10 MODELING SUPPLY CHAIN CONTRACTS: A REVIEW

Andy A. Tsay
Steven Nahmias
Narendra Agrawal

Department of Operations & Management Information Systems
Leavey School of Business
Santa Clara University
Santa Clara, CA 95053

.
.
.
.
.
.
.

Chapter 10: *Quantitative Models for Supply Chain Management*
S. Tayur, R. Ganeshan, and M. Magazine (Eds.)
Kluwer Academic Publishers, MA, 1999, pp. 299-336

10.1 INTRODUCTION

In this review, we summarize model-based research on contracts in the supply chain setting and provide a taxonomy for work in this area. During our discussions it became clear that the field has developed in many directions at once. Furthermore, as we surveyed the literature, it was not obvious what constitutes a contract in this context. While the nomenclature "supply chain management" is relatively new, many of the problems that are addressed are not. In particular, mathematical models for optimizing inventory control have a long history as a significant part of the mainstream of operations research and operations management. Inventory modeling, per se, dates to the early part of the century and the ideas of a Westinghouse engineer named Ford Harris (1915). A natural issue to address first is what is meant by supply chain management (SCM) research and how it relates to the vast body of work constituting classical inventory theory. Modern usage of the term seems to be consistent with the following definition: *a supply chain is two or more parties linked by a flow of goods, information, and funds.* One way to interpret this is that SCM research is essentially the same as multi-echelon inventory theory, introduced by Clark and Scarf (1960) to model logistics problems encountered by the military. Using a dynamic programming formulation, the authors derived the optimal ordering and transshipment policy for a single-product serial system facing independent, identically distributed demand. The optimal policy is characterized by a critical "order-up-to" echelon inventory target for each installation, where echelon inventory counts not only the local stock but also the inventory downstream in the system. On the whole, much of the work on multi-echelon inventory theory relies on relaxing assumptions in the basic Clark-Scarf formulation. Experience thus far suggests that optimal control policies can be calculated only in limited settings. We direct the reader to Federgruen (1993) for a recent review of multi-echelon inventory models. Clark (1972) is also recommended for its excellent coverage of the earlier work. While the basic structure of the systems studied can be similar, we believe that SCM research encompasses a much broader scope of issues than does multi-echelon inventory theory. That is, while multi-echelon inventory theory is primarily about controlling the timing and quantity of material flows, SCM studies this and more. For instance, SCM treats environments in which there are multiple decision makers, which may be different firms or different divisions within a single firm. Behavior that is locally rational can be inefficient from a global perspective (cf. Whang 1995), so attention turns to methods for improving system efficiencies. Some mechanisms for making these improvements are the contractual arrangements surveyed here. These include the reallocation of decision rights, rules for sharing the costs of inventory and stockout, and policies governing pricing to the end-customer or between supply chain partners. Representation of the information structure and rules for information sharing are also important, since the assumption of common information usually made in multi-echelon inventory theory may be inaccurate in real supply chains. SCM also considers the topology of the system; that is, the number of suppliers, distributors or retailers. While traditional multi-echelon inventory theory tends to deal with issues of transportation delays via lead times, modern SCM studies the implications of alternative modes of logistics. One example is the current trend towards the outsourcing of the logistics function to external parties whose comparative advantage derives from scale and focus. Even the product itself may be redesigned to affect supply chain performance improvement. For example, several studies consider the extent to which delaying product differentiation ("postponement") affects volume and inventory pooling efficiencies within the supply chain (Lee et al 1993, Lee 1996). Of course, aspects of these problem areas have been considered in other fields, which suggests a distinguishing characteristic of SCM research: like supply chains themselves, the research questions and methodologies tend to cross traditional functional lines. As suggested above, the management of supply chains consisting of multiple agents with possibly conflicting objectives requires consideration of the relationships among the parties. Recent research explores arrangements not considered in traditional inventory theory. We will refer to these structures as contracts. Because much of the research in this area is quite recent, many of the papers discussed have not yet appeared in the open literature, and are available as working papers only. Our including a paper in this review does not mean we believe it to be completely correct, nor that it will eventually pass peer review and appear in print. We believe this is the first review to try to provide a comprehensive description and classification of the model-based analyses of contracts in supply chain management.

10.2 SCOPE OF THE REVIEW

Before beginning our review, we note that contracts are an important area of study in disciplines other than SCM. Contracts are, of course, a major consideration in law, and the literature in this area is enormous. There is also a substantial literature on contracts in the economics literature (Tirole 1988 and Katz 1989 provide excellent reviews). These bodies of work contribute deep understanding of basic issues of motivation for a broad variety of contractual structures, with particular attention to their legality, enforceability, and ramifications for public policy and social welfare. Furthermore, they are the source of many concepts and techniques which SCM researchers make use of and build upon. What distinguishes SCM contract analysis may be its focus on operational details, requiring more explicit modeling of materials flows and complicating factors such as uncertainty in the supply or demand of products, forecasting and the possibility of revising those forecasts, constrained production capacity, and penalties for overtime and expediting. Due to space limitations, we review only those papers from other fields which are considered foundational to a stream of SCM research. Even after restricting our search in this way, we were faced with the challenge of defining what constitutes a "contracts" paper in the SCM context. This challenge arises because, broadly speaking, all literature on inventory theory could qualify. Certainly, the purchase of materials implies an agreement between two parties even if the behavior and preferences of one party are suppressed. We adopt the more narrow definition that the analysis must explicitly offer guidance in negotiating the terms of the relationship between buyer and seller. This would mean that classical inventory theory, in which the supply parameters (e.g. price, lead time, bounds on order size) are typically treated as exogenous environmental characteristics, would not satisfy this test. Thus, we review in depth only those papers which treat the terms of the contractual relationship as decision variables, or at the very least investigate the behavioral and performance consequences of changing these terms. The specific terms on which we focus are described later in §10.3.3.

10.3 FUNDAMENTALS OF CONTRACT ANALYSIS

In this section we describe the structure of a basic supply chain to provide a context for the subsequent discussion. Specifically, we use this to identify the decisions to be made, the relevant environmental components, and the behavioral dynamics which underlie the contracts discussed.

10.3.1 Supply Chain Structure

Consider a supply chain in which an upstream party (which we refer to as a manufacturer) provides a single product to a downstream party (which we refer to as a retailer), who in turn serves market demand, as shown in Figure 10.1. This scenario could describe the link between any two consecutive nodes in a supply chain, and indeed on occasion the tandem may be referred to as supplier and manufacturer, manufacturer and distributor, or, most generally, supplier and buyer. Researchers commonly make the following simplifying assumptions to render the analysis more tractable. The manufacturer produces (or acquires) the product at a constant unit cost of c and charges the retailer a *wholesale/transfer payment* of $W(Q)$ for Q units. $W(Q)$ may either be exogenous, or a decision variable under the control of one of the parties. The retailer, in turn, sells the product at a price of p per unit. Market demand, denoted as $D(p)$, in reality is both price-sensitive and uncertain. While some models include both these features, it is more common to either take the retail price as fixed and represent market demand as a random variable (as in the operations research literature), or assume a deterministic, downward-sloping demand curve (as in the economics and marketing literatures). In the latter case the retailer's decision is primarily p, whereas in the former it is Q. A simpler underlying structure allows traditional inventory models to treat more complex problem settings including multiple periods, continuous review, and finite and infinite horizons. However, most contract papers assume only a one-period problem (i.e., a newsvendor setting), since the resulting models are often too complex to be amenable to multi-period analysis.

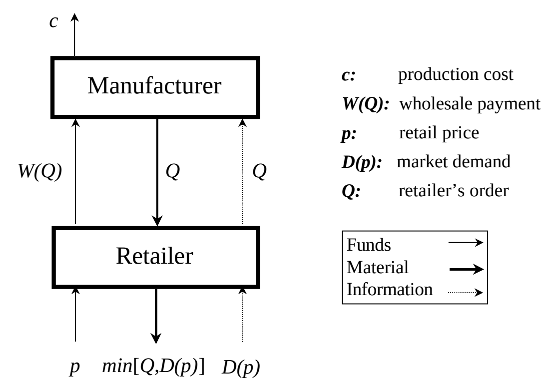

Figure 10.1 The Supply Chain

10.3.2 Purposes of Contracts

To understand what motivates the parties to pursue certain contract structures, consider the simple supply chain pictured in Figure 10.1. In the best of all possible worlds, total expected supply chain profits would be maximized if all decisions are made by a single decision maker with access to all available information. This is referred to as the *first-best case*, and is often associated with *central control*. Denote the resulting expected profit as Π^c in this case. However, in general neither the manufacturer nor the retailer is in a position to control the entire supply chain, and each has his own incentives and state of information. Refer to this as a *decentralized control* structure, with resulting total expected profit Π^d. It is *inefficient* if $\Pi^d < \Pi^c$. Some contracts focus on how Π^d is to be split between the two parties. We call this the *risk-sharing* objective in that it provides a means for the buyer and supplier to share the risks arising from various sources of uncertainty, e.g. market demand, selling price, process yield, product quality, delivery time, and exchange rates. The risk sharing motive is common in the contracts reviewed here. As an example, suppose that the retailer is required to transmit sales forecasts upstream to manufacturers. These forecasts are intended to assist the manufacturer to make capacity and materials purchasing decisions. However, in most cases no commitment is attached to these forecasts. As a result, the manufacturer assumes a large portion of the risk of demand uncertainty. Not only might the retailer cancel orders if demand is lower than anticipated, there is an incentive for the retailer to deliberately inflate forecasts as a form of insurance. Minimum purchase agreements or penalties for returns are often included in contracts to protect the manufacturer against this behavior. Of course, a necessary condition for the adoption of any contractual agreement is that both parties ultimately benefit. Contracts also provide a means for bringing Π^d closer to Π^c, which we call the *system-wide performance improvement* objective. This is also referred to as the *channel coordination* objective (a phrase coined in the marketing literature). Channel coordination may be achieved by first identifying the intra-chain dynamics which cause the inefficiency, then modifying the structure of the relationship to more closely align individual incentives with global optimization. For example, the existence of two separate entities within the channel can lead to *double marginalization*, a well-known cause of supply chain inefficiency. This notion apparently originated in the economics literature (Spengler 1950), where the fundamental decision is usually p, and the tradeoff to be made is between the unit profit margin (which favors higher p) and the size of the market (which favors lower p). If the retailer pays a unit price $t > c$ to the manufacturer, hence creating the two distinct profit margins referred to by the name of this phenomenon, the retailer's choice of p will be consistent with a profit margin of $(p - t)$ rather than the $(p - c)$ that would be appropriate from a central planner's perspective. Double marginalization can also occur when p is fixed, demand is stochastic and Q is the decision variable (more commonly assumed in the inventory management literature). Here, the optimal Q is determined by considering the overstocking and understocking costs. While the consequence of overstocking may be the same whether control is centralized or decentralized, the retailer perceives an overstock cost of $(p - t)$ rather than $(p - c)$. The problem does not go away if the manufacturer controls the order quantity decision, since the manufacturer's cost of overstocking is $(t - c)$. Many of the contractual structures recently studied (cf. Pasternack 1985, Donohue 1996, Tsay 1996, Ha 1997a) attempt to remedy some variant of this basic problem. Contracts also facilitate *long-term partnerships* by delineating mutual concessions that favor the persistence of the business relationship, as well as specifying penalties for non-cooperative behavior. The lengthening of the time horizon may encourage parties to engage in activities that are unfavorable in the short term but have substantial payoffs over time. For example, chip vendor Intel might be willing to consign a large portion of its production of a new generation of microprocessors to a single computer maker, such as Dell, even though the chips would fetch more on the open market. Intel's motivation would be to build a long-term relationship in the hope that Dell would be a volume purchaser for many years to come. In addition to providing a reliable supply for a buyer and demand volume for a supplier, stable partnerships can reduce transactions costs and allow for greater cooperation (e.g. information sharing and collaborative product and process improvement). One effort to study some of these issues is Cohen and Agrawal (1998), which evaluates the impact of various contractual arrangements on total costs incurred by a risk-averse buyer firm. In particular, they analyze the tradeoffs between the flexibility offered by short-term contracts and the fixed investments, improvement opportunities, and price certainty associated with long-term contracts. However, few of the models we have encountered consider the time horizon of the contract as a decision variable. Another important rationale for a contract that is not typically modeled is that it *makes the terms of a relationship explicit.* In fact, in the course of making legally concrete the expectations of each party, a contract can suggest and unambiguously define quantifiable performance metrics which are prerequisite to any systematic process improvement effort. Lead times, on-time delivery rates, and conformance rates are among the metrics commonly specified in supply contracts.

