MANAGEMENT SCIENCE
Vol. 51, No. 2, February 2005, pp. 208–220
ISSN 0025-1909 | EISSN 1526-5501 | 05 | 5102 | 0208

DOI 10.1287/mnsc.1040.0317

An Empirical Analysis of Forecast Sharing in the Semiconductor Equipment Supply Chain

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

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

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

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

We study the demand forecast-sharing process between a buyer of customized production equipment and a set of equipment suppliers. Based on a large data collection we undertook in the semiconductor equipment supply chain, we empirically investigate the relationship between the buyer's forecasting behavior and the supplier's delivery performance. The buyer's forecasting behavior is characterized by the frequency and magnitude of forecast revisions it requests (forecast volatility) as well as by the fraction of orders that were forecasted but never actually purchased (forecast inflation). The supplier's delivery performance is measured by its ability to meet delivery dates requested by the customers. Based on a duration analysis, we are able to show that suppliers penalize buyers for unreliable forecasts by providing lower service levels. Vice versa, we also show that buyers penalize suppliers that have a history of poor service by providing them with overly inflated forecasts.

Key words: forecast sharing; trust; empirical methods; supply chain management; collaborative planning
History: Accepted by William S. Lovejoy, operations and supply chain management; received January 15, 2003. This paper was with the authors 10 months for 2 revisions.

1. Introduction

Sharing demand forecast information has been recognized as a key element in supply chain coordination (Cachon 2001). Over the last decade, companies have engaged in various forecast-sharing practices, including the commonly known Collaborative Planning, Forecasting and Replenishment (CPFR) initiative, which was launched to "create collaborative relationships between buyers and sellers through comanaged processes and shared information."[1] Retailers such as Wal-mart and Best Buy, along with suppliers such as Procter & Gamble and Kimberly-Clark, have all reported substantial benefits from CPFR projects. For example, GlobalNetXchange, a consortium consisting of more than 30 trade partners including Sears, Kroger, Unilever, Procter & Gamble, and Kimberly-Clark, have reported a 5%–20% reduction in inventory costs and an increase in off-the-shelf availability of 2%–12% following the launch of their CPFR program (VICS CPFR Committee 2002).

Despite these success stories, forecast sharing still suffers from several problems in practice. In this article, we analyze two types of problems related to forecast sharing. First, forecasts change and are continually updated as the buyer receives new information about the demand it faces. This problem, which we refer to as *forecast volatility*, raises the question of when the forecast information provided by the buyer is sufficiently accurate to justify the supplier acting on it. A supplier that will act immediately on any given forecast will likely face significant future adjustment and rework costs.

Second, forecasts provide information about what the buyer intends to do in a given future state of the world. These intentions, however, are not verifiable and cannot be enforced. This makes contracting based on shared forecasts extremely difficult. In the absence of a contractual obligation for the buyer to purchase what it has forecasted, the buyer has an incentive to inflate forecasts to assure sufficient supply (*forecast*

Authors are named in reverse alphabetical order. All authors contributed equally.

[1] Website: http://www.cpfr.org.

Figure 1 Forecast Sharing and the Prisoner Dilemma

		Supplier	
		Cooperate (*trust forecast*)	*Do Not Cooperate* (*ignore forecast*)
Buyer	*Cooperate* (*forecast truthfully*)	Buyer forecasts truthfully and supplier trusts the forecast.	Buyer forecasts truthfully, but supplier waits until a firm purchase order is submitted (buyer incurs cost of delay).
	Do Not Cooperate (*inflate forecast*)	Buyer inflates forecast; supplier trusts the inflated forecast (supplier incurs cost of inventory and cancellation).	Buyer inflates forecast, supplier discounts forecasts and waits until firm purchase order is submitted.

inflation; see, e.g., Cachon and Lariviere 2001). Fearing inflated forecasts, the supplier might prefer to delay its actions to a point in time when the buyer is willing to commit to its forecast. This setup shares many similarities with the classical prisoner's dilemma: As is illustrated by Figure 1, both parties can either cooperate (buyer shares forecasts truthfully, and supplier trusts the forecast), achieving the Pareto-optimal outcome, or, as predicted by the one-period equilibrium model, they can decide to act noncooperatively (buyer inflates forecasts; supplier discounts forecast), foregoing the benefits of forecast sharing (shaded cell in Figure 1).

The extent to which the two parties will choose cooperative actions depends on the relevant planning horizon. Most of the existing analytical research on supply chain contracting considers one-shot games (Cachon and Netessine 2003). As demonstrated earlier (Cohen et al. 2003), this single period game induces the buyer to overforecast and the supplier to delay the initiation of a production order. More recently, there has been a growing interest in the supply chain literature (Taylor and Plambeck 2003, Debo 1999, Ren et al. 2004) and beyond (see, e.g., Sommer and Loch 2003 for an application in project management) in the role of trust and reputation in multiperiod games. This study complements this emerging area of research with an empirical foundation. Taking a multiperiod perspective, we demonstrate that both parties consider the outcome of previous periods when deciding whether they should cooperate in the present period.

Our study is grounded on detailed data related to forecast sharing and order fulfillment collected in the semiconductor equipment supply chain. We created a unique proprietary data set, capturing transactions between one buyer and 78 suppliers. Over a period of 2 years we collected data on more than 3,000 orders. This allows us to make the following contributions. First, we show that suppliers in the semiconductor equipment supply chain penalize the buyer for unreliable forecasts by delaying the fulfillment of forecasted orders. Specifically, we show that suppliers that have experienced large amounts of forecast volatility from the buyer are less willing to allocate capacity toward forecasted orders, leading to overproportionally long tool delivery times. Second, we show that suppliers that have been exposed to forecast inflation in the form of excessive order cancellations are less willing to allocate capacity toward forecasted orders, also leading to overproportionally long tool delivery times. Third, we show that the buyer penalizes those suppliers that have not been able to meet prior delivery requests by providing them with overly inflated forecasts. Together with the actions of the supplier, this penalty scheme from the buyer creates a "tit-for-tat" strategy, which is in line with earlier predictions from the economics literature for repeated prisoner dilemma games (e.g., Axelrod 1981, Kreps et al. 1982).

2. Research Setting

Our empirical analysis is based on a proprietary data set that we created in the semiconductor equipment industry. The data set consists of one buyer and a set of 78 suppliers. The buyer in our sample is one of the largest chip manufacturers in the industry and is the most important buyer of semiconductor equipment worldwide. This gives the buyer a substantial amount of power and allows it to implement forecast-sharing agreements that equipment suppliers might not agree to when dealing with smaller equipment buyers. This includes contracts design, information systems implementation, and requests for short delivery lead times. Given the technological complexity of the pieces of equipment requested by the buyer and the large amount of buyer-specific investments that suppliers incur, there exists only one supplier for every piece of equipment (i.e., for any piece of equipment, the buyer is committed to a single-sourcing strategy). While the powerful position of the buyer clearly limits the generalizability of our findings, it is advantageous from a research design perspective, as it holds the forecast-sharing mechanism constant across all 78 suppliers in our sample.

