

Capacity Planning in a General Supply Chain with Multiple Contract Types – Single Period Model

Xin Huang • Stephen C. Graves

Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts, 02139, USA

Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts, 02139, USA

xinhuang@mit.edu • sgraves@mit.edu

A key element to a company's success is its ability to match its supply to uncertain demand by utilizing different types of capacity contracts. We consider the design of a multi-product supply chain, for which each product requires one or more process capabilities and each process can get capacity from multiple resources. We present a single-period model to determine the capacity investments for these resources in the presence of demand uncertainty and option contracts. (We have also extended the model to a multi-period setting in a companion paper) We derive closed-form solutions to two important special cases of the problem and draw managerial insights about the optimal capacity planning strategy. We then develop a stochastic linear programming algorithm to solve the general single period problem and show that our algorithm outperforms alternative algorithms by means of an empirical study. Finally, with the model and algorithm, we study the effects of common processes and option contracts on capacity planning.

Key words: capacity planning, demand uncertainty, multi-stage supply, option contract

1. Introduction

In today's competitive economic environment, customers do not just prefer but demand manufacturers to provide quality products in a timely fashion at competitive prices. To satisfy this requirement, manufacturers need to plan necessary and sufficient capacity to meet market demands. However, capacity planning is a very challenging task for many reasons.

Demand Uncertainty. For most industries, it is very difficult to accurately forecast the demand for new products. In an emerging industry, manufacturers devote substantial efforts to studying the applications and benefits of new technologies. However, when a technology is new, firms have little information on the commercial uptake of new products and, therefore,

have poor forecasts of the product demand. For example, GlobalStar, one of the key players in the emerging mobile satellite services industry during the 1990s, expected between 500,000 and 1,000,000 users in 1999, the first year of its operation; these numbers were confirmed by many other independent analysts. However, the actual number of users was only 100,000, which is significantly lower than the expectation. Because the demand forecast was overly optimistic, the company filed for bankruptcy protection with a debt of 3.34 billion dollars in 2002 after three years of operations [Weck et al. 04].

Demand forecasts for new products can also be inaccurate in existing industries. Customers' tastes and preferences are hard to predict and will change over time. Therefore, the historical demand patterns for an existing product might not always be a good reference for the next generation of products. For example, when Mercedes-Benz first introduced its M-class cars in 1997, it forecasted its annual demand to be about 65,000 vehicles. This forecast was, in fact, too low and the firm expanded its capacity to 80,000 vehicles during 1998-1999, which was also insufficient to meet demand [Van Mieghem 07].

Large Scale. Manufacturers face the difficulty of planning resources for multiple products at the same time. Due to competition and the wide range of applications of a new technology, the manufacturer needs to produce a variety of generic or custom-made products to meet the requirements of its customers. Such variety in products adds complexity to a manufacturer's supply chain. Different products might share common manufacturing processes or use common components. Because of the linkage between the products, the manufacturer needs to plan together its capacity for producing multiple products, as well as possibly multiple generations of a single product. However, finding the right level of capacity for all products at the same time is a large scale problem. A manufacturer, therefore, would benefit from efficient and practical algorithms for solving large scale capacity planning problems.

Option Contracts. In addition to demand uncertainty and large problem size, manufacturers can also benefit from models and tools that can incorporate option contracts into capacity planning. A manufacturer might establish a fixed-price capacity contract with its supplier to rent or reserve a fixed amount of capacity. The manufacturer needs to pay for the capacity whether or not it uses the capacity. In practice, the supplier's cost of capacity might have two components: a fixed cost and a variable cost. For example, equipment costs and the monthly salaries of workers are fixed costs, while power consumption and employee overtime payments are variable costs. An option contract separates these two types of costs.

With option contracts, the manufacturer buys the rights to use a fixed amount of capacity with an upfront fixed payment. If it decides to execute its rights and use the capacity, it needs to pay an exercise price for each unit of capacity that it actually uses.

Option contracts have been in practice for a long time. The manufacturer will often make a deposit to its supplier once both sides agree on a contract. When the supplier delivers the products, the manufacturer will pay the remaining payment. If the manufacturer withdraws from the contract, the deposit will serve as the penalty cost. In these situations, the deposit is equivalent to the upfront payment in an option contract, and the difference between the full payment and deposit is the exercise price.

There are several reasons why both manufacturers and suppliers might prefer an option contract, rather than a fixed-cost contract. For the manufacturer, option contracts can serve as a tool to reduce the risk of committing upfront to a certain amount of capacity at a fixed price. In the context of outsourcing contracts, the manufacturer might want to secure the availability and price of the capacity. However, when demand is lower than expected, committing to a fixed amount of capacity will result in excess capacity. Moreover, if the price of capacity falls, the manufacturer will pay more than its competitors to make the products. Using option contracts can reduce the risk of weak demand and price volatility. For example, Hewlett-Packard has implemented a Procurement Risk Management (PRM) system to utilize option contracts and has realized \$425 million savings in cost over a six-year period [Nagali et al. 08].

From the other side, a supplier can secure higher revenue by taking advantage of option contracts. Since an option contract serves as a hedging tool to protect the downside of its operation, the manufacturer might be willing to pay more for each unit of option capacity, which means that the reservation price plus the exercise price is higher than the fixed-price contract price. Moreover, since the manufacturer bears lower risk, it might purchase more capacity. As a result, the supplier can gain more revenue. Therefore, a method to incorporate option contracts into capacity planning will also be one of the manufacturers' primary interests.

In this paper, we present a mathematical model and tools to help manufacturers plan their capacity under demand uncertainty for a general large-scale supply chain with option contracts. Compared to the existing literature, our model provides a more comprehensive system to study capacity planning. We have developed efficient and practical algorithms to address the following three questions: which suppliers should the manufacturer select,

which types of contracts should it use, and how much capacity should it reserve. Using the model and algorithms, we study the properties of, and draw managerial insights about, the optimal capacity planning strategy. Therefore, our research can help managers to make these complex capacity planning decisions in a more systematic and effective way.

The paper is organized as follows: We review relevant literature in Section 2. In Section 3, we outline a mathematical model for the single period capacity planning problem. We then look at two special cases and derive closed-form solutions for the optimal strategies for these cases in Section 4. After that, in Section 5, we examine several algorithms to solve the general single period capacity problem and show that the one we develop has a better run time through a series of randomly generated test problems. Finally, in Section 6, we discuss the effects of common process and option contract on capacity planning.

2. Related Literature

The research in this paper is related to the literature in three areas: Newsvendor Network and Assembly to Order (ATO) Systems, Option Contracts, and Stochastic Programming.

Newsvendor Network and Assembly to Order (ATO) Systems. Van Mieghem and Rudi [Van Mieghem and Rudi 02] propose a newsvendor network that is closely related to the model that we use. In their model, the authors consider a supply chain that contains multiple products and multiple stocks. The manufacturer consumes the stocks to produce the products through activities. The stocks are subject to inventory constraints and the activities are subject to capacity constraints. They study a joint capacity investment and inventory management problem. The capacity investment decision is made at the beginning of the planning horizon and remains in effect thereafter. They show that the capacity planning problem is concave, and therefore they can apply concave optimization algorithms such as subgradient methods to find the optimal capacity plan.

In contrast to their work, our model allows the manufacturer to establish different types of contracts with its suppliers. These contracts can differ in price and structure (such as fixed-cost contract and option contract). Moreover, their paper focuses on the structure of the optimal inventory replenishment policy, while we focus on how to solve the capacity problem. We discuss different concave optimization algorithms, which include the sub-gradient method suggested by Van Mieghem and Rudi. We show that the algorithm that we propose has a superior performance.

In terms of modelling the supply chain, the model that we propose shares some commonality with the assemble-to-order (ATO) systems in the supply chain literature. An ATO system contains multiple products and multiple components. The system only keeps inventory at the component level. When demand arrives, it will assemble products using the necessary components. ATO systems capture some of the essential characteristics of a real life supply chain, such as common processes (e.g. [Gerchak et al. 88], [Hillier 00], and [Kulkarni et al. 04]) and flexible resources (e.g. [Fine and Freund 90], [Van Mieghem 98], and [Labro 04]). For a detailed survey and discussion of ATO systems, please refer to [Song and Zipkin 03].

There are several major differences between ATO systems and our supply chain capacity model. First, our model has a multi-tier structure that allows both flexible resources and common processes. Second, we incorporate option contracts into the model. Third, our model focuses on capacity planning while ATO systems mainly study inventory policies.

Option Contracts. The consideration of option contracts in supply chains is a more recent research topic. Cheng et al. [Cheng et al. 07] derive the optimal order decision for the manufacturer and the optimal pricing decision for the supplier in a single product, single supplier, and single period supply chain. Yazlali and Frhun [Yazlali and Frhun 06] consider option contracts in a single product, dual supply, and multi-period problem. They use a two-stage decision process: first, the manufacturer reserves capacity for the whole planning horizon by signing a portfolio of contracts; second, it orders from the suppliers based on the contracts. Under certain assumptions on demands and prices, they show that for the second stage problem, a two-level modified base-stock policy is optimal, and, for the first stage, a reserve-up-to policy is optimal. Martinez-de-Albéniz and Simchi-Levi [Martinez-de-Albéniz and Simchi-Levi 05] analyze the optimal option contract for a case of single product and multiple suppliers in the presence of a spot market. In their model, they also adopt a two-stage decision process. The manufacturer decides the quantity and portfolio of contracts at the beginning of the planning horizon. They show that the portfolio selection problem is a concave maximization problem. Fu et al. [Fu et al. 06] examine a single-period procurement problem with option contracts. Their model incorporates random spot prices and demands. They show that option contracts can be very valuable for both the manufacturer and supplier. Nagali et al. [Nagali et al. 08] apply option contracts in HP's procurement risk management; the system that they implemented has realized more than \$425 million cost savings in a six year period. However, they do not provide details on the




Figure 1: A supply chain network with 3 products, 5 processes, and 7 resources.

specific models that are used for evaluating these option contracts.

Compared to the existing literature studying option contracts, we consider a more general supply chain structure that contains multiple products, multiple processes, and multiple resources. However, our model takes the external market conditions as given and does not consider inventory.

Stochastic Programming. Finally, our work is related to the literature studying algorithms for stochastic linear programming. Higle and Sen [Higle and Sen 96] provide an excellent review of how to apply stochastic linear programming to solve large scale capacity planning problems. Higle and Sen [Higle and Sen 91] and Higle and Sen [Higle and Sen 96] propose and summarize several stochastic linear programming algorithms to solve a general capacity planning problem. We adapt some of these techniques in our algorithm for solving our single period capacity planning problem. We show that the algorithm we propose has a better performance than the ones that Higle and Sen suggest through a series of randomly generated test cases.

3. Model

3.1 Mathematical Model

We consider a multi-product and multi-stage supply chain consisting of M products, J processes, and K resources. A sample supply chain network with three products, five processes, and seven resources is given in Figure 1. The production of each product requires a certain amount (possibly zero) of each type of process. The solid links joining products and

processes in Figure 1 signify this relationship. For example, product 1 requires processes 1, 2, and 5. In practice, a process can be either an operation such as assembly, testing, or packaging or a type of material or component or a sub-system that is required to produce the product. A resource provides capacity for one or more processes. The dashed links joining processes and resources in Figure 1 signify that the resource has the capability to deliver the process. For the network given in Figure 1, the firm can get capacity for process 1 from resource 1 or 3 and resource 3 can provide capacity to processes 1, 2, and 4. A resource might be an assembly line with the capability to assemble a single product type. A flexible resource might be an assembly line capable of assembling several different product types. We might also imagine a resource with capability to provide more than one type of process; for instance, a resource might do both assembly and test for a single product type. Without loss of generality, we assume that the production of one unit of product requires one unit of each of its required processes; we also assume that one unit of each process requires one unit of capacity from one of its resource options.

