Information Sharing in a Long-Term Supply Chain Relationship: The Role of Customer Review Strategy

Z. Justin Ren
Operations and Technology Management Department, Boston University School of Management,
Boston, Massachusetts 02215, ren@bu.edu

Morris A. Cohen
Operations and Information Management Department, The Wharton School, University of Pennsylvania,
Philadelphia, Pennsylvania 19104, cohen@wharton.upenn.edu

Teck H. Ho
Marketing Group, Haas School of Business, University of California, Berkeley, California 94720,
hoteck@haas.berkeley.edu

Christian Terwiesch
Operations and Information Management Department, The Wharton School, University of Pennsylvania,
Philadelphia, Pennsylvania 19104, terwiesch@wharton.upenn.edu

In this paper, we study the practice of forecast sharing and supply chain coordination with a game-theoretical model. We find that in a one-shot version of the game, forecasts are not shared truthfully by the customer. The supplier will rationally discount the forecast information in her capacity allocation. This results in Pareto suboptimality for both supply chain parties. However, we show that a more efficient, truth-sharing outcome can emerge as an equilibrium from a long-term relationship. In this equilibrium, forecast information is transmitted truthfully and trusted by the supplier, who in turn allocates the system-optimal capacity. This leaves both the customer and the supplier better-off, compared to the nontruthful-sharing equilibrium.

We identify a multiperiod review strategy profile that supports the truthful-sharing equilibrium. The key element of this strategy is that the supplier computes a scoring index of the customer's behavior that is updated over time and used to evaluate if the customer has sufficient incentive to share his private information truthfully in each transaction of the repeated game. Compared to trigger strategies, review strategies are more tolerant but require diligence and more monitoring effort.

1. Introduction

In November 2004, American Airlines announced that it was deferring the purchase of 54 out of 56 Boeing aircraft originally scheduled for delivery between 2006 and 2010. The financial impact for Boeing as the supplier was estimated to be as high as $2.7 billion (*Wall Street Journal*, November 22, 2004). This example highlights the inherent problem in sharing forecast information in a supply chain. Forecasts provide information about what the buyer intends to do in a given future state of the world. These intentions (often expressed in the form of preliminary purchase orders or soft orders), however, cannot be easily verified or enforced in court.

To overcome the problems associated with forecast sharing, previous research has suggested a set of contracts that align the incentives in the supply chain and induce the buyer to reveal demand information truthfully. Such contracts include two-part tariff and three-part quadratic

contracts (Desai and Srinivasan 1995, Corbett et al. 2004), and capacity commitment and option contracts (Cachon and Lariviere 2001, Ozer and Wei 2006). Despite their theoretical appeal of creating a "win-win" situation for buyer and supplier, empirical and anecdotal evidence shows that simple linear pricing contracts are still the most widely used format (Arrow 1985), possibly due to their simple structure and associated low administrative cost (Cachon and Lariviere 2005). For example, the aircraft industry (see Boeing example above), the semiconductor equipment industry (Cohen et al. 2003), and many other industries (Desai and Srinivasan 1995, Bajari and Tadelis 2001, Iyer and Villas-Boas 2003) still prefer the simple linear price contracts over the more complex coordinating contracts advocated by the supply chain literature. Although one could discount this empirical evidence as irrelevant and simply declare Fortune 500 companies such as Boeing and Intel to be laggards with respect to adopting the latest supply chain research, we offer an alternative explanation.

81

Most prior treatments of forecast sharing have used static or one-shot games. That is, in these models buyers and suppliers transact only once as in a spot-buy market. This eliminates any value of long-term relationships and thereby constrains the involved parties to ignore reputation. However, in practice, many business relationships are long term. Harmonious and mutually successful long-term supplier relationships have been observed in many industries and documented by researchers in strategy and management (e.g., McMillan 1990, Hagen and Choe 1998). In these supplier relationships, both the buyer and supplier are concerned about how their current strategic behavior affects their future interactions, and value long-term cooperation over one-shot business transactions. In fact, some buyers strongly prefer to deal with suppliers on a long-term basis, an observation that has been made repeatedly, especially in the automotive industry (Fujimoto 1999).

Although multiple forms of contracts exist that can achieve credible information sharing, we find that having a repeated relationship can itself also induce truthful information sharing. A long-term relationship between supply chain parties gives each party opportunities to review the credibility of the other party, reward truth telling, punish otherwise, and therefore provides the right incentive for truthful information sharing. In other words, repeated sourcing interactions can promote and form an implicit "relational contract" among supply chain parties that can be viewed as an alternative to complex pricing contracts in achieving supply chain coordination.

The present paper analyzes a forecast-sharing game between one supplier and one customer under information asymmetry. (The customer has superior information about demand compared to the supplier.) There are potentially two sources of inefficiency in this supply chain setting. The first inefficiency results from "double marginalization" (Spengler 1950). It is well known that the supplier operating with a linear wholesale contract would order less than the system-optimal quantity due to the difference between its profit margin and that of the supply chain. The other source of inefficiency results from information asymmetry. Whereas previous literature has identified contracting mechanisms that are able to overcome both sources of inefficiency and achieve supply chain coordination (Cachon and Lariviere 2001, Ozer and Wei 2006), here we restrict the contracting space to the sole usage of linear price contracts (with no other instruments, such as fixed payments or advance purchase orders). We show that under certain conditions an efficient equilibrium can still emerge through the repeated supply chain relationship.

This paper makes the following contributions. First, we are—to the best of our knowledge—the first to establish that truthful information sharing can be achieved in an infinitely repeated forecasting game setting. Our repeated game model hence complements extant literature on supply chain contracting by introducing a long-term perspective (Pyke and Johnson 2003; Taylor and Plambeck 2007a, b). Second, we

expand the strategy space beyond the commonly used trigger strategies and study review strategies in the context of forecast sharing. Whereas trigger strategies act on signals received in one period, review strategies are based on signals collected over a number of periods. The usage of review strategies in studying repeated games with imperfect monitoring is not new, but we are the first to show the applicability of review strategies to a supply chain setting with asymmetric information. Previous applications of review strategies have relied on some unbiased current-period public signal to serve as a monitoring mechanism. However, in our setting the customer not only has private information about the potential market size, but also privately observes actual demand, neither of which the supplier has perfect knowledge of. In this setting, we show that review strategies can still induce truthful information sharing and supply chain coordination in the long run if we utilize signals from *past* periods efficiently.

The remainder of this paper is organized as follows. We survey the relevant literature in §2. After analyzing the one-shot forecasting game (§3), we turn our attention to the analysis of the repeated game (§4) and show that with an incentive-compatible multiperiod review strategy, truthful information sharing and supply chain coordination can be achieved. In §5, we discuss our model limitations as well as several extensions, and conclude.

2. Literature Review

In this section, we briefly review supply chain coordination literature in the following order: one-shot games under perfect information, one-shot games under information asymmetry, and infinitely repeated games.

Previous literature has focused on studying supply chain coordination as a one-shot game. When there is perfect and symmetric information, various types of contracts have been identified that can achieve system coordination. These include buy-back (Pasternack 1985), quantity flexibility (Tsay 1999, Barnes-Schuster et al. 2002), sales rebate (Taylor 2002), quantity premium contract (Tomlin 2003) and revenue-sharing (Cachon and Lariviere 2005) contracts.

When information asymmetry in the supply chain exists, it matters who offers the contract. Cachon and Lariviere (2001) formulate a capacity procurement game, where the customer, who has some private information about demand, offers a signalling contract to the supplier. In order to induce the supplier to provide enough capacity for high demand, a buyer truly expecting high demand has to use more than just a linear price contract to signal its type. For example, the customer can offer a fixed side payment to the supplier in order to separate himself from a low-demand type customer. Whereas Cachon and Lariviere (2001) study such signalling contracting mechanisms in inducing truthful information sharing in a one-shot game, we use repeated games and investigate the effectiveness of just having a long-term relationship on truthful information sharing and supply chain coordination.

