Contracting to Assure Supply: How to Share Demand Forecasts in a Supply Chain

Gérard P. Cachon • Martin A. Lariviere
The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104
Kellogg Graduate School of Management, Northwestern University, Evanston, Illinois 60208

Forecast sharing is studied in a supply chain with a manufacturer that faces stochastic demand for a single product and a supplier that is the sole source for a critical component. The following sequence of events occurs: the manufacturer provides her initial forecast to the supplier along with a contract, the supplier constructs capacity (if he accepts the contract), the manufacturer receives an updated forecast and submits a final order. Two contract compliance regimes are considered. If the supplier accepts the contract under forced compliance then he has little flexibility with respect to his capacity choice; under voluntary compliance, however, he maintains substantial flexibility. Optimal supply chain performance requires the manufacturer to share her initial forecast truthfully, but she has an incentive to inflate her forecast to induce the supplier to build more capacity. The supplier is aware of this bias, and so may not trust the manufacturer's forecast, harming supply chain performance. We study contracts that allow the supply chain to share demand forecasts credibly under either compliance regime.
(*Game Theory; Coordination; Signaling; Asymmetric Information*)

1. Introduction

Each level of a supply chain makes decisions that have ramifications throughout the entire system. The quality of a given decision depends on what the decision maker knows. As a result, the dissemination of accurate information is critical for the supply chain to operate effectively. Credibility is a key factor in the exchange of information: Will and should the receiver of information trust the veracity of the reported information. While credibility is easily established in some cases, it is often more elusive. This is especially true when the informed party has an incentive to distort her message to influence the receiver's actions.

This paper studies the exchange within a supply chain of a particularly important piece of information, demand forecasts. A manufacturer offers a supplier a contract to build capacity for a specialized component for which the supplier is the only source. Along with the contract, the manufacturer provides the supplier an initial demand forecast. Assuming the contract is acceptable, the supplier then builds capacity. After capacity is built the manufacturer observes demand and submits a final order. Finally, the supplier produces as much of the order as possible given the available capacity.

In such a setting, the manufacturer may give the supplier an excessively optimistic initial demand forecast so as to induce him to build more capacity. The manufacturer does not pay to install capacity and strictly prefers having more available in case demand happens to be high. For his part, the supplier is aware of the manufacturer's incentive and may view a rosy forecast with skepticism, building a cautious amount of capacity. Of course, if the demand forecast really is encouraging, the supplier's limited capacity may lead to numerous lost sales, hurting both parties.

0025-1909/01/4705/0629$5.00
1526-5501 electronic ISSN

Many supply chains wrestle with these issues, often unsuccessfully. In the personal computer industry, distributors frequently have better demand information than the manufacturers because they are closer to customers. To better manage their inventories, manufacturers would prefer the most accurate information possible. Unfortunately, they often suspect their distributors of submitting "phantom orders," forecasts of high future demand that do not materialize (Zarley and Damore 1996). Complicating matters, it is difficult to accuse a distributor of lying. A distributor might have truly expected high demand, but random events could still lead to a low demand realization. Since the manufacturers do not trust the orders they receive from the distributors, there effectively is no exchange of information.

The car market may be less volatile than the PC market, but limited supplier capacity constraining downstream production, possibly a result of poor forecast sharing, is still an issue. Even a powerful firm such as General Motors may be held back by its suppliers. In ramping up for the 1994 model year, GM lost nearly two months of production of the Buick Roadmaster because a new supplier could not provide an adequate supply of ashtrays and glove-compartment doors (Suris and Templin 1993). Production of the 1998 Corvette was limited by the transmission supplier's capacity (Mateja 1998). Suppliers may be wise to avoid spending heavily to serve assemblers. In 1999 GM abruptly canceled two new models (supposedly because of poor consumer reactions). Suppliers who had been preparing for the launch suddenly faced getting no return on their investment (Pryweller 1999).

The aerospace industry has similarly been a victim of component capacity shortages and bad forecasts. In 1997, Boeing experienced difficulties in increasing production because of capacity shortages at its 3,000 parts suppliers; company officials acknowledge approximately 500 "notable" part shortages for the 747 alone (Cole 1997a). Even suppliers who attempted to expand capacity did so with some trepidation. One supplier executive, commenting on a major expansion at his firm, observed "We're putting a lot of trust in the Boeing Co." (Cole 1997b). That trust was not necessarily well-deserved. Within a year the Asian financial crisis occurred. Despite initially insisting that the crisis would not have a significant impact on sales, Boeing was forced to cut output substantially (Biddle 1998).

Our analysis begins with the assumption that the initial demand forecast is known to all, which we call the *full information* case. Under full information, the manufacturer seeks to maximize her expected profit subject to gaining the supplier's acceptance. Before evaluating that objective, one must specify how much leeway the supplier has in setting capacity under a given contract. Most of the supply chain contracting literature assumes that the supplier must build enough capacity to satisfy any final order allowed by the contract. Implicitly that assumes that the supplier's capacity decision is verifiable and enforceable by the courts. Failing to fill the manufacturer's final order results in a penalty so stiff, and imposed with such certainty, that not covering the order is not even a consideration.

In reality, enforcement is more complex. There are many determinants of capacity (e.g., worker skills, equipment maintenance, scheduling policies), and they interact in subtle ways. If the supplier does not fill the manufacturer's order, the courts may not be able to distinguish between a supplier who properly reserved capacity but failed to fill the order for reasons beyond his control, and a supplier who willfully ignored the manufacturer's request. Even if verification were possible, it may not be economically viable. If the manufacturer's cost of proving the supplier negligent is sufficiently high relative to her potential reward, the contract may as well be unenforceable; the supplier would recognize that the manufacturer would never pursue the case and consequently would ignore any threat of punishment.

We explore the importance of contract enforcement with two compliance regimes. Under *forced compliance* the supplier is liable for the manufacturer's maximum final order and thus must build sufficient capacity to cover any possible order allowed by the contract. *Voluntary compliance* represents the opposite extreme. Capacity decisions are not verifiable, and the supplier therefore cannot be forced to fill an order. Instead he chooses the capacity that maximizes his profits given the terms of trade. The contract must assure acceptance (as under forced compliance) but must also induce the supplier to build the capacity the manufacturer desires. Voluntary compliance is more complex analytically than forced compliance; the supplier's capacity choice is clear under the latter, but nontrivial under the former.

After the full information case, we consider supply chain performance under *asymmetric information*, assuming that the manufacturer has superior initial demand information. In addition to designing supplier incentives, the manufacturer must now attempt to convey—or "signal"—her demand forecast credibly. Credibility is particularly important if she expects a large market. If he dismisses the manufacturer's forecast as overly optimistic, the supplier will either refuse the contract or build inadequate capacity. While many papers investigate the value of sharing demand information in a supply chain (e.g., Aviv and Federgruen 1998, Cachon and Fisher 1996, Chen 1998, Gavirneni et al. 1999, Moinzadeh 1999), they generally assume that information is always shared truthfully. In our setting, the manufacturer has an incentive to inflate the demand forecast, so the credibility of the forecast is a legitimate concern. We examine how the manufacturer may share her forecast under both compliance regimes.

The next section details the model, and §3 relates it to the literature. The full information case is analyzed in §4 and the asymmetric information case in §5. Section 6 discusses the importance of the model's main assumptions, and §7 concludes. Proofs are in Appendix A.

2. The Model

A manufacturer sells a single product that has uncertain demand, D. Let $F(x)$ be the continuous and differentiable demand distribution and $f(x)$ its density function. $\overline{F}(x) = 1 - F(x)$. We assume that $F(x) = 0$ for all $x \le 0$, and $f(x) > 0$ for all $x > 0$. (The analysis can be extended to any distribution with support of the form $[a, b)$ for $0 \le a < b \le \infty$.)

The manufacturer contracts with a single supplier who is the only source for a customized component. The single source assumption is reasonable when the component has no function other than as part of the manufacturer's product and requires specializes skills and assets to produce. The supplier must install capacity K at a cost of $c_K > 0$ per unit before either party observes demand. Once demand, d, is realized, the manufacturer submits a final order to the supplier, where the contract may limit the set of admissible final orders. The supplier fills the final order subject to component production not exceeding K. Let $S(K)$ be expected sales given an available capacity K:

$$S(K) = E[D - (D - K)^+]$$
$$= \int_0^K x f(x)\,dx + K\overline{F}(K) = K - \int_0^K F(x)\,dx,$$

where $[x]^+ = \max\{x, 0\}$ and the last expression is obtained via integration by parts. It costs the supplier $c_p > 0$ per unit to convert raw capacity into usable components. The manufacturer assembles the component into her product, which sells for a fixed retail price of r per unit, $r > c_K + c_p$. To highlight contracting for component capacity, we normalize to zero all other costs as well as the salvage value of the end product. Both firms are risk neutral.

Note that if one decision maker controlled both the manufacturer and supplier, the system would face a standard newsvendor problem. The objective for the standard integrated system benchmark $\Pi^I(K)$ can thus be written as:

$$\Pi^I(K) = -c_k K + (r - c_p) S(K).$$

Since $S(K)$ is an increasing, concave function, $\Pi^I(K)$ is concave. The optimal capacity choice K^I is the unique solution to

$$S'(K^I) = \overline{F}(K^I) = c_k/(r - c_p). \tag{1}$$

Let Π^I denote $\Pi^I(K^I)$. To confirm that a contract coordinates the decentralized system (i.e., total profits sum to Π^I), it is sufficient to verify that the decentralized system chooses capacity K^I and initiates final production only after demand is observed.

