integer, we can obtain optimal solutions to the primal subproblem (15) and its dual (17) in closed form.

THEOREM 5. *Fix $t \in F$ and let $\mathbf{x} \in \{0,1\}^n$. Let $\mathbf{y}_t$ be defined as*

$$y_{t,\ell} = \begin{cases} 1 & \text{if } \ell = \ell^*, \\ 0 & \text{otherwise}, \end{cases} \tag{18}$$

and let $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ be defined as

$$\alpha_{t,i} = \begin{cases} \max\{0, \max_{\ell \in \mathbf{left}(i)} r_{t,\ell} - r_{t,\ell^*}\}, & \text{if } i \in \mathbf{RP}(\ell^*), \\ 0 & \text{otherwise}, \end{cases} \tag{19}$$

$$\beta_{t,i} = \begin{cases} \max\{0, \max_{\ell \in \mathbf{right}(i)} r_{t,\ell} - r_{t,\ell^*}\}, & \text{if } i \in \mathbf{LP}(\ell^*), \\ 0 & \text{otherwise}, \end{cases} \tag{20}$$

$$\gamma_t = r_{t,\ell^*}. \tag{21}$$

Then: (a) $\mathbf{y}_t$ is a feasible solution for problem (15); (b) $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ is a feasible solution for problem (17); and (c) $\mathbf{y}_t$ and $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ are optimal for problems (15) and (17), respectively.

4.4. Overall Benders algorithm

We conclude Section 4 by summarizing how the results are used. In our overall algorithmic approach below, we first focus on SPLITMIO, as the subproblem can be solved for the formulation when $\mathbf{x}$ is either fractional or binary.

1. *Relaxation phase.* We first solve the relaxed problem (14) using ordinary constraint generation. Given a solution $\mathbf{x} \in [0,1]^n$, we generate Benders cuts by running the primal-dual procedure (Algorithm 1 followed by Algorithm 2).

2. *Integer phase.* In the integer phase, we add all of the Benders cuts generated in the relaxation phase to the integer version of problem (14). We then solve the problem as an integer optimization problem, where we generate Benders cuts for integer solutions using the closed form expressions

in Theorem 4. We add these cuts using *lazy constraint generation*. That is, we solve the master problem using a single branch-and-bound tree, and we check whether the main constraint (14b) of the Benders formulation is violated at every integer solution generated in the branch-and-bound tree.

For ProductMIO, as the subproblem can only be solved when $\mathbf{x}$ is binary, its overall Benders algorithm directly starts with the *Integer Phase*.

5. Numerical Experiments with Synthetic Data

In this section, we examine the formulation strength of SplitMIO and ProductMIO (Section 5.2) and demonstrate the scalability of the Benders decomposition approach(Section 5.3). We focus on synthetically generated instances so that we can scale the problem size. In Appendix C, we report our implementation details and include additional numerical results.

5.1. Background

To test our method, we generate three different families of synthetic decision forest instances, which differ in the topology of the trees and the products that appear in the splits:

1. **T1 forest**. A T1 forest consists of balanced trees of depth d (i.e., trees where all leaves are at depth $d+1$). For each tree, we sample d products $i_1, \ldots, i_d$ uniformly without replacement from $\mathcal{N}$. Then, for every depth $d' \in \{1, \ldots, d\}$, we set the split product $v(t,s)$ as $v(t,s) = i_{d'}$ for every split s that is at depth d'.

2. **T2 instances**. A T2 forest consists of balanced trees of depth d. For each tree, we set the split products at each split iteratively, starting at the root, in the following manner: (a) Initialize $d' = 1$; (b) For all splits s at depth d', set $v(s,t) = i_s$ where i_s is drawn uniformly at random from the set $\mathcal{N} \setminus \cup_{s' \in A(s)} \{v(t,s')\}$, where $A(s)$ is the set of ancestor splits to split s (i.e., all splits appearing on the path from the root node to split s); (c) Increment $d' \leftarrow d' + 1$; and (d) If $d' > d$, stop; otherwise, return to Step (b).

3. **T3 instances**. A T3 forest consists of unbalanced trees with L leaves. Each tree is generated according to the following iterative procedure: (a) Initialize t to a tree consisting of a single leaf; (b) Select a leaf ℓ uniformly at random from $\mathbf{leaves}(t)$, and replace it with a split s and two child leaves ℓ_1, ℓ_2. For split s, set $v(s,t) = i_s$ where i_s is drawn uniformly at random from $\mathcal{N} \setminus \cup_{s' \in A(s)} \{v(t,s')\}$; (c) If $|\mathbf{leaves}(t)| = L$, terminate; otherwise, return to Step (b).

For all three types of forests, we generate the purchase decision $c(t,\ell)$ for each leaf ℓ in each tree t in the following way: for each leaf ℓ, we uniformly at random choose a product $i \in \cup_{s \in \mathbf{LS}(\ell)} \{v(t,s)\} \cup \{0\}$. In words, the purchase decision is chosen to be consistent with the products that are known to be in the assortment if leaf ℓ is reached. Figure 3 shows an example of each type of tree (T1, T2,

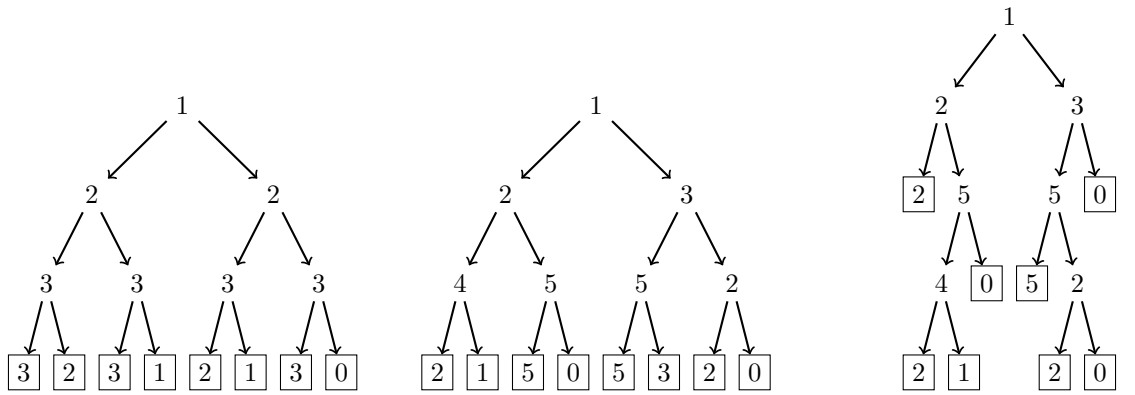

(a) T1 tree ($d=3$). (b) T2 tree ($d=3$). (c) T3 tree ($L=8$).

Figure 3 Examples of T1, T2 and T3 trees.

and T3). Given a forest of any of the three types above, we generate the customer type probability vector $\boldsymbol{\lambda}=(\lambda_t)_{t\in F}$ by drawing it uniformly from the $(|F|-1)-$dimensional unit simplex.

In our experiments, we fix the number of products $n=100$ and vary the number of trees $|F|\in\{50,100,200,500\}$, and the number of leaves $|\mathbf{leaves}(t)|\in\{8,16,32,64\}$. (Note that the chosen values for $|\mathbf{leaves}(t)|$ correspond to depths of $\{3,4,5,6\}$ for the T1 and T2 instances.) For each combination of n, $|F|$ and $|\mathbf{leaves}(t)|$ and each type of instance (T1, T2 and T3) we randomly generate 20 problem instances, where a problem instance consists of a decision forest model and the product marginal revenues $\bar{r}_1,\ldots,\bar{r}_n$. For each instance, the decision forest model is generated according to the process described above and the product revenues are sampled uniformly with replacement from the set $\{1,\ldots,100\}$.

5.2. Experiment #1: Formulation Strength

Our first experiment is to simply understand how the two formulations – SplitMIO and ProductMIO– compare in terms of formulation strength. Recall from Proposition 3 that ProductMIO is at least as strong as SplitMIO. For a given instance and a given formulation $\mathcal{M}$ (one of SplitMIO and ProductMIO), we define the integrality gap $G^{\text{int}}_{\mathcal{M}}\equiv 100\%\times(Z_{\mathcal{M}}-Z^*)/Z^*$, where Z^* is the optimal objective value of the integer problem and $Z_{\mathcal{M}}$ is the optimal objective of the LO relaxation. We consider the T1, T2 and T3 instances with $n=100$, $|F|\in\{50,100,200,500\}$ and $|\mathbf{leaves}(t)|=8$. We restrict our focus to instances with $n=100$ and $|\mathbf{leaves}(t)|=8$, as the optimal value Z^* of the integer problem could be computed within one hour for these instances.

Table 1 displays the average integrality gap of each of the two formulations for each combination of n and $|F|$ and each instance type. From this table, we can see that the integrality gap of both SplitMIO and ProductMIO is in general about 0 to 17%. Note that the difference between ProductMIO and SplitMIO is most pronounced for the T1 instances, as the decision forests in these instances exhibit the highest degree of repetition of products within the splits of a tree. In

Type	$\lvert F\rvert$	$G^{\text{int}}_{\text{SplitMIO}}$	$G^{\text{int}}_{\text{ProductMIO}}$
T1	50	0.9	0.0
T1	100	2.5	0.1
T1	200	5.6	0.2
T1	500	15.8	3.3
T2	50	0.2	0.2
T2	100	1.0	1.0
T2	200	5.4	5.3
T2	500	16.7	16.4
T3	50	0.2	0.2
T3	100	0.5	0.5
T3	200	4.1	3.9
T3	500	14.2	14.0

Table 1 Average integrality gap of SplitMIO and ProductMIO for T1, T2 and T3 instances with $n = 100$, $\lvert\mathbf{leaves}(t)\rvert = 8$.

contrast, the difference is smaller for the T2 and T3 instances, where the trees are balanced but there is less repetition of products within the splits of the tree (as the trees are not forced to have the same product appear on all of the splits at a particular depth).

We also test the tractability of SplitMIO and ProductMIO when they are solved as integer programs (i.e., not as relaxation). We present the results in Section C.1.

5.3. Experiment #2: Benders Decomposition for Large-Scale Problems

In this experiment, we report on the performance of our Benders decomposition approach for solving large scale instances of SplitMIO. We focus on the SplitMIO formulation, as we are able to efficiently generate Benders cuts for both fractional and integral values of $\mathbf{x}$.

We deviate from our previous experiments by generating a collection of T3 instances with $n \in \{200, 500, 1000, 2000, 3000\}$, $|F| = 500$ trees and $|\mathbf{leaves}(t)| = 512$ leaves. As before, the marginal revenue $\bar{r}_i$ of each product i is chosen uniformly at random from $\{1, \ldots, 100\}$. For each value of n, we generate 5 instances. For each instance, we solve the SplitMIO problem subject to the constraint $\sum_{i=1}^{n} x_i = b$, where b is set as $b = \rho n$ and we vary $\rho \in \{0.02, 0.04, 0.06, 0.08, 0.10, 0.12\}$.

We compare three different methods: the two-phase Benders method described in Section 4.4, using the SplitMIO cut results (Section 4.1 and Section 4.3); the divide-and-conquer (D&C) heuristic; and the direct solution approach, where we attempt to directly solve the full SplitMIO formulation using Gurobi. The D&C heuristic is a type of local search heuristic proposed in the product line design literature (see Green and Krieger 1993; see also Belloni et al. 2008). In this heuristic, one iterates through the b products currently in the assortment, and replaces a single product with the product outside of the assortment that leads to the highest improvement in the expected revenue; this process repeats until the assortment can no longer be improved. We choose

the initial assortment uniformly at random from the collection of assortments of size b. For each instance, we repeat the D&C heuristic 10 times, and retain the best solution. We do not impose a time limit on the D&C heuristic. For the Benders approach, we impose a time limit of one hour on the LO phase, and impose a time limit of one hour on the integer phase. For the direct solution approach, we impose a time limit of two hours, in order to be comparable to the Benders approach.

Table 2 reports the performance of the three methods – the Benders approach, the D&C heuristic and direct solution of SplitMIO– across all combinations of n and ρ. We consider several metrics. The metric $G_{\mathcal{M}}$ is defined as $G_{\mathcal{M}} = (Z' - Z_{\mathcal{M}})/Z' \times 100\%$, i.e., it is the optimality gap of the solution obtained by method $\mathcal{M}$ – either the Benders approach, the direct approach or the D&C heuristic – relative to the objective value of the best solution obtained out of the three methods, which is indicated by Z'. Lower values of $G_{\mathcal{M}}$ indicate that the approach tends to deliver solutions that are close to being the best out of the three methods, and a value of 0% implies that the solution obtained by the approach is the best (or tied for the best) out of the three methods. The metric $T_{\mathcal{M}}$ indicates the solution time required for approach $\mathcal{M}$ in seconds. Finally, for the Benders and direct approaches, we compute the optimality gap $O_{\mathcal{M}}$, which is defined as $O_{\mathcal{M}} = (Z_{\text{UB},\mathcal{M}} - Z_{\text{LB},\mathcal{M}})/Z_{\text{UB},\mathcal{M}} \times 100\%$, where $Z_{\text{UB},\mathcal{M}}$ and $Z_{\text{LB},\mathcal{M}}$ are the best upper and lower bounds, respectively, obtained from either the Benders or direct approach after the computation time limit is exhausted. The value reported of each metric is the average over the five replications corresponding to the particular (n, ρ) combination.

Comparing the performance of the Benders approach with the D&C heuristic, we can see that in general, the Benders approach is able to find better solutions than the D&C heuristic. In particular, for larger instances, G_{Benders} is lower than $G_{\text{D\&C}}$ (for example, with $n = 2000$, $\rho = 0.08$, the Benders solution is in general the best one, and the D&C solution has an expected revenue that is 4.6% worse). In addition, from a computation time standpoint, the Benders approach compares quite favorably to the D&C heuristic. While the D&C heuristic is faster for small problems with low n and/or low ρ, it can require a significant amount of time for $n = 2000$ or $n = 3000$. In addition to this comparison against the D&C heuristic, in Appendix C.2, we also provide a comparison of the MIO solutions for the smaller T1, T2 and T3 instances used in the previous two sections against three other heuristic solutions; in those instances, we similarly find that solutions obtained from our MIO formulations can be significantly better than heuristic solutions.