Screening contracts provide an alternative coordination mechanism. Porteus and Whang (1991) study capacity procurement with a screening model, where a supplier who is unsure about the market size offers the contract to a buyer who has the market information. Other screening models include Ha (2001) and Corbett et al. (2004). Ozer and Wei (2006) study a capacity procurement game with asymmetric information similar to Cachon and Lariviere (2001), but they consider both signalling and screening contracts. When the manufacturer can signal her forecast information with an advance purchase, it leads to a signalling game. On the other hand, the less-informed supplier can induce the truthful forecast information from the manufacturer by offering a capacity reservation contract, which is a screening game. Moreover, Ozer and Wei (2006) show that channel coordination is possible when combining a capacity commitment contract with a pay-back contract. In this paper, we also consider channel coordination with information asymmetry. However, we use neither signalling nor screening devices. Instead, we resort to the repeated nature of supplier relationships and show that coordination is achievable in the long term even with a simple linear contract. Other papers that study incentives in information sharing include Lariviere and Padmanabhan (1997), Van Mieghem (1999), Corbett (2001), and Chen (2005). Chen (2002) provides an extensive literature review on information sharing and supply chain coordination.

During recent years, building a long-term supplier relationship has been increasingly recognized as an important aspect of supply chain management. Beth et al. (2003) pointed out that building a relationship is the key in managing supply chains, instead of investing in modern technologies. Pyke and Johnson (2003) study a spectrum of spot-buy and long-term sourcing relationships and identify the key drivers of each type of relationship. Taylor and Plambeck (2007a) study a repeated supply chain relationship in which production cost is uncertain, and therefore ex ante binding contracting is infeasible. They identify an optimal "relationship contract" that maximizes the customer's expected profits. They also demonstrate that better information visibility in monitoring unobserved actions can improve system performance significantly. We develop a qualitatively similar insight in this paper, but in a different setting. Another point of difference is that instead of allowing pregame "settlement" in each stage game, as done in Taylor and Plambeck (2007a), we do not allow for this possibility. Instead, we restrict the contracting space to linear transfer payments after capacity is built and demand is realized. Taylor and Plambeck (2007b) further investigate two forms of relational contracts, where the buyer can either promise to pay a specific price, or can promise to purchase a specific quantity. These contractual arrangements are both easy to implement, and the authors show that which option is optimal to the buyer depends on the production and capacity cost, as well as the discount factor. In our model, the form of transfer payment, i.e., a linear

contract, is predetermined, and our focus is on achieving truthful forecast sharing in a long-term relationship.

In infinitely repeated games, trigger strategies are the most commonly studied type of strategy (Friedman 1986) and are often adopted in the relational contracting literature (e.g., Baker et al. 2002; Levin 2003; Taylor and Plambeck 2007a, 2007b; Plambeck and Taylor 2006). In our setting it can be shown that trigger strategies can induce truthful forecast sharing but not system coordination, because punishment will be triggered in equilibrium due to the fact that the supplier cannot perfectly observe customer demand. Therefore, we seek alternative strategies that not only can induce truthful information sharing, but also system coordination. We find that using a properly defined review strategy can achieve this goal. Review strategies were first studied in economics in the context of the repeated principal-agent model under imperfect monitoring. Radner (1985) studies repeated principal-agent games where the agent's action is not perfectly observed. He shows that by constructing a review strategy the repeated game can achieve asymptotic efficiency. We utilize a similar review strategy in this paper. However, the monitoring mechanism in our model is different. In Radner (1985), the outcomes from the agent's private actions are observable and can serve as a monitoring tool. In our setting, the realized demand sequences are only perfectly observed by the customer and cannot be directly used to evaluate the customer's truthfulness. Instead, the history of forecasts and orders by the customer, which are publicly observable, are used to access the customer's truthfulness.

The game we study also belongs to a class of games called cheap-talk games (see, for example, Crawford and Sobel 1982). Cheap-talk games are different from signaling games in that the signals that are being sent in a cheap-talk game are costless, whereas the signals sent in signaling games must have cost implications (Gibbons 1992, p. 210). When cheap-talk games are played repeatedly, it is generally unclear whether efficient outcomes can be sustained. That is, the well-known folk theorem (Friedman 1971) may not apply in such games. For example, Morris (2001) considers two types of senders with different preferences over receiver's actions, and finds that truth telling is not an equilibrium even in the infinitely repeated game. On the other hand, Stocken (2000) studies the credibility of a firm's manager disclosing private information to an outside investor, and shows that the manager does not have an incentive to truthfully reveal his information in a single-shot game, but can be motivated to do so if such a game is played repeatedly. Different from the abovementioned papers, our paper has an additional source of information asymmetry. The customer not only privately observes the demand-size parameter, he also privately observes the true demand realization, neither of which the supplier observes perfectly.

3. The Model: One-Shot Game

In the supply chain we study, there is one customer (also called the "buyer," referred to as "he") and one supplier

(referred to as "she"). The customer wants to acquire a certain amount of capacity from the supplier before actual demand is realized. We assume that the market demand the customer faces is a scaled random variable $\theta \cdot X$, where X is a nonnegative normal random variable with mean μ and standard deviation σ. Define $\theta \in \Re^+$ as the demand-size parameter. To capture the essence of the uncertain market condition, we assume that θ is a random variable with two possible values "high" and "low:" θ_i, $i = \{h, l\}$. The probability of each happening is $P(\theta_l) = \alpha$, $P(\theta_h) = 1 - \alpha$, $\alpha \in (0, 1)$. The high and low market demand random variables are therefore $D_h = \theta_h X$, and $D_l = \theta_l X$, respectively. Denote their CDF $F_h(.)$ and $F_l(.)$, respectively. It is assumed that θ is independent of X. The distributions of both θ and X are common knowledge.

The supplier's unit capacity cost is c, and she charges the customer a price of r for each unit allocated and utilized. The customer earns a unit revenue of p for each unit of capacity allocated and utilized. When the capacity allocated is not enough to satisfy realized demand, the customer incurs a unit cost of g. This includes potential revenue lost, loss of goodwill, or the additional cost to find alternative sources of supply. On the other hand, if the supplier allocates too much capacity, for those units not purchased by the customer, a unit overage cost of h is incurred. All the cost parameters are public knowledge. For notational convenience, define

$$u(\theta_i, K) = r \min(K, D_i) - h(K - D_i)^+ - cK, \qquad (1)$$

$$v(\theta_i, K) = (p - r) \min(K, D_i) - g(D_i - K)^+. \qquad (2)$$

where $(x)^+ = \max(x, 0)$.

The game begins with a random draw of θ by the customer. The customer observes this realization of θ (either θ_h or θ_l), but the supplier does not. The customer then sends a nonbinding forecast m to the supplier. Here we consider the simplest forecast information set: either $m = H$, or $m = L$. The supplier's strategy is also binary: to trust, or not to trust the customer's forecast. If the supplier trusts the customer's forecast, then she builds capacity according to the received forecast. That is, if she receives a forecast of H from the customer and she trusts the customer, then she would allocate capacity K_h that maximizes her expected profit believing that the market demand random variable would be D_h:

$$K_h = \arg\max_K \mathsf{E}[u(\theta_h, K)]. \qquad (3)$$

Similarly, if the supplier trusts the customer and receives a forecast of L, then she would allocate capacity K_l that maximizes her expected profit, believing that the market demand random variable would be D_l:

$$K_l = \arg\max_K \mathsf{E}[u(\theta_l, K)]. \qquad (4)$$

On the other hand, if the supplier does not trust the customer, then his forecast is considered meaningless. The supplier then allocates capacity K_0 as if there were no meaningful forecast being shared at all:

$$K_0 = \arg\max_K \{(1 - \alpha)\mathsf{E}[u(\theta_h, K)] + \alpha\mathsf{E}[u(\theta_l, K)]\}. \qquad (5)$$

We assume that capacity is publicly observable to the customer. This may not be true in every situation (see Taylor and Plambeck 2007a for a model with unobservable capacity choices), but is reasonable for industries where production volume can easily be observed. For example, we found in our prior research in the semiconductor equipment supply chain, as well as in the aerospace and defense sector, that customers often have on-site personnel (often to support on-site quality assurance and testing) who have detailed knowledge about the current capacity and production plans of their suppliers. After capacity is allocated, the true market demand (i.e., a draw from the underlying distribution of θX) is realized, which we label as d. The customer places an order o. It is easy to see that in the one-shot game, the customer has no incentive to order anything other than the true demand.

