An Internet Appendix to:

The Impact of Duplicate Orders on Demand Estimation and Capacity Investment

Mor Armony
Stern School of Business, New York University, New York, NY 10012

Erica Plambeck
Graduate School of Business, Stanford University, Stanford, CA 94305

Abstract

Motivated by a $2.2 billion inventory write-off by Cisco Systems, we investigate how duplicate orders can lead a manufacturer to err in estimating the demand rate and customers' sensitivity to delay, and to make faulty decisions about capacity investment. We consider a manufacturer that sells through two distributors. If a customer finds that his distributor is out-of-stock, then he will sometimes seek to make a purchase from the other distributor; if the latter is also out-of-stock, the customer will order from both distributors. When his order is filled by one of the distributors, the customer cancels any duplicate orders. Furthermore, the customer cancels all of his outstanding orders after a random period of time.

Assuming that the manufacturer is unaware of duplicate orders, we prove that she will overestimate both the demand rate and the cancellation rate. Surprisingly, failure to account for duplicate orders can cause short-term underinvestment in capacity. However, in long term equilibrium under stable demand conditions the manufacturer overinvests in capacity. Our results suggest that Cisco's write-off was caused by estimation errors and cannot be blamed entirely on the economic downturn. Finally, we provide some guidance on estimation in the presence of double orders.

The majority of the results in our paper apply to the base case where the manufacturer dedicates a fixed fraction of her capacity to each distributor. This Appendix shows that the results remain true when the manufacturer fills distributors' orders FIFO from pooled capacity.

A

A.1 Introduction

To investigate the impact of duplicate orders on parameter estimation and capacity investment, production is modelled as a Poisson process with rate μ. We investigate two plausible allocation rules for finished goods. The base case in the paper is that each distributor i has a fixed portion of the capacity μ_i (with $\mu_1 + \mu_2 = \mu$). This is relevant when the distributors are located in different geographical regions and transportation costs are high, so the manufacturer serves them from different production facilities. (Cisco has regional production facilities, and some resellers will duplicate-order from distributors in different regions to obtain Cisco products.) The base case with $\mu_1 = \mu_2$ also approximates a "fair" division of the output between the distributors. The second case, which is the focus of this appendix, is that distributors' orders are filled FIFO. (Cisco is concerned with fairness.). For this case we assume that duplicate orders are placed in front of each other in the manufacturer's queue with equal probabilities. Our purpose in this appendix is to show that all of the results stated in the paper also hold for the FIFO system, with one minor exception: we have proven Proposition 3 for FIFO only for the special case alpha =1. To underline the equivalence between the results of the main paper to this Appendix's results, we label all sections and propositions the same, preceded by an "A" (for "Appendix").

A.2 Model Formulation

Consider a manufacturer that sells a single product through two independent distributors. For brevity, we will at times use the pronoun 'he' to refer to a distributor and 'she' to refer to the manufacturer. At each distribution center, customers arrive according to a Poisson process with rate λ (which is independent of customer arrivals at the other distribution center) and each customer demands one unit of the product. Let $X_i(t)$ denote the inventory level for distributor i $(i = 1, 2)$ at time t; $[X_i(t)]^- = -\min\{X_i(t), 0\}$ indicates the number of outstanding orders from customers. If a customer arrives when his distributor is out of stock $(X_i(t) \leq 0)$ then with probability α the customer buys immediately from the other distributor (if the other distributor has inventory) or orders the product from *both* distributors. As soon as one of the distributors delivers the product to him, the customer will cancel the duplicate order. With probability $1 - \alpha$ the customer orders from his original distributor only. The customer is impatient; after waiting for a time that is exponentially distributed with rate η, he will cancel all outstanding orders and leave the system without making a purchase. We will denote by $D(t)$ the number of duplicate orders that are outstanding at time t. Clearly, the number of customers waiting for the product at time t is given by $[X_1(t)]^- + [X_2(t)]^- - D(t)$, and if $X_i(t) \geq 0$ for either $i = 1$ or 2, then $D(t) = 0$.

Each distributor follows a base stock policy. In particular, each distributor orders one unit from the manufacturer every time a customer orders a unit from him, and cancels an order with the manufacturer every time a customer cancels an order with him. Let $Y_i(t)$ denote the number of outstanding orders from distributor i to the manufacturer. Then, $Y_i(t) = B - X_i(t)$, where B is the base stock level. The manufacturer does not hold inventory and has a total production capacity of rate μ. When she has outstanding orders from either distributor ($Y_1(t)+Y_2(t) > 0$) the manufacturer makes the product according to a Poisson process with rate μ, and delivers it according to a FIFO rule between the distributors (with simultaneous orders placed in front of each other in the manufacturer's queue with equal probabilities). The manufacturer knows the base stock policy used by the distributors, and can therefore infer the inventory level and the number of customer orders outstanding for each distributor from her own order process $(Y_1(t), Y_2(t))$. Furthermore, the manufacturer knows when a downward transition in $Y_i(t)$ corresponds to an order cancellation and when it corresponds to an order fulfillment, and therefore effectively observes the orders and cancellations made by customers.

