MANAGEMENT SCIENCE
Vol. 50, No. 2, February 2004, pp. 222–238
ISSN 0025-1909 | EISSN 1526-5501 | 04 | 5002 | 0222
DOI 10.1287/mnsc.1030.0190

The Allocation of Inventory Risk in a Supply Chain: Push, Pull, and Advance-Purchase Discount Contracts

Gérard P. Cachon
The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104
{cachon@wharton.upenn.edu, opim.wharton.upenn.edu/˜cachon}

While every firm in a supply chain bears supply risk (the cost of insufficient supply), some firms may, even with wholesale price contracts, completely avoid inventory risk (the cost of unsold inventory). With a push contract there is a single wholesale price and the retailer, by ordering his entire supply before the selling season, bears all of the supply chain's inventory risk. A pull contract also has a single wholesale price, but the supplier bears the supply chain's inventory risk because only the supplier holds inventory while the retailer replenishes as needed during the season. (Examples include Vendor Managed Inventory with consignment and drop shipping.) An advance-purchase discount has two wholesale prices: a discounted price for inventory purchased before the season, and a regular price for replenishments during the selling season. Advance-purchase discounts allow for intermediate allocations of inventory risk: The retailer bears the risk on inventory ordered before the season while the supplier bears the risk on any production in excess of that amount. This research studies how the allocation of inventory risk (via these three types of wholesale price contracts) impacts supply chain efficiency (the ratio of the supply chain's profit to its maximum profit). It is found that the efficiency of a single wholesale price contract is considerably higher than previously thought as long as firms consider both push and pull contracts. In other words, the literature has exaggerated the value of implementing coordinating contracts (i.e., contracts that achieve 100% efficiency, such as buy-backs or revenue sharing) because coordinating contracts are compared against an inappropriate benchmark (often just a push contract). Furthermore, if firms also consider advance-purchase discounts, which are also simple to administer, then the coordination of the supply chain and the arbitrary allocation of its profit is possible. Several limitations of advance-purchase discounts are discussed.

Key words: Nash equilibrium; contracting; coordination; bargaining; newsvendor model
History: Accepted by Fangruo Chen, former department editor; received November 20, 2002. This paper was with the author 2 months for 2 revisions.

1. Introduction

Supply chain management is about matching supply and demand, particularly so with inventory management: Too much supply leads to inefficient capital investment, expensive markdowns and needless handling costs, while too much demand generates the opportunity cost of lost margins. Each situation is the consequence of one of two types of inventory risk: The former is the risk of excessive inventory (inventory risk) while the latter is the risk of insufficient supply (supply risk). Because most supply chains are incapable of perfectly matching supply and demand, all of the firms in a supply chain bear at least some supply risk. But some firms may be able to avoid inventory risk completely.

Consider a supply chain with one supplier and one retailer, and suppose the firms trade with a wholesale price contract. If the retailer orders well in advance of the selling season and the supplier produces just the retailer's order quantity, then the retailer bears all of the supply chain's inventory risk and the supplier bears none. The other extreme is also possible. Suppose the supplier's product is shipped to the retailer on consignment, or the supplier holds the inventory while replenishing the retailer frequently and in small batches during the season. Now the supplier bears essentially all of the supply chain's inventory risk. Advance-purchase discounts generate intermediate allocations of inventory risk. With an advance-purchase discount the retailer bears the inventory risk on inventory purchased at the discount price before the season begins. In anticipation of the retailer's replenishment orders during the season, which trade at the regular wholesale price, the supplier may produce more than the retailer's initial order, thereby bearing the risk on those units.

This paper studies how the allocation of inventory risk influences a supply chain's performance and its division of profit. Two examples from the sporting goods industry help to motivate this research. Trek Inc. manufacturers a wide variety of bicycles, from simple bicycles for kids, to high-end bicycles for enthusiasts, to ultrasophisticated custom bicycles for Lance Armstrong in the Tour de France. They sell exclusively through authorized independent bicycle retailers. Especially in the high-end segment, demand is seasonal and highly uncertain. Yet, despite all of the uncertainties associated with demand, retailers bear essentially no inventory risk: Trek holds inventory in a central warehouse and is willing to ship, at Trek's expense, bicycles to retailers one unit at a time. Thus, retailers carry a few demonstration models, but otherwise only order bicycles when they have a firm customer order. They receive a bicycle from Trek within a few days, assuming it is available at Trek's warehouse. (In this segment customers are generally willing to wait a few days.)

O'Neill Inc. designs and sells apparel and accessories for several water sports (surfing, scuba diving, water skiing, wake boarding, wind surfing, and triathlon), again from the recreational to professional level. The lead time from their manufacturing facility in Thailand is long (3 months), but shipments to retailers are relatively fast from their distribution center in San Diego (1–3 days). They have monthly capacity limits, so they must begin production well in advance of the selling season. O'Neill accepts two kinds of orders from retailers. A "prebook" order is submitted several months before the season starts. Retailers are guaranteed to receive their prebook order, but they also bear the inventory risk on that order (O'Neill does not accept returns). The other type of order is called an "at-once" order: At-once orders are submitted during the selling season for immediate (at-once) delivery, although at-once orders are only filled if inventory is available at O'Neill. Because O'Neill anticipates a substantial amount of at-once orders, O'Neill generally produces more than is prebooked, and bears the inventory risk on that excess production. To encourage retailers to prebook inventory, several years ago O'Neill began offering retailers an advance-purchase discount, which they call a prebook discount. The program has been successful and, in response, they have increased the discount depth on several occasions.

These two examples illustrate that supply chains operate with different inventory risk allocations. Furthermore, the allocation of inventory risk can be adjusted with advance-purchase discounts; the greater the discount, the more risk is shifted to the retailer. A stylized model, which captures some of the features of the industry examples described (but surely not all), is used to study how the allocation of inventory risk impacts supply chain performance. The model has one supplier and one retailer. Total season demand is stochastic with a known distribution function. Due to a long production lead time, the supplier must commit to a production quantity well in advance of the selling season, but due to the relatively short delivery lead time from the supplier to the retailer, the retailer can accept replenishments both before the selling season as well as during the selling season.

The terms of trade between the firms are chosen from three types of wholesale price contracts. The supplier could charge a single wholesale price and not offer at-once orders, i.e., the retailer must prebook inventory and the supplier only produces the retailer's prebook quantity, thereby removing even the possibility of at-once orders. With that "push" contract, all inventory risk is pushed onto the retailer. Lariviere and Porteus (2001) refer to that situation as a "supplier selling to a newsvendor": A single wholesale price is chosen and the retailer is required to purchase inventory before the selling season. Others refer to that practice as "channel stuffing," i.e., the supplier attempts to stuff the retailer with inventory.

In contrast to push, the firms could adopt a pull contract, which also has a single wholesale price but now the supplier charges that wholesale price for both prebook and at-once orders. Instead of prebooking inventory, with a pull contract the retailer pulls inventory from the supplier with at-once orders, thereby leaving the supplier with all inventory risk (as in the Trek example), i.e., the "retailer buys from a newsvendor." There are two other situations that can be represented by a pull contract: Vendor Managed Inventory with consignment inventory (the supplier decides how much inventory to stock at the retailer and owns that inventory), or drop shipping (the supplier holds the inventory and ships directly to consumers, bypassing the retailer).

The third contract option, advance-purchase discounts, blends both push and pull by having two wholesale prices. The prebook wholesale price is lower than the at-once wholesale price, so the retailer may prebook some inventory (bearing the risk on that inventory) and the supplier may produce additional inventory in anticipation of at-once orders (and bears the risk on that additional production), as in the O'Neill example.

The particular contract adopted by the firms is the outcome of some bargaining process. For example, it is possible that one of the firms makes a "take-it-or-leave-it" offer to the other firm, possibly subject to leaving the other firm with some minimum acceptable profit, or, more likely, the firms engage in some alternating offer bargaining process. The bargaining process is intentionally not specified because while the outcome of a particular bargaining process is often quite precise (e.g., a specific single contract), there is no contract that is the outcome of a wide range of bargaining processes.[1] In other words, the

[1] For example, a single outcome is generated if one of the firms is a Stackelberg leader or if the Nash bargaining solution is implemented, but there is little reason to believe that these are the only

likely outcome of the negotiation between the firms is sensitive to what is assumed about how the negotiations will be conducted. Hence, it is valuable to make some prediction regarding the contracts the firms are likely to choose that is independent of the details of the bargaining process. This is done by dividing the set of contracts into two types: contracts that could plausibly be the outcome of the negotiation for some negotiation process, and contracts that should not be observed no matter what negotiation process is used. The former is the set of Pareto contracts: A contract is Pareto if there does not exist an alternative contract such that no firm is worse off and one firm is strictly better off. The latter set is the Pareto-inferior contracts: A contract is Pareto inferior if there exists some other contract that makes one firm better off and no firm worse off. A Pareto-inferior contract should not be chosen by the firms no matter the bargaining process, because, by definition, any Pareto-inferior contract is open to a counteroffer that makes no firm worse off and another firm better off.

If the Pareto set is restrictive (i.e., there exist Pareto-inferior contracts) then it provides a meaningful, albeit not perfect, prediction for the outcome of the contract negotiation. Furthermore, the Pareto set provides a useful bound on a key performance metric, supply chain efficiency (the ratio of the supply chain's expected profit to its maximum expected profit): The Pareto set's minimum efficiency can be achieved without coordination, so the difference between that efficiency and the optimal efficiency (100%) is an upper bound on the value of supply chain coordination activities.

It is also worth mentioning that cooperative game theory is unlikely to be useful in this setting. With cooperative game theory, each potential coalition of players is assigned a number that is the total value that can be achieved by the players in the coalition independent of the activities of the players outside the coalition. In this game there are three possible coalitions: Each firm is on its own or both firms form a single coalition. The case with either firm on its own is not interesting: There is some exogenous value that each firm can achieve independent of the other firm, but this setting provides no insight into what that value could or should be, i.e., it must be treated as an exogenous parameter. If the firms join into a single coalition, then their value equals the supply chain's maximum profit, but cooperative game theory provides no guidance as to how the firms would divide the value assigned to their coalition. Nor is it reasonable to assume that the firms would necessarily achieve 100% efficiency, i.e., there may be some efficiency loss in the bargaining process. (See Cachon and Netessine 2003 for a more extensive treatment on cooperative game theory.)

The next section reviews the related literature. Section 3 describes the model in greater detail. Section 4 evaluates the set of Pareto contracts. Section 5 evaluates three natural extensions to the model. The final section summarizes and discusses the results.