The supply chain structure that we propose for the single period problem is fairly general and can capture different types of interdependency between products, processes, and resources. First, to produce a product requires capacities from all of its processes. Therefore, the capacity levels of different processes of the same product are closely related to each other. Second, different products can share common processes. Third, flexible resources can provide capacity to different processes. These common processes and flexible resources link the capacity planning decisions of different products together. One of our goals is to account for these interdependencies within capacity planning.

In addition to a general supply chain structure, we also consider two alternatives for procuring or reserving capacity for each resource: A firm can reserve capacity on a resource with a fixed-price capacity contract; alternatively a firm can reserve capacity on a resource with an option contract where there is a smaller upfront reservation price and then a variable exercise price for the use of this capacity. For instance, under a fixed-price capacity contract, the price for one unit of capacity is 1 dollar. Under an option contract, the firm might pay a fixed price of 30 cents initially to reserve each unit of the capacity. If the firm decides to use the capacity that it has reserved, it needs to pay another 80 cents per unit. Given these alternatives, the firm wants to determine the amount of each resource to use or reserve, as well as the contracts, so that the resulting supply chain maximizes the firm's expected profit.

We assume that any demand that cannot be filled is lost, and there is no additional

penalty cost for not meeting demand. We also assume a two-stage sequential decision process. In the first stage, the firm determines the types and sizes of the contracts for each resource; in effect the firm decides its capacity plan. In the second stage, demand is realized and the firm decides how to allocate its production capacity to meet demand. To the extent that the firm employs options contracts, it will decide how much of each option to exercise. Also, the firm decides how to utilize the capacity of each flexible resource across the applicable processes.

For naming convention, we use bold letter to indicate a vector. For input parameters, we denote:

A An $J \times M$ matrix such that

$$A(j, m) = \begin{cases} 1, & \text{if product } m \text{ requires process } j; \\ 0, & \text{otherwise.} \end{cases}$$

B An $J \times JK$ matrix such that

$$B(j, (j, k)) = \begin{cases} 1, & \text{if resource } k \text{ can provide capacity to process } j; \\ 0, & \text{otherwise.} \end{cases}$$

H A $K \times JK$ matrix such that

$$H(k, (j, k)) = \begin{cases} 1, & \text{if resource } k \text{ can provide capacity to process } j; \\ 0, & \text{otherwise.} \end{cases}$$

D A vector of random variables, with probability density function, that represents the demand of products. (Vector of size M)

d A realization of random demand *D*. (Vector of size M)

r Unit profit for filling product demand. (Vector of size M)

p Unit price of resources under fixed-price contract. (Vector of size K)

q Unit reservation price of resources under option contract. (Vector of size K)

e Unit exercise price of resources under option contract. (Vector of size K)

Without loss of generality, we assume that for each resource k , $p_k < q_k + e_k$ and $p_k > q_k$.

If $p_k \geq q_k + e_k$, the manufacturer will not use any fixed-price capacity from resource k . Similarly, if $p_k \leq q_k$, the manufacturer will not reserve any option capacity. We also assume that the demand vector is non-negative, i.e. $\mathbf{D} \geq 0$.

For decision variables, we denote:

- z_m Amount of product m that is produced and sold to meet demand. (Scalar)
- $\mathbf{z}$ Amount of products that are produced and sold to meet demand.
(Vector of size M)
- x_{jk} Amount of resource k provided under a fixed-price capacity contract that is
used to provide capacity to process j . (Scalar)
- $\mathbf{x}$ The vector of x_{jk} . (Vector of size JK)
- y_{jk} Amount of resource k provided under an option capacity contract that is
used to provide capacity to process j . (Scalar)
- $\mathbf{y}$ The vector of y_{jk} . (Vector of size JK)
- $\mathbf{c}$ The amount of fixed-price capacity that the firm has reserved. (Vector of size K)
- $\mathbf{g}$ The total amount of capacity, including fixed-price and option capacity,
that the firm has reserved. (Vector of size K)

We now formulate the second stage problem as a single-period production planning problem with the objective to maximize the profit of the firm. We are given the demand realization $\mathbf{d}$ as well as $\mathbf{c}$, the amount of each resource reserved with fixed-price contract, and $\mathbf{g}$, the total amount of each resource reserved. We note that $\mathbf{g} - \mathbf{c}$ is the amount of each resource reserved with an option contract. We have the following linear optimization problem:

$$\begin{aligned}
\pi(\mathbf{c}, \mathbf{g}, \mathbf{d}) = \max_{\mathbf{x}, \mathbf{y}, \mathbf{z}} \quad & \mathbf{r}'\mathbf{z} - \mathbf{e}'H\mathbf{y} \\
s.t. \quad & \mathbf{z} \leq \mathbf{d} \\
& A\mathbf{z} \leq B(\mathbf{x} + \mathbf{y}) \\
& H\mathbf{x} \leq \mathbf{c} \\
& H(\mathbf{x} + \mathbf{y}) \leq \mathbf{g} \\
& \mathbf{x}, \mathbf{y}, \mathbf{z} \geq \mathbf{0}
\end{aligned} \tag{1}$$

The objective function of Problem (1) is the net profit that the manufacturer will gain, given the capacity level $\mathbf{c}$ and $\mathbf{g}$ and demand $\mathbf{d}$. For the second stage problem this is the profit from selling $\mathbf{z}$, net of the additional cost from exercising the option contracts in the amount of $\mathbf{y}$. The first set of constraints restricts the amount of product sold to be less than the demand; we note that $\mathbf{d} - \mathbf{z}$ represents the amount of demand that is not met. The second set of constraints says that the amount of products produced can not exceed the total available capacity; the left hand side is the amount of process capacity required to produce $\mathbf{z}$ and the right hand side is the available process capacity given the allocation decisions $\mathbf{x}$ and $\mathbf{y}$. Finally, the third and fourth sets of constraints assure that the resource availability is not exceeded. The left hand side of the third set represents the resource usage under the fixed-price contract, while the left hand side of the fourth set is the total resource usage for the allocation decisions.




Figure 2: Numerical example: A manufacturer supply chain network containing two laptops, six processes, and six capacity providers.

By solving this optimization problem, we can find the profit maximizing production level for a given demand realization and the given capacity planning decisions. Let $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*)$ be an optimal solution of Problem (1); $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*)$ is a function of $\mathbf{d}$, $\mathbf{c}$, and $\mathbf{g}$. The firm ultimately wants to find the optimal capacity planning strategy under demand uncertainty by solving the following first-stage problem:

$$\begin{aligned}
 \max_{\mathbf{c}, \mathbf{g}} \quad & \Pi(\mathbf{c}, \mathbf{g}, \mathbf{D}) = E[\pi(\mathbf{c}, \mathbf{g}, \mathbf{D})] - \mathbf{p}'\mathbf{c} - \mathbf{q}'(\mathbf{g} - \mathbf{c}) \\
 \text{s.t.} \quad & \mathbf{c} \leq \mathbf{g} \\
 & \mathbf{c}, \mathbf{g} \geq \mathbf{0}
 \end{aligned} \tag{2}$$

The objective function of Problem (2) represents the expected total profit, which is equal to the expected net profit from the second stage, minus the first-stage reservation cost of the capacity. The first set of constraints ensures that the amount of fixed-price capacity reserved is no more than the amount of total capacity reserved.

Proposition 1 $\Pi(\mathbf{c}, \mathbf{g}, \mathbf{D})$ is concave in $(\mathbf{c}, \mathbf{g})$.

Proposition 1 guarantees that every local optimal solution for Problem (2) is a global optimal solution and that the algorithms given in Section 5 will converge. (The proofs of all propositions are in the Appendix 1.)

Table 1: Numerical example: table of product prices and demand information.

	Laptop A	Laptop B
Price (\$)	700	1000
Mean	2200	1000
STD	200	100

Table 2: Numerical example: table of capacity prices.

	Fixed Unit Price	Unit Reservation Price		Unit Exercise Price	
		Case 1	Case 2	Case 1	Case 2
Foundry 1	90	85	10	10	85
Foundry 2	100	80	30	30	80
Foundry 3	200	160	50	50	160
Foundry 4	98	78	28	28	78
CM 1	115	100	25	25	100
CM 2	110	90	30	30	90

3.2 An Example

To illustrate our model, we present a numerical example. A computer manufacturer produces two types of laptop, namely A and B. Laptop A requires three manufacturing processes or inputs: the manufacture or procurement of chipset A, the manufacture or procurement of display A, and Assembly & Testing (A&T). Similarly, each laptop B requires chipset B, display B, and Assembly & Testing.

Laptop A is an entry-level laptop selling at 700 dollars. Laptop B is a mid-range price laptop selling at 1000 dollars. The demand of both laptops follows a normal distribution with their means and standard deviations given in Table 1.

The manufacturer uses contract suppliers to perform the manufacturing processes. It currently has six suppliers from which to choose: Foundry 1, 2, 3, 4 and Contract Manufacturer (CM) 1, 2. The capability of each supplier is given in Figure 2. For instance, Contract Manufacturer 2 (CM 2) is qualified to do the assembly and test for Laptop B, whereas Contract Manufacturer 1 (CM 1) is qualified to do assembly and test for both laptops. Similarly, Foundry 2 is flexible and can produce both chipsets, whereas Foundry 1 (Foundry 4) can only supply Chipset A (Chipset B).

The manufacturer has two ways of contracting with each supplier. The price structure of each supplier for two different scenarios is given in Table 2. For Case 1, the unit reservation price is higher than the unit exercise price. The prices of the resources in Case 2 are the

same as Case 1 except that we swap the unit reservation price and the unit exercise price.

The manufacturer can reserve capacity from each supplier with a fixed-price capacity contract. For instance, Foundry 1 quotes a fixed unit price of \$90. Thus, if the manufacturer were to reserve 200 units of capacity, it would pay Foundry 1 \$1800; Foundry 1 will then commit to provide the manufacturer with upto 200 units of Chipset A over the demand period. To keep things simple, we assume the only cost is the upfront fixed cost of \$1800.

Alternatively the manufacturer can reserve capacity from a supplier with an option contract where there is a smaller upfront fixed cost and then a variable cost for the use of this capacity. For instance, in Case 1, the manufacturer might purchase an option contract with Foundry 3 for 300 units of capacity. The manufacturer would pay Foundry 3 an upfront cost of $300 \times \$160 = \$48,000$ to reserve this capacity. Later, when it needs to make the actual procurement decisions, the manufacturer can decide how much of the capacity to use (up to 300 units) and for what mix of products (i.e., display A or display B). The manufacturer pays an additional \$50 per unit for each unit of capacity that it actually uses. We note that the fixed-price contract is effectively an option contract with a zero exercise price; we don't require the manufacturer to use all of the fixed-capacity, and there is no additional cost for using this capacity.

Given the demand distributions (Table 1), network structure (Figure 2), and cost structures of the suppliers (Table 2), the manufacturer wants to answer the following questions. First, which suppliers should it use? Second, what types of contract should it use for each supplier? Only fixed-price contract? Only option contract? Or Both. Third, how much capacity should it buy?

The firm needs to consider the trade-offs between different factors. First, demand is uncertain and the manufacturer will want to have enough process capacity to meet any demand outcome, up to some level. Second, to deliver a product the manufacturer must have sufficient capacity for all of its processes; having enough chipsets is not very useful if one is short of displays. Third, the resource options vary in terms of cost and flexibility. For instance, the capacity from Foundry 2 is more expensive relative to that from either Foundry 1 or 4; but the capacity at Foundry 2 is flexible as it can produce either display.

