Measuring Imputed Costs in the Semiconductor Equipment Supply Chain

Morris A. Cohen, Teck H. Ho, Justin Z. Ren, and Christian Terwiesch

The Wharton School, University of Pennsylvania

December, 2001

Abstract

We consider the order fulfillment process of a customized capital equipment supplier. Prior to receiving a firm purchase order from the customer, the supplier receives a series of shared forecasts, which are called 'soft orders'. Facing a stochastic internal manufacturing lead-time, the supplier must decide at what time to begin the fulfillment of the order. This decision requires a trade-off between starting too early, leading to potential holding or cancellation cost, and the cost of starting too late, leading to potential loss of goodwill. We collect detailed data of shared forecasts, actual purchase orders, production lead-times, and delivery dates for a supplier-buyer dyad in the semiconductor equipment supply chain. Under the assumption that the supplier acts rationally, optimally balancing the cancellation, holding, and delay costs, we are able to estimate the corresponding cost parameters based on the observed data. Our estimation results suggest that the cost of cancellation is four times higher and the holding cost is two times higher than the delay cost. In other words, the supplier is very conservative when commencing the order fulfillment, which negates the effectiveness of the overall forecast sharing mechanism.

1 Introduction[1]

Many firms selling expensive customized capital goods such as production equipment, commercial aircraft, medical devices, or defense systems, face an order-fulfillment dilemma. On the one hand, their customers expect them to exhibit a high degree of responsiveness, requiring product delivery within an aggressive customer lead-time. On the other hand, the customized nature of these products makes it risky for the supplier to keep sub-systems or even finished goods in inventory, leading to lengthy and stochastic manufacturing lead-times. To resolve this dilemma, suppliers routinely begin procurement or even production based on forecasted orders, as opposed to waiting for firm purchase orders from their customers. Such forecasted orders, sometimes also referred to as soft orders, are based on the customer's purchase intent as typically revealed to the supplier's marketing and sales function.

There exists, however, a substantial difference between a forecasted (soft) order and a firm purchase order. Responding to changes in their technology and market environments, customers may decide to revise the shared forecasted orders, leading to either changes in the requested delivery dates or cancellations of the orders. This raises an interesting question as to how the supplier should respond to the shared forecasts in such a volatile environment. Specifically, the supplier must decide on how to deal with the preliminary nature of a soft order, and operationally, define a point in time (a level of information quality) at which to start the corresponding fulfillment processes. The supplier can start early, facing the risk of the order not materializing (cancellation cost) or the equipment being ready too early (holding cost), or can delay the start until more information becomes available, thereby facing the risk of being late (delay cost). Put differently, the supplier can "Wait and be Late or Rush and be Wrong", neither of which seems all too desirable.

In this article, we report on an empirical study of forecast sharing related to the acquisition of customized production equipment for the manufacturing of semiconductors.

[1]The authors would like to thank the management teamsof the supply-chain dyad who generously provided detailed data, which form the empirical foundation of this study.

Semiconductor equipment perfectly matches the industry characteristics outlined above. Customers, who are responding to the turbulent environment they face in the demands for their end-products, press for short customer lead-times, requiring product delivery within three months or less. At the same time, the complexity and degree of customization of the equipment causes manufacturing lead-times to be long and stochastic, ranging between several months to a whole year.

In this environment, we find that 20% of all soft orders are cancelled, 71% result in changes in the desired delivery dates, and 5% lead to changes in equipment specification. Taking the perspective of a supplier of customized semiconductor production equipment, we develop a formal model addressing the trade-off between the early start of the order fulfillment process (leading to potential cancellation cost and holding cost) and a delay until more information has become available (leading to a potential delay cost due to loss of goodwill).

How this trade-off is resolved depends on the cost structure of the supplier. While traditionally, the supply chain literature has taken these cost parameters as exogenously given and then searched for the optimal operational decision, we take a complementary approach. Based on an empirical observation of the supply chain over time, including detailed data of shared forecasts, actual purchase orders, production lead-times, and delivery dates, and on the assumption that the supplier makes her decisions rationally, we are able to reconstruct the cost parameters that explain the empirical supply chain behavior. Our estimation results suggest that the cost of cancellation is four times higher and the holding costs is two times higher than the delay cost. In other words, the supplier is very conservative when commencing the order fulfillment, which negates the effectiveness of the overall forecast sharing mechanism.

Specifically, this paper makes the following two contributions. First, to the best of our knowledge, this is the first paper to empirically estimate the cost parameters that underlie the existing analytical models of coordination in supply chain management. Our results can be used by buyers who want to quantify to what extent their suppliers respond to the shared forecast information. The 'structural approach' assumes that the manufacturer is rational and hence the estimated parameters are robust in face of policy changes in the presence of rational expectations (Lucas, 1976). Second, our estimated and validated model can be used to rank order ways to improve the effectiveness of an order forecast sharing system in the semiconductor equipment supply chain. Our results suggest that fearing order cancellation, the supplier currently, by and large, views the preliminary forecast information cautiously. Also, a dramatic improvement in lead time can only be accomplished by a simultaneous change in at least 2 of the cost parameters.

The remainder of this article is organized as follows. We begin with a description of the semiconductor equipment industry, including a description of the buyer-supplier relationship (Section 2). After reviewing the relevant literature (Section 3), we present an economic optimization problem that reflects the suppliers decision problem (Section 4). Section 5 describes our empirical study and our econometric framework, followed by our estimation results (Section 6). We further test model fit (Section 7), and conduct a policy scenario analysis (Section 8). Section 9 concludes the paper.

2 Semiconductor Equipment Supply Chain

We conducted a joint research project with a major semiconductor manufacturer and a major equipment supplier. Besides documenting the overall equipment acquisition process as outlined below via in-depth face-to-face interviews, we collected detailed data on the duration of each of the three stages for a total of 100 tool orders.

Demand for semiconductor production equipment is triggered by product demand for the chip supplier's end products, such as micro-processors or memory chips. As illustrated by the recent down-turn in these markets, final demand from PC suppliers exhibits a high degree of uncertainty, leading to adjustments in sales forecasts of large magnitudes within a relatively short amount of time. Market forecasts are done monthly or quarterly at the large semiconductor suppliers for a time period of 2-5 years into the future and are regularly updated following a rolling horizon principle.

These product level demand forecasts are allocated to existing or potentially new fabs, which compare their available tool capacity with the capacity requirements from the forecasts. If the allocated demand is not supported by available capacity, new equipment is ordered. Such mismatches are typically a result of higher peak-load capacity requirements and changes in the technical process specifications of new chip generations.

While the translation of product demand into equipment orders seems relatively straightforward, two factors make this computation extremely complex. First, the semiconductor industry is very capital intense and the capital expenditures for new production equipment are the single largest item on a company's earning statement. For example, the industry leader Intel Inc. spent $5 billion for equipment acquisitions in the year 2000 alone. Given this magnitude of capital expenditure, even minor under-utilization of equipment can have a dramatic financial impact.

Second, while most semiconductor equipment in the world is operated 24 hours a day for 7 days a week, the actual availability of a specific piece of equipment can be substantially lower. Equipment becomes unavailable as a result of machine break-downs, required qualification procedures, engineering trials, preventive maintenance, etc. Moreover, semiconductor manufacturing is a very yield driven process, associated with substantial scrap and the need for rework. While all these factors reduce the average availability, they also introduce substantial variability into the capacity planning, aggravating problems resulting from the demand uncertainty.

2.1 Equipment Acquisition Process

Once this capacity planning process has generated requests for additional pieces of equipment, an elaborate tool acquisition process commences. This process includes three stages, forecast sharing, manufacturing, and installation.

During the forecast sharing stage, the semiconductor manufacturer (Buyer) creates a forecasted order (soft order), which is shared with the equipment manufacturer(supplier) via an on-line collaboration system. This soft order includes the tool's specifications, and the requested delivery date (RDD). This soft order however is merely preliminary information - opposed to a final purchase order - since in the presence of market and capacity uncertainties, the buyer does not want to commit to an order at such an early stage. The buyer can, after getting more information about her market demand and production yields, decide to (a) cancel the order, (b) to move it to another date, (c) or leave the soft order unchanged.

The supplier becomes aware of the buyer's purchase intent, both through on-line information system as well as through direct customer interaction from the sales and marketing department. At some point she needs to initiate the production of the tool, which includes procurement of sub-systems from second tier suppliers and the entry of the order into the production schedule.

The supplier is in a difficult situation, as starting the order too early can lead to holding and cancellation costs while starting the order too late can lead to late-shipment costs. The typical manufacturing lead-time of the supplier ranges between 3 and 5 months. The lead-time exhibits strong variability as a result of differences in product-mix going through the suppliers facility, changes in equipment demand, process generation, and/or uncertainty in lead-times from the second-tier suppliers.

