Opaque Selling of Multiple Substitutable Products with Finite Inventories

Qian Liu,[1] **Yongbo Xiao,**[2*] **Dan Zhang**[3]

[1] Department of Industrial Engineering & Decision Analytics, The Hong Kong University of Science and Technology, Hong Kong

[2]Research Center for Contemporary Management, Key Research Institute of Humanities and Social Sciences at Universities, School of Economics and Management, Tsinghua University, Beijing 100084, China

[3]Leeds School of Business, University of Colorado Boulder, CO 80309, USA

Email: *qianliu@ust.hk, xiaoyb@sem.tsinghua.edu.cn*(∗)*, dan.zhang@colorado.edu*

(Version: November 17, 2021)

Abstract

Opaque selling, in which a seller offers opaque goods (OGs), in addition to physical goods, has been shown to be an effective strategy to segment a market and improve the seller's profit. This paper studies opaque selling with stochastic demand and fixed initial inventories of multiple products, where the seller dynamically controls the product offers and determines the product assignment to fulfill the demand for OGs over time. The problem is formulated as a stochastic dynamic program. Due to the curse of dimensionality, we study the fluid control problem that gives a time-based fluid policy and a stationary probabilistic fulfillment strategy. We show that the fluid policy is asymptotically optimal when the arrival rates and initial inventory level are scaled up linearly. Furthermore, we propose a decomposition heuristic based on the corresponding fluid solution. The decomposition heuristic is shown to provide a tighter upper bound than the fluid control problem. Numerical study on a set of test instances illustrates the performance and efficacy of opaque selling.

Keywords: Opaque Selling; Customer Choice Behavior; Asymptotic Optimality; Approximate Dynamic Programming.

1 Introduction

The rapid development of online selling provides new opportunities for sellers to effectively manage demand and improve revenue. One example is opaque selling under which sellers strategically withhold product information by keeping some characteristics of their products hidden until after purchase (Anderson and Celik, 2019). This selling strategy allows sellers to sell physical goods to brand-loyal customers at higher prices, and to sell "opaque goods" (OGs) to non-loyal customers at discounted prices (Anderson and Xie, 2014). Opaque selling has been widely observed in travel and retail industries. Well-known examples include Hotwire and Priceline which intentionally conceal certain attributes of airline tickets and hotel rooms (e.g., the identity of the service provider) from customers until after purchase. Amazon sells LEGO mini-figure collection series in a "mystery bag" which contains one random mini-figure.

A key driver of introducing OGs is heterogeneous customer preferences for physical goods. For customers who do not have a strong preference for a specific physical good, an OG is attractive because it is usually offered at a discount. By contrast, for customers who have a strong preference for a specific product, an OG is not attractive due to the fulfillment uncertainty. Therefore, opaque selling, if designed properly, has the potential to expand the market without cannibalizing sales of physical goods significantly. By adopting opaque selling, the seller can better match supply with demand and enhance revenue. Specifically, when selling multiple substituable products with fixed inventories in a finite horizon, the seller can dynamically decide which OGs to offer together with some physical goods. Once an OG is purchased, the seller determines which physical good is used to fulfill the demand, depending on the remaining inventories of each physical good and the remaining selling time. For example, the seller may use, as much as possible, a slow-moving physical good to satisfy customers who buy an OG.

With its growing popularity in practice, opaque selling has gained increasing attention in academic research. Much of the literature on opaque selling assumes that there are only two physical goods and a single OG. This is in stark contrast to the practice of multiple OGs in travel and online retail industries. For example, Eurowings, a German low-cost airline, offers the so-called "variable opaque products", for which customers can customize the set of flights for an opaque fare (Post and Spann, 2012). For example, after a customer who plans her summer vacation selects an origin city (e.g., Cologne) and a category (e.g., Sun and Beach) on *www.eurowings.com*, a total of 24 destinations are shown for her to choose from. The customer can tick off the destinations that she does not like by paying a surcharge of 5 euros per destination. In the "grab-bag" program offered by *swimoutlet.com*, a wide variety of grab bags are displayed to online customers. Different grab bags contain a set of swimsuits (typically two to four) with different prints and styles. A buyer chooses the size and the grab bag to purchase and the seller determines the print and style to fulfill the demand.

To illustrate the benefit of opaque selling for the seller, let us consider a stylized example. Three physical goods A, B, and C (e.g., shirts with different colors, or flights with different schedules) with an initial inventory of one unit each are sold over two periods. There are four types of customers whose valuations for the physical goods are denoted by $v_j = (v_j^A, v_j^B, v_j^C)$, $j = 1, 2, 3, 4$. Specifically, $v_1 = (1, 0, 0)$, $v_2 = (0, 1, 0)$, $v_3 = (0, 0, 1)$, and $v_4 = (0.3, 0.3, 0)$. The selling price for each physical product is one, i.e., $r_A = r_B = r_C = 1$. The seller uses physical goods A and B to create an OG, denoted by product D, with a selling price $r_D = 0.3$. Customers are risk neutral and assume equal fulfillment probability for physical goods A and B when purchasing the OG. Then a type-1 customer prefers to purchase the OG over physical good A. This is because the expected surplus from purchasing the OG is $0.5 \times v_1^A + 0.5 \times v_1^B - r_D = 0.2$, which is higher than the zero expected surplus ($v_1^A - r_A = 0$) from purchasing physical good A. Similarly, a type-2 customer prefers the OG over physical good B. A type-3 customer would purchase physical good C only. A type-4 customer would purchases the OG only. There is exactly one customer arrival in each period and the probabilities of belonging to the four types are 0.1, 0.1, 0.3, and 0.5, respectively. We compare the following three control policies:[1]

- Policy 1: offer all the available physical goods but not the OG in both periods. Note that type AB customers are lost. If a physical good is consumed in the first period, demand for this good in the second period will be lost. The total expected revenue is 0.89.
- Policy 2: Offer all the available physical goods and the OG in both periods. Note that the physical goods A and B are not purchased in the first period. If the physical good A or B is sold in the first period, the OG cannot be offered in the second period. The expected total revenue is 0.853.
- Policy 3: Offer all the physical goods in the first period, and all the available physical goods and the OG (if both physical goods A and B are available) in the second period. It can be verified that the expected total revenue is 0.898.

When opaque selling is not adopted, offering all the available goods is the optimal control policy. Interestingly, offering the OG in both periods as in Policy 2 will *decrease* the total revenue. Policy 3 offers the OG only in the second period, making the seller gain a higher revenue than not using opaque selling. Therefore, even though offering the OG all the time is not a good policy, offering it sometimes can be beneficial. This emphasizes the importance of the tactical control of opaque selling. It should be noted that scaling up the inventory and time horizon proportionally will not eliminate the advantage of opaque selling since it expands the market.

When selling multiple products with a fixed inventory each in a finite horizon, the following questions naturally arise: how should the seller dynamically control the product offers, including both physical goods and OGs, over time? How should the seller make the fulfillment decisions for OGs? In other word, which physical good is assigned to fulfill the demand for the OG? To answer these questions, we consider a seller with a fixed capacity/inventory for each physical good prior to a finite selling horizon. Customers arrive sequentially during the horizon. The seller dynamically controls the set of products being offered to customers, and determines which physical good is used to fulfill the demand of OGs. We formulate the seller's problem as a continuous-time, discrete-state, finite-horizon stochastic dynamic program. Due to well-known curse of dimensionality, the optimal policy is computationally infeasible. To overcome this difficulty, we study the corresponding deterministic version of the problem, known as the "fluid control problem." We show that the optimal revenue of the fluid control problem serves as an upper bound for the optimal revenue

[1] The detailed calculation of the revenue under the three policies is shown in Appendix A.

of the stochastic dynamic program. The optimal solution of the fluid control problem can be implemented as a heuristic control policy, which is shown to be asymptotically optimal when both inventory and demand are scaled up. However, the fluid policy is a static policy that does not react to the changing inventory status and time. We further consider a decomposition heuristic that approximates the original dynamic program with a collection of single-dimensional dynamic programs. This method generates a dynamic control policy based on the remaining inventory and remaining time, and outperforms the fluid policy in most problem instances that we have constructed.

Our work contributes to the literature on opaque selling in the following ways. First, we consider multiple OGs that are created using multiple physical goods. Much of the existing literature considers situations with only a single OG, which is a restriction for practical implementation and can lead to substantial missed revenue opportunities. Second, our work contributes to the understanding of the impact of inventory and selling periods on opaque selling, which is understudied in the existing literature. Third, we consider an endogenous fulfillment mechanism under which the seller decides which physical good is used to fulfill the demand. Our dynamic programming model and the ensuing solution methods directly address these questions. Our modelling framework also allows us to investigate the value of opaque selling with multiple OGs. Compared to no opaque selling (i.e., the traditional choice-based revenue management problems), our numerical study shows that offering multiple OGs can substantially improve the seller's revenue. We find that the seller will benefit more from selling multiple OGs when product supply is not tight relative to demand, and when inventory levels of each product are more imbalanced. Because selling multiple OGs enables the seller to fine-tune segmentation of the market and attract more demand, the value of offering multiple OGs becomes more pronounced when there is ample supply. Moreover, opaque selling enables the seller to assign the physical good with high inventory to fulfill the demand of OG, thus it is well expected that opaque selling is more valuable with a large degree of inventory imbalance. Importantly, we show that even offering a single OG can achieve large benefits. This suggests that implementing opaque selling is feasible in many practical settings.

The remainder of this paper is organized as follows. Section 2 provides the literature review. Section 3 introduces the model and formulates the seller's revenue management problem as an MDP. Section 4 studies the fluid control problem, develops a fluid policy, and establishes its asymptotic optimality. Section 5 proposes a decomposition heuristic. Section 6 reports the results of numerical experiments. Finally, Section 7 concludes and discusses future research directions. In this paper, "she" refers to the customer and "he" refers to the seller. All the proofs are presented in the Appendix.

2 Literature Review

Our work falls in the steam of literature on probabilistic/opaque selling. Fay and Xie (2008) first define the probabilistic goods and study the underlying mechanism of probabilistic selling. Using a standard Hotelling model, they investigate why, how, and when probabilistic selling can improve a seller's profit. Shapiro and Shi (2008) consider a system of N hotels, each of which publicly announces its list price and decides whether to participate in opaque selling or not. They show that opaque selling enables service providers to price discriminate between the customers who are sensitive to service characteristics and those who are not. Jiang (2007) adopts opaque selling to a firm with two flights scheduled at different times, and investigates when opaque selling improves profits.

Several authors are interested in understanding the relative merit of probabilistic selling against other selling strategies. Fay and Xie (2010) compare probabilistic selling with advance selling since both strategies deal with buyer uncertainty in valuations. Jerath et al. (2009, 2010) compare opaque selling with last-minute selling through an intermediary. Rice et al. (2014) compare probabilistic selling with markdown selling and examine when probabilistic selling can be more profitable. Elmachtoub and Hamilton (2017) compare opaque selling with two common selling strategies including discriminatory pricing and single pricing under two types of customers including risk-neutral and pessimistic customers.

Quite a few papers examine the impact of probabilistic/opaque selling on a seller's product and inventory management, and customer purchase behavior. For example, Fay (2008) analyzes a game in which two service providers first determine how many units to be allocated to the opaque channel, and then the intermediary chooses the price of the OG. He shows that the opaque channel increases the degree of price rivalry and thus reduces industry profits unless firms have very loyal customers. Fay et al. (2015) examine how probabilistic selling affects a retailer's product mix decision, and find that probabilistic selling can either increase or decrease the degree of product differentiation depending on market heterogeneity. Fay and Xie (2015) uncover a new role that probabilistic selling plays for inventory management. Against the common belief that deferring product assignment will help balance supply and demand for each individual product, they find that early product allocation can be more profitable for the firm. Li et al. (2016) investigate how probabilistic selling affects a firm's profit when consumers display variety seeking behavior. They find that, compared with traditional selling, probabilistic selling improves the firm's profit, but the improvement decreases in the degree of variety seeking.

While most authors consider opaque selling as posted prices by sellers, several authors consider opaque selling through the so-called name-your-own-price (NYOP), where the prices are determined by customers. Cai et al. (2009) investigate how channel structure (single-channel vs. dual-channel) and bidding options (single-bid vs. double-bid) affect the optimal reserve prices under the NYOP policy. Chen et al. (2014) examine the effects of two selling mechanisms including posted-price and name-your-own-price on competing service providers. Huang et al. (2017) study NYOP with two competing suppliers, and explore its impact on the channel structure in a competitive market.

All the aforementioned papers adopt stylized two-period models assuming that all customers arrive simultaneously at the beginning of the selling horizon. They focus on the fundamental insights and managerial implications that probabilistic/opaque selling brings about under various selling environments and market conditions. Our work differs in that we investigate how opaque selling can be implemented practically in a standard network revenue management framework. Specifically, we seek to develop a dynamic control policy for a seller who offers multiple substitutable products with fixed inventories to sequentially arriving customers with heterogeneous valuations. An exception is the work by Xiao and Chen (2014), who consider probabilistic selling of two products over a finite selling period. By devising different scenarios as benchmarks, they focus on the exploration of the demand induction effect (i.e., incremental profit from inducing more customers to purchase), the demand dilution effect (i.e., profit loss when customers switch from physical goods to OGs), and the inventory pooling effect from probabilistic selling. Similar to Xiao and Chen (2014), the seller in our model can dynamically determine the product offers and the assignment of physical goods for fulfillment of OGs. However, we study a more general situation with a menu of physical goods and OGs. Due to the curse of dimensionality, we develop heuristics including a fluid policy and a decomposition heuristic that are easy to implement, and generate insights on opaque selling of multiple products.

Our model and solution borrows ideas from the growing literature on choice-based revenue management and dynamic pricing; see, e.g., Talluri and van Ryzin (2004a, 2004b), Zhang and Cooper (2005), Liu and van Ryzin (2008), Kunnumkal and Topaloglu (2008, 2010), Bront et al. (2009), Zhang and Adelman (2009), Chaneton and Vulcano (2011), Meissner and Strauss (2012), Chen et al. (2017), Hu et al. (2016), and Chen and Chen (2018). The power of heuristic policies from deterministic approximations is well-known in the revenue management literature; see, e.g., Cooper (2002), Maglaras and Meissner (2006), Reiman and Wang (2008), Zhang (2011), and Jasin and Kumar (2012). Particularly, our work shares a close similarity with Jasin and Kumar (2012) in the sense that resources can be randomly consumed by products. However, in their work, the resource consumption of each product follows an exogenously given distribution and it is beyond the seller's control. In contrast, the resource consumption of OGs is endogenously determined by the seller in our paper; that is, the seller decides which resource (physical good) should be allocated to fulfill a request for an OG. In fact, in all the existing work on choice-based revenue management, customers choose among physical goods and thus demand fulfillment is not a decision for the seller. However, we allow the seller to determine which physical good to be assigned to fulfill the demand for OGs. Therefore, the seller dynamically controls not only the set of products being offered but also the assignment of physical goods to fulfill the demand; hence, the decision problem is more involved.

