This article was downloaded by: [216.73.216.198] On: 18 June 2026, At: 02:39
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

INFORMS TutORials in Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Demand Selection and Assignment Problems in Supply Chain Planning

Joseph GeunesYasemin MerzifonluoğluH. Edwin RomeijnKevin Taaffe

To cite this entry: Joseph GeunesYasemin MerzifonluoğluH. Edwin RomeijnKevin Taaffe. Demand Selection and Assignment Problems in Supply Chain Planning. *In* INFORMS TutORials in Operations Research. Published online: 14 Oct 2014; 124-141.
https://doi.org/10.1287/educ.1053.0014

Full terms and conditions of use: https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2005, INFORMS

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

TUTORIALS IN
OPERATIONS RESEARCH
INFORMS 2005

© 2005 INFORMS | ISBN 1-877640-21-2
DOI 10.1287/educ.1053.0014

Demand Selection and Assignment Problems in Supply Chain Planning

Joseph Geunes, Yasemin Merzifonluoğlu, H. Edwin Romeijn
Department of Industrial and Systems Engineering, University of Florida, Gainesville,
Florida 32611 {geunes@ise.ufl.edu, yasemin@ufl.edu, romeijn@ise.ufl.edu}

Kevin Taaffe
Department of Industrial Engineering, Clemson University, Clemson, South Carolina 29634,
taaffe@clemson.edu

Abstract Effective demand planning has recently become recognized as an important source of competitive advantage from both marketing and operations perspectives. The operations literature has often focused on pricing (revenue management) and information-based approaches (e.g., improved forecasting, mechanisms to elicit advance demand information) for integrated demand and production management. We can think of pricing, for example, as an implicit demand selection mechanism that determines the scope and scale of demands a supplier will ultimately serve, and that therefore drives the production requirements (as well as costs and revenues) that the supplier will face. In contrast, a variety of planning and design contexts involve a supplier or supply chain stage explicitly selecting a subset of demands from a collection of potential downstream demand sources. Such decisions on the demands a supply chain stage will serve may arise out of necessity (due to supply capacity limits), may be based strictly on economic considerations, or may be a component of the (assignment) decisions in a larger supply chain network design problem. Incorporating such dimensions of demand selection flexibility within supply chain planning models allows a supplier to best match available resources with downstream requirements, which can create opportunities for enhancing profitability. This chapter discusses a class of optimization models for addressing varying degrees of demand selection flexibility in integrated production and demand planning.

Keywords production planning; demand management; revenue management; economic lot sizing; newsvendor problem

1. Introduction

Recent trends in supply chain management recognize the importance of effective demand management for suppliers. Research on demand management in the operations literature has broadly focused on the mechanisms a supplier can use to influence demand, and how a supplier can best utilize these mechanisms to provide a good match between capacity and demand. In this chapter, we discuss a class of integrated operations planning and demand management problems, for which we provide corresponding planning models. This class of models addresses questions regarding the best levels of demand for a supplier. We begin with single-stage planning problems where the stage has varying degrees of flexibility in determining the set of demands to which it will respond. In addition, we consider multifacility problems that determine the best allocation (or assignment) of downstream demands to upstream facilities. As we will see, the single-stage planning models, while interesting in their own right as combined demand and operations planning models, also serve as subproblems in a decomposition strategy applied to the multifacility problems.

Recent demand management literature in operations discusses mechanisms for affecting both capacity and demand to increase profit. Crandall and Markland [6] studied service industry firms and classified several demand management approaches, including capacity management and demand-influencing strategies. Moodie [24] examined pricing and lead-time negotiation strategies for maximizing long-run net revenue when customers are time and price sensitive and capacity is fixed. Keskinocak et al. [18] considered models for coordinating scheduling and lead-time quotation. In these models, revenues from customers are sensitive to lead times, and the scheduling problem therefore contains a demand management dimension. Iyer et al. [17] use postponement (and an associated customer reimbursement) as a mechanism for managing demand surges under limited capacity. These demand management problems exploit knowledge of customer delivery-timing flexibility to effectively match supply and demand, but do not consider economies of scale often associated with operations costs. Lee et al. [20], on the other hand, consider the benefits of customer delivery-timing flexibility within a multiperiod production-planning context with fixed production setup costs.

Integrated operations and demand-planning problems often draw from two related streams of literature involving yield (or revenue) management and combined inventory planning and pricing models. Yield management is often associated with perishable goods such as airline seats or hotel stays (see, for example, Weatherford and Bodily [36] for a thorough discussion of this line of research), while integrated pricing and inventory planning models often consider influencing uncertain demand for durable physical goods. Gallego and van Ryzin [9] and Petruzzi and Dada [26] provide excellent examples of research in this area, as well as discussions of the broader related literature. In this problem class, pricing decisions serve as a mechanism for effectively setting the best demand levels over time.

Although much of the integrated pricing and inventory literature focuses on using pricing to influence uncertain demand, pricing decisions have also been considered in make-to-order and other effectively deterministic contexts. Here the emphasis lies in managing economies of scale in production and/or capacity in conjunction with demand. Thomas [33] considered the classical uncapacitated economic lot-sizing (UELS) problem with price-sensitive demands and proposed a dynamic programming approach for setting prices in every period. For the problem of setting a single price over the entire horizon in UELS, Kunreuther and Schrage [19] proposed a heuristic method, while Gilbert [13] subsequently provided an exact algorithm. More recently, van den Heuvel and Wagelmans [34] provided a new exact algorithm that applies to a more general set of cost and demand parameters. We also note that Biller et al. [3] incorporated pricing decisions within a capacitated lot-sizing problem, but without fixed production setup costs.

While pricing serves as a key mechanism for influencing demand, the single-stage models discussed in this chapter consider a more direct way of influencing demand through explicit demand selection decisions. That is, these models assume that a set of candidate demands exists from which the supplier is free to select. Such demand selection decisions may be practically relevant in several contexts. For example, if capacity limits exist at a production facility, and total demand outstrips this capacity, a supplier may be forced to turn away some subset of customer demands. As another example, suppose that a firm wishes to determine a set of initial markets to penetrate in the roll-out of a new product, and demand characteristics vary among markets. In both of these examples, the supplier selectively determines some subset of demands it will satisfy to maximize its net profit after operations costs.

Past literature on models that directly select demands from a candidate set is limited, with a few notable exceptions. Charnsirisakskul et al. [5] considered a (deterministic) capacitated order selection model in which customers express acceptable delivery lead times, and the supplier incurs a tardiness cost for late deliveries. They focused on the positive impact of increased customer lead-time flexibility on profits in a capacitated setting without economies of scale in production costs (all of the deterministic models we will consider contain a fixed-charge or general concave production cost structure, resulting in production economies of scale). In a stochastic setting, Carr and Lovejoy [4] proposed an inverse newsvendor model, in which the capacity to satisfy demand is a random variable with known mean and variance. The supplier has available multiple prioritized customer classes, each consisting of a set of customer clusters, and must determine what fraction of each customer cluster within each class it will pursue. The objective is to determine a demand distribution from the opportunity set of available distributions that maximizes expected profit. In the stochastic version of the model that we will consider, capacity is a decision variable, and customer demands must be completely accepted or rejected.

Demand selection serves as the common thread that links the single-stage models we will discuss in this chapter. These models generalize several classical production and inventory planning models to incorporate demand selection. In certain cases, where appropriate, we illustrate a very natural analogous pricing model corresponding to the demand selection version of the model. In addition to the selection dimension of demand management, several of our models consider elements of customer delivery-timing flexibility. We also discuss a set of models in which production is capacity constrained. Such models can serve as valuable benchmarking tools when a supplier has a good understanding of the costs of demand fulfillment as well as customer delivery-timing flexibility.