Finally, the tool is shipped to the corresponding fab, where it is installed and then has to move through an elaborate qualification process before it can produce commercial output. The overall equipment acquisition cycle is illustrated by Figure 1. In total, the equipment acquisition cycle is approximately one year. Some tools, especially in the lithography domain, can take even longer.

2.2 Soft versus Firm Orders

Among other data, we collected the forecasted tool orders as shared by the buyer with the supplier via an on-line collaboration system over an extended period of time. The buyer provided a quarterly forecast of how many tools she planned to acquire in each of the coming seven quarters. These forecasts were updated following a rolling horizon principle. For example, in quarter 4 of 1998 the buyer forecasted her demand for quarters Q1 1999 to Q3 2000. In quarter 1 of 1999 a certain number of tools were purchased and a new forecasts for quarters Q2 1999 to Q4 2000 were placed.

Figure 2 depicts the forecasts as provided by the buyer in the quarters Q1 1999 to Q1 2001. Each of these shared forecasts is a time series consisting of the seven quarters included in the relevant forecast window. Figure 2 also contrasts the forecasts with the actual tool purchases. We can make two interesting observations. First, the forecasts vary widely, both over time (what is forecasted in e.g. Q1 2000 for the time period of Q2 2000 to Q4 2001) as well as from one forecast to the next (e.g. what is forecasted in e.g. Q2 1999 for Q4 1998 vs what is forecasted in Q3 1999 for Q4 1999). The former is a consequence of the lifecycle of process generations and the associated need for capacity expansion. The latter is primarily a result of the large amount of uncertainty in the industry, especially with respect to the height of the peak demand in the product lifecycle.

Second, we see that - on average - the buyer forecasted for more tool purchases than she ended up ultimately committing to. In other words, there are significantly more soft orders than hard orders, leading to numerous order cancellations. This reflects the cost structure of the buyer: forecasting too little can lead to equipment shortages and potential production losses of entire fabs with a substantial negative financial impact. Forecasting too much however, does not cost the buyer anything, as the risk of producing the equipment without having the demand for it is entirely transferred to the supplier.

2.3 Research Goals

While a systemic inflation of forecasts does not necessarily lead to out-of-pocket costs to the buyer, it can have negative implications on the supplier's perception of the buyer's credibility. This in turn can hurt overall supply chain performance, and thereby - albeit indirectly - the buyer.

The importance of credibility in forecast sharing between buyer and supplier can be nicely illustrated by the old parable of "the boy and the wolf". In absence of credibility and trust between the town people and the shepherd boy, the boy will call for help ('cry

WOLF!') in response to the slightest fear. However, the town people - after running out to help the boy just to find out that there was no wolf - will start to discount the cry for help and ultimately they will decide not to provide any help at all, which has fatal consequences for both the boy and the sheep.

The situation where the buyer provides an overly optimistic (high) forecast to the supplier corresponds directly to the boy's cry for help. While the cry itself does not carry any cost to the boy, it requires a substantial effort on the part of the town people. This effort is wasted in case of a false alarm, just like providing a high forecast - while being inexpensive for the buyer - is costly for the supplier because of reserving capacity and second tier procurement decisions. Moreover, just as the town people ultimately decided to not react to the cry, the supplier is likely to discount the shared forecasts and will assign greater importance to cancellation costs and holding costs compared to late shipment cost.

The objective of this article is to measure how the semiconductor equipment supplier we studied perceived the cost of cancellation and holding relative to the cost of late shipment. While previous research on coordination and contracting in supply chains has emphasized the importance of forecast sharing and the risks associated with losing credibility, we provide - to our knowledge - the first empirical study to demonstrate these effects econometrically based on actual supply chain behavior opposed to anecdotal evidence. Our results are of direct managerial relevance in the semiconductor equipment supply chain, as they demonstrate that the current method of forecast sharing as described above has lead to a 'boy and the wolf' effect and presently is by and large ineffective.

3 Related Literature

Lee *et al.* (1997) provide one of the earliest academic discussion of problems related to soft orders and their cancellation. They refer to such orders as "phantom orders", defined as high forecasts of future demand that never materializes, and see them as a key contributor to the bullwhip effect in supply chains.

Problems related to phantom orders and overly optimistic forecasts have frequently made their way into the business news. For example, Zarley and Damore (Computer Reseller News; May 6, 1996) discuss how PC manufacturers suspected that their customers (distributors) placed phantom orders. As a result, these manufacturers frequently produce only a fraction of the quantity specified in the demand forecast. A comparable situation occurred in the cellular phone industry in the 1994 Christmas season. Motorola experienced significant over-ordering by customers concerned with a potential capacity short-fall (Business Week, March 6, 1995).

Similar problems were experienced at Boeing, which had difficulties in increasing its production of 747s due to parts shortages. Boeing's large supplier base apparently did not trust the company's optimistic demand forecast (indicating a strong growth in 1997) and therefore could not fulfill Boeing's large orders. Only one year later though, following the Asian financial crisis, the supplier's conservatism proved to be a wise decision (Cole, Wall Street Journal, June 26 and September 16, 1997; Biddle, Wall Street Journal, June 10, 1998).

Boeing not only experienced the 'boy-and-the-wolf' problem of soft orders with their suppliers, but - at the same time - also with their customers. For example, in May 2001, Northwest Airlines cancelled a soft order of 23 Boeing 737 jets, which led Boeing to fall behind Airbus for the first time (Flight International, June 2001).

In addition to this anecdotal evidence on problems related to forecast sharing, the recent academic literature of supply chain management includes several articles providing game theoretical models of this behavior. Closest to our study, Cachon and Lariviere (2001) distinguish between forecast sharing contracts with forced compliance and voluntary compliance. Under forced compliance, the supplier's behavior can be observed and compared to the contracted behavior with a level of detail sufficient for any deviations to be brought up in court. In this case, all market power rests with the buyer and she will therefore design a contract that maximizes her rent derived from the relationship.

Under voluntary compliance, the supplier's behavior is not fully verifiable. This might be a result of some stochastic element in the environment or within the supplier's operations. In such a situation, courts cannot enforce contract compliance, as the supplier can always argue that she attempted to comply but failed because of some random event.

The situation analyzed in this article is one of voluntary compliance. While there are detailed contracts written between buyer and supplier, they are very hard to enforce. For example, if the supplier is not able to meet the requested delivery date, she could easily find reasons outside her control to explain this. Examples include parts shortages from the 2nd tier supplier, changing specifications from the buyer, and/or other external events[2].

Hariharan and Zipkin (1995) call the time between order arrival and its due dates demand leadtime, and show elegantly that customers sharing such advance demand information is equivalent to a reduction in conventional supply lead time, and thus improves supply chain performance. Gallego and Ozer (1999) have a discrete-time model of advance ordering. Their models all assume that all orders are firm order, i.e., they cannot be cancelled once placed. If such orders can be cancelled, which happens in our study of semiconductor industry, then they are 'soft order', or essentially forecasts, and the effect of sharing such info is less obvious.

A second stream of research that is relevant to our work relates to the sharing of preliminary information, which is a common practice in product development teams, especially in those proceeding concurrently (Terwiesch et al. 2001). Similar to the supplier in our study, who initiates the order fulfillment process prior to receiving a firm purchase order from the customer, development teams frequently begin their work on a new product prior to receiving detailed design specifications from the customer and/or from the market research department.

In this line of research, Loch and Terwiesch (1998) model the situation faced by a concurrent engineering team where an information receiving development activity has to decide how much to rely on the preliminary information provided by the information sender. While on the one hand, the information receiver always wants to start early, in

[2]Other related work includes Chen (2000), Chen (2001), Chu (1992), Corbett (2001), Corbett and Tang (1999), Desai and Srinivasan (1995), Ha (1999), Lariviere and Padmanabhan (1997), and Lariviere and Porteus (2000), Porteus and Whang (1999).

attempt to gain from parallel task execution, this leads to a lower quality of information and thus a higher likelihood of costly rework. When receiving preliminary information, one thereby faces the decision “Rush and be Wrong or Wait and be Late” (Terwiesch et al. 2001). The resulting trade-off is similar to the supplier’s problem we define below as well as to the overall set-up defined in Figure 3.

4 Model Formulation

We use the so-called “structural approach” to estimate unobservable costs of cancellation, holding and delay. We had observed the supply chain dyad, including shared forecasts and final shipments, over an extended period of time. In addition, we had also obtained data about the supplier’s internal manufacturing lead-times. Under the assumption that the supplier would behave rationally in balancing the three cost elements (of holding, cancellation and delay), we wanted to impute values for the costs that explained observed supply chain behavior.

The unit of analysis in both our model and the empirical analysis is an order for a single piece of production equipment. We will first define the supplier’s problem to choose the optimal time to commence work on a given order.