3 Model Formulation

In this section we describe the problem and introduce the model formulation.

3.1 Problem Description

Consider a seller offering m substitutable *physical goods*, indexed by $\mathcal{M} = \{1, 2, \ldots, m\}$, within a finite selling season $[0, T]$. Each arriving customer purchases at most one unit of the physical goods. Besides offering these m physical goods individually, the seller may sell them in the form of an OG. An OG is a subset of all the available physical goods, such that whenever a customer purchases an OG, her demand will be fulfilled by one of the physical goods. Which physical good is used to to fulfill the demand is determined by the seller. Note that each physical good $i \in \mathcal{M}$ can also be viewed as an OG with a subset of only one physical good. To the extent possible, we reserve the use of OGs for products with fulfillment uncertainty, even though a physical good can also be viewed as an OG with no fulfillment uncertainty. A *product* can be either a physical good or an OG throughout the paper.

The seller stocks Q_i units of physical good $i \in \mathcal{M}$ prior to the selling season. Let $Q = (Q_1, Q_2, \ldots, Q_m)$ be the vector of initial inventory levels. We assume the procurement lead time is relatively long such that there is no opportunity for replenishment during the selling season. For simplicity, the salvage value of each physical good is assumed to be zero at the end of the selling season (note that incorporating a salvage value for each physical good will not impact our analysis and major results).

Suppose that the seller offers n products, indexed by $\mathcal{N} = \{1, 2, \ldots, n\}$, using the multiple physical goods. Each product is characterized by a vector (r_j, a_j) where r_j is the selling price of product $j \in \mathcal{N}$ and $a_j := (a_{1j}, a_{2j}, \ldots, a_{mj})$ denotes the physical goods contained in product j. In particular, $a_{ij} = 1$ if physical good i is contained in product j and $a_{ij} = 0$ otherwise. For ease of exposition, let the first m products correspond to the m physical goods; thus $a_{ii} = 1$ and $a_{i'i} = 0$ $\forall i, i' \in \mathcal{M} : i \neq i'$. It is reasonable to assume that r_j is no greater than the maximum selling price of the physical goods contained in product j. This is typically the case in practice since customers usually expect a discount for OGs. That is, we have

$$r_j \leq \max_{i \in \mathcal{M}} r_i \mathbb{I}_{\{a_{ij}=1\}}, \quad \forall j \in \mathcal{N},$$

where $\mathbb{I}_{\{\cdot\}}$ is the indicator function.

The seller makes two-dimensional decisions over the selling horizon, with the objective of maximizing revenue. On the one hand, he controls the subset of products available to customers,

$S \subseteq \mathcal{N}$, over time. We enforce an "honesty" constraint such that a product set S is offered if and only if each physical good contained in S has a positive inventory level. That is, we exclude the possibility that the seller may still offer a set S even if some of the physical goods contained in S are depleted. This kind of untruthful behavior on the seller's part is possible in practice due to the asymmetric information between the seller and customers (as customers typically cannot observe the seller's real time inventory level) and is likely beneficial for the seller. An implication of our modelling choice is that our model provides a lower bound on the benefit of opaque selling; the benefit can increase if the seller's untruthful behavior is accounted for. On the other hand, given that a customer chooses an OG j from the offer set S, the seller determines which physical good that is contained in the OG j for demand fulfillment. Notice that such random product assignment shares some similarity with menus of lotteries well studied in the economics and computer science literature (e.g., Gauthier and Laroque, 2017, Chen et al., 2015). However, the key difference is that the assignment probability is endogenously determined by the seller in our model, while it is randomly drawn in theirs.

Following the revenue management literature, customers are assumed to arrive according to a homogeneous Poisson process with rate λ. At each time t, an arriving customer purchases at most one unit of a product. We assume customers choose among the products following a discrete choice model. The probability of choosing product $j \in S$ is $P_j(S)$ when the offer set is S. Note that the choice probabilities depend only on the offer set, but not on the seller's fulfillment decision. It is possible that customers may adjust their purchase behavior based on the seller's fulfillment decisions, which we do not consider in this work. We assume the choice probabilities follow some regularity assumptions: (i) $P_j(\emptyset) = 0\ \forall j \in S$, and $P_j(S) \geq 0$ if and only if $j \in S$; that is, only offered products can generate demand. (ii) If $S \subseteq S'$, then $\sum_{j=1}^{n} P_j(S) \leq \sum_{j=1}^{n} P_j(S')$, meaning that offering more products will lead to higher total demand across all products (this is the demand induction or market expansion effect). (iii) If $S \subseteq S'$, then $\forall j \in S$, $P_j(S) \geq P_j(S')$, meaning that an added product in the offer set will decrease the demand from existing products (this is the demand dilution effect). The regularity assumptions are implied by the vast majority of parametric choice models considered in the existing literature (e.g., Liu and van Ryzin, 2008; Talluri and van Ryzin, 2004b; Feldman and Topaloglu, 2017). We let $P_0(S)$ denote the no-purchase probability such that $P_0(S) + \sum_{j=1}^{n} P_j(S) = 1$.

Depending on the inventory levels of physical goods and the remaining time, the seller dynamically controls the set of products to be offered at each point in time and the product fulfillment so as to maximize the expected revenue.

3.2 The MDP Formulation and Solution

We now formulate the seller's revenue maximization problem as a Markov decision process (MDP). At any time $t \in [0, T]$, the state of the system is the vector of remaining inventory levels of physical goods, denoted as $y = (y_1, y_2, \dots, y_m)$ where $0 \le y \le Q$. An non-anticipative control policy specifies an offer set over time, with

$$(u, \Theta) = \Big\{ \Big((u^t(S))_{\forall S \subseteq \mathcal{N}}, (\theta_{ij}^t)_{\forall i \in \mathcal{M}, \forall j \in \mathcal{N}} \Big) : 0 \le t \le T \Big\},$$

where $u^t(S)$ denotes whether the set S is offered at time t, i.e.,

$$u^t(S) = \begin{cases} 1, & \text{if the set } S \text{ is offered at time } t, \\ 0, & \text{otherwise,} \end{cases} \quad \forall S \subseteq \mathcal{N},$$

and $(\theta_{ij}^t)_{m \times n}$ represents the control of fulfillment, with

$$\theta_{ij}^t = \begin{cases} 1, & \text{if physical good } i \text{ is assigned to fulfill product } j \text{ at time } t, \\ 0, & \text{otherwise} \end{cases} \quad \forall i \in \mathcal{M}, \forall j \in \mathcal{N}.$$

Note that $\sum_S u^t(S) = 1 \ \forall t$, since at most one product set can be offered at any time. In addition, $\sum_{i \in \mathcal{M}} a_{ij} \theta_{ij}^t = 1 \ \forall t$ and $j \in \mathcal{N}$, as only a single physical good contained in product j should be assigned for fulfillment of any demand. Let $\mathcal{U}$ denote the set of feasible control policies.

Let $\Lambda = \{\Lambda(t) : 0 \le t \le T\}$ denote the process of customer arrivals and T_k denote the arrival time of the kth customer. Under the control policy (u, Θ), the cumulative sales for product $j \in \mathcal{N}$ up to time τ can be written as

$$D_j(\tau, u, \Theta) = \sum_{k=1}^{\Lambda(\tau)} \sum_{S \subseteq \mathcal{N}} u^{T_k}(S) \mathbb{I}_{\{\zeta_k(S) = j\}}, \ \forall \tau \in [0, T], \ j \in \mathcal{N}.$$

In the above, for each k and S, $\{\zeta_k(S) : k \ge 1\}$ is a sequence of independent random variables with distribution $\mathbb{P}\{\zeta_k(S) = j\} = P_j(S)$ for all $j \in \{0\} \cup \mathcal{N}$. The set of all feasible non-anticipating policies during $[t, T]$ must satisfy:

$$\sum_{j \in \mathcal{N}} \Big\{ \int_t^T \theta_{ij}^\tau dD_j(\tau, u, \Theta) \Big\} \le y_i, \ \forall i \in \mathcal{M}. \tag{1}$$

The seller's expected revenue over $[t, T]$ for a given policy (u, Θ) can be expressed as

$$J_{(u,\Theta)}(t, y) = \mathbb{E} \Big\{ \sum_{j \in \mathcal{N}} \int_t^T r_j dD_j(\tau, u, \Theta) \Big\}. \tag{2}$$

The optimal value of (2), denoted by $V(t, y)$, called the *revenue-to-go function*, is the supremum of $J_{(u,\Theta)}(t, y)$ over all feasible policies $(u, \Theta) \in \mathcal{U}$; that is,

$$V(t, y) = \sup_{(u,\Theta) \in \mathcal{U}} J_{(u,\Theta)}(t, y), \tag{3}$$

subject to (1) and with boundary conditions:

$$\begin{aligned} V(T, y) &= 0, \ \forall y, 0 \le y \le Q, \\ V(t, 0) &= 0, \ \forall t \in [0, T]. \end{aligned} \tag{4}$$

As one can expect, the revenue-to-go function has the properties established in Lemma 1. The proof can be shown by replacing jump point processes with their intensities; interested readers are referred to Feng and Gallego (1995).

Lemma 1. *The revenue-to-go function $V(t, y)$ defined in (3) satisfies: (i) $V(t, y)$ is continuous and decreasing in time t, and (ii) $V(t, y)$ is increasing in the inventory level y.*

The seller's problem is to determine the product set available to customers at each point in time and the physical good assignment for fulfillment of OGs so as to maximize the total expected revenue in the finite horizon $[0, T]$:

$$V^* = V(0, Q) = \sup_{(u,\Theta)\in\mathcal{U}} J_{(u,\Theta)}(0, Q). \tag{5}$$

Note that a general feasible policy can be path dependent; i.e., the set being offered and physical good being assigned at time t depends on the entire sales trajectory during the time interval $[0, t)$. However, among all the possible policies $\mathcal{U}$, one can always choose a Markovian policy that is optimal (Theorem 5.5.1 in Puterman, 1994). In the rest of the paper, without loss of generality, we focus on Markovian policies in which the decision of set offering depends only on the current inventory level.

To formulate $V(t, y)$ recursively, let $\mathbb{B}(y)$ be the set of products that can be offered in state $y = (y_1, y_2, \ldots, y_m)$. When $y_i = 0$ for some $i \in \mathcal{M}$, all the products that use physical good i should be excluded from $\mathbb{B}(y)$. Clearly, starting from $\mathbb{B}(Q) = \mathcal{N}$ at time $t = 0$, the set $\mathbb{B}(y)$ shrinks as the inventory level of some physical goods drops to zero. In particular, whenever all the physical goods are sold out, $\mathbb{B}(0)$ reduces to the $\emptyset$ and the seller has to terminate the selling process.

Consider a time interval $[t, t + \Delta t)$ where Δt is sufficiently small such that the probability of more than one customer arrival during this interval is negligible. For a given set $S \subseteq \mathbb{B}(y)$ offered within the interval $[t, t + \Delta t)$, the probability that a customer arrives and purchases product j is $\lambda\Delta t P_j(S)$, $j \in \mathcal{N}$. When the customer purchases product j, the seller earns a revenue of r_j and sells one unit of physical good $i \in \mathcal{M}$ contained in product j according to the fulfillment decision $\hat{\theta} := (\theta^t_{ij})_{m\times n}$. As such, conditioning upon all the possible events during $[t, t + \Delta t)$, we can derive the seller's expected revenue in the following:

$$\begin{aligned} V(t, y|S, \theta) \;&=\; \sum_{j\in\mathcal{N}} \Big\{ \lambda\Delta t P_j(S) \Big[r_j + \sum_{i\in\mathcal{M}, a_{ij}=1} \theta^t_{ij} V(t + \Delta t, y - e_i) \Big] \Big\} \\ &\quad + \Big[1 - \lambda\Delta t \sum_{j\in\mathcal{N}} P_j(S) \Big] V(t + \Delta t, y), \end{aligned}$$

where e_i is the m-vector whose ith component is 1 and all other components are 0.

The maximum expected revenue is achieved by choosing the offer set S and assignment decision θ that maximizes $V(t, y|S, \theta)$; that is,

$$V(t, y) = \max_{S\subseteq\mathbb{B}(y),\theta} V(t, y|S, \theta).$$

Conducting some simple algebraic transformations and letting $\Delta t \to 0$, we obtain the following Hamilton-Jacobi equation:

$$
\begin{aligned}
-\frac{\partial V(t,y)}{\partial t} &= \max_{S\subseteq \mathbb{B}(y),\theta} \sum_{j\in\mathcal{N}} \lambda P_j(S)\Big[r_j - \sum_{i\in\mathcal{M},a_{ij}=1} \theta_{ij}^t\Big(V(t,y)-V(t,y-e_i)\Big)\Big], \\
&= \max_{S\subseteq \mathbb{B}(y)} \sum_{j\in\mathcal{N}} \lambda P_j(S)\Big[r_j - \min_{i\in\mathcal{M},a_{ij}=1} \Delta_i V(t,y)\Big)\Big],
\end{aligned} \tag{6}
$$

where $\Delta_i V(t,y) = V(t,y) - V(t,y-e_i)$. Conditioning upon the arrival time of the next customer, we can show that (6) is equivalent to the following:

$$
V(t,y) = \int_t^T \lambda e^{-\lambda(\tau-t)} g(\tau,y)d\tau, \tag{7}
$$

where

$$
g(\tau,y) \quad := \quad \max_{S\subseteq\mathbb{B}(y)} \Big\{\sum_{j\in\mathcal{N}} P_j(S)\Big[r_j + \max_{i\in\mathcal{M},a_{ij}=1} V(\tau,y-e_i)\Big] + P_0(S)V(\tau,y)\Big\}. \tag{8}
$$

One may wonder whether the revenue-to-go function $V(t,y)$, defined in (6) and (7), exists and is unique. We provide a detailed proof in Appendix B. Even so, it is quite difficult to obtain the seller's optimal control policy, especially for a high dimensional problem. Therefore, in the following section we seek to develop some heuristic policies.

4 The Fluid Policy and Its Asymptotic Optimality

In this section, we study the fluid policy based on the fluid control problem, which is a deterministic version of the original MDP problem. The fluid control problem is introduced in Section 4.1. We develop a fluid policy based on the solution of the fluid control problem in Section 4.2 and show that it is asymptotically optimal for the original MDP problem in Section 4.3.

4.1 The Fluid Control Problem

We define the fluid control problem as follows. Customers arrive continuously at a constant rate λ. The demand rate for product $j \in \mathcal{N}$ when S is offered is given by $\lambda P_j(S)$. Product j consumes the stock of physical good $i \in \mathcal{M}$ at time-varying rate $\hat{\theta}_{ij}^t$. The seller decides the set S to be offered at each point in time t and the product fulfillment to maximize revenue. The associated control policy $(u, \hat{\Theta})$ satisfies

$$
\begin{aligned}
u^t(S) \in [0,1], &\qquad \sum_S u^t(S) = 1, \;\; \forall t \in [0,T], S \subseteq \mathcal{N}, \\
\hat{\theta}_{ij}^t \in [0,1], &\qquad \sum_{i\in\mathcal{M}} a_{ij}\hat{\theta}_{ij}^t = 1, \forall t \in [0,T], i \in \mathcal{M}, j \in \mathcal{N}.
\end{aligned}
$$

To distinguish from the MDP model, let $\mathcal{V}$ be the set of all feasible policies for the fluid control problem.