This chapter brings together a number of recent works with a demand selection theme. The broad scope of problems we will discuss, and space limitations, preclude providing the complete set of our results for all of the models we discuss. We will therefore provide the primary results for those problems we discuss, and cite the sources containing detailed results where appropriate. The remainder of this chapter is organized as follows. Section 2 discusses the generalization of the classical economic order quantity (EOQ) and newsvendor problems to allow for demand selection flexibility. Somewhat surprisingly, the resulting optimization problems are structurally identical, but lead to different interpretations. Section 3 then considers classical dynamic, deterministic production planning models that allow for demand selection. We briefly discuss interesting results on capacity-constrained versions of the problem, including the case in which capacitated overtime and subcontracting options are available. Section 4 provides a framework for multifacility customer assignment problems, for which the models in §§2 and 3 can serve as subproblems. Section 5 concludes with a summary and possible directions for future research.

2. Static Single-Stage Models with Demand Selection

We first study, in §2.1, a generalization of the well-known EOQ model (Harris [15]). In the presence of demand selection flexibility, instead of having a fixed demand rate λ, we have a choice of different markets we can serve, each with its own demand rate and net revenue. Then, in §2.2, we consider a single-stage problem with stochastic market demands in a single-period context, which we call the *selective newsvendor problem*, or SNP. For more details on the models discussed in this section, please see Geunes et al. [12] and Taaffe et al. [31].

2.1. EOQ Model with Market Selection

2.1.1. Basic Model. We begin with the base assumptions of the EOQ model, with a single commodity produced at a single stage, a fixed ordering cost of K, a per-unit order cost of c, and a per-unit holding cost of h per unit time (see Nahmias [25] or Silver et al. [29] for a full discussion of the EOQ model). A set of potential markets I exists (indexed by i), and the producer can choose to supply any subset of markets. Market i provides a net revenue per unit of r_i (net of any variable delivery costs), where we assume without loss of generality that $r_i > 0$ for all $i \in I$. Market i has a deterministic demand rate equal to λ_i (units per unit time). We preindex markets in *decreasing net revenue* (DNR) order, which implies that $r_i > r_j$ for $i > j$ (where we assume without loss of generality that no two markets have identical net revenue values).

Letting y_i denote a binary decision variable equal to 1 if we satisfy the demand of market i, and 0 otherwise, the demand rate the supplier faces equals $\sum_{i\in I}\lambda_i y_i \equiv \lambda^\top y$ (where λ and y are the vectors of λ_i and y_i variables). Thus, the optimal time between orders, $T^*(y)$, is equal to $\sqrt{2K/h\lambda^\top y}$. Substituting this into the minimum average (per unit time) setup and holding-cost equation for the EOQ model results in $\sqrt{2Kh\lambda^\top y}$. The market-selection version of the EOQ problem maximizes net revenue less average holding and ordering costs per unit time, based on order selection decisions, which leads to the following optimization problem.

$$\textbf{[EOQMC]} \qquad \text{maximize} \quad \sum_{i\in I} r_i\lambda_i y_i - \sqrt{2Kh\sum_{i\in I}\lambda_i y_i}$$

$$\text{subject to:} \quad y_i \in \{0,1\} \quad \text{for all } i \in I.$$

Shen et al. [28] showed that problems containing this structure can be solved in $\mathcal{O}(n\log n)$ time (where $n = |I|$) using the following property. If we sort items in nonincreasing order of the ratio of the coefficient of y_i in the linear term to the coefficient of y_i in the square root term, then if $j > i$ and an optimal solution selects item j, an optimal solution exists that selects item i. This indexing approach applied to our EOQMC problem results in the DNR index ordering that we used to preindex the markets in this problem. Therefore, after sorting items in this order, we simply need to evaluate the cost of $n+1$ ordered solutions containing markets $1,\ldots,j$, for $j = 0,\ldots,n$ (with $j = 0$ corresponding to selecting no markets at all). Interestingly, the attractiveness of a market in this context is purely determined by the per-unit net revenue value and is independent of the market demand level. Geunes et al. [12] exercise this model to provide expressions useful in managerial decisions, such as the minimum revenue and demand levels required to enter a new market. They also show that the DNR indexing approach continues to hold under a finite production rate and under market-specific holding-cost parameters, after a holding-cost-adjusted revenue index ordering is applied.

2.1.2. Capacity Constraints.

In addition to using a finite production rate to reflect capacity limits, a supplier might also face limits on its total output per unit time, or on the batch sizes it can produce in a production run (due to, for example, storage space limitations). When total output per unit time is limited by some bound B, we must add the following constraint to the [EOQMC] formulation.

$$\sum_{i\in I}\lambda_i y_i \le B \tag{1}$$

With this additional constraint, the resulting problem becomes a nonseparable, nonlinear, and convex 0–1 knapsack problem. Geunes et al. [12] provide a polynomial-time algorithm for solving the linear relaxation of this problem, using a combination of the DNR index-ordering algorithm and a simple continuous knapsack problem solution. They also provide an asymptotically optimal algorithm (in the number of markets, for a problem class in which the capacity grows linearly in the number of markets) for solving the 0–1 version of the problem. These results continue to hold under a finite production rate assumption.

When a limit of B_L exists on the maximum lot size, the problem becomes slightly trickier. We consider the following constraint:

$$Q \le B_L.$$

Let I^s denote a set of selected markets (i.e., those for which $y_i = 1$), and let λ^s denote the corresponding demand rate (i.e., $\lambda^s = \sum_{i\in I^s}\lambda_i$). Because the average cost per unit time is convex in Q for a given set of markets, if $Q^*(I^s) = \sqrt{2K\lambda^s/h} \le B_L$ holds, then this is the optimal order quantity for the set of markets I^s. If $Q^*(I^s) \ge B_L$, then setting $Q(I^s) = B_L$ provides the best order quantity for the set I^s.

We approach this problem by considering two separate subproblems: one that admits values of y such that the corresponding EOQ value is feasible ($\le B_L$; we will call these EOQ-feasible solutions), and one that admits all values of y but requires a lot size of $Q = B_L$. In the first subproblem, we use the objective function of the EOQMC problem (which assumes an EOQ value is used for the selected markets) and add the constraint

$$\sqrt{2K\lambda^\top y/h} \le B_L \equiv \sum_{i\in I} \lambda_i y_i \le \frac{hB_L^2}{2K},$$

which takes the same form as constraint (1), and we can therefore handle this problem in the same manner discussed at the beginning of this section. We next need to solve the following problem for the case in which $Q = B_L$:

$$\text{maximize} \quad \sum_{i\in I}\left(r_i - \frac{K}{B_L}\right)\lambda_i y_i - \frac{hB_L}{2}: \qquad y_i \in \{0,1\}, \quad i = 1,\ldots,n.$$

The optimal solution to the above problem simply selects all markets i such that $r_i \ge K/B_L$, where K/B_L is the per-unit setup cost (this results because the holding cost is a constant as a result of fixing $Q = B_L$). Let $I^s(B_L)$ denote the selected markets in the above problem. If $Q(I^s(B_L)) \le B_L$, then the best solution with $Q = B_L$ is dominated by the best EOQ-feasible solution; otherwise, the solution $I^s(B_L)$ provides a candidate solution that we compare to the best EOQ-feasible solution obtained by evaluating the DNR ordered solutions.

2.2. The Selective Newsvendor Problem

2.2.1. Basic Model. In this section we consider a single-stage problem with stochastic market demands in a single-period context, which we call the *selective newsvendor problem*, or SNP. For a more in-depth discussion of the SNP, please see Taaffe et al. [31]. The newsvendor problem has a long history in the research literature, particularly in supply chain management applications in the past 20 years (see Hadley and Whitin [14] and Porteus [27]). In line with the basic newsvendor model, we consider a supplier with a single centralized stocking point for a single item with per-unit cost c, a salvage value $v < c$, and shortage cost e. In addition to these basic newsvendor assumptions, we consider a set I of potential markets that the supplier may choose to serve. Market $i \in I$ provides the supplier with a per-unit net revenue r_i, and serving a market requires incurring a fixed market entry cost of S_i.