4.1 The Supplier’s Problem

Consider a time line (see Figure 3) starting at the point in time when the first soft order is received by the supplier ($t = 0$). Associated with this first, preliminary order is a Requested Delivery Date (RDD), which is potentially refined by the buyer over time. At some point in time the uncertainty inherent in the soft order is resolved. Define this point in time as $t = T_N$. At T_N, the tool delivery is requested with a firm delivery date for $t = RDD_N$, or the order is cancelled.

The supplier faces the following problem when deciding about the time T_p at which she begins the fulfillment process on a - potentially soft - order. Specifically, she faces two types of uncertainty, market uncertainty and uncertainty in manufacturing lead-time. Market uncertainty includes the probability that the order is cancelled, which we will label as p, as well as any potential changes in the requested delivery date RDD_N.

From the operations side, there is uncertainty in the manufacturing lead-time. This may result from traditional lead-time variability in a job-shop like production environment, changes in product mix or production volume, and/or from variability in delivery lead-times for sub-systems that are ordered from second tier suppliers.

The supplier must trade-off the cost of beginning too early (cancellation, potential delay) with the risk of producing too late (delay cost). The problem resembles a traditional Newsvendor problem, which in this case occurs in the time domain opposed to the quantity domain. Define cancellation cost, c, as the cost incurred by the supplier per unit of time that an order spends in production and later is cancelled. Define holding cost, h, as the cost incurred by the supplier per unit of time that the tool is produced prior to the date the customer actually needs it (RDD_N). Finally, define delay cost, g, as the cost incurred by the supplier per unit of time that the actual delivery date exceeds the RDD_N, i.e. for the shipment being done late. This cost set-up is illustrated in Figure 3.

With these cost parameters, we can state the following cost minimization problem:

$$\begin{aligned} Min_{T_p} E(Total\ Cost) &= pc\,(T_N - T_P)^+ \\ &\quad +(1-p)\left\{h\left[(RDD_N - T_P) - LT\right]^+ + g\left[LT - (RDD_N - T_P)\right]^+\right\} \end{aligned} \tag{1}$$

where $(.)^+$ denotes $Max(.,0)$. In the equation above, the first term denotes the expected cancellation cost while the second and the third term capture expected holding cost and expected delay cost respectively.

Define $S = RDD_N - LT$, and let the cumulative distribution for the new random variable be $F(S)$. Moreover, let the distribution for T_N be $G(.)$. Equation (1) can be rewritten as:

$$\begin{aligned} Min_{T_p} E(Total\ Cost) &= pc\int_{T_p}^{\infty}(T_N - T_P)\,dG(T_N) \\ &\quad +(1-p)\left[h\int_{T_p}^{\infty}(S - T_P)\,dF(S) + g\int_{-\infty}^{T_p}(T_P - S)dF(S)\right] \end{aligned} \tag{2}$$

It is easy to show that (2) is convex in the decision variable T_p. Thus, there exists a unique cost minimizing starting point T_p^* at which the supplier should begin production of an order, which is characterized by the solution to the first order condition:

$$pcG(T_p^*) + (1-p)(g+h)F(T_p^*) = pc + (1-p)h \tag{3}$$

Define the expected delay time at the optimal decision as

$$V(T_p^*) = E[LT - (RDD_N - T_p^*)]^+ = \int_{-\infty}^{T_p^*} (T_p^* - S)dF(S). \tag{4}$$

4.2 Assumptions and Functional Forms

In order to obtain a solution for the optimal starting point of order fulfillment T_p^*, we will need to empirically reconstruct the cost parameters c, h, and g. We also need to make specific assumptions concerning the underlying distribution function for the arrival time of finalized information, T_N, and for the distribution function underlying S.

For the case of the arrival time of finalized ordering information, T_N, we assume an exponential distribution. Specifically, we assume $G(x) = 1 - e^{-\alpha x}$. Figure 4 compares the actual data that we collected in the semiconductor equipment supply chain with an exponential distribution where $\alpha = 0.21$. The "optimal fit" in Figure 4 appears a reasonable approximation for the oberved data.

Next, we need to choose a distribution function for $S = RDD_N - LT$. Testing separately, we find both RDD_N and LT are asymptotically distributed normal. So a natural assumption would be to assume $S = RDD_N - LT$ is distributed normal. However, such a normality assumption does not lead to closed-form solutions and therefore, in estimation problems similar to ours, the normal distribution is commonly replaced with a Weibull(2, β) distribution. This distribution is also called the Rayleigh distribution.

In our context it is possible for $RDD_N - LT$ to take on negative values, and so we shift the Rayleigh distribution to the right by a constant δ. The exact value of this shift δ will be estimated jointly with the three cost parameters defined above. While the resulting distribution of S no longer is a Weibull distribution, it still has a relatively simple cumulative distribution function, with $F(S) = 1 - e^{-\beta^2(S+\delta)^2}$.

Substituting the functional forms into equation (3) we obtain a simplified first order condition:

$$pce^{-\alpha T_p} + (1-p)(g+h)e^{-\beta^2(T_p+\delta)^2} = (1-p)g \tag{5}$$

In addition to assumptions related to the functional forms of $F(S)$ and $G(x)$, our model is based on three specific assumptions. First, we assume that the starting point of the buyer's order fulfillment is solely determined by the holding, delay, and cancellation costs. This is the standard assumption underlying most models and requires that the supplier has sufficient capacity to begin production whenever she deems it optimal. Our in-depth observation and interaction with the supplier reveal that this is a reasonable approximation because the particular buyer studied is an important customer. Second, we assume the cancellation cost to be proportional to the elapsed time between the initiation of the order fulfillment process and occurrence of the cancellation event. Third, we consider only two types of changes to the order: order cancellation and changes in RDD. We did not consider changes in product specification because they account for less than 5% of the order changes. The tool functions we analyzed have been used in prior generations of technology by the buyer.

In contrast to the assumptions related to the functional forms of $F(S)$ and $G(x)$,which can be validated empirically by means of a goodness of fit test, these additional assumptions will be validated by contrasting observed supply chain behavior with the behavior predicted by our model. This will be the subject of Section 7.

4.3 Optimality Condition and Taylor Approximation

Even under the assumed specific distribution functions for $F(.)$ and $G(.)$, there is no guarantee of the existence of closed-form solutions for T_p^*. We linearize the objective function in the neighborhood of an estimated starting date in order to generate a closed form solution. It can be shown numerically that Taylor series is a good approximation and the solution obtained from the approximation is close to the true optimum value.

Define τ as the time around which we expand the exponential function in the above equation. That is, $e^{-\alpha T_p} = e^{-\alpha\tau} - \alpha e^{-\alpha\tau}(T_p - \tau) + r_G$, and $e^{-\beta^2(T_p+\delta)^2} = e^{-\beta^2(\tau+\delta)^2} - 2\beta^2(\tau+\delta)e^{-\beta^2(\tau+\delta)^2}(T_p - \tau) + r_F$, where r_G and r_F are the residual terms for those two expressions. The resulting first order equation is

$$\begin{aligned} &pc\left[e^{-\alpha\tau} - \alpha e^{-\alpha\tau}(T_p - \tau) + r_G\right] \\ +&(1-p)(g+h)\left[e^{-\beta^2(\tau+\delta)^2} - 2\beta^2(\tau+\delta)e^{-\beta^2(\tau+\delta)^2}(T_p - \tau) + r_F\right] = (1-p)g \end{aligned} \tag{6}$$

For ease of presentation, define $\theta_1 = pc, \theta_2 = (1-p)(g+h)$, and $\theta_3 = (1-p)g$. Note that g, h, and c, our key parameters of interest, can be fully recovered from θ_1, θ_2, and θ_3: $g = \theta_3/(1-p)$; $h = (\theta_2 - \theta_3)/(1-p)$; $c = \theta_1/p$.

Re-arranging terms, while ignoring residuals for a moment, we obtain the following closed-form solution for T_p^*:

$$T_p^* = \frac{\theta_1 e^{-\alpha\tau} + \theta_1\alpha e^{-\alpha\tau}\tau + \theta_2 e^{-\beta^2(\tau+\delta)^2} + 2\theta_2\tau\beta^2(\tau+\delta)e^{-\beta^2(\tau+\delta)^2} - \theta_3}{\theta_1\alpha e^{-\alpha\tau} + 2\theta_2\beta^2(\tau+\delta)e^{-\beta^2(\tau+\delta)^2}} \tag{7}$$

Our approximation approach has a direct analogy in the order fulfillment process of the supplier under study. The supplier has a predefined milestone at which she initiates the manufacturing process. This milestone is then adjusted to reflect order specific considerations, technical tool characteristics, shipment destination, or tool function. This is a natural analogy to Taylor series expansion, where an objective function can be approximated in the neighborhood of a given point, and actions can be adjusted in relation to that given point to improve the objective.