As in many customized capital goods industries, the semiconductor equipment supply chain faces an order-fulfillment dilemma. On the one hand, buyers of equipment expect their suppliers to be responsive and to be able to fulfill orders within a relatively short time. On the other hand, the high value and the customized nature of the product make it risky for the supplier to keep finished products or subsystems in inventory, leading to long and variable manufacturing lead times. Given the integral nature of the equipment, postponement strategies that have been found useful to shorten delivery times and to reduce inventory risks (e.g., Lee 1996) have not yet been implemented in this industry.

To resolve this dilemma, the buyers (producers of microchips) provide their equipment suppliers with order forecasts for 24 months and longer. Unlike

Figure 2 **Changes in Spending Levels in the Semiconductor Industry**

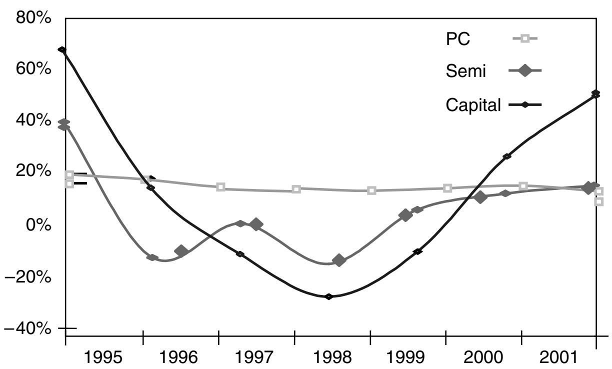

firm purchase orders, such forecasted orders—also referred to as "soft orders"—are a reflection of the buyer's purchase intent and are not legally binding.

Demand for semiconductor production equipment is triggered by the (projected) demand for chips, including microprocessors and memory chips. Given that the demand for chips is in turn generated by the demand for electronic devices, semiconductor equipment makers find themselves at the wrong end of the "bullwhip" (e.g., Lee et al. 1997). They face business cycles that flood them with orders one year and starve them for work the next (see Figure 2). The large chip producers create market forecasts on a monthly or quarterly basis. These forecasts are used to project production capacity needs for the next 2–5 years. Forecasts and capacity plans are updated on the basis of a rolling horizon principle. Chip manufacturers use these product-level demand forecasts, combined with equipment output models, to forecast capacity requirements to both existing and potentially new fabs. If the forecasted capacity requirement is not supported by the size and productivity of the installed equipment base, additional equipment must be ordered. This projected need for additional equipment is shared with equipment suppliers in the form of forecasted (soft) orders, consistent with the principle of forecast sharing and collaborative planning.

The chip manufacturer is unlikely to actually commit to purchase equipment at the time of the initial soft order placement. Over the next two years, the chip manufacturer will obtain new information about demand for chips as well as about the effective capacity of the currently installed equipment base (based on production yields, throughput time, and machine up time). As a result, the chip manufacturer may update the soft order and will usually delay making a firm order (i.e., issue a purchase order) until 3–6 months prior to the projected delivery date. This flexibility of the buyer, which delays a commitment until relatively close to the delivery date, reflects the buyer's strong bargaining position.

During the time between the initial placement of the soft order and the final placement of the purchase order, the buyer and the supplier continue to exchange information. Specifically, the buyer will inform the supplier about changes to the requested delivery date, the location of the fab where the tool will be operating, and other delivery-related information. In contrast to these delivery detail changes, the buyer does not change the specification of the equipment. This reflects the buyer's policy known as "copy exact" (see Terwiesch and Xu 2004 for details), which postulates that every piece of production equipment has to be absolutely identical. In the absence of specification changes, a soft order can be modified in one of the following two ways: (1) The requested delivery date might be moved forward or backward in time, reflecting new information the buyer has about detailed capacity planning at the fab. Given the high capital costs associated with acquiring the equipment, the buyer prefers to delay the requested delivery date rather than receive the equipment earlier than needed and having it be idle. (2) The soft order may be cancelled if market demand is less than initially projected or if existing equipment operates at higher yield levels or at a higher level of productivity. Alternatively, the soft order remains unchanged in the forecast-sharing system. Figure 3 shows the sequence of events for a soft order that is ultimately converted into a firm order.

Table 1 shows an example of four soft orders representative of the type of data we collected. This includes when the soft order was placed, how the requested delivery date changed, and whether the soft order ended up being purchased or being cancelled. Tool #197 has a stable forecast history but was cancelled six months after it was forecast. The requested delivery date for tool #199 changed three times. Tool #316 has a relatively stable forecasting history and was delivered earlier than requested. In contrast, tool #365 has a volatile forecasting history, with its requested delivery date changing widely from as early as 8/16/2000 to as late as 12/30/2000. This order ultimately was delivered almost 2 months later than requested.

Figure 4 shows an aggregation of order forecasts for one specific supplier. Each of the shared forecasts is a

Figure 3 **Events Leading to a Firm Order and Tool Delivery (Noncancellation Case)**

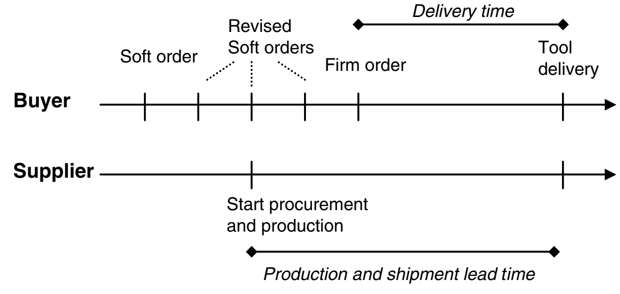

Table 1 **Sample Records**

				Milestones						Forecast history															
ELN	Supplier	Tool type	Site	Date Fcst added	Date Fcst deleted	Date order placed	Date order delivered	Requested LT	Actual LT	RDD 1099	RDD 1199	RDD 1299	RDD 0100	RDD 0200	RDD 0300	RDD 0400	RDD 0500	RDD 0600	RDD 0700	RDD 0800	RDD 0900	RDD 1000	RDD 1100	RDD 1200	
197	KE	LTS	A	9/27/ 1999	5/11/ 2000	—	—	—	—	11/1/ 2000	11/1/ 2000	11/1/ 2000	11/1/ 2000	11/1/ 2000	11/1/ 2000	11/1/ 2000	Cancelled	—	—	—	—	—	—	—	
199	AM	CMP	A	9/27/ 1999	3/9/ 2000	—	—	—	—	6/15/ 2000	6/15/ 2000	9/15/ 2000	9/29/ 2000	9/29/ 2000	9/29/ 2000	Cancelled	—	—	—	—	—	—	—	—	
316	TE	TCY	D	10/4/ 1999	—	5/18/ 2000	1/26/ 2001	10.07	9.04	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	2/23/ 2001	Delivered
365	NI	NSJ	F	10/5/ 1999	—	4/18/ 2000	10/30/ 2000	5.32	6.96	12/30/ 2000	12/30/ 2000	12/30/ 2000	12/30/ 2000	9/14/ 2000	9/14/ 2000	9/14/ 2000	8/16/ 2000	9/14/ 2000	9/14/ 2000	9/14/ 2000	9/14/ 2000	9/14/ 2000	Delivered	—	