2. Literature Review

The newsvendor model provides the backbone for this work. Because there is a considerable amount of related literature, this section focuses on this paper's major points of difference with respect to modeling assumptions and analysis. The papers closest in spirit to this work—because they each study wholesale price contracts in a newsvendor setting—are Lariviere and Porteus (2001), Cachon and Lariviere (2001), Netessine and Rudi (2001a, b), Ferguson et al. (2002), Taylor (2002b), and Özer and Wei (2002).

Lariviere and Porteus (2001) study a model with a "supplier selling to a newsvendor" using a single wholesale price contract, which is equivalent to the "push" contract considered here. However, they assume the supplier chooses the wholesale price, possibly subject to the retailer's participation constraint, and they do not consider pull or advance-purchase discount contracts. In Cachon and Lariviere (2001) a manufacturer purchases from a newsvendor with a wholesale price contract (among other contracts considered), which is equivalent to the "pull" contract considered here. They do not consider push or advance-purchase discount contracts.

Netessine and Rudi (2001a) study a multiperiod newsvendor setting with one supplier and one retailer. Two supply chain strategies are compared: with traditional operations the retailer purchases inventory from the supplier, and with drop shipping the supplier holds inventory (and fills demand directly). The former is a push contract and the latter is a pull contract. They do not consider advance-purchase discounts, nor do they identify the Pareto set of contracts. Netessine and Rudi (2001b) study a single-period model with stochastic demand, multiple retailers, and a single production opportunity. In addition to the traditional and drop-shipping supply chains, they consider a dual-strategy supply chain in which both firms may hold inventory. While those three supply chain structures correspond to the pull, push, and advance-purchase discounts considered here, respectively, their focus is different than in this paper: Among other results, they investigate the relative attractiveness of the traditional and drop shipping supply chains with respect to demand variability and the number of retailers; they characterize equilibrium inventory quantities in the dual-strategy supply

reasonable outcomes of the game. In fact, it is easy to argue that the Stackelberg solutions are unreasonable in practice.

chain; and they numerically evaluate the impact of the drop-shipping markup and number of retailers on supply chain profits with the dual strategy. They find that the supply chain's profit can be higher or lower with drop shipping relative to the traditional strategy. Randall et al. (2002) find empirical evidence in support of the hypothesis that neither strategy dominates in all cases.

Ferguson et al. (2002) build on Ferguson (2003) with a model that is qualitatively similar to the one here. They consider "early-commitment" contracts (push) and "delayed-commitment" contracts (pull), but they do not study advance-purchase contracts, nor do they identify the Pareto set of contracts. Instead, they work with three regimes for choosing the wholesale price: The supplier chooses the price, the manufacturer chooses the price, and the price is chosen to equalize the firms' profits.

Taylor (2002b) studies a manufacturer that can choose to set her contract terms either early or late and characterizes conditions under which either strategy is preferable. In his model the retailer sets the retailer price and asymmetric information is studied. The Pareto set is not identified and advance-purchase discounts are not considered. Özer and Wei (2002) study (essentially) the same model as here and also consider advance-purchase discounts, but they neither identify the set of coordinating advance-purchase discounts nor consider pull contracts. However, they also study asymmetric information.

There are other papers that study wholesale price contracts and inventory management. Cachon and Lariviere (2000) demonstrate that the efficiency of a single wholesale price contract depends on the shape of the retailer's marginal revenue curve. Erhun et al. (2000) and Anand et al. (1999) study wholesale price contracts over multiple periods assuming only the current period wholesale price is fixed. Here, the firms commit up front to the two wholesale prices.

A number of papers study contracting to achieve supply chain coordination in the context of newsvendor models: buy-back contracts (Pasternack 1985), quantity flexibility contracts (Tsay 1999), revenue-sharing contracts (Cachon and Lariviere 2000), sales-rebate contracts (Taylor 2002a, Krishnan et al. 2004), price-discount contracts (Bernstein and Federgruen 2002), and quantity discounts (Cachon 2003, Tomlin 2000). In most of these models it is assumed the supplier's entire production is shipped to the retailer and the retailer never receives more than one replenishment. While these contracts maximize and arbitrarily allocate the supply chain's profit, they all have additional administrative, handling, and monitoring costs that are not explicitly considered. It is possible that these contracts are not implemented in some settings because the additional costs are actually higher than the incremental benefit these contracts provide over simpler contracts, such as wholesale price contracts. The two companies mentioned in the introduction do not use any of those coordinating contracts.

Anupindi and Bassok (1999) consider a one-supplier, multiple-retailer supply chain and analyze performance with two structures: Each retailer carries its own inventory or the retailers pool their inventory. Therefore, their model captures shifts in inventory risk among retailers, but does not address the issue of allocating inventory risk between different levels of the supply chain. Inventory risk allocation among retailers is also studied in models with redistribution of inventory (e.g., Anupindi et al. 2001, Dong and Rudi 2001, Rudi et al. 2001).

Barnes-Schuster et al. (1998), Donohue (2000), Fisher et al. (2001), Fisher and Raman (1996), and Iyer and Bergen (1997) study two-period newsvendor models with demand updating between periods. They do not consider advance-purchase discounts. Iyer and Bergen (1997) compare a structure in which the retailer orders before the demand update (a push contract) with a structure in which the retailer orders after the demand update (quick response, which is not considered here), but in each case the retailer bears all inventory risk. Weng and Parlar (1999), Tang et al. (2001), and McCardle et al. (2002) study the application of advance-purchase discounts between a retailer and his consumers, which is not comparable to the application of advance-purchase discounts between a supplier and a retailer. (For example, in their setting, a retailer uses an advance-purchase discount to gain information about future demand, whereas that motivation is not present in this model.) Several papers consider the use of advance selling to consumers that face uncertainty in their own valuation, which is also not considered in this paper (e.g., Dana 1998, Xie and Shugan 2001).

3. Model

A risk-neutral retailer buys a product from a risk-neutral supplier and sells that product over a single selling season. Demand during the selling season is stochastic: Let $F(x)$ and $f(x)$ be the distribution and density functions of demand, respectively. Assume F is strictly increasing, differentiable, and $F(0) = 0$ (i.e., there is always some demand).[2] Let $g(x)$ be the generalized failure rate $g(x) = xh(x)$, where $h(x)$ is the failure rate, $h(x) = f(x)/(1 - F(x))$. Assume the demand distribution has the strictly increasing generalized failure rate (IGFR) property: $g'(x) > 0$. Many distributions have the IGFR property, including

[2] If $F(0) > 0$ then the marginal analysis provided in this paper must be argumented to included boundary conditions. See footnote 4 for other implications of this assumption.

the Normal, the exponential, the gamma, and the Weibull.[3]

There is one production opportunity, which occurs well before the selling season due to the long production lead time. Let q be the quantity produced. The production cost per unit is c. The retail price p during the selling season is fixed, $p > c$. Units remaining at the end of the season are salvaged for v per unit, $v < c$, no matter which firm salvages the unit. (If there were differences in the salvage value between the retailer and the supplier, then that difference would clearly influence the appropriate allocation of inventory risk. That issue is not explored.)

There are two types of retail orders. The first type is the retailer's "prebook" order, y. The prebook is submitted to the supplier before production begins, which is well in advance of the selling season. As a result, the retailer receives his prebook order at the start of the selling season. The supplier charges the retailer w_1 per unit in the prebook order. The remaining orders are called "at-once" orders. These orders are submitted during the selling season, and, if the supplier has inventory available, an at-once order is received immediately by the retailer. Hence, the retailer submits at-once orders only after running out of the prebook inventory and then only as the retailer incurs demand. The retailer pays the supplier $w_2 \geq w_1$ per unit in an at-once order.

Assume at-once orders do not incur additional shipping or handling costs relative to the prebook order. This is plausible in some cases. For example, the supplier may ship products individually and with the same shipping service whether a unit is ordered as part of the prebook or whether it is an at-once order. (That is the case with Trek bicycles.) Or, the firms may position inventory at the retailer with consignment, in which case an at-once order merely transfers ownership of the inventory without requiring physically moving the inventory (e.g., VMI with consignment). Or, if early season sales are sufficiently informative of total season sales (as observed by Fisher et al. 2001 and Fisher and Raman 1996), then the retailer may be able to submit a single at-once order early in the season (which bears little inventory risk due to the updated forecast) and the shipping cost for that consolidated order may not be substantially larger than for the prebook order. Thus, assuming the retailer bears no risk with at-once orders and at-once orders do not incur additional shipping and handling costs is either an accurate representation or a reasonable approximation of many supply chains. Nevertheless, §5 discusses the implications of relaxing those assumptions.

Both the prebook wholesale price, w_1, and the at-once wholesale price, w_2, are set before the retailer submits his prebook order and both remain in effect until the end of the season (i.e., neither firm attempts to renegotiate w_2 after the prebook order). A pair of wholesale prices $\{w_1, w_2\}$ is referred to as a contract, or the terms of trade between the two firms. As discussed in the introduction, there is no single process assumed by which the firms choose the terms: It is possible the supplier sets the terms with a take-it-or-leave-it offer, possibly subject to giving the retailer at least a minimum acceptable profit (i.e., a participation constraint); or, the retailer sets the terms subject to a profit constraint for the supplier; or, the terms are set by some alternating bargaining process.

There are two extreme types of contracts. With a push contract there is effectively a single wholesale price $\hat{w}_1 < p$, the retailer must prebook inventory (i.e., purchase in advance of the selling season), and there are no opportunities for at-once orders (because the supplier does not produce any more than the retailer's prebook order quantity). (To be consistent with the previous notation, with push $w_2 > p$, so at-once orders are never submitted.) With a pull contract there is effectively a single wholesale price that remains in effect both before and during the selling season, $w_1 = w_2 < p$. With pull the retailer does not prebook inventory, opting instead for only at-once orders. (However, $w_1 = w_2$ does not guarantee the supply chain operates with pull because the retailer may prefer to prebook more inventory than the supplier would produce so that total supply is increased. In effect, the supply chain would then operate in push mode. This issue is discussed later.) The intermediate case between push and pull occurs when there is an advance-purchase discount, $w_1 < w_2 < p$.

4. Analysis

The analysis proceeds along the following sequence. Subsection 4.1 determines the supply chain's optimal production and profit. Subsection 4.2 evaluates the production quantity and the firms' expected profits in push mode. Subsection 4.3 evaluates the production quantity and the firms' expected profits in pull mode. Subsection 4.4 identifies and characterizes the Pareto set of contracts among the push and pull contracts. Subsection 4.5 considers the Pareto set and supply chain coordination with advance-purchase discount contracts.