Comparing the performance of the Benders approach with the direct solution approach, our results indicate two types of behavior. The first type of behavior corresponds to "easy" instances. These are instances with $\rho \in \{0.02, 0.04\}$ for which it is sometimes possible to directly solve SplitMIO to optimality within the two hour time limit. For example, with $n = 2000$ and $\rho = 0.04$, all five instances are solved to optimality by the direct approach. For those instances, the Benders

n	ρ	b	G_{Benders}	$G_{\text{D\&C}}$	G_{Direct}	T_{Benders}	$T_{\text{D\&C}}$	T_{Direct}	O_{Benders}	O_{Direct}
200	0.02	4	0.00	0.00	0.00	14.54	2.86	2060.81	0.00	0.00
200	0.04	8	0.00	0.04	0.61	187.96	4.46	7015.65	0.00	7.72
200	0.06	12	0.11	0.00	0.77	3615.82	7.34	7200.25	8.23	16.71
200	0.08	16	0.70	0.00	2.86	3612.43	9.97	7200.47	17.38	23.39
200	0.10	20	0.47	0.01	6.94	3613.19	15.07	7200.16	22.13	30.38
200	0.12	24	0.67	0.12	7.68	3614.96	18.93	7200.25	27.27	34.03
500	0.02	10	0.00	0.19	0.00	16.71	11.24	184.29	0.00	0.00
500	0.04	20	0.00	0.20	0.08	1640.78	28.93	7200.21	0.48	3.06
500	0.06	30	0.02	0.68	3.57	3630.19	65.75	7200.26	8.10	11.89
500	0.08	40	0.00	0.48	1.05	3623.51	115.36	7200.57	14.26	14.82
500	0.10	50	0.00	1.49	6.88	3625.89	177.26	7200.18	18.26	23.74
500	0.12	60	0.69	0.95	4.11	3631.14	223.61	7200.38	22.94	25.30
1000	0.02	20	0.00	0.29	0.00	18.83	47.62	156.49	0.00	0.00
1000	0.04	40	0.00	1.65	2.76	2823.19	183.33	7200.18	1.05	4.10
1000	0.06	60	0.00	2.48	6.61	3671.12	397.35	7200.18	6.16	12.53
1000	0.08	80	0.00	2.92	7.86	3662.76	687.02	7200.44	10.07	17.39
1000	0.10	100	0.00	2.31	8.65	3670.95	1052.76	7200.17	13.60	20.99
1000	0.12	120	0.00	2.13	5.92	3696.24	1552.76	7200.22	15.77	20.73
2000	0.02	40	0.00	1.36	0.00	14.55	328.25	70.02	0.00	0.00
2000	0.04	80	0.00	3.49	0.00	1774.85	1259.21	1057.28	0.21	0.00
2000	0.06	120	0.00	4.26	21.38	3774.67	2578.07	7200.16	2.42	23.29
2000	0.08	160	0.00	4.63	100.00	3806.42	4431.39	7200.28	5.21	100.00
2000	0.10	200	0.00	5.23	100.00	3925.28	6435.67	7200.17	7.16	100.00
2000	0.12	240	0.74	0.94	100.00	4319.62	10949.83	7200.16	11.91	100.00
3000	0.02	60	0.00	1.97	0.00	16.16	923.12	32.40	0.00	0.00
3000	0.04	120	0.00	4.03	0.00	1883.80	3365.35	1541.40	0.08	0.00
3000	0.06	180	0.00	5.26	0.04	4144.41	7620.89	7200.16	1.81	1.80
3000	0.08	240	0.00	4.55	99.98	4554.74	13624.07	7209.41	4.23	99.99
3000	0.10	300	0.00	4.04	99.98	5013.31	31137.18	7200.38	6.17	99.98
3000	0.12	360	0.32	1.12	99.98	6999.60	43173.66	7216.75	9.83	99.99

Table 2 Comparison of the Benders decomposition approach, the D&C heuristic and direct solution of SplitMIO in terms of solution quality, computation time and optimality gap.

approach is either able to prove optimality (for example, for $n = 200$ and $\rho = 0.04$, $O_{\text{Benders}} = 0\%$) or terminate with a low optimality gap (for example, for $n = 3000$ and $\rho = 0.04$, $O_{\text{Benders}} = 0.08\%$); among all instances with $\rho \in \{0.02, 0.04\}$, the average optimality gap is no more than about 1.05%. More importantly, the solution obtained by the Benders approach is at least as good as the solution obtained after two hours of direct solution of SplitMIO, which can be seen from the fact that G_{Benders} is 0.0% in all of these (n, ρ) combinations.

The second type of behavior corresponds to "hard" instances, which are the instances with $\rho \in \{0.06, 0.08, 0.10, 0.12\}$. For these instances, Gurobi generally struggles to solve the LO relaxation of SplitMIO within the two hour time limit. When this happens, the integer solution returned by Gurobi is obtained from applying heuristics before solving the root node of the branch-and-bound

tree, which is often quite suboptimal. (This often results in integer solutions with an objective value of 0.0, leading to values of O_{Direct} and G_{Direct} that are close to 100%.) In contrast, the Benders approach delivers significantly better solutions; among all of these instances, G_{Benders} is within 1%, indicating that on average the Benders approach is within 1% of the best solution of each instance. In addition, O_{Benders} is generally lower than O_{Direct}, indicating that the Benders approach in general makes more progress towards proving optimality than the direct approach.

Overall, these results suggest that our Benders approach can deliver high quality solutions to large-scale instances of the assortment optimization problem in a reasonable computational time-frame that are at least as good, and often significantly better, than those obtained by the D&C heuristic or those obtained by directly solving the problem using Gurobi.

6. Numerical Experiments with Real Transaction Data

In this section, we examine the performance of the optimal assortment generated by the decision forest model on problem instances calibrated with a real-world transaction dataset. Different from Section 5, the problem instances in this section are of a smaller scale, which help us glean operational insights. We organize the section as follows. Section 6.1 describes the dataset and the pre-processing steps. Section 6.2 compares the performance of the assortments generated according to the decision forest model and the ranking-based model. Section 6.3 is a case study that demonstrates how the decision forest model factors a well-known behavioral anomaly into its assortment decision.

6.1. Background

We consider the IRI Academic Dataset (Bronnenberg et al. 2008), which is comprised of real-world transaction records of store sales for thirty product categories from forty-seven U.S. markets. The same dataset has been used by Jagabathula and Rusmevichientong (2019) to empirically demonstrate the loss of rationality in consumer choice and by Chen and Mišić (2022) to evaluate the predictive performance of the decision forest model.

We follow the literature to pre-process the raw transaction data. We first aggregate items with the same vendor code as a product, a common pre-processing technique in the marketing science community (Bronnenberg and Mela 2004, Nijs et al. 2007). Following the setup in the literature (Jagabathula and Rusmevichientong 2019, Chen and Mišić 2022) and focusing on the data from the first two weeks of the calendar year 2007 due to the data volume, we select the top nine purchased items as the products and combine the remaining items as the outside/no-purchase option. We further transform the transactions into assortment-choice pairs as follows. We first call $\mathcal{T}$ the set of transactions. For each transaction $t \in \mathcal{T}$, we have the following information: the week of the purchase t_{week}, the store where the transaction happened t_{store}, the sold product t_{prod}, and the selling price t_{price}. Let $\mathcal{W}$ and $\mathcal{Z}$ be the nonrepeated collection of $\{t_{\text{week}}\}_{t\in\mathcal{T}}$ and $\{t_{\text{store}}\}_{t\in\mathcal{T}}$, respectively.

For each week $w \in \mathcal{W}$ and store $s \in \mathcal{Z}$, we define $S_{w,s} = \bigcup_{t \in \mathcal{T}} \{t_{\text{prod}} \mid t_{\text{week}} = w, t_{\text{store}} = s\}$ as the set of products that was purchased at least once at store s during week w. The set of transactions $\mathcal{T}$ is thus transformed into the set of assortment-choice pairs as $\{(S_{t_{\text{week}},t_{\text{store}}}, t_{\text{prod}})\}_{t \in \mathcal{T}}$, which will be used for choice model estimation. We further define the marginal revenue $\bar{r}_i$ as the average of the historical prices $(t_{\text{price}})_{t \in \mathcal{T}: t_{\text{prod}} = i}$ for $i \in \mathcal{N}$.

6.2. Results: Improvement in Expected Revenue

We estimate the ranking-based model and the decision forest model from the assortment-choice pairs $\{(S_{t_{\text{week}},t_{\text{store}}}, t_{\text{prod}})\}_{t \in \mathcal{T}}$ in each product category by maximum likelihood estimation (van Ryzin and Vulcano 2014, Chen and Mišić 2022). In particular, for the decision forest model, we follow Chen and Mišić (2022) to warm start the solver by setting the initial solution as the estimated ranking-based model. For simplicity, we set the tree depth limit as three (i.e., leaves can have depth at most four). Typically, the tree depth is determined by cross validation. Chen and Mišić (2022) have reported that trees of depth three lead to good predictive performance and deeper trees tend to over-fit the data.

We further denote S^{RM} and S^{DF} as the optimal assortments under the estimated ranking-based model and the decision forest model, respectively. We use $Z^*_{\mathcal{C}}$ to denote the maximal expected revenue under a given choice model $\mathcal{C}$ and $Z_{\mathcal{C}}(S)$ to denote the expected revenue of assortment S. Obviously, $Z^*_{\text{RM}} = Z_{\text{RM}}(S^{\text{RM}})$ and $Z^*_{\text{DF}} = Z_{\text{DF}}(S^{\text{DF}})$. To discuss the relative performance of assortments S^{RM} and S^{DF}, we define the following two metrics

$$RI_{\text{DF}} = 100\% \times \frac{Z_{\text{DF}}(S^{\text{DF}}) - Z_{\text{DF}}(S^{\text{RM}})}{Z_{\text{DF}}(S^{\text{RM}})} \quad \text{and} \quad RI_{\text{RM}} = 100\% \times \frac{Z_{\text{RM}}(S^{\text{RM}}) - Z_{\text{RM}}(S^{\text{DF}})}{Z_{\text{RM}}(S^{\text{DF}})}.$$

The two metrics measure the relative improvement of the optimal assortment S^{DF} (or S^{RM}) from S^{RM} (or S^{DF}) under the decision forest model (or the ranking-based model). Note that both metrics are non-negative by their definitions. We also define the Hamming distance $\Delta_{\text{H}} = \sum_{i=1}^{n} |\mathbb{I}[i \in S^{\text{DF}}] - \mathbb{I}[i \in S^{\text{RM}}]|$ to measure how different the two assortments are.

Table 3 summarizes the comparison of S^{DF} and S^{RM} in terms of expected revenue under the measures RI_{DF} and RI_{RM}. When the ground truth is the estimated decision forest model, the assortment S^{DF} can outperform S^{RM} up to 32% and with 7% improvement on average in the expected revenue. Meanwhile, when the ground truth is the estimated ranking-based mode, the assortment S^{DF} only performs 3% worse than the optimal assortment S^{RM} on average. Recall that the class of the ranking-based models is equivalent to the class of choice models of random utility maximization (RUM) principle, or, the class of rational choice models (Jagabathula and Rusmevichientong 2019). Table 3 suggests that the optimal assortment generated by the decision forest model can be quite beneficial when customer choice deviates from the rational choice theory,

Category	RI_{DF}	RI_{RM}	Δ_{H}
Beer	32.8	5.9	3
Blades	0.0	0.0	0
Carbonated Beverages	4.8	3.8	4
Cigarettes	2.5	8.3	1
Coffee	26.3	11.0	4
Cold Cereal	6.1	0.3	2
Deodorant	3.7	2.2	3
Diapers	0.0	0.0	0
Facial Tissue	0.2	0.6	1
Frozen Dinners	2.7	5.2	2
Frozen Pizza	10.9	4.0	4
Household Cleaners	16.0	6.6	4
Hotdogs	18.7	1.4	4
Laundry Detergent	0.1	1.9	1
Margarine / Butter	5.3	12.9	1
Mayonnaise	0.0	0.0	0
Milk	5.2	0.7	1
Mustard / Ketchup	5.6	1.1	1
Paper Towels	1.0	2.4	2
Peanut Butter	10.9	5.6	3
Photo	0.0	0.0	0
Salty Snacks	12.0	3.7	2
Shampoo	19.0	1.4	4
Soup	13.0	5.0	3
Spaghetti / Italian Sauce	4.9	3.3	2
Sugar Substitutes	10.7	6.0	1
Toilet Tissue	0.0	0.0	0
Toothbrush	0.0	0.0	0
Toothpaste	2.7	0.9	2
Yogurt	3.8	2.1	1
Average	7.3	3.2	1.9

Table 3 The relative performance of the assortments generated by the decision forest model and the ranking-based model in each product category in the IRI Academic Dataset

and does not lose much even if customers are strictly rational. We also remark that on average, the two assortments S^{DF} and S^{RM} only differ from each other with Hamming distance 1.9, while they have similar sizes: the average sizes of S^{DF} and S^{RM} are 4.9 and 5.3, respectively.

6.3. Case Study: Beer Category

To investigate why we observed high relative improvement with the assortment S^{DF} over S^{RM} in some categories, we look into the Beer category, where the relative improvement is 32.8%.

We first observe that the consumer choice within the Beer category exhibits a notable phenomenon known as *choice overload*, a well-documented behavioral anomaly in the field of marketing science (Iyengar and Lepper 2000, Chernev et al. 2015, Long et al. 2023). This phenomenon

describes the situation in which an abundance of available products can actually deter customers from making a purchase. Choice overload contradicts with the principles of rational choice theory. According to this theory, when a store increases the assortment size, a rational consumer should theoretically be more inclined to make a purchase from the assortment. This is because a larger assortment would increase the chance that the customer finds one product that aligns with her preference. In the context of choice modeling, a rational choice model adheres to the following principle:

$$P(0 \mid S) \geq P(0 \mid S') \quad \text{whenever} \quad S \subseteq S'. \tag{22}$$

Note that equation (22) directly stems from the *regularity property*, which the ranking-based model and other RUM choice models, such as the mixed-MNL model, always adhere to (Rieskamp et al. 2006, Jagabathula and Rusmevichientong 2019). However, real-world scenarios often deviate from equation (22). When provided with more options, customers may become fatigued from the search process or experience a decrease in their confidence in decision-making, leading them to opt not to make a purchase. In the consumer psychology literature, Iyengar and Lepper (2000) conducted a field experiment demonstrating that individuals are more likely to purchase gourmet jams or chocolates when presented with a smaller assortment of size six rather than a more extensive assortment of size 24 or 30.

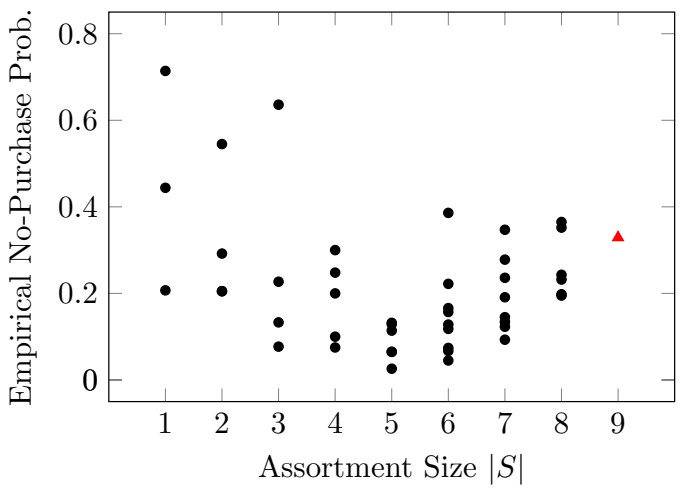

Figure 4 **The choice overload effect in the Beer category. Each scatter point represents the empirical no-purchase probability $\bar{P}(0 \mid S)$ for an historical assortment $S \in \mathcal{S}_{\text{hist}}$.**

Figure 4 visually depicts the prevalence of choice overload within the Beer category of the IRI dataset. Each data point in the figure corresponds to a pair $(|S|, \bar{P}(0 \mid S))$, where $|S|$ represents the assortment size of a historical assortment $S \in \mathcal{S}_{\text{hist}} \equiv \{S_{w,s} \mid w \in \mathcal{W}, s \in \mathcal{S}\}$, and $\bar{P}(0 \mid S)$ represents

the empirical probability of the no-purchase option being chosen when the assortment S was offered. For a better illustration, we have excluded assortments with very few transactions (less than or equal to 10) from Figure 4. In Figure 4, the rightmost data point, denoted by a red triangle, corresponds to the no-purchase probability when the assortment contains all available products ($|S| = 9$). It is important to note that in a scenario where equation (22) holds, indicating the absence of choice overload, each data point should exhibit a higher y-value than the red triangle. However, this condition is met for only 8 out of the 41 data points in the figure, suggesting a violation of equation (22) within the Beer category.