Demand during the selling season in market i is denoted by the random variable D_i, which has probability density function (pdf) $f_i(D_i)$, cumulative distribution function (cdf) $F_i(D_i)$, mean μ_i, and variance σ_i^2. We make two assumptions on market demand probability distributions for analytical tractability. We approximate demand in each market using a normal distribution and assume market demands are statistically independent. Note that the assumption of (approximate) normality is reasonable because we deal with market demands rather than individual orders. Moreover, a vast body of past research employs these assumptions in order to obtain insightful structural results for optimal solutions that can be applied more broadly (see, for example, Eppen [7] and Aviv [2]).

The sequence of events occurs as follows. Prior to the selling season, the supplier simultaneously determines the markets it will enter, along with the order quantity Q, which arrives at the beginning of the selling season. Market demands are then realized; if Q is insufficient to satisfy all demands in the selected markets, then a shortage cost of e is incurred per unit short. This unit shortage cost e may correspond to either an expediting cost for units obtained immediately from an external supplier, or to a backlogging cost for units satisfied at the end of the period through an additional replenishment made to the supplier after realizing demand (for example, at the beginning of the period, the supplier observes demand and schedules a later replenishment to cover any shortfall, which arrives at the end of the period). In either case, the entire demand is ultimately satisfied in each selected market. We will later discuss the implications of assuming that shortages result in lost sales.[1]

As in the previous EOQMC model, we let y denote our binary market selection vector, with ith element y_i, and define μ and ν as corresponding vectors of market demand means and variances. The distribution of demand faced by the supplier is normal with mean $\mu^\top y$ and variance $\nu^\top y$. Let $F_y(\cdot)$ denote the cdf of total demand, and observe that the order quantity that minimizes expected cost is given by $Q^*(y) = F_y^{-1}(\rho)$, where $\rho = (e-c)/(e-v)$. Because market demands are normally distributed, we can express the optimal order quantity as $Q^*(y) = \mu^\top y + z_\rho\sqrt{\nu^\top y}$, where z_ρ is the standard normal variate corresponding to the fraction ρ. Letting $L(z_\rho)$ denote the standard normal loss function (see Nahmias [25]), and defining $k(c,v,e) = (c-v)z_\rho + (e-v)L(z_\rho)$, we can write the minimum expected cost at $Q^*(y)$, which we denote by $C(Q^*(y))$, as

$$c\mu^\top y + k(c,v,e)\sqrt{\nu^\top y}.$$

Observe that if we select market i, then our expected net revenue from this market equals $(r_i - c)\mu_i - S_i \equiv \tilde{r}_i$, and we define $\tilde{r}$ as the vector of expected net revenue values; the supplier's total expected net revenue therefore equals $\tilde{r}^\top y$. Taking expected net revenue less expected cost leads to an expected profit equation of $\tilde{r}^\top y - k(c,v,e)\sqrt{\nu^\top y}$. We therefore solve the following optimization problem to maximize expected net profit.

$$[\textbf{SNP}] \qquad \text{maximize} \quad \sum_{i\in I}\tilde{r}_i y_i - k(c,v,e)\sqrt{\sum_{i\in I}\sigma_i^2 y_i}$$

$$\text{subject to:} \quad y_i \in \{0,1\} \quad \text{for all } i \in I.$$

Interestingly, the SNP model is structurally identical to the EOQMC model, although these two models employ drastically different assumptions. In the SNP, the square root term corresponds to an uncertainty-pooling term, while in the EOQMC problem it has a setup-cost-pooling interpretation. We can, of course, employ the same solution approach, although this again leads to a different interpretation. Indexing markets in nonincreasing order of the coefficient of y_i in the linear term to its coefficient in the square root term leads to a ratio of the form $\tilde{r}_i/\sigma_i^2$, which we call the net revenue to uncertainty (NRU) ratio. To optimally solve [SNP], we therefore sort items in decreasing NRU (DNRU) order[2] and evaluate the expected profit of solution j (containing markets $1,\ldots,j$) for $j = 0,\ldots,n$, retaining the best among these solutions as the optimal solution. The NRU ratio provides an indication of the attractiveness of a market in the selective newsvendor context, and captures the trade-off between revenue and uncertainty. We can think of this ratio as a net-revenue term adjusted for the implicit cost of uncertainty in the market.

Observe that if we assume that shortages result in lost sales, then the resulting critical fractile value ρ becomes a function of the markets we select, and the problem becomes substantially more complex as a result. In the case that all market revenue terms ($\tilde{r}_i$s) are identical, however, then again the resulting ρ does not depend on the selected markets, and by adjusting the $k(c,v,e)$ cost term properly, we can obtain an exact expression for expected

---

[1] Note that the results we obtain here are equally valid for the stationary demand, revenue, and cost data case with an infinite horizon and backlogging, assuming periodic review and zero fixed order cost.

[2] We can show that we can assume without loss of generality that no two markets exist with the same NRU ratio value; see Taaffe et al. [31].

profit that is structurally identical to the SNP (although this reduces to a somewhat trivial case in which markets are indexed in decreasing variance order). If, for example, market revenue parameters are not substantially different, or the fraction ρ is close to one, we can use a similar approach to the lost-sales case with identical revenues to approximate the expected profit equation under lost sales.

Taaffe et al. [31] explore the managerial implications of the SNP model and results by examining the minimum revenue and demand levels required to make a market attractive (all else being equal), as well as the maximum level of uncertainty that would allow a new market to be attractive to the supplier.

2.2.2. Impact of Market Effort. In this section we consider the impacts of efforts by the supplier to influence market demands through some sort of additional market effort, such as advertising, or by providing some additional value-added service or product at no cost to customers. We generically refer to this mechanism as "market effort" and let a_i denote the market effort applied to market i at a per-unit cost of t_i. We assume for this case that markets require some market effort if they are selected, and the minimum demand in market i under no market effort, $\underline{\mu}_i$, provides insufficient revenue to cover the fixed market entry cost, i.e., $S_i > (r_i - c)\underline{\mu}_i$. Here we consider the case in which expected market i demand is a function of market effort a_i, and demand variance is independent of market effort (see Taaffe et al. [31] for an analysis that allows demand variance to be influenced by market effort). In line with much of the marketing literature (see Lilien and Rangaswamy [21]), we assume that expected demand follows an S-shaped curve as a function of market effort. That is, a small amount of market effort can create a sharp increase in market demand, while this growth begins to taper off after a large amount of market effort. Let $\mu_i(a_i)$ denote the expected market i demand as a function of market effort a_i. This function is characterized as follows

$$\mu_i(a_i) = \begin{cases} \mu_i^{(1)}(a_i), & 0 \le a_i \le \alpha_i, \\ \mu_i^{(2)}(a_i), & a_i \ge \alpha_i, \end{cases}$$

where $\mu_i^{(1)}(a_i)$ is convex and increasing, $\mu_i^{(2)}(a_i)$ is concave and nondecreasing, $\mu_i^{(1)}(0) = \underline{\mu}_i$, and $\mu_i^{(1)}(\alpha_i) = \mu_i^{(2)}(\alpha_i)$. We assume here that $\mu_i^{(1)}$ and $\mu_i^{(2)}$ are both everywhere differentiable for $a_i \ge 0$ (see Taaffe et al. [31] for analysis of more general nondifferentiable cases).