Figure 4 Forecasted (Soft) Orders vs. Actual Orders

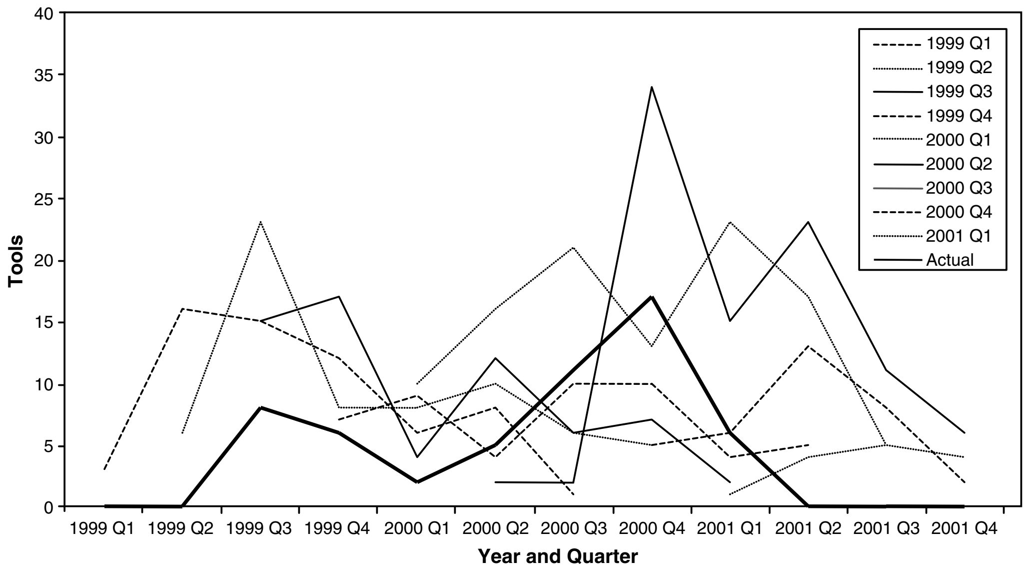

time series consisting of the seven quarters included in the relevant forecast window. For example, in Q2 2000, the buyer provides forecast quantities for the time interval from Q3 2000 to Q1 2002. We observe that forecasts vary widely, both over time (what is forecast in Q1 2000 for the period of Q2 2000 to Q4 2001) as well as from one forecast to the next (e.g., what is forecast in Q4 1998 for Q2 1999 vs. what is forecast in Q1 1999 for Q2 1999). Figure 4 also contrasts the forecasts with the actual tool purchases. On average, the buyer places significantly more soft orders than hard orders, suggesting the use of forecast inflation.

3. Research Objectives and Hypotheses

Our objective was to identify patterns of shared order forecasting that lead to an on-time tool delivery. Given that when a purchase order is placed, its production lead time exceeds the residual time available to the requested delivery date (see Figure 3), an on-time delivery requires that the supplier start working on an order while it was still a forecast (soft order).

Unfortunately, the effectiveness of working with forecast orders can be greatly reduced through noncooperative behavior of either party, buyer or supplier. The buyer can place more soft orders than it anticipates purchasing in the hope that this will secure him production capacity of the supplier. Vice versa, the supplier can discount or even ignore the information provided in the form of a soft order, knowing that it is the single supplier for a specific tool and that it would be legally almost impossible to hold it accountable for a delay. Consequently, the single period game between the buyer and the supplier resembles the traditional "prisoner's dilemma," which is known to have a Pareto inefficient equilibrium (Figure 1).

While playing a game once can lead to mistrust and a noncooperative outcome, the economics literature suggests that playing a game repeatedly can lead to more cooperative outcomes. Specifically, it has been argued that in the repeated game, parties are likely to adopt a "tit-for-tat" strategy, that is, cooperate (the buyer forecasts orders correctly on average and the supplier reacts to the forecast order) as long as the other party does the same and retaliate (the buyer overforecasts and the supplier ignores forecast orders) upon the other party's defection (Axelrod 1981, Kreps et al. 1982). The hypotheses derived below attempt to document that the buyer and supplier indeed follow such a "tit-for-tat" strategy.

The Supplier's Perspective

Consider the perspective of the supplier first. Given that the buyer has the right to change the delivery dates of soft orders and can cancel any open soft order, the supplier carries the risk of commencing production prior to receiving a firm order. However, since the supplier depends on the buyer for business for future technology generations, the supplier is unlikely to completely discount every piece of information he receives from the buyer. Instead, the supplier will evaluate the reputation of the buyer based on prior transactions, rewarding good forecasting behavior with early commencement of the production process and penalizing bad forecasting behavior with delays.

In our context, a buyer's bad forecasting behavior is constituted by two forces, forecast volatility and forecast inflation. Forecast volatility arises as forecast orders are based on preliminary information and made at a point in time that the buyer still faces substantial uncertainty about actual needs for the equipment. This uncertainty is likely to make the forecasts volatile, which in turn makes the supplier reluctant to commit resources to it. Forecast volatility has been analyzed by several prior studies (Heath and Jackson 1994, Graves et al. 1998, Cakanyildirim and Roundy 2002, and Kaminsky and Swaminathan 2001). Cattani and Hausman (2000) show that demand forecasts do not necessarily become more accurate as they are updated. They argue that such forecast churning can cause inefficiencies if the firm reacts to the wrong forecast update. A similar result has been provided by Toktay and Wein (2001). Similar observations have also been made in the coordination and project management literature.[2]

In our research setting, forecast volatility can take one of two forms: order-specific forecast volatility or buyer-specific forecast volatility. With order-specific volatility, we refer to the number of change requests the buyer places for a particular order.[3] In contrast, we label the number of change requests (across orders) the buyer has placed with the supplier as buyer-specific forecast volatility.[4] Buyer-specific forecast volatility thereby captures the recent history of forecast behavior of the buyer.

Hypothesis 1a (Order-Specific Forecast Volatility). *The more the customer changes the requested delivery date of a particular soft order, the more likely this particular order will be delayed.*

Hypothesis 1b (Buyer-Specific Forecast Volatility). *The more the buyer has changed requested delivery dates for soft orders in the past, the more likely it is that the current order will be delayed.*

A second reason why a supplier might not be willing to initiate work for a soft order relates to the perceived probability of order cancellation. Given the complex and capital-intense production process of semiconductor manufacturing, the buyer faces severe costs if the equipment does not arrive on the required delivery date. Late shipments of equipment—and consequently late availability of capacity—can lead to idle time for other equipment in the fab and potentially lost wafer output. Industry observers estimate that a 1-hour delay in installing capacity of a fab is worth in excess of $100,000. This creates an incentive for the buyer to provide overly aggressive forecasts to the supplier, that is, place more soft orders than firm orders. As the real capacity needs of the buyer are unobservable to the supplier, the buyer can always cancel the order and justify such change on information that is not verifiable by the supplier, for example, an unexpected drop in demand or increased production yields from existing equipment. Note that, in contrast to forecast volatility, which would also exist in a vertically integrated firm, forecast inflation reflects an opportunistic (noncooperative) behavior of the buyer.