In the truthful information-sharing situation (the customer shares his forecast truthfully and places his order equal to true demand and the supplier trusts the forecast), the customer's expected profit v^* and the supplier's expected profit u^* are:

$$v^* = (1 - \alpha)\mathsf{E}[v(\theta_h, K_h)] + \alpha\mathsf{E}[v(\theta_l, K_l)]$$
$$u^* = (1 - \alpha)\mathsf{E}[u(\theta_h, K_h)] + \alpha\mathsf{E}[u(\theta_h, K_h)]. \qquad (6)$$

These expected payoffs correspond to the cooperative case. In the noncooperative case, the supplier ignores the customer's forecasts and allocates capacity according to her prior knowledge about θ alone. The expected profits are:

$$v^0 = (1 - \alpha)\mathsf{E}[v(\theta_h, K_0)] + \alpha\mathsf{E}[v(\theta_l, K_0)],$$
$$u^0 = (1 - \alpha)\mathsf{E}[u(\theta_h, K_0)] + \alpha\mathsf{E}[u(\theta_h, K_0)]. \qquad (7)$$

Table 1 summarizes the notation used in the one-shot game.

The following characterizes the equilibrium of the one-shot game: (its proof is provided in the appendix)

PROPOSITION 1. *In the equilibrium of the one-shot game, the supplier ignores the customer's forecast information and allocates capacity K_0 regardless of demand state θ.*

Driving the above results is the customer's underlying incentive to overforecast. Whereas the customer with a high demand wants to reveal the truth to the supplier, the low-demand customer also wants to send a high-demand forecast so that he could have more capacity allocated to him, i.e., he would like to mimic a high-demand customer. Expecting such an incentive to overforecast from the customer, the supplier cannot distinguish the true demand

Table 1.   Notation: One-shot game.

$\theta_i, i = \{h, l\}; \alpha$	Demand-size parameter; probability of low market demand (θ_l)
p, r	Customer's unit revenue, and transfer price to the supplier for capacity allocated and utilized
c, h, g	Supplier's unit capacity cost, overage cost, and shortage cost
$D_i; F_i(.); d$	End-market demand the customer faces; CDF for D_i; actual demand realization from D_i
$m; o$	Customer's forecast; customer's order after observing demand
$K_i; v^*, u^*$	Capacity under cooperation; and the resulting expected profit for customer and supplier
K_0, v^0, u^0	Capacity under noncooperation; and the resulting expected profit for customer and supplier

state, and is better-off ignoring the customer's forecasts. Such a reaction from the supplier results in a nontruthful forecast-sharing equilibrium. A similar result is obtained in Ozer and Wei (2006). After demand is realized, however, the customer always places his order according to true demand $(o = d)$, because ordering any other amount has no benefit and can only hurt his expected profit. This equilibrium is inefficient. It is easy to see that there exists a set of capacity levels K'_i that Pareto dominates K_0:

$$(1 - \alpha)\mathsf{E}[v(\theta_h, K'_h)] + \alpha\mathsf{E}[v(\theta_l, K'_l)] \geq v^0, \tag{8}$$

$$(1 - \alpha)\mathsf{E}[u(\theta_h, K'_h)] + \alpha\mathsf{E}[u(\theta_l, K'_l)] \geq u^0. \tag{9}$$

The capacity levels under the cooperative solution as defined in Equations (3) and (4) belong to this set. Another focal candidate is the supply chain optimal capacity levels K_i^c:

$$K_i^c = \arg\max_K \mathsf{E}[v(\theta_i, K) + u(\theta_i, K)]. \tag{10}$$

Unfortunately, in the one-shot game, none of the those Pareto-dominant capacity levels can happen in equilibrium because of the misaligned incentives induced by information asymmetry and by double marginalization. A natural question is therefore to ask under what circumstances the more desirable outcomes can emerge. In the next section, we investigate the forecast-sharing game in the context of a long-term supply chain relationship.

4. A Long-Term Supply Chain Relationship

We model a long-term supplier relationship as a repeated game between customer and supplier. The repeated game consists of the following features. Let time be indexed by $t = 1, 2, \ldots, \infty$. At each time t, a stage game is played, starting with a random draw of θ, which we denote as θ_t. The stage game is identical to the one-shot game: (1) After privately observing θ_t, the customer's action is to choose which forecast to provide to the supplier, denoted by m_t. (2) The supplier's action is to allocate a capacity level K_t. (3) After capacity K_t has been allocated, true demand d_t is realized. (4) The customer then places the order of o_t to fulfill demand. If $K_t < o_t$, only K_t units of the orders are satisfied, and the remaining unsatisfied orders are lost. If $K_t \geq o_t$, then the entire order is satisfied. (5) The supplier is paid according to a linear pricing contract with a unit rate of r. This concludes the one-stage game. Then time is advanced to $t + 1$, and the game renews, with another independent draw of θ_{t+1}.

The public information at the start of period t consists of all the past forecasts shared $m^t = (m_1, \ldots, m_{t-1})$, orders placed $o^t = (o_1, \ldots, o_{t-1})$, and capacity levels allocated $K^t = \{K_1, \ldots, K_{t-1}\}$. Collectively, this information is called the public history at time t: $h^t = m^t \times o^t \times K^t$, which is fully observed by both parties. In contrast, the private history of the repeated game consists of all the demand signals and demand realizations that the customer privately observed: $h^t_{private} = (\theta_1, \ldots, \theta_{t-1}) \times (d_1, \ldots, d_{t-1})$.

The customer chooses a strategy corresponding to a sequence of functions that map the public history and his private information into forecasts and orders: $h^t \times \theta_t \times d_t \to m_t, o_t, \forall t$. The supplier chooses a strategy corresponding to a sequence of functions that maps the public history into capacity decisions: $h^t \times m_t \to K_t$. Note that neither party's strategies depend on private history $h^t_{private}$, hence, they are public strategies (Fudenberg and Tirole 1991).

Let the discount rate be δ_0. To model an uncertain lifetime of the customer-supplier relationship, suppose at the end of each period that there is a probability p_0 that the relationship will be terminated. Hence, $\delta = \delta_0(1 - p_0)$ is the effective discount factor.

A perfect public equilibrium (PPE) is a profile of public strategies that starting from any period t, given any public history h^t, constitutes a Nash equilibrium from that point onward (Osborne and Rubinstein 1994).

Below we show that a Pareto-efficient outcome can emerge as an equilibrium in the repeated game using a strategy profile called a multiperiod review strategy. This is accomplished by the supplier setting up a scoring system and evaluating the customer's credibility over a number of consecutive periods. Review strategies have previously been studied in the economics literature (Radner 1985, Stocken 2000). In contrast to the widely studied trigger strategies, a review strategy in the context of a repeated game with imperfect monitoring does not prescribe immediate punishment once the customer fails the supplier's credibility test in one period. Instead, the customer is given more chances over a longer period of time. As a result, the likelihood of mistakenly punishing a truth-telling player is smaller. On the other hand, the review strategy needs to be stringent enough so that the customer has no incentive to "coast" once he accumulates high-enough credits. Therefore, the parameters of the review strategy, particularly the

review length and the threshold for punishment, need to be set carefully. Below we explicitly derive a review strategy that achieves both truth telling and improves efficiency for both parties.

Our review strategy works as follows. Time is divided into a sequence of review phases, each having a maximum length of R periods. During each review phase, the customer reveals information truthfully, and the supplier maintains an index (i.e., a "scorecard") I_t and a credibility threshold γ_t. At the end of each period $t \in [1, R]$, after order o_t is placed, the supplier evaluates the truthfulness of the customer with a credibility test, and updates the score I_t.

If the customer reports demand to be high ($m_t = H$), the customer's score is subject to the following two tests. The first test evaluates the proportion of periods in which the customer reported high demand ($M_t = H$) so far in the review period. Because in each period the high-demand scenario happens with probability $1 - \alpha$, in the long run the supplier should observe approximately $1 - \alpha$ proportion of the forecasts be high demand, assuming the customer reports θ_t in each period truthfully. Denote N_t the number of periods up to t in a review phase in which the customer reports high demand. The sampling distribution of α is approximately normal, with mean α and standard deviation $\sqrt{\alpha(1 - \alpha)/(t - N_t)}$. Let z_α be the value associated with a confidence level p_α (for example $z_a = 1.96$ for a confidence level of $p_\alpha = 95\%$). The customer is then judged to pass the first test in period t if the proportion of high demand reported up to t is no greater than $(1 - \alpha) + z_a\sqrt{\alpha(1 - \alpha)/N_t}$. This test is designed to reduce the customer's incentive to excessively overforecast and report high demand, as is the case in the one-shot game.