10.3.3 A Classification Scheme

Having restricted our attention only to papers on "contracts" in the SCM context, we still needed to determine a suitable scheme for classifying the resulting set of papers. This was complicated by the fact that no commonly accepted taxonomy appears to exist, and papers that purport to study contracts in supply chains consider a wide range of problems and issues. After experimenting with several schemes, we decided to classify the literature by contract clauses. These include:

(a) specification of decision rights,
(b) pricing,
(c) minimum purchase commitments,
(d) quantity flexibility,
(e) buyback or returns policies,
(f) allocation rules,
(g) lead time, and
(h) quality.

Several papers were candidates for more than one of these categories. For example, a minimum purchase commitment may also require special pricing terms to attract the buyer. Also, virtually all contracts must explicitly specify decision rights in order to be executable. However, classification in the "specification of decision rights" category suggests that a paper's main thrust is to investigate the relative desirability of alternative configurations of decision authority. Naturally, there is an element of subjectivity in this categorization, and there may be some clauses discussed in the literature or used in practice that simply do not fit well into this classification scheme. However, this classification covers the majority of the research that we have seen. Below we briefly describe the above contract clauses as they relate to the supply chain in Figure 10.1.

(a) Specification of Decision Rights
Here the goal is to achieve specific objectives by reassigning control of the decision variables. For example, although the retailer typically chooses Q and p given the $W(Q)$ specified by the manufacturer, there is an arrangement called

Resale Price Maintenance (RPM) in which the manufacturer is allowed to dictate conditions on the p that the retailer may charge. Likewise, *Quantity Fixing* is in effect when the manufacturer exerts control over Q. The issue of local versus global control may also be interpreted as an issue of decision rights.

(b) Pricing

$W(Q)$ is the component of the contract between the manufacturer and retailer that defines the financial terms of the supply relationship. Typically, $W(Q) = F + tQ$ for constants F and t. $F = 0$ results in linear pricing, perhaps the most commonly assumed pricing structure. A positive F, often referred to as a franchise fee, results in "two-part tariff" pricing. Some researchers also treat more complex pricing schemes, such as quantity discounting.

(c) Minimum Purchase Commitments

Such an agreement requires the retailer to purchase a minimum quantity, either within each single transaction, or cumulatively over a specified time horizon. The manufacturer may reduce $W(Q)$ to provide an incentive to the retailer to agree to this arrangement.

(d) Quantity Flexibility

In a quantity flexibility clause, the quantity the retailer ultimately purchases may deviate from a previous planning estimate, subject to certain constraints and/or financial consequences. To properly represent such a setting requires a stochastic demand model in which some event (such as a forecast update) occurs within the time frame of the model to motivate the exercise of flexibility.

(e) Buyback or Returns Policies

A buyback clause is one that specifies that the retailer may return some or all unsold product to the manufacturer, possibly for only partial credit. Naturally, mismatches between the retailer's purchase and the market demand are only an issue when demand is assumed random.

(f) Allocation Rules

Allocation rules specify how the manufacturer's available stock or production capacity is to be distributed among multiple retailers in a shortage scenario.

(g) Lead Times

The lead time for delivery of product from the manufacturer to the retailer is treated in traditional inventory models as either a fixed constant (one special case being zero), or a realization of a random variable. This clause highlights the possible benefits of adjusting that lead time via a contractual agreement.

(h) Quality

Any supply relationship is premised on the quality of the delivered product. The specific notion of quality may be formalized within the contract.

10.4 REVIEW OF THE LITERATURE

In this section we review the current literature on contracts in supply chains. Each paper will be classified into one of the above categories and briefly summarized.

10.4.1 Specification of Decision Rights

While most of the vast multi-echelon inventory literature presumes central control (e.g. Clark and Scarf 1960, Bessler and Veinott 1966, Eppen and Schrage 1981, Federgruen and Zipkin 1984, Rosling 1989), some recent works have proceeded under the realization that localization of decision-making authority may be necessary for large enterprises operating in complex environments. This creates a different set of challenges. Conflicts in the interests of the different parties may lead to inefficiency. Double marginalization, as described earlier, can be interpreted in this way. Or, in a setting of information asymmetry, a key decision may be under the purview of a party with inferior information. Yet shifting control to the better-informed party may engender opportunistic behavior since self-interested actions might be undetectable. Hence, careful consideration of information and incentives is central to all attempts to improve supply chain performance by reconfiguration of decision rights. We note that the economics literature has contributed a great deal to the understanding of such possibilities, including the phenomena of Resale Price Maintenance (RPM) and Quantity Fixing. Rather than summarize this vast literature, we refer the reader to Katz (1989) for a review. In this section we first describe four papers that propose ways to facilitate the shift from centralized to decentralized control of a supply chain, followed by those which focus on solutions which transfer decision rights among the various independent agents. Lee and Whang (1997) consider the problem of coordinating a decentralized version of the Clark-Scarf serial system. While each site incurs a holding cost, only the site serving the end customer faces a shortage cost. As a result, the upstream sites will carry less buffer inventory than may be best for the system as a whole. Realizing this, the site furthest downstream tends to carry extra inventory, which is inefficient since finished goods usually are most costly to hold. The authors propose rules for performance measurement and accountability, which have the following desirable properties: (i) *cost conservation* (i.e., all costs can be traced to individual sites without the need for subsidies or taxes from the central planner), (ii) *incentive compatibility* (i.e., what is optimal for each individual manager is also optimal for the system as a whole), and (iii) *information decentralizability* (i.e., the rules can be implemented using local information only). Their scheme involves a consignment policy for redistributing inventory carrying costs among the sites, an additional backlog penalty paid to an upstream site by its direct customer, and a shortage reimbursement paid to a downstream site by its direct supplier. No markups are added to the transfer prices. However, implementing this scheme requires common knowledge about the demand distribution. Kumar et al (1996) consider a similar problem but use a different set of control mechanisms for coordination: (i) an internal transfer price and a shortage penalty imposed on the upstream site, and (ii) a service level constraint imposed on the upstream site. For the same model as Lee and Whang (1997), except with delays between the transmission and receipt of orders between sites, Chen (1997) restores system-optimal performance using a measurement scheme based on “accounting inventory”. A division's accounting inventory is its on-hand stock minus backlogs of orders placed by the downstream site, under the assumption that the upstream division is perfectly reliable. Porteus and Whang (1991) describe the conflicts of interest between the manufacturing and marketing divisions of a firm (which can be interpreted loosely as a "supply chain") over the production capacity decision. These conflicts can be reconciled by establishing an internal futures market for capacity. While the discussion in each of these papers concerns the coordination of individual agents within a single organization, we include them here because with merely a slight change in interpretation the mechanisms are equivalent to contracts. We now turn our attention to models in which control of certain decisions is transferred from one independent party to another. The focus of Blair and Lewis (1994) is on designing efficient contracts for a one-manufacturer, one-retailer system, where the manufacturer cannot observe the demand or the retailer's service effort. This introduces both information asymmetry and moral hazard. Market demand (Q) is assumed to be a function of the retail price (p), the retailer's service effort (X) and some exogenous parameter (θ), which is known only to the retailer and reported to the manufacturer at the time of contracting. The retailer's cost for service is assumed to be X. Demand is deterministic, and inventory considerations are ignored. The retailer learns about θ and reports it to the manager (not necessarily truthfully), who in turn offers a menu of contracts from which the retailer selects one. The retailer then chooses his service effort. Assuming full information, first-best performance can be attained when the manufacturer charges a franchise fee and the retailer sets the price, service effort and quantity. However, under information asymmetry and moral hazard, the efficient contracts require some form of resale price maintenance and quantity fixing. The structure is shown to depend upon how the impact of market demand on the retailer's optimal promotion varies with price and quantity, i.e., upon the signs of $\partial^2 X/\partial p \partial \theta$ and $\partial^2 X/\partial Q \partial \theta$. When both these derivatives are zero, there exists a contract that will induce the retailer to select the first-best price, quantity and effort level. Otherwise, the manufacturer must impose a ceiling or floor on the retailer's price, and use quantity fixing/rationing. Ha (1997a) studies the effect of decision rights in the presence of information asymmetry regarding the downstream member's production cost in an environment of demand uncertainty. He considers a two-member channel similar to that of Blair and Lewis (1994), but allows customer demand to be random and price-dependent, modeled as $D = \mu(p) + Y$, where $\mu(p)$ is deterministic and Y is random and independent of p. First, the author argues that under complete information, coordination can be achieved by various contractual arrangements: quantity fixing, franchising, or a returns policy with price fixing. Then, he considers the case in which the downstream member's variable production cost is private information (the upstream member has a probability distribution for this cost). Here, coordination can be achieved with a nonlinear pricing scheme with price fixing, in which the upstream member offers a menu of contracts that specifies the retail price, order quantity and a fixed fee. Narayanan and Raman (1997) focus on the discrepancy between the impacts perceived by the manufacturer and the retailer when the retailer stocks out of a product. Typical retailers carry products that substitute for that of a given manufacturer, so tend to stock less of that product than may be optimal for the manufacturer. In a newsvendor framework with complete information, the authors assume a market demand consisting of a deterministic component (a linear function of the manufacturer's promotional effort) and a component that is uniformly distributed between 0 and 1. The cost of the effort is modeled as the square of the effort level. The retail price is fixed. They compare the optimal stock levels and channel profit for three scenarios: (i) the first-best case (a vertically integrated channel that manufactures and retails the product and its substitutes), (ii) "retailer managed inventory" (RMI), in which the retailer controls the inventory decisions, and (iii) "vendor managed inventory" (VMI), in which the manufacturer has the control. This abstraction of VMI is analogous to what is known in the economics literature as Quantity Fixing. The authors show that VMI may or may not do better than RMI, but VMI's relative performance improves as the manufacturer's influence over demand increases. Finally, Agrawal and Tsay (1998) attempt to generalize the discussion on decision rights by focusing on the effect of intra-organizational goal incongruence on contract efficiency. They consider a single-period model of a supply chain consisting of an independent manufacturer and retailer. The retailer serves a price-sensitive stochastic demand of the form $D = N \cdot g(p)$, where N is random and $g(p)$ is a deterministic downward-sloping function of the retail price, p. Each party seeks to maximize its individual expected profits. However, the operational-level decisions in the retail organization (pricing and stocking level) are delegated by the owner to a manager whose compensation scheme induces behavior inconsistent with this objective. In particular, this manager is interested in maximizing the probability of meeting or exceeding a target profit level specified by the owner. The authors study how this impacts the behaviors and outcomes on both sides of the supply contract. Further, in order to determine the consequences for the welfare of end customers, they define the notion of *expected consumer surplus*. Their analysis evaluates the preferences of the various parties for wholesale and retail prices, inventory levels, and product and customer types. Certain implications for supply chain strategy are suggested: increasing the profitability of a supply chain need not be at the expense of the end customer, and coordination of goals within an organization does not necessarily improve the supply chain efficiency. Industry case studies in which significant supply chain benefits have accrued in conjunction with the reallocation of decision rights include the Vendor Managed Inventory agreements between Proctor & Gamble and WalMart and between Levi Strauss and some of its retail partners, the Efficient Consumer Response initiative in the grocery industry, and the JIT II movement pursued by some manufacturing companies. On the other hand, obstacles to such programs are also well documented, as in the case of Barilla SpA, the world's largest pasta manufacturer. Its Just-in-Time Distribution program was resisted by its distributors as well as its own personnel (Hammond 1994).