To move from the centralized to decentralized setting, we must specify how the parties interact. We assume the game between the manufacturer and the supplier proceeds according to the following sequence of events. First, the manufacturer observes the demand distribution, $F(x)$, which we interpret as her forecast of demand. Next, the manufacturer offers a contract to the supplier along with an initial order quantity for the component. This is a take-it-or-leave-it offer, a reasonable assumption in markets such as automobiles or aerospace in which large downstream players dominate. The supplier accepts any contract with an expected profit no lower than his opportunity cost, which is normalized to zero. Following acceptance, the supplier sets capacity K as allowed by the contract and manufacturer's order. Demand is then observed and final production takes place.

The supplier's willingness to accept the contract and provide capacity depends on his beliefs about the market, i.e., on the *supplier's forecast* of demand. We consider two scenarios. In the full information case, the supplier also observes the demand forecast, $F(x)$, so his forecast coincides with the manufacturer's. In the asymmetric information case, the supplier does not observe the demand forecast. Hence, the supplier must develop his own forecast which may differ from the manufacturer's. That process is described in greater detail in §5.

2.1. Contracts

The manufacturer offers the supplier a contract consisting of firm commitments and options. Let $m \geq 0$ be the number of firm commitments and $o \geq 0$ the number of options. The manufacturer pays the supplier w_m per firm commitment and w_o per option when the supplier accepts the contract. The manufacturer also pays the supplier w_e per option exercised that the supplier actually delivers. Call $m+o$ the manufacturer's initial order and let q be the manufacturer's final order. As their names suggest, firm commitments and options restrict the manufacturer's final order, $m \leq q \leq m+o$. Since the final order is submitted after observing demand, $q = m + [\min\{d-m, o\}]^+$. Thus, $q-m$ is the number of options exercised.

For simplicity, we assume throughout that K never exceeds $m+o$, so the supplier never sets a capacity level above the manufacturer's maximum order. We consider two compliance regimes. Under forced compliance, the supplier who accepts the contract must provide sufficient capacity to cover the manufacturer's maximum order. Hence, $K = m+o$. Under voluntary compliance, this enforcement provision does not hold so the supplier is free to choose $K \leq m+o$. If $K < q$, some units are not delivered.

A special case of our contract is worthy of mention. Consider voluntary compliance and suppose the manufacturer offers $w_m > r$, $w_o = 0$, and $w_e > 0$. Here, firm commitments are prohibitively expensive but an infinite number of options can be purchased for free. The manufacturer's initial order is consequently meaningless, and the manufacturer's final order will equal realized demand d. With this contract the manufacturer only pays the supplier $w_e > 0$ per unit delivered. We term this specialization a *price-only contract* since a single contract parameter will determine the supplier's action. For this case, we refer to the exercise price as the wholesale price and drop the subscript e.

Given a capacity K and firm commitments m, let $P(K, m)$ be the supplier's expected production quantity assuming $K \leq m+o$,

$$P(K, m) = \min\{K, m\} + [S(K) - S(m)]^+.$$

The first term is the production to satisfy the firm commitments and the second term is the production to satisfy exercised options. Suppose the manufacturer purchase m firm commitment and o options while the supplier sets a capacity of K. Then $T(K, m, o)$ is the expected transfer payment from the manufacturer to the supplier,

$$T(K, m, o) = w_m m + w_o o + w_e[S(K) - S(m)]^+.$$

If the supplier accepts the contract, the supplier's expected profit is

$$\pi(K, m, o) = T(K, m, o) - c_K K - c_p P(K, m)$$

and the manufacturer's expected profit is

$$\Pi(K, m, o) = rP(K, m) - T(K, m, o).$$

Note that for notational convenience price terms have been suppressed from the argument list of the above functions.

3. Literature Review

Several contract types have been considered for supply chains with two firms, stochastic demand and a short time horizon. Our contract captures many of these schemes because one can set prices for our contract such that it yields the same transfer payment between the firms. Consider, for example, a buy-back contract (Pasternack 1985, Donohue 1996, Emmons and Gilbert 1998) in which the manufacturer submits an initial order and may return any portion of that order for a partial credit after observing demand. To replicate these payments set $w_m > r$, $w_o > 0$ and $w_e > 0$. The manufacturer buys only options, paying $w_o + w_e$ for each unit and receiving a partial credit equal to w_e for returned units (options that are not exercised). In the quantity-flexibility (QF) contract (Tsay 1999) the manufacturer submits an initial order and may return a portion of that order for a full credit after observing demand. If we again model returns as unexercised options, we can mimic the cash flows of a QF contract by setting $w_m > 0$, $w_o = 0$, and $w_e = w_m$. A backup agreement (Eppen and Iyer 1997) is then a QF contract with a partial return credit: $w_m > 0$, $w_o > 0$, and $w_e = w_m - w_o$. Brown and Lee (1998a,b) study pay-to-delay contracts that generate the same transfer payments as our contract. Barnes-Schuster et al. (1998) study a contract that is similar to ours (i.e., it contains firm commitments and options), but their model includes two periods of demand.

Several papers study price-only contracts. Lariviere and Porteus (2000) study such terms in a setting that is essentially the reverse of ours: An upstream vendor sets the wholesale price for a retailer facing a newsvendor problem. Gerchak and Wang (1999) and Gerchak and Gurnani (1998) consider setting wholesale prices when there are multiple complementary suppliers. Van Mieghem (1999) studies an outsourcing model in which the salvage value of excess capacity is uncertain and the transfer price is set a priori. Caldentey and Wein (1999) consider a longer horizon. Tomlin (1999) works with essentially the same model that we consider but examines a quantity premium contract. The manufacturer pays $w(q)$ for the qth delivered unit, where $w(q)$ is an increasing, piecewise-linear function. The equivalent price in our contract, w_e, is constant.

Our work differs from the supply chain contracting literature mentioned above along two broad dimensions. First, we have the downstream firm setting all contract parameters. It is generally assumed that either the upstream firm makes the contract offer or the contract parameters are exogenous. Second, we consider two compliance regimes, and to the best of our knowledge, we are the only paper to consider both compliance regimes in the same model. Parameter-rich contracts (e.g., QF contracts or backup agreements) usually assume forced compliance while price-only contracts often assume voluntary compliance.

The second dimension that distinguishes our work is how we treat asymmetric information. Previous work in the supply chain literature has not focused on an informed party trying to convey information credibly. Anand and Mendelson (1997) study the location of decision rights when agents in a supply chain are unable to share their information and act as a team. In Cachon and Lariviere (1999), each of several retailers orders from a single supplier with limited capacity, knowing only his own demand forecast but not the forecasts of the other retailers. In that model the supplier chooses an allocation mechanism to divide available capacity when the total quantity ordered exceeds capacity. Ha (1999), Corbett (1999), and Porteus and Whang (1999) consider contracting problems in which the party offering the contract is less informed. Hence their emphasis is on inducing information revelation as opposed to assuring information credibility. Brown (1999) studies information sharing when the firm receiving the information assumes the information is correct. Lee et al. (1997) recognize that phantom orders can be quite detrimental to a supply chain, but they do not explore solutions analytically. There are numerous papers which study forecasting and inventory management in nonstrategic settings (e.g., Aviv 1999, Graves et al. 1998, and Toktay and Wein 1999.)

Our forecast-sharing game falls within a class of economic problems known as signaling models. The field dates from Spence's (1973) analysis of productive workers' signaling their abilities by enduring higher levels of education. More recent research has considered sharing demand information (Chu 1992, Desai and Srinivasan 1995, Lariviere and Padmanabhan

1997) when a channel faces a deterministic demand curve. We show that there are alternative signaling instruments when demand is stochastic.

4. The Full Information Scenario

In the full information scenario, because both firms observe the demand distribution, there is no need for forecast sharing. However, the manufacturer must still design a contract that maximizes her profit given the supplier's anticipated actions.

4.1. Forced Compliance

Under forced compliance, the supplier's capacity decision is trivial. He must set K to cover the largest possible final order, $K = m + o$. The supplier's only meaningful decision is whether to accept the contract, which he does if he expects a nonnegative profit. Contract acceptance is consequently the only constraint the manufacturer must consider, and the analysis begins with characteristics of the manufacturer's expected profit function, $\Pi(K, m, o)$.

Lemma 1. *$\Pi(m+o, m, o)$ is jointly concave in m and o. For fixed prices, the manufacturer's optimal decisions (m^*, o^*) are determined from the following*:

$$\text{If } w_m \geq w_o + w_e, \quad m^* = 0, \quad \bar{F}(o^*) = \frac{w_o}{r - w_e};$$

$$\text{if } w_o + w_e > w_m \geq \frac{w_o r}{r - w_e}, \quad \bar{F}(m^*) = \frac{w_m - w_o}{w_e},$$

$$\bar{F}(o^* + m^*) = \frac{w_o}{r - w_e};$$

$$\text{if } \frac{w_o r}{r - w_e} > w_m, \quad \bar{F}(m^*) = \frac{w_m}{r}, \quad o^* = 0.$$

If the flexibility of options is free ($w_o + w_e \leq w_m$), the manufacturer uses only options. If that flexibility is too expensive ($w_o r/(r - w_e) > w_m$), she forgoes options. For intermediate values, she purchases firm commitments up to the point that the expected cost of an option is equal to the cost of a firm commitment. (If both m^* and o^* are positive, $w_o + w_e\bar{F}(m^*) = w_m$.)