The second test analyzes the actual order quantities o_t. A threshold $\underline{d}_{h,t}$ is set according to the following statistical test. If the customer is reporting truthfully, then the sample mean of the demand realizations from those high-demand periods (θ_h) is normally distributed with mean $\theta_h\mu$ and standard deviation $\sigma\sqrt{\theta_h/N_t}$. Let z_h and p_h be the critical value and probability associated with a certain significance level (for example, $z_h = 1.96$ for a significance level of $p_h = 95\%$). Then $\underline{d}_{h,t} = \theta_h\mu - z_h\sigma\sqrt{\theta_h/N_t}$. The customer passes the second test if the average order size in those reported high-demand periods up to time t exceeds $\underline{d}_{h,t}$. Only after the customer passes both of the above tests can his score be incremented, $I_t = I_{t-1} + 1$.

If the customer reports demand to be low ($m_t = L$), the customer is evaluated with only one test related to the actual order he places. If the customer is ordering truthfully, the sample mean of the demand realizations from those low-demand periods is normally distributed with mean $\theta_l\mu$ and standard deviation $\sqrt{\theta_l}\sigma$. Let z_l and p_l be the critical value and probability associated with a certain significance level (for example, $z_d = 1.96$ for a significance level of $p_l = 95\%$). Then $\bar{d}_{l,t} = \theta_l\mu + z_l\sigma\sqrt{\theta_l/(t - N_t)}$. The customer

passes the test if his average order size in the reported low-demand periods in the review phase up to t is no greater than $\bar{d}_{l,t}$. If this happens, then his score is incremented by one, $I_t = I_{t-1} + 1$. On the other hand, if he fails either test, his score is not updated, $I_t = I_{t-1}$.

The purpose of testing the low-forecast-reporting customer is that under some sample paths of history, the customer can have an incentive to underforecast (i.e., report low demand when in fact observing high demand). This will happen when the customer has cheated in the past, for example, by overforecasting his past demands for a certain number of times. Now his sequence of announced forecasts is in danger of being detected as "not normal." In response, the customer may want to "fake" a randomly generated sequence by underforecast. For this reason, the credibility tests need to be diligent in both high-forecast and low-forecast-reporting customers.

The overall probability that the customer is evaluated as truthful is $S = (1 - \alpha)p_h p_\alpha + \alpha p_l$. This measures the strength of the credibility checking system. The credibility threshold γ_t is defined as $\gamma_t = ((R \cdot S - J)/R)t$, where J is some allowance for measurement error. When $J = 0$, the credibility threshold is simply $\gamma_t = St$.

At the end of each period, the supplier also evaluates whether the customer still has an incentive for truth telling. If he does, the review phase is continued. If the supplier finds that the customer no longer has an incentive to report forecasts truthfully, the review phase is terminated. This can happen when the customer has accumulated a high-enough score before time R and now has an incentive to "coast." Or, the customer can have an incentive to cheat when he has a string of bad luck with a low score so that he has no chance of passing the test during the remaining periods of the review phase. Upon the termination of the review phase, if the supplier score $I_t \geqslant \gamma_t$, the customer passes the review and the games restarts with a new review phase. If $I_t < \gamma_t$, the customer fails the review and the games enters a nontruthful-sharing (punishment) phase that lasts for $M = \chi R$ periods, where $\chi > 0$.

During the review phase, the supplier always trusts the customer, and builds the system-optimal capacity according to the customer's forecasts of market demand. If the supplier deviates from the system-optimal capacity, then the customer will punish such a deviation by stopping to share information truthfully, and consequently the game enters a punishment phase. During the punishment phase, the supplier ignores the customer's forecasts, and the customer does not share forecasts truthfully. Therefore, the game reverts to the one-shot uncommunicative Nash equilibrium for M periods, which can be set long enough so that the supplier has no incentive to unilaterally deviate. The repeated game begins with a review phase, with supplier score $I_0 = 0$. Let the customer's total discounted expected payoff at the beginning of the game be denoted as $V(I_0)$, and let the customer's one-stage payoff associated with the

supply chain optimal capacity level K_f^c (as defined in Equation (10)) be denoted as v^c, i.e., $v^c = (1 - \alpha)\mathsf{E}[v(\theta_h, K_h^c)] + \alpha\mathsf{E}[v(\theta_l, K_f^c)]$.

THEOREM 2. *There exists a discount rate $\underline{\delta}$, a review phase of length $\underline{R}$, and a credibility assessment value $\underline{S}$ such that for $\forall \delta > \underline{\delta}, R > \underline{R}, S > \underline{S}, \varepsilon > 0$, there exists a perfect public equilibrium of the repeated game in which $V(I_0) > v^c - \varepsilon$, provided that the efficient capacity levels K_i^c are Pareto improving for both parties compared to K_0. In this equilibrium, the customer always shares demand information as well as places orders truthfully, and the supplier allocates the system-optimal capacity K_i^c.*

The proof for this theorem is provided in the appendix.

The above theorem not only stipulates that there exist parameters of the game that can support a truthful information-sharing equilibrium, it also states that in such an equilibrium the nontruthful information-sharing punishment phase would almost never be observed even in the presence of asymmetric information. This is because with a review strategy such as defined above, the customer is given an extended period of time to prove his intention of sharing information truthfully. If, for example, the customer had bad luck and had a very low demand realization in one period, but had forecasted high demand in good faith, the customer would not be punished immediately. Instead, his credibility score would suffer, but he would still be regarded as trustworthy and be given more periods to vindicate himself. On the other hand, if the customer intentionally overforecasted demand, he may not be caught immediately, but the credibility tests are designed such that repeated lying would be caught. This is feasible because the supplier has information on the distribution of both demand-size parameters, and those of each of the demand distributions. Therefore, hypothesis tests on sample means can be utilized with repeated samples.

In the long run, the review strategy ensures that the probability that a truth-telling customer fails the review goes to zero as the number of review periods R becomes large (provided also that the discount factor δ is sufficiently high, and the credibility check system S is effective enough). Also, note that as R becomes large, the length of the punishment phase is also longer. This also increases the customer's incentive to share forecast information truthfully. The effect of the inefficiency brought by the longer punishment phase is mitigated by a reduction in the probability of inappropriately punishing a customer who has been sharing forecast information truthfully. As is shown, for a sufficiently long review period, the probability of the customer failing the review when sharing his forecasts truthfully can be bounded arbitrarily close to zero. Hence, the supply chain converges to a fully efficient coordinated equilibrium.

Also, as the discount factor δ increases, future long-term gains from sharing the forecast truthfully, as well as the punishment threat from doing otherwise, become more important. Therefore, the customer is less inclined to

behave myopically and opportunistically. There are a number of factors that affect the discount factor δ. For example, recall that $\delta = \delta_0(1 - p_0)$ is a decreasing function of p_0, the probability that the repeated game may stop at the end of each stage game. Some industries are highly risky, with firms entering and exiting at a fast rate. In such industries, p_0 is high, and therefore firms tend to focus more on near-term payoff. As a result, they tend to ignore the value of long-term relationships. In such environments, our model suggests that truthful information sharing is more difficult to obtain, and a long-term relationship is less likely to emerge.

Review strategies with long review periods R can be observed in practice. For example, Sun Microsystems rates each of its suppliers' performance with a "scorecard" system and reviews the scores with the supplier each quarter. Suppliers are given multiple opportunities to meet the company's requirements. The review system helps Sun to achieve its goal of building long-term relationships with suppliers (Farlow et al. 1996). Companies such as Intel (Minahan 1998a), Chrysler (Minahan 1998b), and Toyota and Mitsubishi (Hagen and Choe 1998) all have similar review practices. Although in those examples scorecards are used to evaluate suppliers, our model suggests that such review systems can also be effective in gauging the customer's truthfulness in information sharing when some information cannot be symmetrically observed.

In our demand model there are two states (high demand or low demand). Therefore, it can be applied to business situations where demand can be characterized as a "boom or bust" cycle. See Terwiesch et al. (2005) for an example in the semiconductor equipment supply chain. In each period the customer gives a forecast about the overall demand. He has an incentive to always predict a strong demand, just as our one-shot model prescribes. However, our model also shows that, in the long run, the supplier is able to tell if the customer is being overconfident by recalling the excessive number of rosy forecasts the customer has been giving.