4.1. Integrated Supply Chain

The integrated supply chain maximizes the sum of the retailer's profit and the supplier's profit. The supply

[3] Lariviere and Porteus (2001) also make extensive use of the IGFR property, but they only assume the generalized failure rate is weakly increasing. The Pareto distribution is the only distribution with a constant generalized failure rate over its entire domain. It can be shown that the results in this section also apply to the Pareto distribution, but the analysis is more cumbersome.

chain's only decision is the production quantity q. The supply chain's expected profit is

$$\Pi(q) = (p-v)S(q) - (c-v)q,$$

where

$$S(q) = q - \int_0^q F(x)\,dx. \tag{1}$$

Hence, the supply chain faces a newsvendor decision: $\Pi(q)$ is concave in q, maximized at q^o, increasing for $q \in [0, q^o]$, and the optimal production quantity satisfies

$$F(q^o) = \frac{p-c}{p-v}. \tag{2}$$

For notational convenience, let $\Pi^o = \Pi(q^o)$. The efficiency of a contract is $\Pi(q)/\Pi^o$.

4.2. Push

With a push contract the retailer only prebooks inventory and pays $\hat{w}_1 < p$ per unit: The supplier sells to a newsvendor that bears all of the supply chain's inventory risk. (Where useful, the notation "ˆ" is used to indicate association with a push contract.) This section evaluates the firms' preferences over the set of push contracts.

With push the supplier's production equals the retailer's prebook, $q = y$, so when discussing push contracts it is notationally convenient to let q also be the retailer's prebook quantity. The retailer's profit is then $\hat{\pi}_r(q, \hat{w}_1) = (p-v)S(q) - (\hat{w}_1 - v)q$. The optimal prebook is implicitly defined by

$$F(q) = \frac{p-\hat{w}_1}{p-v}. \tag{3}$$

Because $F(q)$ is increasing in q, there is a one-to-one relationship between $\hat{w}_1$ and q. Use (3) to solve for $\hat{w}_1$ in terms of q:

$$\hat{w}_1(q) = p - (p-v)F(q).$$

Hence, the analysis can be done in terms of either $\hat{w}_1$ or q: If the analysis is done in terms of q, then it is understood that a push contract with quantity q is in fact a contract with a $\hat{w}_1(q)$ wholesale price and $\hat{w}_2 > p$. However, as is explained in §4.3, there are advantages to working with q rather than $\hat{w}_1$.

To continue, evaluate the retailer's profit in terms of q

$$\hat{\pi}_r(q) = \hat{\pi}_r(q, \hat{w}_1(q)) = (p-v)(1-F(q))(j(q)-q),$$

where

$$j(q) = \frac{S(q)}{1-F(q)}. \tag{4}$$

The retailer's profit is increasing in q:

$$\hat{\pi}_r'(q) = (p-v)f(q)q \ge 0. \tag{5}$$

The supplier's profit is $\hat{\pi}_s(q, \hat{w}_1) = (\hat{w}_1 - c)q$, which, in terms of q, is

$$\hat{\pi}_s(q) = \hat{\pi}_s(q, \hat{w}_1(q)) = (p-v)(F(q^o) - F(q))q. \tag{6}$$

Lariviere and Porteus (2001) demonstrate that $\hat{\pi}_s(q)$ is unimodal in q if the demand distribution has the IGFR property. Let $\hat{q}^*$ be the supplier's most preferred quantity with a push contract, $\hat{q}^* = \arg\max \hat{\pi}_s(q)$. If the supplier were able to choose any wholesale price in push mode then the supplier would choose $\hat{w}_1(\hat{q}^*)$.

4.3. Pull

A pull contract has a single wholesale price, $w_1 = w_2 < p$ and the retailer does not prebook inventory. As a result, the supplier chooses a production quantity q, bears the inventory risk on those units, and q units are available to fill at-once orders. This section evaluates the firms' preferences over the set of pull contracts.

The supplier's expected profit is $\pi_s(q, w_1) = (w_1 - v)S(q) - (c-v)q$. Because $\pi_s(q, w_1)$ is strictly concave in q, the supplier's optimal quantity is implicitly defined by

$$F(q) = \frac{w_1 - c}{w_1 - v}. \tag{7}$$

As with push, there is a one-to-one relationship between the production quantity q and the wholesale price w_1. Use (7) to solve for w_1 in terms of q:

$$w_1(q) = \frac{c - vF(q)}{1-F(q)}. \tag{8}$$

Again, as with push, the analysis of the pull mode can be done in terms of the production quantity q: If q is the contract chosen with pull, then the wholesale price is $w_1(q)$. Working with the production quantity is expositionally cleaner because the supply chain's profit is $\Pi(q)$ no matter whether q is chosen in push mode or pull mode; the same does not hold with the wholesale price.

Use (8) to write the supplier's profit in terms of q:

$$\pi_s(q) = \pi_s(q, w_1(q)) = (p-v)(1-F(q^o))(j(q)-q).$$

The supplier's profit is increasing in q

$$\pi_s'(q) = (p-v)(1-F(q^o))j(q)h(q) \ge 0. \tag{9}$$

According to the following lemma, $j(q)h(q)$ is increasing, so the supplier's profit is also convex. The lemma is also needed for many of the subsequent results.

Lemma 1. *For $q > 0$, $j(q)h(q)$ is increasing, where $j(q) = S(q)/(1-F(q))$ and $h(q)$ is the hazard rate, $f(q)/(1-F(q))$.*

Proof. Differentiate

$$\frac{\partial(j(q)h(q))}{\partial q} = \frac{qg'(q)j(q) + g(q)(j'(q)q - j(q))}{q^2}, \qquad (10)$$

where note that $h(q) = g(q)/q$. By definition, if F is strictly IGFR then $g'(q) > 0$. Thus, (10) is increasing if

$$j'(q)q > j(q). \qquad (11)$$

Given that $j'(q) = 1 + j(q)g(q)/q$, (11) can be written as

$$g(q) > 1 - \frac{q}{j(q)}. \qquad (12)$$

Note that $g(0) \geq (1 - 0/j(0))$ because $\lim_{q\to 0} 1 - q/j(q) = 0$ (from L'Hopital's rule). Thus, the right-hand side of (12) does not begin "above" $g(q)$. (12) then holds for all $q > 0$ if $g(q) = 1 - q/j(q)$ implies $g'(q) > \partial(1 - q/j(q))/\partial q$: If the right-hand side of (12) crosses $g(q)$ at a slower rate than $g(q)$ is increasing, then in fact the right-hand side of (12) does not cross $g(q)$. From differentiation,

$$\frac{\partial}{\partial q}\left(1 - \frac{q}{j(q)}\right) = -\frac{j(q) - qj'(q)}{j(q)^2} = \frac{1}{j(q)}\left(g(q) - \left(1 - \frac{q}{j(q)}\right)\right).$$

Hence, if $g(q) = 1 - q/j(q)$, then the right-hand side of (12) is not increasing and (12) is confirmed. □

The retailer's expected profit is $\pi_r(q, w_1) = (p - w_1)S(q)$, which can be written in terms of q:

$$\pi_r(q) = \pi_r(q, w_1(q)) = (p - v)(F(q^o) - F(q))j(q). \qquad (13)$$

The retailer's profit with pull is similar to the supplier's profit with push with one important difference: The supplier's profit is the product of her margin and her production quantity, whereas the retailer's profit is the product of his margin and his expected sales. Even though this distinction creates a nontrivial analytical challenge, as the supplier's profit is unimodal in q with push, Theorem 2 states that the retailer's profit is unimodal in q with pull.[4]

Theorem 2. *The retailer's profit with a pull contract, $\pi_r(q)$, is concave in q.*

Proof. $\pi_r(q)$ is concave in q if $\pi_r'(q)$ is decreasing in q,

$$\pi_r'(q) = (p - v)[(F(q^o) - F(q))j'(q) - f(q)j(q)] = (p - v)[F(q^o) - (F(q) + j(q)h(q)(1 - F(q^o)))].$$

[4] Cachon and Lariviere (2001) show that the supplier's profit with their pull contract is unimodal in w if the demand distribution has the IFR property, which is somewhat more restrictive than the IGFR property assumed here.

The above is decreasing in q if $j(q)h(q)$ is increasing in q, which Lemma 1 establishes. □

Let q^* be the retailer's most preferred quantity with a pull contract, $q^* = \arg\max \pi_r(q)$. If the retailer were able to choose any wholesale price in pull mode then the retailer would choose $w_1(q^*)$.

While this analysis assumes the retailer does not prebook when $w_1 = w_2$, in fact, it may be in the retailer's interest to prebook: If q is low, the retailer may be better off prebooking more than q units to increase the available supply. ($\pi_r(q)$ is the retailer's actual profit function only if the retailer does not prebook. In general, the retailer's profit function is concave for $y < q$ and concave for $y > q$, but not necessarily globally concave.) In that case the supplier would only produce the retailer's prebook quantity and there would be no inventory available for at-once orders, i.e., the supply chain would effectively operate in push mode. Therefore, the "push challenge" to the pull contract cannot be ignored. (The "push challenge" refers to the possibility that the retailer may prebook inventory even if $w_1 = w_2$, i.e., even if a contract is designed to operate in pull mode, the retailer may challenge that presumption and nevertheless operate in push mode.) However, the next section demonstrates that it is indeed optimal for the retailer to not prebook inventory with the pull contracts in the Pareto set.

4.4. The Pareto Set with Push and Pull Contracts

Push and pull contracts are simple because they each only have a single wholesale price. Hence, it is plausible the firms may restrict attention to just the set of push and pull contracts. Furthermore, given that push contracts provide one extreme allocation of inventory risk, it is natural to balance them with pull contracts which provide the other extreme. Thus, this section identifies the Pareto contracts within the set of just push and pull contracts and investigates its properties (in particular, the minimum efficiency of the Pareto set). Advance-purchase discounts are included in the next section.