The results of Lemma 1 let the manufacturer design her optimal contract (i.e., one that maximizes her profits). From the following theorem, she can coordinate the channel without purchasing firm commitments. Thus decentralization imposes no loss on the system. Better still (at least for her), she captures the supply chain's entire profit.

Theorem 1. *Suppose the manufacturer purchases no firm commitments, purchases K^I options and offers the following prices*:

$$w_o = c_K - \varepsilon S(K^I)/K^I, \quad w_e = \varepsilon + c_p, \quad w_m \geq w_o + w_e$$

for $\varepsilon \in [-c_p, \min\{r - c_p, c_K K^I/S(K^I)\})$. The supplier accepts that contract and just recovers his opportunity cost, $\pi(K^I, m^*, o^*) = 0$. The manufacturer earns the integrated system profit, Π^I.

As in Pasternack (1985), there is a continuum of contracts that coordinate the system. The similarities end there. In Pasternack (1985), the coordinating contracts differ on how the system profit is divided, so the offering party prefers a single coordinating contract. Here, the supplier's profit is zero and the manufacturer's profit is Π^I for any of the coordinating contracts. The contracts differ only in the variability of the firms' earnings, which is irrelevant since both are assumed to be risk neutral.

The manufacturer captures all system profits because the forced compliance regime leaves the supplier with no capacity choice. The terms of trade serve only to allocate the profits; consequently, the supplier is left with his minimally acceptable return. Since the contract allows multiple price parameters, the manufacturer has several degrees of freedom in designing the contract and multiple contracts maximize the manufacturer's objective.

Theorem 1 leaves open the question of whether there are any coordinating contracts that include firm commitments. In fact, there are none. A firm commitment possibly forces the decentralized supply chain to undertake actions that the integrated system would avoid. After observing demand, d, the integrated system never produces more than d, but the decentralized system might produce more if $m > d$. Unless firm commitments result in lower production costs (as in Donohue 1996), such actions are wasteful.

4.2. Voluntary Compliance

Since the supplier is free to choose $K < m + o$ under voluntary compliance, the contract design process must consider both whether the supplier will accept the terms and what capacity he will choose after acceptance. Our analysis begins with the latter.

The supplier's capacity and production costs increase with K, but revenues increase only when $K > m$. Thus, he either builds no capacity ($K = 0$) or selects a capacity greater than the minimum purchase level ($K > m$). In the latter case, $\pi(K, m, o)$ is concave in K, so the supplier's optimal positive capacity level, K', satisfies the first-order condition

$$\bar{F}(K') = \frac{c_K}{w_e - c_p}. \tag{2}$$

K' is the global optimum if $\pi(K', m, o) \geq \pi(0, m, o)$. Since $\pi(K', m, o) - \pi(0, m, o)$ is decreasing in m, there exists an $\bar{m}(K')$ such that $K' > 0$ is the supplier's globally optimal capacity only if $m \leq \bar{m}(K')$. In other words, the supplier builds capacity only when he is not given too many firm commitments.

Now consider the manufacturer's contract design problem. Offering a positive purchase price for options (i.e., $w_o > 0$) has no impact on the supplier's capacity choice, and buying firm commitments can only lead to lower capacity. In other words, buying options and firm commitments transfer wealth to the supplier without increasing capacity. Thus, the manufacturer neither buys firm commitments, $m = 0$, nor pays for options, $w_o = 0$. In effect, voluntary compliance relegates the manufacturer to offering a price-only contract and the only relevant contract parameter is the wholesale price w.

From (2), the supplier builds capacity K if offered the wholesale price $w(K)$,

$$w(K) = \frac{c_K}{\bar{F}(K)} + c_p.$$

We say the manufacturer offers K to the supplier as a shorthand for saying the manufacturer offers wholesale price $w(K)$ with the intention of inducing a capacity of K.

Given that $m = 0$ and $w_o = 0$, the manufacturer's profit function can now be written in terms of the capacity the manufacturer wants the supplier to build,

$$\Pi(K) = (r - w(K))S(K).$$

In what follows, we assume that $w(K)$ is convex (inducing a large capacity is an increasingly expensive proposition) to ensure a well-behaved (concave) profit function for the manufacturer. (Chowdrhry and Jegadeesh (1994) make a similar assumption.) Since

$$w'(K) = (w(K) - c_p)h(K),$$

where $h(K)$ is the failure rate of the demand distribution (Barlow and Proschan 1965), and obvious condition for a convex $w(K)$ is that $h(K)$ is increasing for all K such that $\bar{F}(K) > 0$, a property known as an increasing failure rate (IFR). The normal and the uniform are both IFR, as are the gamma and Weibull subject to parameter restrictions (Barlow and Proschan 1965). IFR, however, is only sufficient. For example, $\bar{F}(x) = x^{-k}$ for $x \geq 1$ and $k \geq 1$ yields a convex $w(K)$ even though its failure rate is always decreasing.

Theorem 2. *If $w(K)$ is convex, $\Pi(K)$ is concave in $K \geq 0$ and the manufacturer's optimal offer is the wholesale price $w(K^*)$, where K^* solves*

$$\bar{F}(K^*) = \frac{c_K}{r - c_p}\left(1 + \frac{f(K^*)}{\bar{F}(K^*)^2}S(K^*)\right). \tag{3}$$

Comparing (1) and (3), it is apparent that the decentralized system provides less capacity than the integrated system ($K^* < K^I$), which means that the supply chain profit is not maximized (i.e., the channel is not coordinated).

The next two theorems relate the manufacturer's profit and the optimal wholesale price to the demand distribution.

Theorem 3. *Consider two demand distributions, F_1 and F_2, such that $\bar{F}_1(x) \geq \bar{F}_2(x)$ for all x. The manufacturer's expected profit is higher in the larger market (i.e., Market 1).*

It is reasonable that a firm earns more in a larger market, but the theorem says nothing about how the optimal wholesale price changes with the demand distribution. For that we impose some additional structure. Suppose demand D is given by

$$D = \delta + \lambda X, \tag{4}$$

where X is a random variable with distribution F and $\delta \geq 0$ and $\lambda > 0$. The distribution of D is then given by

$$F(x|\delta, \lambda) = F\left(\frac{x-\delta}{\lambda}|0, 1\right) = F\left(\frac{x-\delta}{\lambda}\right).$$

We now consider how $w(K^*)$ varies with the scale parameter, λ, and the shift parameter, δ.

Theorem 4. *Suppose that demand is given as in* (4).

(a) *Suppose $\delta = 0$, and consider two values of λ, λ_1 and λ_2. Let K_i and w_i be the optimal capacity and wholesale price for λ_i. Then, $K_1 = \lambda_1 K_2/\lambda_2$ and $w_1 = w_2$.*

(b) *Suppose $\lambda = 1$, and consider two values of δ such that $\delta_1 > \delta_2$. Let K_i and w_i be the optimal capacity and wholesale price for δ_i. Then $K_1 < K_2 + \delta_1 - \delta_2$, and $w_1 < w_2$.*

(c) *There exists a $\bar{\delta}$ such that for all $\delta > \bar{\delta}$, the optimal capacity choice is $K^* = \delta$ and $w(K^*) = c_K + c_p$.*

(d) *Assume that for any λ, δ is chosen to keep the coefficient of variation of the demand distribution constant at ψ: $\delta(\lambda) = \lambda(\sigma/\psi - \mu)$, for $\psi \leq \sigma/\mu$. Then the optimal wholesale price $w(K^*)$ is independent of λ and increasing in ψ.*

Theorem 4 implies that the manufacturer's optimal wholesale price is not directly related to the size of the manufacturer's market. Instead, it is the coefficient of variation of the manufacturer's demand that is critical. As the manufacturer's demand becomes more variable the manufacturer must pay a higher wholesale price to induce the same level of capacity. On the other hand, if the variability of demand is sufficiently small, the manufacturer settles for only serving certain demand of δ and offers only a wholesale price of $c_K + c_p$. These results are generally the mirror opposites of those in Lariviere and Porteus (2000).

5. Asymmetric Information and Sharing Forecasts

We now assume the manufacturer is better informed about demand than the supplier. She should therefore present her forecast when offering her contract. The supplier, however, should not necessarily believe everything he hears. The supplier chooses a larger capacity the larger he believes the market to be and the manufacturer's profit is increasing in the supplier's capacity. Hence, the manufacturer has an incentive to provide a rosy forecast in the hope that the supplier provides more capacity.

A rational supplier consequently views an optimistic forecast skeptically unless it is backed by contract terms that assure its credibility. The contract design problem of a manufacturer expecting high demand has an added degree of complexity. To share—or signal—her information credibly, she must eliminate any incentive for a manufacturer expecting lower demand to "mimic" her. Consequently, a manufacturer anticipating a large market must design terms a pessimistic manufacturer would never want to offer.

5.1. Forecast Sharing Model

We model asymmetric information regarding demand by assuming only the manufacturer knows some parameter θ of the distribution such that for all $\theta' \leq \theta$, $F(x|\theta') > F(x|\theta)$ for $x > 0$ and $F(0|\theta') \geq F(0|\theta)$. Thus, the market is stochastically increasing in θ. The parameterized family (4) is an obvious example of such a distributional family since increasing either the shift or the scale parameter leads to a stochastically larger market.

We work with purely scaled distributions. Let D_θ be demand given $\theta > 0$, where $D_\theta = \theta X$ and X is a random variable with distribution function F. We allow two values for θ, H and L with $H > L$. Let $F_\theta(x) = F(x|\theta)$ for $\theta = H, L$ with similar notation for $f_\theta, \overline{F}_\theta$ and $S_\theta(K)$. Given θ, let Π^I_θ be the integrated channel's optimal profit and let K^I_θ be the integrated system optimal capacity. The manufacturer observes θ before offering her contract and thus knows the true demand distribution (but not realized demand) before the supplier must choose capacity. The observed value of θ is the *type* of the manufacturer so a high type faces demand D_H and a low type faces demand D_L.