The transfer of inventory stocking decision rights to Barilla was perceived by the distributors as a strategic threat. Internally, the sales organization also opposed this program out of fear that it might adversely affect their incentive systems. Modeling of such implementation issues is an area for future research.

10.4.2 Pricing

In most of the traditional inventory models, $W(Q)$ is specified and is not subject to negotiation between the buyer and the seller. More recently, some researchers have considered this cost schedule as a means to modifying the behavior of one or both parties. In particular, they have considered the use of quantity discounting as a coordination mechanism. One stream of such research includes Monahan (1984), Rosenblatt and Lee (1985), Lee and Rosenblatt (1986), and Banerjee (1986). Monahan (1984) considers the economic implication of offering quantity discounts to the (single) buyer from the supplier's point of view. Assuming that the supplier follows a lot-for-lot policy, he shows that a sufficient discount can induce the buyer to order a quantity that increases the supplier's net profit. This modified quantity is related to the buyer's original economic order quantity by a factor that depends only upon the ratio of the fixed ordering costs of the two parties. Lee and Rosenblatt (1986) extend Monahan's model to include the supplier's lot sizing decision, by considering the inventory carrying and fixed costs incurred by the supplier. Instead of the all-units discount schedule featured in both these papers, Rosenblatt and Lee (1985) study a linear discount schedule in the same setting as Lee and Rosenblatt (1986). Interestingly, here the benefits do not all accrue to the supplier (as in Monahan 1984 and Lee and Rosenblatt 1986), but instead both parties can benefit. Bannerjee (1986) takes the perspective of a central decision maker who can jointly optimize the total of both parties' costs. He computes the joint economic lot-size, quantifies the resulting benefit to the buyer and the supplier, and determines the optimal quantity discount schedule. The preceding papers all assume deterministic and price-independent demand, and allow no shortages. The papers reviewed next relax some of these earlier assumptions. A seminal paper considering the role of pricing in channel coordination is Jeuland and Shugan (1983), which appears in the marketing literature. They primarily consider a two-member channel facing deterministic market demand, assumed to be a function of the retailer's selling price (p) and service (s), where the latter costs s. The production cost is $C(Q)$, a function of the manufacturer's product quality (Q). By comparing the optimality conditions with and without coordination, they show that absent coordination, the manufacturer has the motive to raise his price markup above first-best level. At the same time, he will set levels of quality and all other promotional variables that are below the first-best level. The same is true of the retailer. The authors then go on to show how a quantity discount scheme can coordinate the channel, and also share the efficiency gains. However, they acknowledge practical and legal barriers to implementation: in particular, determining the parameters of their scheme will require all costs to be common knowledge of both parties. Moorthy

(1987) offers other schemes which can coordinate the channel of Jeuland and Shugan (1983) and are easier to implement. In particular, he argues that a two-part tariff is superior to quantity discounting because it is simpler, it separates the coordination problem from the profit-sharing problem (coordination is achieved by setting the wholesale price equal to the manufacturer's variable cost, and the lump-sum transfer allows the profits to be split arbitrarily), and leads to fewer legal problems (charging all retailers the same wholesale cost is more consistent with the Robinson-Patman Act). Many subsequent papers in the marketing literature have considered variants of this model. Weng (1995) extends Jeuland and Shugan (1983) by specifically incorporating the mechanisms that determine the relationship between operating decisions (such as order quantities and selling prices) and profits (revenues less inventory and setup costs, calculated using the typical EOQ framework) for a buyer and a supplier, and deriving pricing policies between them that can coordinate the channel's activities. The buyer determines the size of the order placed with the supplier and the selling price charged to customers. Customer demand is a known, deterministic function of that selling price. The supplier, whose profits depend upon the buyer's decisions, controls the transfer price and his own production lot size. All data are assumed to be common knowledge. The focus of the analysis is on determining how to implement a mechanism to divide the additional profits generated through coordination. Under the assumption that the buyer will receive a fixed fraction of the incremental profit, the author shows that a quantity discount for the buyer along with a franchise fee paid to the supplier is sufficient to induce the buyer to make decisions that lead to joint profit maximization. Further, he shows that the form of the quantity discount scheme (all-units vs. incremental quantity discount) is not critical to achieving channel coordination. The dependence of customer demand on price and of operating costs on order quantity are the critical factors. Weng (1997a) also treats a newsvendor model in which the goal is to coordinate the decisions of the manufacturer and the retailer (which he calls the distributor) via pricing. The one-period demand is assumed to follow a phase-type distribution and includes a functional form that is an explicit function of the distributor's selling price. He compares the system performance when the manufacturer and retailer independently choose prices to that when production and ordering is jointly coordinated. He shows how coordination increases system profits and how the magnitude of the increase depends on the various system parameters. Weng (1997b) treats a problem similar to that considered in Weng (1997a), except with somewhat different assumptions regarding the costs and objectives. He assumes a quadratic manufacturer's cost function, and a service level constraint at the retailer level rather than a stockout cost. He derives the optimal system-wide policy and the transfer price between the manufacturer and retailer that achieves that policy. Weng (1997c) includes demand uncertainty in the context examined in Weng (1995). However, a number of simplifying assumptions have been made to facilitate the analysis. A single-period model (newsvendor type) is assumed, and both the transfer price and retail price are exogenous.

The manufacturer produces to order (with a fixed setup cost), and produces a second run at the end of the season to cover unmet demand (at a higher cost). Thus, the only decisions are the buyer's two order quantities. As in Weng (1995), the objective of the paper is to determine a pricing scheme to allocate the additional profits that result from coordination. Assuming an all-units discounting policy, the author shows that coordination increases the number of units produced by the manufacturer and held by the retailer. However, the increase in the system profit may not always be substantial. Corbett and de Groote (1997) consider a two-member channel in a deterministic EOQ-type of environment, where all information is common except the buyer's holding cost. The supplier assumes a probability distribution on this holding cost. The authors compare the supplier's optimal contracts under full information to those under asymmetric information, both with and without coordination. These include contracts with a fixed cost transfer, one-step quantity discount, and with full cost transfer. As expected, all these contracts are equivalent under full information. However, under asymmetric information, only full cost transfer to the buyer is system-optimal. It is also the buyer's preferred scheme, while the supplier is indifferent between it and a one-step quantity discount.

10.4.3 Minimum Purchase Commitments

Traditional inventory theory generally assumes that the buyer can order any quantity from the supplier at any time. Certainly the buyer prefers to avoid any constraints on his ability to meet his own customers' demand in an economical way. However, this may be undesirable from the supplier's point of view for a variety of reasons. For instance, suppose a buyer facing uncertain demand places orders according to an (s, S) policy. The variance of the resulting orders will exceed the demand variance since the buyer makes no order in many periods, waiting until cumulative demand over a period of time is sufficiently large. Amplification of demand variance is referred to as the "bullwhip effect" (see Lee et al 1997), and is undesirable upstream since, generally speaking, the supplier's costs increase with the order variance. This is due, for instance, to the increased need for inventory buffers and/or the more tentative scheduling of machine and labor capacity that may result. Even when the buyer's demand is certain, the supplier's production costs may be lower when the order is larger than what the buyer may consider to be optimal. One response to this conflict is an agreement in which the buyer agrees in advance to accept delivery of at least a certain quantity of stock, either in each individual order or cumulatively over some period of time. Depending on the relative strategic power of the parties, the seller may offer the buyer some forms of inducement, the most natural of which is a lower unit cost on items purchased under the contract. Incorporation of a lower bound on the buyer's purchase quantity is fairly straightforward in many settings, for instance when the problem can be specified in the EOQ framework (cf. Nahmias 1997) or as a single-period newsvendor model (cf. Porteus 1990). We refrain from reviewing such models. Additionally, while some quantity flexibility contracts contain components which resemble minimum purchase commitments, we present them separately since the emphasis in modeling those structures is on the revision of orders. Here we discuss some recent papers that consider the choice to install a minimum purchase agreement in multi-period settings with demand uncertainty. Anupindi and Akella (1993) cast the problem as a stochastic finite horizon model in which the buyer agrees to accept delivery of a fixed quantity of goods in each period. Discounts are given based on the level of this fixed commitment. The buyer can purchase extra units at a price premium. While the supplier guarantees the availability of the previously committed quantities, the additional units might not be delivered immediately. Thus, delivery of the additional units is expensive and uncertain. The authors prove that a modified S-type policy is optimal for this problem. Moinzadeh and Nahmias (1997) treat the same general problem (the minimum commitment per period, Q, is given), but with both fixed and proportional penalties for adjustments, and over an infinite rather than finite horizon. Delivery of the additional units is assured. The authors contend (but do not formally prove) that a type of (s, S) policy is optimal: if the inventory on-hand just prior to a delivery is less than s, adjust so that the total inventory after delivery is S. Assuming normal demand, they develop a diffusion approximation of the system, which allows efficient approximation of both s and S. They show that this type of contract results in order variance lower than under conventional S or (s, S) policies. Hence, the fixed delivery contract serves as a risk sharing mechanism. While the preceding two papers assume a constraint on every period's purchase, the agreement in Bassok and Anupindi (1997a) applies to the cumulative purchase (at least K_N units) over a given planning horizon of N periods. Demand is assumed to be a sequence of independent, identically distributed random variables. The authors prove inductively that the buyer's optimal policy in each period is a modified S policy. Letting (I_n, K_n) be the on-hand inventory and the minimum remaining purchase commitment when n periods are left, the optimal policy is defined by constants S_n and S^M $(S_n < S^M)$ in the following way:

if $I_n < (S_n - K_n)$, then order up to S_n;
if $(S_n - K_n) \leq I_n \leq \left(S^M - K_n\right)$, then order exactly K_n;
if $(S_M - K_n) \leq I_n \leq S_M$, then order up to S_M; and
if $S_M \leq I_n$, then do not order.