For a given initial inventory level Q, the fluid control problem can be formulated as follows:

$$R^* = \max_{(u,\hat{\Theta})\in\mathcal{V}} \Big\{ \sum_S \sum_{j\in\mathcal{N}} \int_0^T \Big(\lambda r_j u^t(S) P_j(S)\Big) dt \Big\}, \tag{9}$$

$$s.t. \sum_S \sum_{j\in\mathcal{N}} \int_0^T \Big(\lambda u^t(S) P_j(S) \hat{\theta}_{ij}^t\Big) dt \le Q_i, \ \forall i \in \mathcal{M}. \tag{10}$$

The constraint (10) says that the quantity of each physical good that is used to fulfill customer demand cannot exceed its corresponding stocking level.

Theoretically, the optimal offer and fulfillment policies are both time-dependent. The following proposition, however, shows that it is optimal for the firm to adopt a stationary (time-independent) fulfillment policy in the fluid control problem.

Proposition 1. *In the optimization problem (9)–(10), there exists an optimal stationary fulfillment policy $\hat{\Theta}^* = (\hat{\theta}_{ij}^*)_{m\times n}$; that is, it is without loss of optimality to take*

$$\hat{\theta}_{ij}^t = \hat{\theta}_{ij}^*, \quad \forall i \in \mathcal{M}, j \in \mathcal{N}, t \in [0, T].$$

As such, one can focus on the fluid solution that has a stationary fulfillment policy $(\hat{\theta}_{ij})_{m\times n}$ by taking $\hat{\theta}_{ij}^t = \hat{\theta}_{ij}$ in (10). Furthermore, for each $S \subseteq \mathcal{N}$, let

$$\tau(S) := \int_0^T u^t(S) dt, \tag{11}$$

be the total time that the set S is offered during the entire time horizon $[0, T]$. The problem (9)–(10) can then be rewritten as:

$$R^* = \max_{(\tau,\hat{\theta})} \sum_S \sum_{j\in\mathcal{N}} \Big(\lambda r_j P_j(S) \tau(S)\Big), \tag{12}$$

$$s.t. \sum_S \sum_{j\in\mathcal{N}} \Big(\lambda P_j(S) \hat{\theta}_{ij} \tau(S)\Big) \le Q_i, \qquad \forall i \in \mathcal{M}, \tag{13}$$

$$\sum_S \tau(S) \le T, \tag{14}$$

$$\sum_{i\in\mathcal{M}} a_{ij} \hat{\theta}_{ij} = 1, \qquad \forall j \in \mathcal{N}, \tag{15}$$

$$\tau(S) \ge 0, \qquad \forall S \subseteq \mathcal{N},$$

$$0 \le \hat{\theta}_{ij} \le 1, \qquad \forall i \in \mathcal{M}, j \in \mathcal{N}. \tag{16}$$

The problem (12)–(16) is a non-linear optimization problem since constraint (13) is nonlinear. However, this problem can be transformed into a linear program by a change of variables. For any (i, j, S), let $X_{ij}(S) = \tau(S)\hat{\theta}_{ij}$ and replace decision variables $\hat{\theta}$ by X. Consequently, problem (12)–(16) transforms to the following form:

$$R^* = \max_{(\tau,X)} \sum_S \sum_{j\in\mathcal{N}} \Big(\lambda r_j P_j(S)\tau(S)\Big), \tag{17}$$

$$s.t. \sum_S \sum_{j\in\mathcal{N}} \Big(\lambda P_j(S) X_{ij}(S)\Big) \le Q_i, \qquad \forall i \in \mathcal{M}, \tag{18}$$

$$\sum_S \tau(S) \le T, \tag{19}$$

$$\sum_{i\in\mathcal{M}} a_{ij} X_{ij}(S) = \tau(S), \qquad \forall j \in \mathcal{N}, S \subseteq \mathcal{N} \tag{20}$$

$$\tau(S) \ge 0, \qquad \forall S \subseteq \mathcal{N},$$

$$0 \le X_{ij}(S) \le \tau(S), \qquad \forall i \in \mathcal{M}, j \in \mathcal{N}, S \subseteq \mathcal{N}. \tag{21}$$

Observe that the program after the change of variables is a linear program and that a solution to the problem (12)–(16) can be constructed from a solution of the transformed program. In general, the number of variables and constraints of the linear program (17)–(21) increases exponentially in the number of products n. Therefore, the size of the problem can grow quickly. In our numerical study, we are able to solve medium sized problems using a commercial linear programming solver. For a large n, specialized algorithms will be required.

Let (τ^*, X^*) be an optimal solution to the problem (17)–(21), and $(\tau^*, \hat{\theta}^*)$ be an optimal solution to the problem (12)–(16). We can show some structural properties of τ^*. First, the sequence of sets being offered does not affect the seller's revenue. Second, if it is optimal to offer the null set $\emptyset$ for some positive time period (i.e., $\tau^*(\emptyset) > 0$), then all of the physical goods must be sold out during $[0, T]$; that is,

$$Q_i = \sum_S \sum_{j\in\mathcal{N}} \Big(\lambda P_j(S)\hat{\theta}^*_{ij}\tau^*(S)\Big), \ \forall i \in \mathcal{M}.$$

In other words, if there exists an over-stocked physical good, i.e.,

$$Q_i > \sum_S \sum_{j\in\mathcal{N}} \Big(\lambda P_j(S)\hat{\theta}^*_{ij}\tau^*(S)\Big)$$

for some $i \in \mathcal{M}$, then we must have $\tau^*(\emptyset) = 0$.

Proposition 2. *There exists an optimal solution τ^* for the fluid control problem (12)–(16), such that for any $i, j \in \mathcal{M}$, $i \ne j$, we have $\tau^*(\{i\}) \times \tau^*(\{j\}) = 0$. That is, for any pair of products, there exists an optimal solution such that at most one of them is offered individually in a specific time period.*

Proposition 2 shows that there exists an optimal solution τ^*, such that among any pair of physical goods (i, j), $\tau^*(\{i\})$ and $\tau^*(\{j\})$ should have, at least, one zero value. The intuition is as follows. Suppose that, for a pair of products, both are offered individually in two non-overlapping time slots. We can construct a solution such that both products are offered together and at most one of them is offered individually with the same sales amount and revenue. The proposition implies the existence of an optimal solution such that, among the set of physical goods, at most one physical good is offered individually for some duration. This is consistent with observations where physicals goods are often offered together in practice.

Given an optimal solution $(\tau^*, \hat{\theta}^*)$, the optimal fluid revenue is

$$R^* = \sum_{S} \sum_{j \in \mathcal{N}} \Big(\lambda r_j P_j(S) \tau^*(S) \Big). \tag{22}$$

We now claim that R^* provides an upper bound on the optimal expected revenue of the MDP problem (5).

Theorem 1. *For any initial inventory level Q, we have $R^* \geq V^*$.*

The fact that a deterministic fluid formulation provides an upper bound on the revenue is a classical result in the revenue management literature (see, e.g., Cooper, 2002; Liu and van Ryzin, 2008; Reiman and Wang, 2008). The intuition is that the uncertainties associated with customer arrival and choice process reduce the expected revenue of the seller. We show the same result holds when opaque selling is considered. Since our model involves OGs, our setup is slightly more involved than the aforementioned papers due to fulfillment uncertainty. Theorem 1 allows us to use the fluid revenue R^* as a benchmark to evaluate the asymptotic optimality of the fluid policy in the following subsections.

4.2 The Fluid Policy

Consider an optimal solution, $(\tau^*, \hat{\theta}^*)$, to the fluid control problem (12)–(16). Suppose, among all the possible offer sets, $S_1, \ldots, S_K$ have positive durations, denoted as $\tau^*(S_\ell)$ for $\ell = 1, \ldots, K$. We divide the selling horizon $[0, T]$ into K intervals by duration $\tau^*(S_\ell)$ for $\ell = 1, \ldots, K$. The optimal fluid revenue R^* is then obtained when the seller offers S_ℓ in the ℓth interval and the fulfillment probability is kept at $\hat{\theta}^*$. Given the initial inventory level $Q = (Q_1, Q_2, \ldots, Q_m)$, let the corresponding remaining inventory level at the end of the ℓth duration be $\mathcal{L}_\ell = (\mathcal{L}_{\ell 1}, \ldots, \mathcal{L}_{\ell m})$. Clearly, we have $\mathcal{L}_0 = Q$. For each $\ell = 1, 2, \ldots, K$ and $i \in \mathcal{M}$:

$$\mathcal{L}_{\ell i} = \mathcal{L}_{(\ell-1)i} - \lambda \tau^*(S_\ell) \sum_{j \in \mathcal{N}} P_j(S_\ell) \hat{\theta}^*_{ij} = Q_i - \lambda \sum_{k=1}^{\ell} \sum_{j \in \mathcal{N}} P_j(S_k) \hat{\theta}^*_{ij} \tau^*(S_k),$$

which is non-increasing in ℓ and non-negative for each i.

In the following we propose a fluid policy for the original MDP model.

Definition 1. (Fluid Policy π) *On the side of offer-set control, order the sets with positive durations in the fluid solution randomly and then offer them sequentially. The set S_ℓ, $\ell = 1, 2, \ldots, K$, is offered until the inventory level of one physical good (say, physical good i) contained in S_ℓ first drops to $\lfloor \mathcal{L}_{\ell i} \rfloor$ or the time duration $\tau^*(S_\ell)$ has elapsed. On the side of product fulfillment, for any customer that chooses product j, assign physical good i for the fulfillment with probability $\hat{\theta}^*_{ij}$.*

In the definition above, $\lfloor x \rfloor$ is the largest integer that does not exceed x; i.e., $\lfloor x \rfloor = \max\{i \in \mathbb{Z} : i \le x\}$. Clearly, the fluid policy π can be easily implemented, without the need for any re-optimization. Given an inventory level at the beginning of the ℓth interval being y, the seller can sell a maximum of $(y_i - \lfloor \mathcal{L}_{\ell i} \rfloor)$ units for each physical good i contained in S_ℓ. This shares some similarities to the class of *allocation* policies studied in the literature (e.g., Cooper, 2002; Reiman and Wang, 2008). However, our policy is more complex than the traditional allocation problems because the problem involves three sources of uncertainties: customer arrivals, customer choices, and product fulfillment (the literature has no uncertainty in product fulfillment). In particular, to ensure the feasibility of the fluid policy, when offering set S_ℓ in the lth interval, we limit the sales volume of each physical good contained in S_ℓ, rather than the accepted demand as in Liu and van Ryzin (2008). This difference arises from the uncertainty of product fulfillment associated with opaque selling.

The fluid policy π defined in Definition 1 is optimal for the fluid control problem, if one relaxes the integer restriction on the inventory level. However, applying the fluid policy π to the original MDP problem may incur some revenue loss compared with R^*. In the following subsection, we will evaluate the performance of the fluid heuristic π and show that the policy is asymptotically optimal when both the inventory level and the demand rate are scaled up proportionally.

4.3 Asymptotic Optimality

We consider a sequence of problems, both stochastic and deterministic, indexed by k ($k = 1, 2, \ldots$), in which the initial inventory level and the arrival rate are scaled up by a factor of k; that is, $Q^k = kQ$ and $\lambda^k = k\lambda$. In the k-scaled problem, the optimal revenues of the MDP problem and the corresponding fluid control problem are V^{k*} and R^{k*}, respectively. Note that the optimal solution to the k-scaled fluid control problem is independent of k because the inventory level and the arrival rate are all scaled up by the same factor k. Therefore, $R^{k*} = kR^*$ where R^* is given by (22). Furthermore, the durations $\tau^*(S_\ell)$, $\ell = 1, 2, \ldots, K$, being used to construct policy π are independent of factor k. We let $t_0 = 0$ and define

$$t_\ell := \sum_{i=0}^{\ell} \tau^*(S_i), \ell = 1, 2, \ldots, K.$$

Then t_ℓ can be interpreted as the time when the set S_ℓ starts to be offered. The corresponding remaining inventory level at time t_ℓ, denoted as $\mathcal{L}^k_\ell$, is $\mathcal{L}^k_\ell = k\mathcal{L}_\ell$. We denote V^k_π as the resulting expected revenue for the k-scaled MDP problem when implementing policy π. With a slight abuse of notation, let $V_\pi^k(t_\ell, y)$ be the expected revenue collected during $(t_\ell, T]$ if policy π is adopted with an inventory level y at time t_ℓ. As such, we have

$$V_\pi^k = V_\pi^k(0, kQ).$$

Clearly, V_π^k is no greater than the optimal revenue V^{k*}, and it follows from Theorem 1 that

$$V_\pi^k \le V^{k*} \le R^{k*}. \tag{23}$$

We call $(V^{k*} - V_\pi^k)$ the revenue loss incurred from adopting the heuristic policy π for the k-scaled MDP problem. The true optimal revenue V^{k*}, however, is difficult to calculate, and thus we consider the fluid gap, defined as $(R^{k*} - V_\pi^k)$. Clearly, for each k, the fluid gap dominates the revenue loss (by Theorem 1). Following the literature, we will evaluate the performance of the policy π by measuring the fluid gap $(R^{k*} - V_\pi^k)$ in an asymptotic regime.

We first formulate the revenue V_π^k and establish its lower bound. Define

$$\mu_i(S) = \lambda \sum_{j\in\mathcal{N}} \hat{\theta}_{ij}^* P_j(S), \quad \forall i \in \mathcal{M}, S \subseteq \mathcal{N},$$

as the *effective* demand rate of physical good i when set S is offered. Note that $\mu_i(S) = 0$ if physical good i is not contained in S. When $\mu_i(S) > 0$, the corresponding *effective* unit revenue of physical good i when set S is offered is defined by

$$p_i(S) = \frac{\lambda}{\mu_i(S)} \sum_{j\in\mathcal{N}} \hat{\theta}_{ij}^* r_j P_j(S), \quad \forall i \in \mathcal{M}, S \subseteq \mathcal{N}.$$

Moreover, we denote

$$\bar{p}(S) = \max_{i\in\mathcal{M}} \left\{ p_i(S) \mathbb{I}_{\{\mu_i(S)>0\}} \right\}, \quad \forall S \subseteq \mathcal{N}$$

as the maximum effective unit revenue of physical goods contained in S, which will be used in the later analysis. (Note that $\bar{p}(\emptyset) = 0$.)