Prior to observing the contract, the supplier assigns the probability $\rho \in (0, 1)$ that the true demand distribution is D_H and the probability $1-\rho$ that the distribution is D_L. The *prior* probability ρ represents the supplier's best assessment of demand before learning any additional information. His initial forecast is a mixture of the two possible distributions and therefore is less accurate than the manufacturer's forecast.

Since the manufacturer knows the true demand distribution, the supplier may possibly learn (or infer) additional information from the contract the manufacturer offers. No information is learned if the supplier expects both types of manufacturers to offer the same contract (and only that contract). This is referred to as a *pooling equilibrium*. However, the supplier does learn information if the supplier expects that certain contracts are offered only by high types and other contracts are offered only by low types. This is referred to as a *separating equilibrium*. Since in a separating equilibrium the different types offer different contracts, the supplier learns the true demand forecast upon observing the offered contract and rationally updates ρ to either 1 or 0. Thus, the manufacturer's forecast is not credible in a pooling equilibrium, but it is credible in a separating equilibrium. Our focus is on forecast sharing, so we concentrate on separating equilibria, but discuss pooling in §6.

A separating equilibrium partitions the space of feasible contracts, Z, into two sets, Z_H and Z_L ($Z = Z_H \cup Z_L$, and $\varnothing = Z_H \cap Z_L$). The set Z is fixed exogenously. For example, we allow contracts with linear prices for firm commitments and options and explicitly rule out nonlinear price schedules. The sets Z_H and Z_L represent the supplier's beliefs regarding the equilibrium actions of both types of manufacturers. If the supplier is offered a contract in Z_θ, then he updates his belief to assume that the manufacturer is type θ, i.e., D_θ is certainly the demand distribution. This *posterior* belief is very important because the supplier evaluates his expected profit from the contract and picks his capacity with the assumption that the demand distribution is certainly D_θ. Hence, we require that the supplier's beliefs be rational. In equilibrium, the supplier's posterior belief should always be correct. Such beliefs are justified if a type θ manufacturer's maximum expected profit from offering a Z_θ contract is greater than her expected profit from offering a Z_τ contract, $\tau \neq \theta$. Only a type θ manufacturer would rationally offer a Z_θ contract.

To evaluate her expected profit from a given action, the manufacturer assumes the supplier behaves as if the manufacturer is a type θ if she offers a contract from Z_θ and as if she were a type $\tau \neq \theta$ if she offers a contract from Z_τ. This approach assures that the manufacturer correctly anticipates how the supplier behaves given his equilibrium beliefs. Formally, for a partition of contracts to represent rational supplier beliefs, we require:

$$\max_{z \in Z_\theta} \Pi_\theta(z, \theta) \geq \max_{z \in Z_{\tau \neq \theta}} \Pi_\theta(z, \tau), \tag{5}$$

where $\Pi_\theta(z, \tau)$ is a type θ manufacturer's expected profit from offering contract z when the supplier assumes demand is D_τ. We make the mild assumption that $\Pi_\theta(z, \tau)$ is strictly increasing in τ for any feasible contract, so any manufacturer is better off the larger the supplier believes the market will be.

Even though the equilibrium specifies the supplier's posterior belief for each contract the supplier could observe, in equilibrium the supplier only observes one of two possible contracts. These contracts, $\{z_L^e, z_H^e\}$, maximize each type's profit within their allowed set of equilibrium contracts

$$z_\theta^e = \arg\max_{z \in Z_\theta} \Pi_\theta(z, \theta),$$

where we assume that z_θ^e is unique for expositional simplicity. Let z_θ^* be the type θ manufacturer's full information optimal contract,

$$z_\theta^* = \arg\max_{z \in Z} \Pi_\theta(z, \theta).$$

Multiple partitions of the feasible contract set satisfying (5) may exist. We focus on the following equilibrium partition. Let Z_H be the solution to the following program.

$$\begin{aligned} Z_H = \arg\max_{z \in Z} & \ \Pi_H(z, H) \\ \text{s.t.} \quad & \Pi_L(z, H) \leq \Pi_L(z_L^*, L) \\ & \Pi_H(z, H) \geq \max_{z \in Z} \Pi_H(z, L) \end{aligned} \tag{6}$$

Let Z_L be all other feasible contracts. In this equilibrium partition, the supplier assumes the manufacturer is a high type if he is offered contract Z_H but assumes the manufacturer is a low type if any other contract is offered. It is straightforward to confirm that this partition satisfies (5). The low-type manufacturer can do no better than offering z_L^* if she is taken to be a low type, so the first constraint ensures that a low-type manufacturer is not willing to choose the contract in Z_H. The second constraint ensures the high-type manufacturer prefers to offer the Z_H contract than to offer her best contract if she were taken to be a low type.

In this equilibrium, the low-type manufacturer offers the supplier z_L^* and the high-type manufacturer offers Z_H. In fact, it must be that the low-type manufacturer offers the supplier z_L^* in any equilibrium. The supplier never assumes a high type offers z_L^* because then the low type would surely be willing to mimic and z_L^* is certainly the low type's best offer if she is taken to be a low type. On the other hand, there are equilibrium partitions which do not include the solution to (6). These equilibria, however, do not survive reasonable refinements on the supplier's beliefs.[1] Finally, there are equilibria partitions that add additional contracts to Z_H, but these contracts would not be observed in practice because the high-type manufacturer would be better off offering the solution to (6). Hence, even though other separating equilibrium exist, it is sufficient to consider just our proposed equilibrium.

Since z_L^* is identified in the full information section, it only remains to identify Z_H for each compliance regime.

5.2. Forced Compliance

From §4 there are contracts that coordinate the channel and allow the manufacturer to earn the integrated channel profit. Clearly the high-type manufacturer would prefer signaling with one of those contracts. Therefore, let the contracts defined in Theorem 1 be the set of feasible contracts, i.e., the type θ manufacturer purchases K_θ^I options and offers

$$w_o(\theta) = c_K - \frac{\varepsilon}{K_\theta^I} S_\theta(K_\theta^I), \quad w_e(\theta) = \varepsilon + c_p,$$
$$w_m(\theta) > w_o(\theta) + w_e(\theta).$$

for $\varepsilon \in [-c_p, \hat{\varepsilon}(\theta))$ where $\hat{\varepsilon}(\theta) = \min\{r - c_p, c_K K_\theta^I / S_\theta(K_\theta^I)\}$. Let $\Pi_\theta(K, \varepsilon)$ be a type θ manufacturer's expected profit from offering one of these contracts assuming the supplier accepts the contract. (We drop the supplier's belief from the notation because a supplier who accepts the contract is required to build K regardless of his belief.)

Any of the high-type coordinating contracts satisfies the second constraint in (6) because the high-type manufacturer can do no better than earn Π_H^I. The first constraint in (6) requires that the low-type manufacturer does not wish to offer a high-type contract. In equilibrium the low type earns Π_L^I,

$$\Pi_L^I = (r - c_p) S_L(K_L^I) - K_L^I c_K, \tag{7}$$

but when she offers a high-type contract she earns

$$\Pi_L(K_H^I, \varepsilon) = (r - c_p) S_L(K_H^I) - K_H^I c_K + \varepsilon (S_H(K_H^I) - S_L(K_H^I)) \tag{8}$$

for a feasible ε. Her mimicking profit is increasing in ε, so there is a unique $\bar{\varepsilon}$ that leaves the low type indifferent between truthfully reporting her forecast (and earning Π_L^I) and mimicking the high-type manufacturer,

$$\bar{\varepsilon} = \frac{c_K(K_H^I - K_L^I) - (r - c_p)(S_L(K_H^I) - S_L(K_L^I))}{S_H(K_H^I) - S_L(K_H^I)}.$$

We now show how the high type can convey her forecast credibly under forced compliance.

Theorem 5. *There exists a separating equilibrium under forced compliance in which the high-type manufacturer offers a contract specified in Theorem 1 with* $\min\{\hat{\varepsilon}(\theta), \bar{\varepsilon}\} > \varepsilon \geq -c_p$. *The high-type manufacturer's expected profit is* Π_H^I *and the supplier's expected profit is zero.*

Forced compliance is powerful. It allows a high-type manufacturer to signal for free. She must choose from a (possibly) restricted subset of the contracts available under full information, but the limitation is trivial since she is indifferent across all of them. Note that the contract $w_o = c_K$ and $w_e = c_p$ is always credible. These terms are tantamount to taking over the supplier; the manufacturer dictates exactly what capacity should be chosen and sets compensation so that the supplier earns zero for all demand realizations.

5.3. Voluntary Compliance

Under voluntary compliance, a high type can no longer dictate the supplier's capacity choice. The contract she offers must consequently play two roles, convincing the supplier that her high demand forecast is genuine and inducing him to build adequate capacity.

[1] See Kreps (1990). These refinements argue that it is unreasonable to assume the solution to (6) would be offered by a low type; therefore that solution should not be included in Z_L.