Define the "push Pareto set" to be the Pareto set among just the push contracts. Given that the supplier's profit is unimodal in q and the retailer's profit is increasing in q, it follows immediately that the push Pareto set is $[\hat{q}^*, q^o]$ and the set of Pareto-inferior contracts is $[0, \hat{q}^*) \cup (q^o, \infty)$. ($q > q^o$ is clearly always Pareto inferior, so no further attention is given to those contracts.) This can be seen graphically in Figure 1. (All of the numerical results reported in this section, including Figures 1 and 2 and Tables 1 and 2, are constructed assuming demand has a Gamma distribution with mean 10, $p = 10$, and $v = 0$. While absolute outcomes depend on the chosen salvage value,

Figure 1 **Profit with Push and Pull Contracts**

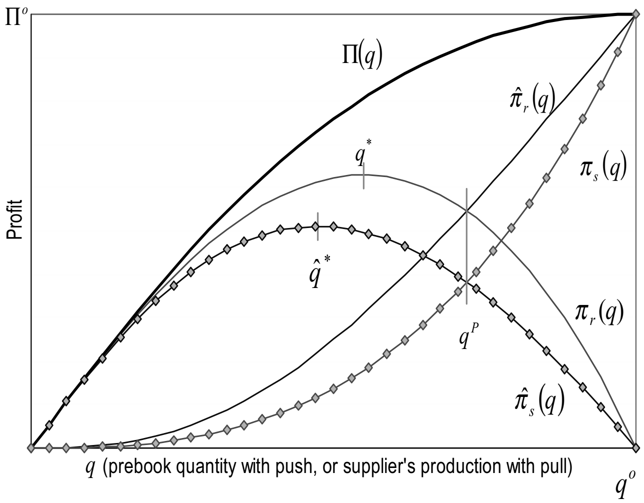

Note. Demand follows a Gamma distribution with mean 10, the coefficient of variation $2^{-1/2} \approx 0.707$, $p = 10$, $c = 4$, and $v = 0$.

the relative data reported here do not.) Given that the supply chain's profit is increasing on the interval $[0, q^o]$, $\hat{q}^*$ yields the minimum efficiency of the push Pareto set, i.e., if the supplier chooses the contract in push mode then the supplier chooses the least efficient contract in the push Pareto set. As Lariviere and Porteus (2001) argue, if the supplier chooses the contract subject to a retailer participation constraint, then the supplier chooses the smallest q to satisfy the retailer's minimum acceptable profit. Because that quantity is greater than $\hat{q}^*$, the supply chain's efficiency would increase relative to the $\hat{q}^*$ contract. Lariviere and Porteus (2001) conclude that increasing retailer bargaining power (in the form of a higher minimum acceptable profit) increases supply chain efficiency.

The "pull Pareto set" is the Pareto set among just the pull contracts (and assume all pull contracts survive the push challenge, which is later confirmed). Given that the retailer's profit is unimodal in q and the supplier's profit is increasing in q, it follows that the pull Pareto set is $[q^*, q^o]$ and the set of Pareto-inferior contracts is $[0, q^*)$. Again, this result can be seen for one example in Figure 1. Analogous to the push Pareto set, the retailer's preferred contract in the Pareto set, q^*, is the least efficient contract. Thus, increasing the supplier's bargaining power, in the form of a higher minimum acceptable profit, would increase supply chain efficiency when the retailer is given the privilege of choosing the contract. However, according to Theorem 3, the push and pull Pareto sets are not the same: the minimum efficiency of the pull Pareto set is higher than the push Pareto set.

Figure 2 **Supply Chain Efficiency,** $\Pi(q)/\Pi^o$

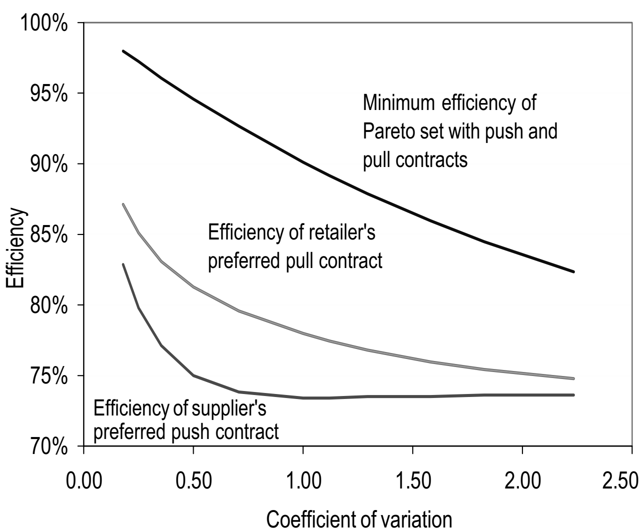

Note. Demand follows a Gamma distribution with mean 10, the coefficient of variation $2^{-1/2} \approx 0.707$, $p = 10$, $c = 5$, and $v = 0$.

Theorem 3. *The retailer's maximum profit with pull is greater than the supplier's maximum profit with push*: $\pi_r(q^*) > \hat{\pi}_s(\hat{q}^*)$, *the inventory in the supply chain is greater*, $q^* > \hat{q}^*$ *and the supply chain's profit* (*and efficiency*) *is higher*, $\Pi(q^*) > \Pi(\hat{q}^*)$.

Proof. For any q, the retailer earns more in pull mode than the supplier in push mode if $\pi_r(q) > \hat{\pi}_s(q)$, which, from (13) and (6), can be written as $j(q) > q$. From (4) and (1), the condition $j(q) > q$ can be written as

$$qF(q) > \int_0^q F(x)\,dx,$$

which holds given F is increasing. It follows that $\pi_r(q^*) \geq \pi_r(\hat{q}^*) > \hat{\pi}_s(\hat{q}^*)$. To demonstrate $q^* > \hat{q}^*$, begin with the supplier's first-order condition in push mode,

$$\hat{\pi}_s'(\hat{q}^*) = (p-v)(F(q^o) - F(\hat{q}^*) - f(\hat{q}^*)\hat{q}^*) = 0,$$

which implies

$$F(q^o) - F(\hat{q}^*) = f(\hat{q}^*)\hat{q}^*. \tag{14}$$

Substitute (14) into $\pi_r'(\hat{q}^*)$,

$$\begin{aligned}\pi_r'(\hat{q}^*) &= (p-v)[(F(q^o) - F(\hat{q}^*))j'(\hat{q}^*) - f(\hat{q}^*)j(\hat{q}^*)]\\ &= (p-v)f(\hat{q}^*)[\hat{q}^* j'(\hat{q}^*) - j(\hat{q}^*)].\end{aligned}$$

As shown in Lemma 1, $\hat{q}^* j'(\hat{q}^*) > j(\hat{q}^*)$, so $\pi_r'(\hat{q}^*) > 0$. Given that π_r is unimodal, it follows that $q^* > \hat{q}^*$. It is easy to confirm that $q^o > q^*$, hence $\Pi(q^*) > \Pi(\hat{q}^*)$. □

Theorem 3 indicates the pull mode is attractive to the supply chain relative to the push mode. It remains to identify the Pareto set when both contract types are considered. The next lemma identifies the existence of a useful quantity, q^P: the retailer's profit with the q^P pull contract equals his profit with the q^P push contract and the supplier's profit with the q^P pull contract also equals her profit with the q^P push contract, as is displayed in Figure 1.

Table 1 Supplier's Optimal Push Contract and the Impact of Switching to the Supplier's Preferred Pareto Contract That Leaves the Retailer No Worse Off

		Supplier's optimal push contract, $\hat{q}^*$		Supplier's preferred Pareto contract, q', subject to $\hat{\pi}_r(\hat{q}^*) \le \pi_r(q')$		
σ/μ	$\frac{p-c}{p-v}$	Supplier's profit share $\frac{\hat{\pi}_S(\hat{q}^*)}{\Pi^o}$ (%)	Efficiency $\frac{\hat{\pi}_S(\hat{q}^*)+\hat{\pi}_r(\hat{q}^*)}{\Pi^o}$ (%)	Supplier's profit increase $\frac{\pi_S(q')-\hat{\pi}_S(\hat{q}^*)}{\hat{\pi}_S(\hat{q}^*)}$ (%)	Supplier's profit share $\frac{\pi_S(q')}{\Pi^o}$ (%)	Efficiency $\frac{\pi_S(q')+\pi_r(q')}{\Pi^o}$ (%)
0.30	0.75	64	76	37	89	100
0.30	0.50	67	78	32	89	100
0.30	0.25	70	80	28	90	100
0.75	0.75	48	72	58	76	100
0.75	0.50	51	74	50	77	99
0.75	0.25	53	75	45	77	99
1.50	0.75	41	72	62	67	98
1.50	0.50	43	74	54	65	97
1.50	0.25	44	74	50	65	96

Note. Demand follows a Gamma distribution with mean 10, $p = 10$, and $v = 0$.

LEMMA 4. *The following hold*:
(i) *there exists a unique* q' *such that* $\pi_r(q') = \hat{\pi}_r(q')$,
(ii) *there exists a unique* q'' *such that* $\pi_s(q'') = \hat{\pi}_s(q'')$,
(iii) q^P *is the unique maximizer of* $\pi_r(q) - \hat{\pi}_s(q)$,
(iv) $q^P = q' = q''$, *and*
(v) $q^P > q^*$.

PROOF. (i) It is sufficient to show that $\pi_r(q) - \hat{\pi}_r(q)$ is unimodal in q, which holds if there is at most one q such that $\pi'_r(q) - \hat{\pi}'_r(q) = 0$:

$$\pi'_r(q) - \hat{\pi}'_r(q) = (p-v)(1-F(q)) \cdot \left(1 - \left[g(q) + \frac{1-F(q^o)}{1-F(q)}(1+j(q)h(q))\right]\right).$$

Because $\pi'_r(0) - \hat{\pi}'_r(0) > 0$, to obtain the needed result it is sufficient to show that the bracketed term is increasing. The first term in the bracket is increasing and the second term is increasing because $j(q)h(q)$ is increasing (from Lemma 1).

(ii) Because total supply chain profit depends only on q, $\Pi(q) = \pi_r(q) + \pi_s(q) = \hat{\pi}_r(q) + \hat{\pi}_s(q)$ which implies $\pi_s(q) - \hat{\pi}_s(q) = \pi_r(q) - \hat{\pi}_r(q)$. Thus, $\pi_s(q) - \hat{\pi}_s(q)$ is unimodal in q because $\pi_r(q) - \hat{\pi}_r(q)$ is unimodal in q. Furthermore, $q' = q''$.