The selective newsvendor problem with market effort (SNPM) can be formulated as follows.

$$[\textbf{SNPM}] \qquad \text{maximize} \quad \sum_{i\in I}[(r_i - c)\mu_i(a_i) - t_i a_i - S_i]y_i - \sqrt{\sum_{i\in I}\sigma_i^2 y_i}$$

$$\text{subject to:} \quad y_i \in \{0,1\} \quad \text{for all } i \in I,$$

$$a_i \ge 0 \quad \text{for all } i \in I$$

For any given market-selection vector y, the market effort decisions decompose by market. Finding the optimal amount of market effort in a market requires solving

$$\begin{aligned} \text{maximize} \quad & (r_i - c)\mu_i(a_i) - t_i a_i - S_i \\ \text{subject to:} \quad & a_i \ge 0. \end{aligned}$$

Let $\hat{a}_i^{(2)}$ denote a value of $a_i \ge \alpha_i$ such that $d\mu_i^{(2)}(a_i)/da_i = t_i/(r_i - c)$. Taaffe et al. [31] show that two candidates exist for an optimal solution to the above market i effort problem: $a_i = 0$ and $a_i = \hat{a}_i^{(2)}$ (if no such value of $\hat{a}_i^{(2)}$ exists, then $a_i = 0$ is the only candidate solution). Letting a_i^* denote the optimal value of a_i for the market i effort subproblem, and defining $\tilde{r}_i^a = (r_i - c)\mu_i(a_i^*) - t_i a_i^* - S_i$, we can replace the coefficient of y_i in the SNPM formulation with this value of $\tilde{r}_i^a$, drop the last constraint set, and the resulting problem is identical to the original SNP. Our DNRU ratio follows for assessing the attractiveness of markets, where the net-revenue parameter $\tilde{r}_i^a$ is now adjusted to reflect the optimal level of market effort and corresponding expected demand level.

Taaffe et al. [31] consider the implications of an overall marketing budget, which may not allow applying the optimal (unconstrained) market effort in each market. The resulting problem becomes a nonlinear integer programming model for which they provide a branch-and-bound algorithm, which they showed can quickly solve problems involving up to 50 markets.

3. Dynamic Single-Stage Models with Demand Selection

Manufacturers often face an environment in which demands for a good arrive in the form of discrete orders for a good. When production capacity is limited, a manufacturer may be forced to turn some orders away. In other contexts, it may be in the supplier's best interest to reject orders regardless of capacity levels, depending on the economics of production and how much the customer is willing to pay for the good. This section introduces a class of single-product finite-horizon planning models to handle such order selection and rejection decisions under production economies of scale and different customer reservation prices (in each case we consider a horizon of length T, where time periods are indexed by t). In certain cases, as we discuss in the following subsection, the resulting problem is equivalent to a pricing problem, where a (possibly) unique price is set in every period, which determines the aggregate demand level for the period. We first consider contexts in which capacity is effectively unlimited and customers require delivery in a prespecified period, if the supplier elects to satisfy the order. Section 3.2 generalizes this model to situations in which customers allow delivery-time flexibility (in a manner similar to the lot-sizing problem with delivery-time windows discussed in Lee et al. [20]). In §3.3, we consider the model under time-invariant production capacities, and §3.4 generalizes this model to allow for capacity-constrained overtime options and subcontracting. For each of these cases we cite our primary results and refer the reader to additional works that study the problems in greater detail.

3.1. Uncapacitated Order Selection Problem

Consider a supplier to whom customers submit orders for production of a product. The jth order in period t requests d_{jt} units of the good, and prior to the production horizon, the supplier must decide whether or not to accept the order, which provides a per-unit revenue of r_{jt}. Defining $J(t)$ as the number of orders in period t and letting y_{jt} denote the fraction of order j in period t that the supplier chooses to satisfy, the supplier's total revenue over the horizon becomes $\sum_{t=1}^{T}\sum_{j=1}^{J(t)} r_{jt}d_{jt}y_{jt}$. As in the majority of prior production-planning literature, we assume that the supplier faces a fixed production setup cost of S_t if a production setup is performed in period t, and that a variable production cost of c_t is incurred per unit produced. Let z_t denote a binary variable equal to 1 if the supplier sets up in period t, and 0 otherwise. A holding cost of h_t is incurred for each unit remaining in inventory at the end of period t, where I_t denotes the inventory at the end of period t. Letting $D(t,T) = \sum_{s=t}^{T}\sum_{j=1}^{J(t)} d_{st}$ and denoting x_t as the total production amount in period t, we can formulate the supplier's *uncapacitated order selection problem* (UOSP) as follows.

$$[\textbf{UOSP}] \quad \text{maximize} \quad \sum_{t=1}^{T}\left(\sum_{j=1}^{J(t)} r_{jt}d_{jt}y_{jt} - S_t z_t - c_t x_t - h_t I_t\right)$$

$$\text{subject to:} \quad \sum_{j=1}^{J(t)} d_{jt}y_{jt} + I_t = x_t + I_{t-1}, \qquad t = 1,\ldots,T, \tag{2}$$

$$x_t \le D(t,T) z_t, \qquad t = 1,\ldots,T, \tag{3}$$

$$
\begin{aligned}
&x_t, I_t \geq 0, && t = 1, \dots, T,\\
&0 \leq y_{jt} \leq 1, && t = 1, \dots, T,\ j = 1, \dots, J(t),\\
&z_t \in \{0, 1\}, && t = 1, \dots, T
\end{aligned}
$$

The objective function maximizes total revenue less production and holding costs, while constraint (2) ensures inventory balance. Constraint (3) forces production to zero if no setup occurs; otherwise, production is limited only by the maximum remaining demand.

It is straightforward to show that an optimal solution for the UOSP exists such that (i) all orders are completely accepted or rejected (i.e., $y_{jt} = 0$ or 1 for all j and t) and (ii) the zero-inventory production property holds (i.e., $I_{t-1}x_t = 0$ for all t). Using these properties, Geunes et al. [11] illustrate how to solve the UOSP using an acyclic longest-path graph structure that is nearly identical to the one used to solve the well-known Wagner and Whitin [35] economic lot-sizing problem. This graph contains $T + 1$ nodes and an arc emanating from each node to every higher-numbered node. The profit of an arc (t, t') equals the maximum profit possible if we consider using the setup in period t to satisfy demands in periods $t, \dots, t' - 1$. An order's net profit will be included in this profit value if (i) its unit revenue exceeds the associated variable production and holding cost incurred, and (ii) the total net profit (total revenue less variable production and holding costs) on the arc exceeds the setup cost in period t (otherwise the arc is not profitable and receives a zero profit value). The profit value of an arc (t, t') then equals the maximum between the total net profit minus S_t, and zero. Finding the longest path in this graph solves the UOSP. If $J_{\max} = \max_{t=1,\dots,T} J(t)$, then the worst-case complexity of this longest-path approach is $\mathcal{O}(J_{\max}T^2)$.

Geunes et al. [11] consider an analogous pricing problem that can be described as follows. Because orders in a period will be selected in nondecreasing order of their revenue (r_{jt}) values instead of viewing the total revenue as the sum of individual revenues from orders, we can view it as a market's piecewise-linear and concave total revenue curve. In other words, if we index orders in each period in nondecreasing revenue order, then a solution that selects orders $1, \dots, k$ in period t provides a corresponding total revenue of $\sum_{j=1}^{k} r_{jt}d_{jt}$, and the total revenue graphed as a function of the total demand satisfied in the period provides a piecewise-linear revenue curve. We might alternatively view this curve as the total revenue available in the market at a unit price of $p_{kt} = \sum_{j=1}^{k} r_{jt}d_{jt} / \sum_{j=1}^{k} d_{jt}$. Using this interpretation, d_{jt} corresponds to the width of the jth segment of the revenue curve, and r_{jt} corresponds to the slope. Thus, the UOSP is equivalent to a pricing problem when market revenue can be approximated by nondecreasing piecewise-linear concave function of the total demand satisfied (which corresponds to a demand curve that is decreasing in price). The fact that an optimal solution exists for the UOSP in which each order is either fully satisfied or rejected implies that the optimal amount of demand satisfied in the pricing problem will occur at one of the breakpoints of the piecewise-linear revenue curve. When a single price is offered to the entire market, the resulting model is equivalent to the one proposed by Thomas [33]. Geunes et al. [11] also consider a general concave revenue function, different market response curves in each period, and piecewise-linear concave production cost functions.