5 Structural Estimation Approach

We observed the semiconductor equipment supply chain at the interface of the buyer and the supplier for a total of 18 months, beginning in January 1999. For this time period, we created a complete history of forecast sharing via direct access to an on-line collaboration system. The buyer of the supply chain updates the on-line collaboration system on a monthly basis, providing the latest forecasts in the form of soft orders.

The system enables us to follow every individual equipment order from its initiation as a soft order to the completion of the installation and qualification procedure. For every order, we collected information on its total order lead time (from the time an order enters the system to the time it is fulfilled), the process technology generation that the equipment was designed to support, the function of the equipment, and the destination (geographic location of the fab) to which the equipment is to be delivered. As we collected the data over an extended period of time, we were also able to observe information about order cancellations and requested changes to the delivery date. This first database provided an accurate description of forecast sharing between buyer and supplier. The total number of orders in the sample was 100.

We complemented this first database which emphasizes forecast sharing with a second database. This second database included information about the detailed operations of the supplier, which were unobservable to the buyer. Specifically, it included data on historical manufacturing lead-times for equipment that were supplied by the manufacturer to the buyer. Our prior interviews had identified that lead times varied with the function of the equipment, its destination, and the process technology generation it was designed to support. Accordingly, the corresponding data were collected in addition to the manufacturing lead-time itself.

5.1 Estimation Procedure

Assume that α and β, the parameters for $G(.)$ and $F(.)$ respectively, vary across orders and this variation can be explained by a set of independent variables. Let the observations corresponding to the individual pieces of equipment be indexed as $i = 1, 2, \ldots, I$, and explanatory variables be indexed as $j = 1, 2, \ldots, J$. The explanatory variables are derived from the first (forecast sharing) database.

Let any given $\alpha_i > 0$ consist of a "base rate" γ_0 , and an order-specific term which can be explained by a set of values of $x_{ij}, j = 1, \ldots, J$. In other words, α_i can be written as:

$$\alpha_i = \exp(\gamma_0 + \gamma_1 x_{i1} + ... + \gamma_J x_{iJ}) \tag{8}$$

This approach of making a distribution parameter a function of independent variables is common practice in empirical marketing as well as operations research models (See for example Duenyas 1995 for a case where the customer arrival rate is explained through various attributes of the customer). Similarly we define $\beta_i > 0$ as:

$$\beta_i = \exp(\rho_0 + \rho_1 x_{i1} + ... + \rho_J x_{iJ}) \tag{9}$$

Our notation can be further simplified by labeling the explanatory variables in vector form as $X_i = (1, x_{i1}, ..., x_{iK})'$, a $J \times 1$ vector for α_i and a $J \times 1$ vector for β_i . That is

$$\begin{aligned} \alpha_i &= \exp(\gamma X_i) \\ \beta_i &= \exp(\rho X_i) \end{aligned} \tag{10}$$

Note, that γ and ρ are the respective parameters vectors of dimension $1 \times J$ to be estimated.

Finally let ε_i be the error term for T_p, which follows the standard assumption that errors are identically and independently distributed $N(0, \sigma)$.

Thus the equation to be estimated can be written as:

$$\begin{aligned} T_{p,i} &= T_{p,i}(\theta_1, \theta_2, \theta_3, \alpha_i = \exp(\gamma X_i), \beta_i = \exp(\rho X_i), \delta) + \varepsilon_i \\ &= \frac{\theta_1 e^{-\exp(\gamma X_i)\tau} + \theta_1 \exp(\gamma X_i) e^{-\exp(\gamma X_i)\tau}\tau + \theta_2 e^{-\exp^2(\rho X_i)(\tau+\delta)^2}}{\theta_1 \exp(\gamma X_i) e^{-\exp(\gamma X_i)\tau} + 2\theta_2 \exp^2(\rho X_i)(\tau+\delta) e^{-\exp^2(\rho X_i)(\tau+\delta)^2}} \\ &\quad + \frac{2\theta_2 \tau \exp^2(\rho X_i)(\tau+\delta) e^{-\exp^2(\rho X_i)(\tau+\delta)^2} - \theta_3}{\theta_1 \exp(\gamma X_i) e^{-\exp(\gamma X_i)\tau} + 2\theta_2 \exp^2(\rho X_i)(\tau+\delta) e^{-\exp^2(\rho X_i)(\tau+\delta)^2}} + \varepsilon_i \end{aligned} \tag{11}$$

One difficulty that arises in estimating the parameters is that data on $T_{p,i}$ are not available. While we know the completion date for each order, we cannot observe the specific manufacturing lead-time for this piece of equipment, which could have helped in deriving its order starting time. However, utilizing our second database on Supplier's manufacturing lead-times, we are able to estimate the manufacturing lead-time for the individual orders. This allows us to employ a two-step estimation procedure for the estimation of the relative cost parameters θ_1, θ_2, and θ_3, as well as for other parameters.

The idea behind this two-step approach is to utilize the observed finish time of the piece of equipment in manufacturing, that we denoted as FT, and recognize that:

$$FT = T_p + LT \tag{12}$$

So instead of estimating based on T_p, we transform the dependent variable to be FT by adding LT to both sides of equation (11).

Our first-step regression is on LT using the dataset on manufacturing lead-times, which can be written as the following linear regression:

$$LT_i = \eta Y_i + e_i \tag{13}$$

where Y_i indicates the set of variables that influence manufacturing lead-times, η is the parameter vector to be estimated, and e is the residual vector. After obtaining the estimates for η, we can obtain estimates of LT for the tools whose data will be used in estimate equation (12), which we denote as $\widehat{LT}$. That is $\widehat{LT}_i = \eta Y_i$.

In our second step, we can substitute LT with $\widehat{LT}$ in equation (12), and proceed to estimate the following nonlinear system:

$$\begin{aligned} FT_i &= T_{p,i} + \widehat{LT} + \varepsilon_i \\ &= T_{p,i}(\theta_1, \theta_2, \theta_3, \alpha_i = \exp(\gamma X_i), \beta_i = \exp(\rho X_i), \delta) + \eta Y_i + \varepsilon_i \end{aligned}$$

Because the equation is non-linear in the parameters to be estimated, the Maximum Likelihood Estimation (MLE) method is applied. We invoke the assumption that the finishing times of orders are statistically independent. Thus, the joint likelihood function L takes the form:

$$L = \left(\frac{1}{\sqrt{2\pi\sigma^2}}\right)^N \exp\left\{-(1/2\sigma^2)\sum_{i=1}^{N}[FT_i - T_{p,i}(\theta_1, \theta_2, \theta_3, \alpha_i = \exp(\gamma X_i), \beta_i = \exp(\rho X_i), \delta) - \eta Y_i]^2\right\}$$

From this equation we establish estimates for $\widehat{\theta}_1$, $\widehat{\theta}_2$, $\widehat{\theta}_3$, $\widehat{\gamma}$, $\widehat{\rho}$, $\widehat{\delta}$, and $\widehat{\sigma}$ such that the likelihood of the observed sample is maximized. Murphy and Topel (1985) show that the estimated parameters using the above 2-step estimation procedure are consistent and asymptotically normal. Since we have a reasonable sample size of 100, we assume that the parameter estimates are consistent and normal.

6 Estimation Results

We first perform a linear Least Squares regression using model specified by equation (13) to obtain η, coefficient estimates for LT. The variables associated are:

- Lead time: Dependent variable, in months;
- Tool process generation: At the time of study, one new generation tool was in its peak manufacturing volume. We code it as a binary variable named NEW;
- Tool type: Tools of function CMP (Chemical Mechanical Planarization) are considered premium tools, and are thus identified as one of the key drivers of manufacturing lead time. We code it as a binary variable PREMIUM, which takes on value of 1 if it is of such function, and zero otherwise representing other standard tools;
- Tool destination: It is also speculated that tools ordered from and delivered to development fabs have a shorter lead time. So we include this binary variable as well, and call it DEVELOP.

With a constant term included, the total number of explanatory variables is 4, i.e., $J = 4$. Model specification as well as estimation results are reported in Table 1.

After we calculate fitted value of LT, using the coefficient estimates obtained above, we perform maximum likelihood estimation to estimate θ_1, θ_2, and $\theta_3, \gamma, \rho, \delta$, and σ. using the dataset constructed from the observed forecast sharing. The explanatory variables are similar to the ones introduced in the regression on LT, namely, process generation

(NEW), tool function (PREMIUM), and tool destination (DEVELOP), plus a new variable, CHRDD, on forecast changes, which measures the number of RDD changes that an order has had since its initial forecast.

Because MLE is a nonlinear method, results can sometimes be sensitive to starting points. In order to make sure the parameter estimates are not locally optimal, we explored the parameter space by trying out various sets of starting values[3]. For each parameter, 12000 sets of starting values are tried and the estimation result with the highest likelihood is reported. Specifically, we randomly generate 12000 sets of starting values for each of the parameter from a normal distribution with mean 0 and standard deviation 3, which covers the usual range of the parameters from our observation. θ_3 is fixed at 0.700 (or equivalently $g = 1.000$) for identification. We also restrict g, h and c to be non-negative, by restricting the values of θ_1, θ_2, and θ_3 . Table 2 reports the parameter estimates that yielded the maximum likelihood.