The presence of choice overload in the data has a significant impact on the estimation outcomes of both the ranking-based model and the decision forest model, leading to divergent assortment decisions. To begin, we define $\bar{P}_{\mathcal{C}}(0 \mid m)$ as the average no-purchase probability when presented with an assortment of size m,

$$\bar{P}_{\mathcal{C}}(0 \mid m) = \frac{\sum_{S \subseteq \mathcal{N} \,:\, |S|=m} P_{\mathcal{C}}(0 \mid S)}{|\{S \subseteq \mathcal{N} \,:\, |S| = m\}|}. \tag{23}$$

Here, $P_{\mathcal{C}}(0 \mid S)$ represents the no-purchase probability given assortment S under a choice model $\mathcal{C}$. We consider two choice models, the decision forest model (DF) and the ranking-based model (RM), both estimated from data. Figure 5 presents the average no-purchase probability curves, $\bar{P}_{\mathcal{C}}(0 \mid m)$, for these two models, with the error bars representing the standard deviation. Recall that the ranking-based model adheres to the regularity property and strictly follows equation (22). Consequently, the resulting no-purchase probability curve (depicted in blue) does not exhibit the choice overload effect; it consistently decreases as we expand the assortment size. In contrast, the decision forest model is not bound by the regularity property and has the capacity to learn the choice overload effect from the data. The resulting no-purchase probability curve (shown in red) indicates that customers are indeed more inclined to make a purchase as the assortment initially expands. However, after reaching an assortment size of $|S| = 6$, the choice overload effect becomes evident. Customers become less motivated, and we observe an increase in the no-purchase probability when $|S| \geq 7$.

The extent to which the decision forest model and the ranking-based model capture the phenomenon of choice overload from the data significantly influences their downstream assortment decisions. It is important to note that within the dataset, the marginal revenues $\bar{r}_i$ of the products exhibit low variation, ranging between 8.14 and 10.83 USD (keeping in mind that beer is often sold in packs of six). Since prices are quite uniform and well above zero, the optimal strategy is *not* about making customers choose a specific product but rather about making them buy *any* product from the offered assortment, i.e., reducing the no-purchase probability.

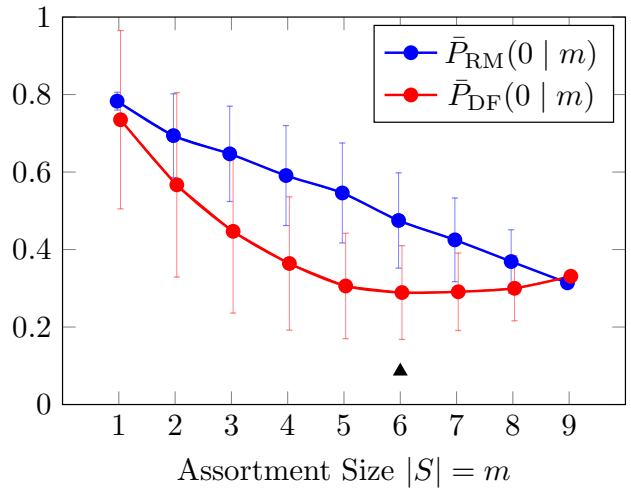

Figure 5 **The average no-purchase probability of the two choice models under different assortment sizes (with the error bars as the standard deviation). The black triangle represents that $P_{\mathsf{DF}}(0 \mid S^{\mathsf{DF}}) = 8.6\%$.**

For the ranking-based model, the optimal assortment S^{RM} includes all products, resulting in $|S^{\mathrm{RM}}| = 9$, since this is the only way it can minimize the no-purchase probability. In contrast, the decision forest model, having captured the choice overload effect from the data, takes a different approach. It examines assortments of size six and identifies one that decreases the no-purchase probability from $P_{\mathrm{DF}}(0 \mid S^{\mathrm{RM}}) = 33.1\%$ to $P_{\mathrm{DF}}(0 \mid S^{\mathrm{DF}}) = 8.6\%$ (see the black triangle in Figure 5). This strategic choice results in a remarkable 32% improvement in expected revenue, measured according to the decision forest model. This highlights the advantage of the assortment planning approach proposed in this paper: by leveraging the decision forest model's ability to capture consumer choice patterns, businesses can tailor their product offerings to effectively capitalize on consumers' departures from strictly rational purchasing behavior.

7. Conclusions

In this paper, we have developed a mixed-integer optimization methodology for solving the assortment optimization problem when the choice model is a decision forest model. This methodology allows a firm to find optimal or near optimal assortments given a decision forest model, which is valuable due to the ability of the decision forest model to capture non-rational customer behavior. We developed two formulations, connected them with the optimization models in the literature, and proposed a large-scale methodology based Benders decomposition. We showed that for one of our formulations it is possible to solve the primal and dual subproblems in the Benders decomposition using a greedy algorithm when the master solution is fractional, and that for both formulations it is possible to solve primal and dual subproblems in closed form when the master solution is binary. Using synthetic data, we demonstrated the scalability of our models and solution methods. We

further used a real-world dataset to showcase how consumers' behavioral anomalies can be learned and capitalized under the framework proposed in this paper.

References

A. Aouad, V. Farias, R. Levi, and D. Segev. The approximability of assortment optimization under ranking preferences. *Operations Research*, 66(6):1661–1669, 2018.

A. Aouad, V. Farias, and R. Levi. Assortment optimization under consider-then-choose choice models. *Management Science*, 2020.

A. Belloni, R. Freund, M. Selove, and D. Simester. Optimizing product line designs: Efficient methods and comparisons. *Management Science*, 54(9):1544–1552, 2008.

G. Berbeglia. The generalized stochastic preference choice model. *arXiv preprint arXiv:1803.04244*, 2018.

D. Bertsimas and V. V. Mišić. Exact first-choice product line optimization. *Operations Research*, 67(3): 651–670, 2019.

D. Bertsimas and R. Weismantel. *Optimization over integers*, volume 13. Dynamic Ideas Belmont, 2005.

J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.

B. J. Bronnenberg and C. F. Mela. Market roll-out and retailer adoption for new brands. *Marketing Science*, 23(4):500–518, 2004.

B. J. Bronnenberg, M. W. Kruger, and C. F. Mela. Database paper—the IRI marketing data set. *Marketing science*, 27(4):745–748, 2008.

J. J. M. Bront, I. Méndez-Díaz, and G. Vulcano. A column generation algorithm for choice-based network revenue management. *Operations research*, 57(3):769–784, 2009.

K. D. Chen and W. H. Hausman. Mathematical properties of the optimal product line selection problem using choice-based conjoint analysis. *Management Science*, 46(2):327–332, 2000.

Y.-C. Chen and V. V. Mišić. Decision forest: A nonparametric approach to modeling irrational choice. *Management Science*, 68(10):7065–7791, 2022.

A. Chernev, U. Böckenholt, and J. Goodman. Choice overload: A conceptual review and meta-analysis. *Journal of Consumer Psychology*, 25(2):333–358, 2015.

I. Contreras, J.-F. Cordeau, and G. Laporte. Benders decomposition for large-scale uncapacitated hub location. *Operations Research*, 59(6):1477–1490, 2011.

J.-F. Cordeau, F. Furini, and I. Ljubić. Benders decomposition for very large scale partial set covering and maximal covering location problems. *European Journal of Operational Research*, 275(3):882–896, 2019.

A. Désir, V. Goyal, D. Segev, and C. Ye. Constrained assortment optimization under the markov chain–based choice model. *Management Science*, 66(2):698–721, 2020.

F. Echenique and K. Saito. General Luce model. *Economic Theory*, 68(4):811–826, 2019.

J. Feldman, A. Paul, and H. Topaloglu. Assortment optimization with small consideration sets. *Operations Research*, 67(5):1283–1299, 2019.

J. B. Feldman and H. Topaloglu. Revenue management under the markov chain choice model. *Operations Research*, 65(5):1322–1342, 2017.

M. Fischetti, I. Ljubić, and M. Sinnl. Redesigning benders decomposition for large-scale facility location. *Management Science*, 63(7):2146–2162, 2017.

A. Flores, G. Berbeglia, and P. Van Hentenryck. Assortment and price optimization under the two-stage luce model. *arXiv preprint arXiv:1706.08599*, 2017.

G. Gallego and H. Topaloglu. *Assortment Optimization*, pages 129–160. Springer New York, 2019.

G. Gallego, R. Ratliff, and S. Shebalov. A general attraction model and sales-based linear program for network revenue management under customer choice. *Operations Research*, 63(1):212–232, 2015.

M. R. Garey and D. S. Johnson. *Computers and intractability.* W. H. Freeman New York, 1979.

P. E. Green and A. M. Krieger. Conjoint analysis with product-positioning applications. *Handbooks in operations research and management science*, 5:467–515, 1993.

J. Huchette and J. P. Vielma. A combinatorial approach for small and strong formulations of disjunctive constraints. *Mathematics of Operations Research*, 44(3):793–820, 2019.

S. S. Iyengar and M. R. Lepper. When choice is demotivating: Can one desire too much of a good thing? *Journal of personality and social psychology*, 79(6):995, 2000.

S. Jagabathula and P. Rusmevichientong. The limit of rationality in choice modeling: Formulation, computation, and implications. *Management Science*, 65(5):2196–2215, 2019.

R. Kivetz, O. Netzer, and V. Srinivasan. Extending compromise effect models to complex buying situations and other context effects. *Journal of Marketing Research*, 41(3):262–268, 2004.

X. Long, J. Sun, H. Dai, D. Zhang, J. Zhang, Y. Chen, H. Hu, and B. Zhao. The choice overload effect in online retailing platforms. *Available at SSRN 3890056*, 2023.

M. Lubin and I. Dunning. Computing in operations research using julia. *INFORMS Journal on Computing*, 27(2):238–248, 2015.

R. D. McBride and F. S. Zufryden. An integer programming approach to the optimal product line selection problem. *Marketing Science*, 7(2):126–140, 1988.

I. Méndez-Díaz, J. J. Miranda-Bront, G. Vulcano, and P. Zabala. A branch-and-cut algorithm for the latent-class logit assortment problem. *Discrete Applied Mathematics*, 164:246–263, 2014.

V. V. Mišić. Optimization of tree ensembles. *Operations Research*, 68(5):1605–1624, 2020.

V. R. Nijs, S. Srinivasan, and K. Pauwels. Retail-price drivers and retailer profits. *Marketing Science*, 26 (4):473–487, 2007.

R. Rahmaniani, T. G. Crainic, M. Gendreau, and W. Rei. The benders decomposition algorithm: A literature review. *European Journal of Operational Research*, 259(3):801–817, 2017.

J. Rieskamp, J. R. Busemeyer, and B. A. Mellers. Extending the bounds of rationality: Evidence and theories of preferential choice. *Journal of Economic Literature*, 44(3):631–661, 2006.

R. P. Rooderkerk, H. J. Van Heerde, and T. H. A. Bijmolt. Incorporating context effects into a choice model. *Journal of Marketing Research*, 48(4):767–780, 2011.

C. Schön. On the optimal product line selection problem with price discrimination. *Management Science*, 56(5):896–902, 2010.

M. Sumida, G. Gallego, P. Rusmevichientong, H. Topaloglu, and J. Davis. Revenue-utility tradeoff in assortment optimization under the multinomial logit model with totally unimodular constraints. *Management Science*, 2020.

K. Talluri and G. Van Ryzin. Revenue management under a general discrete choice model of consumer behavior. *Management Science*, 50(1):15–33, 2004.

G. van Ryzin and G. Vulcano. A market discovery algorithm to estimate a general class of nonparametric choice models. *Management Science*, 61(2):281–300, 2014.

J. P. Vielma and G. L. Nemhauser. Modeling disjunctive constraints with a logarithmic number of binary variables and constraints. *Mathematical Programming*, 128(1-2):49–72, 2011.

J. P. Vielma, S. Ahmed, and G. Nemhauser. Mixed-integer models for nonseparable piecewise-linear optimization: Unifying framework and extensions. *Operations Research*, 58(2):303–315, 2010.

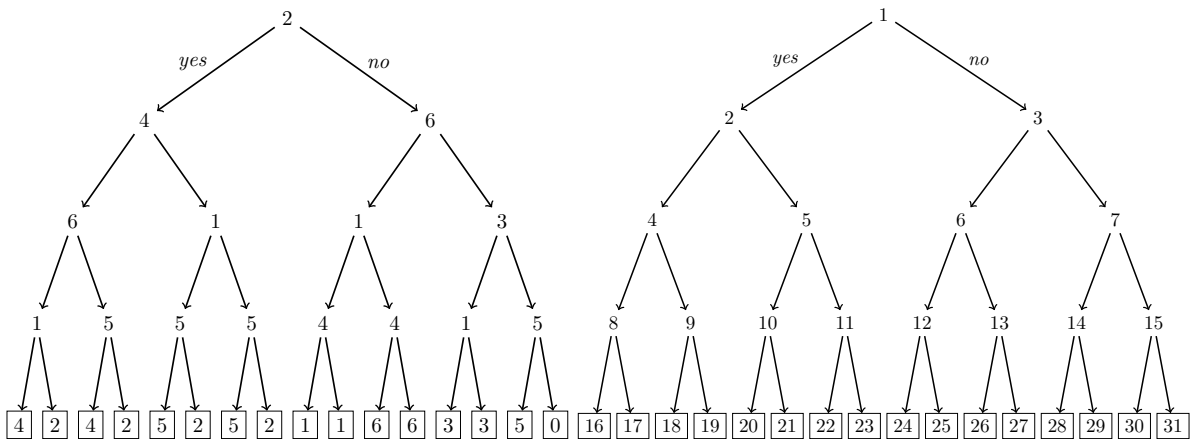

(a) Purchase decision tree. (b) Indexing of nodes in tree.

Figure EC.1 Tree used in example of SplitMIO primal-dual algorithms.

Appendix A: Example of Benders algorithms for SplitMIO

In this section, we provide a small example of the primal-dual procedure (Algorithms 1 and 2) for solving the SplitMIO subproblem. Suppose that $n=6$, and that $\mathbf{x}=(x_1,\ldots,x_6)=(0.62,0.45,0.32,0.86,0.05,0.35)$. Suppose that $\bar{\mathbf{r}}=(\bar{r}_1,\ldots,\bar{r}_6)=(97,72,89,50,100,68)$. Suppose that the purchase decision tree t has the form given in Figure EC.1a; in addition, suppose that the splits and leaves are indexed as in Figure EC.1b. For example, in the latter case, 8 corresponds to the split node that is furthest to the bottom and to the left, while 30 corresponds to the second leaf from the right.

We first run Algorithm 1 on the problem, which carries out the steps shown below in Table EC.1. For this execution of the procedure, we assume that the following ordering of leaves (encoded by τ) is used:

$$20,22,30,24,25,28,29,17,19,21,23,26,27,16,18,31.$$

After running the procedure, the primal solution $\mathbf{y}=(y_{16},\ldots,y_{31})$ follows that $y_{17}=y_{24}=0.35$, $y_{28}=0.15$, $y_{20}=y_{22}=y_{30}=0.05$ and all other components are zero. Here we drop the index t to simplify the notation. The event set is $\mathcal{E}=\{A_{10},A_{11},A_{15},A_3,B_1,C\}$, and the function $f:\mathcal{E}\to\textbf{leaves}$ is defined as $f(A_{10})=20$, $f(A_{11})=22$, $f(A_{15})=30$, $f(A_3)=24$, $f(B_1)=28$, and $f(C)=17$.