Our intent is to develop a model and algorithms to help the manufacturer to answer these questions and understand the trade-offs.

For this example, we report the optimal solution to Problem (2) in Table 3. For Case 1, the manufacturer should use all six suppliers, use only a fixed-price contract from Foundry 1,

Table 3: Numerical example: results.

	Fixed-Price Capacity		Option Capacity		Total Capacity	
	Case 1	Case 2	Case 1	Case 2	Case 1	Case 2
Foundry 1	1977	1667	0	491	1977	2158
Foundry 2	364	378	0	0	364	378
Foundry 4	757	823	79	0	836	823
Foundry 3	3023	2871	154	488	3177	3359
CM 1	2341	2115	0	370	2341	2485
CM 2	774	805	62	69	836	874

Foundry 2, and CM 1, and use both types of contract from the other suppliers. This example also shows that the sums of total capacity for Foundry 1, 2, and 4, the total capacity for Foundry 3, and the sum of total capacity for CM 1 and CM 2 are equal to each other. Foundry 1, 2, and 4 provide capacity for the chipsets; Foundry 3 provides capacity for the display; CM 1 and 2 provide capacity for the A&T. Since to produce a product requires all three processes, the total capacity reserved for these processes is the same.

We can draw similar conclusion for Case 2; as the option contracts are more favorable for case 2, we see that we use both more option capacity and more total capacity relative to Case 1.

This example also illustrates the complexity of the optimal strategy. We expect that the manufacturer will reserve more capacity in Case 2, since the unit reservation prices of all resources are lower than Case 1. However, from the optimal solutions we can see that the manufacturer should not reserve any option capacity for Foundry 4 in Case 2, while in Case 1 it should reserve 79 units of option capacity. This is due to the interdependency between Foundry 1, 2, and 4. In Case 2, since the unit reservation price for Foundry 1 is much lower than the reservation price for Foundry 4, Foundry 1 has a more attractive option contract. Therefore, the manufacturer should buy more option contract from Foundry 1 and less option contract from Foundry 4.

4. Two Special Cases

We first study two special cases of the problem: a single product with dedicated resources and a single process with dedicated resource.




Figure 3: A supply chain network with single product and dedicated resources.

4.1 Special Case I: Single Product with Dedicated Resources

Let us consider the first special case where the network contains one product and dedicated resources. Figure 3 shows such a supply chain with a single product that requires three processes and each process has a dedicated resource that provides capacity for it. As there is exactly one resource for each process, we will view these as synonymous and will use the terms interchangeably. For this special case, we can derive a closed-form representation for the optimal capacity planning strategy. Without loss of generality, we assume that we number the resources such that

$$\frac{p_i - q_i}{e_i} \leq \frac{p_j - q_j}{e_j}, \text{ if } i < j \quad (3)$$

The ratio $\frac{p_i - q_i}{e_i}$ is non-negative and less than 1. The bigger the ratio is, the more attractive the option contract.

Proposition 2 *Assume D is a continuous random variable and $r > \sum_j e_j$. For a supply chain with single product and dedicated resources, a capacity planning strategy $(\mathbf{c}, \mathbf{g})$ is optimal iff there exists an integer $1 \leq \psi \leq J + 1$ such that all of the following conditions are satisfied:*

$$g_j = g, \forall j \quad (4)$$

$$Pr(D > c_j) = \frac{p_j - q_j}{e_j}, \forall j \geq \psi \quad (5)$$

$$c_j = g, \forall j < \psi \quad (6)$$

$$Pr(D > g) = \frac{\sum_{j=1}^{\psi-1} p_j + \sum_{j=\psi}^J q_j}{r - \sum_{j=\psi}^J e_j} \quad (7)$$

$$\frac{p_\psi - q_\psi}{e_\psi} > \frac{\sum_{j=1}^{\psi-1} p_j + \sum_{j=\psi}^J q_j}{r - \sum_{j=\psi}^J e_j} \text{ if } \psi \leq J \quad (8)$$

From Proposition 2, we can make a number of observations. First we see that all processes reserve the same amount of total capacity, given by g . Since producing the product requires all processes, reserving more capacity for some but not all processes is a waste as the excess capacity can never be used.

Second, we can interpret the optimal planning strategy in terms of the newsboy problem. To determine g , suppose we know how to partition the resources based on whether or not they will buy an option contract. Namely, we assume for resources $1, \dots, \psi - 1$, we only invest in a fixed-price contract, while for resources $\psi, \dots, J$, we invest in both a fixed-price contract and an option contract. Then in a newsboy context, the cost of overage is given by $C_o = \sum_{j=1}^{\psi-1} p_j + \sum_{j=\psi}^J q_j$, which equals the upfront investment to reserve the last unit of capacity. This is the incremental cost when demand falls below g . The underage cost is $C_u = r - \sum_{j=1}^{\psi-1} p_j - \sum_{j=\psi}^J (q_j + e_j)$, which equals the incremental profit net of the costs for all of the resources. This is the lost profit when demand exceeds g . Thus the critical ratio for determining g is given by (7), namely the traditional critical ratio for the newsboy:

$$Pr(D > g) = \frac{C_o}{C_o + C_u} = \frac{\sum_{j=1}^{\psi-1} p_j + \sum_{j=\psi}^J q_j}{r - \sum_{j=\psi}^J e_j},$$

where we assume that we are given the partition of resources.

Now to get insight into how to construct the partition, we consider each resource independently. Suppose we were to buy both a fixed-price contract and an option contract for resource j , subject to the fact that the total capacity is fixed at g . We wish to determine how much to buy of the fixed-price contract. For resource j , the overage cost is $C_{o,j} = p_j - q_j$ as this represents the upfront premium that is paid for the fixed-price contract relative to the option contract, and equals the amount that would be lost if this capacity is not needed. The underage cost is $C_{u,j} = q_j + e_j - p_j$ which is equal to the cost premium to serve demand from the option contract relative to the fixed-price contract. Thus, the critical ratio for determining the size of the fixed-price contract for resource j is given by:

$$Pr(D > c_j) = \frac{C_{o,j}}{C_{o,j} + C_{u,j}} = \frac{p_j - q_j}{e_j}$$

which corresponds to (5). If this equation suggests buying more than g units of capacity, then we should not buy an option contract for resource j and we should reduce its fixed-price contract to g . In effect, this is what is enforced by Equation (6) and (8). Finally sorting the resources, as prescribed by (3), provides a simple way to find the partition.

Also, we note from condition (5) that for those processes that do buy an option contract, the optimal fixed-price capacity is independent of r , the profit of the product. We also observe that from condition (7) that the optimal total capacity is independent of the prices of the fixed-price contract for the resources for which we buy option contracts.

For each process, the optimal strategy has a similar structure to that given by [Martinez-de-Albéniz and Simchi-Levi 05]. They study the replenishment policy and portfolio selection strategy for a single product that has a single process in the presence of a spot market. In their model, there are multiple option contracts available for the single process. For a single period model, they give a closed-form solution to the portfolio selection (capacity investment) problem. Our result for the optimal level of fixed-price capacity is similar to the result that they have for their single period problem. However, in our model, the manufacturer needs to acquire the capacity for multiple processes at the same time. Therefore, our results for the optimal level of total capacity differ from those in Martinez-de-Albeniz and Simchi-Levi, especially with regard to the partition property for separating the processes between those that use an option contract and those that do not.

A supply chain with a single product and dedicated resources is a practical case. Proposition 2 provides a closed-form solution for the optimal capacity planning strategy for this class of supply chain. This proposition not only reveals some interesting insights of the optimal strategy but also provides an effective way to find the optimal strategy.

We can also extend the special case to a more general setting, where each process can have multiple dedicated resources. We can obtain a closed-form representation of the optimal capacity for each resource, given the total capacity g . We can then apply a convex search algorithm to find the optimal total capacity g^* . We refer the reader to [Huang 08] for the details.

4.2 Special Case II: Single Process with Dedicated Resource

We will now consider the second special case: a network with a single process and a single dedicated resource. An example of such a network is given in Figure 4. Without loss of generality, we assume $r_i \geq r_j$ if $i \leq j$. For this class of supply chains, we can also obtain a




Figure 4: A supply chain network with single process and dedicated resource.

closed-form solution of the optimal capacity planning strategy, which is given in the following proposition:

Proposition 3 *Assume $\mathbf{D}$ is a continuous random vector and $r_i > e$ for all i . For a supply chain network with a single process and dedicated resource, an optimal capacity planning strategy (c, g) satisfies one of the following two sets of conditions:*

$$\text{Set 1 :} \quad Pr \left(\sum_{i=1}^M D_i > c \right) = \frac{p-q}{e} \quad (9)$$

$$q + ePr \left(\sum_{i=1}^M D_i > g \right) - \sum_{i=1}^M \left[r_i Pr \left(\sum_{j=1}^i D_j > g > \sum_{j=1}^{i-1} D_j \right) \right] = 0 \quad (10)$$

$$c < g \quad (11)$$

$$\text{Set 2 :} \quad p - \sum_{i=1}^M \left[r_i Pr \left(\sum_{j=1}^i D_j > g > \sum_{j=1}^{i-1} D_j \right) \right] = 0 \quad (12)$$

$$c = g \quad (13)$$

The first (second) set of conditions is the necessary and sufficient conditions for an optimal planning strategy with (without) the purchase of an option contract. If it is better for the manufacturer to reserve a positive amount of option capacity, the optimal policy will have similar structure to that found in Proposition 2 for the process using an option contract. Equation (9) is the same as for Proposition 2. For Equation (10), the first two terms are the incremental cost for increasing the size of the option contract; the third term is the incremental profit from increasing the size of the option contract. The first order condition just equates the incremental cost with the incremental benefit, under the assumption that

we use the option contract. Similarly, if it is better not to use an option contract, the incremental cost (profit) for increasing the size of the fixed-price contract is given in the first (second) term of Equation (12). Equating them gives us the first order condition. Therefore, to find the optimal capacity plan, we first solve (9) to find $\mathbf{c}$ and (10) to find $\mathbf{g}$. If $\mathbf{c} \leq \mathbf{g}$, then this is the optimal capacity plan. Otherwise, we solve (12) to find the optimal capacity plan with $\mathbf{c} = \mathbf{g}$.

By examining the optimal strategy given in Proposition 3, we see that if it is optimal for the firm to reserve option capacity, the optimal fixed-price capacity is independent of the profits of the products and the optimal total capacity is independent of the price of the fixed-price capacity. We have observed a similar property for the class of a supply chain with single product and dedicated resources.

5. Solving the Single Period Capacity Planning Problem

Unlike the special cases we have studied in the previous sections, it is very difficult to derive a closed-form solution for the optimal capacity planning strategy in a general supply chain setting. Therefore, in this section, we will study different algorithms for solving the general single period capacity planning problem (2) and compare their performances.

5.1 Sampling

Through the rest of this paper, we will use sampling to model demand uncertainty. Given any probability or empirical distribution of the demand, we randomly draw a set of demand realizations and denote this set by $\mathcal{S}$. In effect, we will model the given demand distribution by the sample; that is, we assume demand comes from a discrete distribution defined on the sample space, where each sample point is equally likely to occur. Let L be the size of the sample set. In this section, we will give some guidelines for picking the number L .