Since under full information the manufacturer is relegated to offering a price-only contract, we begin with that structure. Some additional notation is needed to account for the information asymmetry between the players. Let $w_\theta(K) = c_K/\overline{F}_\theta(K) + c_p$. Thus, if the manufacturer offers the wholesale price $w_\theta(K)$ and the supplier believes for sure that demand is D_θ, then the supplier constructs capacity K. Let $\Pi_\theta(K, \tau)$ be a type θ manufacturer's expected profit with a price-only contract if the supplier believes the manufacturer is type τ and therefore offering the wholesale price $w_\tau(K)$:

$$\Pi_\theta(K, \tau) = (r - w_\tau(K))S_\theta(K).$$

As with full information, we assume $w_\tau(K)$ is convex. As a result, $\Pi_\theta(K, \tau)$ is convex in K for all θ and τ. Let $K_\theta(\tau)$ be the type θ manufacturer's optimal capacity offer (in a price-only contract) given that the supplier believes the manufacturer is type τ:

$$K_\theta(\tau) = \arg\max_K \Pi_\theta(K, \tau).$$

We first consider whether the high type can signal with her optimal full information wholesale price, $w_H(K_H^*)$. From Theorem 4, $w_H(K_H^*) = w_L(K_L^*)$, so that contract clearly does not provide a credible signal; a low-type manufacturer would happily pay her full information price to receive a high type's capacity. Something must be added to the contract. Suppose the high type also commits to paying the supplier a lump sum, A, upon contract acceptance. To be credible, A must be sufficiently large that a low type would never want to pay it, the first constraint in (6),

$$\Pi_L(K_H^*, H) - A \le \Pi_L(K_L^*, L),$$

and the high type must be willing to pay the lump sum, the second constraint in (6),

$$\Pi_H(K_H^*, H) - A \ge \Pi_H(K_H(L), L).$$

By the next lemma, there indeed exists a feasible lump sum. Hence, there exists a separating equilibrium in which the high-type manufacturer offers $w_H(K_H^*)$ and the minimum lump sum that satisfies the above constraints.

Lemma 2. $\Pi_H(K_H^*, H) - \Pi_H(K_H(L), L) > \Pi_L(K_H^*, H) - \Pi_K(K_L^*, L)$.

The amount A can literally be a lump sum, or the manufacture can purchase K_H^* options at a price of $w_o = A/K_H^*$; under voluntary compliance purchasing options has no impact on the supplier's capacity decision, making them strategically equivalent to a fixed payment.

Although it permits forecast sharing, a lump sum is an expensive *signaling instrument* because it reduces the profits of both types at the same rate. Information is communicated only because the high-type manufacturer has deeper pockets. From the high type's perspective, the ideal signaling instrument is very painful (i.e., costly) to a low-type manufacturer but essentially harmless (i.e., costless) to herself. The wholesale price is one instrument that has some of the differential impact we seek. At a constant wholesale price, the marginal value of additional capacity is greater for a high type because that manufacturer has a higher probability of actually using the extra capacity. In fact, at least some signaling with the wholesale price is always better than signaling with just a lump sum.

Theorem 6. *The high type's profit is higher when she signals by requesting $K > K_H^*$ and possibly offering a lump sum than when she offers only a lump sum.*

The wholesale price signal is effective because requesting a little bit more than K_H^* is free to a high-type manufacturer (since $\partial\Pi_H(K_H^*, H)/\partial K = 0$) but expensive to the low-type manufacturer (since $\partial\Pi_L(K_H^*, H)/\partial K < 0$, i.e., $K_L(H) < K_H^*$). Thus, ordering more capacity than K_H^* is at least initially a cheaper way to purchase credibility than a lump sum payment. Interestingly, it may not be wise to purchase credibility only with capacity. At a sufficiently high capacity level it may be more expensive, on the margin, for a high-type manufacturer to purchase more capacity than a low type because inducing a higher capacity requires paying a higher wholesale price for all units purchased.

Firm commitments are also more effective than a lump sum. To explain, define $\Pi_\theta(K, m, A)$ as the expected profit of a type θ manufacturer who the supplier believes is type H when the manufacturer offers lump sum A and buys m firm commitments at $w_m = w_H(K)$:

$$\Pi_\theta(K, m, A) = rS_\theta(K) - w_H(K)(S_\theta(K) - S_\theta(m) + m) - A, \quad m \le \bar{m}_H(K).$$

(Recall, the upper bound on m ensures that the supplier actually builds some capacity.) At first glance, firm commitments appear to be merely an upfront promise to pay the supplier some amount. However, whereas a lump sum is costly for all demand realizations, a firm commitment is *ex post* costless if demand exceeds m. That is more likely for the high type, so it is always cheaper for her to make a firm commitment than to pay a lump sum.

To see that firm commitments are a better way to signal let $\Pi_\theta^x = \partial\Pi_\theta(K, m, A)/\partial x$ for $x = m, A$ and for $0 < m < \bar{m}_H(K)$, and consider:

$$\frac{\Pi_L^m}{\Pi_H^m} = \frac{F_L(m)}{F_H(m)} > 1 = \frac{\Pi_L^A}{\Pi_H^A}, \tag{9}$$

where the inequality follows from first-order stochastic dominance. We compare the ratio of the marginal rates because the high type is interested in dissipating a fixed amount of profit for the low type at a minimal cost to herself. The relationship in (9) implies that firm commitments are a more efficient signaling instrument than lump sum transfers since they provide a greater "bang for the buck." A dollar spent by the high type on a lump sum costs a mimicking low type a dollar. If, however, the high type agrees to firm commitments to the point that she increases her own costs by a dollar, then the costs of a mimicking low type go up by more than a dollar.

Theorem 7. *Suppose the high type can signal her information with a capacity request, firm commitments and/or a lump sum payment. The following characterize her optimal contract*:

(a) *The high-type manufacturer requests more capacity than she desires, i.e.,* $K > K_H^*$.

(b) *If the optimal contract has* $m < \bar{m}_H(K)$, *then there is no lump sum payment,* $A = 0$.

(c) *If the optimal contract has* $0 < m$, *the corresponding capacity and wholesale price are less than they would be if the manufacturer could not offer firm commitments.*

Since the high-type manufacturer always signals by distorting the capacity she induces, $K > K_H^*$, signaling is never free under voluntary compliance. However, while the need to purchase forecast credibility lowers the high-type manufacturer's profit, signaling with a lump sum and/or capacity always increases total supply chain profit because more capacity is built. (The optimal signaling capacity is less than K_H^I, since $w_H(K_H^I) = r$, and greater than K_H^*.) In this example, reducing the bargaining power of one player (by requiring the manufacturer to signal) improves the performance of the supply chain. (Lariviere and Porteus (2000) and van Mieghem (1999) make similar observations in different settings.) If firm commitments are used to signal, the consequences for supply chain profits are ambiguous; more capacity is provided but the supply chain may commit to wasteful production.

The third point in the theorem indicates that a manufacturer using firm commitments to signal pays a lower wholesale price than one that does not signal with firm commitments. It generally has been observed that firm commitments allow a firm to obtain per-unit discounts (e.g., Anupindi and Bassok 1999), but in our model those discounts are an endogenous result rather than an exogenous assumption. However, the theorem does not state whether the optimal contract always includes firm commitments. We conjecture that there exist examples in which it is sufficient to purchase increased capacity. Firm commitments may be cheaper than lump sum payments, but they are not free, whereas purchasing additional capacity is initially free. Nevertheless, firm commitments are a very attractive instrument.

5.4. An Example with Exponential Demand

We provide an illustrative example of the previous section's results. Demand is exponentially distributed with mean $\theta_H = 10$ for the high type and $\theta_L = 5$ for the low. Assume $r = 1, c_K = 0.1$, and $c_p = 0.1$. Figure 1 displays expected profit as a function of the induced capacity under a price-only contract and full information for a low and a high-type manufacturer, $\Pi_L(K, L)$ and $\Pi_H(K, H)$, respectively. The high type earns a higher profit at any capacity level and prefers to induce more capacity. Also displayed is the profit

Figure 1 Manufacturer's Profit as a Function of Induced Capacity with Exponential Demand Under Voluntary Compliance

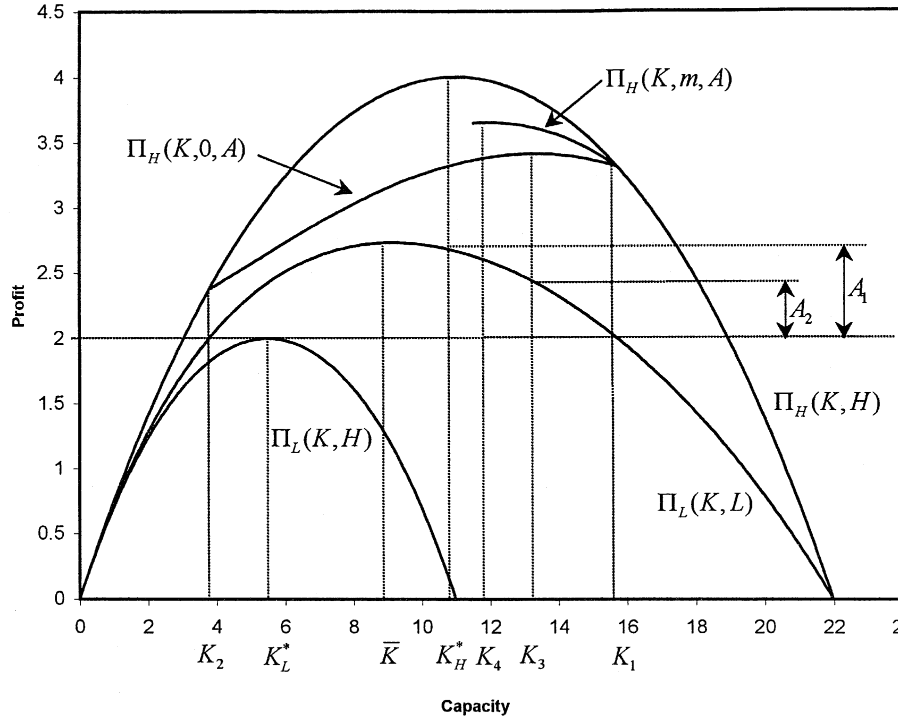

function of a mimicking low type with a price-only contract, $\Pi_L(K, H)$, which is maximized at $K_L(H)$.