Based on the estimated values for θ_1, θ_2, and θ_3, we can find our cost parameters for delay cost, g, holding cost, h, and cancellation cost, c. Specifically, we obtain the following values: $g = 1.00, h = 4.28, c = 2.07$. These results suggest that the manufacturer weights the holding and cancellation costs about four and two times higher than cost of loss of goodwill.

Next, in order to test the effect of explanatory variables on the model, we perform four cases of nested models:

1. Tool Destination: DEVELOP

   The corresponding hypothesis is $H_{0,1} : \gamma_2 = 0$, and $\rho_2 = 0$.

2. Forecast Change: CHRDD

   The corresponding hypothesis is $H_{0,2} : \gamma_3 = 0$, and $\rho_3 = 0$.

3. Tool Process Generation: NEW

[3] We use GAUSS, a matrix programming language, to perform the estimation task. The algorithm for maximum likelihood estimation is BFGS.

The corresponding hypothesis is $H_{0,3} : \gamma_4 = 0$, and $\rho_4 = 0$.

4. Tool type: PREMIUM

   The corresponding hypothesis is $H_{0,4} : \gamma_5 = 0$, and $\rho_5 = 0$.

The results for the above sub models are reported in Table 3. Likelihood ratio tests show that each of these sub models is rejected (see χ^2 test statistic in each table). Thus, all the variables identified are essential for the analysis, and cannot be eliminated.

7 Model Validation

One way to validate our modeling approach and the corresponding estimation results is to compare the predicted order finish time $\widehat{FT}$, with actual order time FT. The predicted order finish time is calculated as:

$$\begin{aligned}
\widehat{FT}_i &= \widehat{T}_{p_i} + \widehat{LT} \qquad (14) \\
&= \frac{\widehat{\theta}_1 e^{-\exp(\widehat{\gamma}X_i)\tau} + \widehat{\theta}_1 \exp(\widehat{\gamma}X_i) e^{-\exp(\widehat{\gamma}X_i)\tau}\tau + \widehat{\theta}_2 e^{-\exp^2(\widehat{\rho}X_i)(\tau+\widehat{\delta})^2}}{\widehat{\theta}_1 \exp(\widehat{\gamma}X_i) e^{-\exp(\widehat{\gamma}X_i)\tau} + 2\widehat{\theta}_2 \exp^2(\widehat{\rho}X_i)^2(\tau+\widehat{\delta}) e^{-\exp^2(\widehat{\rho}X_i)(\tau+\widehat{\delta})^2}} \\
&\quad \frac{+2\widehat{\theta}_2\tau \exp^2(\widehat{\rho}X_i)(\tau+\widehat{\delta}) e^{-\exp^2(\widehat{\rho}X_i)(\tau+\widehat{\delta})^2} - \widehat{\theta}_3}{\widehat{\theta}_1 \exp(\widehat{\gamma}X_i) e^{-\exp(\widehat{\gamma}X_i)\tau} + 2\widehat{\theta}_2 \exp^2(\widehat{\rho}X_i)(\tau+\widehat{\delta}) e^{-\exp^2(\widehat{\rho}X_i)(\tau+\widehat{\delta})^2}} + \eta Y_i
\end{aligned}$$

Figure 5 plots the actual finish times FT against the estimated finish times $\widehat{FT}$. The graph shows an overall good fit. An Ordinary Least Square (OLS) regression of FT against $\widehat{FT}$ yields the following results:

$$FT = \underset{(0.98)}{0.70} + \underset{(0.11^*)}{0.92} * \widehat{FT}$$

*:Significant at 1% level.

R^2=40%; Number of observations=100

This regression line is also shown in Figure 5. The results show a statistically non-significant intercept, but a statistically significant slope close to identity, which provides a formal validation of our estimation approach.

8 Policy Scenario Simulation

The mathematical model outlined in Section 4, combined with the parameter estimations for cancellation cost, holding cost, and late shipment cost, enable us to analyze several scenarios corresponding to modifications of cost parameters.

For example, the buyer we interacted with was interested in the question of what would be the impact of a financial late shipment fee on the timeliness of deliveries. Economic intuition suggests that such penalty would increase the late shipment cost for the supplier, thereby encouraging her to commence production earlier (smaller T_p^*). However, we can go one step further than this. Based on our analytical results and the empirical data, we can recompute the expected shipping delay with any given late shipment cost parameter. This is depicted by Figure 6, which shows the relationship between the cost parameter g and the expected slippage, as defined in Equation (4), averaged over all orders and expressed in month. Currently, $g = 1$ and the corresponding slippage is a little less than one month. Increasing late shipment cost from $g = 1$ to $g = 2$ translates into a .1 month reduction in late shipment.

Next, consider the impact of holding cost h on the expected shipping delay (Figure 7). Again, the status quo corresponds to $h = 4$ and a delay of .95 months. Now consider what happens if holding costs are cut in half. For example, the buyer could accept the equipment prior to the specified requested dock date (and of course, also pay for it earlier). This would lead to a .1 month reduction in shipping delay. While, obviously the financial burden of capital cost now rest with the buyer (who pays for the equipment and then leaves it idle up to the time of actual need), the buyer might still be better off, as shipment delays can put the production of entire fabs at risk.

Similarly, Figure 8a and 8b investigate the relationship between changes in cancellation cost c and cancellation probability p. A reduction in cancellation cost could be achieved, if the buyer would take over some of the cost incurred by the seller in the case of cancellation (e.g. procurement cost). Alternatively, the buyer could develop product specifications with more standardized components in it, which would allow the supplier to reuse entire sub-assemblies for another customer after receiving notice of cancellation.

Interestingly, none of the changes outlined above results in a truly dramatic reduction in late shipment. The reason for this lies in the complex trade-off that the supplier faces when deciding upon the optimal time to start working on an order, T_p^*. This trade-off not only involves two forces, but rests on a subtle balance between three forces. Hence, even a large improvement along one dimension will only lead to a small change in the supplier's decision - and thereby the expected delay - as the other two forces are still unchanged. Consequently, large changes in expected delay can only be achieved by changing at least two of the cost parameters jointly (opposed to changing them one at a time). This is illustrated in Figure 9. We see that reductions in holding cost and in cancellation costs actually complement each other, opposed to acting as substitutes. A 50% reduction in both of them would reduce the expected delay by as much as 75%, while each of the two changes implemented individually would only achieve a 10% reduction each.

9 Discussion and Conclusion

Our results indicate that the supplier fears order cancellations, making her averse to commencing order fulfillment based on soft orders. This results from the fact that the supplier's effort, including procurement of components and the actual building of the equipment, is very customer specific. We also find that the supplier perceives holding cost resulting from an early completion of the order to be high, but perceives the cost of shipping late to be small.

The large emphasis on early completion cost relative to late completion cost clearly does not mirror the overall cost for the supply chain. If the tool is finished early, it remains at the supplier's plant and only traditional inventory holding costs are incurred. However, a late shipment of the tool can lead to idle time and lost output at the buyer's fab, which is associated with substantial margin losses which are magnitudes larger than the holding cost for a piece of equipment. This suggests there is a lack of coordination in the supply chain which can lead to sub-optimal performance.

In this environment, supply chain performance could be improved, if the buyer would be willing to share some of the holding cost. One operational way of doing this would be if the buyer accepted the tool delivery for some time window prior to the RDD. This would reduce the expected holding cost for the supplier and thereby move the optimal starting point T_P^* forward in time. In general, given the high degree of customization demanded by the buyer, supply chain performance also could potentially be improved if the buyer would be ready to share some of the risk of cancellation. This would have two beneficial effects. First, it would reduce the supplier's cancellation cost, moving the optimal starting point T_P^* forward in time. Second, it would make the forecast more credible, and thereby rebuild some of the trust missing in the system. In presence of a cancellation fee, phantom orders become costly to the buyer, allowing the supplier to have more confidence in the soft orders. In the long-term, the supply chain would also benefit from a more platform-based design of the production equipment. The more components that are shared across customized designs, the easier it becomes for the supplier to redirect a half-finished piece of equipment that was initially built for customer A to customer B. This would also reduce the cost of cancellation for the supplier.

The sharing of demand forecast information does not have a positive value for the supplier, who distrusts these data and delays the production start for the equipment. In our situation of voluntary compliance (the buyer can not monitor the supplier), financial incentives are needed to make signals related to the forecast credible (Cachon and Lariviere 2001). As the buyer, who is in control over the design of the coordinating mechanisms in the supply chain, does not incur any financial loss in case of a cancellation, such credible signaling is not possible. This would change if the buyer would pay some cancellation fee (potentially as a function of time). While having the right of free cancellation obviously is attractive to the buyer and potentially saves her some direct out-of-pocket cost, the buyer pays a (much higher) price indirectly, resulting from long tool delivery lead-times generated by the supplier's response to the current system structure.