Now consider the k-scaled MDP problem. For a given ℓ ($\ell = 1, \ldots, K$), suppose the seller starts with inventory level $\lfloor k\mathcal{L}_{\ell-1} \rfloor$ at time $t_{\ell-1}$, and fluid policy π is adopted. For ease of presentation, let

$$\mathbb{A}_\ell^k = \left\{ i \middle| \lfloor k\mathcal{L}_{\ell i} \rfloor < \lfloor k\mathcal{L}_{(\ell-1)i} \rfloor, i \in \mathcal{M} \right\}$$

be the set of physical goods that are used to fulfill demand when policy π is adopted during $(t_{\ell-1}, t_\ell]$. In the following, we formulate $V_\pi^k(t_{\ell-1}, \lfloor k\mathcal{L}_{\ell-1} \rfloor)$ via dynamic programming.

Within each time interval $[t_{\ell-1}, t_\ell)$, for each physical good $i \in \mathbb{A}_\ell^k$, the problem is equivalent to the one facing a homogeneous Poisson process, denoted as $N_{\ell i}^k = \{N_{\ell i}^k(t) : t \ge 0\}$, with rate being

$k\mu_i(S_\ell)$ and unit revenue being $p_i(S_\ell)$, until one physical good, say i, first runs out of its *inventory quota*, $\lfloor k\mathcal{L}_{(\ell-1)i}\rfloor - \lfloor k\mathcal{L}_{\ell i}\rfloor$. With a slight abuse of notation, let $N_{\ell i}^k(t,t')$ denote the accumulated arriving customers of process $N_{\ell i}^k$ during the time interval $(t,t']$. Consider an arbitrary sample path of demand when the policy π is applied during $(t_{\ell-1},t_\ell]$. Then either none of the physical goods within the set $\mathbb{A}_\ell^k$ runs out of its quota, or some physical good runs out of its quota before time t_ℓ. We define

$$\Gamma_{\ell i}^k := \inf\Big\{t : N_{\ell i}^k(t_{\ell-1},t) = \lfloor k\mathcal{L}_{(\ell-1)i}\rfloor - \lfloor k\mathcal{L}_{\ell i}\rfloor\Big\},\ i\in\mathbb{A}_\ell^k.$$

Then $\Gamma_{\ell i}^k$ is the time that the inventory level of physical good i first drops to $\lfloor k\mathcal{L}_{\ell i}\rfloor$ from $\lfloor k\mathcal{L}_{(\ell-1)i}\rfloor$.

Naturally, if each physical good $i\in\mathbb{A}_\ell^k$ has an unlimited quantity to sell, the expected revenue from policy π during $(t_{\ell-1},t_\ell]$ would be

$$\begin{aligned}\mathbb{E}\Big\{\sum_{i\in\mathbb{A}_\ell^k}\Big[p_i(S_\ell)N_{\ell i}^k(t_{\ell-1},t_\ell)\Big]\Big\} &= \sum_{i=1}^m\Big(p_i(S_\ell)\tau^*(S_\ell)k\mu_i(S_\ell)\Big)\\ = \quad k\lambda\sum_{i\in\mathcal{M}}\sum_{j\in\mathcal{N}}\Big(\tau^*(S_\ell)P_j(S_\ell)\hat{\theta}_{ij}^*r_j\Big) &= k\lambda\sum_{j\in\mathcal{N}}\Big(\tau^*(S_\ell)r_jP_j(S_\ell)\Big),\end{aligned}$$

where we use the fact that $\mu_i(S_\ell)=0\ \forall i\notin\mathbb{A}_\ell^k$. However, given the inventory quota, we know that the actual expected revenue during $(t_{\ell-1},t_\ell]$ is less than the value above. $\forall i\in\mathbb{A}_\ell^k$, define

$$\Delta R_{\ell i} \quad := \quad \mathbb{E}\Big\{\sum_{j\in\mathbb{A}_\ell^k}\Big[\bar{p}(S_\ell)N_{\ell j}^k(\Gamma_{\ell i}^k,t_\ell)\mathbb{I}_{\{\Gamma_{\ell i}^k\le t_\ell,\Gamma_{\ell i}^k\le\Gamma_{\ell i'}^k,\ \forall i'\in\mathbb{A}_\ell^k,i'\neq i\}}\Big]\Big\},$$

which is an upper bound for the expected revenue during $(t_{\ell-1},t_\ell]$ after physical good $i\in\mathbb{A}_\ell^k$ first runs out of its inventory quota. Since the inventory level at time t_ℓ is at least $\lfloor k\mathcal{L}_\ell\rfloor$, we have

$$V_\pi^k(t_{\ell-1},\lfloor k\mathcal{L}_{\ell-1}\rfloor)\ge k\lambda\sum_{j\in\mathcal{N}}\Big(\tau^*(S_\ell)r_jP_j(S_\ell)\Big)-\sum_{i\in\mathbb{A}_\ell^k}\Delta R_{\ell i}+V_\pi^k(t_\ell,\lfloor k\mathcal{L}_\ell\rfloor). \tag{24}$$

Notice that (24) establishes a lower bound for the revenue function $V_\pi^k(t_{\ell-1},\lfloor k\mathcal{L}_{\ell-1}\rfloor)$. We immediately have the following Proposition.

Proposition 3. *For any ℓ, $\ell=1,2,\ldots,K$, we have*

$$\lim_{k\to\infty}\Big\{\frac{V_\pi^k(t_{\ell-1},\lfloor k\mathcal{L}_{\ell-1}\rfloor)-V_\pi^k(t_\ell,\lfloor k\mathcal{L}_\ell\rfloor)}{k}\Big\}\ge\lambda\sum_{j\in\mathcal{N}}\Big(\tau^*(S_\ell)r_jP_j(S_\ell)\Big). \tag{25}$$

The lower bound for the revenue V_π^k can be obtained by adding up all of the K inequalities shown in (25) together. Using this lower bound, we can then construct a lower bound for the fluid gap, and establish the major result about the asymptotic optimality of the fluid policy in Theorem 2.

Theorem 2. *If policy π is adopted for the k-scaled problem, the fluid gap is of lower order k; that is,*

$$\lim_{k\to\infty}\left\{\frac{R^{k*}-V_{\pi}^{k}}{k}\right\}=0. \tag{26}$$

Consequently,

$$\lim_{k\to\infty}\left\{\frac{V^{k*}-V_{\pi}^{k}}{V^{k*}}\right\}=0. \tag{27}$$

Note that inequality (27) follows directly from (26) because the fluid gap dominates the revenue loss. Inequality (27) implies that the relative revenue loss (compared with the MDP revenue) when applying the fluid policy π, approaches zero when the scale of the problem k is sufficiently large. In other words, the seller can almost achieve the optimal revenue using the proposed heuristic fluid policy for systems with large demand rates and inventory levels. The result implies that when there is a sufficiently high inventory level for each physical good and a sufficiently large number of arriving customers, the benefit of dynamic control diminishes.

Before concluding this section, we remark that, based on the optimal solution τ^* of the fluid control problem, one may develop other heuristic policies as well. For example, motivated by the PAC policy of Jasin and Kumar (2012), one may develop a probabilistic offering control (POC) policy as follows. At any time, the seller is suggested to offer each feasible set S with probability $\tau^*(S)/T$. In a POC policy, the sets are offered more uniformly over time and may lower the variance of the collected revenue. Following a similar analysis, we can show that the POC policy is also asymptotically optimal when the scale of the problem k is sufficiently large.

5 A Decomposition Approximation

Although the fluid policy π is asymptotically optimal, it is crude and may not perform well for small to medium-sized problems that are far from the asymptotic regime. In particular, the fluid policy is static and lacks the flexibility to match supply with demand dynamically. It also does not specify the order in which each set S is offered. One can attempt to resolve the fluid control problem frequently to adjust to changes in remaining inventory and remaining time; however, this is computationally complex when the number of substitutable products is large. To overcome these problems, we propose a decomposition approximation of the problem to generate a dynamic control policy.

Consider the fluid problem (17) – (21). Let $w^*=(w_1^*,w_2^*,\ldots,w_m^*)$ be the vector of dual prices associated with the inventory constraint (18). Intuitively, w^* provides an estimate of the marginal value of the inventories.

We follow the dynamic programming decomposition approach proposed in the literature (i.e., Chapter 3 of Talluri and van Ryzin, 2004a; Liu and van Ryzin, 2008; and Zhang, 2011) and decompose the multi-dimensional DP into a collection of single-dimensional DPs. Specifically, let

$$V(t, y) \approx V_i(t, y_i) + \sum_{l \in \mathcal{M}, l \neq i} w_l^* y_l, \quad \forall i \in \mathcal{M}, \tag{28}$$

Then, for any physical good i with $y_i > 0$, we have the following estimate for the marginal expected revenue:

$$V(t, y) - V(t, y - e_i) \approx V_i(t, y_i) - V_i(t, y_i - 1) := \Delta V_i(t, y_i),$$

and for any $l \neq i$, $l \in \mathcal{M}$, we have

$$V(t, y) - V(t, y - e_l) \approx w_l^*.$$

Note that $V_i(t, y_i)$ is a nonlinear approximation of the value of physical good i and $w_\ell^* y_\ell$ is a linear approximation of the value of physical good ℓ. Applying the approximation (28) on the Hamilton-Jacobi equation (6) leads to

$$-\frac{\partial V_i(t, y_i)}{\partial t} = \max_{S \subseteq \mathcal{N}} \left\{ \sum_{j \in \mathcal{N}} \lambda P_j(S) \Big[r_j - \min \Big(\Delta V_i(t, y_i) | a_{ij} = 1, \min_{l \neq i, l \in \mathcal{M}, a_{\ell j} = 1} w_l^* \Big) \Big] \right\}, \quad \forall y_i \geq 1 \tag{29}$$

with boundary conditions

$$V_i(T, y_i) = 0, \ \ \forall y_i \geq 0,$$
$$V_i(t, 0) = 0, \ \ \forall t \in [0, T].$$

Note that (29) is a one-dimensional MDP and can be solved efficiently. We apply this approximation for each physical good and obtain a set of value functions $V_i(t, y_i)$, $i \in \mathcal{M}$.

Results in the linear programming based approximation dynamic programming literature imply that (28) gives an upper bound to the optimal total expected revenue for each i. This stems from the fact that (28) is a feasible solution of the linear programming formulation of the dynamic program; see e.g., Proposition 2 in Zhang and Adelman (2009). We summarize this result in the following proposition without proof.

Proposition 4. *For each physical good $i \in \mathcal{M}$, we have*

$$R^* \geq V_i(0, Q_i) + \sum_{l \in \mathcal{M}, l \neq i} w_l^* Q_l \geq V_0(Q).$$

It follows that $R^* \geq \bar{R}_{ADP} \geq V_0(Q)$, where

$$\bar{R}_{ADP} := \min_{i \in \mathcal{M}} \left\{ V_i(0, Q_i) + \sum_{l \in \mathcal{M}, l \neq i} w_l^* Q_l \right\}.$$

Proposition 4 shows that the decomposition heuristic gives an upper bound for each physical good $i \in \mathcal{M}$, which is tighter than the upper bound R^* from the fluid approximation. Taking a minimum over i gives the smallest upper bound, $\bar{R}_{ADP}$, which we call the decomposition bound in our numerical study. Proposition 2(ii) in Zhang and Adelman (2009) states the result for the choice-based network revenue management problem without opaque selling. Still, the same proof can be applied to show the result of Proposition 4.

These approximate value functions can be further combined to form an approximation of the MDP revenue function $V(t, y)$ as

$$V(t, y) \approx \sum_{i \in \mathcal{M}} V_i(t, y_i), \quad \forall t \in [0, T], y \in [0, Q].$$

Using this additive approximation, we propose the decomposition policy below to control the sets being offered over time.

Definition 2. (Decomposition Policy) *If the inventory level is y at time t, the seller offers set S^*, which is obtained by solving the following optimization problem:*

$$S^* = \underset{S \subseteq \mathbb{B}(y)}{\operatorname{argmax}} \left\{ \sum_{j \in S} \lambda P_j(S) \Big[r_j - \min_{i \in \mathcal{M}, a_{ij}=1} \Big(\Delta V_i(t, y_i) \Big) \Big] \right\}. \tag{30}$$

On the fulfillment side, assign product i^* to fulfill the demand of product j:

$$i^* = \underset{i \in \mathcal{M}, a_{ij}=1, y_i > 0}{\operatorname{argmin}} \Big\{ \Delta V_i(t, y_i) \Big\}. \tag{31}$$

In the following section, we will evaluate the performance of the decomposition policy by conducting numerical experiments.

6 Numerical Studies

In this section, we numerically evaluate the performance of the fluid policy and the decomposition policy in order to show the profit potential of opaque selling. For ease of illustration, we consider the Salop Circular Choice Model to characterize the choice behavior of customers. Unlike the original Salop's circle model (Salop, 1979) where customers' valuations are proportional to their arc lengths to the products, we assume that customers are uniformly distributed on the interior of the circle. So customers are distributed over a disc instead of along the circle. This allows us to introduce additional heterogeneity among the customer population. Specifically, the m physical goods are located, uniformly, over a circle of unit radius; see Figure 1 for an illustration with $m = 4$ physical goods where physical goods are represented by points A, B, C, and D. Each potential customer is represented by a point that is uniformly distributed within the circle. For a customer located at point P, her utility for a physical good $i \in \{A, B, C, D\}$ is

$$U_i(P) = v_i - t|Pi|,$$

where v_i is the maximal utility of physical good i, t is the unit transportation cost, and $|Pi|$ is the distance between P and i.

We assume customers are risk-neutral towards the uncertain fulfillment. In addition, if a customer chooses to buy an OG, she anticipates to receive each physical good contained in the OG with an equal probability. For example, the customer P's expected surplus from buying an OG j that contains physical goods A, B, and C is

$$V_j(P) = \sum_{i \in \{A,B,C\}} \frac{1}{3}(v_i - t|Pi|) - r_j,$$

where r_j is the price of OG j. Given an offer set, each customer always chooses the product that provides the highest expected surplus. Of course, if the net expected surpluses of all the products are negative, the customer simply leaves empty-handed.

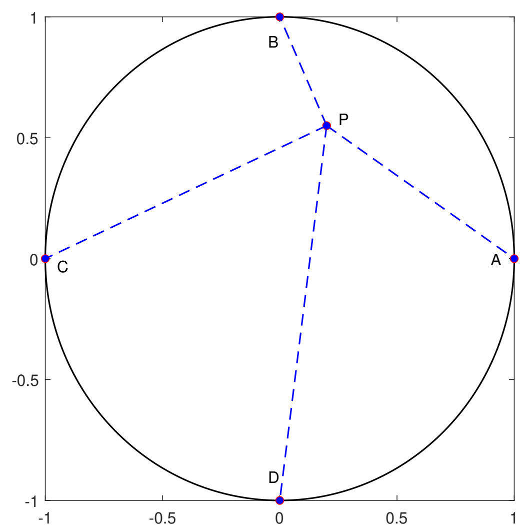

Figure 1: The Salop Circular Choice Model with $m = 4$

Note that it is quite difficult to analytically calculate the choice probability $P_j(S)$. Therefore, we compute choice probabilities via Monte Carlo simulation. Our numerical studies consist of two sets of experiments. In the first set, we consider $m = 3$ physical goods and multiple OGs. In the second set, we consider multiple physical goods $m \geq 3$ but a single OG only.