For the high type to signal, she must offer a contract that the low type does not want to mimic. At the high type's optimal price-only contract, K_H^*, the mimicking low type earns $A_1 = 2/3$ more than if she were to report her forecast truthfully. Thus, the high-type manufacturer must offer a contract that dissipates at least A_1 in profit of the mimicking low type. An immediate option is to offer the optimal price-only contract with a lump-sum payment A_1, yielding an expected profit of 3.33 for the high type.

A cheaper signal is available. Suppose the high type chooses to signal only with her capacity choice (equivalently, her wholesale price). One option is to offer K_1. As proven in Theorem 7, $K_1 > K_H^*$. While K_1 avoids a lump-sum payment altogether, the combination of offering K_3 and a payment of A_2 with no fixed commitments (labeled as $\Pi_H(K, 0, A)$) does better. For $K > K_3$ the high-type manufacturer's expected profit falls at a faster rate than the low type's profit, making a lump-sum payment preferable because it lowers each type's profit at the same rate.

If the high-type manufacturer is willing to use firm commitments, she can do better still. With capacity K_4, $m = 2.65$ and no lump-sum the manufacturer earns an expected profit of 3.65, while her best signal without fixed commitments earns 3.41. (Her best signal without a lump sum, K_1, earns 3.29.) It is possible to show that the optimal fixed commitment does not violate the $\bar{m}_H(K)$ constraint, and consequently no lump-sum transfers is part of the contract. (Only profits that do not violate that constraint are displayed as $\Pi_H(K, m, A)$.)

6. Discussion

We now consider the implications of altering some of our model's assumptions.

A Second Source. Suppose there is a second source for the component that can provide an unlimited quantity at a cost of c_2 per unit. It is natural to assume that the second source is less desirable than the primary source, $c_2 > c_K + c_p$, but still a viable source, $r > c_2$. When $d \leq K$, the manufacturer purchases d

units from her primary supplier. When $d > K$, K units are purchased from the primary source and $d - K$ units are purchased from the second source. Thus, the manufacturer always fills demand d and earns, at a minimum, $\underline{\Pi} = (r - c_2)E[D]$ in expected profit. The supply chain earns an additional $c_2 - (c_K + c_p)$ per unit in profit on the first K units demanded but zero additional profit per unit for demand greater than K. That is the same profit profile as a supply chain with only one supplier and a "retail price" equal to c_2. Thus, the single supplier model can be used to analyze a supply chain with two suppliers: the two source supply chain has a "retail price" of c_2 and a manufacturer that enjoys the a priori profit $\underline{\Pi}$ (instead of a zero a priori profit in the single supplier case).

No Forecast Sharing. With forced compliance, the high-type manufacturer can signal for free, so credible forecast sharing can be expected. With voluntary compliance, the high type can signal at some cost. If being taken as a low type is the alternative to signaling, we have shown that the high type prefers to signal. But what if the supplier assumes the observed contract provides no information regarding the manufacturer's type? The supplier's posterior beliefs about the chance of demand being high is equal to his prior ρ. If a high type does not signal, the supplier might treat her as an average manufacturer, which may not be so bad, especially if the supplier has a high prior assessment of high demand. Thus, there might exist one or more pooling equilibria in which the supplier assumes that both types offer the same terms. The analysis of such equilibria is complex. (Interested readers are referred to Kreps 1990.) Since our focus is on the exchange of demand forecasts, we defer the analysis of pooling equilibria to future research.

Multiple Demand States. It is possible to extend the model to include multiple demand states, $\Theta = \{\theta_1, \theta_2, \ldots, \theta_n\}$, where $\theta_i < \theta_j$ for all $i < j$, but the analysis is significantly more complex. Consider the highest type, θ_n. To signal successfully, this manufacturer would have to convince the supplier that none of the other types, $\theta_1, \ldots, \theta_{n-1}$, would be willing to mimic that type's contract. θ_1 will almost surely be easier to eliminate from consideration than θ_{n-1}. Hence, it is possible that the highest type would separate herself from extremely low types but would pool with relatively high types. In other words, imperfect signals must be introduced into the analysis. We leave the details to future research.

Imperfect Demand Signal. Suppose the manufacturer does not observe demand after capacity is built but rather receives an imperfect signal of demand that is more accurate than her initial forecast. Now the manufacturer faces a newsvendor problem constrained by the supplier's capacity decision when choosing her final order quantity. The manufacturer's optimal final order is stochastic and increasing in the observed demand signal. Final production should be still delayed until after observing the demand signal, making firm commitments unattractive under full information. Hence, the model is analytically more complex, but qualitatively unchanged.

Compliance Regime. While most of the contracting literature assumes forced compliance, we have also studied the opposite extreme of voluntary compliance. An alternative is a hybrid compliance regime. Each side begins with an estimate of the cost to enforce the contract in court and a probability assessment of success. Voluntary compliance assigns probability zero, whereas forced compliance assigns probability one. A hybrid mechanism would assume an intermediate value. As long as success in court is not assured, the manufacturer will have to create an incentive structure that induces the supplier to provide the desired capacity. In this sense, we regard forced compliance as a "knife-edge" assumption since it completely eliminates the supplier's capacity decision.

Other Demand Distributions. We have chosen to work with scaled distributions since scaled families are both general and tractable. All that is essential for an interesting problem, however, is that a high-type manufacturer's demand first order stochastically dominates a low type's. Shifted distributions with $D_\theta = \theta + X$, for $\theta \geq 0$ and X a nonnegative random variable, are another natural alternative. The only significant change to the analysis would be that signaling under voluntary compliance would either always or never include firm commitments because demand is certain to be greater than θ. Recall that firm commitments must be less than a bound $\bar{m}_H(K)$. If $L < \bar{m}_H(K)$, than the high type can increase the costs of a mimicking low type at no cost to herself by choosing $m \in (L, \min\{\bar{m}_H(K), H\})$; firm commitments are then always part of a separating equilibrium. On the other hand, if $L \geq \bar{m}_H(K)$, any level of firm commitments that is costly to the low type results in the supplier providing no capacity. Thus they are never offered.

7. Conclusion

This paper demonstrates that (1) the contract compliance regime significantly impacts the analysis and outcome of the supply chain contracting game we consider and (2) it is always in the interest of a manufacturer with a high demand forecast to share her forecast with the supplier, but sharing her forecast credibly may be costly (in the sense that her profit would be higher if full information prevailed).

The supply chain contracting literature often assumes (either implicitly or explicitly) forced compliance. While possibly appropriate in some settings, forced compliance is not an innocuous assumption. Under forced compliance, a downstream leader (our manufacturer) essentially dictates all operating decisions because the supplier has no leeway in setting capacity once he accepts the contract. In some sense forced compliance violates the original premise for studying supply chain contracting: that no one firm controls all supply chain actions. Under forced compliance the following firm is left only with veto power (whether to accept or reject the contract) but that power buys the firm nothing. The contract designer sets the terms of the contract to give the follower only his minimally acceptable return. Since that is a relatively easy task, numerous contracts satisfy the contract designer's objective. In contrast, under voluntary compliance neither firm controls all supply chain decisions. Consequently, the analysis of the game is more complex, and the power of the contract designer is greatly diminished.

The power of forced compliance is particularly evident under asymmetric information. In that case, despite the supplier's uncertainty with respect to the demand distribution and the manufacturer's incentive to present an overly optimistic forecast, the manufacturer is nevertheless able to share her information credibly and still expropriate all supply chain profit. In other words, asymmetric information presents no obstacle to the manufacturer's dominance under forced compliance. Under voluntary compliance the manufacturer is still able to share her forecast credibly, just not for free. Interestingly, firm commitments, which are part of several popular contracts (e.g., quantity flexibility contracts, backup agreements and pay-to-delay contracts), are a particularly effective signaling instrument even though a manufacturer should never offer them under full information. Firm commitments are undesirable because they restrict the system's ability to respond to evolving information. The supply chain should maintain the flexibility to defer the final production decision until after the manufacturer observes demand. We conclude that in our setting firm commitments are not useful for aligning incentives but are useful for communicating information.

While there has been a substantial and growing research in operations management that investigates the value of sharing demand information within a supply chain, much of the literature assumes truthful information is always exchanged. That is a reasonable premise in some settings. For example, when demand forecasts are constructed with past sales data, the demand forecast can be transferred by exchanging verifiable sales data. However, for products with no reliable past sales data (e.g., new products), forecasts must be generated using multiple data sources and managerial judgement. In those cases the credible exchange of demand forecasts is a significant issue. This model is a first step to understand how demand forecasts can be shared in such an environment. We predict that there will be high demand for additional research in this area.

Really.

Acknowledgments

The authors would like to thank Larry Wein and Paul Zipkin, as well as the seminar participants at The Kenan-Flagler Business School, University of North Carolina, The David Eccles School of Business, University of Utah, the 1997 MIT Summer Camp and the Supply Chain Roundtable sponsored by Eindhoven University. The helpful comments of the associate editor and reviewers are graciously acknowledged. The previous version of this paper was titled "Contracting to Assure Supply or What did the Supplier Know and When did He Know It?"