Forecast inflation has been analyzed by Lee et al. (1997), Celikbas et al. (1999), and Cachon and Lariviere (2001). While these models are based on one-shot games, there has been a growing interest in the role of trust and reputation in supply chains from a multiperiod perspective (Taylor and Plambeck 2003, Debo 1999, Ren et al. 2004, Cachon and Netessine 2003). These studies, directly or indirectly, fit the repeated prisoner's dilemma framework outlined in Figure 1 and hence predict that the supplier will penalize the buyer for order cancellations by providing longer delivery times.

Hypothesis 1c (Forecast Inflation). *Past soft-order cancellations prolong current order lead time. That is, the more frequently the buyer has cancelled soft orders in the past, the more likely it is for the supplier to delay production, which leads to longer order lead time.*

Cancelled orders are especially costly to the supplier while operating at full capacity, as in such cases the cancellation costs include not only costs of inventory and procurement, but also the cost of lost business. We therefore extend our hypothesis as follows:

Hypothesis 1d (Forecast Inflation in Economic Upturn). *The delay from order cancellation is more severe during an economic upturn.*

The Buyer's Perspective

While cooperation from the supplier's perspective means reacting to the forecasted orders provided by the buyer, cooperation from the buyer's perspective means providing realistic estimates for the forecasted orders. To the extent that buyer and supplier indeed follow a tit-for-tat strategy, the buyer will react to noncooperative behavior of the supplier by acting noncooperatively itself.

In the buyer's eyes, noncooperative supplier behavior is characterized by late deliveries of equipment. Although the action of the supplier itself is not observable to the buyer, the buyer can estimate supplier cooperation based on delivery dates; everything

[2] See Krishnan et al. (1997), Loch and Terwiesch (1998), and Roemer et al. (2000) for models of sharing preliminary information in which the manager of an information-receiving task needs to decide when it is willing to commit resources to information supplied by other, concurrently executed, tasks.

[3] Consider, for example, a supplier in February 2001 that received a soft order in May 2000 with an initially requested delivery date of July 2001. However, between May 2000 and February 2001, the soft order has been modified (e.g., pushed out) multiple times.

[4] Consider, again, a supplier that has received a soft order in May 2000 with a requested delivery date of July 2001. In January 2001, the supplier considers initiating the order-fulfillment process. Yet, from prior experience with the same buyer, the supplier knows that in more than half of the cases the buyer has delayed the requested delivery date up to five months from the initially requested delivery date.

else being equal, a supplier with late equipment deliveries is more likely to have engaged in noncooperative behavior than a supplier that has delivered on time.

Once the buyer has decided to punish a supplier, it can do so by placing soft orders and then cancelling them overproportionally often compared to the case of cooperation. In absence of forced compliance (Cachon and Lariviere 2001), this is the only punishment mechanism available to the buyer during the interaction with the supplier for this tool generation. We therefore hypothesize—

Hypothesis 2 (Forecast Inflation). *Past delivery delays lead to an increase in future cancellations.*

4. Model Specification

We model the evolution of a soft order to a firm order and ultimately to a delivered piece of equipment in the form of a two-stage process. The first stage captures the fact that soft orders can either end up as firm orders, that is, the buyer places an order, or be cancelled. A firm order will see a delivery time that consists of the elapsed time between the placement of the firm order and its arrival at the customer's fab. These two stages are summarized by Figure 5.

Let (s, j) denote the index of the jth soft order the buyer places with supplier s. We use a logit formulation to describe the probability that this soft order is transformed into a firm order:

$$\Pr_{s,j}(\text{firm order}) = \frac{1}{1+\exp(x_{s,j}\beta)}, \tag{1}$$

where $x_{s,j}$ is a vector of explanatory variables and β is a parameter vector of appropriate dimensionality. Since any soft order will either be transformed into a firm order or be cancelled, the probability of cancellation is

$$\begin{aligned}\Pr_{s,j}(\text{cancel}) &= 1-\Pr_{s,j}(\text{firm order})\\ &= \frac{\exp(x_{s,j}\beta)}{1+\exp(x_{s,j}\beta)}.\end{aligned} \tag{2}$$

On placement, a firm order will experience a strictly positive delivery lead time. We model the duration between the buyer's placing a firm order and its delivery by the supplier using a hazard rate model (Cox 1972). Using the hazard rate as a dependent variable rather than the actual delivery lead time has several advantages. First, durations may have a nonnormal distribution. Restricted to being positive, they are often skewed. Thus the normality assumption of standard regression is violated. Second, hazard rate models should be chosen instead of standard regression analysis when working with survival data (Helsen and Schmittlein 1993). In our case, performing a regression analysis on only those soft orders that have been delivered would lead to a right-censoring of the data, as many of the soft orders we traced were not yet delivered at the end of our data collection. Finally, hazard models are also capable of capturing interesting dynamics of durations, such as the change in hazard rate over time, which can lead to additional insights in the underlying dynamics of the order fulfillment process.

Figure 5 **Two-Step Model**

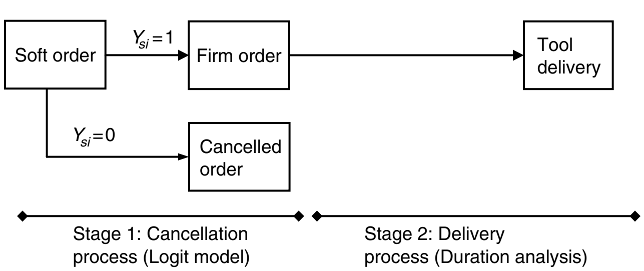

Despite their advantages, standard hazard rate models require that observations be independent of each other. This may be reasonable in the context of a medical lifetime study, yet in a manufacturing environment like the one we study, the lead time of one order is likely to be positively correlated with the lead time of the subsequent order at the same supplier. Such correlation reflects congestion effects: A long lead time for one particular order will increase the probability of the next order in the production pipeline also experiencing a long lead time. Consequently, the independence assumption is violated and a refined model specification is needed.