To completely describe the system dynamics, it remains to specify the sequence in which customer orders are filled. We will assume that each distributor knows which of his customers have placed a duplicate order, and gives priority to serving these customers (to avoid losing a sale to the other distributor). The assumption that distributors can identify double orders is plausible because any customer that double orders has an incentive to reveal this to the distributors, in order to shorten his lead time. Furthermore, software for channel management enables distributors to share information in real time about customer identity and purchasing behavior. Note that under this FIFO rule $(X_1(t), X_2(t), D(t))$ is not a Markov chain. One needs to specify the sequence in which orders were made with the manufacturer, in order to fully describe the dynamics of the system. To this end, let $(D; C) = (d; c(1), c(2), ..., c(n))$ be the state descriptor, with d the number of outstanding duplicate orders, n the total number of outstanding orders (counting double orders twice), and $c(k) \in \{1, 2\}$ the distributor that made the order which is currently in position k in the manufacturer's queue ($k = 1, ..., n$). Then, under these two sequencing and prioritizing assumptions $(D; C)$ is a continuous time Markov chain.

Finally, to guarantee that the Markov process $(D; C)$ is ergodic, we assume that the number of dedicated orders at each distributor and the number of duplicate orders are bounded by a very large number M; that is, $X_i^- - D \leq M$ for $i = 1, 2$, and $D \leq M$. Throughout, we omit the time index t whenever we refer to the whole process, and write $t = \infty$ to refer to the process in steady state.

A.3 ML Estimation when $\alpha = 0$ *(in the manufacturer's opinion)*

For some arbitrary long time horizon $T > 0$, let $N_i(T)$ be the number of customers that have ordered from distributor i, $i = 1, 2$, and let $Z_i(T)$ be the number of these orders that have been cancelled.

Proposition A.1. *For the system with* $\alpha = 0$*, the maximum likelihood estimators for* λ *and* η *are given by:*

$$\widehat{\lambda}(T) = \frac{N_1(T) + N_2(T)}{2T}$$

and

$$\widehat{\eta}(T) = \frac{Z_1(T) + Z_2(T)}{\int_0^T [Y_1(t) - B]^+ + [Y_2(t) - B]^+ dt} = \frac{Z_1(T) + Z_2(T)}{\int_0^T X_1^-(t) + X_2^-(t) dt},$$

respectively.

Proof: This problem can be viewed as estimating the transition rate parameters in a continuous time Markov chain (CTMC). In particular, since there are no double orders $C = (c(1), c(2), ..., c(n))$ is a CTMC with generator (transition rate matrix) Q, which satisfies:

1. **Incoming order:** $Q(C, (C, i)) = \lambda$, for $i = 1, 2$ (an order made by distributor i).
2. **Order cancellation:** $Q((C, i, \tilde{C}), (C, \tilde{C})) = \eta 1_{\{Y_i(\tilde{C}) \geq B\}}$, where $Y_i(\tilde{C})$ is the number of orders in $\tilde{C}$ made by distributor i.
3. **Order completion and delivery:** $Q((i, C), C) = \mu$ (the head-of-line order delivered to distributor i).

Note that there is a certain overlap between 2. and 3., whenever the cancelled order is at the head of the line. We assume that the manufacturer is able to distinguish between these two types of transitions. For a given state C, we count the number of transitions out of this state during the time interval $[0, T]$, including $N(C; T)$ arrivals, $Z(C; T)$ order cancellations and $E(C; T)$ service completions. In addition, let $\tau(C; T)$ denote the total amount of time during this interval that the state of the system is C. Thus, the likelihood function given the observation of $Y(\cdot)$ can be written as follows:

$$\begin{aligned} \pounds(\lambda,\eta) &= \prod_C \exp(-\{(2\lambda+\mu 1_{\{n(C)>0\}}+\eta\sum_{i=1,2}[Y_i(C)-B]^+)\tau(C;T)\}) \\ &\quad \cdot(2\lambda)^{N(C;T)}(\eta\sum_{i=1,2}[Y_i(C)-B]^+)^{Z(C;T)}\mu^{E(C;T)} \\ &= \exp(-\{2\lambda T+\eta\sum_C\sum_{i=1,2}[Y_i(C)-B]^+\tau(C;T)\})\lambda^{N(T)}\eta^{Z(T)}K, \end{aligned}$$

where K stands for a constant that involves only terms that are not a function of λ or η. One can easily see that the values of λ and η that maximize $\pounds(\lambda,\eta)$ are $\widehat{\lambda}(T)$ and $\widehat{\eta}(T)$ as given in the statement of the proposition. □