3.2. Uncapacitated Order Selection with Delivery-Timing Flexibility

Instead of requiring a fixed delivery period, in certain contexts customers might permit flexibility in the time period in which an order is delivered (Lee et al. [20] and Charnsirisakskul et al. [5] considered this type of delivery-timing flexibility in production planning contexts). For this case, we let J denote the total set of orders available to the supplier over the horizon, where order j requests d_j units. Full delivery of order j in period t provides a total revenue of R_{jt}, and delivery of any fraction y_{jt} of the total quantity requested provides proportional revenue of $R_{jt}y_{jt}$. We formulate the uncapacitated order selection problem with delivery-timing flexibility (UOSP_{DTF}) as follows.

$$\begin{array}{lll}
[\mathbf{UOSP_{DTF}}] \quad \text{maximize} & \displaystyle\sum_{t=1}^{T}\left(\sum_{j=1}^{J} R_{jt}y_{jt} - S_t z_t - c_t x_t - h_t I_t\right) & \\
\text{subject to:} & \displaystyle\sum_{j=1}^{J} d_{jt}y_{jt} + I_t = x_t + I_{t-1}, & t = 1,\ldots,T, \\
& x_t \le D(1,T)z_t, & t = 1,\ldots,T, \\
& \displaystyle\sum_{t=1}^{T} y_{jt} \le 1, & j = 1,\ldots,J, \\
& x_t, I_t, y_{jt} \ge 0, & t = 1,\ldots,T,\ j = 1,\ldots,J, \\
& 0 \le y_{jt} \le 1, & t = 1,\ldots,T,\ j = 1,\ldots,J, \\
& z_t \in \{0,1\}, & t = 1,\ldots,T
\end{array}$$

Merzifonluoğlu and Geunes [22] show that an optimal solution for the UOSP_{DTF} exists where (i) all orders are either completely satisfied or rejected, (ii) production for each satisfied order occurs within only one period, and (iii) full delivery of each accepted order occurs in a single period (in addition to the zero-inventory production property). Despite these properties, they show that when the revenue (R_{jt}) parameters can take arbitrary values, the UOSP_{DTF} generalizes the uncapacitated facility location problem, which implies that the UOSP_{DTF} is NP-hard (see Garey and Johnson [10]).

In practical contexts, however, we might expect that each customer order would have a preferred delivery period, and that the supplier's revenues might be nondecreasing in time prior to this preferred period, and nonincreasing thereafter. To reflect this possibility, suppose that order j has an associated preferred delivery period of t_j^p, and that the supplier receives a revenue of r_j per unit, regardless of the delivery period. In addition, suppose the supplier incurs a per-unit revenue loss of h_{jt}^c for each unit delivered in a period $t < t_j^p$ (this lost revenue may reflect a customer holding cost for early delivery). Similarly, the supplier loses a per-unit revenue equal to b_{jt}^c for units delivered in a period $t > t_j^p$ (this parameter reflects a unit backlogging cost). In addition, we make the following cost assumptions:

(1) $c_t + h_t \ge c_{t+1}$, for $t = 1,\ldots,T-1$.
(2) $h_{jt}^c \ge h_{j,t+1}^c$, for all $j \in J$ and $t < t_j^p$.
(3) $c_t + b_{jt}^c \ge c_{t-1} + b_{j,t-1}^c$, for all $j \in J$ and $t > t_j^p$.

The above assumptions imply that if we select order j, it is always beneficial to use either the latest setup prior to period t_j^p, or the earliest setup following period t_j^p. These assumptions are often referred to as equivalent to *nonspeculative motives* for holding inventory (or backlogging; see, for example, Lee et al. [20]). Under these assumptions, we can solve UOSP_{DTF} in polynomial time using an acyclic longest-path approach similar to the one discussed in the prior section. Instead of having $T+1$ nodes, the graph now contains $T+2$ nodes, where an arc of the form $(0,t')$ implies that no setup is performed prior to period t'. Given an arc (t,t') in the graph, for each order $j : t \le t_j^p \le t'$, we consider the net revenue obtained by assigning the order to the setups in periods t and t'. If the higher of these two values exceeds the associated variable cost, then the order's net contribution to profit is included in the arc's profit value (otherwise, the order is "rejected" if the longest path in the graph traverses this arc). Unlike the approach in the previous section, the profit value of an arc equals the maximum variable profit (revenue less production, holding, and backorder costs) minus the setup cost in period t (whereas in the previous section we took the maximum between this value and zero), which allows us to accurately account for the total setup cost incurred on any path in the graph. The resulting worst-case complexity of this algorithm is $\mathcal{O}(|J|T^2)$. Merzifonluoğlu and Geunes [22] also provide a dual-based heuristic solution approach for the general case in which cost and revenue parameters do not follow these specialized nonspeculative motive assumptions.

3.3. Equal-Capacity Order Selection Problem

This section briefly discusses the implications of finite production capacity levels in dynamic deterministic order selection problems. In particular, we return to the model of §3.1 under time-invariant production capacities. For a discussion of heuristic approaches for the model under time-varying capacities, please see Taaffe and Geunes [30]. The model we consider is the same as the UOSP, except that we replace constraint (3) with the constraint

$$x_t \le C z_t, \quad t = 1, \ldots, T,$$

where C denotes the production capacity in every period. We refer to the resulting problem as the equal-capacity order selection problem (ECOSP). Our solution approach for this problem builds on the results obtained by Florian and Klein [8], who provided a polynomial-time algorithm for the lot-sizing problem with equal capacities. Observe that for any set of y_{jt} values, which determine the sequence of demands the supplier will satisfy, the resulting problem is an equal-capacity lot-sizing problem (ECLSP). The structural properties of optimal solutions for the ECLSP will thus apply to optimal solutions of the ECOSP for any given choice of y_{jt} values. The remaining difficulty then lies in determining a manageable characterization of the orders that will be selected in an optimal solution.

Our approach for this problem relies on the optimality of *capacity-constrained sequences* for the ECLSP. Given a *regeneration interval* (i.e., a sequence of periods $\tau, \ldots, \tau'$ such that $I_\tau = I_{\tau'} = 0$ and $I_t > 0$ for $t = \tau + 1, \ldots, \tau' - 1$), a capacity-constrained sequence is one in which production output x_t equals either 0 or C in all periods from $\tau + 1$ to $\tau' - 1$, except for at most one of these periods. If we are required to satisfy a total demand of $D(\tau, \tau')$ in periods $\tau + 1, \ldots, \tau' - 1$, then for the period in the capacity-constrained sequence in which we produce neither 0 nor C, we must produce ϵ, where $\epsilon = D(\tau, \tau') (\text{mod } C)$. Given $D(\tau, \tau')$, we also know that the total number of setups that must occur in periods $\tau + 1, \ldots, \tau' - 1$ must equal $k + 1$, where $k = \lfloor D(\tau, \tau')/C \rfloor$. Florian and Klein [8] show that given $D(\tau, \tau')$, the subproblem to determine the optimal production plan for regeneration interval (τ, τ') can be solved in $\mathcal{O}(T^2)$ time.

Geunes et al. [11] provide results that allow considering a limited number of demand values for every potential regeneration interval. That is, given a potential regeneration interval (τ, τ'), efficiently solving ECOSP requires consideration of a manageable number of possible $D(\tau, \tau')$ values. To do this, we define an adjusted revenue parameter, $\rho_{jt} = r_{jt} + \sum_{s=t}^{T} h_s$. Given a regeneration interval (τ, τ'), Geunes et al. [11] show that if $\rho_{jt} \ge \rho_{it'}$ (with both t and t' between τ and τ'), then if an optimal regeneration interval solution exists with $y_{it'} > 0$, an optimal solution exists with $y_{jt} = 1$ (equivalently, if $y_{jt} < 1$, then $y_{it'} = 0$). This implies that within any regeneration interval we will have at most one order that is partially satisfied. In addition, Geunes et al. [11] show that when an order is partially satisfied within a regeneration interval, then all production levels must be at 0 or C. Similarly, when one period within a regeneration interval exists in which production is at neither 0 nor C, all orders within the regeneration interval will be either fully satisfied or completely rejected.