The impact of unit cost discounts on the buyer's total cost is described. The computations show that the advantages of the unit discount drop off quickly on approaching a particular quantity, which depends on the parameters of the problem. The buyer should never commit to more than this quantity. Cachon and Lariviere (1997a) consider the problem faced by a manufacturer who faces uncertain demand for a component that is purchased from an external supplier. The manufacturer must offer a contract to the supplier for production capacity that must be built before the uncertain demand is observed. Obviously, the manufacturer would like the supplier to build ample capacity, but wishes to avoid the cost of excess. Under the simplifying assumption that the demand is a Bernoulli random variable, the authors show that termination fee contracts are equivalent to minimum purchase contracts in this setting. Under the former, in addition to a unit cost, the manufacturer pays a cancellation fee per unit for all units not purchased if he takes delivery of fewer than the agreed-upon number of units. Under the latter, the manufacturer guarantees a minimum purchase amount and pays a penalty per unit if he takes delivery of fewer than the guaranteed number of units. Further, the appropriate contract terms depend upon both the information available to the parties as well as the enforceability of the contract. When the demand distribution is known to all, the manufacturer offers such contracts only if he is able to force compliance and verify the supplier's capacity choice. Since such contracts require the manufacturer to share some risk, they simply add to his costs in this case. The manufacturer should, therefore, simply offer a price-only contract. However, when the manufacturer is privately informed about the demand distribution, he may offer a cancellation fee or a minimum purchase contract to convey information credibly, even when he must rely on the supplier's voluntary compliance.

10.4.4 Quantity Flexibility

A Quantity Flexibility (QF) clause defines terms under which the quantity a buyer ultimately obtains may deviate from a previous planning estimate. The conditions can include limits on the range of allowable changes, pricing rules, or both. The motivation for each party to agree to such a clause depends on the nature of the alternative. The benefit to the buyer is clear if the clause ensures a degree of flexibility where previously there was none. This implies a willingness to pay a higher cost on all units and/or for each unit of deviation from the initial estimate. Indeed, some premium would be necessary to compensate the supplier for the increased exposure to demand risk. On the other hand, some buyers enjoy the strategic power to affect flexibility without the authority of a formal agreement. As noted earlier, when a buyer's estimate does not entail enforceable commitment, buyers commonly overstate their intended purchase, only to refuse undesired product later on. In this context, the clause is a way to encourage the buyer to forecast and plan more deliberately and honestly. In exchange, the supplier might need to provide a price break to give the buyer an incentive to participate. Either way, a QF clause has risk-sharing intent, and the hope is that the agreement can make both parties better off. Typical research questions to ask of QF settings include: (1) how should the buyer behave (i.e. forecast and purchase) given the available flexibility, (2) how should the supplier behave given the flexibility promised the buyer, (3) what would be the cost or benefit to each party of changes to parameters of the agreement, e.g. a flexibility bound or pricing term. Attempts to address these questions rigorously for the various types of QF clauses share a number of modeling challenges. Because the exercise of flexibility implies reconsideration of a prior decision, even the simplest model requires at least two decisions on the part of the buyer for each purchase: there is an initial inventory decision, and then a revision conditional on whatever new information about demand should become available. The first decision must properly anticipate the second for all contingencies, since the two are linked by the terms of the QF clause. Dimensionality of the decision space is an obstacle to analytical solution, as is the machinery required to represent updating of a stochastic demand. For these reasons, the available models tend to be very stylized, e.g. a selling season represented by a single random variable with a distribution amenable to Bayesian updating. Furthermore, most models focus only on the buyer's perspective, since the supplier's problem is even more complex. As is well known from multi-echelon inventory theory, fairly simple multi-level systems using relatively straightforward operating policies often yield demand processes at upper levels which are analytically quite difficult to characterize (cf. Schwarz 1981). Here, the informational dynamics of market demand are already an issue, and the buyer's ability to revise purchase quantities only adds complexity. Below we describe efforts to model contracts containing variants of the QF structure. "Backup agreements," which have been observed in the fashion apparel industry, are the focus of Eppen and Iyer (1997). Such agreements are parametrized by (ρ, c). Prior to the selling season, the buyer (e.g. a mail-order retailer) commits to y units for the season, and takes immediate ownership of $(1-\rho)\,y$ at unit price c. After observing the first two weeks of sales data (approximately 10% of total sales), which is used to perform a Bayesian update on the prior distribution of the total season's demand (modeled as a negative binomial random variable), the buyer can order up to the remaining ρy for the original price and receive quick delivery. There is a penalty cost of b for any of the backup units not purchased. The analysis suggests that the buyer's optimal strategy has "order-up-to" structure, meaning that each of the two decisions has some target threshold. The second is obtained as a critical fractile of the distribution of remaining demand conditional on the early sales, while the first does not have a simple solution. Some comparative statics properties of these thresholds are provided. While the bulk of the analysis focuses on the buyer's side, this paper reports that for certain parameter combinations, the backup agreement contract can improve profits for both parties relative to a setting with no backup agreement. In Anupindi and Bassok (1995), the supply contract commits the buyer to purchase at least a given total quantity of a single product over a finite time horizon. A certain additional volume is also available at the same price, beyond which a higher price is charged. This paper derives the purchase policy that minimizes the buyer's total costs (purchase, holding, and backorder) under stationary and random demand, which is shown to be of modified order-up-to structure. The critical values are related to those obtained by solving a standard finite horizon model with no commitments and a single price. Bassok and Anupindi (1995) consider forecasting and purchasing behavior in an arrangement in which the buyer initially forecasts its period-by-period purchases over a T-period horizon, then may revise each period's purchase one time within specified percentage bounds. Demands are non-stationary but independent. Because information about that demand never changes, the adjustment in period n is a response to the demand realizations and purchases made in periods 1 through $(n-1)$. With linear holding and shortage costs, and complete backordering from month to month, their objective is to solve for: (i) the purchase estimates prior to period 1, and (ii) each period's actual purchase. The authors discovered this problem to be very complex, and ultimately proposed a heuristic policy, constructed numerically. No special structure to this policy was reported. Milner and Rosenblatt (1997) analyze a setting in which the buyer places orders for two periods, and may then adjust the second order after observing demand in the first period. This differs from the contract in Bassok and Anupindi (1995) in that there is per-unit penalty for any adjustments. They describe the optimal behavior of the buyer, both in the initial orders and the subsequent adjustment. As in Bassok and Anupindi (1995), beliefs about the second period's demand do not change between the time of the initial order and the adjustment, so what drives the use of flexibility in period 2 is purely the discrepancy between the purchase and the demand outcome in period 1. The optimal adjustment is characterized by a range $[L, U]$, whose endpoints are simple functions of the cost parameters and the demand distributions. If the pre-adjustment inventory position on entering the second period falls in this interval, no adjustment should be made; otherwise, adjust to get to the closest boundary of the interval. Closed forms are not available for the optimal initial orders, but some structural properties and comparative statics results are presented, all of which are consistent with intuition. Finally, parameter combinations are derived which characterize the buyer's preference for either the flexible contract, a non-flexible contract, or no contract at all (assumed to allow adjustment without financial consequence). The supplier's preferences are not considered. Barnes-Schuster et al (1998) examine the use of options for supplier capacity as a means of affecting flexibility for the buyer. Here, a selling season is divided into two periods of correlated demands. Excess demand is backordered after period 1, and lost after period 2. Prior to the first period, the buyer places firm orders for both periods, and also purchases options which reserve additional supply in period 2. After observing the period 1 demand, the buyer has the prerogative to exercise some number of the options (at an additional fee), or let them expire, thereby losing the reservation fee but avoiding any additional cost that would have been incurred had actual production been commissioned initially. The supplier is obligated to position raw materials to the maximum buyer request (firm orders + options), but may convert these to finished product at two different points in time at different costs. A cheaper mode of production is available only immediately after the buyer's initial orders, while more costly production, if necessary, occurs after the buyer has exercised any options. The supplier dictates the price terms of the contract (wholesale price, option price, exercise price). Examining the efficiency of the contract form, this paper concludes that linear prices cannot coordinate the channel in a way that offers the supplier positive profits, and then proposes various quantity discount schemes that achieve this goal. The cost ramifications for both parties under linear pricing are characterized numerically. Finally, the decision of when to update demand is analyzed numerically by dividing the time horizon into several subintervals. In contrast to the more common assumption in which the procurement price of a required component is either a decision variable or an exogenous parameter, in Li and Kouvelis (1997) what motivates a buyer's desire for flexibility is uncertainty in that price (modeled as a geometric Brownian motion with drift). The main thrust of this paper is to evaluate the value of flexibility in quantity and time, and determine when each form might be the more desirable contract structure conditional on the cost parameters, the design of the risk-sharing price arrangement, and the dynamics of the anticipated price fluctuations. In the model setting, the buyer must procure exactly D units by time T under pricing terms specified in the supply contract at time 0, and is interested in selecting a favorable flexibility arrangement to minimize his own expected purchasing and inventory costs. Possible arrangements that are considered in various combinations are: (1) a "time-inflexible contract", in which the buyer must state up front the purchase times, (2) a "time-flexible contract", in which the buyer may observe price movements and decide dynamically when to buy, and (3) a "quantity flexible contract", in which the buyer chooses a Q at time 0 that entails commitment to purchase an amount in the window $[(1-\alpha)\,Q, Q]$ for a given α. The analysis determines the timing and quantity of purchases that minimize the buyer's expected net present value of purchase and holding costs. Given a single supplier that offers no quantity flexibility and sells at exactly the market price, the buyer's optimal strategy is shown to be invariant to the existence of any time flexibility. The buyer will purchase the full amount either at time 0 or time T, the choice of which depends on the holding cost, the discount rate, the drift parameter of the price process, and the value of T. However, if the contract price is distinguished from the market price, for example through an agreement that transfers some of the fluctuation risk to the supplier, then the optimal strategy must be obtained computationally. This is done via a discretization of the geometric Brownian motion into a random-walk type of process. Quantity flexibility is analyzed in the two-supplier case, where the suppliers are assumed to exist in different markets with separate but correlated prices. Tsay (1996) models the incentives of both the supplier and the buyer in a setting in which the buyer first estimates a purchase quantity for a given selling season, the supplier then commits to production, and finally the buyer makes his actual purchase (which may differ from the estimate) in light of updated information about a stochastic market demand. System inefficiency can result in such an environment because, as noted earlier, the buyer might be expected to inflate the initial estimate. The root cause of this behavior can be linked to the phenomenon of double marginalization. This paper considers a QF contract which couples the buyer's commitment to purchase no less than a certain percentage below the initial estimate with the supplier's guarantee to deliver up to a certain percentage above. There is no penalty for making adjustments within the defined range. In conjunction with an appropriately chosen unit price, this contract structure is shown to be able to allocate the costs of market demand uncertainty so as to coordinate the individually motivated supplier and buyer to the system-wide optimal outcome. The structural properties of efficient QF contracts are characterized, as is the mechanism for allocating efficiency gains. Bassok and Anupindi (1997b) analyze the buyer's side of the percentage flexibility contract studied in Tsay (1996), generalizing to an ongoing supply relationship in which planning for multiple future periods is performed in a rolling-horizon fashion. At each period, the buyer makes a purchase, and also provides estimates of purchases to be made in subsequent periods. The contract defines percentage limits on how these estimates may be revised from one planning iteration to the next. Demand is assumed to be independent and stationary, with known distributions. Shortages are completely backordered, and holding and backordering costs are linear. Because the decision space becomes substantially more complex, a heuristic forecast and purchase algorithm is proposed and defended by simulation analysis. Simulation experiments are also used to characterize the value to the buyer of a certain amount of flexibility and the variability of buyer purchases as a function of flexibility. Tsay and Lovejoy (1998) consider the rolling-horizon QF contract in a multi-echelon setting, allowing for non-stationary demand with information updating. At period t, the buyer states to the supplier the vector