We believe that these results are of substantial interest, both from an academic as well as from a direct managerial perspective. On the academic side, these results provide the first econometric evidence of problems related to forecast sharing. While there is a rapidly growing stream of research following Lee *et al.* (1997), no previous study could empirically demonstrate the existence of the coordination problems analyzed analytically. Our results provide both an empirical foundation for this important stream of research as well as quantifications (estimates) of some of the most important parameters.

From a managerial perspective, our results demonstrate that information sharing by itself is not sufficient to build superior supply chain performance. Our results were presented to senior executives at both buyer and supplier and have started several projects related to overcoming some of the credibility problems we uncovered. In particular the buyer, who owns the forecast sharing mechanism, is now working on implementing a new system with the goal of making her forecasts more credible.

Finally, we believe that our work not only serves as the empirical foundation for much of the contracting research, but that it also provides a fruitful starting point for future research. A larger empirical study could analyze how the cost parameters that we estimated change over time. For example, one would expect that the cancellation of an order directly increased the supplier's perception of cancellation cost in the subsequent period. Another interesting research opportunity relates to how the forecast is shared. Similar to the field of concurrent product development discussed in the literature review, where there has been a recent trend towards set-based - opposed to point-based - information exchange, the buyer could provide multiple scenarios of demand to the buyer or could even share a confidence interval. This would be consistent with established supply chain concepts such as minimum purchase commitments and its effect on forecast credibility would be interesting to study both analytically as well as empirically.

10 References

Cachon, G. P. and M. Lariviere. 2001. Contracting to assure supply: how to share demand forecasts in a supply chain. Management Science 47 (5), 629-646.

Chen, F. 2000. Salesforce incentives, market information, and production/inventory planning. Working paper, Columbia Business School.

Chen, F. 2001. Market segmentation, advanced demand information, and supply chain performance. Manufacturing & Seroice Operations Management 3 (1), 53-67.

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

Corbett, C. 2001. Stochastic Inventory Systems in a Supply Chain with Asymmetric Information. Operations Research, forthcoming.

Corbett, C. and C. Tang. "Designing Supply Contracts: Contract Type and Information Asymmetry." Chapter 9 in Quantitative Models for Supply Chain Management, S. Tayur, R. Ganeshan and M. Magazine (eds.), Kluwer Academic Publishers, Boston, pp. 270-297, 1999.

Desai, P. S. and K. Srinivasan. 1995. Demand signaling under unobservable effort in franchising: Linear and nonlinear price contracts. Management Science 41 (10), 1608-1623.

Duenyas, I., "Single Facility Due Date Setting with Multiple Customer Classes", Management Science 41(4), 608-619, April 1995.

Gallego, G. and O. Ozer. 1999. Integrating replenishment decisions with advance demand information. Working paper, Columbia University and Stanford University.

Ha, AY. 1999. Supplier-buyer contracting: Asymmetric cost information and cutoff level policy for buyer participation, Naval Research Logistics 48 (1): 41-64

Hariharan, R. and P. Zipkin. 1995. Customer order information, lead times, and inventories. Management Science 41 (10), 1599-1607.

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

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

Lee, H., P. Padmanabhan, and S. Whang. 1997a. The bullwhip effect in supply chains. Sloan Management Review 38,93-102.

Lee, H., P. Padmanabhan, and S. Whang. 1997b. Information distortion in a supply chain: The bullwhip effect. Management Science 43, 546-558.

Loch, C. H., C. Terwiesch, "Communication and Uncertainty in Concurrent Engineering", Management Science, 44(8), August 1998.

Lucas, R. E. "Econometric Policy Evaluation: A Critique" in Bruneer, K. and A. H. Metzler (eds.), The Phillips Curve and Labor Markets, North-Holland, Amsterdam, 1976.

Murphy, K.M., Topel, R., "Estimation and Inference in Two-Step Econometric Models," Journal of Business and Economic Statistics, October 1985.

Porteus, E., S. Whang. 1999. Supply chain contracting: Non-recurring engineering charge, minimum order quantity, and boilerplate contracts. Working paper, Stanford University, Stanford, CA. Cachon, G. P., Lariviere, M. A., "Contracting to assure supply: How to share demand forecasts in a supply chain", Management Science 47 (5), 629-646, May 2001.

Tang, C., K. Rajaram, J. Ou, and A. Alptekinoglu, "The Benefits of Advance Booking Programs: Model and Analysis," Management Science, forthcoming.

Terwiesch, C, A. De Meyer, C. H. Loch, "A Framework for Exchanging Preliminary Information in Concurrent Development Processes", Organization Science, forthcoming.

Table 1. Estimation of Lead Time

Coefficient	Estimate	Standard Error	P-value
Constant	4.4922	0.1053	0.0000
DEVELOP	-0.4124	0.2650	0.1225
NEW	-1.2645	0.2877	0.0000
PREMIUM	0.4844	0.2868	0.0940

R^2=30%. Number of Observation=115.

Table 2. Parameter Estimate: The Complete Model

Variable	Coefficient	Estimate	Standard Error	P-value
Transformed Cancellation Cost	θ_1	0.6215	0.3749	0.0039
Transformed Holding and Delay costs	θ_2	3.4757	0.0509	0.0000
Lead Time Shift Parameter	δ	0.0000	0.1980	0.5000
Standard Deviation	σ	1.9059	0.0000	0.0000
Constant	γ_1	-51.9441	0.0973	0.0000
DEVELOP	γ_2	15.5830	0.0701	0.0000
CHRDD	γ_3	-6.9420	0.0840	0.0000
NEW	γ_4	-7.6690	0.0700	0.0000
PREMIUM	γ_5	6.0040	0.0785	0.0000
Constant	ρ_1	17.2660	0.0699	0.0000
DEVELOP	ρ_2	-9.9809	0.0982	0.0000
CHRDD	ρ_3	-7.3851	0.1566	0.0000
NEW	ρ_4	-6.9859	0.0711	0.0000
PREMIUM	ρ_5	7.3510	0.0695	0.0000

$g = 1.0000, h = 4.2800, c = 2.0715$. LL=-206.5000. Number of Observation=100.

Table 3. Summary: General and Nested Models

Model	Log-Likelihood	χ^2-test statistic	P-value
General Model	-206.5000		
Nested 1: DEVELOP	-217.6000	22.2000	0.0000
Nested 2: CHRDD	-218.0000	23.0000	0.0000
Nested 3: NEW	-213.7480	14.4960	0.0007
Nested 4: PREMIUM	-214.4000	15.8000	0.0004