5. Discussion and Conclusion

In this paper, we study the practice of forecast sharing and supply chain coordination with a game-theoretical model. We find that in a one-shot version of the game, forecasts are not shared truthfully by the customer. The supplier will rationally discount the forecast information in her capacity allocation. This results in Pareto suboptimality for both supply chain parties. However, we show that a more efficient truth-sharing outcome can emerge as an equilibrium from a long-term relationship. In this equilibrium, forecast information is transmitted truthfully, and trusted by the supplier who in turn allocates the system-optimal equilibrium. This leaves both customer and supplier better off, compared to the nontruthful-sharing equilibrium.

We identify a multiperiod review strategy profile that supports the truthful-sharing equilibrium. The key element

of this strategy is that the supplier computes a scoring index of the customer's behavior that is updated over time and used to evaluate whether the customer has sufficient incentive to share his private information truthfully in each transaction of the repeated game. Compared to single-period trigger strategies, multiperiod review strategies are more tolerant, but require diligence and more monitoring effort. This is also consistent with what we observe in a coordinated supply chain. For example, Toyota, when working with its suppliers, does not immediately sever the business relationship with any of its suppliers when it has quality issues or delivery problems that fail to meet Toyota's standards. Instead, the company dispatches its own staff to the supplier's site and works with their suppliers for extended periods of time to help them correct the problem. Only under extreme circumstances does Toyota drop any supplier from its supplier base (Liker 2004). A similar observation is made in the semiconductor equipment industry (Cohen et al. 2003).

The implication of our findings is that when supply chain parties have a long-term ongoing relationship, or expect to do so, they can achieve coordination even though incentives are not aligned in the short run. In the case of forecast sharing, the customer has an incentive to overforecast if the supplier relationship is of short-term nature. However, when the parties interact repeatedly, the threats of retaliation in response to bad behavior and the promise of cooperation in response to good behavior in the current period can have significant impact on future payoffs. Concerns about credibility and payoff in the long run can thus yield cooperation and coordination. This "relational contract" is desirable especially when actions are not perfectly observable and cannot be contracted or enforced.

Our model has limitations. First, not every supply chain relationship is long term. Our model also shows that not every long-term supply chain relationship can achieve truthful forecast sharing and coordination. In order to achieve cooperation in those supplier relationships, other coordinating mechanisms such as explicit contracting can help. For example, in our one-stage model, the customer can pay the customer a capacity reservation fee. After demand is realized, the customer can decide how much capacity to actually utilize, or it can pay a higher price for any additional unit of capacity requested. Similar monetary contracts have been shown to be effective by, for example, Brown and Lee (1998) and Ozer and Wei (2006). There are only two demand states (θ_h and θ_l) in our model. An extension to multiple states will be meaningful because the review strategy for distinguishing each state from other multiple potential states will be more sophisticated.

There are a few possible extensions of our model. In our current setting we assume that the supplier's capacity choice is perfectly observable. This needs not be the case. It can be shown that a truth-sharing long-term equilibrium can be upheld even when the supplier's capacity choice is not perfectly observable. In the equilibrium the customer

can use a trigger strategy that will punish the supplier once the supplier is not able to fill all the customer's orders. Once the supplier's capacity is reserved, any underage in delivery (i.e., the portion of the customer's orders that cannot be delivered due to insufficient capacity) can serve as an informative signal on the true capacity level reserved by the supplier. Therefore, the customer can probabilistically test the credibility of the supplier. Of course, in order for this equilibrium to be sustained, the signal-to-noise ratio has to be high enough.

In our current model, the demand signal θ_i is independent across time periods. It can be extended to include a more general Markov-modulated demand case. When θ_i follows a two-state (i.e., high demand or low demand) Markov chain, one can still compute the sample mean and standard deviation of $S_n = \sum_{i=1}^{n} \theta_i$. Because θ_i is ergodic and stationary, a version of the central limit theorem stills applies (Durrett 1996, p. 421). Therefore, the supplier still reviews the truthfulness of the customer by testing the proportion of time that high demand is reported.

In our review strategy, the supplier checks the proportion of time high-demand situations are reported, as well as the mean order quantities in both high- and low-demand times to ensure that the customer has no room for sharing forecasts untruthfully. Such a review strategy needs not be unique. For example, the supplier can use the Kolmogorov-Smirnov test in each period to evaluate whether the data points fit the target distribution well, or the supplier can employ a Bayesian procedure in testing normality. However, an effective review strategy should at its minimum contain tests on both the demand-size parameter and the order quantities under each demand scenario. This is because the customer, in deviating from truth telling, can potentially lie about both. A natural review strategy is to look at some of the sufficient statistics in the sequences of random variables, such as the proportion of times high demand is reported, and the mean of reported demands, which this paper adopts. They are natural because they are the first moments of the sequences of random variables. In fact, that is how the frequentist statistical theory on hypothesis testing works. Our conjecture is that the more complex the review strategy is, the more effective it may be.

One can also generalize other aspects of the model. It would also be interesting to generalize the supply chain structure to one with one supplier and multiple customers, and see how the dynamics change in a long-term relationship, especially how competition among customers affects the equilibrium outcome. We hope to address these modeling issues in our future research.

Appendix: Proofs

PROOF OF PROPOSITION 1. Suppose the supplier trusts the customer and always reserves capacity according to his forecast. The high-type customer ($\theta = \theta_h$) will surely send H to reserve more capacity. However, the low-type customer also has an incentive to forecast H even though he

receives low demand $\theta = \theta_l$, because his profit is strictly increasing in capacity reserved. The supplier's expected profit would be higher if she would just build $K_0 \in [K_l, K_h]$, because

$$(1 - \alpha)\mathsf{E}[(p - r)\min(K_h, D_h) - g(D_h - K_h)^+]$$
$$+ \alpha\mathsf{E}[(p - r)\min(K_h, D_l) - g(D_l - K_h)^+]$$
$$\leqslant (1 - \alpha)\mathsf{E}[(p - r)\min(K_0, D_h) - g(D_h - K_0)^+]$$
$$+ \alpha\mathsf{E}[(p - r)\min(K_0, D_l) - g(D_l - K_0)^+]$$

by the definition of K_0 (see Equation (5)). Therefore, trusting the customer's forecast is not a best response for the supplier. Given that the supplier's best response is to ignore the forecasts, then it does not matter what the customer forecasts, because his assigned capacity level is fixed at K_0. After demand is realized, the customer always placed its order according to the true demand realization, because ordering any other amount has no benefit, and can only hurt his own expected profit. $\square$

The following is for the proof of the main theorem (Theorem 2), but first we define some notation:

$v(I_0)$ is the customer's normalized discounted expected profit at the beginning of the game;

$v(I_{R-n}^x)$ is the customer's normalized discounted expected profit when there are n periods remaining in the review phase and the value of the trust index at period $t = R - n$ is x;

$q = R \cdot S - J$, where J is some allowable margin of error (whose exact value is to be defined in the latter part of the proof). This is the threshold below which the customer fails the review at the end of the review phase of length R. Also note that $\gamma_t = (q/R)t$.

Now we outline the main idea of the proof of the main theorem. It involves the following steps:

(1) First, note that players do not have an incentive to deviate from the equilibrium strategies.

(2) We show that the customer will fail the review when at any $t < R$, $I_t < \max(0, t - (R - q))$, i.e., when I_t is such that the customer will fail the review for sure. For example, suppose $R = 30$, $q = 24$, $t = 10$. If at $t = 10$, I_t is less than $10 - (30 - 24) = 4$, then the customer has no hope of reaching the threshold, which is 24. This is because although there are $30 - 10 = 20$ periods to go, even in the best-possible case where the customer index is incremented in each of the 20 remaining periods, it will still be less than $4 + 20 = 24$. From this example, we see that when $I_t \geqslant \max(0, t - (R - q))$, then the customer may still have a chance to pass the review. A sufficient condition is to show that the customer has an incentive to share forecasts truthfully when $I_{R-n} = q - n$, i.e., when there are n periods left, and the customer has to have his credibility score incremented in each of the remaining periods in order to pass the test at the end of the review phase. We will prove this in Lemma 2.

(3) There exist parameters of the game such that the customer will share forecasts truthfully at the beginning of the game, which will be proven in Lemma 3.