Systematic Error in Maximum Likelihood Estimation

Now suppose that some customers double order ($\alpha > 0$) unbeknown to the manufacturer, who uses the MLE for the system with $\alpha = 0$. To compute the resulting systematic error in estimation, some additional notation is needed. The superscript '0' will indicate that the distributor is out of stock and the superscript '1' will indicate that the corresponding distributor has items in stock. The first superscript will refer to distributor 1, and the second to distributor 2. For example, $N_i^{01}(T)$ denotes the number of orders placed with distributor i up to time T, when distributor 1 is out of stock, and distributor 2 has some items in inventory immediately prior to the arrival. Similarly, $N_i^{00}(T)$ denotes the number of orders placed with distributor i while both distributors are out of stock (immediately before the customer arrives). $N_i^{10}(T)$ and $N_i^{11}(T)$ are defined in an analogous fashion. Also, let $\tau^{00}(T)$ be the total time the system spends in states where both distributors are out of stock during the corresponding time interval, and let $P^{00} = P(X_1(\infty) \leq 0, X_2(\infty) \leq 0)$ be the steady state probability that both distributors are out of stock. Finally, recall that $D(\infty)$ is the steady state (random) number of duplicate orders in the system and EX_i^- is the expected backlog level at distributor i in steady state.

The next proposition characterizes the systematic error, establishing that the manufacturer will overestimate the demand rate and the reneging rate. The statement of the proposition as well as its proof are completely analogous to the statement and proof of Proposition 2 which appear in the paper. For completeness, we present them in their entirety here.

Proposition A.2. *Suppose that customers double order with positive probability ($\alpha > 0$) but the manufacturer uses the MLE for the system with $\alpha = 0$ given in Proposition A.1. Then, the systematic error in estimating the demand rate is given by*

$$\widehat{\lambda} - \lambda = \lambda\alpha P^{00} > 0,$$

and the systematic error in estimating the reneging rate is given by

$$\widehat{\eta} - \eta = \frac{\mu P(D(\infty) > 0)}{EX_1^- + EX_2^-} > 0,$$

where $\widehat{\lambda} = \lim_{T\to\infty} \widehat{\lambda}(T)$ *and* $\widehat{\eta} = \lim_{T\to\infty} \widehat{\eta}(T)$.

Proof: Given the notation introduced above, the systematic error in estimating the demand rate is

$$\begin{aligned}
\widehat{\lambda} - \lambda &= \lim_{T\to\infty} \widehat{\lambda}(T) - \lambda = \lim_{T\to\infty} \frac{N_1(T) + N_2(T)}{2T} - \lambda \\
&= \lim_{T\to\infty} \frac{N_1^{00}(T) + N_2^{00}(T)}{2\tau^{00}(T)} \frac{\tau^{00}(T)}{T} \\
&\quad + \frac{[N_1^{01}(T) + N_2^{01}(T)] + [N_1^{10}(T) + N_2^{10}(T)] + [N_1^{11}(T) + N_2^{11}(T)]}{2(T - \tau^{00}(T))} \frac{T - \tau^{00}(T)}{T} - \lambda \\
&= (\lambda + \lambda\alpha)P^{00} + \lambda(1 - P^{00}) - \lambda = \lambda\alpha P^{00},
\end{aligned}$$

where the last equality follows from the strong law of large numbers (SLLN) for renewal processes. The above equalities indicate that the systematic error is strictly positive because double orders are counted as true customer arrivals. In addition, they imply that as the production capacity μ increases to ∞, the systematic error goes to 0.