Let us assume that we have selected a set of demand samples $\mathcal{S}$ with size L . Denote $\Pi_L(\mathbf{c}_L, \mathbf{g}_L, \mathbf{D}_L)$ to be the objective function value of Problem (2) by replacing the expectation over the original demand distribution with the average over the L sample points and $\Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)$ be the corresponding maximum objective function value of Problem (2). We would like to find a bound on the probability that $|\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)|$ is smaller than

a positive scalar ϵ . We will give two bounds based on two different inequalities: Hoeffding inequality [Hoeffding 63] and Chernoff inequality [Wikipedia 08].

We assume that for any given $\mathbf{c}$, $\mathbf{g}$, and $\mathbf{d}$, we can identify a lower and upper bound on the expected net profit, $\pi_{min} \leq \pi(\mathbf{c}, \mathbf{g}, \mathbf{d}) \leq \pi_{max}$. In practice, π_{min} and π_{max} can be the minimum and maximum profit that the manufacturer can gain. Then, we can derive the following bound using Hoeffding Inequality.

$$Pr(|\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)| \leq \epsilon) \geq 1 - 4 \exp\left(\frac{-L\epsilon^2}{2(\pi_{max} - \pi_{min})^2}\right). \quad (14)$$

The derivation of Equation (14) is given in Appendix 2. Equation (14) suggests a guideline to pick a suitable sample size given the knowledge of the bound on the expected profit.

Similarly, if the manufacturer has an estimate of the maximum standard deviation, σ_π , of the expected profit, $\Pi(\mathbf{c}, \mathbf{g}, \mathbf{D})$, and knows that the expected profit is bounded, it can bound the quantity $|\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)|$ using the Chernoff inequality.

$$Pr(|\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)| \leq \epsilon) \geq 1 - 4 \exp\left(\frac{-L\epsilon^2}{4\sigma^2}\right). \quad (15)$$

Finally, because we use sampling to model the uncertainty, we do not make any assumption on the distribution of demand. In practice, the manufacturer can generate the demand samples from some probability distributions or from the demand history.

5.2 Sub-gradient Method

Van Mieghem and Rudi [Van Mieghem and Rudi 02] suggest a sub-gradient algorithm to solve a different but similar single period capacity planning problem. In their model, the firm can only reserve fixed-price capacity but not option capacity. The main purpose of their paper is to study the properties of the optimal planning strategies. They prove the necessary and sufficient conditions of the optimal solution and briefly mention that the problem can be solved using a sub-gradient algorithm. Since our single period problem is similar, we can develop a similar sub-gradient algorithm to our model.

We first consider the sub-gradients of Problem (2). For each demand realization $\mathbf{d}$, let $\lambda(\mathbf{c}, \mathbf{g}, \mathbf{d})$ be the associated dual variables of constraints $H\mathbf{x} \leq \mathbf{c}$ and $\gamma(\mathbf{c}, \mathbf{g}, \mathbf{d})$ be the associated dual variables of constraints $H(\mathbf{x} + \mathbf{y}) \leq \mathbf{g}$ in Problem (1). Then the sub-gradients of the objective function of Problem (2) are $\nabla_{\mathbf{c}}\Pi = E[\lambda(\mathbf{c}, \mathbf{g}, \mathbf{D})] - \mathbf{p} + \mathbf{q}$ and $\nabla_{\mathbf{g}}\Pi = E[\gamma(\mathbf{c}, \mathbf{g}, \mathbf{D})] - \mathbf{q}$. We omit the proof since it is very similar to the one given in

[Van Mieghem and Rudi 02]. By Proposition 1, $\Pi(\mathbf{c}, \mathbf{g}, \mathbf{D})$ is concave in $(\mathbf{c}, \mathbf{g})$. Thus, the first order conditions are the necessary conditions for optimality. Therefore, we can use a sub-gradient method to find the optimal solution [Van Mieghem and Rudi 02].

Sub-gradient Algorithm:

Step 0: Set $s = 0$. We start with a given initial feasible solution $(\mathbf{c}^0, \mathbf{g}^0)$.

Step 1: For capacity strategy $\mathbf{c}^s$ and $\mathbf{g}^s$, solve the linear program (1) and find the associated dual variables $\lambda(\mathbf{c}^s, \mathbf{g}^s, \mathbf{d})$ and $\gamma(\mathbf{c}^s, \mathbf{g}^s, \mathbf{d})$ numerically for each sample demand vector $\mathbf{d}$. Take

$$\sum_{\mathbf{d} \in \mathcal{S}} \frac{\lambda(\mathbf{c}^s, \mathbf{g}^s, \mathbf{d})}{L} \text{ and } \sum_{\mathbf{d} \in \mathcal{S}} \frac{\gamma(\mathbf{c}^s, \mathbf{g}^s, \mathbf{d})}{L}$$

as unbiased estimates of $E[\lambda(\mathbf{c}^s, \mathbf{g}^s, \mathbf{D})]$ and $E[\gamma(\mathbf{c}^s, \mathbf{g}^s, \mathbf{D})]$, and use them to compute estimates of the sub-gradient $\nabla_{\mathbf{c}^s} \Pi$ and $\nabla_{\mathbf{g}^s} \Pi$.

Step 2: If $|\nabla_{\mathbf{c}^s} \Pi|$ and $|\nabla_{\mathbf{g}^s} \Pi|$ are smaller than some tolerance level, then stop. Otherwise, adjust capacity in the direction of the sub-gradients, $\mathbf{g}^{s+1} = \mathbf{g}^s + \xi \nabla_{\mathbf{g}^s} \Pi$ and $\mathbf{c}^{s+1} = \min \{\mathbf{c}^s + \xi \nabla_{\mathbf{c}^s} \Pi, \mathbf{g}^{s+1}\}$ where ξ is some step-size (or perform a line-search). Set $s = s + 1$ and return to step 1.

At each iteration, step 1 of the sub-gradient algorithm will solve L linear programs where L is the number of sample demand points that are used to estimate the sub-gradients. The computational requirements at each step can be very large depending upon the number of sample points. Therefore, if the sub-gradient method requires many iterations to converge, the algorithm will take a long time to run. We refer the reader to [Huang 08] for a detailed discussion of the sub-gradient algorithm. Hagle and Sen [Hagle and Sen 96] show that the supporting hyperplane algorithm outperforms the sub-gradient algorithm for solving large scale stochastic linear problems. We discuss this algorithm next.

5.3 Regular Supporting Hyperplane Algorithm

Another algorithm that uses the sub-gradient is the Supporting Hyperplane Algorithm suggested by [Veinott 67]. Let us consider a new problem:

$$\begin{aligned} \min \quad & f \\ \text{s.t.} \quad & f + E[\pi(\mathbf{c}, \mathbf{g}, \mathbf{D})] - \mathbf{p}'\mathbf{c} - \mathbf{q}'(\mathbf{g} - \mathbf{c}) \geq 0 \\ & \mathbf{c} \leq \mathbf{g} \end{aligned} \tag{16}$$

We can show that $(\mathbf{c}^*, \mathbf{g}^*)$ solves Problem (2) iff $(\mathbf{c}^*, \mathbf{g}^*, f^*)$ solves Problem (16) with

$$f^* + E[\pi(\mathbf{c}^*, \mathbf{g}^*, \mathbf{D})] - \mathbf{p}'\mathbf{c}^* - \mathbf{q}'(\mathbf{g}^* - \mathbf{c}^*) = \mathbf{0}.$$

We use the supporting hyperplane algorithm to solve Problem (16).

We assume we can identify upper and lower bounds on f , $\mathbf{c}$, and $\mathbf{g}$. Let $\mathbf{c}_{upper}(\mathbf{c}_{lower})$ and $\mathbf{g}_{upper}(\mathbf{g}_{lower})$ be the upper (lower) bounds on the fixed-price and total capacities. Let $f_{upper}(f_{lower})$ be the upper (lower) bound of f . Let

$$V^0 = \{(\mathbf{c}, \mathbf{g}, f) : \mathbf{c} \in [\mathbf{c}_{lower}, \mathbf{c}_{upper}], \mathbf{g} \in [\mathbf{g}_{lower}, \mathbf{g}_{upper}], f \in [f_{lower}, f_{upper}]\}$$

and set $s = 0$. The algorithm consists of the following steps:

Regular Supporting Hyperplane Algorithm:

Step 1: Solve the linear program of minimizing f , subject to $(\mathbf{c}, \mathbf{g}, f) \in V^s$, and let $(\mathbf{c}^s, \mathbf{g}^s, f^s)$ be the optimal solution. If

$$f^s + E[\pi(\mathbf{c}^s, \mathbf{g}^s, \mathbf{D})] - \mathbf{p}'\mathbf{c}^s - \mathbf{q}'(\mathbf{g}^s - \mathbf{c}^s) \geq -\varepsilon.$$

where ε is a small positive number chosen by the user, stop. Otherwise, go to step 2.

Step 2: Use the simulation method given in the sub-gradient algorithm to calculate the sub-gradient $\nabla_{\mathbf{c}^s}\Pi$ and $\nabla_{\mathbf{g}^s}\Pi$. Add a linear constraint to the set V^s :

$$f + \Pi(\mathbf{c}^s, \mathbf{g}^s, \mathbf{D}) + [(\mathbf{c}, \mathbf{g}) - (\mathbf{c}^s, \mathbf{g}^s)]'(\nabla_{\mathbf{c}^s}\Pi, \nabla_{\mathbf{g}^s}\Pi) \geq 0 \quad (17)$$

where $\Pi(\mathbf{c}^s, \mathbf{g}^s, \mathbf{D})$ is a constant, which equals $E[\pi(\mathbf{c}^s, \mathbf{g}^s, \mathbf{D})] - \mathbf{p}'\mathbf{c}^s - \mathbf{q}'(\mathbf{g}^s - \mathbf{c}^s)$. Let the new set be V^{s+1} . Set $s = s + 1$ and go to step 1.

Geometrically, the supporting hyperplane method approximates the function $\Pi(\mathbf{c}, \mathbf{g}, \mathbf{D})$ with hyperplanes. To construct the initial constraint set V^0 , one can set $\mathbf{c}_{lower}$ and $\mathbf{g}_{lower}$ to be $\mathbf{0}$, $\mathbf{c}_{upper}$ and $\mathbf{g}_{upper}$ to be maximal capacity requirement to fill all demand, f_{lower} to be the objective value of any feasible strategy, and f_{upper} to be the maximal profit that the firm can achieve. At each step, the algorithm adds a new supporting hyperplane to the constraint set, based on the sub-gradient from the last solution (supporting point). It then uses all the sub-gradients that it has calculated so far to find the next supporting point. By the nature of the supporting hyperplane algorithm, it does not require a starting point or a

step size. Finally, at each step $-f^s$ is an upper bound for $\Pi(\mathbf{c}^*, \mathbf{g}^*, \mathbf{D})$. Therefore, the ε in the stopping criterion in step 1 is an upper bound for $|\Pi(\mathbf{c}^s, \mathbf{g}^s, \mathbf{D}) - \Pi(\mathbf{c}^*, \mathbf{g}^*, \mathbf{D})|$.

Even though we expect the supporting hyperplane algorithm in general to have a better convergence rate compared to sub-gradient method, it suffers from the problem of slow start. Note that at each iteration, the algorithm needs to solve L linear programs to find the supporting hyperplane where L is the number of samples. At the beginning, the supporting point is likely to be far away from the optimal solution. It might not be necessary to construct an accurate supporting hyperplane at points that are far away from the optimum using all samples, since these hyperplanes are only used to find an approximate location of the next supporting point. As the algorithm proceeds, the supporting points get closer and closer to optimum, and we need more accurate supporting hyperplanes. Since the regular supporting hyperplane algorithm uses all samples at each iteration, it wastes computational effort at the beginning and therefore has a slow start problem. To overcome this problem, we can adapt a variation of the regular supporting hyperplane algorithm from large-scale stochastic linear programming to solve Problem (2) ([Higle and Sen 91] and [Higle and Sen 96]). We describe this algorithm in the next section.