Appendix A. Proofs

PROOF OF LEMMA 1. At $K = m+o$, $\partial^2\Pi/\partial\sigma^2 = \partial^2\Pi/\partial m\partial o = -(r-w_e)f(m+o)$ and $\partial^2\Pi/\partial m^2 = \partial^2\Pi/\partial\sigma^2 - w_e f(m)$. The Hessian of $\Pi(m+0, m, o)$ is negative definite and $\Pi(m+o, m, o)$ is concave, making first-order conditions sufficient:

$$\partial\Pi/\partial m = (r-w_e)\overline{F}(m+o) - w_m + w_e\overline{F}(m) = 0,$$

$$\partial\Pi/\partial o = -w_o + (r-w_e)\overline{F}(m+o) = 0.$$

Solving jointly yields $\overline{F}(m^*) = (w_m - w_o)/w_e$ and $\overline{F}(o^*+m^*) = w_o/(r-w_e)$. The boundary conditions are found from considering $\overline{F}(m^*) = 0$ and $\overline{F}(m^*) > \overline{F}(o^*+m^*)$. □

PROOF OF THEOREM 1. Supply chain capacity is K^I and final production does not begin until demand is known so supply chain profit is Π^I. w_o and w_e are chosen such that the supplier's profit is zero while the limits on ε assure that $w_o \geq 0$ and $r > w_e \geq 0$. □

PROOF OF THEOREM 2. From direct differentiation:

$$\Pi''(K) = (r - w(K))S''(K) - 2w'(K)S'(K) - w''(K)S(K) < 0.$$

First-order conditions are sufficient and yield (3). □

PROOF OF THEOREM 3. Let K_2^* be the manufacturer's optimal quantity in the second market and $w_i(K)$ be the price that induces K in market $i = 1, 2$. Note that $\overline{F_1}(K) \geq \overline{F_2}(K)$ implies that $w_1(K) \leq w_2(K)$. Since $\int_0^K F_2(x)\,dx \geq \int_0^K F_1(x)\,dx$ for all K, the supplier could choose to induce K_2^* but pay less and have higher expected sales. □

PROOF OF THEOREM 4. We prove part (b) since the proof for (a) is similar. Consider marginal revenue at $K_2 + \delta_1 - \delta_2$ given δ_1:

$$r\overline{F}(K_2 + \delta_1 - \delta_2 | \delta_1) = r\overline{F}(K_2 - \delta_2) = r\overline{F}(K_2|\delta_2).$$

Marginal revenue equals marginal cost at K_1. Note $S(K|\delta) = S(K-\delta) + \delta$. Let $MC_i(K)$ be the manufacturer's marginal cost given δ_i at capacity K.

$$\begin{aligned} &MC_1(K_2+\delta_1-\delta_2)\\ &\quad= c_K\left(1 + \frac{f(K_2+\delta_1-\delta_2|\delta_1)}{\overline{F}(K_2+\delta_1-\delta_2|\delta_1)^2} S(K_2+\delta_1-\delta_2|\delta_1)\right)\\ &\qquad + c_p\overline{F}(K_2+\delta_1-\delta_2|\delta_1)\\ &\quad= c_K\left(1 + \frac{f(K_2-\delta_2)}{\overline{F}(K_2-\delta_2)^2}(S(K_2-\delta_2)+\delta_1)\right) + c_p\overline{F}(K_2-\delta_2)\\ &\quad= MC_2(K_2) + c_K(\delta_1-\delta_2). \end{aligned}$$

Thus, at $K_2+\delta_1-\delta_2$ marginal revenue given δ_1 is less than marginal cost. As revenue is concave and costs are convex, K must be reduced. The new optimum will be at a lower fractile of demand than before, so the wholesale price must be less.

For (c), suppose not. That would require that marginal revenue at δ is always greater than marginal cost. That is: $r > c_K(1+f(0)\delta) + c_p$, which does not hold for all δ.

Finally, for (d), the demand distribution is given by $D = \delta(\lambda) + \lambda X = \lambda\widetilde{X}$, where $\widetilde{X} = \sigma/\psi - \mu + X$. Demand can thus be expressed as a member of the family defined by (4) with $\delta = 0$. By (a), the optimal wholesale price is independent of λ. If we consider $\widetilde{X}$, it also fits within our family with $(\delta, \lambda) = (\sigma/\psi - \mu, 1)$. By (b), the wholesale price is decreasing in δ, which would make it increasing in ψ. □

PROOF OF THEOREM 5. Note that the denominator of $\bar{\varepsilon}$ is positive. For the numerator:

$$\begin{aligned} &c_K(K_H^I - K_L^I) - (r-c_p)(S_L(K_H^I) - S_L(K_L^I))\\ &\quad= c_K(K_H^I - K_L^I) - (r-c_p)\int_{K_L^I}^{K_H^I} \overline{F}_L(x)\,dx\\ &\quad> c_K(K_H^I - K_L^I) - (r-c_p)(K_H^I - K_L^I)\overline{F}_L(K_L^I)\\ &\quad= 0. \end{aligned}$$

where the inequality follows from the fact that $F_L(x)$ is an increasing function and the last equality follows from the definition of K_L^I. By construction, $\Pi_L(K_H^I, \varepsilon) < \Pi_L^I$ for all $\varepsilon < \bar{\varepsilon}$, so a low-type manufacturer prefers revealing her type truthfully. Since the high type is offering a coordinating contract, her profit is Π_H^I, which of course is better than what she could earn if she choose a low-type contract. □

PROOF OF LEMMA 2. Since $\Pi_\theta(K, \tau)$ is concave in K, from the envelope theorem,

$$\frac{\partial\Pi_\theta(K_\theta(\tau), \tau)}{\partial\theta\partial\tau} = -\frac{\partial S_\theta(K_\theta(\tau))}{\partial\theta}\frac{\partial w_\tau(K_\theta(\tau))}{\partial\tau} > 0.$$

The above is positive because for a given K expected sales is increasing in θ and the required wholesale price is decreasing in τ. The above implies that

$$\Pi_H(K_H^*, H) - \Pi_H(K_H(L), L) > \Pi_L(K_L(H), H) - \Pi_L(K_L^*, L),$$

which implies the lemma's result. □

PROOF OF THEOREM 6. Define

$$g(x, \theta) = \frac{S(x|\theta)}{S'(x|\theta)} = \frac{x - \int_0^x F(x|\theta)\,dx}{\overline{F}(x|\theta)}.$$

It is clear that $g(x, \theta)$ is increasing in x. Now show that $g(x, \theta)$ is convex in x:

$$\frac{\partial^2 g(x, \theta)}{\partial x^2} = \frac{F_\theta'(x)}{\overline{F}_\theta(x)} + \int_0^x \overline{F}(y)\,dy\left(\frac{2F_\theta'(x)^2}{\overline{F}_\theta(x)^3} + \frac{F_\theta''(x)}{\overline{F}_\theta(x)^2}\right) > 0,$$

where the sign follows because the first term is clearly positive and the latter term is positive if $w_\theta(K)$ is convex (which we have assumed). Since $g(x, \theta)$ is convex and increasing in x,

$$g(x, \theta) < x\frac{\partial g(x, \theta)}{\partial x}. \tag{A1}$$

With a scaled demand distribution, $g(x, \theta) = \theta_g(x/\theta, 1)$. Thus

$$\frac{\partial g(x, \theta)}{\partial\theta} = g(x/\theta, 1) - (x/\theta)\frac{\partial g(x/\theta, 1)}{\partial x} < 0,$$

where the inequality follows from (A1).

Now demonstrate that $K_L^* < K_L(H) \leq K_H^*$. Concavity implies that $K_L(H)$ is determined by the first-order condition:

$$\frac{\partial \Pi_L(K_L(H), H)}{\partial K} = (r - w_H(K_L(H)))S_L'(K_L(H)) - w_H'(K_L(H))S_L(K_L(H)) = 0.$$

For a scaled distribution, $w_L'(K) > w_H'(K)$ and $w_L(K_L^*) > w_H(K_L^*)$. Hence,

$$\frac{\partial \Pi_L(K_L(H), H)}{\partial K} > (r - w_L(K_L^*))S_L'(K_L^*) - w_L'(K_L^*)S_L(K_L^*) = 0,$$

implying that $K_L(H) > K_L^*$. If $K_L(H) > K_H^*$, then

$$\frac{(r - w_H(K_H^*))}{w_H'(K_H^*)} > \frac{S_L(K_H^*)}{S_L'(K_H^*)}, \tag{A2}$$

but $(r - w_H(K_H^*))/(w_H'(K_H^*)) = (S_H(K_H^*))/(S_H'(K_H^*))$ so (A2) would contradict that $g(K, \theta)$ is decreasing in θ.

Given that $K_L(H) \leq K_H^*$, it is costless for a high type to increase her capacity offer above K_H^*, but costly to the low type, therefore increasing her capacity offer by some amount is a more effective signal than just using a lump sum.