We now run Algorithm 2, which carries out the steps shown in Table EC.2. We obtain the dual solution $(\boldsymbol{\alpha},\boldsymbol{\beta},\gamma)$ as follows. $\gamma=72$. $\boldsymbol{\alpha}=(\alpha_1,\ldots,\alpha_{15})$ follows that $\alpha_{10}=\alpha_{11}=28$, $\alpha_{15}=11$, and all other components are zero. $\boldsymbol{\beta}=(\beta_1,\ldots,\beta_{15})$ follows that $\beta_1=17$ and other components are zero.

The feasibility of the dual solution is visualized in Figure EC.2. The colored bars correspond to the different dual variables; a colored bar appears multiple times when the variable participates in multiple dual constraints. The height of the black lines for each leaf indicates the value of r_ℓ, while the total height of the colored bars at a leaf corresponds to the value $\gamma+\sum_{s\in\mathbf{LS}(\ell)}\alpha_s+\sum_{s\in\mathbf{RS}(\ell)}\beta_s$ (the left hand side of the dual constraint (13b). For each leaf, the total height of the colored bars exceeds the black line, which indicates that all dual constraints are satisfied.

Iteration	Values of q_C and $q_{A,B}$	Steps
$\ell = 20$	$q_C = 1.0$, $q_{A,B} = 0.05$	Set $y_{20} \leftarrow 0.05$ A_{10} event
$\ell = 22$	$q_C = 0.95$, $q_{A,B} = 0.05$	Set $y_{22} \leftarrow 0.05$ A_{11} event
$\ell = 30$	$q_C = 0.90$, $q_{A,B} = 0.05$	Set $y_{30} \leftarrow 0.05$ A_{15} event
$\ell = 24$	$q_C = 0.85$, $q_{A,B} = 0.35$	Set $y_{24} \leftarrow 0.35$ A_3 event
$\ell = 25$	$q_C = 0.5$, $q_{A,B} = 0.0$	Set $y_{25} \leftarrow 0.0$
$\ell = 28$	$q_C = 0.5$, $q_{A,B} = 0.15$	Set $y_{28} \leftarrow 0.15$ B_1 event
$\ell = 29$	$q_C = 0.35$, $q_{A,B} = 0.0$	Set $y_{29} \leftarrow 0.0$
$\ell = 17$	$q_C = 0.35$, $q_{A,B} = 0.35$	Set $y_{17} \leftarrow 0.35$ C event **break**

Table EC.1 Steps of primal procedure (Algorithm 1).

Phase	Calculation
Initialization	$\alpha_s \leftarrow 0$, $\beta_s \leftarrow 0$ for all s
Set γ	$\gamma \leftarrow r_{17} = 72$
Loop: $d = 1$	$\beta_1 \leftarrow r_{28} - \gamma = 89 - 72 = 17$
Loop: $d = 2$	$\alpha_3 \leftarrow r_{24} - \gamma - \beta_1 = 97 - 72 - 17 = 8$
Loop: $d = 4$	$\alpha_{10} \leftarrow r_{20} - \gamma = 100 - 72 = 28$
	$\alpha_{11} \leftarrow r_{22} - \gamma = 100 - 72 = 28$
	$\alpha_{15} \leftarrow r_{30} - \gamma - \beta_1 = 100 - 72 - 11 = 11$

Table EC.2 Steps of dual procedure (Algorithm 2).

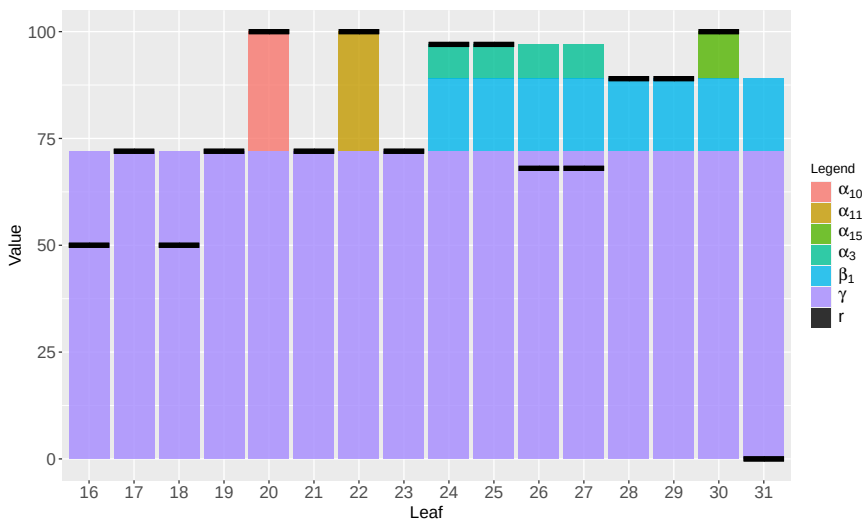

Figure EC.2 Visualization of feasibility of dual solution in the SplitMIO example.

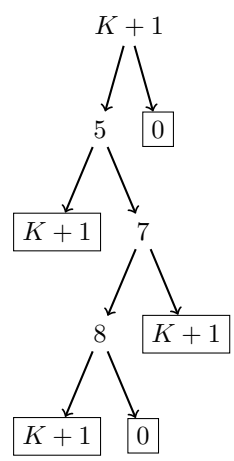

Figure EC.3 Tree for clause $x_5 \vee \neg x_7 \vee x_8$.

The objective value of the primal solution is $\sum_{\ell \in \textbf{leaves}} r_\ell \cdot y_\ell = r_{20} \times y_{20} + r_{22} \times y_{22} + r_{30} \times y_{30} + r_{24} \times y_{24} + r_{28} \times y_{28} + r_{17} \times y_{17} = 87.5$. The objective value of the dual solution is $\gamma + (1 - x_2) \times \beta_1 + x_6 \times \alpha_3 + x_5 \times \alpha_{10} + x_5 \times \alpha_{11} + x_5 \times \alpha_{15} = 72.0 + 0.55 \times 17 + 0.35 \times 8 + 0.05 \times 28 + 0.05 \times 28 + 0.05 \times 11 = 87.5$

Appendix B: Proofs

B.1. Proof of NP-Hardness

We show that MAX 3SAT problem reduces to the decision forest assortment optimization problem. In the MAX 3SAT problem, one has K binary variables, $x_1, \ldots, x_K$, and is given a Boolean formula of the form $c_1 \wedge c_2 \wedge \cdots \wedge c_M$, where $\wedge$ denotes "and". Each clause is a disjunction involving three literals, where a literal

is either one of the binary variables or its negation, and the literals involve distinct binary variables. For example, a clause could be $x_5 \vee \neg x_7 \vee x_8$, where $\vee$ denotes "or". The MAX 3SAT problem is to find the assignment of the variables $x_1, \ldots, x_K$ so as to maximize the number of clauses $c_1, \ldots c_T$ that are true.

Given an instance of the MAX 3SAT problem, we show how the problem can be transformed into an instance of the decision forest assortment optimization problem (2). Consider an instance of problem (2) with $n = K+1$ products. Each of the first K products corresponds to one of the binary variables; the last $(K+1)$-th product is necessary to ensure that the revenue of the assortment can correspond to the number of satisfied clauses. Assume that the marginal revenues of the products are set so that $\bar{r}_1 = \cdots = \bar{r}_K = 0$, and $\bar{r}_{K+1} = 1$. For each clause $m \in \{1, \ldots, M\}$, introduce a tree t_m which is constructed as follows:

1. Set the root node of the tree to be a split node involving product $K+1$. The right child of the root node is a leaf node with 0 (the no-purchase option) as the purchase decision. (The left child node will be defined in the next step.)

2. For the first literal, create a split at the left child node of the root, with the corresponding binary variable's index as the product (e.g., if the literal is x_7 or $\neg x_7$, the split node has product 7). If the literal is the binary variable itself (i.e., x_k), we set the left child node of the split to be a leaf node with product $K+1$ as the purchase decision. Otherwise, if the literal is the negation of the binary variable (i.e., $\neg x_k$), we set the right child node of the split to be a leaf node with product $K+1$ as the purchase decision.

3. For the other child node of the split created in Step 2, repeat Step 2 with the second literal.

4. For the other child node of the split created in Step 3, repeat Step 2 with the third literal.

5. Lastly, for the other child node of the third split node in Step 4 corresponding to the third literal, set the purchase decision to be 0.

Figure EC.3 visualizes the structure of the tree for the example clause $x_5 \vee \neg x_7 \vee x_8$. Applying the above procedure for each clause results in a forest F consisting of M trees. Lastly, we set the probability distribution $\boldsymbol{\lambda}$ by setting $\lambda_t = 1/M$ for each tree $t \in F$.

We note that in the resulting instance of problem (2), any optimal assortment must include produce $K+1$: due to the structure of the trees, the expected revenue is exactly equal to 0 if $K+1$ is not included in the assortment, but by including $K+1$ and including/excluding products from $\{1, \ldots, K\}$ in accordance with one of the clauses, one can obtain an expected revenue of at least $1/M$. (For example, if the set of clauses includes the clause $x_5 \vee \neg x_7 \vee x_8$ shown in Figure EC.3, then any assortment S that includes products 5 and 8 and does not include product 7 automatically has an expected revenue of at least $1/M$.)

Note that given an optimal assortment $S \subseteq \mathcal{N} = \{1, \ldots, K+1\}$, we immediately obtain an assignment $\mathbf{x}$ for the MAX 3SAT problem by setting $x_i = \mathbb{I}\{i \in S\}$. The revenue obtained from each tree t_m corresponding to clause m, which is given by $\sum_{j=1}^{K+1} \bar{r}_j \mathbb{I}\{\hat{A}(t_m, S) = j\}$, is exactly 1 if the assignment $\mathbf{x}$ that corresponds to S satisfies clause m, and 0 otherwise; this holds because the optimal assortment must include product $K+1$. Since each tree t_m has a probability of $1/M$, the expected revenue of the assortment S therefore corresponds to the fraction of the M clauses that are satisfied by the assignment $\mathbf{x}$. Thus, it follows that an assignment $\mathbf{x}$ that corresponds to an optimal assortment S is an optimal solution of MAX 3SAT. Since the MAX 3SAT problem is NP-Complete (Garey and Johnson 1979), it follows that problem (2) is NP-Hard. $\square$

B.2. Proof of Proposition 2

The feasible region $\mathcal{F}_{\text{SplitMIO}}$ of the LO relaxation of problem (3) is the set of $(\mathbf{x}, \mathbf{y})$ solutions to the following system of inequalities:

$$\sum_{\ell \in \textbf{leaves}} y_\ell \leq 1, \tag{EC.1a}$$

$$\sum_{\ell \in \textbf{leaves}} -y_\ell \leq -1, \tag{EC.1b}$$

$$\sum_{\ell \in \textbf{left}(s)} y_\ell - x_{v(s)} \leq 0, \quad \forall\, s \in \textbf{splits}, \tag{EC.1c}$$

$$\sum_{\ell \in \textbf{right}(s)} y_\ell + x_{v(s)} \leq 1, \quad \forall\, s \in \textbf{splits}, \tag{EC.1d}$$

$$x_i \leq 1, \quad \forall\, i \in \mathcal{N}, \tag{EC.1e}$$

$$x_i \geq 0, \quad \forall\, i \in \mathcal{N}, \tag{EC.1f}$$

$$y_\ell \geq 0, \quad \forall\, \ell \in \textbf{leaves}. \tag{EC.1g}$$

(Since $|F| = 1$, we drop the index t to simplify the notation.) In the above, note that the unit sum constraint on $\mathbf{y}$ has been re-written as a pair of inequalities, and that all constraints from problem (3) have been re-arranged to have the variables on one side. This system of inequalities can be written compactly as

$$\mathbf{A} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \leq \mathbf{b}, \qquad \mathbf{x}, \mathbf{y} \geq \mathbf{0}. \tag{EC.2}$$

To show that $\mathcal{F}_{\text{SplitMIO}}$ is integral, we will prove that the matrix $\mathbf{A}$ is totally unimodular. We do so using the following standard characterization of total unimodularity (see Bertsimas and Weismantel 2005):

Proposition EC.1 **(Corollary 3.2 of Bertsimas and Weismantel 2005)**. *A matrix* $\mathbf{A}$ *is totally unimodular if and only if each collection* Q *of rows of* $\mathbf{A}$ *can be partitioned into two parts so that the sum of the rows in one part minus the sum of the rows in the other part is a vector with entries only 0, +1, and -1.*

To simplify our notation, we will work with algebraic expressions in terms of $\mathbf{x}$ and $\mathbf{y}$ rather than rows of the matrix $\mathbf{A}$. We have the following four primitive expressions:

$$A(s), s \in \textbf{splits}: \quad \sum_{\ell \in \textbf{left}(s)} y_\ell - x_{v(s)}, \tag{EC.3}$$

$$B(s), s \in \textbf{splits}: \quad \sum_{\ell \in \textbf{right}(s)} y_\ell + x_{v(s)}, \tag{EC.4}$$

$$C(i), i \in \mathcal{N}: \quad x_i, \tag{EC.5}$$

$$D(1): \quad \sum_{\ell \in \textbf{leaves}} y_\ell, \tag{EC.6}$$

$$D(2): \quad \sum_{\ell \in \textbf{leaves}} -y_\ell. \tag{EC.7}$$

Thus, a collection of rows Q of the matrix $\mathbf{A}$ can be viewed as a collection of each of the four types of expressions above:

$$S_A \subseteq \textbf{splits}, \tag{EC.8}$$

$$S_B \subseteq \textbf{splits}, \tag{EC.9}$$

$$S_C \subseteq \mathcal{N}, \tag{EC.10}$$

$$S_D \subseteq \{1, 2\}. \tag{EC.11}$$

To establish the condition in Proposition EC.1, we need to show that given S_A, S_B, S_C, S_D, we can partition these expressions into two groups R_+ and R_- such that the difference of the two groups,

$$\sum_{e \in R_+} e - \sum_{e \in R_-} e = \sum_{i \in \mathcal{N}} v_i x_i + \sum_{\ell \in \textbf{leaves}} w_\ell y_\ell, \tag{EC.12}$$

is such that each $v_i \in \{-1, 0, +1\}$ and each $w_\ell \in \{-1, 0, +1\}$. We proceed in several steps.

Step 1. We begin by assigning the A and B expressions. Before doing so, we require some additional notation. Define $d^* = \max_{s \in S_A \cup S_B} d(s)$, where $d(s)$ is the depth of split s and we assume the depth of the root node is 1. Let us define the sets $S_A(d) = \{s \in S_A \mid d(s) = d\}$ and $S_B(d) = \{s \in S_B \mid d(s) = d\}$ for each depth $d \in \{1, \ldots, d^*\}$. These are the sets of splits in S_A and S_B, respectively, that are at a particular depth. Let us also define $\mathbf{LD}(s)$ and $\mathbf{RD}(s)$ to be the sets of splits in $S_A \cup S_B$ that are to the left and right, respectively, of split $s \in S_A \cup S_B$. (The splits in $\mathbf{LD}(s)$ are all those that can be reached by proceeding to the left child of split s; similarly, $\mathbf{RD}(s)$ is the set of splits reachable by going to the right of split s.) Finally, define $\sigma : S_A \cup S_B \to \{-1, +1\}$ to be a mapping that is specified according to the following procedure:

for $d = 1, \ldots, d^*$ **do**
  for $s \in S_A(d)$ **do**
    Set $\sigma(s') = (-1)\sigma(s)$ for $s' \in \mathbf{LD}(s)$
  for $s \in S_B(d)$ **do**
    Set $\sigma(s') = (-1)\sigma(s)$ for $s' \in \mathbf{RD}(s)$

Now, assign the A and B expressions as follows:

- Assign $A(s)$ to R_+ for each $s \in S_A$ with $\sigma(s) = +1$;
- Assign $A(s)$ to R_- for each $s \in S_A$ with $\sigma(s) = -1$;
- Assign $B(s)$ to R_+ for each $s \in S_B$ with $\sigma(s) = +1$;
- Assign $B(s)$ to R_- for each $s \in S_B$ with $\sigma(s) = -1$.