5.4 Stochastic Supporting Hyperplane Algorithm

To address the slow start problem of the regular supporting hyperplane algorithm, we adapt the technique suggested by [Higle and Sen 91] and [Higle and Sen 96]. In their algorithms, instead of using all samples at each step, they incrementally increase the number of sample points at each iteration.

Stochastic Supporting Hyperplane Algorithm:

Step 0: Set up the initial V^0 as for the regular supporting hyperplane algorithm. Set $s = 0$ and the initial demand sample set $\mathcal{S}^0 = \emptyset$.

Step 1: Set $s = s + 1$. Randomly generate a demand observation ω^s independent of any previously generated observations. Let $\mathcal{S}^s = \mathcal{S}^{s-1} \cup \omega^s$. Construct the s^{th} supporting hyperplane using the same method given in step 2 of the regular supporting hyperplane algorithm. Define the s^{th} supporting hyperplane at s^{th} iteration to be:

$$f + \alpha_s^s + (\beta_s^s)' \mathbf{c} + (\zeta_s^s)' \mathbf{g} \geq 0$$

where $\alpha_s^s = \Pi(\mathbf{c}^s, \mathbf{g}^s, \mathbf{D}) - (\mathbf{c}^s, \mathbf{g}^s)'(\nabla \mathbf{c}^s \Pi, \nabla \mathbf{g}^s \Pi)$, $\beta_s^s = \frac{1}{s} \sum_{\mathbf{d} \in \mathcal{S}^s} \lambda(\mathbf{c}^s, \mathbf{g}^s, \mathbf{d}) - \mathbf{p} + \mathbf{g}$, and $\zeta_s^s = \frac{1}{s} \sum_{\mathbf{d} \in \mathcal{S}^s} \gamma(\mathbf{c}^s, \mathbf{g}^s, \mathbf{d}) - \mathbf{q}$.

Step 2: Update the coefficients of all previously generated supporting hyperplane:

$$\alpha_t^s = \frac{s-1}{s}\alpha_t^{s-1} + \frac{1}{s}U, \quad \beta_t^s = \frac{s-1}{s}\beta_t^{s-1} - \frac{\mathbf{p}}{s} + \frac{\mathbf{q}}{s}, \quad \zeta_t^s = \frac{s-1}{s}\zeta_t^{s-1} - \frac{\mathbf{q}}{s},$$

where U is an upper bound on $\pi(\mathbf{c}, \mathbf{g}, \mathbf{D})$, for $t = 1, \dots, s-1$.

Step 3: Find the next supporting point using the same method given in step 1 of regular supporting hyperplane algorithm. If the algorithm does not terminate, go to step 1.

We derive the update rules in step 2 in Appendix 3. Note that the supporting hyperplane constructed at iteration s uses s samples. Therefore, the supporting hyperplanes from different iterations use different numbers of sample points. The updating rules in step 2 modify the previously generated supporting hyperplanes to incorporate this difference. For details of stochastic supporting hyperplane algorithm, such as its convergence property, please refer to [Higle and Sen 96]. The stochastic supporting hyperplane algorithm addresses the problem of slow start by incrementally increasing the size of the sample set by one at each iteration. Even though in general the algorithm will take more iterations to converge, the average computational effort required in each iteration is less than the regular supporting hyperplane algorithm. As a result, the performance of the algorithm increases significantly as we will see in Section 5.6. However, adding one demand sample at each step means that the algorithm needs to solve one more linear program for all future iterations. For our problem, the computational requirement increases very quickly as the number of iterations increases. Therefore, we have developed another algorithm based on the stochastic supporting hyperplane algorithm to solve Problem (2).

5.5 Stochastic Supporting Hyperplane Algorithm with Pre-solve Routine

The new algorithm contains two stages. We first choose a small subset of the sample set and use the regular supporting hyperplane method to construct an initial polyhedra V^0 . We then use the stochastic hyperplane supporting algorithm to find the optimal solution. We now outline this algorithm:

Stochastic Supporting Hyperplane with Pre-solve Routine:

Stage I: Pick a subset $\bar{\mathcal{S}} \subset \mathcal{S}$, solve the problem with the regular supporting hyperplane algorithm described in Section 5.3. Let $\bar{V}$ be the final polyhedra of the master LP.

Stage II: Set $V^0 = \bar{V}$ and use the stochastic supporting hyperplane algorithm described in Section 5.4 to solve the capacity planning problem.

In stage I, the algorithm takes advantage of the fast convergence rate of the regular supporting hyperplane algorithm but with a reduced computational requirement at each iteration by using a small sample size. We expect the solution from stage I to be close to the optimum. The algorithm then uses the stochastic supporting hyperplane algorithm to refine the solution. Since the second stage problem starts with a good starting point and initial constraint set, we expect that the stochastic supporting hyperplane algorithm should converge faster compared to starting from scratch.

5.6 Algorithm Run Time Comparisons

After examining different algorithms for solving the single period capacity planning problem, we now discuss their run time performances. We use a free linear program solver, GNU Linear Programming Kit (GLPK 4.11), for all of the test cases. This solver is slower than a commercial LP solver such as CPLEX. However, the computational tests presented here show the relative performance comparison of the algorithms. The test machine is an IBM x40 laptop with a 1.29 GHz Intel Pentium M CPU and 760 MB of memory. All the tests were written in the C++ programming language and performed in a Windows XP environment.

We consider a supply chain with 15 products, 30 processes, and 30 resources. We generate random test cases according to the following rules. The demand of each product follows a normal distribution with mean uniformly distributed between 100 and 120 and standard deviation 10. The profit of each product is uniformly distributed between 150 and 170. The price of fixed-price capacity, p_k , is uniformly distributed between 9 and 12. The cost of option capacity, q_k , is uniformly distributed between 1 and p_k . The exercise cost of option capacity is set to $p_k \times 1.1 - q_k$. A link joins a product and process with probability 0.2 (e.g. $Pr(A(j, m) = 1) = 0.2$) and a link joins a process and a resource with probability 0.2 (e.g. $Pr(B(j, (j, k)) = 1) = 0.2$). We check whether the supply chain generated is connected or not. If not, we repeat the generation process until we have a connected supply chain. In each case, we set the sample size to be 500, and we use the same 500 sample demands for all algorithms.

We first convert one of the test cases (2) into a deterministic linear program and solve it with the LP solver. The deterministic linear program has 182,560 variables and 197,530

Table 4: Run time comparison statistics.

	Regular	Stochastic	Pre-solve	Pre-solve/Stoc.	Pre-solve/Reg.
Average	406.7	143.8	52.9	43.35%	13.23%
STD	138.5	93.2	20.1	16.14%	2.53%
Min	174.0	43.0	29.0	15.47%	8.20%
Max	794.0	513.0	130.0	83.72%	17.96%

constraints. It takes the algorithm 10 hours 36 minutes 12 seconds to find an optimal solution. This run time is significantly slower than the other algorithms, as we will see. Higle and Sen [Higle and Sen 96] show that the supporting hyperplane algorithm outperforms the sub-gradient algorithm for solving large scale stochastic linear problems. Therefore, we will focus on comparing the performances of the three types of supporting hyperplane algorithms presented above.

For the three supporting hyperplane algorithms, we set the terminating error percentage to be less than 1%. We set the sample size to be 100 for the stage I of the stochastic supporting hyperplane algorithm with pre-solve routine. For 40 randomly generated test cases, we report the detailed results in Appendix 4 and the summary statistics of the runtime comparisons in Table 4.

The average runtime of the algorithm using pre-solve for these test cases is 13.23% of the average runtime of the regular supporting hyperplane algorithm and 43.35% of the average runtime of the stochastic supporting hyperplane algorithm. For the maximum improvement, the runtime of the algorithm with pre-solve routine is 8.20% of the runtime of the regular algorithm and 15.47% of the runtime of the stochastic algorithm. We also note that the average runtime of the algorithm with pre-solve routine is 43 seconds. The runtime will be much faster if we use a commercial LP solver. Therefore, the algorithm is a practical solution approach for these problems.

6. Common Process and Option Capacity

In our model, there are three types of flexibility that the manufacturer can use to cope with the demand uncertainty: common processes, flexible resources, and option contracts. In this section, we will discuss the effects of using common processes and option contracts through a series of examples. Finally, we will draw some managerial insights into how to use these flexibilities.




Figure 5: a) A supply chain with two products and four processes. b) Replacing the process 2a and 2b in (a) with a common process with the same price.

Table 5: The benefits of using common processes and option contracts.

Profit	Margin	S1	S2	S3	S4	S2 vs. S1	S3 vs. S1	S4 vs. S1
66	1.54%	3,849	3,849	3,849	3,996	0.00%	0.00%	3.81%
67	3.08%	4,584	4,584	4,585	4,804	0.00%	0.04%	4.81%
68	4.62%	5,329	5,329	5,355	5,629	0.00%	0.49%	5.63%
69	6.15%	6,085	6,085	6,158	6,472	0.00%	1.20%	6.36%
70	7.69%	6,850	6,850	6,979	7,326	0.00%	1.87%	6.94%
80	23.08%	14,902	15,039	15,858	16,319	0.92%	6.41%	9.50%
90	38.46%	23,418	23,842	25,222	25,704	1.81%	7.70%	9.76%
100	53.85%	32,268	32,900	34,799	35,287	1.96%	7.84%	9.36%
110	69.23%	41,286	42,131	44,471	44,967	2.05%	7.71%	8.92%
150	130.77%	78,399	79,902	83,580	84,141	1.92%	6.61%	7.32%

Consider a supply chain given in Figure 5(a) that contains two products and four processes. Each process has a dedicated resource (not shown in figure) and we will view them as synonymous and use the terms interchangeably.

Both products have the same unit profit. The unit prices of the fixed-price capacity are $[p_1, p_{2a}, p_{2b}, p_3] = [10, 50, 50, 10]$. The demand of each product follows normal distribution with mean and standard deviation: $E[D_1] = 502$; $\sigma(D_1) = 99$; $E[D_2] = 496$; $\sigma(D_2) = 99$. To study the effects of common processes and option contracts, we will consider the following four scenarios:

1. The optimal capacity strategy for the supply chain given in Figure 5(a).
2. The same problem in scenario 1 except that we replace processes 2a and 2b with a common process with the same price. The supply chain after the replacement is given in Figure 5(b).
3. The same problem in scenario 1 except that we add an option contract to process 2a




Figure 6: Comparing the benefits of using common process and option capacity.

and 2b. The option contract has a unit reservation price 5 and unit exercise price 50.

4. We combine scenario 2 and 3.

We will compare the change to the maximum expected profit in the four scenarios as we increase the unit profits for both products from 66 to 150. The results are given in Table 5. For each choice of product profit, we define the profit margin to be $(\text{profit} - 60)/60$, where 60 corresponds to the product cost using fixed-price contracts. We use the maximum expected profit in scenario 1 as the reference point. We then quantify the benefit gained in the other scenarios as the percentage increase in profit compared to the reference. We have plotted the benefits versus profit margin in Figure 6. From this example, we have the following observations:

1. *The benefits of using common process and option contract are small when the profit margin is low. The benefits increase and then decrease as the profit margin increases.*

The benefit of common process comes from risk pooling. In this example, when the profit margin is very low, the loss from excess dedicated capacity is larger than the gain from the risk pooling on the common process. Therefore, the benefit of using the common process is small.