$$\{r(t)\} = [r_0(t), r_1(t), r_2(t), ...],$$

where

$r_0(t)$ = actual purchase in period t
$r_j(t)$ = estimate of purchase to be made in period $(t + j)$, for each $j \geq 1$.

The QF contract between the buyer and the supplier is parametrized by (α, ω), where $\alpha = [\alpha_1, \alpha_2, ...]$ and $\omega = [\omega_1, \omega_2, ...]$. This describes how much flexibility the buyer enjoys in revising $\{r(t)\}$ going forward in time. Specifically, for each t and $j \geq 1$:

$$[1 - \omega_j]\, r_j(t) \leq r_{j-1}(t+1) \leq [1 + \alpha_j]\, r_j(t).$$

$\{r(t)\}$ is the only information available to the supplier concerning the buyer's purchases, and any revisions are allowed provided that they observe the stated bounds. The analysis provides heuristics based on Open-Loop-Feedback-Control logic indicating how the buyer should construct $\{r(t)\}$ in light of the statistics of market demand and the flexibility parameters, as well as how the supplier should behave (i.e. submit orders and forecasts to its own upstream supplier, with whom there may be a separate QF contract) in order to fulfill its contractual commitment to support the buyer's order sequence. Attention is then given to how the flexibility characteristics of the system impact the inventory and service patterns, as well as how order variability propagates along a multi-level supply chain. Simulation experiments suggest that QF contracts can dampen the transmission of order variability throughout the chain, thus potentially retarding the well-known "bullwhip effect". This paper also addresses the issue of contract specification. A nonlinear programming formulation is provided for answering the channel coordination question: if the supplier and buyer are both independently managed units of the same firm, how might a central planner specify the internal contract to achieve good joint performance? Also, an investigation of the value of flexibility to a buyer indicates that flexibility increases in value as the market environment becomes more volatile, and that this value observes a principle of diminishing returns. Anupindi and Bassok (1997) is the only paper among those reviewed here which analyzes flexibility contracts in a multi-product context. The contract requires the buyer to commit to a minimum cumulative dollar value of purchases during a specified time horizon to be eligible for a percentage discount off regular prices. The discount is available only for purchase volumes up to a certain fraction above the minimum commitment, with the regular price charged for any additional purchases. Again, because of problem complexity, a heuristic algorithm is used by the buyer. Numerical studies indicate that a policy which commits to a total dollar value that is the sum of the mean dollar volumes for the individual products performs relatively well. Moreover, the flexibility to increase purchases at the lower price is not particularly critical. This is because risk pooling across demands in different periods for various products already provides sufficient flexibility.

10.4.5 Buyback or Returns Policies

A buyback or return clause establishes who bears responsibility for unsold inventory, and to what extent. One can make an analogy between buyback clauses and QF clauses, in that both structures lay out ground rules for the buyer to end up with an amount that is potentially less than his prior estimate. A subtle difference is that buybacks generally take place after demand has been observed, whereas QF-style order reductions may be executed while demand uncertainty remains. Nevertheless, many of the modeling issues discussed in the previous section also apply here. Analytical treatment of buybacks first appeared in the marketing literature, in Pasternack (1985). In this paper a manufacturer produces a commodity for sale to a retailer. The item has a relatively short shelf or demand life, and the retailer places only one order with the manufacturer. The manufacturer sets the wholesale price and the market selling price is fixed, so the only decision for the retailer is the order quantity. Using a single-period, newsvendor-style model, the author characterizes potential inefficiencies within this channel that are essentially due to double marginalization, and finds that neither a policy of allowing for unlimited returns at full credit nor one which allows for no returns is efficient. He determines that coordination of the channel can be achieved by a buyback clause which allows full return at a partial refund, and that the efficient prices (wholesale and buyback) can be set in a way that guarantees Pareto improvement. A key result is that the channel-coordinating prices are *independent of the market demand distribution*. This is significant in that the manufacturer need not know the market demand distribution in order to implement an efficient contract, although this remains necessary in order to properly value and allocate efficiency gains in a way that will insure the retailer's participation. In the economics literature, Kandel (1996) extends Pasternack (1985) by modeling price-sensitivity in end customer demand, with a general downward-sloping expression for the expected demand. Hence, the retailer must also set the retail price. The author confirms that Pasternack's results proving the inefficiency of certain decision structures generalize to this setting, and proposes two arrangements under which coordination can be achieved. The first assumes that the manufacturer can additionally impose a resale price maintenance contract on the retailer, in which case a consignment agreement (full returns for full credit) is efficient. This allows the manufacturer to impose the channel-optimal retail price, and leads the retailer to choose the channel-optimal quantity. Here, the retailer makes zero profit and the manufacturer retains all profits of the coordinated channel. Alternatively, the manufacturer can charge a wholesale price of exactly the marginal production cost, rendering the retailer's decision problem identical to that of the integrated channel. The profits of a coordinated channel accrue entirely to the retailer in this case. The author concludes with conjectures regarding the impacts of risk-aversion of either party, consideration of additional variables such as the manufacturer's choice of product quality and the retailer's promotions effort, information asymmetry, the elasticity of market demand, and the stability of market demand over time. Emmons and Gilbert (1998) generalize Pasternack (1985) in much the same way as does Kandel (1996), however assuming a specific multiplicative form of demand model. One result that is enabled by the functional assumption is that for a given transfer price, the offer to buy back excess stock tends to increase the total profits of the channel. Donohue (1996) studies buyback contracts within two different two-stage production environments. In the first, the supplier offers the product for delivery at two different lead times. The buyer commits in advance to a quantity of the long-lead time item at a given wholesale price. After revising his demand forecast to reflect information gathered prior to the season, the buyer can place an additional order for short-lead time delivery at a different wholesale price. At the end of the season, the manufacturer takes back any unsold items at a third price. The author first finds that the buyer's optimal ordering policy has order-up-to structure, then determines the three price parameters that will result in the same system profit as the optimal centralized solution. Similar analysis is performed for the second setting, an assemble-to-order context in which some critical component must be prepositioned. Whereas coordination in the first environment entails a sort of minimum purchase commitment in the initial purchase, in the second environment there is an option-like arrangement which communicates a maximum purchase commitment. As in Pasternack (1985), the prices that coordinate the channel turn out to be independent of the distribution of market demand. There is also some discussion of how the efficiency gains might be split between the buyer and manufacturer, and the implications of the contract structure for the variability in each party's profits.

10.4.6 Allocation Rules

Allocation issues arise when multiple retailers compete for a product that is rendered scarce by some limit on the manufacturer's production capacity or stock availability. Such concerns are not considered in most of the papers in this review since they model only supply chains with a single retailer and a manufacturer with unlimited capacity. However, as suggested in Lee et al (1997), the possibility of rationing by the manufacturer can induce competition between retailers and, therefore, lead to strategic behavior. In particular, retailers will tend to inflate their orders, which distorts the flow of information. While insightful, Lee et al (1997) does not model the effect of alternative allocation policies. We are aware of only two papers on supply contracts that consider the design of allocation policies, as described below. Cachon and Lariviere (1996) model a single-period, single-supplier, multi-retailer supply chain where the supplier's production capacity is limited and each retailer's stocking level is private information. The wholesale price is fixed. In case of shortage, the supplier's capacity is allocated using some fairly general allocation scheme, which is required only to be (i) "efficient", meaning that the available capacity is never wasted, (ii) "insured", i.e. if a retailer desires a positive quantity of stock, he will receive at least some stock, and (iii) "individually responsive", i.e. if a retailer wants more stock, he gets more as long as capacity is available. The main result is that Pareto optimal allocation mechanisms are easily manipulated by the retailers. At the same time, truth inducing mechanisms lead to system inefficiency. The authors also show that some mechanisms may lead the supplier to choose a higher capacity. Cachon and Lariviere (1997b) consider a one-supplier, two-retailer supply chain in a two-period context. The supplier's production capacity and the wholesale price to the retailer are fixed, but retailers can set their own price. Demand in the second period can be high or low, and is a simple linear function of price. Common information is assumed. In case of shortage, capacity is allocated using a publicly known allocation scheme, of which they consider two versions: "even allocation" and "turn-and-earn". "Even allocation" means that the available capacity is divided evenly among all retailers placing orders. In "turn-and-earn", allocation is a function of past sales (for example, a retailer with higher past sales may receive a more favorable allocation). The main result of their analysis is that under turn-and-earn, the supplier's profits will increase but the retailers' profits may not. Since the retailers are identical, they end up just selling greater volume at a lower price to protect their allocation. While both papers illustrate the effect of allocation on supply chain behavior and performance, they offer little by way of specifying what allocation policies might be optimal. Indeed, this is a difficult problem even when system control is centralized, as has become evident from extensive studies in a variety of contexts. Some examples include single-warehouse/multi-retailer settings (e.g. Eppen 1979, Eppen and Schrage 1981, Jonsson and Silver 1986, Jackson 1988), manufacturing systems with component commonality (e.g. Baker 1985, Gerchak and Henig 1989, Agrawal and Cohen 1997), inventory systems with different customer classes (e.g. Nahmias and Demmy 1981, Cohen et al 1988, Ha 1997b), and queuing systems with multiple servers or customer classes (e.g. Gilbert and Weng 1997). To understand the source of the difficulty, consider an assembly system with component commonality. The optimal allocation policy must depend on the inventory levels of all items in the system, which is a function of the joint distribution of demand for all end items. This makes the problem analytically and computationally complex. Also, in practice, when capacity shortages occur, suppliers may vary pricing to mitigate the effects of the problem and buyers may turn to alternate sources of supply. The models described in this section do not allow such possibilities.