(4) In order to show that the probability of failing the review when the players use the specified review strategies can be made arbitrarily close to zero, we resort to an upper bond. We show that such a probability is less than or equal to the probability of failing the review when the customer shares information truthfully and is evaluated at the end of period R only, and is said to fail the review if $I_R < q$ and pass the review if $I_R \geqslant q$. This is achieved in Lemma 4.

(5) Finally, we show that the characterized equilibrium yields almost full revelation of the customer's private information.

LEMMA 1. *Suppose the customer is always truthful in telling. For all n such that $1 \leqslant n \leqslant q$,*

$$v(I_{R-n}^{q-n}) = (1 - \delta)v^c \sum_{i=0}^{n-1}(S\delta)^i + v(I_0)(S\delta)^n$$
$$+ [(1 - \delta^M)v^0 + \delta Mv(I_0)] \sum_{i=1}^{n}(1 - S)\delta^i S^{i-1}.$$

PROOF. By induction. When $n = 1$, the customer either passes the review or the supply chain enters the punishment phase at the end of the period. Therefore,

$$v(I_{R-1}^{q-1}) = (1 - \delta)v^c + \delta Sv(I_0)$$
$$+ \delta(1 - S)[(1 - \delta^M)v^0 + \delta^M v(I_0)].$$

Suppose the induction hypothesis holds for n.

$$v(I_{R-(n+1)}^{q-(n+1)}) = (1 - \delta)v^c + \delta Sv(I_{R-n}^{q-n})$$
$$+ \delta(1 - S)[(1 - \delta^M)v^0 + \delta^M v(I_0)].$$

Now we show that it is true for $n + 1$. Substituting in the induction hypothesis yields:

$$v(I_{R-(n+1)}^{q-(n+1)}) = (1-\delta)v^c + \delta S\left\{(1-\delta)v^c \sum_{i=0}^{n-1}(S\delta)^i + v(I_0)(S\delta)^n\right.$$
$$\left. + ((1-\delta^M)v^0 + \delta^M v(I_0))\sum_{i=1}^{n}(1-S)\delta^i S^{i-1}\right\}$$
$$+ \delta(1-S)[(1-\delta^M)v^0 + \delta^M v(I_0)]$$
$$= (1-\delta)v^c \sum_{i=0}^{n-1}(S\delta)^i + v(I_0)(S\delta)^{n+1}$$
$$+ [(1-\delta^M)v^0 + \delta^M v(I_0)]$$
$$\cdot \sum_{i=1}^{n+1}(1-S)\delta^i S^{i-1}. \quad \square$$

LEMMA 2. *The customer will only fail the review when $I_t < \max(0, t - (R - q))$, i.e., when I_t is such that the customer will fail the review with probability 1.*

PROOF. To show that the customer will truthfully share his private information when I_t is such that $I_t \geqslant \max(0, t - (R-q))$, it is sufficient to show that the customer has an incentive to release his private information truthfully when $I_{R-n} = q - n$. This result is established in the following claim.

CLAIM Let $I_{R-n} = q - n$ and consider n such that $1 \leqslant n \leqslant q$. There exists $\underline{\delta}'$ such that for all $\delta \in [\underline{\delta}', 1)$, there exists an $\underline{S}$ and $\underline{R}'$ such that for all $S > \underline{S}$ and $R > \underline{R}'$ the customer truthfully reveals his private information.

PROOF. Consider a one-shot deviation when the customer observes θ_l. Define the following:

β is the conditional probability that the supplier will assess the customer's truthful forecast of $\theta = \theta_l$ and orders to be truthful;

η is the conditional probability that the supplier will assess the customer's behavior to be truthful when the customer observes $\theta = \theta_l$ and follows an optimal deviation strategy;

$v'(\theta_l)$ is the expected one-stage payoff when the customer observes θ_l and follows his optimal deviation;

$\Delta = v'(\theta_l) - v(\theta_l, K_l^c)$;

λ denotes the probability of failing the review conditional on being at the beginning of the game;

For any R, $I_{R-n} = q - n$ and any n such that $1 \leqslant n \leqslant q$, the customer would share information truthfully if:

$$(1 - \delta)v'(\theta_l) + \delta \eta v(I_{R-(n-1)} = q - (n-1)) + (1 - \eta)$$
$$\cdot [\delta(1 - \delta^M)v^0 + \delta^{M+1}v(I_0)]$$
$$\leqslant (1 - \delta)v(\theta_l, K_l^c) + \delta \beta v(I_{R-(n-1)} = q - (n-1))$$
$$+ (1 - \beta)[\delta(1 - \delta^M)v^0 + \delta^{M+1}v(I_0)].$$

Rearranging terms, we have

$$\Delta \leqslant \frac{(\beta - \eta)}{1 - \delta}\delta[v(I_{R-(n-1)} = q - (n-1))$$
$$- ((1 - \delta^M)v^0 + \delta^M v(I_0))].$$

Using Lemma 1, which gives the customer's expected profit when $I_{R-n} = q - n$ for all $1 \leqslant n \leqslant q$, and substituting for $v(I_{R-(n-1)} = q - (n-1))$ yields the condition:

$$\frac{(\beta - \eta)}{1 - \delta}\delta\left\{(1 - \delta)v^c\sum_{i=0}^{n-2}(S\delta)^i + v(I_0)(S\delta)^{n-1} + ((1 - \delta^M)v^0\right.$$
$$\left. + \delta^M v(I_0))\left(\sum_{i=1}^{n-1}(1 - S)\delta^i S^{i-1} - 1\right)\right\} \geqslant \Delta.$$

Using L'Hospital's rule, we have as $\delta \to 1$,

$$\frac{(\beta - \eta)}{-1}\left\{v(I_0)S^{n-1} + v(I_0)\left(\sum_{i=1}^{n-1}(1 - S)S^{i-1} - 1\right)\right.$$
$$\left. - v^c\sum_{i=0}^{n-2}S^i + (n-1)v(I_0)S^{n-1} + v'(I_0)S^{n-1}\right.$$

$$+ (-Mv^0 + Mv(I_0) + v'(I_0))\left(\sum_{i=1}^{n-1}(1 - S)S^{i-1} - 1\right)$$
$$\left. + v(I_0)(1 - S)\sum_{i=1}^{n-1}iS^{i-1}\right\} \geqslant \Delta$$

Using the fact that $\sum_{i=1}^{n-1}iS^{i-1} = (1 - S^n - nS^{n-1}(1 - S))/(1 - S)^2$, the above is simplified to be

$$(\beta - \eta)\left\{(v^c - v(I_0))\left(\frac{1 - S^{n-1}}{1 - S}\right) + S^{n-1}M(v(I_0) - v^0)\right\} \geqslant \Delta.$$

As prescribed in the review strategy, fix $M = \chi R$, $\chi > 0$. For large n, $S^{n-1}M \to 0$ and $(1 - S^{n-1})/(1 - S) \to 1/(1 - S)$. Thus, for all $\delta \in [\underline{\delta}', 1)$, there exists an $\underline{S}$ and an $\underline{R}'$ such that for all $S > \underline{S}$ and $R > \underline{R}'$, $((\beta - \eta)/(1 - S))[v^c - v(I_0)]$ will exceed Δ for all $I_{R-n} = q - n$. Hence, the customer's incentive constraint is violated and, consequently, fails the review only if $I_t < \max(0, t - (R-q))$. □

LEMMA 3. Let $I_0 = 0$. There exists a $\underline{\delta}''$ such that for all $\delta \in [\underline{\delta}'', 1)$, there exists an $\underline{R}''$ such that for $R > \underline{R}''$, the customer will truthfully share his private information for at least the first period of the review phase.

PROOF. Consider the following sufficient condition for the customer's truth-telling constraint for hold: The passing of the review phase is deferred until I_t is such that $I_{R-n} \geqslant q$. Because $v(I_0) \in [v^0, v^c]$, the truth-telling constraint will hold if the sufficient condition is imposed and the customer still forecasts truthfully.

Recall from the claim proven above that the failing of the review occurs only when $I_t < \max(0, t - (R-q))$. Therefore, as $\delta \to 1$, if

$$(\beta - \eta)\left\{(v^c - v(I_0))\left[\sum_{i=0}^{q-2}\binom{R-q+i}{R-q}S^i(1 - S)^{R-q}\right.\right.$$
$$\left. - \sum_{i=0}^{R-q-1}\binom{q-1+i}{q-1}(1 - S)^iS^{q-1}\right]$$
$$\left. + \binom{R-1}{R-q}S^{q-1}(1 - S)^{R-q}M(v(I_0) - v^0)\right\}$$
$$> \Delta,$$

then the customer will truthfully share his private information at the start of the review phase.