The benefit of using the option contract is to reduce the overage cost (from the excess capacity) when the demand is low. When the profit margin is very low, the reduced overage cost is still too large compared to the underage cost (from the unfilled demands). Therefore, the option contract is also not very effective.

When the profit margin is high, the underage cost is more significant compared to the overage cost. The manufacturer is willing to bear the cost of excess capacity to avoid not filling the demand. Therefore, any savings from reducing the common process capacity or reducing the cost of excess capacity become less significant when the profit margin is high.

2. *Using an option contract with small reservation price is more effective than using common process.*

In Figure 6, the red line representing the benefit of using option contract is always above the blue line, which represents the benefit of using common process. We assume that the option contract has a price that is 10% of the fixed-price capacity. The effective price of the option capacity, which is the sum of reservation price and exercise price, is 10% more than the price of the fixed-price capacity. In this case, the option capacity can help the manufacturer to increase the expected profit by as much as 7.84% in the case where the profit margin is 53.83%, which is much larger than the 1.96% profit increase realized by using common processes. This example shows that an option contract with a small upfront price can be very effective. As the reservation price increases, the benefit of using an option contract decreases. For example, when the profit margin is 83.83%, if the reservation price is 5, the total net profit that the manufacturer can gain is 34,799. We then hold $q + e = 55$ and increase q from 5 to 25, the profit drops to 32,915. The increase in the expected profit compared to scenario 1 is 2.00%, which is slightly better than the 1.96% increase achieved from using a common process. If we further increase the reservation price to 40, the profit drops to 32,287. The percentage increase in profit compared to scenario 1 is only 0.06%.

Table 6: The effects of common processes and option capacity: the optimal solutions of different strategies.

	c_1	c_{2a}	c_{2b}	c_3	g_1	g_{2a}	g_{2b}	g_3
Combined	437	760	-	426	437	863	-	426
Only option	406	378	367	397	406	406	397	397
Only common	384	763	-	378	384	763	-	378

3. *When the profit margin is low, the strategy of using common process and an option contract with small reservation price together can gain extra benefits compared to using these two strategies separately.*

Implementing both strategies, using the common process and the option contract, is better than just deploying one of them. Furthermore, there is a synergy at low profit margins in that the manufacturer gains an extra benefit by combining these two strategies. In Figure 6, the green line is the sum of the benefit of using just the common process plus the benefit of just using option contract. When the profit margin is less than 40%, the benefit of combining two strategies, represented by the purple line, is bigger than the green line. The gap between these two lines is the extra benefit from the synergy of these two strategies. Moreover, this extra benefit can be significant. For example, when the profit margin is 4.62%, using the common process and the option contract can increase the profit by 0% and 0.49% respectively. However, the combined strategy achieves a 5.63% profit increase. To understand the reasons behind this phenomenon, we consider the optimal solution of the different strategies in Table 6.

In the combined strategy, the manufacturer first buys more capacity and second uses a larger fraction of option capacity. The effectiveness of an option contract depends on two factors: the price structure and the variability of the demand. After replacing the process 2a and 2b with a common process, the standard deviation of the demand for the common process is larger than the standard deviation for each of the original dedicated processes. Therefore, using the common process increases the effectiveness of the option contract. On the other hand, the option contract makes the capacity for the common process more flexible. As a result, using the option contract also amplifies the effectiveness of the risk pooling effect. Therefore, the combined strategy achieves a much higher percentage increase of profit.

7. Conclusion

In this paper, we present a model to study single period capacity planning in a general supply chain with multiple contract types. The model incorporates several important factors such as uncertain product demands, common processes, flexible resources, and option contracts that the manufacturer needs to consider when it does its capacity planning. For two special supply chain structures, we develop closed-form representations of the optimal capacity plan. For a general supply chain, we develop and test an efficient algorithm to find the optimal capacity plan. Manufacturers can use the findings in this paper to plan their capacity systematically. Furthermore, they can also examine the effects of different planning strategies such as using common processes and option contracts.

A natural extension of this paper is to generalize the single period problem to a multi-period setting. When planning for multi-periods, the manufacturer needs to consider the durations of the contracts. Therefore, to maximize its expected profit, the manufacturer first needs to select the contracts with suitable durations and then to decide the size of each contract. However, since the multi-period capacity planning problem is not separable, there are exponentially many combinations of contracts that the firm can choose. As a result, the multi-period problem is much more complicated than the single period problem. We have formulated a model to study multi-period capacity planning based on the work presented in this paper; we have also developed a heuristic solution procedure for finding multi-period capacity plans. Many of the results and algorithms that we develop for the single period problem play important roles in the multi-period case. We cover the results and findings for the multi-period capacity planning problem in the sequel of this paper ([Huang and Graves 08]).

Appendix 1: Proofs

Proof: Proof of Proposition 1:

Let $(\mathbf{c}^1, \mathbf{g}^1)$ and $(\mathbf{c}^2, \mathbf{g}^2)$ be two feasible capacity planning strategies. Let λ be a scalar that $0 < \lambda < 1$. Then, capacity planning strategy

$$(\mathbf{c}^3, \mathbf{g}^3) = (\lambda \mathbf{c}^1 + (1 - \lambda) \mathbf{c}^2, \lambda \mathbf{g}^1 + (1 - \lambda) \mathbf{g}^2)$$

is also feasible. For any demand realization $\mathbf{d}$, let $(\mathbf{x}^i, \mathbf{y}^i, \mathbf{z}^i)$ be an optimal solution of Problem (1) given capacity planning strategy $(\mathbf{c}^i, \mathbf{g}^i)$. Fix a scalar $\lambda \in [0, 1]$ and consider

the production level

$$(\mathbf{x}^3, \mathbf{y}^3, \mathbf{z}^3) = (\lambda \mathbf{x}^1 + (1 - \lambda) \mathbf{x}^2, \lambda \mathbf{y}^1 + (1 - \lambda) \mathbf{y}^2, \lambda \mathbf{z}^1 + (1 - \lambda) \mathbf{z}^2).$$

We can verify that $(\mathbf{x}^3, \mathbf{y}^3, \mathbf{z}^3)$ is a feasible solution of Problem (1) given demand realization $\mathbf{d}$ and strategy $(\mathbf{c}^3, \mathbf{g}^3)$. Therefore

$$\begin{aligned} \pi(\mathbf{c}^3, \mathbf{g}^3, \mathbf{d}) &\geq \pi(\mathbf{c}^3, \mathbf{g}^3, d, \mathbf{x}^3, \mathbf{y}^3, \mathbf{z}^3) \\ &= \mathbf{r}' \mathbf{z}^3 - \mathbf{e}' H \mathbf{y}^3 \\ &= \lambda (\mathbf{r}' \mathbf{z}^1 - \mathbf{e}' H \mathbf{y}^1) + (1 - \lambda) (\mathbf{r}' \mathbf{z}^2 - \mathbf{e}' H \mathbf{y}^2) \\ &= \lambda \pi(\mathbf{c}^1, \mathbf{g}^1, \mathbf{d}) + (1 - \lambda) \pi(\mathbf{c}^2, \mathbf{g}^2, \mathbf{d}) \end{aligned}$$

Therefore, $\pi(\mathbf{c}, \mathbf{g}, \mathbf{d})$ is concave in $(\mathbf{c}, \mathbf{g})$ for any given $\mathbf{d}$. Since taking expectation will maintain concavity, $E[\pi(\mathbf{c}, \mathbf{g}, \mathbf{D})]$ is concave in $(\mathbf{c}, \mathbf{g})$. Therefore

$$\begin{aligned} \Pi(\mathbf{c}^3, \mathbf{g}^3, \mathbf{D}) &= E[\pi(\mathbf{c}^3, \mathbf{g}^3, \mathbf{D})] - \mathbf{p}' \mathbf{c}^3 - \mathbf{q}' (\mathbf{g}^3 - \mathbf{c}^3) \\ &\geq \lambda \left(E[\pi(\mathbf{c}^1, \mathbf{g}^1, \mathbf{D})] - \mathbf{p}' \mathbf{c}^1 - \mathbf{q}' (\mathbf{g}^1 - \mathbf{c}^1) \right) \\ &\quad + (1 - \lambda) \left(E[\pi(\mathbf{c}^2, \mathbf{g}^2, \mathbf{D})] - \mathbf{p}' \mathbf{c}^2 - \mathbf{q}' (\mathbf{g}^2 - \mathbf{c}^2) \right) \\ &= \lambda \Pi(\mathbf{c}^1, \mathbf{g}^1, \mathbf{D}) + (1 - \lambda) \Pi(\mathbf{c}^2, \mathbf{g}^2, \mathbf{D}). \end{aligned}$$

Therefore, $\Pi(\mathbf{c}, \mathbf{g}, \mathbf{D})$ is concave in $(\mathbf{c}, \mathbf{g})$. **Q.E.D.**

Proof: Proof of Proposition 2:

By Proposition 1, $\Pi(\mathbf{c}, \mathbf{g}, \mathbf{D})$ is concave in both $\mathbf{c}$ and $\mathbf{g}$. Therefore, the first order necessary conditions will also be sufficient conditions for optimality. If $r > \sum_j e_j$, for given capacity plan $(\mathbf{c}, \mathbf{g})$ and product demand d , the maximal profit is as follows:

$$\pi(\mathbf{c}, \mathbf{g}, d) = r \min\{d, \min_j \{g_j\}\} - \sum_{j=1}^J [\min\{d, g_j\} - \min\{d, c_j\}] e_j \quad (18)$$

We can show by contradiction that under the optimal planning strategy g_j is the same for all j . Thus, we let

$$g_j = g, \forall j,$$

and we can rewrite Problem (2) as

$$\begin{aligned} \arg \min_{\mathbf{c}, \mathbf{g}} \quad & -\mathbf{E} \left[r \min\{D, g\} - \sum_{j=1}^J (c_j p_j + (g - c_j) q_j) - \sum_{j=1}^J [\min\{D, g\} - \min\{D, c_j\}] e_j \right] \\ \text{s.t.} \quad & c_j \leq g, \forall j \end{aligned}$$

Since the constraints are linearly independent, the lagrange multipliers exist. Then, we can consider its Lagrange function

$$L(g, \mathbf{c}, \mu) = -\mathbf{E}_D \left[r \min\{D, g\} - \sum_{j=1}^J (c_j p_j + (g - c_j) q_j) - \sum_{j=1}^J [\min\{D, g\} - \min\{D, c_j\}] e_j \right] + \sum_{j=1}^J \mu_j (c_j - g) \quad (19)$$

By the first order necessary conditions, we have

$$\frac{\partial L}{\partial c_j} = p_j - q_j - Pr(D > c_j) e_j + \mu_j = 0 \quad (20)$$

$$\frac{\partial L}{\partial g} = \sum_{j=1}^J q_j + Pr(D > g) \left(\sum_{j=1}^J e_j - r \right) - \sum_{j=1}^J \mu_j = 0 \quad (21)$$

From Equation (20) we have

$$Pr(D > c_j) = \frac{p_j - q_j + \mu_j}{e_j}, \quad \forall j$$

Then we will have two cases:

Case 1: There exists a process i such that $c_i < g$. Define i to be the process with the smallest index such that $c_i < g$. If $c_i < g$ then, $\mu_i = 0$. Therefore,

$$Pr(D > c_i) = \frac{p_i - q_i}{e_i}$$

Now, assume that there exist a $j > i$ such that $c_j = g$. Since

$$\frac{p_j - q_j}{e_j} \geq \frac{p_i - q_i}{e_i}$$

and $\mu_j \geq 0$,

$$Pr(D > c_j) = \frac{p_j - q_j + \mu_j}{e_j} \geq \frac{p_i - q_i}{e_i} = Pr(D > c_i)$$