For this assignment of the expressions in S_A and S_B, every y_ℓ coefficient in $\sum_{e \in R_+} e - \sum_{e \in R_-} e$ will be either 0 or +1. This follows because the sets of left and right leaves $\mathbf{left}(s)$ and $\mathbf{right}(s)$ are nested. In particular, if $s' \in \mathbf{LD}(s)$, then we will have that $\mathbf{left}(s') \subseteq \mathbf{left}(s)$ and $\mathbf{right}(s') \subseteq \mathbf{left}(s)$. Similarly, if $s' \in \mathbf{RD}(s)$, then we will have that $\mathbf{left}(s') \subseteq \mathbf{right}(s)$ and $\mathbf{right}(s') \subseteq \mathbf{right}(s)$.

In addition, by the assumption that there is at most one split s in **splits** such that $v(s) = i$, we are also guaranteed that the coefficient of every x_i in $\sum_{e \in R_+} e - \sum_{e \in R_-} e$ will be $\{-1, 0, +1\}$. In particular, if s is in both S_A and S_B (i.e., we were given the expression $A(s)$ and $B(s)$), then observe that by the procedure for setting σ above, we are guaranteed to assign both $A(s)$ and $B(s)$ to the same set (they cannot be assigned to different sets). This means that the coefficients of $x_{v(s)}$ in $A(s)$ and $B(s)$ will cancel out, leaving $x_{v(s)}$ with a coefficient of 0.

Step 2. We next assign the C expressions. After Step 1, we are guaranteed that the coefficient of each x_i in $\sum_{e \in R_+} e - \sum_{e \in R_-} e$ is 0, -1 or +1. Since each $C(i)$ expression involves only one variable (x_i), it is straightforward to assign these expressions to R_+ and R_- to ensure that every variable's coefficient in $\sum_{e \in R_+} e - \sum_{e \in R_-} e$ is 0, -1 or +1. For completeness, we give the procedure below – for each $i \in S_C$:

- If $v(s) \neq i$ for all splits $s \in \textbf{splits}$, then x_i does not appear in any A or B expressions and its coefficient after Step 1 is just 0; thus, $C(i)$ can be arbitrarily assigned to R_+ or R_-.
- If there exists an $s \in S_A \cup S_B$ such that $v(s) = i$, then:

—If $s \in S_A \cup S_B$, then $A(s)$ and $B(s)$ were both assigned to R_+ and R_-, and so the coefficient of x_i will be 0 due to cancellation; thus, $C(i)$ can again be arbitrarily assigned to R_+ or R_-.

—If $s \in S_A$, $s \notin S_B$, and $\sigma(s) = +1$, then the coefficient of x_i is -1 after Step 1; thus, $C(i)$ should be assigned to R_+.

—If $s \in S_A$, $s \notin S_B$, and $\sigma(s) = -1$, then the coefficient of x_i is +1 after Step 1; thus, $C(i)$ should be assigned to R_-.

—If $s \notin S_A$, $s \in S_B$, and $\sigma(s) = +1$, then the coefficient of x_i is +1 after Step 1; thus, $C(i)$ should be assigned to R_-.

—If $s \notin S_A$, $s \in S_B$, and $\sigma(s) = -1$, then the coefficient of x_i is -1 after Step 1; thus, $C(i)$ should be assigned to R_+.

Step 3. Lastly, we assign the D expressions. This step is also straightforward:

- If $S_D = \{1,2\}$, then assign $D(1)$ and $D(2)$ to R_+; since $D(1)$ is just the negative of $D(2)$, the two expressions will cancel out, and the expression $\sum_{e \in R_+} e - \sum_{e \in R_-} e$ will remain unchanged.
- If $S_D = \{1\}$, then assign $D(1)$ to R_-; since the coefficient of each y_ℓ is 0 or +1 after Step 2, this will ensure that the coefficient of each y_ℓ is either -1 or 0.
- If $S_D = \{2\}$, then assign $D(2)$ to R_-; since the coefficient of each y_ℓ is 0 or +1 after Step 2, this will ensure that the coefficient of each y_ℓ is either -1 or 0.

After completing Step 3, we have assigned all of the expressions in S_A, S_B, S_C, S_D to the sets R_+ and R_- in a way that each expression is assigned to exactly one of the two sets, and no expression is unassigned. Moreover, the difference of the two expressions, $\sum_{e \in R_+} e - \sum_{e \in R_-} e$, is such that the coefficient of every x_i and y_ℓ variable is in $\{0, -1, +1\}$. By Proposition EC.1, this establishes that the matrix $\mathbf{A}$ is totally unimodular. We now employ another standard result:

Proposition EC.2 **(Theorem 3.1(b) of Bertsimas and Weismantel 2005)**. *Let* $\mathbf{A}$ *be an integer matrix. The matrix* $\mathbf{A}$ *is totally unimodular if and only if the polyhedron* $P(\mathbf{b}) = \{\mathbf{x} \in \mathbb{R}^n_+ \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}\}$ *is integral for all* $\mathbf{b} \in \mathbb{Z}^m$ *for which* $P(\mathbf{b}) \neq \emptyset$.

To use this result, we simply have to establish that the feasible region of the polyhedron defined in (EC.1) is nonempty. To do so, we explicitly construct a feasible solution to (EC.1). Let r be the root node of the tree. Set $x_i = 0.5$ for all $i \in \mathcal{N}$. Fix any leaf $\ell' \in \textbf{left}(r)$ and any leaf $\ell'' \in \textbf{right}(r)$, and set $y_{\ell'} = 0.5$, $y_{\ell''} = 0.5$, and $y_\ell = 0$ for all $\ell \in \textbf{leaves} \setminus \{\ell', \ell''\}$. It is straightforward to verify that this solution satisfies the system of inequalities (EC.1): the y_ℓ's sum to 1 and are nonnegative by construction, and each $x_i \in [0,1]$ by construction. For constraints (EC.1c) and (EC.1d), note that since ℓ' and ℓ'' are on opposite sides of the root node, it is impossible for ℓ' and ℓ'' to both belong to $\textbf{left}(s)$ and $\textbf{right}(s)$ for any split s; armed with this fact, it is straightforward to establish the two constraints.

Since we have established that $\mathbf{A}$ is totally unimodular and that the set $\mathcal{F}_{\text{SplitMIO}}$ is nonempty, invoking Proposition EC.2 concludes the proof. □

B.3. Proof of Proposition 4

Define $\Delta^{\mathbf{leaves}} = \{\mathbf{y} \in \mathbb{R}^{|\mathbf{leaves}|} \mid \sum_{\ell \in \mathbf{leaves}} y_\ell = 1; y_\ell \geq 0, \forall\, \ell \in \mathbf{leaves}\}$ to be the $(|\mathbf{leaves}|-1)$-dimensional unit simplex. In addition, for any $S \subseteq \mathbf{leaves}$, define $Q(S) = \{\mathbf{y} \in \Delta^{\mathbf{leaves}} \mid y_\ell \leq 0 \text{ for } \ell \in \mathbf{leaves} \setminus S\}$. We write the combinatorial disjunctive constraint over the ground set $\mathbf{leaves}$ as $\text{CDC}(\mathbf{leaves}) = \bigcup_{\ell \in \mathbf{leaves}} Q(\{\ell\})$. Consider now the optimization problem

$$\underset{\mathbf{y}}{\text{maximize}} \left\{ \sum_{\ell \in \mathbf{leaves}} r_\ell y_\ell \;\middle|\; \mathbf{y} \in \text{CDC}(\mathbf{leaves}) \right\}. \tag{EC.13}$$

We will re-formulate this problem into a mixed-integer optimization problem. To do this, we claim that $\text{CDC}(\mathbf{leaves})$ can be written as the following pairwise independent branching scheme:

$$\bigcup_{\ell \in \mathbf{leaves}} Q(\{\ell\}) = \bigcup_{i=1}^{n} (Q(L_i) \cup Q(R_i)), \tag{EC.14}$$

where $L_i = \{\ell \in \mathbf{leaves} \mid \ell \in \mathbf{left}(s) \text{ for some } s \text{ with } v(s) = i\}$ and $R_i = \{\ell \in \mathbf{leaves} \mid \ell \in \mathbf{right}(s) \text{ for some } s \text{ with } v(s) = i\}$. Note that (EC.14) is equivalent to the statement

$$\bigcup_{\ell \in \mathbf{leaves}} \{\ell\} = \bigcup_{i=1}^{n} ((\mathbf{leaves} \setminus L_i) \cup (\mathbf{leaves} \setminus R_i)). \tag{EC.15}$$

To establish (EC.15), it is sufficient to prove the following equivalence:

$$\{\ell\} = \bigcap_{i \in I(\ell)} (\mathbf{leaves} \setminus R_i) \cap \bigcap_{i \in E(\ell)} (\mathbf{leaves} \setminus L_i) \cap \bigcap_{\substack{i \in \mathcal{N}: \\ i \notin I(\ell) \cup E(\ell)}} (\mathbf{leaves} \setminus L_i), \tag{EC.16}$$

where $I(\ell) = \{i \in \mathcal{N} \mid \ell \in \bigcup_{s:v(s)=i} \mathbf{left}(s)\}$ and $E(\ell) = \{i \in \mathcal{N} \mid \ell \in \bigcup_{s:v(s)=i} \mathbf{right}(s)\}$.

We now prove (EC.16). *Equation* (EC.16), $\subseteq$ *direction:* For $i \in I(\ell)$, we have that $\ell \in \bigcup_{s:v(s)=i} \mathbf{left}(s)$. This means that there exists $\bar{s}$ such that $\ell \in \mathbf{left}(\bar{s})$ and $v(\bar{s}) = i$. Since $\ell \in \mathbf{left}(\bar{s})$, this means that $\ell \notin \mathbf{right}(\bar{s})$ (a leaf cannot be to the left and to the right of any split). Moreover, ℓ cannot be in $\mathbf{right}(s)$ for any other s with $v(s) = i$, because this would mean that product i appears more than once along the path to leaf ℓ, violating Assumption 1. Therefore, $\ell \in \mathbf{leaves} \setminus R_i$ for any $i \in I(\ell)$.

For $i \in E(\ell)$, we have that $\ell \in \bigcup_{s:v(s)=i} \mathbf{right}(s)$. This means that there exists a split $\bar{s}$ such that $\ell \in \mathbf{right}(\bar{s})$ and $v(\bar{s}) = i$. Since $\ell \in \mathbf{right}(\bar{s})$, we have that $\ell \notin \mathbf{left}(\bar{s})$. In addition, ℓ cannot be in $\mathbf{left}(s)$ for any other s with $v(s) = i$. Therefore, $\ell \in \mathbf{leaves} \setminus L_i$ for any $i \in E(\ell)$.

Lastly, for any $i \notin I(\ell) \cup E(\ell)$, note that product i does not appear in any split along the path from the root of the tree to leaf ℓ. Therefore, for any s with $v(s) = i$, it will follow that either $\mathbf{left}(s) \subseteq \mathbf{left}(s')$ for some s' for which $\ell \in \mathbf{right}(s')$, or $\mathbf{left}(s) \subseteq \mathbf{right}(s')$ for some s' for which $\ell \in \mathbf{left}(s')$ – in other words, there is a split s' such that every leaf in $\mathbf{left}(s)$ is to one side of s' and ℓ is on the other side of s'. This means that ℓ' cannot be in $\mathbf{left}(s)$ for any s with $v(s) = i$, or equivalently, $\ell \in \mathbf{leaves} \setminus L_i$ for any $i \notin I(\ell) \cup E(\ell)$.

Equation (EC.16), $\supseteq$ *direction:* We will prove the contrapositive $\{\ell' \in \mathbf{leaves} \mid \ell' \neq \ell\} \subseteq \bigcup_{i \in I(\ell)} R_i \cup \bigcup_{i \in E(\ell)} L_i \cup \bigcup_{\substack{i \in \mathcal{N}: \\ i \notin I(\ell) \cup E(\ell)}} L_i$. A straightforward result (see Lemma EC.1 from Mišić 2020) is that $\{\ell' \in \mathbf{leaves} \mid \ell' \neq \ell\} = \bigcup_{s:\ell \in \mathbf{left}(s)} \mathbf{right}(s) \cup \bigcup_{s:\ell \in \mathbf{right}(s)} \mathbf{left}(s)$. Thus, if $\ell' \neq \ell$, then we have that $\ell' \in \mathbf{right}(s)$ for some s

such that $\ell \in \mathbf{left}(s)$, or $\ell' \in \mathbf{left}(s)$ for some s such that $\ell \in \mathbf{right}(s)$. Let $i^* = v(s)$. In the first case, since $\ell \in \mathbf{left}(s)$, we have that $i^* \in I(\ell)$, and we thus have

$$\mathbf{right}(s) \subseteq \bigcup_{s':v(s')=i^*} \mathbf{right}(s') = R_{i^*} \subseteq \bigcup_{i \in I(\ell)} R_i \cup \bigcup_{i \in E(\ell)} L_i \cup \bigcup_{\substack{i' \in \mathcal{N}: \\ i' \notin I(\ell) \cup E(\ell)}} L_i.$$

In the second case, since $\ell \in \mathbf{right}(s)$, we have that $i^* \in E(\ell)$, and thus we have

$$\mathbf{left}(s) \subseteq \bigcup_{s':v(s')=i^*} \mathbf{left}(s') = L_{i^*} \subseteq \bigcup_{i \in I(\ell)} R_i \cup \bigcup_{i \in E(\ell)} L_i \cup \bigcup_{\substack{i' \in \mathcal{N}: \\ i' \notin I(\ell) \cup E(\ell)}} L_i.$$

This establishes the validity of the pairwise independent branching scheme (EC.15). Thus, a valid formulation for problem (EC.13) (see formulation (9) in Vielma et al. 2010, formulation (14) in Vielma and Nemhauser 2011 and formulation (13) in Huchette and Vielma 2019) is

$$\begin{aligned}
\underset{\mathbf{x},\mathbf{y}}{\text{maximize}} \quad & \sum_{\ell \in \mathbf{leaves}} r_\ell y_\ell && \text{(EC.17a)}\\
\text{subject to} \quad & \sum_{\ell \in \mathbf{leaves}} y_\ell = 1, && \text{(EC.17b)}\\
& \sum_{\ell \in L_i} y_\ell \le x_i, \quad \forall\, i \in \mathcal{N}, && \text{(EC.17c)}\\
& \sum_{\ell \in R_i} y_\ell \le 1 - x_i, \quad \forall\, i \in \mathcal{N}, && \text{(EC.17d)}\\
& y_\ell \ge 0, \quad \forall\, \ell \in \mathbf{leaves}, && \text{(EC.17e)}\\
& x_i \in \{0,1\}, \quad \forall\, i \in \mathcal{N}. && \text{(EC.17f)}
\end{aligned}$$

Observe that, by the definition of L_i and R_i, formulation (EC.17) is identical to ProductMIO when $|F| = 1$. By invoking Theorem 1 from Vielma et al. (2010) with appropriate modifications, we can assert that formulation (EC.17) is integral. Therefore, when $|F| = 1$, formulation ProductMIO is always integral. □

B.4. Proof of Theorem 1

Proof of part (a) (feasibility): By definition, the solution produced by Algorithm 1 produces a solution $\mathbf{y}_t$ that satisfies the left and right split constraints (12c) and (12d). With regard to the nonnegativity constraint (12e), we can see that at each stage of Algorithm 1, the quantities $x_{v(t,s)} - \sum_{\ell \in \mathbf{left}(s)} y_{t,\ell}$, $1 - x_{v(t,s)} - \sum_{\ell \in \mathbf{right}(s)} y_{t,\ell}$ and $1 - \sum_{\ell \in \mathbf{leaves}(t)} y_{t,\ell}$ never become negative; thus, the solution $\mathbf{y}_t$ produced upon termination satisfies the nonnegativity constraint (12e).