10.4.7 Lead Times

In classical inventory models, delivery lead time is either fixed (at zero in some cases), or a realization of a random variable. Additionally, some researchers have allowed the buyer a choice from multiple (but fixed) lead times that result from multi-sourcing or the option to expedite material from a single supplier, two practices frequently observed in industry. For example, Hausman et al (1994) build on the efforts of Daniel (1963) and Fukuda (1964) and others in analyzing a buyer's optimal purchasing behavior given an exogenous menu of various lead time and unit cost combinations. Lawson and Porteus (1996) perform similar analysis on a multi-echelon system under central control. Agrawal et al (1998) model the contracting of capacity with suppliers which differ in commitment dates as well as other structural features, from the perspective of a retailer of multiple products. We do not present these works in detail since they take the characteristics of the supply base as given, and therefore fall beyond the scope of what we define to be "contracts" papers. Below we discuss papers that consider issues that arise when the terms of timing in delivery are control variables. Barnes-Schuster et al (1997) consider a system composed of a supplier and one or more buyers in which the supplier faces a known production lead time (l_p), while the lead time for delivery to the buyer (l_d) is a decision variable. The buyer faces a stationary periodic review problem for which a static base-stock inventory policy is known to be optimal, and the appropriate base-stock level is an increasing function of l_d. The supplier faces an analogous problem, except that increasing l_d reduces its required safety stock. This paper shows that in the single-buyer case the l_d that is optimal from the system's point of view is either zero, which has the supplier hold all the system safety stock, or equal to $(l_p + 1)$, in which case the supplier produces to order and hence holds no safety stock. Conditions on the cost and demand parameters are provided that determine which lead time to use. In the case of multiple buyers with identical cost parameters, the supplier should hold the safety stock for buyers with "sufficiently low" standard deviations of demand, while the remaining buyers hold their own. Iyer and Bergen (1997) model Quick Response (QR), a movement in the fashion apparel industry, in a manufacturer-retailer supply chain. Achieved by any of a number of process improvements, QR is simplified to mean lead time reduction, i.e. a delay of the point at which the supply chain makes its quantity commitment. The retailer benefits since its orders may then be placed under an improved state of information (as modeled within a Bayesian updating framework), yet the manufacturer can actually be made worse off. Specifically, since this manufacturer is assumed to operate in pure make-to-order mode, its payoff is determined once the retailer orders, regardless of how the uncertain market demand eventually resolves. Unaffected by any overage risk, the manufacturer naturally prefers high retailer orders, even if this includes large amounts of safety stock that never get sold. Thus the manufacturer may resist efforts that will reduce lead-times precisely because improved forecasts enable the retailer to reduce its safety stock position. According to the authors, this may explain the various side-agreements which have been observed to accompany QR efforts, such as commitments to higher service levels to the end customer, higher wholesale prices, or volume commitments across multiple products. These mechanisms all work by forcing the retailer to buy and/or pay more than it would with QR alone, enough so that the manufacturer's original profit is preserved. The key conclusion is that Pareto-improving contractual combinations do exist whereby the implementation of lead time reduction can proceed with the blessings of both parties. Grout and Christy (1993) discuss the incentives faced by a supplier in quoting a delivery time to a buyer, and the implications for the likelihood of on-time delivery. During the contracting process associated with the one-time purchase of some item, the buyer offers a lump-sum bonus of B for on-time delivery, and the supplier in turn specifies a delivery time A. The supplier's risk is due to uncertainty in the production time. If the supplier completes production prior to time A, he collects B but incurs a cost of α per unit of time for holding the item until delivery. If instead a production delay causes tardiness, the supplier incurs a cost of β per unit of time late. Given this structure, the supplier's optimal A turns out to be a critical fractile of the distribution of the random production time, reflecting the relative values of α, β and B. Anticipating this, the buyer specifies B so as to minimize his own expected shortage cost (incurred at a rate of δ per unit of time the delivery is late) and expected bonus payment. The analysis focuses on the role of the bonus by comparing the case of $B = 0$ to a contract in which the buyer sets the bonus with channel coordination in mind. Indeed, $B = 0$ leads to inefficient performance since the supplier's decision does not take the buyer's shortage cost into account. A B which achieves the first-best outcome is shown to exist, which leads to a recommendation regarding the make-or-buy decision: compare the expected bonus paid to an independent supplier under the first-best B against the cost of vertical integration, and choose the cheaper option. Moinzadeh and Ingene (1993) call attention to the supplier's perspective on the design of what amounts to a dual lead-time supply arrangement. The supplier carries two different products which are partially substitutable. Good 1 is held in inventory, hence offers immediate delivery if in stock. Good 2 is available by special order only, imposing a one-period delay on customers. Each product is the first choice of some segment of the population, as characterized by two stationary Poisson processes. (While the discussion is placed in a context of a population of individual consumers, the overall demand pattern could potentially describe the purchase preferences of a single downstream organization.) The net demand for each item is computed by examining the customer's 3 possible reactions to unavailability of the stocked good: (i) "walk", meaning that the sale is lost altogether, (ii) "wait", which imposes a backorder cost, or (iii) "switch" by special-ordering good 2 instead. The rate of occurrence of each of these is known, with an interpretation based on consumer utility. The analysis considers the supplier's problem in maximizing its long-run profit (discounted over an infinite horizon), assuming that good 1 has fixed price and is managed according to a base-stock inventory policy. The decisions are the base-stock level of good 1 (R) and the markup on good 2 (m_2). For a given m_2, the optimal R is shown to have newsvendor-style structure, which balances the margin on good 1 against the marginal cost of being out of stock on good 1. The latter includes lost sales and backorders as in a traditional model, but also the possible benefit that accrues when customers switch to good 2. Naturally, the relative magnitudes of these factors depend on m_2. In fact, if a sufficient fraction of good 1 customers show a tendency to either switch or wait (which can be induced by the choice of m_2), the supplier may prefer not to stock good 1 at all. More generally, the profit maximizing strategy can involve setting a price that encourages switching in order to reduce holding costs for the zero lead-time item. A numerical method for setting m_2 under additional structural assumptions is presented, as is some discussion of how this parameter might be used to compensate customers for the dissatisfaction that results from their not obtaining their first-choice item.

10.4.8 Quality

The papers discussed thus far are concerned primarily with the timing and quantity of material flows and the associated financial transfers. However, any supply relationship is premised on the quality of the delivered product. This may be formalized by conditions of the contract. The economics literature has an extensive history of modeling supply settings in which product quality is a management choice. The representation of quality is relatively abstract, treating it as a product attribute which has a positive effect on both sales volume and production cost. This is typically encoded in a deterministic demand curve that is downward-sloping in price and shifts upward with quality, in conjunction with a production cost function that increases in both volume and quality. Attention then turns to characterizing the decisions that are optimal for a given market structure, and commenting on the consequences for profits and social welfare. Early examples of this approach include Spence (1975) and Dixit (1979). Similar efforts appeared later in the marketing literature, one example being the treatment of non-price variables by Jeuland and Shugan (1983). Because of the generality of this structure, virtually identical models have also been observed with "service" or "advertising" taking the place of the quality parameter in the formulation. A great deal of insight into supplier and buyer behavior results from relaxing the assumption that product quality is common knowledge to all parties. In many models in economics, the buyer must attempt to infer in advance the true quality of the product. This may be based on signals conveyed by other terms of the contract, such as the selling price or the supplier's willingness to offer a warranty. Or, if transactions recur over time, the buyer may rely on previous experiences with that supplier's products. In turn, the supplier may incur the cost of initially offering high quality in order to establish a reputation that will lead to repeat business. See Chapter 2 of Tirole (1988) for a textbook treatment of this body of literature, which relies heavily on game theoretic constructs. Quality has also long been of concern to the inventory management community, and, by extension, the SCM community. The models in this area examine the production process in more detail and therefore tend to have a much more concrete notion of quality, operationalizing it primarily in one of two ways: (1) as a probability that a particular item is defective or non-conforming, or (2) as a yield rate (either deterministic or stochastic). Many inventory models treat the process quality capability as an exogenous variable and then determine the appropriate lot-sizing behavior (see Yano and Lee 1995 for an extensive review). Others consider the choice of the quality level, but more from the vantage point of a single organization contemplating how to design its internal practices in light of its own costs of quality. For example, Porteus (1986) modifies a basic EOQ setting to consider the manufacturer's option to invest in the improvement of process quality (defined as a probability of going "out of control") as a way to manage the cost of reworking defective lots. Starbird (1994) examines how a risk-averse supplier's choice of quality level depends on the acceptance sampling method used by the buyer, and Starbird (1997a) performs similar analysis for an expected-cost-minimizing supplier with EOQ-style setup and holding cost concerns. Models which explore the negotiation for quality are much less common in the SCM literature, and many of the phenomena described earlier (e.g. signaling behavior) are usually beyond the scope of the analysis. This is because, as noted, the relative emphasis on operational-level details obstructs the in-depth consideration of issues such as information asymmetry. Below we describe some efforts to examine the motivations that determine the quality terms in supply chain relationships. Reyniers and Tapiero (1995) use a simple game-theoretic formulation of a supplier-producer channel to examine the impact of contract structure on the supplier's quality and the producer's inspection practices, and the implications for the quality of the end product. The supplier chooses one of two production methods, indexed by i, which differ in output quality (modeled as the probability of defect, p_i); the production cost T_i is higher for the process with higher quality. The producer may choose to perfectly inspect the incoming item at cost of m, or incorporate it directly into the end product. The supply contract stipulates that if the inspection reveals a defect, the supplier pays the producer $(C + \Delta\pi)$, where C covers the cost of repair and $\Delta\pi$ represents any additional rebate. If a defective input reaches the end customer, a failure cost of R is incurred with certainty. A fraction α of this is paid by the supplier, the rest by the producer. All parties are risk-neutral, all parameters are common knowledge, and the game is played only once. This results in a simple 2x2 matrix representation of the payoffs, which can be analyzed for Nash equilibria in the standard way. The game has an equilibrium in mixed strategies, whose structure depends on the value of $\Delta T/\Delta p$ $(= (T_1 - T_2)/(p_1 - p_2))$ and the other cost parameters in the following intuitive ways: (i) the probability that the producer uses inspection is increasing in $\Delta T/\Delta p$, (ii) the probability that the supplier provides low quality is increasing in m, and (iii) the quality of the end product is decreasing in α and m, and increasing in $\Delta T/\Delta p$. The "value of cooperation", which is the value of moving from independent to joint decision-making given the contract structure, is shown to be decreasing in $\Delta\pi$ for both parties, and decreasing in α for the producer. Tagaras and Lee (1996) focus on a manufacturer who has the option to increase the quality of an input material by paying a higher unit procurement cost. (The menu of cost-quality combinations may come from multiple differentiated vendors, or could also be different offerings from a single vendor.) The input item is defective with probability p, and the manufacturer's production process fails with probability q. Costs are assigned to the various root causes which might lead to a defective final product: r_1 if the input is defective, r_2 if the manufacturing process is at fault, and r_{12} if both apply. The manufacturer's expected cost per unit processed is then $\phi(p) = p(1-q)r_1 + q(1-p)r_2 + pqr_{12}$. This is compared against $C(p)$, the unit cost charged by the vendor for a defect rate of p ($C(p) = c(1-p)$ and $C(p) = c(1-p)^2$ are considered). Incoming inspection is also a possibility (at a unit cost of a, with any necessary rework costing r_i). The analysis reveals that the manufacturer's proper choice of vendor quality depends not only on the vendor's price, but also on the capability of the process using the item as an input (the structure of the dependence reflects the relative magnitudes of r_1, r_2, r_{12}). Contrary to a view that is popular in the modern quality movement, under some circumstances the manufacturer is better off buying lower quality inputs at a lower cost, because the value of high quality inputs is negated by internal process problems. Starbird (1997b) examines supplier buyer behavior in a model which features a careful accounting of the quality-related costs categorized by Joseph Juran (cf. Juran and Gryna 1988). Prevention and internal failure (e.g. scrap, rework) costs are incurred by the supplier, and appraisal (inspection) and external failure costs (warranty, replacement costs, etc.) are paid by the buyer. The buyer, who faces a deterministic market demand, procures from the supplier in lot sizes (L) that are economically chosen to minimize the expected cost of ordering, holding, purchasing, inspection, and external failure. The supplier minimizes the sum of expected setup, holding, manufacturing, prevention, and scrap costs by choosing a production lot size (Q) and a quality level (ϕ, the probability that an individual item is defective). The probability that a procurement lot will be accepted by the buyer is determined by the acceptance sampling rule, and has the binomial form $P_A(\phi) = \sum_{d=0}^{c} \binom{n}{d} (1-\phi)^d \phi^{n-d}$ where n is the sample size and c is the acceptance number ($c \leq 1$ simplifies the analysis). The author characterizes the resulting non-cooperative Nash equilibrium outcome, which demonstrates that the independence of the parties may lead the supplier to choose a quality level that is either higher or lower than the quality level that would arise under cooperation.