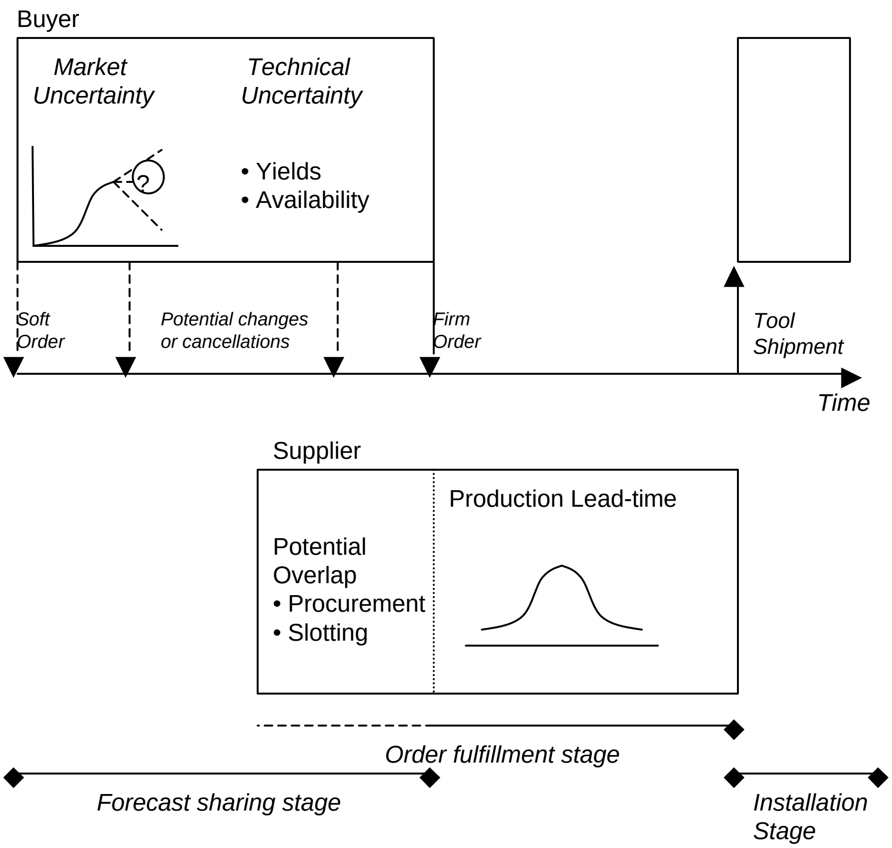

Figure 1: The equipment acquisition cycle

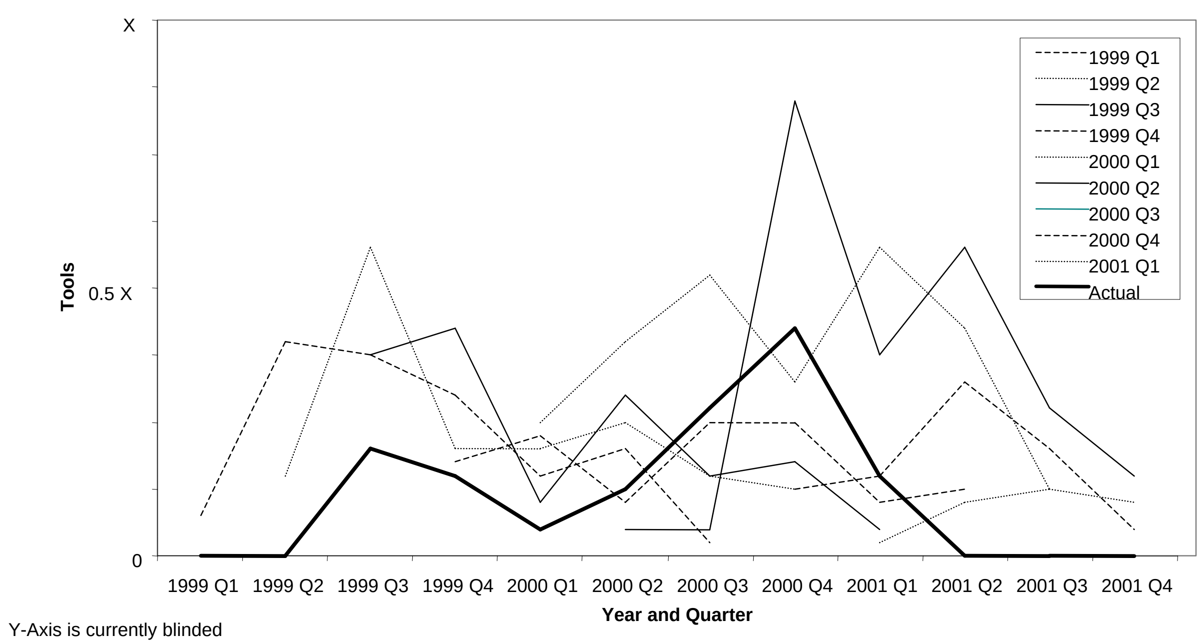

Figure 2: Forecasted (soft) orders versus actual orders

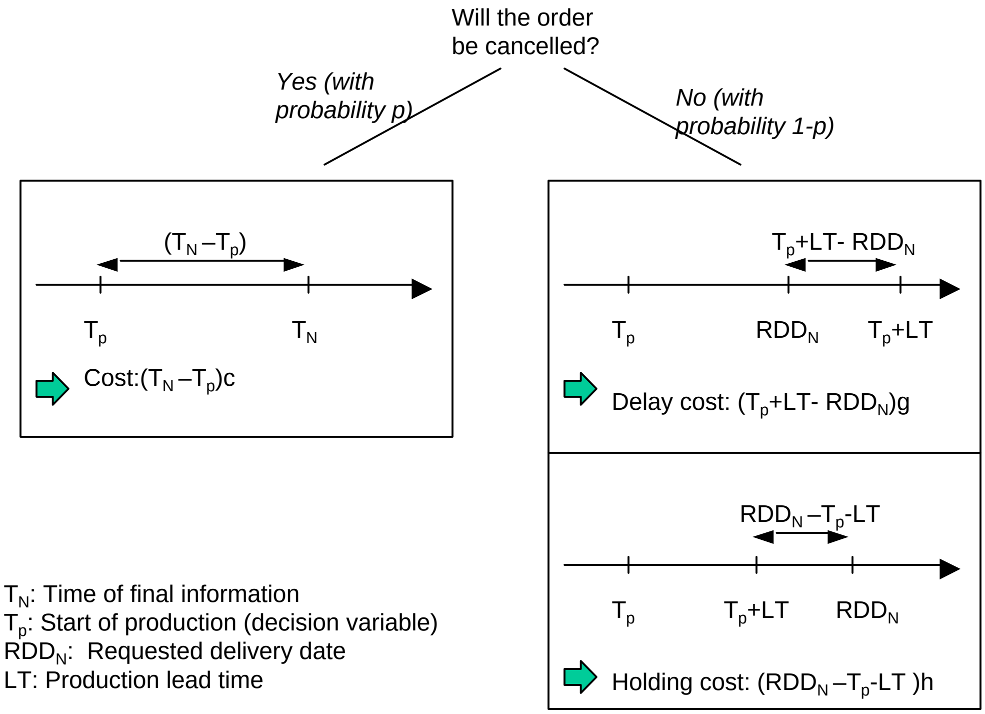

Figure 3: Variable definitions and basic cost drivers

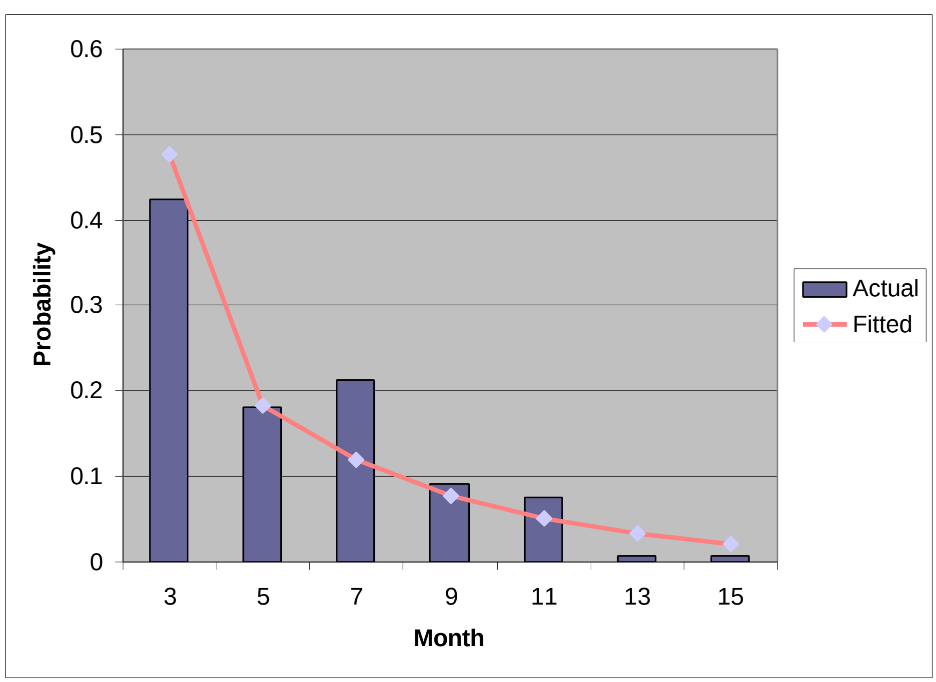

Figure 4: Distribution of G(.)