Let (s, i) denote the index of the ith firm order at supplier s, and let I_s be the number of firm orders received by supplier s. Define random variables $T_{s,i}$ as the logarithm of the duration between the placement of the firm order (s, i) and the delivery date of the equipment. Let $t_{s,i}$ be the realizations of these random variables. In our estimation, we assume that the congestion at the supplier can be captured via a first-order correlation between lead times for this supplier: Two orders close together (or with overlapping lead times) will be more tightly correlated than two orders that are far apart in time. Specifically, define the hazard rate of one completed order (s, i) conditional upon the completion time of the preceding order to the same supplier (i.e., the order that the supplier received directly before order i), $(s, i-1)$, as

$$h(t_{s,i} \mid t_{s,i-1}) = h_0(t_{s,i} \mid t_{s,i-1}) \cdot \exp(z_{s,i}\alpha), \tag{3}$$

where $h_0(t_{s,i} \mid t_{s,i-1})$ is the correlated baseline hazard function, $z_{s,i}$ is a vector of explanatory variables, and α is a parameter vector of appropriate dimensionality. According to Cox (1972), the baseline hazard function is

$$h_0(t_{s,i} \mid t_{s,i-1}) = \frac{f(t_{s,i} \mid t_{s,i-1})}{1-F(t_{s,i} \mid t_{s,i-1})}, \tag{4}$$

where $f(t_{s,i} \mid t_{s,i-1})$ ($F(t_{s,i} \mid t_{s,i-1})$) is the conditional normal density (distribution) function for the ith order at supplier s, given that the lead time of the preceding order to the same supplier (i.e., order $s, i-1$) is $t_{s,i-1}$. Given identical marginal means μ and standard deviations σ, as well as a correlation coefficient ρ for the unconditional bivariate normal distribution, it follows that $T_{s,i} \mid (T_{s,i-1} = t_{s,i-1}) \sim N(\mu + \rho(t_{s,i-1} - \mu), \sigma^2(1-\rho^2))$. In order to formally test to what extent the log-normal distribution indeed represents the delivery durations in our sample, we performed both a Kolmogorov-Smirnov test as well as a traditional Chi square test (see Law and Kelton 1991 for details). Both tests supported our assumption—that is, the hypothesis of log-normality could not be rejected. The importance of the correlation coefficient, ρ, will become apparent in the estimation results of our model.

Define an indicator variable $r_{s,i} = 0$ if the duration is censored (i.e., the firm order was not completed at the time of our data collection), and $r_{s,i} = 1$ if it is not censored. Then the likelihood contribution, that is, the probability of observing duration $t_{s,i}$ conditional upon it being firm ordered, is (Kalbfleisch and Prentice 1980):

$$\Pr(t_{s,i} \mid t_{s,i-1}) = [f(t_{s,i} \mid t_{s,i-1})]^{r_{s,i}}[1 - F(t_{s,i} \mid t_{s,i-1})]^{1-r_{s,i}}.$$

Given our assumption of first-order correlation, we can write the likelihood contribution of observing the vector $(t_{s,1}, \ldots, t_{s,I_s})$ of delivery times at supplier s as

$$\Pr_s(t_{s,1}, \ldots, t_{s,I_s}) = \Pr(t_{s,1}) \cdot \Pr(t_{s,2} \mid t_{s,1}) \cdot \cdots \cdot \Pr(t_{s,I_s} \mid t_{s,I_s-1}). \quad (5)$$

Finally, we obtain the log-likelihood function of the complete two-stage model:

$$LL(\alpha, \beta, \mu, \sigma, \rho) = \sum_s \left\{ \left[\sum_j \ln(\Pr_{s,j}(\text{firm order})) + \ln(\Pr_{s,j}(\text{cancel})) \right] + \ln(\Pr_s(t_{s,1}, \ldots, t_{s,I_s})) \right\} \quad (6)$$

5. Construct Definition

Over the period from September 1999 to July 2001 we collected data on all soft and firm orders the buyer placed with his 78 equipment suppliers, leading to a total of 3031 observations. The econometric model specified above uses two dependent variables. For the first stage, the dependent variable is binary, with a value of 1 denoting that the soft order was converted into a firm order and a value of 0 denoting a cancellation. In total, 53.2% of the soft orders were converted into firm orders. For the second stage, the dependent variable is the duration between the placement of the firm order and the delivery of the equipment to the buyer's fab.

In addition to these dependent variables, our hypotheses include the following set of explanatory variables. For a given soft order, we measured order-specific volatility (ORDER_VOLA) as the amount of due date change (forward or backward in time) that this soft order has experienced prior to becoming a firm order. In other words, we added up the absolute value of all due date changes this soft order experienced. For example, a soft order that was initially placed for May 2002, moved forward to March 2002, and finally moved back to June 2002 would have a score of $2 + 3 = 5$ months. Similarly, we measured buyer-specific volatility (BUYER_VOLA) for a given soft order as the average amount of due date change (forward or backward in time) across all soft orders the buyer submitted to the supplier within the last three months prior to this soft order. Both, BUYER_VOLA and ORDER_VOLA, are measured in months. BUYER_VOLA ranged between 0 and 16.4 months, with an average of 3.76 months. In our data set, ORDER_VOLA ranged from 0 to 51.2 months. The average was—coincidentally—also 3.76 months. Forecast inflation was measured by comparing the number of soft order cancellations over the past three months to the total number of (soft and firm) orders. The corresponding ratio, which we label as CANCEL, can be interpreted as the probability of order cancellation.

We measured the overall economic conditions by including the industry's book-to-bill ratio, as defined and tracked by Semiconductor Equipment and Materials International. It is defined as a ratio of the three-month moving average bookings to the three-month moving average shipments for the North American semiconductor equipment industry. This statistic characterizes the relative balance of supply and demand in the industry. If the ratio is larger than 1, demand exceeds current supply. We defined a binary variable, BOOK_BILL, that was equal to 1 if demand exceeds supply (indicating an economic upturn) and 0 otherwise. Finally, we measured the past delivery performance of the supplier for a given soft order as the total delay across all tool deliveries that occurred within the last six months of this soft order. The mean value of this variable, which we label as PAST_LATE, was 0.14 month.

In addition to the variables relating directly to our hypotheses, we included several control variables in our analysis. First, we included a binary variable DEV_FAB to indicate if the corresponding tool was requested by a development fab. Development fabs play a crucial role in the development of new equipment technologies and thereby order tools only at the very beginning of the tool's product lifecycle. About 19% of the tools in our sample were ordered for a development fab. We expected tools for development fabs to take longer compared to tools shipped to high volume manufacturing facilities.

A second tool characteristic reflected differences between the traditional 8-inch wafer technology and the new 12-inch technology. A binary variable NEW_TECH was set equal to 1 if the corresponding tool was based on 12-inch technology. Since mid 1999, fabs have been gradually shifting toward using wafers of 12-inch diameter, which leads to a much higher number of chips on a wafer and consequently improved productivity. Roughly 10% of tool orders in our sample were for the new 12-inch technology. Tools for the 12-inch technology were expected to require longer lead times compared to tools based on 6- or 8-inch technologies.

Third, we used the variable TOOL_PRICE to reflect the price of the tool as stated in the contract between buyer and supplier. Prices for tools in our sample averaged around $1.4 million but in some cases went as high as $10 million per tool. We expected expensive tools to have longer lead times, reflecting that expensive tools are typically based on more complex technologies. Fourth, we defined a binary variable FOREIGN indicating if a tool was requested for a non-US fab. Production in these fabs, all of which are owned by the buyer, was managed locally and our interviews suggested differences between the behavior of fabs in the United States and abroad. About 16% of the tools were for non-US fabs.