Proof of Theorem 7. The problem of the high-type manufacturer can be written as:

$$\begin{aligned} \max \quad & \Pi_H(K, m, A) \\ \text{s.t.} \quad & \Pi_L(K, m, A) \leq \Pi_L(K_L^*, L) \\ & m < \bar{m}(K), \end{aligned}$$

where the first constraint ensures that the low type would not wish to offer this contract and the second ensures that the supplier will actually build the requested capacity. Consider the corresponding Lagrangian:

$$Y = \Pi_H(K, m, A) + \lambda_1(\Pi_K(K_L^*, L) - \Pi_L(K, m, A)) + \lambda_2(\bar{m}(K) - m).$$

Let $\Pi_\theta^x = \partial \Pi_\theta(K, m, A)/\partial x$ for $x = K, m, A$. Necessary conditions are then

$$\frac{\partial Y}{\partial K} = \Pi_H^K - \lambda_1 \Pi_L^K + \lambda_2 \frac{\partial \bar{m}(K)}{\partial K} \leq 0 \quad K \geq 0 \quad \text{and} \quad K\frac{\partial Y}{\partial K} = 0 \tag{A3}$$

$$\frac{\partial Y}{\partial m} = \Pi_H^m - \lambda_1 \Pi_L^m - \lambda_2 \leq 0, \quad m \geq 0, \quad \text{and} \quad m\frac{\partial Y}{\partial m} = 0 \tag{A4}$$

$$\frac{\partial Y}{\partial A} = 1 - \lambda_1 \leq 0, \quad A \geq 0, \quad \text{and} \quad A\frac{\partial Y}{\partial A} = 0. \tag{A5}$$

Let K^* be the optimal capacity choice.

We first rule out $K^* < K_H^*$. $K_L(H) \leq K^* < K_H^*$ is never optimal since lowering capacity hurts the high type and actually benefits the low type. Now consider $K^* < K_L(H)$. It is possible to show that if $g(x, \theta)$ is decreasing in θ, then $\Pi_H^K > \Pi_L^K > 0$ for all $K^* < K_L(H)$, i.e., cutting back on capacity is always marginally more expensive for a high type than a low type. Consequently, $K^* < K_H^*$ cannot be optimal: increasing A is equally expensive for both types and thus is a cheaper instrument.

Now consider whether $K^* = K_H^*$. Suppose $m < \bar{m}(K_H^*)$, so $\lambda_2 = 0$ when $K = K_H^*$. Since we have assumed that the credibility constraint is violated with $A = 0$, it must be that $A > 0$ or $m > 0$. Suppose $A > 0$. From (A5), if $A > 0$, $\lambda_1 = 1$. But then (A3) yields $\Pi_L^K > 0$. However, $K_L(H) < K_H^*$, so $K^* \neq K_H^*$. Suppose $A = 0$ and $m > 0$ (and still $m < \bar{m}(K_H^*)$). From (A4), $\lambda_1 > 0$, so again (A3) yields $\Pi_L^K > 0$, i.e., $K^* \neq K_H^*$. Now suppose $m = \bar{m}(K_H^*)$, so $\lambda_2 > 0$. Note that $\partial \bar{m}(K)/\partial K > 0$. Therefore, evaluated at $K = K_H^*$, $\partial Y/\partial K > 0$, and again it must be that $K^* > K_H^*$.

Now show that $A > 0$ and $m < \bar{m}(K)$ cannot be optimal. Since $A > 0$, $\lambda_1 = 1$. Since $m < \bar{m}(K)$, $\lambda_2 = 0$. Thus, from (A4), $\Pi_H^m = \Pi_L^m$, but that cannot hold with $m > 0$ due to strict stochastic dominance. While that does hold at $m = 0$, it is easy to see that it must be a local minimum. □

References

Anand, H., H. Mendelson. 1997. Information and organization for horizontal multimarket coordination. *Management Sci.* **43**(12) 1609–1628.

Anupindi, R., Y. Bassok. 1999. Supply contracts with quantity commitments and stochastic demand. S. Tayur, R. Ganeshan, M. Magazine, eds. *Quantitative Models for Supply Chain Management*. Kluwer Academic Publishers, London, U.K., 197–234.

Aviv, Y. 1999. The effect of forecasting capability on supply chain coordination. Working paper, Olin School of Business, Washington University, St. Louis, MO.

——, A. Federgruen. 1998. The operational benefits of information sharing and vendor managed inventory (VMI) programs. Working paper, Olin School of Business, Washington University, St. Louis, MO.

Barlow, R. E., F. Proschan. 1965. *Mathematical Theory of Reliability*. John Wiley & Sons, New York.

Barnes-Schuster, D., Y. Bassok, R. Anupindi. 1998. Coordination and flexibility in supply contracts with options. Working paper, University of Chicago, Chicago, IL.

Biddle, F. 1998. Boeing to cut 747 output 30% in 1999 and to curtail production of its 777. *Wall Street Journal* (June 10).

Brown, A. 1999. A coordinating supply contract under asymmetric demand information: Guaranteeing honest information sharing. Working paper, Vanderbilt University, Nashville, TN.

——, H. Lee. 1998a. Optimal "pay-to-delay" capacity reservation with application to the semi-conductor industry. Working paper, Vanderbilt University, Nashville, TN.

——, ——. 1998b. The win-win nature of options based capacity reservation arrangements. Working paper, Vanderbilt University, Nashville, TN.

Cachon, G., M. Fisher. 1996. Supply chain inventory management and the value of shared information. *Management Sci.* Forthcoming.

——, ——. 1999. Capacity choice and allocation: strategic behavior and supply chain performance. *Management Sci.* **45**(8) 1091–1108.

Caldentey, R., L. Wein. 1999. Analysis of a production-inventory system. Working paper, Massachusetts Institute of Technology, Boston, MA.

Chen, F. 1998. Echelon reorder points, installation reorder points, and the value of centralized demand information. *Management Sci.* **44**(12) S221–S234.

Chowdrhry B., N. Jegadeesh. 1994. Pre-tender offer share acquisition strategy in takeovers. *J. Financial Quant. Anal.* **29**(1) 117–129.

Chu, W. 1992. Demand signaling and screening in channels of distribution. *Marketing Sci.* **11**(4) 327–347.

Cole, J. 1997a. Boeing, pushing for record production, finds parts shortages, delivery delays. *Wall Street Journal* (June 26).

——. 1997b. Boeing suppliers are feeling the heat as jet maker pushes to boost output. *Wall Street Journal* (September 16).

Corbett, C. 1999. Stochastic inventory systems in a supply chain with asymmetric information. Working paper, The Anderson School, University of California at Los Angeles, Los Angeles, CA.

Desai, P., K. Srinivasan. 1995. A franchise management issue: Demand signaling under unobservable service. *Management Sci.* **41**(10) 1608–1623.

Donohue, K. 1996. Supply contracts for fashion goods: optimizing channel profits. Working paper, The Wharton School, University of Pennsylvania, Philadelphia, PA.

Emmons, H., S. M. Gilbert. 1998. Returns policies in pricing and inventory decisions for catalogue goods. *Management Sci*. **44**(2) 276–283.

Eppen, G., A. Iyer. 1997. Backup agreements in fashion buying—the value of upstream flexibility. *Management Sci.* **43**(11) 1469–1484.

Gavirneni, S., R. Kapuscinski, S. Tayur. 1999. Value of information in capacitated supply chains. *Management Sci.* **45**(1) 16–24.

Gerchak, Y., H. Gurnani. 1998. Coordination in decentralized assembly systems with uncertain component yield. Working paper, University of Waterloo, Waterloo, Ontario, Canada.

——, Y. Wang. 1999. Coordination in decentralized assembly systems with random demand. Working paper, University of Waterloo, Waterloo, Ontario, Canada.

Graves, S., D. Kletter, W. Hetzel. 1998. A dynamic model for requirements planning with application to supply chain optimization. *Oper. Res.* **46** S35–S49.

Ha, A. 1999. Supplier-buyer contracting: asymmetric cost information and cut-off level policy for buyer participation. Working paper, Yale School of Management, New Haven, CT.

Kreps, D. 1990. *A Course in Microeconomic Theory*. Princeton University Press, Princeton, NJ.

Lariviere, M. A., V. Padmanabhan. 1997. Slotting allowances and new product introductions. *Marketing Sci.* **16** 112–128.

——, E. L. Porteus. 2000. Selling to the Newsvendor. Working paper, Fuqua School of Business, Duke University, Durham, NC.

Lee, H., P. Padmanabhan, S. Whang. 1997. Information distortion in a supply chain: The bullwhip effect. *Management Sci.* **43**(4) 546–558.

Mateja, J. 1998. Sold out: Corvette's 6-speed manual more popular than expected. *Chicago Tribune* (April 12).

van Mieghem, J. 1999. Coordinating investment, production and subcontracting. *Management Sci.* **45**(7) 954–971.

Moinzadeh, K. 1999. A multi-echelon inventory system with information exchange. Working paper, University of Washington, Seattle, WA.

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

Porteus, E., S. Whang. 1999. Supply chain contracting: Nonrecurring engineering charge, minimum order quantity, and boilerplate contracts. Working paper, Stanford University, Stanford, CA.

Pryweller, J. 1999. Delays jolt GM suppliers: Some part makers invest millions. *Automotive News* (June 21).

Spence, A. 1973. Job market signaling. *Quart. J. Econom.* **87** 355–374.

Suris, O., N. Templin. 1993. GM production problems hurt '94 model sales. *Wall Street Journal* (October 5).

Toktay, L. B., L. Wein. 1999. Analysis of a forecasting-production-inventory system with stationary demand. Working paper, INSEAD, Fontainebleau, France.

Tomlin, B. 1999. Short life cycle capacity decisions in supply chains with independent agents: The value of quantity premiums. Working paper, MIT, Boston, MA.

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

Zarley C., K. Damore. 1996. Backlogs plague HP – Resellers place phantom orders to get more products. *Computer Reseller News* (May 6).

Accepted by Christopher S. Tang; received September 1997. This paper was with the authors 18 months for 2 revisions.