6.1 Multiple OGs ($m = 3$, $n = 7$)

In this subsection, we consider a problem with three physical goods and seven products (i.e., $m = 3$ and $n = 7$). Note that any combination of physical goods is considered a product, and the number of possible offer sets is $2^n = 128$. The base parameters are set as follows. The selling period is normalized to $T = 1$ and the arrival rate of potential customers is set at $\lambda = 400$. The unit transportation cost is set at $t = 0.8$. The initial inventory level of physical goods, the selling price of each product, and each physical good's maximal utility v_j, are all listed in Table 1. Note that we first consider a symmetric scenario in which the physical goods are identical in initial inventory level and maximal utility. In the pricing scheme, motivated by the pricing mechanism of Eurowings, the price is decreased by 0.2 for each additional physical good added to an OG.

Table 1: Base Parameters with $m = 3$ Physical Goods and $n = 7$ Products

Physical Good, i	Product, j							Q_i	v_i
	1	2	3	4	5	6	7		
1	1	0	0	1	1	0	1	60	2
2	0	1	0	1	0	1	1	60	2
3	0	0	1	0	1	1	1	60	2
Price, r_j	1.5	1.5	1.5	1.3	1.3	1.3	1.1		

We conduct a simulation study to evaluate the performance of regular selling (with no OGs) and opaque selling under different demand loads. For the fluid policy, we solve the fluid control problem (17)–(21) and offer each set S for a duration $\tau^*(S)$. For the decomposition policy, we use the dual bid-prices from the linear program (17)–(21), and calculate the decomposed value functions following (29). A time and inventory dependent control policy is computed following (30). We compute the mean and standard deviation (Std.) for each policy over 1000 simulations. To reduce the variance for policy comparison, the same set of randomly-generated customer arrivals are used to simulate different policies. The mean and standard deviation of the seller's revenue under regular selling and the two policies are reported in Table 2. The optimal fluid revenue R^*, the decomposition bound $\bar{R}_{ADP}$, and the relative revenue improvement of the fluid and decomposition policies over regular selling are also reported in the table. Note that the improvement (i.e., "Impv.") column is measured based on the mean revenues.

Under the parameters for our simulation study, roughly 12% of arriving customers will choose to buy each physical good under regular selling. Therefore, when $\lambda \geq 500$, almost all the physical goods can be sold at the full price during the selling horizon under the fluid control problem. Therefore, one expects that offering OGs adds little value when $\lambda \geq 500$. Table 2 confirms this intuition: the revenue improvement from opaque selling under both fluid policy and decomposition policy is more significant when demand is relatively low. This should be contrasted with the numerical findings in the choice-based network revenue management literature. For example, Liu and van Ryzin (2008) show that the decomposition heuristic performs better when the load factor

Table 2: Simulation Results for Different Size of Potential Customers ($m = 3$, $n = 7$)

λ	R^*	$\bar{R}_{ADP}$	Regular Selling		Fluid Policy			Decomposition Policy		
			Mean	Std.	Mean	Std.	Impv.	Mean	Std.	Impv.
300	243.4	242.9	159.3	14.0	222.2	11.0	39.44%	227.4	11.9	42.73%
310	245.5	245.0	164.7	14.2	222.8	11.6	35.28%	231.9	11.9	40.79%
320	247.6	247.1	170.1	15.2	223.5	12.1	31.39%	236.4	11.6	38.96%
330	249.7	249.3	175.0	14.9	225.0	12.0	28.54%	240.7	10.4	37.52%
340	251.9	251.4	180.7	15.9	228.7	13.6	26.52%	244.4	10.1	35.21%
350	254.0	253.5	187.1	15.1	228.5	12.9	22.14%	248.5	8.7	32.84%
360	256.1	255.6	192.5	16.2	230.3	12.4	19.62%	251.2	7.7	30.52%
370	258.2	257.7	196.5	16.1	231.5	11.6	17.83%	253.1	6.4	28.79%
380	260.4	259.8	201.7	16.0	234.0	11.4	16.01%	255.1	5.2	26.44%
390	261.7	260.9	208.5	15.8	246.5	10.6	18.21%	256.8	5.8	23.13%
400	262.4	261.7	212.6	15.5	246.1	10.9	15.75%	257.9	4.8	21.33%
410	263.2	262.4	217.5	15.8	243.5	11.0	11.95%	259.0	4.1	19.06%
420	263.9	263.2	222.8	15.8	248.9	10.8	11.69%	259.2	5.0	16.32%
430	264.6	263.9	225.9	15.9	247.6	11.2	9.60%	259.8	5.0	14.98%
440	265.4	264.6	232.2	14.6	247.0	11.0	6.38%	261.2	4.0	12.49%
450	266.1	265.3	236.3	14.9	248.3	11.4	5.07%	261.8	4.2	10.78%
460	266.8	266.0	240.9	14.6	250.6	11.2	4.04%	262.8	3.8	9.07%
470	267.5	266.7	245.3	13.2	252.4	10.5	2.90%	263.7	3.2	7.54%
480	268.3	267.3	248.1	12.8	252.6	11.3	1.78%	264.3	3.2	6.50%
490	269.0	268.0	251.6	12.5	254.0	11.5	0.98%	264.9	3.1	5.29%
500	269.7	268.6	254.7	11.4	255.1	11.0	0.18%	265.4	3.1	4.20%

is high. Our result can be explained as follows. With relatively low demand, the seller should try to sell to each arriving customer by offering OGs. On the contrary, when the demand is high, the supply becomes tight. Then the seller should sell each product at a high price. Since an OG is sold at a discount, one can expect that fewer OGs should be offered by the seller when the demand is high.

Table 2 also uncovers two interesting observations regarding the relative performance of the fluid policy and the decomposition policy. On the one hand, when the demand load is low, opaque selling under either the fluid policy or the decomposition policy leads to a significant improvement over regular selling, which stems from the demand induction effect. On the other hand, Table 2 shows that when the demand is medium, the decomposition policy leads to a significant improvement over the fluid policy. For example, when $\lambda = 370$, the decomposition policy can improve revenue by 10.96%. However, the advantage of the decomposition policy decreases when the demand load is low or high. For example, when $\lambda = 300$, the decomposition policy only gains 3.29% improvement over the fluid policy. Furthermore, Table 2 shows that when the demand load is high, the standard deviation of simulated revenue under the decomposition policy is much lower than that under regular selling and opaque selling with a fluid policy. This implies that the decomposition policy performs quite robustly when the demand is relatively high.

Next, we investigate how the degree of inventory imbalance among different physical goods affects the performance of the fluid policy and the decomposition policy. To do so, we fix $\lambda = 450$ and the inventory level of physical good 2 at $Q_2 = 60$. We fix the total inventory level of physical goods 1 and 3 at $Q_1 + Q_3 = 120$ and alter the value of (Q_1, Q_3). Different inventory levels for physical goods 1 and 3 represent different degrees of inventory imbalance. For example, an extremely low (or high) value of Q_1 implies that physical good 1 is understocked (or overstocked) and physical good 3 is overstocked (or understocked). Because the two physical goods are symmetric, we only report the results for $Q_1 \leq Q_3$ in Table 3. We observe that the optimal fluid revenue, the mean revenue from regular selling, the mean revenue from the fluid policy, and the mean revenue from the decomposition policy all increase in Q_1 when $Q_1 \leq Q_3$. This implies that the seller in general earns lower profits when the inventory levels for physical goods become more imbalanced. Note that when $Q = (60, 60, 60)$, opaque selling can only generate a small incremental revenue. However, Table 3 shows that the fluid policy and the decomposition policy can improve the seller's revenue substantially when the inventory levels are more imbalanced. This can be attributed to the fulfillment flexibility rendered by opaque selling. In particular, the seller has the flexibility to fulfill the demand of opaque goods using the physical good that has a high inventory. Furthermore, the decomposition policy outperforms the fluid policy quite significantly in all the experiments. For example, when $Q = (10, 60, 110)$, the decomposition policy can improve revenue from the fluid policy by 15.24%, which is much higher than the improvement of 5.72% with $Q = (60, 60, 60)$. This confirms that the decomposition policy adds more value compared with the fluid policy when there is a higher degree of inventory imbalance.

Table 3: Simulation Results for Different Inventory Levels

Q	R^*	$\bar{R}_{ADP}$	Regular Selling		Fluid Policy			Decomposition Policy		
			Mean	Std.	Mean	Std.	Impv.	Mean	Std.	Impv.
(10,60,110)	253.9	253.1	173.0	13.7	224.1	15.2	29.54%	250.5	4.4	44.79%
(15,60,105)	255.9	255.1	181.6	13.9	231.4	14.4	27.43%	252.7	4.2	39.18%
(20,60,100)	258.0	257.1	188.4	13.7	236.1	13.5	25.33%	254.5	4.0	35.09%
(25,60,95)	260.0	259.1	196.6	13.4	241.2	12.3	22.70%	256.5	3.4	30.47%
(30,60,90)	261.4	260.7	204.7	13.5	242.5	12.4	18.46%	257.0	4.8	25.56%
(35,60,85)	262.4	261.8	210.9	13.7	243.2	12.0	15.31%	258.4	4.2	22.52%
(40,60,80)	263.4	262.8	218.4	13.8	243.0	12.5	11.25%	260.0	3.6	19.01%
(45,60,75)	264.4	263.8	225.0	13.7	244.9	11.8	8.86%	260.7	3.6	15.90%
(50,60,70)	265.4	264.7	231.9	13.8	251.1	10.6	8.30%	261.3	4.0	12.71%
(55,60,65)	266.1	265.2	235.8	14.7	249.8	11.4	5.96%	261.9	3.9	11.11%
(60,60,60)	266.1	265.3	236.3	14.9	248.3	11.4	5.07%	261.8	4.2	10.78%

6.2 A Single OG (i.e., $n = m + 1$)

The fluid solution suggests that only a limited number of offer sets should be offered in the fluid control problem, even if the number of possible offer sets is large. This intuition is quite useful when the number of products is large, because the number of offer sets is given by 2^n, which grows exponentially in n. This motivates us to evaluate the performance of opaque selling with only a limited number of OGs. In the second set of experiments, besides selling the physical goods, we consider a single OG, which consists of all the physical goods. That is, the number of products is $n = m + 1$.

We first consider $m = 3$ physical goods, with all the parameters being the same as those of Table 1, except that products $j = 4, 5, 6$ are excluded. The simulation results with different λ values are reported in Table 4. The observations are almost identical to those of Table 2 with $n = 7$ products. That is: (i) opaque selling under either the fluid policy or the decomposition policy can lead to a significant improvement over regular selling, especially when the demand load is low, and (ii) the decomposition policy may outperform the fluid policy significantly. Interestingly, even though only a single OG is allowed, the performance of opaque selling is only moderately worse than the case with four OGs. For example, when $\lambda = 400$, restricting to a single OG, reduces revenue by 3.85% (from 246.1 to 236.6) under the fluid policy, and by 3.94% (from 257.9 to 250.3) under the decomposition policy, compared with considering the full set of OGs.

Table 4: Simulation Results for Different Size of Potential Customers ($m = 3$, $n = 4$)

λ	R^*	$\bar{R}_{ADP}$	Regular Selling		Fluid Policy			Decomposition Policy		
			Mean	Std.	Mean	Std.	Impv.	Mean	Std.	Impv.
300	237.5	237.3	159.7	15.3	220.9	11.0	38.33%	235.0	4.4	47.15%
310	239.1	238.9	165.9	14.5	222.9	10.6	34.34%	236.8	4.2	42.70%
320	240.6	240.4	170.9	14.9	224.2	10.8	31.21%	238.1	4.3	39.38%
330	242.1	242.0	174.9	15.0	225.7	10.5	29.04%	239.3	4.3	36.84%
340	243.7	243.6	180.7	14.8	227.3	10.2	25.74%	241.0	4.2	33.33%
350	245.2	245.2	186.3	15.8	228.9	11.0	22.86%	242.5	4.4	30.21%
360	246.7	246.7	193.1	15.2	231.4	10.4	19.85%	244.6	4.2	26.68%
370	248.3	248.3	197.2	16.4	232.0	10.9	17.67%	245.6	4.5	24.56%
380	249.8	249.9	201.9	16.1	233.3	11.2	15.52%	247.0	4.6	22.34%
390	251.4	251.4	207.4	15.8	234.9	11.4	13.23%	248.7	4.3	19.92%
400	252.9	253.0	213.2	15.7	236.6	10.7	10.99%	250.3	4.4	17.44%
410	254.4	254.6	217.3	15.4	238.2	10.8	9.62%	251.6	4.4	15.81%
420	256.0	256.1	223.5	15.6	240.6	11.1	7.62%	253.4	4.4	13.36%
430	257.5	257.7	228.2	15.3	242.0	10.8	6.08%	254.9	4.4	11.73%
440	259.0	259.2	232.1	14.8	243.2	10.8	4.79%	256.2	4.2	10.36%
450	260.6	260.8	235.6	14.4	244.5	11.2	3.76%	257.2	4.3	9.16%
460	262.1	262.4	240.3	13.8	246.2	11.0	2.46%	258.7	4.2	7.67%
470	263.7	263.9	244.8	13.6	248.7	11.3	1.58%	260.1	4.1	6.25%
480	265.2	265.4	247.7	13.1	250.1	11.4	0.95%	261.0	3.9	5.38%
490	266.7	267.0	250.9	12.7	252.3	11.5	0.55%	262.0	3.8	4.41%
500	268.3	268.5	254.2	11.5	254.7	11.0	0.20%	263.0	3.6	3.45%

Next, we vary the number of physical goods. The total inventory is fixed at 180 and the inventory level of each physical good is taken to be $180/m$. The regular selling price of each physical good is set at 1.5, and the OG is priced at 1.1. All the other parameters are the same as those in subsection 6.1. A sample of the results are reported in Table 5. The numerical findings are similar to those of Table 2: opaque selling shows a larger advantage than regular selling when the demand is relatively low; and the decomposition policy offers a significant improvement over the fluid policy when the demand is relatively low. Table 5 also shows that the value of opaque selling diminishes when the number of physical goods increases. For example, when $m = 6$, the fluid policy and the decomposition policy show only a small profit improvement over regular selling for $\lambda = 300$, whereas the improvements under the fluid policy and the decomposition policy are both above 38% for $m = 3$ physical goods. This result can be partially explained by the Salop Circular Choice Model considered in our study. As more physical goods are added to the circle, the seller provides more options for customers to choose from. Therefore, it is less likely for a customer to leave empty-handed under regular selling. Consequently, the demand induction effect from selling OGs diminishes and the incremental revenue from opaque selling decreases. Despite of this effect, Table 5 shows that selling a single OG can still improve revenue substantially when the demand is relatively low.