(iii) $\pi_r(q) - \hat{\pi}_s(q)$ has a unique maximum if $\pi_r(q) - \hat{\pi}_s(q)$ is unimodal, which holds if there is at exactly one $q > 0$ such that $\pi'_r(0) - \hat{\pi}'_s(0) > 0$ and $\pi'_r(q) - \hat{\pi}'_s(q) = 0$:

$$\pi'_r(q) - \hat{\pi}'_s(q) = (p-v)\big[(F(q^o) - F(q))(j'(q) - 1) - f(q)(j(q) - q)\big] \quad (15)$$

which simplifies to

$$\pi'_r(q) - \hat{\pi}'_s(q) = (p-v)f(q)q\left[1 - j(q)\frac{1-F(q^o)}{q(1-F(q))}\right]. \quad (16)$$

Table 2 Retailer's Optimal Pull Contract and the Impact of Switching to the Retailer's Preferred Pareto Contract That Leaves the Supplier No Worse Off

		Retailer's optimal pull contract, q^*		Retailer's preferred Pareto contract, q', subject to $\pi_s(q^*) \le \hat{\pi}_s(q')$		
σ/μ	$\frac{p-c}{p-v}$	Retailer's profit share $\frac{\pi_r(q^*)}{\Pi^o}$ (%)	Efficiency $\frac{\pi_r(q^*)+\pi_S(q^*)}{\Pi^o}$ (%)	Retailer's profit increase $\frac{\hat{\pi}_r(q')-\pi_r(q^*)}{\pi_r(q^*)}$ (%)	Retailer's profit share $\frac{\hat{\pi}_r(q')}{\Pi^o}$ (%)	Efficiency $\frac{\hat{\pi}_S(q')+\hat{\pi}_r(q')}{\Pi^o}$ (%)
0.30	0.75	77	84	21	93	100
0.30	0.50	74	83	23	90	100
0.30	0.25	73	83	24	90	100
0.75	0.75	65	76	36	88	99
0.75	0.50	60	77	36	82	99
0.75	0.25	57	77	38	79	99
1.50	0.75	55	71	51	83	99
1.50	0.50	50	74	46	73	97
1.50	0.25	46	74	46	68	96

Note. Demand follows a Gamma distribution with mean 10, $p = 10$, and $v = 0$.

We have $\pi_r'(0) - \hat{\pi}_s'(0) > 0$ because

$$\lim_{q \searrow 0} \frac{1 - j(q)(1 - F(q^o))}{q(1 - F(q))} = F(q^o) > 0.$$

There is one $q > 0$ that satisfies $\pi_r'(q) - \hat{\pi}_s'(q) = 0$ if

$$\frac{j(q)}{q(1 - F(q))} \tag{17}$$

is increasing in q. Differentiate (17) and rearrange terms:

$$\frac{(1 - F(q))[j'(q)q - j(q)] + j(q)f(q)q}{(1 - F(q))^2 q^2}. \tag{18}$$

From Lemma (1), $j'(q)q > j(q)$, so (18) is positive and the needed condition is confirmed.

(iv) It has already been shown that $q' = q''$. Now show $q' = q^P$. After rearranging terms,

$$\pi_r(q) - \hat{\pi}_r(q) = (p - v)q(1 - F(q))\left[1 - \frac{j(q(1 - F(q^o)))}{q(1 - F(q))}\right]. \tag{19}$$

A comparison of (16) with (19) reveals if $\pi_r'(q) - \hat{\pi}_s'(q) = 0$ then $\pi_r(q) - \hat{\pi}_r(q) = 0$. Hence, $q' = q^P$.

(v) From Theorem 3, $\hat{q}^* < q^*$; from Lariviere and Porteus (2001), $\hat{\pi}_s(q)$ is unimodal in q; and from Theorem 2, $\pi_r(q)$ is unimodal in q. Hence, $\pi_r'(\hat{q}^*) - \hat{\pi}_s'(\hat{q}^*) > 0$ (because $\hat{\pi}_s'(\hat{q}^*) = 0$) and $\pi_r'(q^*) - \hat{\pi}_s'(q^*) > 0$ (because $\pi_r'(q^*) = 0$ and $\hat{\pi}_s'(q^*) < 0$). Given that $\pi_r(q) - \hat{\pi}_s(q)$ is unimodal in q (from part (iii) of this lemma) and is increasing at both q^* and $\hat{q}^*$, it follows that $q^P > q^*$. □

According to the next lemma, all of the pull contracts with $q \geq q^P$ survive the push challenge.

Lemma 5. *All pull contracts with $q \geq q^P$ survive the push challenge, i.e., the retailer prefers to prebook zero inventory and depend on the supplier's production for at-once orders rather than to prebook any positive amount.*

Proof. In §4.5 it is shown that there exists a fixed q_s such that the supplier's optimal production is q_s if $y \leq q_s$, otherwise the supplier's optimal production is y. Hence, if the retailer does not prebook, $y = 0$, then the supplier's production is q_s. If $w_1 = w_2$, then clearly $y = 0$ is better for the retailer than $y = q_s$: total availability is the same in either case, but the retailer bears no inventory risk with the former. The issue is whether there exists some prebook $y' > q_s$ such that the retailer's profit is higher with y' than with $y = 0$. (If the retailer prebooks $y' > q_s$, then the supplier's optimal production is y'.) It is sufficient to show that for all $y \geq q_s$ the retailer's optimal prebook quantity is $y = q_s$, i.e., the retailer never wishes to prebook more than q_s.

Given that $w_1(q)$ is increasing in q, $\hat{w}_1(q)$ is decreasing in q and $w_1(0) < \hat{w}_1(0)$, it follows that there exists a unique $\bar{w}$ such that $\bar{w} = w_1(\bar{q}) = \hat{w}_1(\bar{q})$ for some $\bar{q}$. Define $\hat{w}_1^{-1}(w)$ as the inverse function of $\hat{w}_1(q)$ and $w_1^{-1}(w)$ as the inverse function of $w_1(q)$. It follows that $\hat{w}_1^{-1}(w) < w_1^{-1}(w)$ for all $w > \bar{w}$. Note that $\hat{w}_1^{-1}(w_1)$ is the retailer's optimal push quantity and $w_1^{-1}(w_1)$ is the supplier's optimal production quantity with pull. From Lemma 4, $\pi_r(q^P) = \hat{\pi}_r(q^P)$, which implies $w_1(q^P) > \bar{w}$: If the retailer's profit is the same with pull quantity q^P and push quantity q^P, then the push wholesale price must be lower than the pull wholesale price to compensate the retailer for the inventory risk with push. Hence, for all w such that $w_1^{-1}(w) \geq q^P$, it follows that $\hat{w}_1^{-1}(w) < w_1^{-1}(w)$: the retailer's optimal push quantity is less than the supplier's quantity with pull, so those pull contracts survive the push challenge. □

Interestingly, given $q^P > q^*$ (from Lemma 4), it can be shown numerically that the retailer's optimal pull contract may not survive the push challenge, i.e., if the wholesale price is chosen to maximize the retailer's profit in pull mode, the retailer may nevertheless be better off prebooking inventory, thereby operating in push mode.

The primary use for q^P is revealed in the next theorem.

Theorem 6. *The Pareto set among the single wholesale price contracts (push and pull contracts) includes all pull contracts with $q \in [q^P, q^o]$ and all push contracts with $q \in [q^P, q^o]$.*

Proof. Two insights reduce the set of contracts that need to be considered. First, any contract in the Pareto set must belong to either the pull Pareto set or the push Pareto set: By definition, pull contracts that are not in the pull Pareto set are Pareto dominated by some pull contract in the pull Pareto set, and the analogous argument holds for the push contracts. Second, a contract q can only be Pareto dominated by some $q' > q$ because $\Pi(q)$ is increasing in q: $\Pi(q') < \Pi(q)$ implies that some firm must be worse off with q' than q.

Now consider the contracts in the pull Pareto set, $q \in [q^*, q^o)$. (To follow along with the proof it is helpful to refer to Figure 1.) First show that the contracts in $[q^*, q^P)$ are not in the Pareto set and then show that the remaining contracts, $[q^P, q^o]$, are in the Pareto set. Begin with the $q \in [q^*, q^P)$ subset of contracts. Because $\pi_r(q)$ is decreasing for $q \in [q^*, q^P)$ (by Theorem 2), $\pi_r(q^P) = \hat{\pi}_r(q^P)$ (from Lemma 4), and $\hat{\pi}_r(q)$ is increasing (from 5), for any $q \in [q^*, q^P)$ there exists a q' such that $\pi_r(q) = \hat{\pi}_r(q')$. But $\pi_s(q)$ is increasing in q (from 9), so $\pi_s(q) < \hat{\pi}_s(q')$. Hence, the Pareto set does not include any of those contracts. The pull contracts with $q \in [q^P, q^o]$ remain. Given that $\pi_s(q') > \hat{\pi}_s(q')$ for all $q' > q^P$ (from Lemma 4), for any $q \in [q^P, q^o]$ there does not exist a $q' > q$ such that neither firm is worse off, i.e., all pull contracts with $q \in [q^P, q^o]$ are in the Pareto set.

An analogous argument is used to demonstrate that only the push contracts with $q \in [q^P, q^o]$ are in the Pareto set. The remaining issue is whether the pull contracts survive the push challenge, which is confirmed by Lemma 5. □

It is now possible to present several characteristics of this Pareto set. To begin, the supplier's optimal push contract, $\hat{q}^*$, and the retailer's optimal pull contract, q^*, are neither in the Pareto set (because $q^P > q^* > \hat{q}^*$). Hence, the supplier's optimal selling to the newsvendor contract is actually Pareto inferior: The supplier can do better without making the retailer worse off with some pull contract. (Note, this result does depend on the assumption of risk neutrality because, by switching from push to pull, the supplier faces more uncertainty in profit.) Table 1 presents data on how much better the supplier can do. The data demonstrate that the supplier can do much better than the optimal push (i.e., selling to the newsvendor) contract: the supplier's profit can increase between 28% to 50%. Furthermore, supply chain efficiency improves dramatically: efficiency increases from 72%–80% to 96%–100%. Table 2 presents comparable data for the retailer.