These results imply that there are $\mathcal{O}(J_{\max} T)$ candidate demand vectors that we must consider within any regeneration interval. Given that we can determine the optimal production plan for a given demand vector and regeneration interval in $\mathcal{O}(T^2)$ time, the time required to solve a regeneration interval subproblem (considering all candidate demand vector values) is bounded by $\mathcal{O}(J_{\max} T^3)$. The number of possible regeneration intervals is $\mathcal{O}(T^2)$, which implies that the ECOSP can be solved in $\mathcal{O}(J_{\max} T^5)$ time in the worst case. Note that this result improves the worst-case complexity found in Geunes et al. [11] by an order of magnitude, based on a recent analysis and improvement of the algorithm.

The analogous pricing problem follows from our discussion in §3.1. That is, within each period we sort orders in nonincreasing order of revenue values (r_{jt}s), and take these values as the slopes of the piecewise-linear and concave total revenue function (as a function of total demand satisfied). The width of the jth segment of the revenue function in period t equals d_{jt}. Observe that, unlike the UOSP pricing analog, under capacity constraints it is possible that an optimal solution exists such that the optimal amount of total demand satisfied in a period occurs at a value that is between breakpoints of the piecewise-linear revenue curve.

3.4. ECOSP with Capacitated Overtime and Subcontracting

Firms with internal production capacity limits often have options for exceeding that capacity through, for example, the use of overtime and/or subcontracting. Such options in order selection contexts may permit a supplier to accept a greater amount of total demand to increase total revenue. In this section, we generalize the ECOSP to address the trade-off between the additional capacity costs associated with overtime and subcontracting, and the potential for increasing revenue by using these additional sources of capacity. Atamtürk and Hochbaum [1] considered the dynamic economic lot-sizing problem with subcontracting options, and our approach is structurally similar to theirs, although they did not consider an order selection (or the analogous pricing problem) dimension, or the availability of overtime as a source of additional capacity.

Overtime options are often viewed as an alternative form of subcontracting from a modeling perspective, although in practice, distinct differences typically exist. For example, a firm can choose to completely subcontract without using any internal capacity, whereas overtime is only used after exhausting regular internal capacity. Moreover, if regular-time production capacity is limited, then some associated limit on overtime capacity will also likely exist, and the costs of regular-time and overtime production will contain a similar structure (with the overtime cost rate being higher than the regular-time cost rate). The model we present in this section captures these subtle but important differences between overtime and subcontracting costs.

We generalize the cost model used in the ECOSP by allowing the production cost function in period t, which we denote by $p_t(x_t)$, to be concave and nondecreasing in x_t (as opposed to the fixed-charge structure we employed in the previous section), where x_t is now defined as the total production from internal sources (regular and overtime production) in period t. Letting v_t and ε_t denote the regular and overtime production output in period t, respectively, we have $x_t = v_t + \varepsilon_t$. We define the function $o_t(\varepsilon_t)$ as an *incremental* cost function for units produced during overtime, where $o_t(\varepsilon_t)$ is also concave and nondecreasing. The amount of *overtime capacity* in a period in this model equals C, and the amount of *regular-time capacity* equals κC, where κ is a positive integer. The amount of available overtime capacity is therefore some fraction $1/\kappa$ of the regular capacity in a period (the model and results presented here also apply when available overtime capacity is a positive integer multiple of regular capacity).

The total production cost is therefore equal to $p_t(x_t)$ when $x_t \leq \kappa C$, and $p_t(x_t) + o_t(\varepsilon_t)$ when $C < x_t \leq \kappa C$. The regular production cost function $p_t(x_t)$ accounts for the basic production cost structure for using internal capacity, while the incremental overtime cost function $o_t(\varepsilon_t)$ accounts for the incremental cost of units produced during overtime hours. Thus, the structure of the total production cost function in a period depends on the amount of regular internal capacity C (Merzifonluoğlu et al. [23] consider additional model extensions in which additional regular capacity is available to the supplier at a cost, and this capacity level therefore becomes a decision variable). In effect, because economies of scale in production exist, the total cost of a unit of overtime production depends on how much regular-time production is available.

In addition to overtime options, we assume that an external subcontractor is available for producing the good, and define $g_t(s_t)$ as the subcontracting cost function in period t, where s_t is the amount subcontracted in the period (we assume that no limit exists on the amount of subcontracting available in any period). We formulate the ECOSP with overtime and subcontracting options (ECOSP_{OS}) as follows.

$$
\begin{aligned}
[\mathbf{ECOSP_{OS}}]\quad \text{minimize}\quad & \sum_{t=1}^{T}(p_t(v_t+\varepsilon_t)+o_t(\varepsilon_t)+g_t(s_t)-h_tI_t)-\sum_{t=1}^{T}\sum_{j=1}^{J(t)} r_{jt}d_{jt}y_{jt} \\
\text{subject to:}\quad & \sum_{j=1}^{J(t)} d_{jt}y_{jt}+I_t=v_t+\varepsilon_t+s_t+I_{t-1}, && t=1,\dots,T, \\
& v_t\le \kappa C, && t=1,\dots,T, \\
& \varepsilon_t\le C, && t=1,\dots,T, \\
& 0\le y_{jt}\le 1, && t=1,\dots,T,\ j=1,\dots,J(t), \\
& v_t,\varepsilon_t,s_t,I_t\ge 0, && t=1,\dots,T
\end{aligned}
$$

The ECOSP_{OS} minimizes a concave function over a polyhedron, which implies that an extreme-point optimal solution exists. We next provide a characterization of the extreme points of the above polyhedron, which allows us to provide a polynomial-time solution approach. The following lemma paves the way for solving a set of shortest-path subproblems in order to determine the optimal order selection and production-planning decisions for a regeneration interval.

Lemma 1 (Merzifonluoğlu et al. [23]). *In an extreme-point solution for* $[ECOSP_{\text{OS}}]$*, a regeneration interval can have at most one period* t *with* $0<v_t<\kappa C$ *(fractional regular-time production), at most one period* t *with either* $0<\varepsilon_t<C$ *(fractional overtime production) or* $s_t>0$ *(subcontracting), or at most one period with* $0<y_{jt}<1$ *(partial demand satisfaction), but not any of these simultaneously.*

Lemma 1 implies that given a potential regeneration interval (t,t'), we need to consider four types of solutions:

(1) One period containing fractional regular production with overtime production always at 0 or C, zero subcontracting, and all orders fully satisfied or rejected;

(2) One period containing fractional overtime production with regular production always at 0 or κC, zero subcontracting, and all orders fully satisfied or rejected;

(3) One period containing subcontracting, with regular-time (overtime) production always at 0 or κC (0 or C), and all orders fully satisfied or rejected;

(4) One period containing a fractionally satisfied order with regular-time (overtime) production always at 0 or κC (0 or C), and zero subcontracting.