Fifth, about 8.5% of the tools in our sample were reused tools, that is, tools that were initially built based on an older technology and then upgraded to be usable for the latest process technologies. Such upgrades, also referred to as converted tools, require that the tool's critical components be replaced. A binary variable CONVERTED is equal to 1 if the tool has been converted at least once. Converted tools are expected to have shorter lead times.

Sixth, and finally, we needed to control for the lead time requested by the buyer when writing a purchase order to the supplier (REQ_LEADT). The fact that a tool with a long requested lead time takes longer until it is delivered had nothing to do with our research focus on forecast sharing. It is the deviation from this requested lead time that was of interest to us. The average requested lead time was about 5 months.

6. Estimation Results

To test our hypotheses, we specified and estimated a sequence of five models. The specifications as well as the parameter estimates are reported in Table 2. Model 1 contains a constant and the control variables DEV_FAB, FOREIGN, TOOL_PRICE, CONVERTED, NEW_TECH, and BOOK_BILL, and—for the duration analysis only—the requested lead time REQ_LEADT. The effect of the control variables are as predicted.

All models indicate that the correlation coefficient between subsequent orders to the same supplier is

Table 2 **Estimation Results**

Model parameters		Model 1	Model 2	Model 3	Model 4	Model 5
β	Constant	−0.001 (0.0001)	−0.001 (0.0001)	−0.001 (0.0001)	−0.001 (0.0001)	−0.020 (0.0001)
	DEV_FAB	−1.387 (0.0009)	−1.387 (0.0009)	−1.387 (0.0009)	−1.387 (0.0009)	−1.389 (0.0002)
	FOREIGN	0.6829 (0.0004)	0.6829 (0.0004)	0.6829 (0.0004)	0.6829 (0.0004)	0.8224 (0.0003)
	TOOL_PRICE	0.0857 (0.0001)	0.0857 (0.0001)	0.0857 (0.0001)	0.0857 (0.0001)	0.0934 (0.0003)
	CONVERTED	−0.5290 (0.0002)	−0.5290 (0.0002)	−0.5290 (0.0002)	−0.5290 (0.0002)	−0.5686 (0.0003)
	NEW_TECH	0.5442 (0.0018)	0.5442 (0.0018)	0.5442 (0.0018)	0.5442 (0.0018)	0.5238 (0.0002)
	BOOK_BILL	−0.009 (0.0002)	−0.009 (0.0002)	−0.009 (0.0002)	−0.009 (0.0002)	−0.019 (0.0001)
	PAST_LATE					0.190 (0.0003)
α	Constant	1.043 (0.0034)	1.317 (0.0041)	1.360 (0.0036)	1.329 (0.0042)	1.329 (0.0042)
	DEV_FAB	−0.075 (0.0022)	−0.155 (0.0025)	−0.117 (0.0030)	−0.138 (0.0031)	−0.138 (0.0031)
	FOREIGN	0.456 (0.0015)	0.415 (0.0041)	0.422 (0.0041)	0.403 (0.0041)	0.403 (0.0041)
	TOOL_PRICE	−0.109 (0.0006)	−0.103 (0.0006)	−0.091 (0.0006)	−0.093 (0.0006)	−0.093 (0.0006)
	CONVERTED	0.299 (0.0029)	0.246 (0.0031)	0.325 (0.0032)	0.324 (0.0043)	0.324 (0.0043)
	NEW_TECH	−0.369 (0.0013)	−0.413 (0.0015)	−0.347 (0.0041)	−0.335 (0.0043)	−0.335 (0.0043)
	BOOK_BILL	−0.147 (0.0004)	−0.200 (0.0021)	−0.212 (0.0022)	−0.070 (0.0004)	−0.070 (0.0004)
	REQ_LEADT	−0.127 (0.0021)	−0.145 (0.0004)	−0.146 (0.0028)	−0.147 (0.0027)	−0.147 (0.0027)
	CANCEL			−1.022 (0.0144)	−0.491 (0.0153)	−0.491 (0.0153)
	CANCEL * BOOK_BILL				−2.347 (0.0304)	−2.347 (0.0304)
	BUYER_VOLA		−0.036 (0.0006)	−0.032 (0.0006)	−0.031 (0.0005)	−0.031 (0.0005)
	ORDER_VOLA		−0.026 (0.0003)	−0.031 (0.0003)	−0.029 (0.0003)	−0.029 (0.0003)
μ		1.682 (0.0005)	1.718 (0.0012)	1.718 (0.0013)	1.722 (0.0013)	1.722 (0.0013)
ρ		0.172 (0.0005)	0.172 (0.0047)	0.167 (0.0044)	0.167 (0.0006)	0.167 (0.0006)
LL (In sample)		−2,825.808	−2,817.200	−2,814.833	−2,812.167	−2,807.576
LL (Out of sample)		−2,682.620	−2,678.003	−2,674.764	−2,668.180	−2,659.400

significant and positive. The actual estimates range between $\rho = 0.167$ and $\rho = 0.172$. This significant correlation captures the effect of congestion in the buyer's production facility: If the nth order from a given supplier is experiencing a longer-than-average lead time, chances are that the $(n+1)$st order will also be delayed. Thus, our extension of the traditional duration analysis to include first-order correlation was indeed necessary.

Consider Hypothesis 1a (Order-Specific Forecast Volatility) and Hypothesis 1b (Supplier-Specific Forecast Volatility) first. As shown by Model 2, forecast volatility indeed leads to longer delivery duration, as indicated by the negative coefficient of BUYER_VOLA and ORDER_VOLA. Moreover, comparing the log-likelihood of Model 2 to that of Model 1, we find that adding these variables improves the explanatory power of the model. This is indicated by the significant likelihood ratio test as reported in Table 3.

Interestingly, we observe that BUYER_VOLA's impact (ranging from −0.031 to −0.036) is stronger than that of ORDER_VOLA (ranging from −0.026 to −0.031), which suggests that the long-run effect of supplier reputation is more profound than the short-term effect of changing a single order. Based on the relationship between the hazard rate and the expected lead time, we obtain the marginal effect on lead time of an increase of BUYER_VOLA. Each month of delivery date change results in an average of 0.25 month of additional delay. Thus, for every month the buyer changes the requested delivery date of an order, it will experience a 0.25-month increase in expected lead time. A one-month increase in the average change in requested delivery date will lead to a 0.16-month increase in expected lead time.

Model 3 indicates that an increase in cancellation (CANCEL) will lead to a significant decrease in the hazard rate, which is in line with Hypothesis 1c. Moreover, as shown by Model 4, the business cycle, as indicated by the book-to-bill ratio (BOOK_BILL), has a strong interaction effect with the forecast inflation measure CANCEL, confirming Hypothesis 1d. During a business upturn (BOOK_BILL = 1), the delaying effect of CANCEL increases drastically (from −0.491% to −2.838% in elasticity across models). This confirms our hypothesis that cancellations prolong delivery times more profoundly during an economic upturn.