Table 5: Simulation Results for Different Number of Physical Goods

m	λ	R^*	$\bar{R}_{ADP}$	Regular Selling		Fluid Policy			Decomposition Policy		
				Mean	Std.	Mean	Std.	Impv.	Mean	Std.	Impv.
3	260	231.4	231.0	137.5	13.2	214.2	10.4	55.80%	228.7	4.0	66.3%
3	280	234.5	234.2	149.5	14.0	217.7	10.5	45.61%	232.0	4.2	55.1%
3	300	237.5	237.3	159.7	15.3	220.90	11.0	38.33%	235.0	4.4	47.15%
4	260	244.6	244.1	183.8	15.2	223.7	10.9	21.67%	238.7	6.1	29.9%
4	280	248.8	248.4	198.7	16.2	228.8	11.1	15.14%	243.2	6.5	22.4%
4	300	253.1	252.6	212.3	15.8	232.9	11.5	9.69%	247.6	6.2	16.7%
5	260	254.9	254.4	220.2	14.9	233.1	11.8	5.90%	248.3	6.3	12.8%
5	280	260.3	259.7	234.2	13.8	238.9	11.9	2.02%	253.6	5.1	8.3%
5	300	265.6	265.0	245.6	12.0	246.3	11.4	0.30%	257.6	4.0	4.9%
6	260	258.0	257.6	227.2	13.7	234.1	11.6	3.05%	250.1	5.0	10.1%
6	280	263.7	263.3	240.9	12.8	242.2	11.9	0.52%	254.6	4.2	5.7%
6	300	269.1	267.9	250.9	11.2	251.0	11.2	0.02%	257.6	6.4	2.7%

7 Concluding Remarks

This paper studies opaque selling in a revenue management setting where a seller offers multiple substitutable products in a finite horizon. In addition to selling each physical good, the seller creates virtual products, called OGs, to offer over time. The seller dynamically controls the offer set and fulfillment of OGs over time to maximize the expected total revenue. Because the optimal control policy is too complex to characterize, we study the corresponding fluid control problem and develop a time-based fluid policy. We show that this heuristic policy is asymptotically optimal; that is, its relative revenue loss approaches zero as both demand rate and inventory become sufficiently large. In order to adjust the control policy with the remaining inventory and remaining selling time, we propose a decomposition method where the original MDP is decomposed into a collection of single-dimensional MDPs. We show that decomposition method provides a tighter upper bound than the fluid control problem for the MDP problem. Our numerical results show that opaque selling with either the fluid policy or the decomposition policy, has great potential to improve the seller's revenue, especially when the demand is relatively low and/or when the seller has imbalanced inventory for the physical goods. Particularly, the decomposition policy performs much better (in terms of the expected revenue and the standard deviation) than the fluid policy when the demand is medium and low. Importantly, this benefit can be achieved with a relatively small collection of OGs, potentially reducing the burden of implementing the opaque selling strategy.

Product design is typically not considered as a part of revenue management decision, which tends to be tactical in nature. In the vast majority of revenue management literature, products are exogenously determined. When opaque selling is considered, a large number of virtual products can be offered. Our work shows that designing and controlling the virtual products can bring great benefits to sellers. Creating virtual products is an example of innovative selling strategies with outsized revenue gains. Another example is flexible products (Gallego and Phillips, 2004, Petrick et al., 2012). The main difference between flexible products and OGs is the timing of product fulfillment decision; in our model, it is determined immediately after purchase, while product assignment for flexible products is determined later (e.g., at the end of the booking horizon). We hope to see more work along this line in the future.

There are several future research directions. First, prices of OGs are given in our work. It is an interesting direction to study the dynamic pricing decision with endogenous fulfillment. Second, our numerical experiments show that opaque selling with a single OG can significantly boost the revenue under certain conditions. However, whether a small set of OGs suffices remains a question to be investigated. Third, consumer behavior modeling is an important aspect of the literature on opaque and/or probabilistic selling. How to incorporate various consumer behavior such as strategic behavior (e.g., Jerath et al., 2009), bounded rationality (e.g., Huang and Yu, 2014) and collaboration among strategic consumers (e.g., Levina et al., 2015), into our work is another fruitful avenue of future research. Furthermore, the decomposition heuristic considered in our paper is motivated by Zhang and Adelman (2009). We point out that there are several related approaches to solving high-dimensional dynamic programs, such as Lagrangian relaxation (Adelman and Mersereau, 2008; Topaloglu, 2009; Kunnumkal and Talluri, 2016). Even though we do not consider Lagrangian relaxation in the current paper, it is a direction that deserves investigation to tackle the curse of dimensionality for the problem in this paper.

Acknowledgement

We are grateful to the the Associate Editor and three anonymous referees for their valuable comments that have helped improve the paper significantly. The first author is supported by the Hong Kong Research Grants Council (Grant No. 16206618). The second author is partially supported by the National Natural Science Foundation of China (Grant No. 72125002) and National Key R&D Program of China (Grant No. 2020AAA0103800).

References

[1] Adelman, D., and A.J. Mersereau. 2008. Relaxations of Weakly Coupled Stochastic Dynamic Programs. *Operations Research.* 56(3): 712-727.

[2] Anderson, S.P. and L. Celik. 2019. Opaque Selling. Working paper, University of Virginia. Available at *https://economics.virginia.edu/sites/economics.virginia.edu/files/anderson/Opaque190515.pdf.*

[3] Anderson, C.K. and X. Xie. 2014. Pricing and Market Segmentation Using Opaque Selling Mechanisms. *European Journal of Operational Research.* 233(1): 263-272.

[4] Bront, J.M., I. Méndez-Díaz, and G. Vulcano. 2009. A Column Generation Algorithm for Choice-Based Network Revenue Management. *Operations Research.* 57(3): 769-784.

[5] Cai, G., X. Chao, and J. Li. 2009. Optimal Reserve Prices in Name-Your-Own-Price Auctions with Bidding and Channel Options. *Production and Operations Management.* 18(6): 653-671.

[6] Chaneton, J.M., and G. Vulcano. 2011. Computing Bid Prices for Revenue Management Under Customer Choice Behavior. *Manufacturing & Service Operations Management.* 13(4): 452-470.

[7] Chen, M. and Z.L. Chen. 2018. Robust Dynamic Pricing with Two Substitutable Products. *Manufacturing & Service Operations Management.* 20(2): 249-268.

[8] Chen, R.R., E. Gal-Or, and P. Roma. 2014. Opaque Distribution Channels for Competing Service Providers: Posted Price vs. Name-Your-Own-Price Mechanisms. *Operations research.* 62(4): 733-750.

[9] Chen, X., H. Peng, and Z. Hu. 2017. Efficient Algorithms for the Dynamic Pricing Problem with Reference Price Effect. *Management Science.* 63(12): 4389-4408.

[10] Chen, X., I. Diakonikolas, A. Orfanou, D. Paparas, X. Sun, and M. Yannakakis. 2015. On the Complexity of Optimal Lottery Pricing and Randomized Mechanisms. *Working Paper.* Columbia University, New York.

[11] Cooper, W.L. 2002. Asymptotic Behavior of An Allocation Policy for Revenue Management. *Operations Research.* 50(4): 720-727.

[12] Elmachtoub, A.N. and M.L. Hamilton. 2017. The Power of Opaque Products in Pricing. *Working Paper*, Department of Industrial Engineering and Operations Research, Columbia University, New York, NY, 10027.

[13] Fay, S. 2008. Selling an Opaque Product through an Intermediary: The Case of Disguising One's Product. *Journal of Retailing*. 84(1): 59-75.

[14] Fay, S. and J. Xie. 2008. Probabilistic Goods: An Innovative Way of Selling Products and Services. *Marketing Science*. 27(4): 674-690.

[15] Fay, S. and J. Xie. 2010. The Economics of Buyer Uncertainty: Advance Selling vs. Probabilistic Selling. *Marketing Science*. 29(6): 1040-1057.

[16] Fay, S. and J. Xie. 2015. Timing of Product Allocation: Using Probabilistic Selling to Enhance Inventory Management. *Management Science*. 61(2): 474-484.

[17] Fay, S., J. Xie, and C. Feng. 2015. The Effect of Probabilistic Selling on the Optimal Product Mix. *Journal of Retailing*. 91(3): 451-467.

[18] Feldman, J.B. and H. Topaloglu. 2017. Revenue Management under the Markov Chain Choice Model. *Operations Research*. 65(5): 1322-1342.

[19] Feng, Y.Y. and G. Gallego. 1995. Optimal Starting Times for End-of-Season Sales and Optimal Stopping Times for Promotional Fares. *Management Science*. 41(8): 1371-1391.

[20] Gallego, G. and R. Phillips. 2004. Revenue Management of Flexible Products. *Manufacturing & Service Operations Management*. 6(4): 321-337.

[21] Gauthier, S. and G. Laroque. 2017. Redistribution by means of lotteries. *Journal of Economic Policy*. 169: 707-716.

[22] Hu, Z., X. Chen, and P. Hu. 2016. Dynamic Pricing with Gain-Seeking Reference Price Effects. *Operations Research*. 64(1): 150-157.

[23] Huang, X., G. Sosic., and G. Kersten. 2017. Selling Through Priceline? On the Impact of Name-Your-Own-Price in Competitive Market. *IIE Transactions*. 49(3): 304-319.

[24] Huang, T. and Y. Yu. 2014. Sell Probabilistic Goods? A Behavioral Explanation for Opaque Selling. *Marketing Science*. 33(5): 743-759.

[25] Jasin, S. and S. Kumar. 2012. A Re-Solving Heuristic with Bounded Revenue Loss for Network Revenue Management with Customer Choice. *Mathematics of Operations Research*. 37(2): 313-345.

[26] Jerath, K., S. Netessine, and S.K. Veeraraghavan. 2009. Selling to Strategic Customers: Opaque Selling Strategies. In: S. Netessine and C.S. Tang (eds), *Consumer-Driven Demand and Operations Management Models*, International Series in Operations Research & Management Science. Springer Science+Business Media. Volume 131, Part 3, 253-300.

[27] Jerath, K., S. Netessine, and S.K. Veeraraghavan. 2010. Revenue Management with Strategic Customers: Last-minute Selling and Opaque Selling. *Management Science*. 56(3): 430-448.

[28] Jiang, Y. 2007. Price Discrimination with Opaque Products. *Journal of Revenue and Pricing Management*. 6(2): 118-134.

[29] Kubrusly, C.S. 2015. Essentials of Measure Theory. 1st Edition. Springer.

[30] Kunnumkal, S., and K. Talluri. 2016. On a Piecewise-Linear Approximation for Network Revenue Management. *Mathematics of Operations Research*. 41(1): 72-91.

[31] Kunnumkal, S. and H. Topaloglu. 2008. A Refined Deterministic Linear Program for the Network Revenue Management Problem with Customer Choice Behavior. *Naval Research Logistics*. 55(6): 563-580.

[32] Kunnumkal, S. and H. Topaloglu. 2010. A New Dynamic Programming Decomposition Method for the Network Revenue Management Problem with Customer Choice Behavior. *Production & Operations Management*. 19(5): 575-590.

[33] Levina, T., Y. Levin, J. McGill, and M. Nediak. 2015. Strategic Consumer Cooperation in a Name-Your-Own-Price Channel. *Production and Operations Management*. 24(12): 1883-1900.

[34] Li, Q., H. Xu, and P. Yu. 2016. Probabilistic Sales in a Variety Seeking Market. Working Paper. Hong Kong University of Science and Technology.

[35] Liu, Q. and G. van Ryzin. 2008. On the Choice-Based Linear Programming Model for Network Revenue Management. *Manufacturing & Service Operations Management*. 10(2): 288-310.

[36] Maglaras, C. and J. Meissner. 2006. Dynamic Pricing Strategies for Multiproduct Revenue Management Problems. *Manufacturing & Service Operations Management*. 8(2): 136-148.

[37] Meissner, J. and A.K. Strauss. 2012. Network Revenue Management with Inventory-sensitive Bid Prices and Customer Choice. *European Journal of Operational Research*. 216(2): 459-468.

[38] Petrick, A., C. Steinhardt, J. Gönsch, and R. Klein. 2012. Using Flexible Products to Cope with Demand Uncertainty in Revenue Management. *OR Spectrum*. 34(1): 215-242.

[39] Post, D. and M. Spann. 2012. Improving Airline Revenues with Variable Opaque Products: Blind Booking at Germanwings. *Interfaces*, 42(4): 329-338.

[40] Puterman, M.L. 1994. Markov Decision Processes : Discrete Stochastic Dynamic Programming. New York : John Wiley & Sons. New York, USA.

[41] Reiman, M.I. and Q. Wang. 2008. An Asymptotically Optimal Policy for a Quantity-based Network Revenue Management Problem. *Mathematics of Operations Research*. 33(2): 257-282.

[42] Rice, D. H., S. Fay, and J. Xie. 2014. Probabilistic Selling vs. Markdown Selling: Price Discrimination and Management of Demand Uncertainty in Retailing. *International Journal of Research in Marketing*. 31: 147-155.

[43] Salop, S.C. 1979. Monopolistic Competition with Outside Goods. *The Bell Journal of Economics*. 10(1): 141-156.

[44] Shapiro, D. and X. Shi. 2008. Market Segmentation: The Role of Opaque Travel Agencies. *Journal of Economics & Management Strategy*. 17(4): 803-837.

[45] Talluri, K.T. and G.J. van Ryzin. 2004a. The Theory and Practice of Revenue Management (International Series in Operations Research & Management Science, 68). Kluwer Academic Publishers, New York, USA.

[46] Talluri, K.T. and G.J. van Ryzin. 2004b. Revenue Management under a General Discrete Choice Model of Consumer Behavior. *Management Science*. 50(1): 15-33.

[47] Topaloglu, H. 2009. Using Lagrangian Relaxation to Compute Capacity-Dependent Bid Prices in Network Revenue Management. *Operations Research*. 57(3): 637-649.

[48] Xiao, Y.B. and J. Chen. 2014. Evaluating the Potential Effects from Probabilistic Selling of Seasonal Items. *Naval Research Logistics*. 61(8): 604-620.

[49] Zhang, D. 2011. An Improved Dynamic Programming Decomposition Approach for Network Revenue Management. *Manufacturing & Service Operations Management*. 13(1): 35-52.

[50] Zhang, D. and D. Adelman. 2009. An Approximate Dynamic Programming Approach to Network Revenue Management with Customer Choice. *Transportation Science*. 43(3): 381-394.

[51] Zhang, D. and W.L. Cooper. 2005. Revenue Management for Parallel Flights with Customer-Choice Behavior. *Operations Research*. 53(3): 415-431.

Appendix A: Calculation of the Expected Revenues in the Stylized Example

Under Policy 1: In the first period, all the physical goods A, B and C are offered.

1. If a type-1 (or type-2) customer arrives in the first period (with probability 0.1), the physical good A (or B) is sold and the revenue is 1. In the second period, the available physical goods B and C (or A and C) are offered and the expected revenue is 0.4.
2. If a type-3 customer arrives in the first period (with probability 0.3), the physical good C is sold and the revenue is 1. In the second period, the available physical goods A and B are offered and the expected revenue is 0.2.
3. If a type-4 customer arrives in the first period (with probability 0.5), no good is sold in the first period. In the second period, all the physical goods A, B and C are offered, and the expected revenue is 0.5.