For each of these types of regeneration intervals, and given any potential regeneration interval (t,t'), we can determine the optimal regeneration interval solution of each type in polynomial time. This solution relies on using the adjusted revenue parameters ρ_{jt}, described in the previous section, to measure the relative preferences of orders within any regeneration interval. For regeneration interval types (1)–(3), because we know that all orders are either fully satisfied or rejected, we have $\mathcal{O}(J_{\max}T)$ candidate demand vectors for the regeneration interval demand $D(t,t')$. Given any one of these candidate demand vectors, for each type of regeneration interval, we can determine a discrete number of possible cumulative production quantities throughout the regeneration interval. We use a layered graph with $t'-t$ layers (one for each period within the regeneration interval, plus a start node), where a node exists in each layer for every possible cumulative production value. Arcs connect each layer to the next-highest layer based on feasible cumulative production transitions between periods. We then solve a shortest-path problem on this network to determine the best regeneration interval solution for each type. Merzifonluoğlu et al. [23] show that each of these shortest-path graphs is solvable in polynomial time, and we must solve $\mathcal{O}(J_{\max}T)$ of these shortest-path problems within each possible regeneration interval. For regeneration interval type (4), although at most one order may be satisfied at a fractional value, we know that cumulative production in a regeneration interval must be an integer multiple of C. Given an integer multiple of C, we can then determine which orders (demand) to satisfy based on filling the capacity with orders in nonincreasing order of the ρ_{jt} values. Merzifonluoğlu et al. [23] show that the worst-case complexity for solving the regeneration interval subproblem is $\mathcal{O}(\kappa^2 J_{\max}T^4)$, which leads to the following theorem.

Theorem 1 (Merzifonluoğlu et al. [23]). *A polynomial-time algorithm exists for the* $ECOSP_{OS}$ *with worst-case complexity of* $\mathcal{O}(\kappa^2 J_{\max}T^6)$.

4. Multifacility Demand Assignment Models

Our focus thus far has dealt exclusively with single-stage models. This section illustrates the role these models might play in a larger tactical distribution network design context. We consider two-echelon problems containing an upstream supply echelon and a downstream demand echelon. The upstream echelon might, for example, consist of production/distribution (P/D) facilities, while the downstream echelon might correspond to retail sites. Our model takes the supplier perspective, considering upstream production and inventory holding costs, plus any additional costs (such as transportation costs) associated with meeting retailer demands. We assume the supplier has a number of geographically dispersed P/D facilities, each of which can produce a given product. The supplier faces demands from a number of retail locations for the product and wishes to determine which source facility will serve each retail site. Our approach assumes that the supplier uses a *single-sourcing* strategy, with each retail site being served by (or assigned to) a single upstream P/D facility. While some distribution systems might employ transhipments in extraordinary situations, as a primary strategy, this single-sourcing approach is often found in practice because it reduces coordination complexity and decreases information systems integration requirements. Our model therefore applies to contexts where such single-sourcing is the primary operating approach. Due to space limitations, we provide a sketch of our overall solution approach for these problems. For more details on this approach, and in particular how it can be applied using a cost model of the form discussed in §§2.1 and 2.2, please see Huang et al. [16] and Taaffe et al. [32].

4.1. Model Formulation

Our modeling approach begins with a general set-partitioning formulation that partitions a set of n retailers into m subsets, where m is the number of P/D facilities. After formulating this set-partitioning problem, we discuss a column generation approach in which the column-pricing problem brings us back to the single-stage models that have served as the primary focus of this chapter. We begin by letting K_i denote the number of subsets of retailers that can be assigned to facility i, for $i = 1, \ldots, m$. We define $\gamma_{i\cdot}^k$ as an indicator vector corresponding to the kth feasible subset of retailers at facility i, where $\gamma_{ij}^k = 1$ if retailer j is included in facility i's kth feasible subset. The function $H_i(\gamma_{i\cdot}^k)$ (where $H_i: \{0,1\}^n \to \mathbb{R} \cup \{+\infty\}$) provides the cost of assigning the kth feasible subset to facility i, and we let y_i^k equal 1 if this assignment is made, and 0 otherwise. Constraints may be incorporated into this cost function (H_i) by assigning a cost of $+\infty$ to any infeasible assignment. At this point we do not provide any further specification of the $H_i(\cdot)$ functions, although we will return to this later when we discuss the pricing problem. For now, we simply assume that given $\gamma_{i\cdot}^k$, we can compute $H_i(\gamma_{i\cdot}^k)$. We formulate our set-partitioning (SP) formulation as follows.

$$
\begin{array}{lll}
[\mathbf{SP}] \quad \text{minimize} & \displaystyle\sum_{i=1}^{m}\sum_{k=1}^{K_i} H_i(\gamma_{i\cdot}^k) y_i^k & \\
\text{subject to:} & \displaystyle\sum_{i=1}^{m}\sum_{k=1}^{K_i} \gamma_{ij}^k y_i^k = 1, & j = 1,\dots,n, \\
& \displaystyle\sum_{k=1}^{K_i} y_i^k = 1, & i = 1,\dots,m, \\
& y_i^k \in \{0,1\}, & i = 1,\dots,m,\ k = 1,\dots,K_i
\end{array}
$$

The objective function minimizes the sum of the assignment costs at each facility, while the constraints ensure that every retailer is assigned to a facility and that some subset is assigned to each facility. Our solution approach begins with some subset $\widehat{K}_i \subset K_i$ $(i = 1,\dots,m)$ that admits a feasible solution to the linear programming relaxation of [SP], which we call LP(SP). After solving LP(SP) with the subset of columns, we then solve a pricing problem for each facility i. If the optimal solution value to the pricing problem for facility i is negative, the resulting new column is added to the set $\widehat{K}_i$, and we repeat the process. If all m pricing problems produce nonnegative solution values, then we have an optimal solution to LP(SP). If this solution is integral, then we have an optimal solution for [SP]. Otherwise, we apply a branch-and-price algorithm to solve [SP] to optimality (see Huang et al. [16] for details on this branch-and-price algorithm). The next section discusses the pricing problem for LP(SP).

4.2. The Pricing Problem

Let u_j^* and v_i^* denote optimal dual vectors corresponding to the constraints in LP(SP) for some set of $\widehat{K}_i$ values. The pricing problem for facility i is written as

$$
\begin{array}{ll}
[\mathrm{PP}_i] \quad \text{minimize} & \displaystyle H_i(z) - \sum_{j=1}^{n} u_j^* z_j + v_i^* \\
\text{subject to:} & z_j \in \{0,1\}, \quad j = 1,\dots,n,
\end{array}
$$

where an optimal solution $z^* = (z_1^*,\dots,z_n^*)^\top$ provides a new member for the set $\widehat{K}_i$ when $H_i(z^*) - \sum_{j=1}^{n} u_j^* z_j^* + v_i^* < 0$. If we compare the pricing problem formulation [PP$_i$] to the [EOQMC] formulation of §2.1 or the [SNP] formulation of §2.2, we see that through a proper definition of the $H_i(\cdot)$ function, the pricing problem becomes equivalent to these problems. Taken in the context of these models, the z_j variables correspond to market (or retailer) selection variables, and the dual price u_j^* provides either an additional revenue term for selecting market (retailer) j (if $u_j^* > 0$) or a cost penalty (if $u_j^* < 0$). Thus, if the facility cost (and revenue) structures are consistent with the [EOQMC] or [SNP] model, we have efficient methods for solving the pricing problem within a larger distribution network design problem. Huang et al. [16] directly address the case in which facility costs take the EOQ cost structure, and provide computational test results using the branch-and-price procedure, which they show can solve problems of reasonably large size (up to 10 facilities and 50 retailers) in acceptable computing time (1–4 minutes). Taaffe et al. [32] provide motivating contexts and a solution approach when supplier cost structures are consistent with those of the [SNP] model.

Because we have defined these cost functions $H_i(\cdot)$ very generally, and by incorporating the proper cost structures and constraints in the definition of these functions, the pricing problem for each facility i can in principle also incorporate the class of dynamic order selection problems covered in §3. In this case, instead of the downstream demand echelon consisting of facilities, it might consist of a pool of time-phased (or time-flexible) orders that the supplier firm must satisfy, where facility cost structures are consistent with those defined in §3. One caveat applies, however, when considering the order selection models of §3 in general. In the absence of production capacities, we know that all orders will be either fully accepted or rejected by a facility. Our capacity-constrained models, however, allowed for partial satisfaction of orders. If orders must be completely accepted or rejected by facilities under capacity constraints (which would likely be the case in the distribution network design problem), the resulting complexity of the pricing problems (which then generalize the NP-hard knapsack problem), might preclude the column generation approach (which then needs to be embedded in a branch-and-price algorithm) as a viable solution method. However, the results in Huang et al. [16] indicate that high-quality solutions can often still be found in reasonable time in such cases. Further research is required to determine situations in which capacity-constrained facilities with all-or-nothing order satisfaction requirements would be amenable to this solution method.