Table 3 **Likelihood Ratio Test**

	Model 1	Model 2	Model 3	Model 4	Model 5
LL	−2,825.808	−2,817.200	−2,814.833	−2,812.167	−2,807.576
LR		17.216 (Model 2 vs. Model 1)	4.734 (Model 3 vs. Model 2)	5.332 (Model 4 vs. Model 3)	9.182 (Model 5 vs. Model 4)
d.f.		2	1	1	1
p value		0.000	0.030	0.021	0.002

Our results suggest an increasingly delaying impact of CANCEL on the delivery time. Moreover, the state of the economy, represented by the book-to-bill ratio, aggravates such negative impact drastically. The impact from each additional percentage increase in CANCEL ranges from 7.6 days (CANCEL = 0%) to 14.1 days (CANCEL = 45%) during an economic downturn. The impact becomes substantially more profound during an economic upturn, ranging from a 19.5-day (CANCEL = 0%) delay to a delay of 91.2 days (CANCEL = 45%). Thus, a 1 percentage point increase in cancellation frequency leads to an increase of 1.59 days in delivery duration.

Finally, Model 5 tests the hypothesized effect of prior late shipments on the cancellation probability. Based on the significant coefficient of PAST_LATE in Model 5, we find also Hypothesis 2 supported. The coefficient of 0.190 indicates that a one-week lateness in previous shipments will increase the likelihood of future order cancellations by 19 percentage points. This complements the tit-for-tat perspective to the repeated buyer interaction discussed in the Introduction.

7. Model Validation

To validate the robustness of our results with respect to our construct definition, we used alternative measures for buyer volatility (BUYER_VOLA) and cancellation probability (CANCEL). In addition to measuring these constructs based on the last three months as defined above, we varied the "memory" of these variables to six and nine months. Similarly, for the past shipment delays from the supplier (PAST_LATE), we used a time window of three and nine months. All our findings reported in Table 2 remained structurally unchanged.

To validate the robustness of our results with respect to our sample composition, we ran our analysis with and without the converted tools. Again, all results of Table 2 remained structurally unchanged.

To test the validity of our logit model (first stage), we calculated its ability to correctly predict if a soft order would become a firm order as opposed to being cancelled. Our logit model predicts more than 70% of the binary outcomes correctly, which is in line with previous applications of logit models.

To test the validity of our duration analysis (second stage), we performed a May-Hosmer test. The test is based on a comparison of the observed number of deliveries with the expected number of deliveries as predicted by the duration analysis (see May and Hosmer 1998). The test first requires calculating the estimated risk score $z\hat{\alpha}$ for each observation and then grouping the subjects into subgroups indexed $g = 1, \ldots, G$. For each subgroup, we compute and compare the observed and the expected number of uncensored deliveries. A large p-value (typically

Table 4 May-Hosmer Test

	Actual	Expected	z Score	p Value
Decile of risk score				
1	2	0.218	3.811	0.000
2	1	0.348	1.106	0.269
3	3	2.142	0.586	0.558
4	23	18.245	1.113	0.266
5	48	53.912	−0.805	0.421
6	177	169.569	0.571	0.568
7	545	575.146	−1.257	0.209
8	607	629.92	−0.913	0.361
9	259	249.264	0.617	0.537
10	51	30.188	3.788	0.000

greater than 10%) accepts the hypothesis that there is no significant difference between the observed number of deliveries and the expected number of deliveries and therefore indicates a good model prediction. The test results are reported in Table 4.

We observe that our model performs well except for the first and the last decile. The first decile is not of significance because the corresponding subgroup only contains two observations. The 10th decile has 51 observed deliveries, compared to 31 predicted deliveries. This is due to the fact that the risk score subgroup contains observations with unusually large risk scores, and our model fails to predict those outliers. For the other groups, which contain 96.5% of the observations in our sample, the test results show that our model predicts well, with p-values all greater than 10%.

The overall model fit is visualized by plotting the actual observed durations against the fitted durations (Figure 6). Toward this end, we increase the number of subgroups to 100. A perfect model fit would lead to points lying on the 45-degree line in the graph. The points obtained from our model are overall close to the 45-degree line, indicating a good fit. This is formalized by the following regression analysis:

$$\text{Predicted} = -0.28 + 0.94^{*} \times \text{Observed}.$$

*indicates significance at 0.1% level. $R^2 = 90.5\%$.

Despite this good fit, it should be emphasized that our empirical findings might not directly generalize to other supply chain settings. The strong buyer, the fast-changing technology, and the complexity of the orders clearly differentiate the semiconductor equipment supply chain from many other industrial settings. Empirical research in other industries is needed to overcome this limited generalizability.

8. Conclusion

Forecast sharing has the potential to dramatically improve supply chain performance. Yet, as demonstrated by our research findings, a supply chain might not be able to achieve the potential performance improvements from forecast sharing. From the perspective of the supplier, the forces that prevent effective forecast sharing are forecast volatility and forecast inflation. Forecast volatility arises because forecasts are based on preliminary information and made at a point in time at which the equipment buyer still faces substantial uncertainty about the market demand for chips as well as of the capacity of the currently installed production equipment. As the buyer is exposed to additional information, it updates its forecasts to the supplier. While always sharing the latest information with the supply chain seems like a reasonable behavior for the buyer, frequent updates of information are perceived as disturbing from the perspective of the supplier. As we showed with respect to Hypothesis 1a, the supplier views a soft order that has been changed already multiple times as less reliable than a soft order that has not yet been changed. Consequently, the supplier is not willing to allocate production capacity to this soft order. Hypothesis 1b demonstrates that frequent changes to one soft order have externalities on how the supplier views future soft orders. Specifically, the more a buyer changes the requested delivery dates for equipment, the more the supplier will wait for the forecasts to stabilize when considering subsequent soft orders.

Forecast inflation can occur in the semiconductor equipment supply chain, as the buyer has an incentive to create overly aggressive forecasts. Forecast inflation is facilitated in this setting because shared forecasts are not verifiable and thus the supplier will never be able to validate whether actual inflation occurred. However, as we demonstrate in conjunction with Hypothesis 1c, frequent forecast inflation can

Figure 6 Actual vs. Fitted Durations

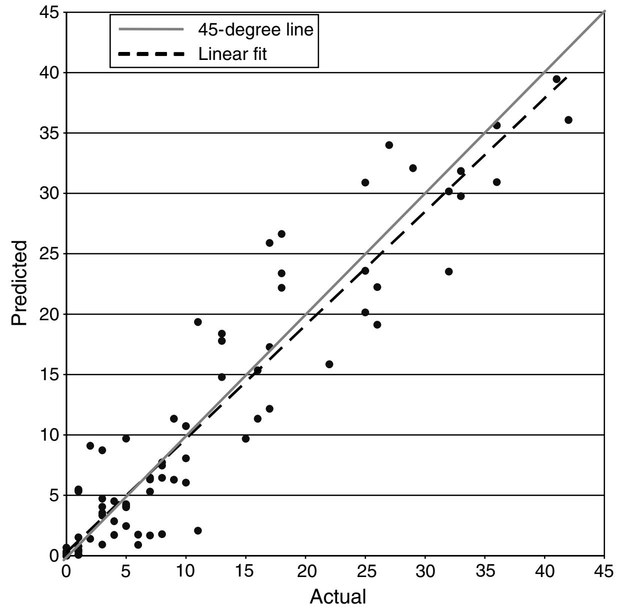

hurt the buyer in the long run. This penalty for past cancellations is especially severe during an economic upturn, during which the supplier has many other profitable opportunities to use its production capacity (Hypothesis 1d).