The left-hand side of the inequality has three major components. The first one, $(\beta - \eta)(v^c - v(I_0))\sum_{i=0}^{q-2}\binom{R-q+i}{R-q}S^i(1 - S)^{R-q}$, represents the gain from avoiding failing the review before period R. The review is terminated before period R if the customer has had $R - q$ assessments where the trust index is not incremented. In this case, for each assessment where the trust index *is* incremented the customer avoids termination for an additional period up to the beginning of the final review period, i.e.,

$R - 1$. The histories S^i and $(1 - S)^{R-q}$ can be arranged in $\binom{R-q+i}{R-q}$ combinations. Notice that for $i = q - 2$, the summand becomes $\binom{R-2}{R-q}S^{q-2}(1-S)^{R-q}$; this product multiplied by $(\beta - \eta)$ represents the probability of failing the review phase in period $R - 1$.

Second, $(\beta - \eta)(v^c - v(I_0))\sum_{i=0}^{R-q-1}\binom{q-1+i}{q-1}(1-S)^i S^{q-1}$ represents the marginal loss from passing the review before period R. The review is passed before period R if the customer has had q assessments where the trust index has been incremented before period R. If $I_t < q$, then the customer is assessed to have reported dishonestly and the passing of the review is delayed for an additional period. The passing of the review can be delayed for at most $R - q - 1$ periods prior to the beginning of the final period. These histories $(1 - S)^i$ and S^{q-1} can be arranged in $\binom{q-1+i}{q-1}$ combinations.

And, finally, $(\beta - \eta)\binom{R-1}{R-q}S^{q-1}(1 - S)^{R-q}M(v(I_0) - v^0)$ represents the marginal gain from avoiding the punishment phase in period R. This event occurs for a history of $R - q$ assessments where the trust index is not incremented and q assessments where it is incremented. The histories S^{q-1} and $(1 - S)^{R-q}$ can be arranged in $\binom{R-1}{R-q}$ combinations.

We now show that the constraint can always be satisfied if R is sufficiently large and if $M = \chi R$, where $\chi > 0$. Fix $M = \chi R$ and let $J = \tau R^\rho$ where $\tau > 0$, $\frac{1}{2} < \rho < 1$ for the remainder of the analysis. Recall that $q = R \cdot S - J$. Then let $R \to \infty$, and after applying Stirling's formula

$$\lim_{R\to\infty} \frac{R!}{\sqrt{2\pi}R^{(R+1/2)}e^{-n}} = 1$$

we get

$$(\beta - \eta)\binom{R-1}{R-q}S^{q-1}(1 - S)^{R-q}M(v(I_0) - v^0)$$

$$= \frac{\sqrt{R}(\beta - \eta)\chi(v(I_0) - v^0)}{\sqrt{2\pi}S(1 - S)} \to \infty.$$

Similarly, it can be shown that $\binom{R-q+i}{R-q}S^i(1 - S)^{R-q} < \binom{R-q+i+1}{R-q}S^{i+1}(1 - S)^{R-q}$ for all $i = 0, 1, 2, \ldots, q - 3$, and also that, as $R \to \infty$,

$$\binom{R-2}{R-q}S^{q-2}(1-S)^{R-q} = \frac{R \cdot S - 1}{S(R-1)}\frac{1}{\sqrt{2\pi R \cdot S(1 - S)}} \to 0.$$

Thus as $R \to \infty$, $\sum_{i=0}^{q-2}\binom{R-q+i}{R-q}S^i(1 - S)^{R-q} \to 0$, and

$$\sum_{i=0}^{R-q-1}\binom{q-1+i}{q-1}(1 - S)^i S^{q-1} \to 0.$$

Because Δ is bounded and the inequality is strict, it follows that for $R > \underline{R}''$ and $\delta \in [\underline{\delta}'', 1)$, the customer will truthfully share his private information and the supplier will trust the customer's forecast for at least the first period of the review phase. $\square$

Lemma 4. *The probability of failing the review when the players use the specified review strategies is less than or equal to the probability of failing the review when the customer reports in good faith and is evaluated at the end of period R only, and is said to fail the review if $I_R < q$ and pass the review if $I_R \geq q$.*

Proof. Because once the customer passes the review the review phase is terminated, it follows that the probability of failing the review under the specified review strategies (i.e., the customer can pass the review before the end of period R and fail the review if for any $t < R$, $I_t < \max(0, t - (R - q)))$ is less than or equal to the probability of failing the review if the customer can pass the review only at the end of period R, but can fail the review before period R if for any $t < R$, $I_t < \max(0, t - (R - q))$. To see this, choose some history that results in the review being passed in the former review specification but failed in the latter specification.

If the customer reports demand in good faith, then the probability of failing the review if the customer can pass the review only at the end of period R but can fail the review before period R if for any $t < R$, $I_t < \max(0, t - (R - q))$ equals the probability of failing the review when the customer is evaluated at period R only, and is said to fail the review if $I_R < q$ and pass the review if $I_R \geq q$. To prove this, consider: If a history occurs such that $I_R < q$, then it follows that there exists a history for the game such that $I_t < \max(0, t - (R - q))$. Conversely, if there exists a history such that $I_t < \max(0, t - (R - q))$, then $I_{t+(R-t)} < \gamma_t + (R - t) = q$. However, $I_{t+R-t} \leq I_t + (R - t) < q$. Thus, $I_R < q$. $\square$

Proof of Theorem 2. First, we note that neither the customer nor the supplier has an incentive to deviate. Given that the customer would share forecasts and order truthfully, the supplier's best response is to trust and build K_i^c. Therefore the supplier does not have a strict incentive to deviate from the characterized equilibrium strategies.

For the customer, the equilibrium is constructed so that if the truth-telling constraints are no longer satisfied, the review will be terminated, and if $I_t \geq \gamma_t$, the customer passes the review, or if $I_t < \gamma_t$, the customer fails the review. Thus, the customer follows a sequentially rational strategy and has no incentive to deviate.

Next, we show that under the specified strategy, the customer would almost always share forecasts and place orders truthfully. This is done with the following steps:

(1) The customer will only fail the review when $I_t < \max(0, t - (R - q))$, i.e., when I_t is such that the customer will fail the review for sure. A sufficient condition of the above is to show that the customer has an incentive to share forecasts truthfully when $I_{R-n} = q - n$, which we have proven in Lemma 2.

(2) From Lemma 3, there exist parameters of the game such that the customer will share forecasts truthfully at the beginning of the game.

(3) We now show that the characterized equilibrium yields almost full revelation of the customer's private information. Because the review phase may terminate earlier than at the end of period R, the following sufficient conditions are imposed to obtain a lower bound on the customer's normalized discounted expected profit: Assume that during the review phase the customer reports in good faith, but if he fails, then he fails the review at the earliest possible period, i.e., $t = R - q$; and assume that the customer reports in good faith and if he passes, then he passes the review at the earliest possible time period denoted as t'; recall that $v(I_0) \in [v^0, v^c]$ and $t' > 0$. The Markovian property of this repeated game yields a lower bound on the customer's normalized discounted expected profit, $v(I_0)$:

$$v(I_0) \geqslant \lambda \left[(1 - \delta) \sum_{i=1}^{R-q} \delta^{i-1} v^c + \delta^{R-q} (1 - \delta^M) v^0 \right.$$
$$\left. + \delta^{R-q+M} v(I_0) \right]$$
$$+ (1 - \lambda) \left[(1 - \delta) \sum_{i=1}^{t'} \delta^{i-1} v^c + \delta^{t'} v(I_0) \right].$$

Solving for $v(I_0)$ yields:

$$v(I_0) \geqslant \frac{(1 - \delta^{R-q}) v^c \lambda + \delta^{R-q} (1 - \delta^M) v^0 \lambda + (1 - \delta^{t'}) v^c (1 - \lambda)}{1 - \delta^{R-q+M} \lambda - \delta^{t'} (1 - \lambda)}.$$

$$(11)$$

Substitute for $M = \chi R$ and observe that $\lim_{R \to \infty} (M/(R - q)) = \chi/(1 - S)$. Because $t' > 0$, it follows that the left-hand side of the inequality in expression (11) converges uniformly in R as δ approaches 1 to

$$\frac{(v^c + (M/((R - q)) v^0) \lambda + (t'/(R - q)) v^c (1 - \lambda)}{((R - q + M)/(R - q)) \lambda + (t'/(R - q))(1 - \lambda)}.$$

Lemma 4 shows that the probability of failing the review when the players use the specified review strategies is less than or equal to the probability of failing the review when the customer reports in good faith and is evaluated at period R only, and is said to fail the review if $I_R < q$ and pass the review if $I_R \geqslant q$. This upper bound is denoted $\bar{\lambda}(I_R < R \cdot S - J)$. If the customer reports in good faith, then I_R is a binomial random variable. Invoking Chebychev's inequality implies that $\bar{\lambda}(I_R < R \cdot S - J) \leqslant R \cdot S(1 - S)/J^2$. If $J = \beta R^\rho$, where $\beta > 0$ and $\frac{1}{2} < \rho < 1$, then $\bar{\lambda}(I_R < R \cdot S - J)$ approaches zero as R increases without limit; because λ is bounded from above by $\bar{\lambda}(I_R < R \cdot S - J)$, $\lambda \to 0$.