The only constraint that remains to be verified is constraint (12b), which requires that $\mathbf{y}_t$ adds up to 1. Observe that it is sufficient for a C event to occur during the execution of Algorithm 1 to ensure that constraint (12b) is satisfied. We will show that a C event must occur during the execution of Algorithm 1.

We proceed by contradiction. For the sake of a contradiction, let us suppose that a C event does not occur during the execution of the algorithm. Note that under this assumption, for any split s, it is impossible that the solution $\mathbf{y}_t$ produced by Algorithm 1 satisfies $x_{v(t,s)} = \sum_{\ell \in \mathbf{left}(s)} y_{t,\ell}$ and $1 - x_{v(t,s)} = \sum_{\ell \in \mathbf{right}(s)} y_{t,\ell}$, as this would imply that $\sum_{\ell \in \mathbf{left}(s)} y_{t,\ell} + \sum_{\ell \in \mathbf{right}(s)} y_{t,\ell} = x_{v(t,s)} + 1 - x_{v(t,s)} = 1$; by the definition of Algorithm 1, this would have triggered a C event at one of the leaves in $\mathbf{left}(s) \cup \mathbf{right}(s)$.

Thus, this means that at every split, either $\sum_{\ell \in \mathbf{left}(s)} y_{t,\ell} < x_{v(t,s)}$ or $\sum_{\ell \in \mathbf{right}(s)} < 1 - x_{v(t,s)}$. Using this property, let us identify a leaf ℓ^* using the following procedure:

Procedure 3:
1. Set $j \leftarrow \mathbf{root}(t)$.
2. If $j \in \mathbf{leaves}(t)$, terminate with $\ell^* = j$; otherwise, proceed to Step 3.
3. If $\sum_{\ell \in \mathbf{left}(j)} y_{t,\ell} < x_{v(t,j)}$, set $j \leftarrow \mathbf{leftchild}(j)$; otherwise, set $j \leftarrow \mathbf{rightchild}(j)$.
4. Repeat Step 2.

Note that by our observation that at most one of the left or right split constraints can be satisfied at equality for any split s, Procedure 3 above is guaranteed to terminate with a leaf ℓ^* such that:

$$y_{t,\ell^*} \leq \sum_{\ell \in \mathbf{left}(s)} y_{t,\ell} < x_{v(t,s)}, \quad \forall\, s \in \mathbf{splits}(t) \text{ such that } \ell^* \in \mathbf{left}(s),$$
$$y_{t,\ell^*} \leq \sum_{\ell \in \mathbf{right}(s)} y_{t,\ell} < 1 - x_{v(t,s)}, \quad \forall\, s \in \mathbf{splits}(t) \text{ such that } \ell^* \in \mathbf{right}(s).$$

However, this is impossible, because Algorithm 1 always sets each leaf $y_{t,\ell}$ to the highest value it can be without violating any of the left or right split constraints; the above conditions imply that y_{t,ℓ^*} could have been set higher, which is not possible. We thus have a contradiction, and it must be the case that a C event occurs.

Proof of part (b) (extreme point): To show that $\mathbf{y}_t$ is an extreme point, let us assume that $\mathbf{y}_t$ is not an extreme point. Then, there exist feasible solutions $\mathbf{y}_t^1$ and $\mathbf{y}_t^2$ different from $\mathbf{y}_t$ and a weight $\theta \in (0,1)$ such that $\mathbf{y}_t = \theta \mathbf{y}_t^1 + (1-\theta)\mathbf{y}_t^2$. Let ℓ^* be the first leaf checked by Algorithm 1 at which $y_{t,\ell^*} \neq y_{t,\ell^*}^1$ and $y_{t,\ell^*} \neq y_{t,\ell^*}^2$. Such a leaf must exist because $\mathbf{y}_t \neq \mathbf{y}_t^1$ and $\mathbf{y}_t \neq \mathbf{y}_t^2$, and because $\mathbf{y}_t$ is the convex combination of $\mathbf{y}_t^1$ and $\mathbf{y}_t^2$. Without loss of generality, let us further assume that $y_{t,\ell^*}^1 < y_{t,\ell^*} < y_{t,\ell^*}^2$.

By definition, Algorithm 1 sets $y_{t,\ell}$ at each iteration to the largest it can be without violating the left split constraints (12c) and the right split constraints (12d), and ensuring that $\sum_{\ell \in \mathbf{leaves}(t)} y_{t,\ell}$ does not exceed 1. Since $y_{t,\ell^*}^2 > y_{t,\ell^*}$, and since $\mathbf{y}_t^2$ and $\mathbf{y}_t$ are equal for all leaves checked before ℓ^*, this implies that $\mathbf{y}_t^2$ either violates constraint (12c), violates constraint (12d), or is such that $\sum_{\ell \in \mathbf{leaves}(t)} y_{t,\ell} > 1$. This implies that $\mathbf{y}_t^2$ cannot be a feasible solution, which contradicts the assumption that $\mathbf{y}_t^2$ is a feasible solution. □

B.5. Proof of Theorem 2 (SplitMIO dual is BFS)

Proof of part (a) (feasibility): Before we prove the result, we first establish a helpful property of the events that are triggered during the execution of Algorithm 1.

LEMMA EC.1. *Let $s_1, s_2 \in \mathbf{splits}(t)$, $s_1 \neq s_2$, such that s_2 is a descendant of s_1. Suppose that $e_1 = A_{s_1}$ or $e_1 = B_{s_1}$, and that $e_2 = A_{s_2}$ or $e_2 = B_{s_2}$. If e_1 and e_2 occur during the execution of Algorithm 1, then $r_{t,f(e_1)} \leq r_{t,f(e_2)}$.*

Proof: We will prove this by contradiction. Suppose that we have two splits s_1 and s_2 and events e_1 and e_2 as in the statement of the lemma, and that $r_{t,f(e_2)} < r_{t,f(e_1)}$. This implies that leaf $f(e_1)$ is checked before leaf $f(e_2)$. When leaf $f(e_1)$ is checked, the event e_1 occurs, which implies that either the left split constraint (12c) becomes tight (if $e_1 = A_{s_1}$) or the right split constraint (12d) becomes tight (if $e_1 = B_{s_1}$) at s_1. In either case, since s_2 is a descendant of s_1, the leaf $f(e_2)$ must be contained in the left leaves of split s_1 (if $e = A_{s_1}$) or the right leaves of split s_1 (if $e_1 = B_{s_1}$). Thus, when leaf $f(e_2)$ is checked, the event e_2 cannot occur, because q_{s_1} in Algorithm 1 will be zero (implying that $q_{A,B} = 0$), and so s^* cannot be equal to s_2 because s_1 is a shallower split that attains the minimum of $q_{A,B} = 0$. □

To establish that $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ is feasible for the SplitMIO dual subproblem (13), we will first show that the $\alpha_{t,s}$ variables are nonnegative.

Fix $s \in \mathbf{splits}(t)$. If $A_s \notin \mathcal{E}$, then $\alpha_{t,s} = 0$, and constraint (13c) is satisfied. If $A_s \in \mathcal{E}$, then consider the split $\tilde{s} = \arg\min_{s'} [\{d(s') \mid s' \in \mathbf{LS}(f(A_s)),\ d(s') < d,\ A_{s'} \in \mathcal{E}\} \cup \{d(s') \mid s' \in \mathbf{RS}(f(A_s)),\ d(s') < d,\ B_{s'} \in \mathcal{E}\}]$, where we recall that $d = d(s)$ is the depth of split s. In words, $\tilde{s}$ is the shallowest split (i.e., closest to the root) along the path of splits from the root node to split s such that either an $A_{\tilde{s}}$ event occurs or a $B_{\tilde{s}}$ event occurs for split $\tilde{s}$. There are three possible cases that can occur here, which we now handle.

Case 1: $\tilde{s} \in \mathbf{LS}(f(A_s))$. In this case, $A_{\tilde{s}} \in \mathcal{E}$, and we have

$$\begin{aligned}
\alpha_{t,s} &= r_{t,f(A_s)} - \left[\sum_{s' \in \mathbf{LS}(f(A_s)):\ d(s') < d,\ A_{s'} \in \mathcal{E}} \alpha_{t,s'} + \sum_{s' \in \mathbf{RS}(f(A_s)):\ d(s') < d,\ B_{s'} \in \mathcal{E}} \beta_{t,s'} + \gamma_t \right] \\
&= r_{t,f(A_s)} - \left[\alpha_{t,\tilde{s}} + \sum_{s' \in \mathbf{LS}(f(A_s)):\ d(s') < d(\tilde{s}),\ A_{s'} \in \mathcal{E}} \alpha_{t,s'} + \sum_{s' \in \mathbf{RS}(f(A_s)):\ d(s') < d(\tilde{s}),\ B_{s'} \in \mathcal{E}} \beta_{t,s'} + \gamma_t \right] \\
&= r_{t,f(A_s)} - \left[\alpha_{t,\tilde{s}} + \sum_{s' \in \mathbf{LS}(f(A_{\tilde{s}})):\ d(s') < d(\tilde{s}),\ A_{s'} \in \mathcal{E}} \alpha_{t,s'} + \sum_{s' \in \mathbf{RS}(f(A_{\tilde{s}})):\ d(s') < d(\tilde{s}),\ B_{s'} \in \mathcal{E}} \beta_{t,s'} + \gamma_t \right] = r_{t,f(A_s)} - r_{t,f(A_{\tilde{s}})} \geq 0,
\end{aligned}$$

where the first step follows by the definition of $\alpha_{t,s}$ in Algorithm 2; the second step follows by the definition of $\alpha_{t,\tilde{s}}$ as the deepest split for which an A or B event occurs that is at a depth lower than s; the third step by the fact that the left splits and right splits of $f(A_{\tilde{s}})$ at a depth below $d(\tilde{s})$ are the same as the left and right splits of $f(A_s)$ at a depth below $d(\tilde{s})$; and the fourth step follows from the definition of $\alpha_{t,\tilde{s}}$ in Algorithm 2. The inequality follows by Lemma EC.1.

Case 2: $\tilde{s} \in \mathbf{RS}(f(A_s))$. In this case, $B_{\tilde{s}} \in \mathcal{E}$, and analogously to Case 1, we have:

$$\begin{aligned}
\alpha_{t,s} &= r_{t,f(A_s)} - \left[\sum_{s' \in \mathbf{LS}(f(A_s)):\ d(s') < d,\ A_{s'} \in \mathcal{E}} \alpha_{t,s'} + \sum_{s' \in \mathbf{RS}(f(A_s)):\ d(s') < d, B_{s'} \in \mathcal{E}} \beta_{t,s'} + \gamma_t \right] \\
&= r_{t,f(A_s)} - \left[\beta_{t,\tilde{s}} + \sum_{s' \in \mathbf{LS}(f(A_s)):\ d(s') < d(\tilde{s}),\ A_{s'} \in \mathcal{E}} \alpha_{t,s'} + \sum_{s' \in \mathbf{RS}(f(A_s)):\ d(s') < d(\tilde{s}),\ B_{s'} \in \mathcal{E}} \beta_{t,s'} + \gamma_t \right] \\
&= r_{t,f(A_s)} - \left[\beta_{t,\tilde{s}} + \sum_{s' \in \mathbf{LS}(f(B_{\tilde{s}})):\ d(s') < d(\tilde{s}),\ A_{s'} \in \mathcal{E}} \alpha_{t,s'} + \sum_{s' \in \mathbf{RS}(f(B_{\tilde{s}})):\ d(s') < d(\tilde{s}),\ B_{s'} \in \mathcal{E}} \beta_{t,s'} + \gamma_t \right] = r_{t,f(A_s)} - r_{t,f(B_{\tilde{s}})} \geq 0.
\end{aligned}$$

Case 3: $\tilde{s}$ is undefined because the underlying sets are empty. In this case, $\alpha_{t,s} = r_{t,f(A_s)} - \gamma_t$, and we have $\alpha_{t,s} = r_{t,f(A_s)} - \gamma_t = r_{t,f(A_s)} - r_{t,f(C)} \geq 0$, where the inequality follows because $f(C)$ is the last leaf to be checked before Algorithm 1 terminates, and thus it must be that $r_{t,f(A_s)} \geq r_{t,f(C)}$. This establishes that $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ satisfy constraint (13c). Constraint (13d) can be shown in an almost identical fashion; for brevity, we omit the steps. We thus only need to verify constraint (13b). Let $\ell \in \mathbf{leaves}(t)$. Here, there are four mutually exclusive and collectively exhaustive cases to consider.

Case 1: $r_{t,\ell} \leq r_{t,f(C)}$. In this case we have $\sum_{s \in \mathbf{LS}(\ell)} \alpha_{t,s} + \sum_{s \in \mathbf{RS}(\ell)} \beta_{t,s} + \gamma_t \geq \gamma_t = r_{t,f(C)} \geq r_{t,\ell}$.

Case 2: $r_{t,\ell} > r_{t,f(C)}$ and $\ell = f(A_s)$ for some $s \in \mathbf{splits}(t)$. In this case, we have

$$\sum_{s' \in \mathbf{LS}(\ell)} \alpha_{t,s'} + \sum_{s' \in \mathbf{RS}(\ell)} \beta_{t,s'} + \gamma_t \geq \alpha_{t,s} + \sum_{\substack{s' \in \mathbf{LS}(\ell): \\ d(s') < d(s),\ A_{s'} \in \mathcal{E}}} \alpha_{t,s'} + \sum_{\substack{s' \in \mathbf{RS}(\ell): \\ d(s') < d(s),\ B_{s'} \in \mathcal{E}}} \beta_{t,s'} + \gamma_t = r_{t,f(A_s)} = r_{t,\ell},$$

where the first step follows by the nonnegativity of $\alpha_{t,s'}$ and $\beta_{t,s'}$ for all s', and the second step by the definition of $\alpha_{t,s}$ in Algorithm 2.