While q^* and $\hat{q}^*$ are always Pareto dominated by some contract that shifts the risk allocation in the supply chain and changes the wholesale price, remarkably, those contracts can be Pareto dominated without changing the wholesale price, i.e., merely shifting the risk allocation is sufficient to create a Pareto improvement. (The result is important given that Iyer and Bergen 1997 report there can be strong resistance in industry to changing the wholesale price.) It can be shown numerically that q^* is Pareto dominated by a push contract with the same wholesale price ($w_1(q^*)$) if the coefficient of variation is less than 0.55, and $\hat{q}^*$ is Pareto dominated by a pull contract with the same wholesale price ($\hat{w}_1(\hat{q}^*)$) if the coefficient of variation is less than 0.64. To appreciate this result, suppose the supply chain currently operates with the supplier's preferred push contract, $\hat{q}^*$ (i.e., the supplier's optimal selling to the newsvendor contract). If the coefficient of variation is not too high (but with a threshold of 0.64, neither must it be very low), then the retailer could suggest that the firms switch to Vendor Managed Inventory with consignment (i.e., the supplier chooses the stocking level at the retailer and owns that inventory) and both firms can be better off even if the wholesale price is unchanged, i.e., the retailer can transfer his entire inventory risk onto to the supplier, without changing the wholesale price, and the supplier can be better off even if the wholesale price is the supplier's optimal push contract! (See Bernstein et al. 2002 for a different motivation for Vendor Managed Inventory with consignment.) However, the threshold for which this is a Pareto-improving change decreases if the supplier is risk averse.

Now consider supply chain efficiency. The minimum efficiency of the Pareto set is greater than the minimum efficiency of either the push Pareto set or the pull Pareto set (because $\Pi(q)$ is increasing in q and $q^P > q^* > \hat{q}^*$). Thus, supply chain performance improves if the firms consider both push and pull contracts rather than just considering one or the other type. (This complements the conclusion in Netessine and Rudi 2001a, b that no single supply chain configuration dominates.) Figure 2 reports on the supply chain's minimum efficiency with a Pareto contract. (These data are with a critical ratio of 0.5.) The minimum efficiency of the Pareto set decreases with the coefficient of variation, but the gap relative to the supplier's best push contract, $\hat{q}^*$, or the retailer's best pull contract, q^*, is substantial: With a coefficient of variation of 0.50, efficiency increases from 75% with $\hat{q}^*$, and 81% with q^*, to 95% with q^P. Furthermore, since q^P is only a lower bound on the Pareto set's efficiency, it is likely that the firms will agree to a Pareto contract with even higher efficiency.

The Pareto set also allows each firm to earn any share of the supply chain's optimal profit. Figure 1 illustrates this result. For $q \in [q^P, q^o]$ $\hat{\pi}_s(q)$ is decreasing from $\hat{\pi}_s(q^P)$ to 0 and $\pi_s(q)$ is increasing from $\pi_s(q^P)$ to Π^o, and the analogous result applies to the retailer. In contrast, with just push contracts the supplier's profit is bounded by $\hat{\pi}_s(\hat{q}^*) < \Pi^o$ and with pull contracts the retailer's profit is bounded by $\pi_r(q^*) < \Pi^o$.

As illustrated in Figure 1, in the Pareto set the supplier prefers any pull contract over any push contract and the retailer prefers any push contract to any pull contract. Thus, assuming a Pareto contract is chosen, a firm always earns a higher profit if the firm bears the supply chain's inventory risk. In other words, counter to intuition, a firm should not negotiate with the objective of getting the other firm to bear more inventory risk.

Now suppose the supplier is allowed to choose a contract subject to the constraint that the retailer's profit is no less than some fixed threshold, $\bar{\pi}_r$ (i.e., a participation constraint with a minimum acceptable profit). If $\bar{\pi}_r \le \pi_r(q^P)$ then the supplier chooses a pull contract such that $\pi_r(q) = \bar{\pi}_r$, otherwise the supplier chooses a push contract such that $\hat{\pi}_r(q) = \bar{\pi}_r$. Therefore, the supplier chooses a push contract only if the retailer's minimum acceptable profit is quite high. Furthermore, as the retailer's minimum acceptable profit increases in the range $[0, \pi_r(q^P)]$, supply chain efficiency decreases (because $\pi_r(q)$ is decreasing for

$q \in [q^P, q^o]$). If $\bar{\pi}_r$ is taken as a proxy for the retailer's bargaining power, then increasing the retailer's bargaining power may very well decrease supply chain efficiency. This result contrasts with Lariviere and Porteus (2001) because they only consider push contracts, in which case increasing the retailer's bargaining power always improves supply chain efficiency.

To summarize, when the firms consider both push and pull contracts (i.e., when they are willing to allocate inventory risk to either the retailer or the supplier) then supply chain efficiency improves (often substantially) and the qualitative characteristics of the Pareto set are different than if only push contracts or pull contracts are considered. In particular, the supplier's preferred push contract (the selling to the newsvendor contract) is Pareto inferior and may even be Pareto dominated by the pull contract with the same wholesale price. Furthermore, both push and pull contracts are in the Pareto set, suggesting that neither allocation of inventory risk dominates the other.

4.5. Advance-Purchase Discount Contracts

The objective of this section is to identify and characterize the Pareto set among the push, pull, and advance-purchase discount contracts. The supplier surely produces at least the retailer's prebook, $q \geq y$, so define the supplier's profit function,

$$\pi_s(y, q) = (w_1 - v)y + (w_2 - v)(S(q) - S(y)) - (c - v)q: \tag{20}$$

the first term is the incremental revenue (above the salvage value) from the prebook quantity, the second term is the incremental revenue from at-once sales and the last term is the net production cost. The supplier's profit is concave in q, and the optimal production is $\max\{y, q_s\}$, where q_s satisfies

$$F(q_s) = \frac{w_2 - c}{w_2 - v}. \tag{21}$$

Hence, the supplier's production is independent of w_1 and y whenever $y < q_s$, otherwise the supplier just produces the retailer's prebook (i.e., the supply chain operates in push mode).

The retailer's optimal prebook quantity depends on whether the retailer anticipates $y < q$ or $y = q$. In the latter case the retailer operates as if in push mode. In the former case the retailer's profit is

$$\pi_r(y, q) = -(w_1 - v)y + (p - v)S(y) + (p - w_2)(S(q) - S(y)),$$

where the first term is the net loss on the prebook quantity, the second term is the incremental revenue on prebook inventory and the last term is incremental revenue on at-once orders. $\pi_r(y, q)$ is concave in y, so let $y_r = \arg\max \pi_r(y, q)$, where

$$F(y_r) = \frac{w_2 - w_1}{w_2 - v}. \tag{22}$$

Hence, if $y < q$, the retailer's optimal prebook is independent of the supplier's production. It is now possible to identify the Pareto set of contracts.

Theorem 7. *The Pareto set among the push, pull, and advance-purchase discount contracts includes all advance-purchase discount contracts with $w_2 = p$ and $c \leq w_1 \leq w_2$. The supply chain efficiency of all contracts in the Pareto set is 100% (i.e., the Pareto set coordinates the supply chain) and any division of the supply chain's profit is achievable.*

Proof. If $w_2 = p$, then, from (21) and (2), $q_s = q^o$, i.e., the supplier produces the supply chain optimal production quantity if the retailer prebooks less than q^o. Therefore, if $y < q^o$, then the supply chain's profit is Π^o, efficiency is 100% and, by definition, the supply chain is coordinated. Furthermore, the retailer's profit decreases from Π^o to 0 as w_1 increases from c to p. However, it remains to confirm that the retailer indeed prebooks less than q^o when $w_2 = p$. $\hat{q}^*$ is the retailer's optimal prebook, where $\hat{q}^*$ satisfies (3), conditional on prebooking at least q^o. However, a comparison of (3) with (21) reveals that $\hat{q}^* \leq q^o$ when $w_2 = p$: It is never optimal for the retailer to prebook more than q^o when $w_2 = p$. Given that this set of advance-purchase discounts coordinates the supply chain and arbitrarily allocates its profit, it is easy to confirm that all other contracts with less than 100% efficiency are Pareto dominated. (Strictly speaking, $w_1 = w_2 = p$ may be considered a pull contract and not an advance-purchase discount, and the push contract with $\hat{w}_1 = c$ and $\hat{w}_2 > p$ is also in the Pareto set, and equivalent to the contract $w_1 = c$, $w_2 = p$.) □

The theorem indicates that supply chain coordination and the arbitrary allocation of its profit is achievable with wholesale price contracts. (The surprising result is that arbitrary allocation is possible, because it is well known that marginal cost pricing coordinates the supply chain.) In particular, coordination is accomplished with a subset of the advance-purchase discounts: The retailer earns no profit on at-once orders and the retailer's profit is increasing in the discount depth $(w_2 - w_1)$. The supply chain is coordinated because the supplier has the marginal incentive to provide the correct amount of capacity to the supply chain $(w_2 = p)$ and the retailer has the correct incentive to fill as much demand as the supply chain would. The arbitrary allocation of profit is achieved because there exists a distributive contract term, w_1, that does not distort the firms' incentives to adopt the optimal actions (the supplier's quantity supplied and the retailer's filling of demand). Thus, advance-purchase discounts can be included into the list of newsvendor coordinating contracts (e.g., buybacks, revenue sharing, etc.). As with all coordinating contracts, there are limitations to coordination with advance-purchase discounts which are discussed in the next section. But advance-purchase discounts, like the noncoordinating push and pull contracts, are simple to administer, which is a very attractive feature.

5. Extensions

This section considers three natural extensions: §5.1 includes additional shipping and handling costs for at-once orders; §5.2 allows the retailer to exert costly effort to increase sales; and §5.3 introduces inventory risk with at-once orders due to residual uncertainty. In each case push becomes relatively more attractive than pull. Because advance-purchase discounts are a blend of push and pull, these extensions also reduce their effectiveness. In particular, advance-purchase discounts no longer coordinate the supply chain. However, the main analytical result of each section demonstrates that advance-purchase discounts Pareto dominate a single wholesale price available throughout the season (pull contracts).[5]

5.1. Shipping Costs for At-Once Orders

Consider the model from §3, but now the supplier incurs an additional shipping and handling cost, τ, for each unit in an at-once order. τ does not affect actions and profits with push contracts because there are no at-once orders with push. τ also does not affect the supply chain optimal actions or profits because the optimal supply chain ships all production to the retailer before the season begins to avoid this additional cost. However, τ creates a problem for any contract that physically pulls inventory from the supplier to the retailer in response to at-once orders. Those contracts incur a cost that the optimal supply chain avoids, hence, they cannot coordinate the supply chain. The remainder of this section demonstrates that a Pareto improvement on a pull contract is always possible by pushing at least some inventory to the retailer before the season begins with an advance-purchase discount.

The retailer is clearly never worse off with an advance-purchase discount ($w_1 < w_2$) relative to a pull contract ($w_1 = w_2$) for a fixed w_2. The issue is whether the supplier can benefit from an advance-purchase discount.

[5] These results do depend on the assumption that demand is positive, $F(0) = 0$. If $F(0) > 0$, then the retailer may prebook inventory only if the prebook discount is above a positive threshold. In that case the supplier incurs a cost to induce the retailer to prebook his first unit, which implies that an advance-purchase discount may not be Pareto improving.