Hence, for any $\varepsilon > 0$, there exists a discount factor $\underline{\delta}$ such that for all $\delta \in [\underline{\delta}, 1)$, there exists an $\underline{R}$ such that for all $R > \underline{R}$, $V(I_0) > v^c - \varepsilon$. This provides a lower bound on the customer's expected profits. Moreover, λ can be made arbitrarily close to zero by choosing R sufficiently large.

Acknowledgments

References

Arrow, K. J. 1985. The economics of agency. J. W. Pratt, R. J. Zeckhauser, eds. *Principals and Agents*. Harvard Business School Press, Cambridge, MA.

Bajari, P., S. Tadelis. 2001. Incentives versus transaction costs: A theory of procurement contracts. *RAND J. Econom.* **32**(3) 387–407.

Baker, G., R. Gibbons, K. J. Murphy. 2002. Relational contracts and the theory of the firm. *Quart. J. Econom.* **117**(1) 39–84.

Barnes-Schuster, D., Y. Bassok, R. Anupindi. 2002. Coordination and flexibility in supply contracts with options. *Manufacturing Service Oper. Management* **4**(3) 171–207.

Beth, S., D. N. Burt, W. Copacino, C. Gopal, H. L. Lee, R. P. Lynch, S. Morris. 2003. Supply chain challenges—Building relationships. *Harvard Bus. Rev.* **81**(7) 64–73.

Brown, A., H. Lee. 1998. Optimal "pay to delay" capacity reservation with application to the semiconductor industry. Working paper, Stanford University, Stanford, CA.

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

Cachon, G. P., M. A. Lariviere. 2005. Supply chain coordination with revenue-sharing contracts: Strengths and limitations. *Management Sci.* **51**(1) 30–44.

Chen, F. 2002. Information sharing and supply chain coordination. A. G. de Kok, S. Graves, eds. *Handbooks in Operations Research and Management Science: Supply Chain Management*. North-Holland, Amsterdam.

Chen, F. 2005. Salesforce incentives, market information, and production/inventory planning. *Management Sci.* **51**(1) 60–75.

Cohen, M. A., T. H. Ho, Z. J. Ren, C. Terwiesch. 2003. Measuring imputed cost in the semiconductor equipment supply chain. *Management Sci.* **49**(12) 1653–1670.

Corbett, C. J. 2001. Stochastic inventory systems in a supply chain with asymmetric information: Cycle stocks, safety stocks, and consignment stock. *Oper. Res.* **49**(4) 487–500.

Corbett, C. J., D. Zhou, C. S. Tang. 2004. Designing supply contracts: Contract type and information asymmetry. *Management Sci.* **50**(4) 550–559.

Crawford, V. P., J. Sobel. 1982. Strategic information-transmission. *Econometrica* **50**(6) 1431–1451.

Desai, P. S., K. Srinivasan. 1995. Demand signalling under unobservable effort in franchising: Linear and nonlinear price contracts. *Management Sci.* **41**(10) 1608–1623.

Durrett, R. 1996. *Probability: Theory and Examples*, 2nd ed. Duxbury Press, Belmont, CA.

Farlow, D., G. Schmidt, A. Tsay. 1996. Supplier management at Sun Microsystems (A). Stanford Graduate School of Business, Stanford, CA.

Friedman, J. W. 1971. Non-cooperative equilibrium for supergames. *Rev. Econom. Stud.* **38**(113) 1–12.

Fudenberg, D., J. Tirole. 1991. *Game Theory*. MIT Press, Cambridge, MA, ixxx, 579.

Fujimoto, T. 1999. *The Evolution of a Manufacturing System at Toyota.* Oxford University Press, New York.

Gibbons, R. 1992. *Game Theory for Applied Economists.* Princeton University Press, Princeton, NJ.

Ha, A. Y. 2001. Supplier-buyer contracting: Asymmetric cost information and cutoff level policy for buyer participation. *Naval Res. Logist.* **48**(1) 41–64.

Hagen, J. M., S. Choe. 1998. Trust in Japanese interfirm relations: Institutional sanctions matter. *Acad. Management Rev.* **23**(3) 589–600.

Iyer, G., J. M. Villas-Boas. 2003. A bargaining theory of distribution channels. *J. Marketing Res.* **40**(1) 80–100.

Lariviere, M. A., V. Padmanabhan, 1997. Slotting allowances and new product introductions. *Marketing Sci.* **16**(2) 112–128.

Levin, J. 2003. Relational incentive contracts. *Amer. Econom. Rev.* **93**(3) 835–857.

Liker, J. K. 2004. *The Toyota Way: 14 Management Principles from the World's Greatest Manufacturer.* McGraw-Hill, New York.

McMillan, J. 1990. Managing suppliers: Incentive systems in Japanese and United States industries. *California Management Rev.* **32**(4) 38–55.

Minahan, T. 1998a. Purchasing needs to do more than measure. *Purchasing* **124**(1) 59.

Minahan, T. 1998b. Is partnering a sham? *Purchasing* **124**(2) 61.

Morris, S. 2001. Political correctness. *J. Political Econom.* **109**(2) 231–265.

Osborne, M. J., A. Rubinstein. 1994. *A Course in Game Theory.* MIT Press, Cambridge, MA, xv, 352.

Ozer, O., W. Wei. 2006. Strategic commitment for optimal capacity decision under asymmetric forecast information. *Management Sci.* **52**(8) 1238–1257.

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

Plambeck, E. L., T. A. Taylor. 2006. Partnership in a dynamic production system with unobservable actions and noncontractible output. *Management Sci.* **52**(10) 1509–1527.

Porteus, E. L., S. Whang. 1991. On manufacturing marketing incentives. *Management Sci.* **37**(9) 1166–1181.

Pyke, D. F., E. M. Johnson. 2003. Sourcing strategies and supplier relationships: Alliances vs. eProcurement. T. Harrison, H. Lee, J. Neale, eds. *The Practice of Supply Chain Management: Where Theory and Application Converge.* Kluwer Academic Publishing, Boston, 77–89.

Radner, R. 1985. Repeated principal-agent games with discounting. *Econometrica* **53**(5) 1173–1198.

Spengler, J. J. 1950. Vertical integration and anti-trust policy. *J. Political Econom.* **58** 347–352.

Stocken, P. C. 2000. Credibility of voluntary disclosure. *RAND J. Econom.* **31**(2) 359–374.

Taylor, T. A. 2002. Supply chain coordination under channel rebates with sales effort effects. *Management Sci.* **48**(8) 992–1007.

Taylor, T. A., E. L. Plambeck. 2007a. Supply chain relationships and contracts: The impact of repeated interaction on capacity investment and procurement. *Management Sci.* **53** 1577–1593.

Taylor, T. A., E. L. Plambeck. 2007b. Simple relational contracts for capacity investment: Price-only versus quantity commitment. *Manufacturing Service Oper. Management* **9**(1) 94–113.

Terwiesch, C., Z. J. Ren, T. H. Ho, M. A. Cohen. 2005. An empirical analysis of forecast sharing in the semiconductor equipment supply chain. *Management Sci.* **51**(2) 208–220.

Tomlin, B. 2003. Capacity investments in supply chains: Sharing the gain rather than sharing the pain. *Manufacturing Service Oper. Management* **5**(4) 317–333.

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

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