Case 3: $r_{t,\ell} > r_{t,f(C)}$ and $\ell = f(B_s)$ for some $s \in \mathbf{splits}(t)$. By similar logic as case 2, we have

$$\sum_{s' \in \mathbf{LS}(\ell)} \alpha_{t,s'} + \sum_{s' \in \mathbf{RS}(\ell)} \beta_{t,s'} + \gamma_t \geq \beta_{t,s} + \sum_{\substack{s' \in \mathbf{LS}(\ell): \\ d(s') < d(s),\ A_{s'} \in \mathcal{E}}} \alpha_{t,s'} + \sum_{\substack{s' \in \mathbf{RS}(\ell): \\ d(s') < d(s),\ B_{s'} \in \mathcal{E}}} \beta_{t,s'} + \gamma_t = r_{t,f(B_s)} = r_{t,\ell}.$$

Case 4: $r_{t,\ell} > r_{t,f(C)}$ and ℓ is not equal to $f(A_s)$ or $f(B_s)$ for any split s. In this case, when leaf ℓ is checked by Algorithm 1, the algorithm reaches line 17 where s^* is determined and e is set to either A_{s^*} or B_{s^*}, and it turns out that e is already in $\mathcal{E}$. If $e = A_{s^*}$, then this means that leaf $f(A_{s^*})$ was checked before leaf ℓ, and that $r_{t,\ell} \leq r_{t,f(A_{s^*})}$. We thus have

$$\begin{aligned}\sum_{s \in \mathbf{LS}(\ell)} \alpha_{t,s} + \sum_{s \in \mathbf{RS}(\ell)} \beta_{t,s} + \gamma_t &\geq \alpha_{t,s^*} + \sum_{s \in \mathbf{LS}(\ell):\ d(s) < d(s^*),\ A_s \in \mathcal{E}} \alpha_{t,s} + \sum_{s \in \mathbf{RS}(\ell):\ d(s) < d(s^*),\ B_s \in \mathcal{E}} \beta_{t,s} + \gamma_t \\ &= \alpha_{t,s^*} + \sum_{\substack{s \in \mathbf{LS}(f(A_{s^*})): \\ d(s) < d(s^*),\ A_s \in \mathcal{E}}} \alpha_{t,s} + \sum_{\substack{s \in \mathbf{RS}(f(A_{s^*})): \\ d(s) < d(s^*),\ B_s \in \mathcal{E}}} \beta_{t,s} + \gamma_t = r_{t,f(A_{s^*})} \geq r_{t,\ell},\end{aligned}$$

where the first equality follows because ℓ and $f(A_{s^*})$, by virtue of being to the left of s^*, share the same left and right splits at depths lower than $d(s^*)$. Similarly, if $e = B_{s^*}$, then the leaf $f(B_{s^*})$ was checked before ℓ, which means that $r_{t,\ell} \leq r_{t,f(B_{s^*})}$; in this case, we have

$$\begin{aligned}\sum_{s \in \mathbf{LS}(\ell)} \alpha_{t,s} + \sum_{s \in \mathbf{RS}(\ell)} \beta_{t,s} + \gamma_t &\geq \beta_{t,s^*} + \sum_{s \in \mathbf{LS}(\ell):\ d(s) < d(s^*),\ A_s \in \mathcal{E}} \alpha_{t,s} + \sum_{s \in \mathbf{RS}(\ell):\ d(s) < d(s^*),\ B_s \in \mathcal{E}} \beta_{t,s} + \gamma_t \\ &\geq \beta_{t,s^*} + \sum_{\substack{s \in \mathbf{LS}(f(B_{s^*})): \\ d(s) < d(s^*),\ A_s \in \mathcal{E}}} \alpha_{t,s} + \sum_{\substack{s \in \mathbf{RS}(f(B_{s^*})): \\ d(s) < d(s^*),\ B_s \in \mathcal{E}}} \beta_{t,s} + \gamma_t = r_{t,f(B_{s^*})} \geq r_{t,\ell}.\end{aligned}$$

We have thus shown that $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ is a feasible solution to the SplitMIO dual subproblem (13).

Proof of part (b) (extreme point): To prove this, we will use the equivalence between extreme points and basic feasible solutions, and show that $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ is a basic feasible solution of problem (13).

Define the sets $L_A = \{\ell \in \mathbf{leaves}(t) \mid \ell = f(A_s) \text{ for some } s \in \mathbf{splits}(t)\}$ and $L_B = \{\ell \in \mathbf{leaves}(t) \mid \ell = f(B_s) \text{ for some } s \in \mathbf{splits}(t)\}$. Consider the following system of equations:

$$\sum_{s \in \mathbf{LS}(\ell)} \alpha_{t,s} + \sum_{s \in \mathbf{RS}(\ell)} \beta_{t,s} + \gamma_t = r_{t,\ell}, \quad \forall \ell \in L_A, \tag{EC.18}$$

$$\sum_{s \in \mathbf{LS}(\ell)} \alpha_{t,s} + \sum_{s \in \mathbf{RS}(\ell)} \beta_{t,s} + \gamma_t = r_{t,\ell}, \quad \forall \ell \in L_B, \tag{EC.19}$$

$$\sum_{s \in \mathbf{LS}(f(C))} \alpha_{t,s} + \sum_{s \in \mathbf{RS}(f(C))} \beta_{t,s} + \gamma_t = r_{t,f(C)}, \tag{EC.20}$$

$$\alpha_{t,s} = 0, \quad \forall\ s \text{ such that } A_s \notin \mathcal{E}, \tag{EC.21}$$

$$\beta_{t,s} = 0, \quad \forall\ s \text{ such that } B_s \notin \mathcal{E}. \tag{EC.22}$$

Observe that each equation corresponds to a constraint from problem (13) made to hold at equality. In addition, we note that there are $|L_A| + |L_B| + 1 + (|\mathbf{splits}(t)| - |L_A|) + (|\mathbf{splits}(t)| - |L_B|) = 2|\mathbf{splits}(t)| + 1$ equations, which is exactly the number of variables. We will show that the unique solution implied by this system of equations is exactly the solution $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ that is produced by Algorithm 2.

In order to establish this, we first establish a couple of useful results.

LEMMA EC.2. *Suppose that $e \in \mathcal{E}$, $\ell = f(e)$ and $e = A_s$ or $e = B_s$ for some $s \in \mathbf{splits}(t)$. Then: (a) $A_{s'} \notin \mathcal{E}$ for all $s' \in \mathbf{LS}(\ell)$ such that $d(s') > d(s)$; and (b) $B_{s'} \notin \mathcal{E}$ for all $s' \in \mathbf{RS}(\ell)$ such that $d(s') > d(s)$.*

Proof of Lemma EC.2: We will prove this by contradiction. Without loss of generality, let us suppose that there exists an $A_{s'}$ event in $\mathcal{E}$ where $s' \in \mathbf{LS}(\ell)$ and $d(s') > d(s)$. (The case where there exists an $B_{s'}$ event in $\mathcal{E}$ where $s' \in \mathbf{RS}(\ell)$ and $d(s') > d(s)$ can be shown almost identically.)

Since $A_{s'} \in \mathcal{E}$, consider the leaf $\ell' = f(A_{s'})$. There are now two possibilities for when Algorithm 1 checks leaf ℓ':

1. **Case 1**: Leaf ℓ' is checked after leaf ℓ. In this case, in the iteration of Algorithm 1 corresponding to leaf ℓ', it will be the case that $q_s = 0$ because the left constraint (12c) at split s (if $e = A_s$) or the right constraint (12d) at split s (if $e = B_s$) became tight when leaf ℓ was checked. As a result, $q_{A,B} = 0$ in the iteration for leaf ℓ'. This implies that s' cannot be the lowest depth split that attains the minimum q_s value of $q_{A,B}$, because $q_s = 0$, and s has a depth lower than s', which contradicts the fact that the $A_{s'}$ event occurred.

2. **Case 2**: Leaf ℓ' is checked before leaf ℓ. In this case, consider the value of q_s when leaf ℓ is checked in Algorithm 1.

If $q_s > 0$, then there is immediately a contradiction, because $q_{s'} = 0$ when leaf ℓ is checked (this is true because the left split constraint (12c) at s' became tight after leaf ℓ' was checked), and thus it is impossible that $s^* = s$. If $q_s = 0$, then this implies that $x_{v(t,s)} = 0$. This would imply that $q_s = 0$ when leaf ℓ' was checked, which would imply that s^* cannot be s' when leaf ℓ' is checked because s is at a lower depth than s'.

Thus, in either case, we arrive at a contradiction, which completes the proof. □

LEMMA EC.3. *Suppose $\ell = f(C)$. Then: (a) $A_{s'} \notin \mathcal{E}$ for all $s' \in \mathbf{LS}(\ell)$; and (b) $B_{s'} \notin \mathcal{E}$ for all $s' \in \mathbf{RS}(\ell)$.*

Proof of Lemma EC.3: We proceed by contradiction. Suppose that A_s occurs for some $s \in \mathbf{LS}(\ell)$ or that B_s occurs for some $s \in \mathbf{LS}(\ell)$; in the former case, let $e = A_s$, and in the latter case, let $e = B_s$. Let $\ell' = f(e)$. Then ℓ' must be checked before ℓ by Algorithm 1, since the algorithm always terminates after a C event occurs. Consider what happens when Algorithm 1 checks leaf ℓ:

1. **Case 1**: $q_C > 0$. This is impossible, because if e occurs, then q_s when leaf ℓ is checked would have to be 0, which would imply that $q_{A,B} < q_C$ and that a C event could not have occurred when ℓ was checked.

2. **Case 2**: $q_C = 0$. This is also impossible, because it implies that the unit sum constraint (12b) was satisfied at an earlier iteration, which would have triggered the C event at a leaf that was checked before ℓ.

We thus have that $A_{s'}$ does not occur for any $s' \in \mathbf{LS}(\ell)$ and $B_{s'}$ does not occur for any $s' \in \mathbf{RS}(\ell)$. □

With these two lemmas in hand, we now return to the proof of Theorem B.5 (b). Observe now that by using Lemmas EC.2 and EC.3 and using equations (EC.23) and (EC.24), the system of equations (EC.18)-(EC.22) is equivalent to

$$\alpha_{t,s} + \sum_{s' \in \mathbf{LS}(f(A_s)):\, d(s') < d(s),\, A_{s'} \in \mathcal{E}} \alpha_{t,s'} + \sum_{s' \in \mathbf{RS}(\ell):\, d(s') < d(s),\, B_{s'} \in \mathcal{E}} \beta_{t,s'} + \gamma_t = r_{t,f(A_s)}, \quad \forall s \text{ such that } A_s \in \mathcal{E},$$

$$\beta_{t,s} + \sum_{s' \in \mathbf{LS}(\ell):\, d(s') < d(s),\, A_{s'} \in \mathcal{E}} \alpha_{t,s'} + \sum_{s' \in \mathbf{RS}(\ell):\, d(s') < d(s),\, B_{s'} \in \mathcal{E}} \beta_{t,s'} + \gamma_t = r_{t,f(B_s)}, \quad \forall s \text{ such that } B_s \in \mathcal{E},$$

$$\gamma_t = r_{t,f(C)},$$

$$\alpha_{t,s} = 0, \quad \forall\ s \text{ such that } A_s \notin \mathcal{E}, \tag{EC.23}$$

$$\beta_{t,s} = 0, \quad \forall\ s \text{ such that } B_s \notin \mathcal{E}. \tag{EC.24}$$

Thus the solution implied by this system of equations is exactly the solution produced by Algorithm 2, establishing that $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ is a basic feasible solution of problem (13), and thus an extreme point. □

B.6. Proof of Theorem 3 (SplitMIO primal and dual are optimal)

To prove that the $\mathbf{y}_t$ and $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ produced by Algorithms 1 and 2 are optimal for their respective problems, we show that they satisfy complementary slackness for problems (12) and (13):

$$\alpha_{t,s} \cdot \left(x_{v(t,s)} - \sum_{\ell \in \mathbf{left}(s)} y_{t,\ell} \right) = 0, \quad \forall\, s \in \mathbf{splits}(t), \tag{EC.25}$$

$$\beta_{t,s} \cdot \left(1 - x_{v(t,s)} - \sum_{\ell \in \mathbf{right}(s)} y_{t,\ell} \right) = 0, \quad \forall\, s \in \mathbf{splits}(t), \tag{EC.26}$$

$$y_{t,\ell} \cdot \left(\sum_{s \in \mathbf{LS}(\ell)} \alpha_{t,s} + \sum_{s \in \mathbf{RS}(\ell)} \beta_{t,s} + \gamma_t - r_{t,\ell} \right) = 0, \quad \forall\, \ell \in \mathbf{leaves}(t). \tag{EC.27}$$

Condition (EC.25): If $\alpha_{t,s} = 0$, then the condition is trivially satisfied. If $\alpha_{t,s} > 0$, then this implies that $A_s \in \mathcal{E}$. This means that the left split constraint (12c) at s became tight after leaves $f(A_s)$ was checked, which implies that $\sum_{\ell \in \mathbf{left}(s)} y_{t,\ell} = x_{v(t,s)}$ or equivalently, that $x_{v(t,s)} - \sum_{\ell \in \mathbf{left}(s)} y_{t,\ell} = 0$, which again implies that the condition is satisfied.

Condition (EC.26): This follows along similar logic to condition (EC.25), only that we use the fact that $\beta_{t,s} > 0$ implies that a B_s event occurred and that the right split constraint (12d) at s became tight.

Condition (EC.27): If $y_{t,\ell} = 0$, then the condition is trivially satisfied. If $y_{t,\ell} > 0$, then either $\ell = f(C)$, or $\ell = f(A_s)$ for some split $s \in \mathbf{LS}(\ell)$, or $\ell = f(B_s)$ for some split $s \in \mathbf{RS}(\ell)$. In any of these three cases, as shown in the proof of part (b) of Theorem 2, the dual constraint (13b) holds with equality for any such leaf ℓ. Thus, we have that $\sum_{s \in \mathbf{LS}(\ell)} \alpha_{t,s} + \sum_{s \in \mathbf{RS}(\ell)} \beta_{t,s} + \gamma_t - r_{t,\ell} = 0$, and the condition is again satisfied.

As complementary slackness holds, $\mathbf{y}_t$ is feasible for the primal problem (12) (Theorem B.4), and $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ is feasible for the dual problem (13) (Theorem 2). Thus they are optimal for their respective problems. □

B.7. Proof of Theorem 4 (SplitMIO primal and dual are closed form solvable for binary x)

Proof of part (a): Observe that by construction, $\mathbf{y}_t$ automatically satisfies the unit sum constraint (12b) and the nonnegativity constraint (12e). We thus need to verify constraints (12c) and (12d). For constraint (12c), observe that for any split $s \notin \mathbf{LS}(\ell^*)$, it must be that $\ell^* \notin \mathbf{left}(s)$. Thus, we will have $\sum_{\ell \in \mathbf{left}(s)} y_{t,\ell} = 0$, which means that constraint (12c) is automatically satisfied, because the right hand side $x_{v(t,s)}$ is always at least 0. On the other hand, for any split $s \in \mathbf{LS}(\ell^*)$, we will have that $x_{v(t,s)} = 1$, and that $\sum_{\ell \in \mathbf{left}(s)} y_{t,\ell} = \sum_{\ell \in \mathbf{left}(s): \ell \neq \ell^*} y_{t,\ell} + y_{t,\ell^*} = 1$, which implies that constraint (12c) is satisfied. Similar reasoning can be used to establish that constraint (12d) holds. This establishes that $\mathbf{y}_t$ is indeed a feasible solution of problem (12).