Hence, the total expected revenue collected over two periods is:

$$0.1 \times (1 + 0.4) \times 2 + 0.3 \times (1 + 0.2) + 0.5 \times (0 + 0.5) = 0.89.$$

Under Policy 2: In the first period, all the physical goods A, B and C, and the OG are offered.

1. If a type-1 (or type-2) customer arrives in the first period (with probability 0.1), the OG is sold (either A or B is used to fulfill the demand) and the revenue is 0.3. In the second period, the available physical goods A and C (or B and C) are offered, and the expected revenue is 0.4.
2. If a type-3 customer arrives in the first period (with probability 0.3), the physical good C is sold and the revenue is 1. In the second period, the physical goods A, B and the OG are offered, and the expected revenue is $0.7 \times 0.3 = 0.21$.

3. If a type-4 customer arrives in the first period (with probability 0.5), either A or B is used to fulfill the demand. In the second period, the physical goods A and C (or B and C) are offered, and the expected revenue is 0.4.

Hence, the total expected revenue collected over two periods is:

$$0.1 \times (0.3 + 0.4) \times 2 + 0.3 \times (1 + 0.21) + 0.5 \times (0.3 + 0.4) = 0.853.$$

Under Policy 3: In the first period, all the physical goods A, B and C are offered.

1. If a type-1 (or type-2) customer arrives in the first period (with probability 0.1), the physical good A (or B) is sold and the revenue is 1. In the second period, the available physical goods B and C (or A and C) are offered and the expected revenue is 0.4.
2. If a type-3 customer arrives in the first period (with probability 0.3), the physical good C is sold and the revenue is 1. In the second period, the available physical goods A and B, and the OG are offered, and the expected revenue is $0.7 \times 0.3 = 0.21$.
3. If a type-4 customer arrives in the first period (with probability 0.5), no good is sold. In the second period, all the physical goods A, B, C and the OG are offered, and the expected revenue is $0.3 \times 1 + 0.7 \times 0.3 = 0.51$.

Hence, the total expected revenue collected over two periods is:

$$0.1 \times (1 + 0.4) \times 2 + 0.3 \times (1 + 0.21) + 0.5 \times (0 + 0.51) = 0.898.$$

Appendix B: Existence and Uniqueness of $V(t, y)$

The revenue-to-go function and the optimal policy satisfy the Hamilton-Jacobi equation (6) with boundary conditions (4). In the following, we will propose a constructive method to show the existence of $V(t, y)$. To do so, consider a sequence of problems called (P_k) where $k = 0, 1, 2, \ldots$. For problem (P_k), we artificially impose a constraint on the original MDP problem – only the first k arriving customers (or less, if that is all that arrives) are considered as potential buyers, and the $(k+1)$th or subsequent arrivals are ignored. Denote the value function of problem (P_k) as $V_k(t, y)$ associated with state y at time t.

Obviously, when $k = 0$, $V_0(t, y) = 0 \ \forall (t, y)$, $0 \le t \le T$, $0 \le y \le Q$. Given that $V_k(t, y)$ is already known, $V_{k+1}(t, y)$ can be determined by the following recursion:

$$V_{k+1}(t, y) = \int_t^T \lambda e^{-\lambda(\tau - t)} g_k(\tau, y) ds, \tag{32}$$

where

$$g_k(\tau, y) = \max_{S \subseteq \mathbb{B}(y)} \Big\{ \sum_{j \in \mathcal{N}} P_j(S) \Big[r_j + \max_{i \in \mathcal{M}, a_{ij} = 1} V_k(\tau, y - e_i) \Big] + P_0(S) V_k(\tau, y) \Big\}. \tag{33}$$

In addition, a terminal condition $V_{k+1}(T, y) = 0$ is required to well define $V_{k+1}(t, y)$.

The value function $V_k(t, y)$ constructed above has some nice structural properties, which help us obtain one of our main results about the unique existence of the revenue-to-go function.

Proposition 5. *For any state (t, y), $0 \le y \le Q$ and $0 \le t \le T$, the value function of problem (P_k) has the following structural properties:*

(i) $V_k(t,y)$ *is increasing in* k*;*

(ii) $V_k(t,y)$ *is bounded for each* k*, in particular,* $V_k(t,y) \le \bar{r}\sum_{i=1}^m y_i$*, where* $\bar{r} := \max_{j\in\mathcal{N}} r_j$ *is the maximum selling price of all the products.*

(iii) $V_k(t,y)$ *has a limit when* k *goes to infinity.*

Proof. We prove (i) and (ii) by induction.

(i) Because $V_0(t,y) = 0$, it is trivial that $V_1(t,y) \ge 0$ from (32). Therefore, $V_1(t,y) \ge V_0(t,y)$. Suppose $V_{k+1}(t,y) \ge V_k(t,y)$ for $k \ge 0$ and $\forall y$. Consider $V_{k+2}(t,y)$. By the induction hypothesis, we can easily show that $g_{k+1}(\tau,y) \ge g_k(\tau,y)$ $\forall \tau \in [t,T]$ and $V_{k+2}(t,y) \ge V_{k+1}(t,y)$ follows immediately by (32).

(ii) When $k = 0$, we have $V_0(t,y) = 0 \le \bar{r}\sum_{i=1}^m y_i$. Suppose $V_k(t,y) \le \bar{r}\sum_{i=1}^m y_i$ for $k \ge 0$ and $\forall y$. Because $r_j \le \bar{r}$ $\forall j = 1,2,\ldots,n$, one can check that,

$$r_j + \max_{i\in\mathcal{M},a_{ij}=1} V_k(t,y-e_i) \le r_j + \max_{i\in\mathcal{M},a_{ij}=1} \bar{r}\Big(\sum_{\ell=1}^m y_\ell - 1\Big) \le \bar{r}\sum_{i=1}^m y_i.$$

Substituting the above inequality into (33), we can show that $g_k(\tau,y) \le \bar{r}\sum_{i=1}^m y_i$. Consequently,

$$V_{k+1}(t,y) \le \int_t^T \Big(\lambda e^{-\lambda(\tau-t)}\bar{r}\sum_{i=1}^m y_i\Big)d\tau = [1-e^{-\lambda(T-t)}]\bar{r}\sum_{i=1}^m y_i < \bar{r}\sum_{i=1}^m y_i.$$

As a bounded monotone sequence converges, (iii) follows directly from (i) and (ii). □

Proposition 5 means that for any state (t,y), the value function $V_k(t,y)$, $k = 0,1,2,\ldots$, is an increasing and bounded sequence. By the monotone convergence theorem (e.g., Kubrusly, 2015), we know that the value function $V_k(t,y)$ converges when k is sufficiently large. In particular, the following theorem shows that $V_k(t,y)$ actually converges to the revenue-to-go function of the original MDP problem as k goes to infinity.

Theorem 3. *For any state* (t,y)*, with* $0 \le y \le Q$ *and* $0 \le t \le T$*, we have:*

$$V(t,y) = \lim_{k\to\infty} V_k(t,y). \tag{34}$$

Proof. (a) We first show that $\forall k \ge 0$, $V(t,y) \ge V_k(t,y)$. This is obviously true for $k = 0$. Suppose the inequality holds for $k \ge 0$. Then $g_k(\tau,y) \le g(\tau,y)$ $\forall(\tau,y)$. By comparing the value functions (8) and (33), we have

$$V_{k+1}(t,y) = \int_t^T \lambda e^{-\lambda(\tau-t)} g_k(\tau,y)d\tau \le \int_t^T \lambda e^{-\lambda(\tau-t)} g(\tau,y)ds = V(t,y).$$

As the above inequality holds for any sufficiently large k, we conclude that

$$\lim_{k\to\infty} V_k(t,y) \le V(t,y). \tag{35}$$

(b) Let $J^k_{(u,\Theta)}(t,y)$ be the expected revenue for problem (P_k) under a given control policy (u,Θ). Clearly,

$$J^k_{(u,\Theta)}(t,y) \le \sup_{(u,\Theta)\in\mathcal{U}} J^k_{(u,\Theta)}(t,y) = V_k(t,y).$$

Let $J_{(u,\Theta)}(t,y,\omega)$ be the revenue for the original MDP model under the policy (u,Θ) on sample path ω, $J^k_{(u,\Theta)}(t,y,\omega)$ be the revenue for problem (P_k) under the policy (u,Θ) on sample path ω, and $A(\omega)$ be the total number of arrivals in $[t, T]$ on sample path ω. Then,

$$\begin{aligned} J_{(u,\Theta)}(t,y) &= \mathbb{E}[J_{(u,\Theta)}(t,y,\omega)] \\ &= \mathbb{E}[J_{(u,\Theta)}(t,y,\omega)\mathbb{I}_{\{A(\omega)>k\}}] + \mathbb{E}[J_{(u,\Theta)}(t,y,\omega)\mathbb{I}_{\{A(\omega)\le k\}}] \\ &= \mathbb{E}[J_{(u,\Theta)}(t,y,\omega)\mathbb{I}_{\{A(\omega)>k\}}] + \mathbb{E}[J^k_{(u,\Theta)}(t,y,\omega)\mathbb{I}_{\{A(\omega)\le k\}}]. \end{aligned}$$

Similarly,

$$\begin{aligned} J^k_{(u,\Theta)}(t,y) &= \mathbb{E}[J^k_{(u,\Theta)}(t,y,\omega)] \\ &= \mathbb{E}[J^k_{(u,\Theta)}(t,y,\omega)\mathbb{I}_{\{A(\omega)>k\}}] + \mathbb{E}[J^k_{(u,\Theta)}(t,y,\omega)\mathbb{I}_{\{A(\omega)\le k\}}]. \end{aligned}$$

Therefore,

$$\begin{aligned} J_{(u,\Theta)}(t,y) &= J^k_{(u,\Theta)}(t,y) + \mathbb{E}[J_u(t,y,\omega)\mathbb{I}_{\{A(\omega)>k\}}] - \mathbb{E}[J^k_{(u,\Theta)}(t,y,\omega)\mathbb{I}_{\{A(\omega)>k\}}] \\ &\le J^k_{(u,\Theta)}(t,y) + \mathbb{E}[|J_{(u,\Theta)}(t,y,\omega)|\mathbb{I}_{\{A(\omega)>k\}}] + \mathbb{E}[|J^k_{(u,\Theta)}(t,y,\omega)|\mathbb{I}_{\{A(\omega)>k\}}]. \end{aligned}$$

Clearly, $|J_{(u,\Theta)}(t,y,\omega)| \le \bar{r}A(\omega)$ and $|J^k_{(u,\Theta)}(t,y,\omega)| \le \bar{r}A(\omega)$. Therefore,

$$\begin{aligned} J_{(u,\Theta)}(t,y) &\le J^k_{(u,\Theta)}(t,y) + 2\bar{r}\mathbb{E}[A(\omega)\mathbb{I}_{\{A(\omega)>k\}}] \\ &\le V_k(t,y) + 2\bar{r}\mathbb{E}[A(\omega)\mathbb{I}_{\{A(\omega)>k\}}]. \end{aligned} \tag{36}$$

Note that the above inequality holds for any $k \ge 0$; i.e., it also holds when k is sufficiently large. By letting $k \to \infty$ in (36), we have

$$J_{(u,\Theta)}(t,y) \le \lim_{k\to\infty} \left\{ V_k(t,y) + 2\bar{r}\mathbb{E}[A(\omega)\mathbb{I}_{\{A(\omega)>k\}}] \right\} = \lim_{k\to\infty} V_k(t,y),$$

where we have used the fact that $\lim_{k\to\infty}\mathbb{E}[A(\omega)\mathbb{I}_{\{A(\omega)>k\}}] = 0$ because $\mathbb{E}[A(\omega)] < \infty$. Given that (u, Θ) is an arbitrary policy, we have

$$V(t,y) = \sup_{(u,\Theta)\in\mathcal{U}} J_{(u,\Theta)}(t,y) \le \lim_{k\to\infty} V_k(t,y). \tag{37}$$

By (35) and (37), it must be true that $V(t,y) = \lim_{k\to\infty} V_k(t,y)$. This completes the proof. □

Theorem 3 suggests a way to construct a solution for the MDP problem from the sequence of problems (P_k). At any state (t, y), starting from $k = 0$, we can obtain $V_{k+1}(t, y)$ using the recursive formulae (32)–(33) until k is sufficiently large. Then $V_{k+1}(t, y)$ is a good approximation for the revenue-to-go function, $V(t, y)$, for the MDP problem. Note that the limit of $V_k(t, y)$ as k goes to infinity is uniquely determined in the way the sequence of problems (P_k) is constructed. Therefore, the optimal revenue function for the MDP problem not only exists, but is unique.

Appendix C: Proofs

Proof of Proposition 1. For any feasible solution $(u, \hat{\Theta})$ of problem (9)–(10), let

$$\hat{\theta}_{ij} = \frac{\sum_S \int_0^T [u^t(S)P_j(S)\hat{\theta}^t_{ij}]dt}{\sum_S \int_0^T [u^t(S)P_j(S)]dt}, \quad \forall i \in \mathcal{M}, j \in \mathcal{N}.$$

Clearly, we have $0 \le \hat{\theta}_{ij} \le 1$. In particular, for any $j \in \mathcal{N}$:

$$\sum_{i\in\mathcal{M}} a_{ij}\hat{\theta}_{ij} = \frac{\sum_S \int_0^T \Big(u^t(S)P_j(S)\sum_{i\in\mathcal{M}}[a_{ij}\hat{\theta}_{ij}^t]\Big)dt}{\sum_S \int_0^T [u^t(S)P_j(S)]dt} = 1.$$

Therefore, $(\hat{\theta}_{ij})_{m\times n}$ is a feasible solution for the fluid control problem. Consequently, one can always transform any time-varying fulfillment policy into an equivalent policy that is independent of time. □

Proof of Proposition 2. Consider an arbitrary pair of (i,j), such that $i,j \in \mathcal{M}$, $i \neq j$. Recall that $\tau^*(\{i\})$ (or $\tau^*(\{j\})$) denotes the length of time that a single physical good i (or j) is offered. For any optimal solution in which $\tau^*(\{i\})$ and $\tau^*(\{j\})$ are both positive, let

$$\delta_3 = \min\Big(\frac{P_i(\{i\})}{P_i(\{i,j\})}\tau^*(\{i\}), \frac{P_j(\{j\})}{P_j(\{i,j\})}\tau^*(\{j\})\Big),\ \delta_1 = \frac{P_i(\{i,j\})}{P_i(\{i\})}\delta_3, \text{ and } \delta_2 = \frac{P_j(\{i,j\})}{P_j(\{j\})}\delta_3.$$

Clearly, $\delta_1 = \tau^*(\{i\})$ or $\delta_2 = \tau^*(\{j\})$ will hold. We can easily show that

$$\delta_1 + \delta_2 = \Big(\frac{P_i(\{i,j\})}{P_i(\{i\})} + \frac{P_j(\{i,j\})}{P_j(\{j\})}\Big)\delta_3 \ge \Big(\frac{P_i(\{i,j\})}{P_i(\{i,j\})+P_j(\{i,j\})} + \frac{P_j(\{i,j\})}{P_i(\{i,j\})+P_j(\{i,j\})}\Big)\delta_3 = \delta_3,$$

where the inequality follows from the demand dilution effect.