Theorem 8. *For a fixed w_2 and $\tau > 0$, if the retailer does not prebook when there is no advance-purchase discount (i.e., when $w_1 = w_2$) then there exists some Pareto-improving advance-purchase discount, i.e., profit increases for both the retailer and the supplier.*

Proof. If the retailer prebooks his entire season supply when $w_1 = w_2$ then an advance-purchase discount is meaningless, i.e., the firms act as if they operate with a push contract with price w_1. Thus, assume w_2 is sufficiently high that the retailer does not prebook. Using the results from §4.5, the retailer's profit function, $\pi_r(y, q)$, and the retailer's optimal prebook quantity, y_r, remain the same because the supplier is assumed to incur the extra shipping cost, τ. The supplier's profit function, $\pi_s(y, q)$, can be used if the at-once wholesale price is replaced with $w_2 - \tau$ in (20) to reflect the supplier's net revenue. This adjustment naturally reduces the supplier's optimal production q_s (see 21). However, if the retailer prebooks less than the supplier's production, the retailer's optimal prebook quantity is still independent of the supplier's production and the supplier's optimal production is still independent of the retailer's prebook quantity and w_1. Therefore, let $y_r(w_1)$ be the retailer's optimal prebook quantity as a function of w_1, assuming $y_r(w_1) < q_s$. From (22), $y_r(w_1) > 0$ for all $w_1 < w_2$, where recall $F(0) = 0$ is assumed. We now have

$$\begin{aligned}\frac{d\pi_s(y_r(w_1), q)}{dw_1} &= \frac{\partial \pi_s}{\partial w_1} + \frac{\partial \pi_s}{\partial y_r}\frac{\partial y_r(w_1)}{\partial w_1}\\ &= y_r(w_1) - \frac{(w_1 - v) - (w_2 - \tau - v)(1 - F(y_r(w_1)))}{(w_2 - v)f(y_r(w_1))},\end{aligned}$$

where from the implicit function theorem, $\partial y_r(w_1)/\partial w_1 = -[(w_2 - v)f(y_r(w_1))]^{-1}$. As $w_1 \to w_2$ we have $y_r(w_1) \to 0$ and

$$\lim_{w_1 \to w_2} \frac{d\pi_s(y_r(w_1), q)}{dw_1} = -\frac{\tau}{(w_2 - v)f(0)} < 0.$$

Hence, π_s is decreasing in w_1 as $w_1 \to w_2$, so $w_1 < w_2$ is surely better than $w_1 = w_2$. □

Advance-purchase discounts help the supplier relative to a pull contract because the supplier need only offer a small discount to induce the retailer to prebook some inventory and by doing so the supplier avoids the almost certain transportation cost on that inventory. While the theorem confirms that pull contracts are not in the Pareto set, the Pareto set no longer is composed of just one contract type. Push contracts suffer from insufficient supply while advance-purchase discounts (with their pull component) suffer from excessive transportation cost. Hence, the Pareto set contains both types of contracts, with their share determined by their relative inefficiencies (i.e., push gains more share as τ increases).

5.2. Retail Effort

The push strategy can also be described as "channel stuffing," i.e., with push the supplier stuffs the retailer with all of the supply chain's inventory. Proponents of this approach often argue that a "stuffed retailer is a captive retailer": A retailer gives the supplier's product a significant amount of marketing attention when the retailer has a significant amount of inventory to sell. Hence, the concern with any contract that reduces the retailer's inventory liability (i.e., any contract with a pull component) is that the retailer will not be sufficiently motivated to increase sales. This section formalizes this argument in favor of push contracts over pull contracts.

Consider the model described in §3 (with $\tau = 0$) but now total demand during the season is ed, where d is the base demand rate and $e \geq 1$ is the amount of effort the retailer exerts to increase sales. Sales effort is costly: Let $\phi(e)$ be the cost of effort level e and assume $\phi(1) = 0$, $\phi'(e) > 0$, and $\phi''(e) > 0$. At the time of the prebook order the base demand rate is unknown with distribution function F, but the realized base demand rate d is observed at the start of the selling season (before the at-once orders). The retailer chooses e after observing the base demand rate d but before submitting the at-once order. Thus, the retailer's effort choice depends on the observed base demand, the retailer's prebook amount and the amount of inventory available from at-once orders. Note, with this effort model the effectiveness of one unit of effort depends on the base demand, i.e., it is easier to sell a "hit" than a "dog." This model is reasonable if early season sales provides an accurate assessment of total season sales. With a push contract the retailer sells $\min\{ed, y\}$ units during the season whereas the retailer sells $\min\{ed, q\}$ units when at-once orders are allowed.

There are number of other papers that study retail effort (e.g., Netessine and Rudi 2001a), but almost all of them assume effort is chosen before demand is realized. Krishnan et al. (2001) is an exception. Their model of effort is identical to this one with the exception that they allow the cost function to depend also on the base demand, d.

Due to the inclusion of retail effort, the integrated supply chain's analysis is different than in §4. The supply chain's profit function after observing demand d is

$$\Pi(q, e, d) = \min\{ed, q\}(p - v) + vq - \phi(e).$$

Let $e_o(d)$ be the optimal effort level, $\phi'(e_o(d)) = d(p - v)$, without a supply constraint. The retailer's total unconstrained sales, $e_o(d)d$, is increasing in d. Therefore, define $d_o(q)$ to be the demand level such that $e_o(d_o(q)) = q$, i.e., if $d \in [d_o(q), q)$ then the optimal effort increases sales to exactly q. The supply chain's profit is now

$$\Pi(q) = -c + \int_0^{d_o(q)} f(x)\Pi(q, e_o(x), x)\, dx + \int_{d_o(q)}^{q} f(x)\Pi(q, q/x, x)\, dx + (1 - F(q))q(p - v).$$

$\Pi(q)$ is strictly concave in q, so let q^o be the unique optimal production quantity.

With a push contract $\Pi(y, e, d)$ is the retailer profit function after observing demand: The retailer exerts the supply chain optimal effort because inventory is a sunk cost to the retailer when choosing effort, just as it is for the supply chain. Therefore, supply chain efficiency with push is less than 100% due to the limited availability of inventory ($y < q^o$) but not due to a lack of retail effort.

With a pull contract the retailer's profit after observing demand is

$$\pi_r(q, e, d) = \min\{ed, q\}(p - w_2) - \phi(e).$$

Define $e_r(d)$ as the retailer's optimal effort if there is no supply constraint, $\phi'(e_r(d)) = d(p - w_2)$. The retailer's total sales with pull is increasing in d, so define $d_r(q)$ such that $e_r(d_r(q)) = q$. It is easy to see that $e_r(d) < e_o(d)$ whenever $w_2 > v$: If the at-once price is above the salvage value then the retailer's optimal effort with pull is less than the supply chain's optimal effort. Hence, supply chain efficiency with pull is less than 100% for two reasons: limited availability and less than optimal retail effort. It follows that any contract with a pull component cannot coordinate the supply chain. Netessine and Rudi (2001a) find a comparable result: In their model, drop shipping dampens a retailer's incentive to invest in customer-acquisition costs.

A pull contract is particularly vulnerable to situations in which demand is significantly lower than supply: The supply chain may exert effort to sell at least a portion of the inventory, but with pull the retailer is unconcerned with the supplier's stockpile. This suggests a little bit of push, in the form of an advance-purchase discount, could help pull.

Theorem 9. *In the model with sales increasing effort, for a fixed w_2, if the retailer does not prebook when there is no advance-purchase discount (i.e., when $w_1 = w_2$) then there exists some Pareto-improving advance-purchase discount.*

Proof. Only an outline of the proof is provided, with details available from the author. Similar to Theorem 8, w_2 must be sufficiently high so that the retailer does not prebook his entire supply, i.e., the retailer does not operate as if a push contract is in effect. Next, show that the retailer prebooks a positive amount (less than the supplier's production) for any $w_1 < w_2$. Furthermore, that prebook quantity depends on w_1 but does not depend on q. Now show that the derivative of the supplier's profit function is negative at $w_1 = w_2$, i.e., a discount increases the supplier's profit. □

While advance-purchase discounts no longer coordinate the supply chain (because of the pull component), as in the transportation cost model, this does not mean that advance-purchase discounts should be ignored. The Pareto set surely contains a blend of push contracts and advance-purchase discount contracts: The push contracts provide optimal effort, but are less effective at providing adequate supply, while the reverse holds for the advance-purchase discounts.

5.3. Inventory Risk with At-Once Orders

At-once orders occur during the season in §3, and only after demand is observed. Hence, there is no residual uncertainty or inventory risk associated with the at-once orders. However, if the lead time to deliver at-once orders is sufficiently long or if shipping costs require the consolidation of at-once orders, then the retailer may only be able to make a single at-once order at the start of the selling season. The retailer's demand information at the time of the at-once order may be more precise than at the time of the prebook order, but the retailer may nevertheless face residual uncertainty when submitting his at-once order. This is the issue in this section.

Suppose between the prebook and the at-once order the retailer receives information, which is also called a signal, regarding demand. Let ξ be the signal observed, $\xi \in [\xi_l, \infty)$, let G be the distribution function for ξ and let $F(q, \xi)$ be the distribution of demand conditional on the demand signal. Both distribution functions are continuous, increasing and differentiable. Furthermore, assume $\partial F(q, \xi)/\partial \xi < 0$ (i.e., the demand distribution is stochastically increasing in the signal) and $F(0, \xi) = 0$ for all ξ (no matter the signal, there surely is some demand). The distribution function before the demand signal remains F, where

$$F(q) = \int_{\xi_l}^{\infty} F(q, x) g(x)\, dx.$$

The quality of the signal can range from meaningless ($F(q, \xi) = F(q)$ for all ξ) to perfect ($F(q, \xi) = 1$ if $q \geq \xi$, otherwise $F(q, \xi) = 0$), where the intermediate cases are referred to as noisy signals. The models in the previous sections have a perfect signal. With a meaningless signal the distinction between push and pull vanishes: The retailer faces the same inventory risk with the prebook order as with the at-once order, so the retailer submits a single order at the lowest wholesale price. In other words, with a meaningless signal the supply chain can only operate with a push contract. The remainder of this section considers noisy signals.

Actions and profits with a push contract are unaffected by the presence of a noisy signal: With a push contract there are no actions taken after the signal is observed (the only action is the prebook quantity, which is before the signal is observed). Hence, the analysis of the push contracts remains the same even in the presence of a noisy signal.