10.5 CONCLUSION

In this paper we have reviewed some recent efforts to study how the design of contracts affects supply chain behavior and performance. We believe this to be a very important and challenging field of research. The scope of issues that this literature has addressed thus far has been limited, as the variety and complexity of contracts used in practice do not lend themselves easily to mathematical modeling. This is one reason why so many of the papers reviewed here are done in the single-period "newsvendor" setting. Our review suggests several opportunities for future research. Following the evolution of inventory theory, these analyses could naturally benefit from extension to production systems of greater structural complexity. This could include consideration of multiple planning periods, a larger number of products, or multi-layered and branching supply "networks" in which each party might have contracts with several others. Other issues worthy of further attention derive from the multi-party aspect of real systems, a few of which are described below. Despite recent advances in information technology and trends towards sharing information with supply chain partners (cf. Kumar 1996, Verity 1996, Lee et al 1997), information asymmetry remains a key feature of real supply relationships. However, virtually all multi-player models in this review rely at some level on common knowledge of all parameters. One of the difficulties of including information asymmetry is that the analysis must then consider the multiple points of view. Since the probability structure of an uncertain event may be perceived differently by the parties to the contract, there may be disagreement in the calculation of expected profits. Hence, the notions of efficiency and optimality are not clearly defined. Addressing this issue might require a substantially more complex informational structure. This can include, for example, conditions regarding the buyer's beliefs about the supplier's beliefs, and vice versa. A related concern is the assumption of risk neutrality, as the notion of efficiency again becomes unclear once the various players are allowed to have different objective functions. Further, decisions made by individuals are often motivated by a variety of complex incentive and compensation schemes instituted by their own firms. Not only are such schemes difficult to formulate analytically in the kind of models described here, they are often unknown to the other contracting parties. Nevertheless, contracts may be a way to reconcile the differing preferences the parties may have towards the uncertainty in the outcomes. Another deficiency in the current literature is the lack of attention to competition, either between multiple buyers or multiple suppliers. Buyers that share a common supplier and compete in the same consumer market might behave in a way that obstructs their competitors' access to suppliers. In turn, the supplier might consider playing the buyers off one another to obtain price or purchase commitments. Multiple suppliers to a common buyer might need to alter their price, service, lead time, or flexibility offerings in light of the competitive environment. The kinds of contracts discussed here could play a role in structuring such relationships so as to improve efficiency and/or reallocate the risks. While the stylized models we have reviewed offer many insights into supply chain behavior, they fail to address a variety of issues that become relevant to actual implementation. For example, as mentioned earlier, these models ignore many of the legal, public policy and social issues associated with contracts. Further, contractual arrangements can substantially affect the roles of particular individuals within any organization. If the affected parties feel threatened, or if their incentive systems are not appropriately modified, implementation of such contracts could be far from successful. Another practical concern that is unattended to by the existing SCM contract literature is that of prescribing how the benefits from coordination ought to be divided among the parties, a decision which might require extensive bargaining and negotiations. Clearly, there are opportunities to integrate the existing literature with the substantial body of knowledge from the field of Game Theory. We believe that these issues offer a rich set of possibilities for future research on contracts in supply chains, and look forward to the interesting work in this area that the current literature will spawn.

Acknowledgments: The authors would like to thank the following individuals for helpful comments: Karen Donohue, Steve Gilbert, Warren Hausman, Ananth Iyer, Marty Lariviere, Chung-Lun Li, Kamran Moinzadeh, Dave Pyke, Sridhar Tayur, and Jin Whang.

References

Agrawal, N. and M.A. Cohen, "Optimal Material Control in an Assembly Environment with Component Commonality," Working Paper, Leavey School of Business, Santa Clara University, 1997.

Agrawal, N., S.A. Smith, and A.A. Tsay, "Multi-Vendor Sourcing in a Retail Supply Chain," Working Paper, Leavey School of Business, Santa Clara University, 1998.

Agrawal, N. and A.A. Tsay, "The Impact of Intrafirm Performance Incentives on Supply Chain Contracts," Working Paper, Leavey School of Business, Santa Clara University, 1998.

Anupindi, R. and R. Akella, "An Inventory Model with Commitments," Working Paper, Northwestern University, 1993.

Anupindi, R. and Y. Bassok, "Analysis of Supply Contracts with Total Minimum Commitment and Flexibility," *Proceedings of the 2nd International Symposium in Logistics*, University of Nottingham, U.K., 1995.

Anupindi, R. and Y. Bassok, "Approximations for Multiproduct Contracts with Stochastic Demands and Business Volume Discounts: Single Supplier Case," Working Paper, Northwestern University, 1997. To appear in *IIE Transactions*.

Baker, K.R., "Safety Stocks and Component Commonality," *Journal of Operations Management*, **6**, 1 (1985), 13-22.

Bannerjee, A., "A Joint Economic-Lot-Size Model for Purchaser and Vendor," *Decision Science*, **17** (1986), 292-311.

Barnes-Schuster, D., Y. Bassok, and R. Anupindi, "The Effect of Delivery Lead Time in a Two-Stage Decentralized System," Working Paper, Graduate School of Business, University of Chicago, 1997.

Barnes-Schuster, D., Y. Bassok, and R. Anupindi, "Supply Contracts with Options: Flexibility, Information and Coordination," Working Paper, Graduate School of Business, University of Chicago, 1998.

Bassok, Y. and R. Anupindi, "Analysis of Supply Contracts with Forecasts and Flexibility," Working Paper, Northwestern University, 1995.

Bassok, Y. and R. Anupindi, "Analysis of Supply Contracts with Total Minimum Commitment," *IIE Transactions*, **29**, 5 (1997a), 373-381.

Bassok, Y. and R. Anupindi, "Analysis of Supply Contracts with Commitments and Flexibility," Working Paper, Northwestern University, 1997b.

Bessler, S.A. and A.F. Veinott, "Optimal Policy for a Dynamic Multi-Echelon Inventory Model," *Naval Research Logistics Quarterly*, **13**, 4 (1966), 355-389.

Blair, B.F. and T.R. Lewis, "Optimal Retail Contracts with Asymmetric Information and Moral Hazard," *Rand Journal of Economics*, **25** (1994), 284-296.

Cachon, G.P. and M.A. Lariviere, "Capacity Choice and Allocation: Strategic Behavior and Supply Chain Performance," Working Paper, The Fuqua School of Business, Duke University, 1996.

Cachon, G.P. and M.A. Lariviere, "Contracting to Assure Supply or What Did the Supplier Know and When Did He Know It?", Working Paper, The Fuqua School of Business, Duke University, 1997a.

Cachon, G.P. and M.A. Lariviere, "Capacity Allocation Using Past Sales: When to Turn-and-Earn," Working Paper, The Fuqua School of Business, Duke University, 1997b.

Chen, F., "Decentralized Supply Chains Subject to Information Delays," Working Paper, Graduate School of Business, Columbia University, 1997.

Clark, A.J, "An Informal Survey of Multi-Echelon Inventory Theory," *Naval Research Logistics Quarterly*, **19** (1972), 621-650.

Clark, A.J. and Scarf, H., "Optimal Policies for a Multiechelon Inventory Problem," *Management Science*, **6** (1960), 475-490.

Cohen, M.A and N. Agrawal, "An Analytical Comparison of Long and Short Term Contracts," Working Paper, Department of OPIM, The Wharton School, University of Pennsylvania, 1998.

Cohen, M.A., P.R. Kleindorfer, and H.L. Lee, "Service Constrained (s,S) Inventory Systems with Priority Demand Classes and Lost Sales," *Management Science*, **34**, 4 (1988), 482-499.

Corbett, C.J. and X. de Groote, "Integrated Supply-Chain Lot Sizing Under Asymmetric Information," Working Paper, UCLA, 1997.

Daniel, K.H., "A Delivery-Lag Inventory Model With Emergency," in H.E. Scarf, D.M. Gilford and M.W. Shelly (Eds.), *Multistage Inventory Models and Techniques*, Stanford University Press, Stanford, CA, 1963.