We now construct another solution, denoted by $(\tilde{\tau}, \hat{\theta}^*)$, which differs from $(\tau^*, \hat{\theta}^*)$ in that $\tilde{\tau}(\{i\}) = \tau^*(\{i\}) - \delta_1$, $\tilde{\tau}(\{j\}) = \tau^*(\{j\}) - \delta_2$, and $\tilde{\tau}(\{i,j\}) = \tau^*(\{i,j\}) + \delta_3$ ($\tilde{\tau}(S) = \tau^*(S)$ for all other non-empty set S). One can easily check that the solution $\tilde{\tau}$ is also feasible for the fluid program (12)–(16), and the corresponding fluid revenue is the same as that under the solution τ^*. Therefore, the solution $\tilde{\tau}$, in which $\tilde{\tau}(\{i\}) \times \tilde{\tau}(\{j\}) = 0$, is also optimal for problem (12)–(16). □

Proof of Theorem 1. Define process $Y = \{Y(t) : 0 \le t \le T\}$ where $Y(t) = (Y_1(t), \ldots, Y_m(t))$ represents the vector of inventory state at time t. Consider an arbitrary Markovian policy (u, Θ) for the MDP problem (5), let $u = \{u^t = f_1(t, Y(t-)) : 0 \le t \le T\}$ where $u^t(S) = f_1(S; t, Y(t-))$ be the control of set $S \subseteq \mathcal{N}$, and let $\Theta = \{\Theta^t = f_2(t, Y(t-) : 0 \le t \le T)$ where $\theta_{ij}^t = f_2(i,j;t,Y(t-))$ be the control of product assignment for $i \in \mathcal{M}, j \in \mathcal{N}$. (Note that $Y(t-)$ represents the left limit of $Y(t)$ at time t.) Define a solution, denoted as $(\hat{\tau}, \hat{X})$, for the fluid control problem (17)–(21), where

$$\begin{aligned}
\hat{\tau}(S) &:= \mathbb{E}\Big\{\int_0^T f_1(S;t,Y(t-))dt\Big\} = \int_0^T \mathbb{E}[f_1(S;t,Y(t-))]dt,\ \forall S \subseteq \mathcal{N},\\
\hat{X}_{ij}(S) &:= \mathbb{E}\Big\{\int_0^T f_1(S;t,Y(t-))f_2(i,j;t,Y(t-))dt\Big\}\\
&= \int_0^T \mathbb{E}[f_1(S;t,Y(t-))f_2(i,j;t,Y(t-))]dt,\ \forall S \subseteq \mathcal{N}, i \in \mathcal{M}, j \in \mathcal{N}.
\end{aligned}$$

It's not difficult to show that the solution $(\hat{\tau}, \hat{X})$ defined above satisfies conditions (18)–(21) considering the fact that the Markovian policy (u, Θ) satisfies condition (1).

Given that customers arrive according to a Poisson process, we know that process $\hat{D}(u,\Theta) = \{\hat{D}(t,u,\Theta) : 0 \le t \le T\}$ is a martingale, where

$$\hat{D}_j(t,u,\Theta) = D_j(t,u,\Theta) - \int_0^t \lambda \sum_S [u^t(S)P_j(S)]ds,\ \ t \ge 0,\ j \in \mathcal{N}.$$

As such,

$$\begin{aligned}
J_{(u,\Theta)}(0,Q) &= \mathbb{E}\Big\{\sum_{j\in\mathcal{N}}\int_0^T r_j dD_j(\tau,u,\Theta)\Big\} \\
&= \mathbb{E}\Big\{\sum_{j\in\mathcal{N}}\int_0^T r_j\lambda\sum_S[u^t(S)P_j(S)]dt\Big\} \\
&= \sum_S\sum_{j\in\mathcal{N}}\int_0^T \lambda r_j\mathbb{E}[f_1(S;t,Y(t-))]P_j(S)dt.
\end{aligned}$$

Clearly, we have $J_{(u,\Theta)}(0,Q) \leq R^*$ because $J_{(u,\Theta)}(0,Q)$ equals to the revenue function of the fluid control problem (17)–(21) under solution $(\hat{\tau},\hat{X})$. By letting the Markovian policy (u,Θ) be the optimal solution to the MDP problem, we arrive at the conclusion that $V^* \leq R^*$ accordingly. □

To prove Proposition 3, we first present the following technical lemma.

Lemma 2. *Let X^k be the time at which a homogeneous Poisson process with intensity $k\mu$ first reaches level kq, $\mu, q > 0$, $k = 1, 2, \ldots$. We then have*

$$X^k = \frac{q}{\mu} + \frac{1}{\mu}\sqrt{\frac{q}{k}}Z^k, \tag{38}$$

where $Z^k \xrightarrow{d} Z$, the standard normal random variable, when k goes to infinity.

Proof. For each fixed $k \geq 1$, let $\{X^{k,j} : j \geq 1\}$ be a sequence of *i.i.d.* exponential random variables with mean $\frac{1}{k\mu}$. Then we have

$$X^k = \sum_{j=1}^{kq} X^{k,j}.$$

Define

$$Z^{k,j} = \sqrt{\frac{k\mu^2}{q}}\Big(X^{k,j} - \frac{1}{k\mu}\Big).$$

Then $\{Z^{k,j}, k = 1, 2, \ldots, j = 1, 2, \ldots, kq\}$ is a triangular array that satisfies the following properties:

(1) For each k, $Z^{k,1}$, $Z^{k,2}$, …, $Z^{k,kq}$ are mutually independent;

(2) $\mathbb{E}[Z^{k,j}] = 0$ for each $k \geq 1$ and $j \leq kq$.

(3) $\sum_{j=1}^{kq}\mathbb{E}[(Z^{k,j})^2] = 1$.

Since $X^{k,j} \sim \text{Poisson}(\frac{1}{k\mu})$, we have

$$\mathbb{E}[X^{k,j}] = \frac{1}{k\mu},\ \mathbb{E}[(X^{k,j})^2] = \frac{2}{(k\mu)^2},\ \mathbb{E}[(X^{k,j})^3] = \frac{6}{(k\mu)^3}.$$

Therefore,

$$\begin{aligned}
\mathbb{E}\Big[\Big|X^{k,j} - \frac{1}{k\mu}\Big|^3\Big] &= \mathbb{E}\Big[\Big(X^{k,j} - \frac{1}{k\mu}\Big)^3\Big] + 2\int_0^{\frac{1}{k\mu}}(\frac{1}{k\mu} - x)^3 k\mu e^{-k\mu x}dx \\
&= \frac{12e^{-1} - 2}{(k\mu)^3}.
\end{aligned}$$

As a result, the following Lyapunov condition holds:

$$\lim_{k\to\infty}\Big\{\sum_{j=1}^{kq}\mathbb{E}[|Z^{k,j}|^3]\Big\}=\lim_{k\to\infty}\Big\{kq\Big(\frac{k\mu^2}{q}\Big)^{3/2}\times\frac{12e^{-1}-2}{(k\mu)^3}\Big\}=0.$$

By Lyapunov's Theorem (see Durrett 2010, p.113), we conclude that $Z^k=\sum_{j=1}^{kq}Z^{k,j}$ converges in distribution to a standard normal random variable. Thus (38) is derived. □

Proof of Proposition 3. We first characterize the structure of $\Gamma^k_{\ell i}$. By its definition, clearly, $\Gamma^k_{\ell i}-t_{\ell-1}$ is a Gamma random variable with shape parameter $\lfloor k\mathcal{L}_{(\ell-1)i}\rfloor-\lfloor k\mathcal{L}_{\ell i}\rfloor$ and scale parameter $\frac{1}{k\mu_i(S_\ell)}$ (i.e., a sum of $\lfloor k\mathcal{L}_{(\ell-1)i}\rfloor-\lfloor k\mathcal{L}_{\ell i}\rfloor$ *i.i.d.* exponential random variables with mean $\frac{1}{k\mu_i(S_\ell)}$). By the technical Lemma 2, we can express $\Gamma^k_{\ell i}$ as follows:

$$\Gamma^k_{\ell i}=t_{\ell-1}+\frac{\lfloor k\mathcal{L}_{(\ell-1)i}\rfloor-\lfloor k\mathcal{L}_{\ell i}\rfloor}{k\mu_i(S_\ell)}+\frac{Z^k_{\ell i}}{k\mu_i(S_\ell)}\sqrt{\lfloor k\mathcal{L}_{(\ell-1)i}\rfloor-\lfloor k\mathcal{L}_{\ell i}\rfloor},\ i\in\mathbb{A}^k_{\ell i},\tag{39}$$

where $Z^k_{\ell i}$ converges in distribution to a standard normal distribution Z, denoted as $Z^k_{\ell i}\xrightarrow{d}Z$. Equation (39) shows that $\Gamma^k_{\ell i}$ comprises two parts, a deterministic component representing the expected time it takes for the inventory level of product i to first reach the target level $\lfloor k\mathcal{L}_{\ell i}\rfloor$, and a random component denoted by $\Delta\Gamma^k_{\ell i}$:

$$\Delta\Gamma^k_{\ell i}=\frac{Z^k_{\ell i}}{k\mu_i(S_\ell)}\sqrt{\lfloor k\mathcal{L}_{(\ell-1)i}\rfloor-\lfloor k\mathcal{L}_{\ell i}\rfloor}.$$

For a sufficiently large k, we have:

$$\lim_{k\to\infty}\Big\{\Delta\Gamma^k_{\ell i}\Big\}=\lim_{k\to\infty}\Big\{\frac{Z^k_{\ell i}}{\sqrt{k}\mu_i(S_\ell)}\sqrt{\mathcal{L}_{(\ell-1)i}-\mathcal{L}_{\ell i}}\Big\}\xrightarrow{d}\lim_{k\to\infty}\Big\{\frac{Z}{\sqrt{k}\mu_i(S_\ell)}\sqrt{\mathcal{L}_{(\ell-1)i}-\mathcal{L}_{\ell i}}\Big\}.$$

That is, $\Delta\Gamma^k_{\ell i}$ is of order $\frac{1}{\sqrt{k}}$ and is therefore equal to zero with probability 1. As a result, we have

$$\lim_{k\to\infty}\{\Gamma^k_{\ell i}\}\xrightarrow{d}t_{\ell-1}+\frac{\mathcal{L}_{(\ell-1)i}-\mathcal{L}_{\ell i}}{\mu_i(S_\ell)}=t_{\ell-1}+\tau^*(S)=t_\ell$$

with probability 1.

To prove inequality (25), we only need to show that $\Delta R_{\ell i}$ is of a lower order k; that is,

$$\lim_{k\to\infty}\Big\{\frac{\Delta R_{\ell i}}{k}\Big\}=0,\ i\in\mathbb{A}^k_\ell.\tag{40}$$

Recall that $\Delta R_{\ell i}$ corresponds to the scenario in which the inventory level of physical good i contained in S_ℓ first drops to $\lfloor k\mathcal{L}_{\ell i}\rfloor$ before time t_ℓ and prior to all other physical goods. By its formulation, we have

$$\begin{aligned}\Delta R_{\ell i}&\leq\mathbb{E}\Big\{\sum_{j\in\mathbb{A}^k_\ell}\Big[\bar{p}(S_\ell)N^k_{\ell j}(\Gamma^k_{\ell i}\wedge t_\ell,t_\ell)\Big]\Big\}\\&=\sum_{j\in\mathcal{N}}\Big\{\bar{p}(S_\ell)k\mu_j(S_\ell)\mathbb{E}\Big[t_\ell-\Gamma^k_{\ell i}\wedge t_\ell\Big]\Big\}=k\sum_{j\in\mathcal{N}}\Big\{\bar{p}(S_\ell)\mu_j(S_\ell)\mathbb{E}\Big[t_\ell-\Gamma^k_{\ell i}\Big]^+\Big\}.\end{aligned}$$

Therefore,

$$\lim_{k\to\infty}\Big\{\frac{\Delta R_{\ell i}}{k}\Big\}\leq\lim_{k\to\infty}\Big\{\sum_{j\in\mathcal{N}}\bar{p}(S_\ell)\mu_j(S_\ell)\mathbb{E}\Big[t_\ell-\Gamma^k_{\ell i}\Big]^+\Big\}=\sum_{j\in\mathcal{N}}\Big(\bar{p}(S_\ell)\mu_j(S_\ell)\Big)\lim_{k\to\infty}\Big\{\mathbb{E}\Big[-\Delta\Gamma^k_{\ell i}\Big]^+\Big\}=0,$$

where we have used $Z_{\ell i}^k \xrightarrow{d} Z$, $\sup_k \mathbb{E}[-Z_{\ell i}^k]^2 = 1$, and $\mathbb{E}[-Z]^+ = \frac{1}{\sqrt{2\pi}}$. This completes the proof of (40).

□

Proof of Theorem 2. Starting from time $t_0 = 0$ and inventory level being $k\mathcal{L}_0 = kQ$, we have

$$V_\pi^k = V_\pi^k(t_0, k\mathcal{L}_0) = \sum_{\ell=1}^{K} \Big[V_\pi^k(t_{\ell-1}, \lfloor k\mathcal{L}_{\ell-1} \rfloor) - V_\pi^k(t_\ell, \lfloor k\mathcal{L}_\ell \rfloor) \Big] + V_\pi^k(t_n, \lfloor k\mathcal{L}_n \rfloor).$$

By Proposition 3, we have

$$\lim_{k\to\infty} \left\{ \frac{V_\pi^k}{k} \right\} \geq \lim_{k\to\infty} \left\{ \sum_{\ell=1}^{K} \sum_{j\in\mathcal{N}} \Big(\lambda\tau^*(S_\ell) r_j P_j(S_\ell) \Big) + \frac{V_\pi^k(t_K, \lfloor k\mathcal{L}_n \rfloor)}{k} \right\} \geq \sum_{\ell=1}^{K} \sum_{j\in\mathcal{N}} \Big(\lambda\tau^*(S_\ell) r_j P_j(S_\ell) \Big) = R^*,$$

from which we arrive at (26), because $R^{k*} = kR^*$ and $V_\pi^k \leq V^{k*} \leq R^{k*}$. Consequently,

$$\lim_{k\to\infty} \left\{ \frac{V^{k*} - V_\pi^k}{V^{k*}} \right\} \leq \lim_{k\to\infty} \left\{ \frac{R^{k*} - V_\pi^k}{R^{k*}} \right\} = \frac{1}{R^*} \lim_{k\to\infty} \left\{ \frac{kR^* - V_\pi^k}{k} \right\} = 0,$$

which arrives at (27). □