Pull contracts are affected by the quality of the signal because with a pull contract the retailer's at-once order occurs after the signal is observed. With a perfect signal the retailer's optimal at-once order is the same as the supply chain's optimal at-once order, i.e., order enough inventory to cover the observed demand. With a noisy signal the retailer's at-once action is no longer optimal: the retailer's optimal at-once order raises his inventory upto y_2, where

$$F(y_2, \xi) = \frac{p - w_2}{p - v},$$

whereas the supply chain's optimal at-once order ships all inventory remaining at the supplier to the retailer (the cost of inventory is sunk, so the inventory should be shipped to the retailer in the hope of earning an additional sale). Hence, supply chain coordination cannot be achieved with a contract that has a pull component in the presence of a noisy signal. (Coordination requires $w_2 = v$, but then the supplier surely earns a negative profit.) Again, the solution to the problem with pull is to introduce some push with an advance-purchase discount.

Theorem 10. *In the model with residual uncertainty for at-once orders, for a fixed w_2, if the retailer does not prebook when there is no advance-purchase discount (i.e., when $w_1 = w_2$) then there exists some Pareto-improving advance-purchase discount.*

Proof. The outline of the proof is identical to the outline in Theorem 9. A detailed proof is available from the author. □

6. Discussion

This paper studies a supply chain with a long production lead time but fast replenishments between the supplier and the retailer. Hence, while there is a single production opportunity, as is often assumed in the supply chain contracting literature, there are potentially multiple opportunities to transfer inventory between the supplier and the retailer, which introduces the possibility of distinguishing between early (prebook) and late (at-once) orders both in terms of the wholesale price charged as well as the demand information available to the firms. Depending on the contractual terms, the supply chain's inventory risk can be allocated to either the retailer (push contract) or the supplier (pull contract) or shared between them (advance-purchase discount).

This research demonstrates that the allocation of inventory risk matters for supply chain efficiency even if firms are risk neutral. For example, consider a "supplier selling to a newsvendor" situation in which the retailer bears the inventory risk and the supplier has chosen her optimal wholesale price. While that contract is always Pareto inferior if the firms are willing to change the wholesale price, an even more remarkable result demonstrates that contract can be Pareto inferior without changing the wholesale price: Merely shifting the inventory risk from one firm to another can improve supply chain efficiency and increase profit at both firms. Furthermore, if the firms are willing to share inventory risk via advance-purchase discounts, then supply chain coordination is achievable with any division of the supply chain's profit.

Just as there are limitations with other contractual forms, there are limitations with advance-purchase discounts. Due to the pull component in the advance-purchase discount, those contracts no longer coordinate the supply chain when there are extra shipping costs for at-once orders, when the retailer can exert costly effort to increase sales or when there is inventory risk even with at-once orders. Nevertheless, advance-purchase discounts should not be ignored because they Pareto dominate the single wholesale price pull contract: Pushing at least some inventory onto the retailer can be good for all of the firms in the supply chain.

It is also interesting to speculate on other extensions of the model. With multiple retailers pull becomes more attractive than push due to risk pooling at the supplier but pull also introduces the issue of inventory rationing (see Netessine and Rudi 2001b). Advance purchase discounts lose the power to coordinate the supply chain if the retailer sets the retail price because then the supplier cannot extract all of the retailer's margin on at-once orders. However, with a second (albeit more expensive) production opportunity advance-purchase discounts may coordinate the supply chain and leave the retailer with a positive margin on at-once orders. (The supplier's initial production decision no longer depends on the retail price, but rather on the difference between the cost of late versus early production.) Finally, in some cases the retailer may have better demand information than the supplier even before the supplier must commit to her production decision. In that case the advance-purchase discount could be used to communicate information from the retailer to the supplier, thereby improving the supplier's production decision. (With a realistic model of that situation the supplier must not be able to perfectly infer the retailer's demand information from the retailer's prebook order, otherwise an arbitrarily small discount would be sufficient to communicate the retailer's information. Hence, a proper analysis of that model is beyond the scope of this research.)

Although a number of results are developed in this research with respect to advance-purchase discounts, the most significant and surprising result comes from identifying the Pareto set among the simplest contracts: the minimum efficiency of the single wholesale price contracts (push and pull contracts) is much higher than previously thought; e.g., efficiency of 95% even with a coefficient of variation of 0.50. (Of course, this assumes the mitigating effects on pull contracts discussed in §5 are either not present or not strong.) Furthermore, since that is only a bound, it is quite likely the firms would agree to a contract with even higher efficiency. Hence, the incremental value of any coordinating contract, which almost surely costs more to implement, monitor, and administer than a single wholesale price contract, has been exaggerated.

Acknowledgments
For their helpful consultation, thanks is extended to the reviewers, the associate editor, Fangruo Chen, Martin Lariviere, Mark Ferguson, Serguei Netessine, Terry Taylor, and the MSOM Conference seminar participants at Cornell University. The original title of this paper was "The allocation of inventory risk and advanced purchase discounts in a supply chain."

References

Anand, K., R. Anupindi, Y. Bassok. 1999. Strategic inventories in procurement contracts. Working paper, University of Pennsylvania, Philadelphia, PA.

Anupindi, R., Y. Bassok. 1999. Centralization of stocks: Retailers vs. manufacturer. *Management Sci.* **45**(2) 178–191.

Anupindi, R., Y. Bassok, E. Zemel. 2001. A general framework for the study of decentralized distribution systems. *Manufacturing Service Oper. Management* **3**(4) 349–368.

Barnes-Schuster, D., Y. Bassok, R. Anupindi. 1998. Coordination and flexibility in supply contracts with options. Working paper, University of Southern California, Los Angeles, CA.

Bernstein, F., A. Federgruen. 2002. Coordination mechanisms for supply chains under price- and service competition. Working paper, Duke University, Durham, NC.

Bernstein, F., F. Chen, A. Federgruen. 2002. Vendor managed inventories and supply chain coordination: The case with one supplier and competing retailers. Working paper, Duke University, Durham, NC.

Cachon, G. 2003. Supply chain coordination with contracts. S. Graves, T. de Kok, eds. *The Handbook of Operations Research and Management Science: Supply Chain Management.* Kluwer, Amsterdam, The Netherlands.

Cachon, G., M. Lariviere. 2000. Supply chain coordination with revenue sharing: Strengths and limitations. *Management Sci.* Forthcoming.

Cachon, G., M. Lariviere. 2001. Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Sci.* **47**(5) 629–646.

Cachon, G., S. Netessine. 2003. Game theory in supply chain analysis. D. Simchi-Levi, S. D. Wu, Z. M. Shen, eds. *Supply Chain Analysis in the eBusiness Era*. Kluwer Academic Publishers, Amsterdam, The Netherlands. Forthcoming.

Dana, J. 1998. Advance-purchase discounts and price discrimination in competitive markets. *J. Political Econom.* **106**(2) 395–422.

Dong, L., N. Rudi. 2001. Who benefits from transshipment? Exogenous vs. endogenous wholesale price. *Management Sci.* Forthcoming.

Donohue, K. 2000. Efficient supply contracts for fashion goods with forecast updating and two production modes. *Management Sci.* **46**(11) 1397–1411.

Erhun, F., P. Keskinocak, S. Tayur. 2000. Spot markets for capacity and supply chain coordination. Working paper, Carnegie Mellon, Pittsburgh, PA.

Ferguson, M. 2003. When to commit in a serial supply chain with forecast updating. *Navel Res. Logist.* **50**(8) 917–936.

Ferguson, M., G. DeCroix, P. Zipkin. 2002. When to commit in a multi-echelon supply chain with partial information updating. Working paper, Duke University, Durham, NC.

Fisher, M., A. Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Oper. Res.* **44** 87–99.

Fisher, M., K. Rajaram, A. Raman. 2001. Optimizing inventory replenishment of retail fashion products. *Manufacturing Service Oper. Management* **3**(3) 230–241.

Iyer, A., M. Bergen. 1997. Quick response in manufacturer-retailer channels. *Management Sci.* **43**(4) 559–570.

Krishnan, H., R. Kapuscinski, D. Butz. 2004. Coordinating contracts for decentralized supply chains with retailer promotional effort. *Management Sci.* Forthcoming.

Lariviere, M., E. Porteus. 2001. Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing Service Oper. Management* **3**(4) 293–305.

McCardle, K., K. Rajaram, C. Tang. 2002. Advance booking discounts under retail competition. Working paper, University of California at Los Angeles, Los Angeles, CA.

Netessine, S., N. Rudi. 2001a. Supply chain structures on the Internet: Marketing-operations coordination under drop shipping. Working paper, University of Pennsylvania, Philadelphia, PA.

Netessine, S., N. Rudi. 2001b. Supply chain choice on the Internet. Working paper, University of Pennsylvania, Philadelphia, PA.

Özer, O., W. Wei. 2002. Strategic commitment for optimal capacity decision under asymmetric forecast information. Working paper, Stanford University, Stanford, CA.

Pasternack, B. 1985. Optimal pricing and returns policies for perishable commodities. *Marketing Sci.* **4**(2) 166–176.

Randall, R., S. Netessine, N. Rudi. 2002. Inventory structure and internet retailing: An empirical examination of the role of inventory ownership. Working paper, University of Utah, Salt Lake City, UT.

Rudi, N., S. Kapur, D. Pyke. 2001. A two-location inventory model with transhipment and local decision making. *Management Sci.* **47**(12) 1668–1680.

Tang, C., K. Rajaram, A. Alptekinoglu, J. Ou. 2001. The benefits of advance booking discount programs: Model and analysis. *Management Sci.* Forthcoming.

Taylor, T. 2002a. Coordination under channel rebates with sales effort effect. *Management Sci.* **48**(8) 992–1007.

Taylor, T. 2002b. Sale timing in a supply chain: When to sell to the retailer. Working paper, Columbia University, New York.

Tomlin, B. 2000. Capacity investments in supply chain: Sharing-the-gain rather than sharing-the-pain. *Manufacturing Service Oper. Management*. Forthcoming.

Tsay, A. 1999. Quantity-flexibility contract and supplier-customer incentives. *Management Sci.* **45**(10) 1339–1358.

Weng, K., M. Parlar. 1999. Integrating early sales with production decisions: Analysis and insights. *IIE Trans. Scheduling Logist.* **31** 1051–1060.

Xie, J., S. Shugan. 2001. Electronic tickets, smart cards and online prepayments: When and how to advance sell. *Marketing Sci.* **20**(3) 219–243.