Dixit, A.K., "Quality and Quantity Competition," *Review of Economic Studies*, **46** (1979), 587-599.

Donohue, K.L, "Supply Contracts for Fashion Goods: Optimizing Channel Profit," Working Paper, Department of OPIM, The Wharton School, University of Pennsylvania, 1996.

Emmons, H. and S.M. Gilbert, "Note: The Role of Returns Policies in Pricing and Inventory Decisions for Catalogue Goods, " *Management Science*, **44**, 2 (1998), 276-283.

Eppen, G.D., "Effects of Centralization on Expected Costs in a Multi-Location Newsboy Problem," *Management Science*, **25**, 5 (1979), 498-501.

Eppen, G.D. and A.V. Iyer, "Backup Agreements in Fashion Buying-The Value of Upstream Flexibility," *Management Science*, **43** (1997), 1469-1484.

Eppen, G.D. and L. Schrage, "Centralized Ordering Policies in a Multi-Warehouse System with Lead Times and Random Demand," In L.B. Schwarz, (Ed.), *Multi-Level Production/Inventory Control Systems: Theory and Practice*, Netherlands: North-Holland Publishing Co., 1981, 51-67.

Federgruen, A., "Centralized Planning Models for Multi-Echelon Inventory Systems Under Uncertainty," in S.C. Graves, A.H.G. Rinnooy Kan and P.H. Zipkin (Eds.), *Handbooks in Operations Research and Management Science, Vol. 4 (Logistics of Production and Inventory)*, Elsevier Science Publishing Company B.V., Amsterdam, The Netherlands, 1993, 133-173.

Fukuda, Y., "Optimal Policies for the Inventory Problem with Negotiable Leadtime," *Management Science*, **10** (1964), 690-708.

Gerchak, Y. and M. Henig, "Component Commonality in Assemble-To-Order Systems: Models and Properties," *Naval Research Logistics Quarterly*, **36** (1989), 61-68.

Gilbert, S.M. and Z.K. Weng, "Incentive Effects Favor Non-Consolidating Queues in a Service System: The Principal-Agent Perspective," Working Paper, Weatherhead School of Management, Case Western Reserve University, 1997.

Grout, J.R. and D.P. Christy, "An Inventory Model of Incentives for On-Time Delivery in Just-in-Time Purchasing Contracts," *Naval Research Logistics*, **40** (1993), 863-877.

Ha, A.Y., "Supply Contract for a Short-Life-Cycle Product with Demand Uncertainty and Asymmetric Cost Information," Working Paper, Yale School of Management, 1997a.

Ha, A.Y., "Inventory Rationing in a Make-to-Stock Production System with Several Demand Classes and Lost Sales," *Management Science*, **43**, 8 (1997b), 1093-1103.

Hammond, J.H., "Barilla SpA (A)," HBS Case 9-694-046, Harvard Business School Publishing, Boston, MA, 1994.

Hausman, W.H., H.L. Lee, and L. Zhang, "Optimal Ordering for an Inventory System with Dual Lead Times," Working Paper, Department of Industrial Engineering-Engineering Management, Stanford University, 1994.

Harris, F.W., *Operations And Cost (Factory Management Series)*, Shaw, Chicago, 1915.

Iyer, A. and M.E. Bergen, "Quick Response in Manufacturer-Retailer Channels," *Management Science*, **43**, 4 (1997), 559-570.

Jackson, P.L., "Stock Allocation in a Two-Echelon Distribution System or 'What to Do Until Your Ship Comes In'," *Management Science*, **34**, 7 (1988), 880-895.

Jeuland, A.P. and S.M. Shugan, "Managing Channel Profits," *Marketing Science*, **2** (1983), 239-272.

Jonsson, H. and E.A. Silver, "Overview of a Stock Allocation Model for a Two-Echelon Push System Having Identical Units at the Lower Echelon," In S. Axsater, C. Schneeweiss and E.A. Silver, (Eds.), *Multi-Stage Production*

Planning and Inventory Control: Lecture Notes in Economics and Mathematical Systems, 266, Berlin: Springer-Verlag, 1981, 44-49.

Juran, J.M. and F.M. Gryna, *Juran's Quality Control Handbook*, 4th Edition, McGraw-Hill, New York, 1988.

Kandel, E., "The Right to Return," *Journal of Law and Economics*, **39** (1996), 329-356.

Katz, M.L., "Vertical Contractual Relations," in R. Schmalensee and R.D. Willig (Eds.), *Handbook of Industrial Organization: Volume I*, Elsevier Science Publishers B.V., New York, NY, 1989.

Kumar, A., P.S. Giridharan, and R. Akella, "Control Structures, Performance and Incentives: An Operations Management Perspective," Working Paper, Department of Engineering Economic Systems, Stanford University, 1996.

Kumar, N., "The Power of Trust in Manufacturer-Retailer Relationships," *Harvard Business Review*, **74** (1996), 92-106.

Lawson, D.G. and E.L. Porteus, "Dynamic Lead Time Management," Research Paper #1409, Graduate School of Business, Stanford University, 1996.

Lee, H.L., "Effective Inventory and Service Management Through Product and Process Redesign," *Operations Research*, **44** (1996), 151-159.

Lee, H.L., C. Billington, and B. Carter, "Hewlett-Packard Gains Control of Inventory and Service Through Design for Localization," *Interfaces*, **23** (1993), 1-11.

Lee, H.L., V. Padmanabhan, and S. Whang, "Information Distortion in a Supply Chain: The Bullwhip Effect," *Management Science*, **43** (1997), 546-558.

Lee, H.L. and M.J. Rosenblatt, "A Generalized Quantity Discount Pricing Model to Increase Supplier's Profits," *Management Science*, **30** (1986), 1179-1187.

Lee, H.L. and S. Whang, "Decentralized Multi-Echelon Supply Chains: Incentives and Information," Working Paper, Graduate School of Business, Stanford University, 1997.

Li, C. and P. Kouvelis, "Flexible and Risk-Sharing Supply Contracts Under Price Uncertainty," Working Paper, Olin School of Business, Washington University, 1997.

Milner, J.M. and M.J. Rosenblatt, "Two-Period Supply Contracts: Order Adjustments and Penalties," Working Paper, Olin School of Business, Washington University, 1997.

Moinzadeh, K. and C. Ingene, "An Inventory Model of Immediate and Delayed Delivery," *Management Science*, **39**, 5 (1993), 536-548.

Moinzadeh, K. and S. Nahmias, "Adjustment Strategies for a Fixed Delivery Contract," Working Paper, University of Washington, 1997.

Monahan, J.P., "A Quantitative Discount Pricing Model to Increase Vendor Profits," *Management Science*, **30** (1984), 720-726.

Moorthy, K.S., "Managing Channel Profits: Comment," *Marketing Science*, **6**, 4 (1987), 375-379.

Nahmias, S., *Production and Operations Analysis*, Irwin, Homewood, IL, 1997.

Nahmias, S. and W.S. Demmy, "Operating Characteristics of an Inventory System with Rationing," *Management Science*, **27** (1981), 1236-1244.

Narayanan, V.G. and A. Raman, "Assignment of Stocking Decision Rights under Incomplete Contracting," Working Paper, Harvard Business School, 1997.

Pasternack, B.A., "Optimal Pricing and Returns Policies for Perishable Commodities," *Marketing Science*, **4** (1985), 166-176.

Porteus, E.L., "Stochastic Inventory Theory," in D.P. Heyman and M.J. Sobel (Eds.), *Handbook in Operations Research and Management Science, Volume 2 (Stochastic Models)*, Elsevier Science Publishers B.V., 1990, 605-652.

Porteus, E.L. and S. Whang, "On Manufacturing/Marketing Incentives," *Management Science*, **37**, 9 (1991), 1166-1181.

Reyniers, D.J. and C.S. Tapiero, "The Delivery and Control of Quality in Supplier-Producer Contracts," *Management Science*, **41**, 10 (1995), 1581-1589.

Rosenblatt, M.J. and H.L. Lee, "Improving Profitability with Quantity Discounts under Fixed Demand," *IIE Transactions*, **17**, 4 (1985), 388-395.

Rosling, K., "Optimal Inventory Policies for Assembly Systems Under Random Demands," *Operations Research*, **37** (1989), 565-579.

Schwarz, L.B., *Multi-Level Production/Inventory Control Systems: Theory and Practice*, North-Holland Publishing Company, New York, NY, 1981.

Spence, A.M., "Monopoly, Quality, and Regulation," *Bell Journal of Economics*, **6** (1975), 417-429.

Spengler, J.J, "Vertical Restraints and Antitrust Policy," *Journal of Political Economy*, **58** (1950), 347-352.

Starbird, S.A., "The Effect of Acceptance Sampling and Risk Aversion on the Quality Delivered by Suppliers," *Journal of the Operational Research Society*, **45**, 3 (1994), 309-320.

Starbird, S.A., "Acceptance Sampling, Imperfect Production, and the Optimality of Zero Defects," *Naval Research Logistics*, **44** (1997a), 515-530.

Starbird, S.A., "Supply Chain Management, Joint Cost Optimization, and Quality," Working Paper, Leavey School of Business, Santa Clara University, 1997b.

Tagaras, G. and H.L. Lee, "Economic Models for Vendor Evaluation with Quality Cost Analysis," *Management Science*, **42**, 11 (1996), 1531-1543.

Tirole, J., *The Theory of Industrial Organization*, The MIT Press, Cambridge, MA, 1988.

Tsay, A.A, "The Quantity Flexibility Contract and Supplier-Customer Incentives," Working Paper, Leavey School of Business, Santa Clara University, 1996.

Tsay, A.A. and W.S. Lovejoy, "Quantity Flexibility Contracts and Supply Chain Performance," Working Paper, Leavey School of Business, Santa Clara University, 1998. To appear in *Manufacturing & Service Operations Management*.

Verity, J.W., "Clearing the Cobwebs from the Stockroom," *Business Week*, October 21, 1996, p.140.

Weng, Z.K., "Channel Coordination and Quantity Discounts," *Management Science*, **41** (1995), 1509-1522.

Weng, Z.K., "Pricing and Ordering Strategies in Manufacturing and Distribution Alliances," *IIE Transactions*, **29** (1997a), 681-692.

Weng, Z.K., "Coordination Strategies for Aligning Divergent Interests in a Manufacturing and Distribution Supply Chain," Working Paper, School of Business, University of Wisconsin-Madison, 1997b.

Weng, Z.K., "A Decision Framework for the Manufacturer to Coordinate with the Buyer," Working Paper, School of Business, University of Wisconsin-Madison, 1997c.

Whang, S., "Coordination in Operations: A Taxonomy," *Journal of Operations Management*, **12** (1995), 413-422.

Yano, C.A. and H.L. Lee, "Lot Sizing with Random Yields: A Review," *Operations Research*, **43**, 2 (1995), 311-334.