This implies $c_j \leq c_i$. However,

$$g = c_j \leq c_i < g,$$

and this is a contradiction. Therefore,

$$c_j \begin{cases} < g, & \text{if } j \geq i; \\ = g, & \text{if } j < i. \end{cases} \quad (22)$$

We note that the second part of Equation (22) follows because we chose i to be the smallest index such that $c_i < g$. Moreover, for all j such that $c_j < g$, $\mu_j = 0$ and

$$Pr(D > c_j) = \frac{p_j - q_j}{e_j}$$

If we let $\psi = i$, we have shown that conditions (5) and (6) hold. Since $\mu_j = 0$ for all $j \geq i$, from Equation (21) we have

$$\sum_{j=1}^J q_j + Pr(D > g) \left(\sum_{j=1}^J e_j - r \right) - \sum_{j=1}^{\psi-1} \mu_j = 0. \quad (23)$$

From Equation (20) and (22), we have

$$\mu_j = -p_j + q_j + Pr(D > g)e_j, \quad \forall j < \psi. \quad (24)$$

Then by Equation (23) and (24), we can re-express Equation (21) as

$$\sum_{j=1}^J q_j + Pr(D > g) \left(\sum_{j=1}^J e_j - r \right) + \sum_{j=1}^{\psi-1} (p_j - q_j - Pr(D > g)e_j) = 0$$

Simplifying the equation above we get,

$$\sum_{j=1}^{\psi-1} p_j + \sum_{j=\psi}^J q_j + Pr(D > g) \left(\sum_{j=\psi}^J e_j - r \right) = 0$$

Therefore,

$$Pr(D > g) = \frac{\sum_{j=1}^{\psi-1} p_j + \sum_{j=\psi}^J q_j}{r - \sum_{j=\psi}^J e_j}$$

which is condition (7). Finally, since $c_i < g$ for all $i \geq \psi$, then

$$\frac{p_i - q_i}{e_i} = Pr(D > c_i) > Pr(D > g) = \frac{\sum_{j=1}^{\psi-1} p_j + \sum_{j=\psi}^J q_j}{r - \sum_{j=\psi}^J e_j}$$

which shows condition (8) holds for $i = \psi$.

Case 2: $c_j = g$ for all j . By Equation (20) and (21) we have

$$\sum_{j=1}^J q_j + Pr(D > g) \left(\sum_{j=1}^J e_j - r \right) + \sum_{j=1}^J (p_j - q_j - Pr(D > g)e_j) = 0.$$

This implies

$$Pr(D > g) = \frac{\sum_{j=1}^J p_j}{r}.$$

For this case, $\psi = J + 1$ and we can verify that all of the conditions are satisfied.

Q.E.D.

Proof: Proof of Proposition 3:

By Proposition 1, it will be sufficient to show that one of the two sets of conditions is in fact the first order necessary condition for Problem (2). For any $(c, g, \mathbf{d})$, we can write

$$\begin{aligned} \pi(c, g, \mathbf{d}) = & \min\{d_1, g\}r_1 + \sum_{i=2}^M \left(\max \left\{ 0, g - \sum_{j=1}^{i-1} d_j \right\} r_i - \max \left\{ 0, g - \sum_{j=1}^i d_j \right\} r_i \right) \\ & - \left[\min \left\{ \sum_{i=1}^M d_i, g \right\} - \min \left\{ \sum_{i=1}^M d_i, c \right\} \right] e. \end{aligned}$$

Therefore, we can write Problem (2) as follows

$$\begin{aligned} \arg \min_{c, g} \quad & -\mathbf{E}[\pi(c, g, \mathbf{D})] + pc + q(g - c) \\ \text{s.t.} \quad & c \leq g. \end{aligned}$$

Since lagrange multiplier exists, we can write the lagrange function

$$L(c, g, \mathbf{D}) = -\mathbf{E}[\pi(c, g, D)] + pc + q(g - c) + \mu(c - g).$$

From the first order necessary condition, we get

$$\frac{\partial L}{\partial c} = p - q - ePr \left(c < \sum_{i=1}^M D_i \right) + \mu = 0 \quad (25)$$

and

$$\begin{aligned} \frac{\partial L}{\partial g} = & -r_1Pr(g < D_1) + \sum_{i=2}^M \left(r_i \left(-Pr \left(g > \sum_{j=1}^{i-1} D_j \right) + Pr \left(g > \sum_{j=1}^i D_j \right) \right) \right) \\ & + q + ePr \left(g < \sum_{i=1}^M D_i \right) - \mu = 0 \quad (26) \end{aligned}$$

Case 1: If $c < g$, then $\mu = 0$. Then, Equation (25) implies condition (9) and Equation (26) implies condition (10). Therefore, the first order necessary condition is equivalent to the first set of conditions.

Case 2: If $c = g$, then from Equation (25) and (26) we have

$$\begin{aligned} & -r_1Pr(g < D_1) + \sum_{i=2}^M \left(r_i \left(-Pr \left(g > \sum_{j=1}^{i-1} D_j \right) + Pr \left(g > \sum_{j=1}^i D_j \right) \right) \right) \\ & + q + ePr \left(g < \sum_{i=1}^M D_i \right) + p - q - ePr \left(g < \sum_{i=1}^M D_i \right) = 0 \end{aligned}$$

Simplifying the equation above will get condition (12). Therefore, the first order necessary condition is equivalent to the second set of conditions.

Q.E.D.

Appendix 2: Bounds on Sampling Error

Bound Based on Hoeffding Inequality

Assume that we have a lower and an upper bound on the expected net revenue,

$$\pi_{min} \leq \pi(\mathbf{c}, \mathbf{g}, \mathbf{d}) \leq \pi_{max}.$$

By Hoeffding Inequality, we have

$$Pr(|\pi(\mathbf{c}, \mathbf{g}, \mathbf{D}) - \pi_L(\mathbf{c}, \mathbf{g}, \mathbf{D})| \geq \epsilon) \leq 2 \exp \left(\frac{-L\epsilon^2}{2(\pi_{max} - \pi_{min})^2} \right), \quad \epsilon > 0. \quad (27)$$

Therefore,

$$Pr(|\pi(\mathbf{c}_L^*, \mathbf{g}_L^*, \mathbf{D}) - \pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*, \mathbf{D})| \geq \epsilon) \leq 2 \exp \left(\frac{-L\epsilon^2}{2(\pi_{max} - \pi_{min})^2} \right), \quad \epsilon > 0 \quad (28)$$

and

$$Pr(|\pi(\mathbf{c}^*, \mathbf{g}^*, \mathbf{D}) - \pi_L(\mathbf{c}^*, \mathbf{g}^*, \mathbf{D})| \geq \epsilon) \leq 2 \exp \left(\frac{-L\epsilon^2}{2(\pi_{max} - \pi_{min})^2} \right), \quad \epsilon > 0. \quad (29)$$

Since $(\mathbf{c}^*, \mathbf{g}^*)$ is the optimal solution of problem Π and $(\mathbf{c}_L^*, \mathbf{g}_L^*)$ is the optimal solution of problem Π_L , we have

$$\Pi(\mathbf{c}^*, \mathbf{g}^*) \geq \Pi(\mathbf{c}_L^*, \mathbf{g}_L^*) \quad (30)$$

and

$$\Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*) \geq \Pi_L(\mathbf{c}^*, \mathbf{g}^*). \quad (31)$$

From Equation (31), we have

$$\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*) \leq \Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}^*, \mathbf{g}^*). \quad (32)$$

Similarly, from Equation (30), we have

$$\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*) \geq \Pi(\mathbf{c}_L^*, \mathbf{g}_L^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*). \quad (33)$$

By Eqs. (28), (29), (32) and (33), we get

$$\begin{aligned} & Pr(|\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)| \leq \epsilon) \\ & \geq Pr(|\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}^*, \mathbf{g}^*)| \leq \epsilon \text{ AND } |\Pi(\mathbf{c}_L^*, \mathbf{g}_L^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)| \leq \epsilon) \\ & = Pr(|\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}^*, \mathbf{g}^*)| \leq \epsilon) + Pr(|\Pi(\mathbf{c}_L^*, \mathbf{g}_L^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)| \leq \epsilon) \\ & \quad - Pr(|\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}^*, \mathbf{g}^*)| \leq \epsilon \text{ OR } |\Pi(\mathbf{c}_L^*, \mathbf{g}_L^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)| \leq \epsilon) \end{aligned}$$

$$\begin{aligned}
&\geq \Pr(|\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}^*, \mathbf{g}^*)| \leq \epsilon) + \Pr(|\Pi(\mathbf{c}_L^*, \mathbf{g}_L^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)| \leq \epsilon) - 1 \\
&\geq 2 \left(1 - 2 \exp \left(\frac{-L\epsilon^2}{2(\pi_{\max} - \pi_{\min})^2} \right) \right) - 1 \\
&= 1 - 4 \exp \left(\frac{-L\epsilon^2}{2(\pi_{\max} - \pi_{\min})^2} \right)
\end{aligned} \tag{34}$$

Bound Based on Chernoff Inequality

If the manufacturer has an estimate of the maximum standard deviation, σ_π , of the expected profit, $\pi(\mathbf{c}, \mathbf{g}, \mathbf{D})$, and knows that the expected profit is bounded, it can bound the quantity $|\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)|$ using the Chernoff inequality.

By Chernoff inequality, we have

$$\Pr(|\pi(\mathbf{c}_L^*, \mathbf{g}_L^*, \mathbf{D}) - \pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*, \mathbf{D})| \geq \epsilon) \leq 2 \exp \left(\frac{-L\epsilon^2}{4\sigma^2} \right), \quad \epsilon > 0$$

and

$$\Pr(|\pi(\mathbf{c}^*, \mathbf{g}^*, \mathbf{D}) - \pi_L(\mathbf{c}^*, \mathbf{g}^*, \mathbf{D})| \geq \epsilon) \leq 2 \exp \left(\frac{-L\epsilon^2}{4\sigma^2} \right), \quad \epsilon > 0.$$

Following a similar argument given above, we have

$$\Pr(|\Pi(\mathbf{c}^*, \mathbf{g}^*) - \Pi_L(\mathbf{c}_L^*, \mathbf{g}_L^*)| \leq \epsilon) \geq 1 - 4 \exp \left(\frac{-L\epsilon^2}{4\sigma^2} \right). \tag{35}$$

Appendix 3: Deriving the Updates Rules in the Stochastic Supporting Hyperplane Algorithm

In this appendix, we look at the update rules in step 2 of the stochastic supporting hyperplane algorithm given in Chapter 2. The derivation is based on the method given by [Higle and Sen 96].