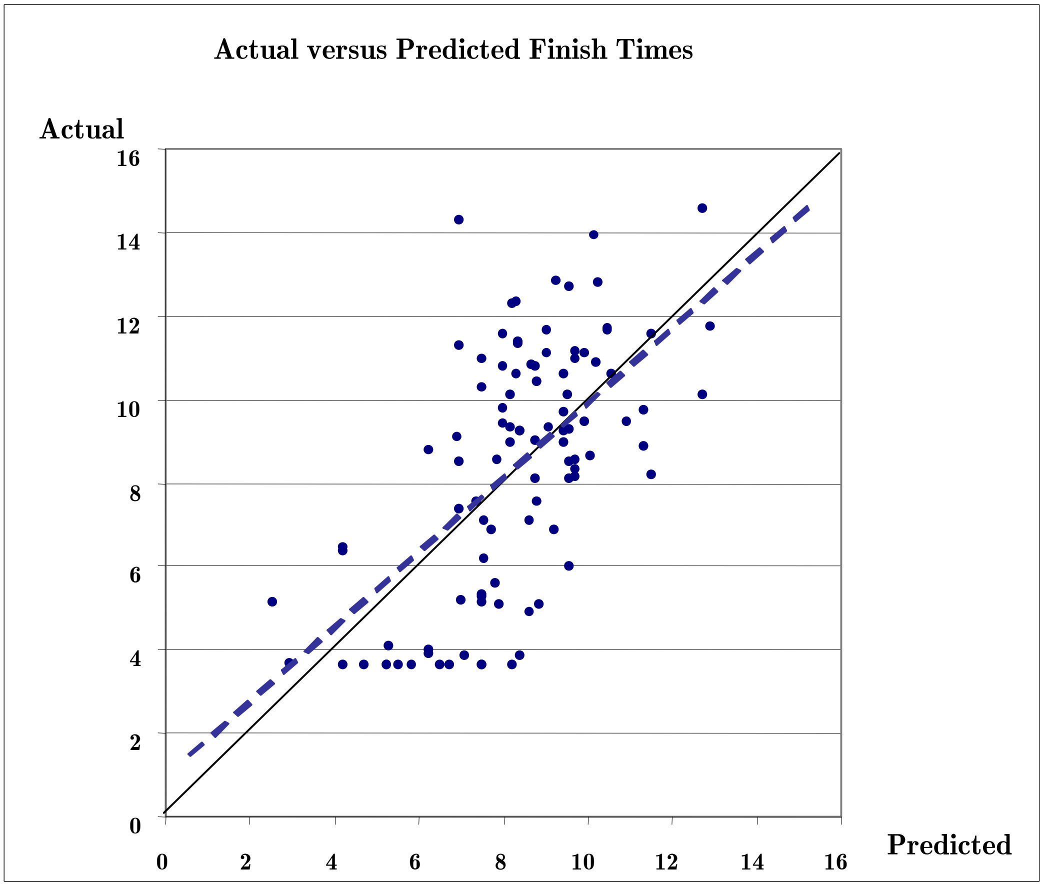

Figure 5: Model Validation

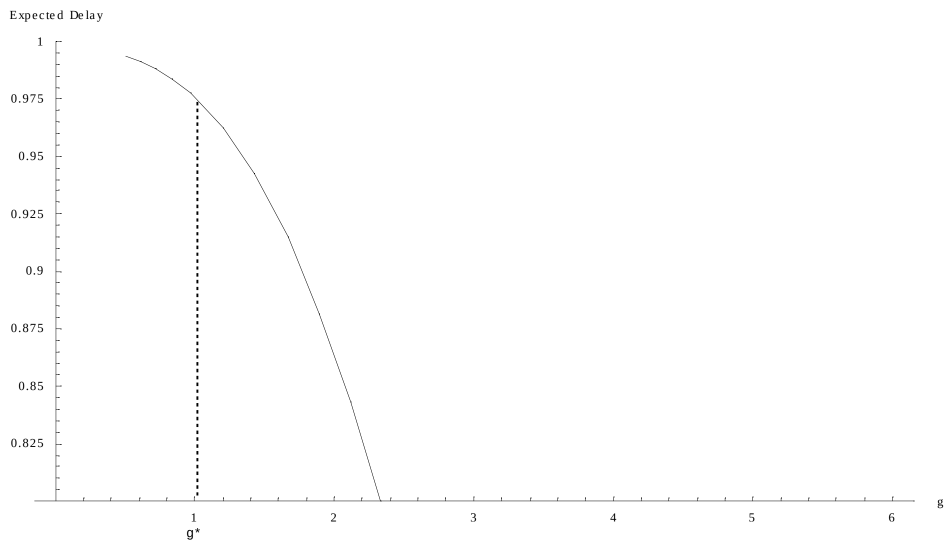

Figure 6: Impact of delay cost parameter *g* on expected delay

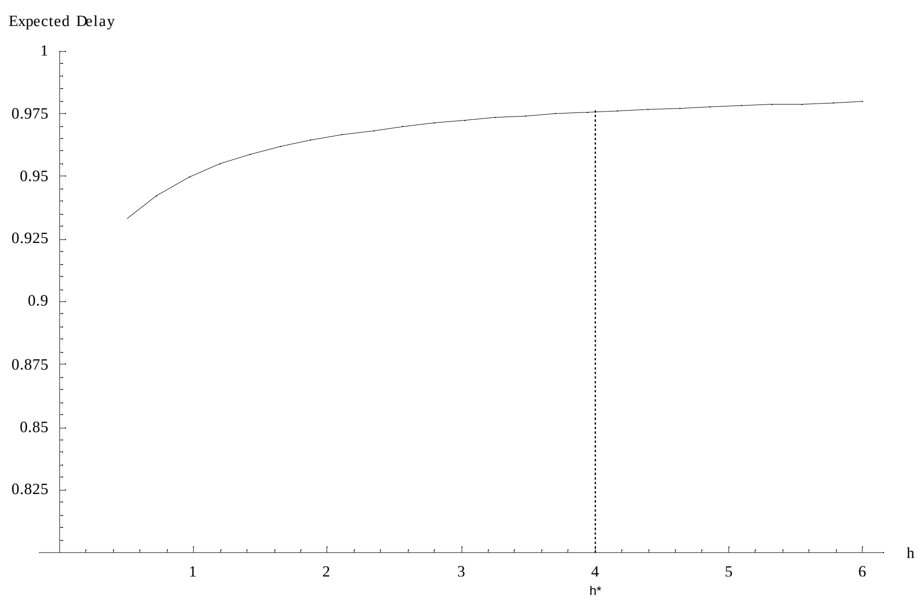

Figure 7: Impact of holding cost parameter *h* on expected delay

(a)

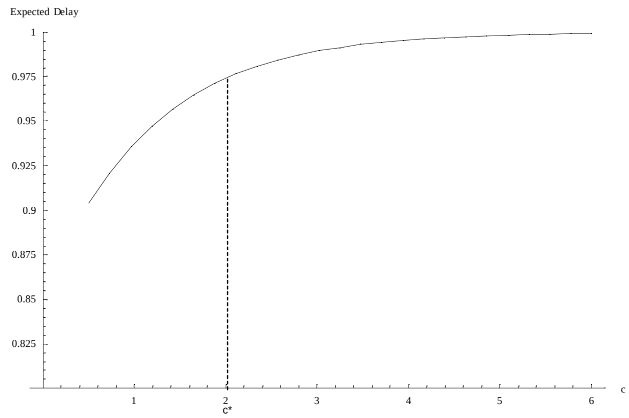

(b)

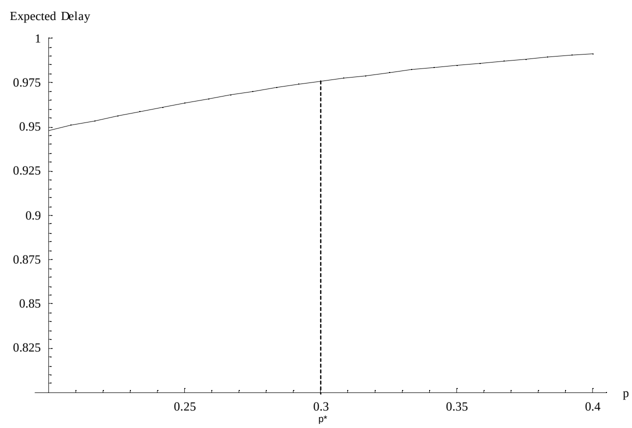

Figure 8: Impact of cancellation cost *c* (a) and probability *p* (b) on expected delay

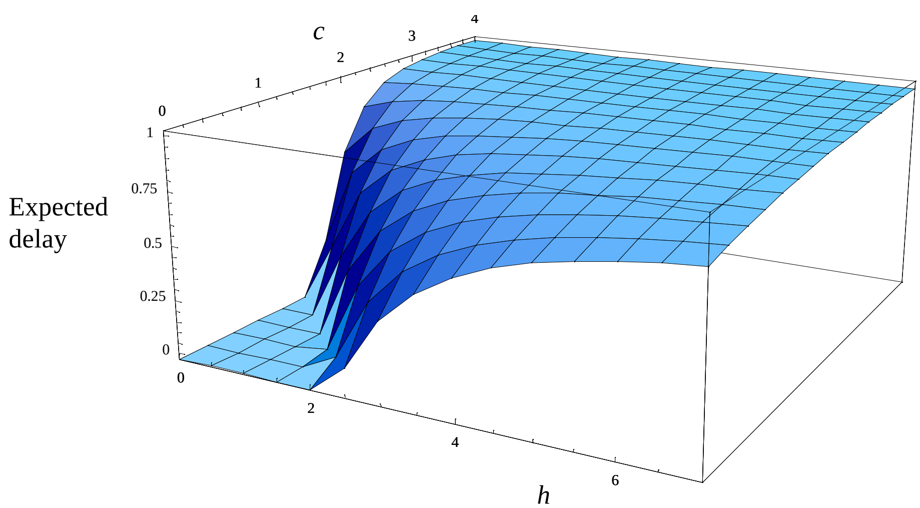

Figure 9: Joint impact of cancellation cost *c* and holding cost *h* on expected delay