5. Concluding Remarks and Future Research Directions

This chapter provided an overview of a collection of related models that address a supplier's ability to influence the collective demands to which it will respond. The majority of these models apply to single-stage, single-product contexts, and generalize classical inventory models to address elements of demand selection. While we posed the problem in the context of market (or order) selection decisions, we can view such selection decisions as an explicit mechanism for shaping demand characteristics. For certain cases, we discussed a different interpretation of the models in which pricing served as an implicit mechanism for effectively selecting demand levels. As the previous section illustrated, these demand selection problems also arise as pricing subproblems in distribution network design problems. The solution approaches we have developed for the single-stage models can therefore serve as a subroutine for solving these larger network design models.

As we noted in the previous section, the dynamic models discussed in §3 can be, in principle incorporated as subproblems in a distribution network design problem that accounts for economies of scale at facilities that must collectively respond to a firm's set of orders over multiple periods. Exploring models under which this solution method is efficient for this class of problems serves as a potential avenue for further research. When the context requires that all orders are either accepted in full or completely rejected, further methodologies must be explored that can handle the added complexity this implies. Our distribution network design approach can also be extended to address more comprehensive supply chain cost structures, where economies of scale exist in transportation between echelons, and inventory costs at both echelons of the supply chain are incorporated.

Acknowledgments

This research was partially supported by NSF Grants DMI-0355533 and DMI-0322715.

References

[1] A. Atamtürk and D. S. Hochbaum. Capacity acquisition, subcontracting, and lot sizing. *Management Science* 47(8):1081–1100, 2001.

[2] Y. Aviv. The effect of collaborative forecasting on supply chain performance. *Management Science* 47(10):1326–1343, 2001.

[3] S. Biller, L. M. A. Chan, D. Simchi-Levi, and J. Swann. Dynamic pricing and the direct-to-customer model in the automotive industry. *Electronic Commerce Research* 5(2):309–334, 2005.

[4] S. Carr and W. Lovejoy. The inverse newsvendor problem: Choosing an optimal demand portfolio for capacitated resources. *Management Science* 46(7):912–927, 2000.

[5] K. Charnsirisakskul, P. M. Griffin, and P. Keskinocak. Order selection and scheduling with leadtime flexibility. *IIE Transactions* 36(7):697–707, 2004.

[6] R. E. Crandall and R. E. Markland. Demand management—Today's challenge for service industries. *Production & Operations Management* 5(2):106–120, 1996.

[7] G. Eppen. Effects of centralization on expected costs in a multi-location newsboy problem. *Management Science* 25(5):498–501, 1979.

[8] M. Florian and M. Klein. Deterministic production planning with concave costs and capacity constraints. *Management Science* 18(1):12–20, 1971.

[9] G. Gallego and G. van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science* 40(8):999–1020, 1994.

[10] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, 1979.

[11] J. Geunes, H. E. Romeijn, and K. Taaffe. Requirements planning with pricing and order selection flexibility. *Operations Research*. Forthcoming.

[12] J. Geunes, Z.-J. Shen, and H. E. Romeijn. Economic ordering decisions with market choice flexibility. *Naval Research Logistics* 51(1):94–112, 2004.

[13] S. M. Gilbert. Coordination of pricing and multi-period production for constant priced goods. *European Journal of Operational Research* 114:330–337, 1999.

[14] G. Hadley and T. M. Whitin. *Analysis of Inventory Systems.* Prentice Hall, Englewood Cliffs, NJ, 1963.

[15] F. W. Harris. How many parts to make at once. *Factory, The Magazine of Management* 10:135–136,152, 1913.

[16] W. Huang, H. E. Romeijn, and J. Geunes. The continuous-time single-sourcing problem with capacity constraints and expansion opportunities. *Naval Research Logistics* 52(3):193–211, 2005.

[17] A. V. Iyer, V. Deshpande, and Z. Wu. A postponement model for demand management. *Management Science* 49(8):983–1002, 2003.

[18] P. Keskinocak, R. Ravi, and S. Tayur. Scheduling and reliable lead time quotation for orders with availability intervals and lead-time sensitive revenues. *Management Science* 47(2):264–279, 2001.

[19] H. Kunreuther and L. Schrage. Joint pricing and inventory decisions for constant priced items. *Management Science* 19(7):732–738, 1973.

[20] C. Y. Lee, S. Çetinkaya, and A. P. M. Wagelmans. A dynamic lot sizing model with demand time windows. *Management Science* 47(10):1384–1395, 2001.

[21] G. L. Lilien and A. Rangaswamy. *Marketing Management and Strategy: Marketing Engineering Applications*, 2nd ed. Prentice Hall, Englewood Cliffs, NJ, 2003.

[22] Y. Merzifonluoğlu and J. Geunes. Uncapacitated production and location planning models with demand fulfillment flexibility. *International Journal of Production Economics*. Forthcoming.

[23] Y. Merzifonluoğlu, J. Geunes, and H. E. Romeijn. Production planning models with variable demand and capacity. Working paper, University of Florida, Gainesville, FL, 2005.

[24] D. R. Moodie. Demand management: The evaluation of price and due date negotiation strategies using simulation. *Production & Operations Management* 8(2):151–162, 1999.

[25] S. Nahmias. *Production and Operations Analysis*, 5th ed. McGraw-Hill Irwin, Boston, MA, 2005.

[26] N. P. Petruzzi and M. Dada. Pricing and the newsvendor problem: A review with extensions. *Operations Research* 47(2):183–194, 1999.

[27] E. L. Porteus. Stochastic inventory theory, Ch. 2. D. P. Heyman and M. J. Sobel, eds. *Stochastic Models. Handbooks in Operations Research and Management Science*, Vol. 2. North-Holland, Amsterdam, The Netherlands, 1990.

[28] Z.-J. Shen, C. Coullard, and M. Daskin. A joint location-inventory model. *Transportation Science* 37(1):40–55, 2003.

[29] E. A. Silver, D. F. Pyke, and R. Peterson. *Inventory Management and Production Planning and Scheduling*, 3rd ed. John Wiley & Sons, New York, 1998.

[30] K. Taaffe and J. Geunes. Models for integrated order selection and requirements planning decisions under limited production capacity. P. M. Pardalos, A. Migdalas, and G. Baourakis, eds. *Supply Chain and Finance*. Series on Computers and Operations Research, Vol. 2. World Scientific, River Edge, NJ, 2004.

[31] K. Taaffe, J. Geunes, and H. E. Romeijn. Integrated marketing and operational decisions: The selective newsvendor problem. Working paper, University of Florida, Gainesville, FL, 2005.

[32] K. Taaffe, J. Geunes, and H. E. Romeijn. Capacity acquisition and stochastic customer demand assignment in a network of facilities. *Proceedings of the 2005 Annual IE Research Conference*, Atlanta, GA, 2005.

[33] J. Thomas. Price-production decisions with deterministic demand. *Management Science* 16(11):747–750, 1970.

[34] W. van den Heuvel and A. P. M. Wagelmans. A polynomial time algorithm for a deterministic joint pricing and inventory model. *European Journal of Operational Research*. Forthcoming.

[35] H. Wagner and T. Whitin. Dynamic version of the economic lot size model. *Management Science* 5:89–96, 1970.

[36] L. R. Weatherford and S. E. Bodily. A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking and pricing. *Operations Research* 40(5):831–844, 1992.