In step s , if we use all the sample points in the set $\mathcal{S}^s$ to construct the supporting hyperplane at $(\mathbf{c}^{s-1}, \mathbf{g}^{s-1})$, the constraint to be added as given in (17) is

$$\begin{aligned}
&f + \frac{1}{s} \sum_{\mathbf{d} \in \mathcal{S}^s} \pi(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}) - \mathbf{p}' \mathbf{c}^{s-1} - \mathbf{q}' (\mathbf{g}^{s-1} - \mathbf{c}^{s-1}) \\
&+ [(\mathbf{c}, \mathbf{g}) - (\mathbf{c}^{s-1}, \mathbf{g}^{s-1})]' \left(\frac{1}{s} \sum_{\mathbf{d} \in \mathcal{S}^s} \lambda(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}) - \mathbf{p} + \mathbf{q}, \frac{1}{s} \sum_{\mathbf{d} \in \mathcal{S}^s} \gamma(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}) - \mathbf{q} \right) \geq 0
\end{aligned}$$

Simplifying the equation above, we get

$$f + \frac{1}{s} \sum_{\mathbf{d} \in \mathcal{S}^s} \pi(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d})$$

$$\begin{aligned}
& -(\mathbf{c}^{s-1}, \mathbf{g}^{s-1})' \left(\frac{1}{s} \sum_{\mathbf{d} \in \mathcal{S}^s} \lambda(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}), \frac{1}{s} \sum_{\mathbf{d} \in \mathcal{S}^s} \gamma(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}) \right) \\
& +(\mathbf{c}, \mathbf{g})' \left(\frac{1}{s} \sum_{\mathbf{d} \in \mathcal{S}^s} \lambda(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}) - \mathbf{p} + \mathbf{q}, \frac{1}{s} \sum_{\mathbf{d} \in \mathcal{S}^s} \gamma(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}) - \mathbf{q} \right) \geq 0
\end{aligned}$$

If we separate the terms associated with ω^s , which is the demand sample generated in step s , we get

$$\begin{aligned}
& f + \frac{s-1}{s} \left[\frac{1}{s-1} \sum_{\mathbf{d} \in \mathcal{S}^{s-1}} \pi(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}) \right. \\
& -(\mathbf{c}^{s-1}, \mathbf{g}^{s-1})' \left(\frac{1}{s-1} \sum_{\mathbf{d} \in \mathcal{S}^{s-1}} \lambda(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}), \frac{1}{s-1} \sum_{\mathbf{d} \in \mathcal{S}^{s-1}} \gamma(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}) \right) \\
& \left. +(\mathbf{c}, \mathbf{g})' \left(\frac{1}{s-1} \sum_{\mathbf{d} \in \mathcal{S}^{s-1}} \lambda(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}) - \mathbf{p} + \mathbf{q}, \frac{1}{s-1} \sum_{\mathbf{d} \in \mathcal{S}^{s-1}} \gamma(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \mathbf{d}) - \mathbf{q} \right) \right] \\
& + \frac{1}{s} \left[\pi(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \omega^s) - (\mathbf{c}^{s-1}, \mathbf{g}^{s-1})' (\lambda(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \omega^s), \gamma(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \omega^s)) \right. \\
& \left. +(\mathbf{c}, \mathbf{g})' (\lambda(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \omega^s) - \mathbf{p} + \mathbf{q}, \gamma(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \omega^s) - \mathbf{q}) \right] \geq 0
\end{aligned} \tag{36}$$

By the definitions of α , β , and ζ , Equation (36) can be written as

$$\begin{aligned}
& f + \frac{s-1}{s} \left[\alpha_{s-1}^{s-1} + (\mathbf{c}, \mathbf{g})' (\beta_{s-1}^{s-1}, \zeta_{s-1}^{s-1}) \right] + \frac{1}{s} (\mathbf{c}, \mathbf{g})' (-\mathbf{p} + \mathbf{q}, -\mathbf{q}) \\
& + \frac{1}{s} \left[\pi(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \omega^s) - (\mathbf{c}^{s-1}, \mathbf{g}^{s-1})' (\lambda(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \omega^s), \gamma(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \omega^s)) \right. \\
& \left. +(\mathbf{c}, \mathbf{g})' (\lambda(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \omega^s), \gamma(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \omega^s)) \right] \geq 0
\end{aligned} \tag{37}$$

Now if we set $(\mathbf{c}, \mathbf{g}) = (\mathbf{c}^{s-1}, \mathbf{g}^{s-1})$, the last three terms of the Equation (37) will be

$$\frac{1}{s} \left[\pi(\mathbf{c}^{s-1}, \mathbf{g}^{s-1}, \omega^s) \right]$$

If we replace this with its upper bound U , we get

$$f + \frac{s-1}{s} \left[\alpha_{s-1}^{s-1} + (\mathbf{c}, \mathbf{g})' (\beta_{s-1}^{s-1}, \zeta_{s-1}^{s-1}) \right] + \frac{1}{s} (\mathbf{c}, \mathbf{g})' (-\mathbf{p} + \mathbf{q}, -\mathbf{q}) + \frac{1}{s} U \geq 0 \tag{38}$$

Equation (38) uses the hyperplane at $(\mathbf{c}^{s-1}, \mathbf{g}^{s-1})$ in step $s-1$ to construct a relaxation of the supporting hyperplane at the same point in step s . A similar relaxation is applied to all the hyperplanes generated before step s .

Appendix 4: Algorithms Runtime Comparison Results

Test Case	Regular	Stochastic	Pre-solve	Pre-solve/Stoc.	Pre-solve/Reg.
1	465	112	51	45.45%	10.97%
2	380	128	44	34.38%	11.58%
3	373	140	67	47.86%	17.96%
4	678	152	93	61.18%	13.72%
5	791	513	130	25.34%	16.43%
6	679	392	79	20.15%	11.63%
7	259	119	40	33.61%	15.44%
8	239	43	36	83.72%	15.06%
9	315	125	45	36.00%	14.29%
10	260	75	36	48.00%	13.85%
11	473	104	71	68.27%	15.01%
12	262	72	47	65.28%	17.94%
13	259	58	37	63.79%	14.29%
14	534	147	71	48.30%	13.30%
15	314	278	43	15.47%	13.69%
16	386	100	51	51.00%	13.21%
17	464	115	50	43.48%	10.78%
18	442	163	44	26.99%	9.95%
19	293	110	40	36.36%	13.65%
20	554	144	68	47.22%	12.27%
21	534	215	72	33.49%	13.48%
22	267	62	34	54.84%	12.73%
23	299	117	41	35.04%	13.71%
24	340	79	39	49.37%	11.47%
25	274	80	44	55.00%	16.06%
26	334	58	44	75.86%	13.17%
27	231	60	29	48.33%	12.55%
28	410	102	50	49.02%	12.20%
29	423	140	45	32.14%	10.64%
30	553	114	50	43.86%	9.04%
31	311	103	43	41.75%	13.83%
32	451	136	37	27.21%	8.20%
33	488	220	43	19.55%	8.81%
34	472	179	51	28.49%	10.81%
35	440	214	79	36.92%	17.95%
36	515	295	77	26.10%	14.95%
37	174	99	31	31.31%	17.82%
38	552	211	55	26.07%	9.96%
39	484	103	74	71.84%	15.29%
40	294	74	34	45.95%	11.56%

Table 7: Run time (in seconds) comparison of Regular Supporting Hyperplane algorithm, Stochastic Supporting Hyperplane algorithm, and Stochastic Supporting Hyperplane algorithm with Pre-solve Routine.

Acknowledgments

This research has been funded by the Singapore MIT Alliance (SMA).

References

- [Cheng et al. 07] Cheng, F., Ettl, M., Lin, G. Y., Schwarz, M., and Yao, D. D., “Flexible Supply Contracts via Options”, working paper.
- [Fine and Freund 90] Fine, C. H. and Freund, M., “Optimal Investment in Product-Flexible Manufacturing Capacity,” *Management Science*, Vol. 36, No. 4, 1990, pp. 449-466.
- [Fu et al. 06] Fu, Q., Lee, C. Y., and Teo, C. P., “Procurement Risk Management using Options: Random Spot Price and the Portfolio Effect”, working paper, 2006.
- [Gerchak et al. 88] Gerchak, Y., Magazine, M. J., Gamble, A. B., “Component Commonality with Service Level Requirements”, *Management Science*, Vol. 34, No. 6, 1988, 753-760.
- [Higle and Sen 91] Higle, J. and Sen, S., “Stochastic Decomposition: An Algorithm for Two Stage Linear Programs with Resource,” *Mathematics of Operations Research*, Vol. 16, No. 3, 1991. pp. 650-669.
- [Higle and Sen 96] Higle, J. L. and Sen, S., “Stochastic Decomposition: A Statistical Method for Large Scale Stochastic Linear Programming”, Springer, 1996.
- [Hillier 00] Hillier, M. S., “Component Commonality in Multiple-period Assemble-to-order Systems”, *IIE Transactions*, Vol. 32, No. 8, 2000, pp. 755-766.
- [Hoeffding 63] Hoeffding, W., “Probability inequalities for sums of bounded random variables”, *Journal of the American Statistical Association*, Vol. 58, No. 301, 1963, pg. 1330.
- [Huang 08] Huang, X., “Capacity Planning in a General Supply Chain with Multiple Contract Types”, Ph.D. Thesis, Massachusetts Institute of Technology, 2008.
- [Huang and Graves 08] Huang, X. and Graves, S. C., “Multiple-Period Capacity Planning in a General Supply Chain with Outsourcing Contracts”, working paper, 2008.

- [Kulkarni et al. 04] Kulkarni, S. S., Magazine, M. J., and Raturi, A. S., “Risk Pooling Advantages of Manufacturing Network Configuration”, *Production and Operations Management*, Vol. 13, No. 2, 2004, pp. 186-199.
- [Labro 04] Labro, E., “The Costs Effects of Component Commonality: A Literature Review Through a Management-Accounting Lens”, *Manufacturing & Service Operations Management*, Vol. 6, No. 4, 2004, pp. 358-367.
- [Martinez-de-Albéniz and Simchi-Levi 05] Martinez-de-Albéniz, V. and Simchi-Levi, D., “A Portfolio Approach to Procurement Contracts,” *Production and Operations Management*, Vol. 14, No. 1, 2005, pg. 90-114.
- [Nagali et al. 08] Nagali, V., Hwang, J., Sanghera, D., Gaskins, M., Pridgen, M., Thurston, T., Mackenroth, P., Scholler, P., and Shoemaker G., “Procurement Risk Management (PRM) at Hewlett-Packard Company”, *Interfaces*, Vol. 38, No.1, January-February, 2008, pp. 51-60.
- [Song and Zipkin 03] Song, J. S. and Zipkin, P., “Supply Chain Operations: Assemble-to-Order Systems,”, *Chapter 11 in Handbooks in Operations Research and Management Science*, Vol. 30, edited by T. De Kok, and S. Graves. Amsterdam: North-Holland, 2003.
- [Van Mieghem and Rudi 02] Van Mieghem, J. A. and Rudi, N., “Newsvendor Networks: Inventory Management and Capacity Investment with Discretionary Activities,” *Manufacturing and Service Operations Management*, Vol. 4, No. 4, 2002, pg. 313-335.
- [Van Mieghem 98] Van Mieghem, J. A., “Investment Strategies for Flexible Resources,” *Management Science*, Vol. 44, No. 8, August 1998, pp. 1071-1078.
- [Van Mieghem 07] Van Mieghem, J. A., “Operations Strategy: Practices and Principles”, Dynamic Ideas, 2007.
- [Veinott 67] Veinott, A. F., “The Supporting Hyperplane Method for Unimodal Programming,” *Operations Research*, Vol. 18, 1967, pp. 404-413.
- [Weck et al. 04] Weck, O., Neufville, R., and Chaize, M., “Staged Deployment of Communications Satellite Constellations in Low Earth Orbit,” *Journal of Aerospace Computing, Information, and Communication*, Vol. 1, 2004, pg. 119-136.

[Yazlali and Frhun 06] Yazlali, Ö. and Frhun, F., “Managing Demand Uncertainty with Dual Supply Contracts,” Working paper, 2006.

[Wikipedia 08] Wikipedia, http://en.wikipedia.org/wiki/Chernoff%27s_inequality.