Proof of part (b): By construction, $\alpha_{t,s} \geq 0$ and $\beta_{t,s} \geq 0$ for all $s \in \mathbf{splits}(t)$, so constraints (13c) and (13d) are satisfied. To verify constraint (13b), fix a leaf $\ell \in \mathbf{leaves}(t)$. If $\ell \neq \ell^*$, then either $\ell \in \mathbf{left}(s')$ for some $s' \in \mathbf{RS}(\ell^*)$ or $\ell \in \mathbf{right}(s')$ for some $s' \in \mathbf{LS}(\ell^*)$. If $\ell \in \mathbf{left}(s')$ for some $s' \in \mathbf{RS}(\ell^*)$, then

$$\sum_{s: \ell \in \mathbf{left}(s)} \alpha_{t,s} + \sum_{s: \ell \in \mathbf{right}(s)} \beta_{t,s} + \gamma_t \geq \alpha_{t,s'} + \gamma_t \geq \max_{\ell' \in \mathbf{left}(s')} r_{t,\ell'} - r_{t,\ell^*} + r_{t,\ell^*} \geq r_{t,\ell}$$

where the first inequality follows because $\ell \in \mathbf{left}(s')$ and the fact that all $\alpha_{t,s}$ and $\beta_{t,s}$ variables are nonnegative; the second follows by how the dual solution is defined in the statement of the theorem; and the third by the definition of the maximum. Similarly, if $\ell \in \mathbf{right}(s')$ for some $s' \in \mathbf{LS}(\ell^*)$, then we have

$$\sum_{s:\ell\in\mathbf{left}(s)} \alpha_{t,s} + \sum_{s:\ell\in\mathbf{right}(s)} \beta_{t,s} + \gamma_t \geq \beta_{t,s'} + \gamma_t \geq \max_{\ell'\in\mathbf{right}(s')} r_{t,\ell'} - r_{t,\ell^*} + r_{t,\ell^*} \geq r_\ell - r_{t,\ell^*} + r_{t,\ell^*} = r_{t,\ell}.$$

Lastly, if $\ell = \ell^*$, then we automatically have $\sum_{s:\ell^*\in\mathbf{left}(s)} \alpha_{t,s} + \sum_{s:\ell^*\in\mathbf{right}(s)} \beta_{t,s} + \gamma_t \geq \gamma_t = r_{t,\ell^*}$. Thus, we have established that constraint (13b) is satisfied for all leaves ℓ, and thus $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ as defined in the statement of the theorem is a feasible solution of the dual (13).

Proof of part (c): To establish that the two solutions are optimal, by weak duality it is sufficient to show that the two solutions attain the same objective values in their respective problems. For the primal solution $\mathbf{y}_t$, it is immediately clear that its objective is r_{t,ℓ^*}. For the dual solution $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$, we have

$$\sum_{s\in\mathbf{splits}(t)} \alpha_{t,s} x_{v(t,s)} + \sum_{s\in\mathbf{splits}(t)} \beta_{t,s}(1 - x_{v(t,s)}) + \gamma_t = \sum_{s\in\mathbf{RS}(\ell^*)} \alpha_{t,s} x_{v(t,s)} + \sum_{s\in\mathbf{LS}(\ell^*)} \beta_{t,s}(1 - x_{v(t,s)}) + \gamma_t = \gamma_t = r_{t,\ell^*},$$

where the first step follows because $\alpha_{t,s} = 0$ for $s \notin \mathbf{RS}(\ell^*)$ and $\beta_{t,s} = 0$ for $s \notin \mathbf{LS}(\ell^*)$; the second step follows by the fact that $x_{v(t,s)} = 0$ for $s \in \mathbf{RS}(\ell^*)$ and $x_{v(t,s)} = 1$ for $s \in \mathbf{LS}(\ell^*)$; and the final step follows just by the definition of γ_t. This establishes that $\mathbf{y}_t$ and $(\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t, \gamma_t)$ are optimal for their respective problems. □

B.8. Proof of Theorem 5 (ProductMIO primal and dual are closed form solvable for binary x)

The proof follows a similar argument to the proof of Theorem 4. Due to space limitations, we omit the proof.

Appendix C: Additional Numerical Results

Our experiments were implemented in the Julia programming language, version 1.8.4 (Bezanson et al. 2017) and executed on Amazon Elastic Compute Cloud (EC2) using a single instance of type `m6a.48xlarge` (AMD EPYC 7R13 processor with 2.95GHz clock speed, 192 virtual CPUs and 768 GB memory). All LO and MIO formulations were solved using Gurobi version 10.0.1 and modeled using the `JuMP` package (Lubin and Dunning 2015), with a maximum of four cores per formulation.

C.1. Experiment: Tractability

In this experiment, we seek to understand the tractability of SplitMIO and ProductMIO when they are solved as integer problems (i.e., not as relaxations). For a given instance and a given formulation $\mathcal{M}$, we solve the integer version of formulation $\mathcal{M}$. Due to the large size of some of the problem instances, we impose a computation time limit of 1 hour for each formulation. We record $T_{\mathcal{M}}$, the computation time required for formulation $\mathcal{M}$, and we record $O_{\mathcal{M}}$ which is the final optimality gap, and is defined as $O_{\mathcal{M}} = 100\% \times (Z_{\mathrm{UB},\mathcal{M}} - Z_{\mathrm{LB},\mathcal{M}})/Z_{\mathrm{UB},\mathcal{M}}$, where $Z_{\mathrm{UB},\mathcal{M}}$ and $Z_{\mathrm{LB},\mathcal{M}}$ are the best upper and lower bounds, respectively, obtained at the termination of formulation $\mathcal{M}$ for the instance. We test all of the T1, T2 and T3 instances with $n = 100$, $|F| \in \{50, 100, 200, 500\}$ and $|\mathbf{leaves}(t)| \in \{8, 16, 32, 64\}$. Table EC.3 displays the average computation time and average optimality gap of each formulation for each combination of n, $|F|$ and $|\mathbf{leaves}(t)|$. For ease of notation in the table, we abbreviate SplitMIO as S-MIO and ProductMIO as P-MIO, and write $N_L \equiv |\mathbf{leaves}(t)|$.

From this table, we can see that for the smaller instances, SplitMIO requires more time to solve than ProductMIO. For larger instances, where the computation time limit is exhausted, the average gap obtained by ProductMIO tends to be lower than that of SplitMIO.

Type	$\|F\|$	N_{L}	$O_{\text{S-MIO}}$	$O_{\text{P-MIO}}$	$T_{\text{S-MIO}}$	$T_{\text{P-MIO}}$	Type	$\|F\|$	N_{L}	$O_{\text{S-MIO}}$	$O_{\text{P-MIO}}$	$T_{\text{S-MIO}}$	$T_{\text{P-MIO}}$
T1	50	8	0.0	0.0	0.0	0.0	T2	200	8	0.0	0.0	0.6	0.6
T1	50	16	0.0	0.0	0.1	0.0	T2	200	16	0.0	0.0	747.7	673.6
T1	50	32	0.0	0.0	0.5	0.1	T2	200	32	9.8	9.7	3600.0	3600.0
T1	50	64	0.0	0.0	2.4	0.5	T2	200	64	17.1	16.3	3600.1	3600.0
T1	100	8	0.0	0.0	0.0	0.0	T2	500	8	0.0	0.0	171.6	167.9
T1	100	16	0.0	0.0	0.8	0.1	T2	500	16	14.1	13.8	3600.0	3600.0
T1	100	32	0.0	0.0	13.5	1.6	T2	500	32	23.4	23.0	3600.1	3600.1
T1	100	64	0.0	0.0	479.0	65.1	T2	500	64	28.7	28.1	3600.1	3600.1
T1	200	8	0.0	0.0	0.1	0.0	T3	50	8	0.0	0.0	0.0	0.0
T1	200	16	0.0	0.0	25.0	3.0	T3	50	16	0.0	0.0	0.1	0.1
T1	200	32	0.7	0.0	2330.2	421.5	T3	50	32	0.0	0.0	0.6	0.6
T1	200	64	9.8	5.3	3600.1	3600.0	T3	50	64	0.0	0.0	4.1	3.7
T1	500	8	0.0	0.0	4.6	1.4	T3	100	8	0.0	0.0	0.0	0.0
T1	500	16	5.0	1.2	3600.0	2951.7	T3	100	16	0.0	0.0	0.8	0.7
T1	500	32	15.9	11.8	3600.1	3600.0	T3	100	32	0.0	0.0	29.0	26.4
T1	500	64	20.9	17.4	3600.1	3600.1	T3	100	64	0.9	0.5	2600.8	2206.5
T2	50	8	0.0	0.0	0.0	0.0	T3	200	8	0.0	0.0	0.4	0.4
T2	50	16	0.0	0.0	0.1	0.1	T3	200	16	0.0	0.0	76.8	79.1
T2	50	32	0.0	0.0	0.5	0.5	T3	200	32	5.3	5.1	3600.0	3600.0
T2	50	64	0.0	0.0	24.0	24.6	T3	200	64	13.5	12.5	3600.1	3600.0
T2	100	8	0.0	0.0	0.0	0.0	T3	500	8	0.0	0.0	75.9	72.9
T2	100	16	0.0	0.0	1.5	1.4	T3	500	16	8.8	8.8	3600.0	3600.0
T2	100	32	0.0	0.0	245.2	234.9	T3	500	32	21.0	20.3	3600.1	3600.0
T2	100	64	4.8	4.4	3600.0	3600.0	T3	500	64	28.9	27.9	3600.2	3600.1

Table EC.3 Comparison of final optimality gaps and computation times for SplitMIO and ProductMIO.

C.2. Comparison to Heuristic Approaches

In this experiment, we compare the performance of the two formulations to three different heuristic approaches:

1. LS: A local search heuristic, which starts from the empty assortment, and in each iteration moves to the neighboring assortment which improves the expected revenue the most. The neighborhood of assortments consists of those assortments obtained by adding a new product to the current assortment, or removing one of the existing products from the assortment. The heuristic terminates when there is no assortment in the neighborhood of the current one that provides an improvement.

2. LS10: This heuristic involves running LS from ten randomly chosen starting assortments. Each assortment is chosen uniformly at random from the set of 2^n possible assortments. After the ten repetitions, the assortment with the best expected revenue is retained.

3. ROA: This heuristic involves finding the optimal revenue ordered assortment. More formally, we define $S_k = \{i_1, \ldots, i_k\}$, where $i_1, \ldots, i_n$ corresponds to an ordering of the products so that $r_{i_1} \geq r_{i_2} \geq \cdots \geq r_{i_n}$, and we find $\arg\max_{S \in \{S_1, \ldots, S_n\}} R^{(F,\boldsymbol{\lambda})}(S)$.

Type	$\|F\|$	N_L	$G_{\text{S-MIO}}$	$G_{\text{P-MIO}}$	G_{LS}	G_{LS10}	G_{ROA}	Type	$\|F\|$	N_L	$G_{\text{S-MIO}}$	$G_{\text{P-MIO}}$	G_{LS}	G_{LS10}	G_{ROA}
T1	50	8	0.0	0.0	8.1	2.1	26.8	T2	200	8	0.0	0.0	3.8	0.8	23.1
T1	50	16	0.0	0.0	9.8	4.0	31.2	T2	200	16	0.0	0.0	5.5	2.8	24.2
T1	50	32	0.0	0.0	9.0	5.2	27.5	T2	200	32	0.2	0.2	7.0	4.5	24.7
T1	50	64	0.0	0.0	10.0	7.2	28.6	T2	200	64	1.0	0.4	8.2	5.5	23.5
T1	100	8	0.0	0.0	3.9	2.4	26.0	T2	500	8	0.0	0.0	2.0	0.5	15.6
T1	100	16	0.0	0.0	6.6	3.7	26.2	T2	500	16	0.1	0.1	3.8	1.5	16.3
T1	100	32	0.0	0.0	8.5	6.3	25.6	T2	500	32	0.5	0.4	4.9	1.8	15.1
T1	100	64	0.0	0.0	9.9	6.6	24.5	T2	500	64	1.1	0.9	4.5	1.9	14.3
T1	200	8	0.0	0.0	3.5	1.3	19.9	T3	50	8	0.0	0.0	13.2	3.8	33.1
T1	200	16	0.0	0.0	5.5	3.1	21.7	T3	50	16	0.0	0.0	14.4	5.2	34.9
T1	200	32	0.0	0.0	7.8	4.6	22.3	T3	50	32	0.0	0.0	12.6	5.0	33.7
T1	200	64	0.3	0.0	7.6	6.3	19.5	T3	50	64	0.0	0.0	13.9	8.1	33.0
T1	500	8	0.0	0.0	1.5	0.3	14.6	T3	100	8	0.0	0.0	8.0	1.9	30.2
T1	500	16	0.0	0.0	3.7	1.6	15.3	T3	100	16	0.0	0.0	10.0	3.1	32.6
T1	500	32	0.3	0.0	5.3	2.4	15.9	T3	100	32	0.0	0.0	10.4	3.8	32.0
T1	500	64	0.6	0.1	4.9	3.5	13.8	T3	100	64	0.0	0.0	10.0	4.5	31.0
T2	50	8	0.0	0.0	13.8	3.2	31.5	T3	200	8	0.0	0.0	4.0	1.1	25.8
T2	50	16	0.0	0.0	11.6	4.9	32.2	T3	200	16	0.0	0.0	6.5	2.5	26.5
T2	50	32	0.0	0.0	10.0	5.8	31.1	T3	200	32	0.1	0.1	7.6	3.3	26.9
T2	50	64	0.0	0.0	11.6	7.0	30.4	T3	200	64	0.3	0.1	9.2	4.1	24.1
T2	100	8	0.0	0.0	5.5	1.9	28.1	T3	500	8	0.0	0.0	2.6	0.4	15.8
T2	100	16	0.0	0.0	8.2	4.0	30.8	T3	500	16	0.0	0.0	4.3	1.1	16.5
T2	100	32	0.0	0.0	8.9	4.8	31.3	T3	500	32	0.5	0.5	5.3	1.3	16.0
T2	100	64	0.0	0.0	11.6	6.6	27.1	T3	500	64	0.9	0.4	5.5	1.3	16.3

Table EC.4 Comparison of SplitMIO and ProductMIO against the heuristics LS, LS10 and ROA.

We compare these heuristics against the best integer solution obtained by each of our two MIO formulations, leading to a total of five methods for each instance. We measure the performance of the solution corresponding to approach $\mathcal{M}$ using the metric $G_{\mathcal{M}}$, which is defined as $G_{\mathcal{M}} = 100\% \times (Z' - Z_{\mathcal{M}})/Z'$, where Z' is the highest lower bound (i.e., integer solution) obtained from among the two MIO formulations and the three heuristics, and $Z_{\mathcal{M}}$ is the objective value of the solution returned by approach $\mathcal{M}$.

Table EC.4 shows the performance of the five approaches – SplitMIO, ProductMIO, LS, LS10 and ROA – for each family of instances. The gaps are averaged over the twenty instances for each combination of instance type, $|F|$ and $|\mathbf{leaves}(t)|$. Again, for ease of notation in the table, we abbreviate SplitMIO as S-MIO and ProductMIO as P-MIO, and write $N_L \equiv |\mathbf{leaves}(t)|$. We can see from this table that in general, for the small instances, the solutions obtained by the MIO formulations are either the best or close to the best out of the five approaches, while the solutions produced by the heuristic approaches are quite suboptimal. For cases where the gap is zero for the MIO formulations, the gap of LS ranges from 1.5% to 13.9%; the LS10 heuristic improves on this, due to its use of restarting and randomization, but still does not perform as well as the MIO solutions (gaps ranging from 0.3 to 8.1%). For the larger instances, where the gap of the MIO solutions is larger, LS and LS10 still tend to perform worse. Across all of the instances, ROA achieves much higher gaps than all of the other approaches (ranging from 13.8 to 34.9%). Overall, these results suggest that *the very general structure of the decision forest model poses significant difficulty to standard heuristic approaches*, and highlight the value of using exact approaches over inexact/heuristic approaches to the assortment optimization problem.