As does the supplier, which penalizes the buyer for inflated forecasts through longer delivery times, the buyer provides more aggressive forecasts to those suppliers that have failed to deliver previous orders on time (Hypothesis 2). This follows the logic of the repeated prisoner's dilemma game and establishes that both buyer and supplier apply a tit-for-tat strategy.

Our empirical research findings and our multiperiod framework of forecast sharing open up interesting opportunities for future research. First, we believe additional research is needed to analyze supply chain coordination in repeated game settings. While repeated games have been extensively studied in the economics literature, most of the contracting research in operations management has taken a rather static perspective, ignoring effects of trust building and reputation.

Second, one needs to overcome the forecast volatility problem. Currently, forecasts provided by the buyer do not acknowledge that they are based on preliminary information and are likely to change. Thus, while the buyer shares the expected outcome for a particular equipment order in the form of a best guess, it does not relay information reflecting possible alternative outcomes as well as the probabilities that such alternative outcomes occur. The supplier in turn perceives the almost unavoidable iterations as an indication that the shared forecasts are of low quality and consequently is not willing to commit resources based on this information. Recent research related to the information sharing in teams outlines alternative approaches to this (Terwiesch et al. 2002), including the concept of sharing information in the form of sets, which are gradually narrowed over time, rather than sharing information in the form of points, which "jump around" in an unpredictable fashion. In our setting, set-based information sharing could be based on quantities ("We will order between 5 and 10 tools this year") or requested delivery times ("We need this soft order between June and December"). Addressing some of the concerns related to trust and reputation raised by the present study, the buyer initiated a fundamental redesign of the forecast-sharing mechanism, which included providing information to the suppliers about forecasted orders in the form of intervals.

While new information technologies have enabled firms involved in a supply chain to gain insight into the planning processes of other firms, our findings demonstrate that there remain substantial organizational barriers preventing firms from fully achieving the benefits of forecast sharing and collaborative planning.

References

Axelrod, R. 1981. The emergence of cooperation among egoists. *Amer. Political Sci. Rev.* **75**(2) 306–318.

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

Cachon, G. P., M. A. Lariviere. 1999. Capacity choice and allocation: Strategic behavior and supply chain performance. *Management Sci.* **45**(8) 1091–1108.

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., S. Netessine. 2003. Game theory in supply chain analysis. D. Simchi-Levi, S. D. Wu, Z.-J. (Max) Shen, eds. *Supply Chain Analysis in the eBusiness Era*. Kluwer, Boston, MA, Forthcoming.

Cakanyildirim, M., R. O. Roundy. 2002. Sedfam: Semiconductor demand forecast accuracy model. *IIE Trans.* **34**(5) 449–465.

Cattani, K., W. Hausman. 2000. Why are forecast updates often disappointing? *Manufacturing Service Oper. Management* **1**(2) 119–127.

Celikbas, M., J. G. Shanthikumar, J. M. Swaminathan. 1999. Coordinating production quantities and demand forecasts through penalty schemes. *IIE Trans.* **31**(9) 851–864.

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

Cox, D. R. 1972. Regression models and life-tables. *J. Roy. Statist. Soc. Ser. B—Statist. Methodology* **34**(2) 187–200.

Debo, L. 1999. Repeatedly selling to an impatient newsvendor when demand fluctuates. Working paper, GSIA, Carnegie Mellon University.

Graves, S., D. B. Kletter, W. B. Hetzel. 1998. A dynamic model for requirements planning with application to supply chain optimization. *Oper. Res.* **46**(3) S35–S49.

Heath, D. C., P. L. Jackson. 1994. Modeling the evolution of demand forecasts with application to safety stock analysis in production distribution systems. *IIE Trans.* **26**(3) 17–30.

Helsen, K., D. C. Schmittlein. 1993. Analyzing duration times in marketing—Evidence for the effectiveness of hazard rate models. *Marketing Sci.* **12**(4) 395–414.

Kalbfleisch, J. D., R. L. Prentice. 1980. *The Statistical Analysis of Failure Time Data*. Wiley, New York.

Kaminsky, P., J. M. Swaminathan. 2001. Utilizing forecast band refinement for capacitated production planning. *Manufacturing Service Oper. Management* **3**(1) 68–81.

Kreps, D. M., P. Milgrom, J. Roberts, R. Wilson. 1982. Rational cooperation in the finitely repeated prisoners-dilemma. *J. Econom. Theory* **27**(2) 245–252.

Krishnan, V., S. D. Eppinger, D. E. Whitney. 1997. A model-based framework to overlap product development activities. *Management Sci.* **43**(4) 437–451.

Law, A., W. D. Kelton. 1991. *Simulation Modeling and Analysis.* McGraw Hill Science, New York.

Lee, H. L. 1996. Effective inventory and service management through product and process redesign. *Oper. Res.* **44**(1) 151–159.

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

Loch, C. H., C. Terwiesch. 1998. Communication and uncertainty in concurrent engineering. *Management Sci.* **44**(8) 1032–1048.

May, S., D. W. Hosmer. 1998. A simplified method of calculating an overall goodness-of-fit test for the cox proportional hazards model. *Lifetime Data Anal.* **4**(2) 109–120.

Ren, Z. J., M. A. Cohen, T. H. Ho, C. Terwiesch. 2004. Sharing forecast information in a long-term supply chain relationship. Working paper, Boston University, Boston, MA.

Roemer, T. A., R. Ahmadi, R. H. Wang. 2000. Time-cost trade-offs in overlapped product development. *Oper. Res.* **48**(6) 858–865.

Sommer, S. C., C. H. Loch. 2003. Incomplete incentive contracts under ambiguity and complexity. Working paper, INSEAD, Fontainebleau, France.

Taylor, T. A., E. L. Plambeck. 2003. Supply chain relationships and contracts: The impact of repeated interaction on capacity investment and procurement. Working paper, Stanford University, Palo Alto, CA.

Terwiesch, C., C. H. Loch, A. De Meyer. 2002. Exchanging preliminary information in concurrent engineering: Alternative coordination strategies. *Organ. Sci.* **13**(4) 402–419.

Terwiesch, C., Y. Xu. 2004. The copy-exactly ramp-up strategy: Trading-off learning with process change. *IEEE Trans. Engrg. Management* **51**(1) 70–84.

Toktay, L. B., L. M. Wein. 2001. Analysis of a forecasting-production-inventory system with stationary demand. *Management Sci.* **47**(9) 1268–1281.

Wedel, M., W. Kamakura. 1998. *Market Segmentation: Conceptual and Methodological Foundations*. Klumer Academic Publishers, Boston, MA.