Calculation of the systematic error in the estimator for η is more involved. Because distributors prioritize double orders, whenever $D(t) > 0$ each service completion is coupled with an order cancellation. Hence, one cancellation is seen whenever a non-duplicate order is cancelled or a service completion occurs for a duplicate order. The resulting rate of *one* cancellation at time t is $\eta(X_1^-(t) + X_2^-(t) - 2D(t)) + \mu 1_{\{D(t)>0\}}$. *Two* simultaneous cancellations will be observed with rate $\eta D(t)$. Let $Z_i(x_1, x_2, d; T)$ be the number of order cancellations during the time interval $[0, T]$, when immediately prior to the cancellation $(X_1(t), X_2(t), D(t)) = (x_1, x_2, d)$. We deal first with the numerator of the expression for $\widehat{\eta}(T)$.

$$\begin{aligned}
&\lim_{T\to\infty} \frac{Z_1(T) + Z_2(T)}{T} \\
&= \lim_{T\to\infty} \sum_{x_1 \le B} \sum_{x_2 \le B} \sum_{d \ge 0} \frac{Z_1(x_1, x_2, d; T) + Z_2(x_1, x_2, d; T)}{\tau(x_1, x_2, d; T)} \frac{\tau(x_1, x_2, d; T)}{T} \\
&= \sum_{x_1 \le B} \sum_{x_2 \le B} \sum_{d \ge 0} (\eta(x_1^- + x_2^- - 2d) + \mu 1_{\{d>0\}} + 2\eta d) P(X_i(\infty) = x_i, i = 1, 2, D(\infty) = d) \\
&= \eta(EX_1^- + EX_2^-) + \mu P(D(\infty) > 0),
\end{aligned}$$

where the second equality follows from the SLLN for renewal processes. The denominator of the expression for $\widehat{\eta}(T)$ is simpler to analyze. Specifically,

$$\lim_{T\to\infty} \frac{\int_0^T X_1^-(t) + X_2^-(t)dt}{T} = EX_1^- + EX_2^-,$$

from ergodicity. Hence,

$$\widehat{\eta} - \eta = \lim_{T\to\infty} \widehat{\eta}(T) - \eta = \lim_{T\to\infty} \frac{(Z_1(T) + Z_2(T))/T}{(\int_0^T X_1^-(t) + X_2^-(t)dt)/T} - \eta = \frac{\mu P(D(\infty) > 0)}{EX_1^- + EX_2^-}. \qquad \square$$

Sensitivity Analysis of Systematic Errors

Overestimation of λ and η occurs because the manufacturer fails to recognize the potential for duplicate orders. One might therefore expect that as the production capacity increases, the systematic error will decrease because there is less opportunity for double ordering to occur. In this section, we prove that as expected, the systematic error in the estimator for λ decreases with μ for $\alpha = 1$. In the body of the paper we have a proof of this assertion for general values of $0 < \alpha \leq 1$. However, this proof fails if the manufacturer allocates FIFO between the distributors. Therefore, in this appendix, which is concerned with FIFO, we prove the proposition for the special case where $\alpha = 1$, which turns out to be rather simple.

Proposition A.3. *The systematic error in estimating the demand rate λ is decreasing in μ when the double ordering probability is $\alpha = 1$.*

Proof: From Proposition A.2., the systematic error in the demand rate is $\widehat{\lambda} - \lambda = \lambda\alpha P^{00} = \lambda\alpha P(X_1(\infty) \leq 0, X_2(\infty) \leq 0)$. We will prove that $P(X_1(\infty) \leq 0, X_2(\infty) \leq 0)$, the steady-state probability that both distributors are out-of-stock, decreases with μ.

Note that when $\alpha = 1$ the total number of outstanding orders $Y = Y_1 + Y_2 - D$ (counting double orders only once) is a birth and death process with the following transition rates:

1. **Birth:** $Q(Y, Y+1) = 2\lambda$.

2. **Death:** $Q(Y, Y-1) = \mu 1_{\{Y>0\}} + \eta[Y - 2B]^+$.

In this case, the steady-state probability that both distributors are out-of-stock satisfies: $P(X_1(\infty) \leq 0, X_2(\infty) \leq 0) = P(Y(\infty) \geq 2B)$. Hence, it is sufficient to show that $P(Y(\infty) \geq 2B)$ is decreasing in μ. We show that not only is this steady-state probability decreasing in μ, but, in fact, Y is *path-wise* decreasing in μ. We prove this using sample-path coupling and uniformization.

First, consider the uniformized discrete time Markov chain with one step transition probabilities equal to the corresponding transition rates of the continuous time Markov chain Y, divided by $v \geq 2\lambda + \mu + 2M\eta$. Transitions from a state to itself are allowed in order to ensure that the transition probabilities sum up to 1. The steady state distribution of the uniformized discrete time chain is identical to that of the continuous time Markov chain.

Second, let $\mu_L < \mu_H$, and denote by Y^L the state of the system when $\mu = \mu_L$. Similarly, denote by Y^H the corresponding state descriptor when $\mu = \mu_H$. We use *sample-path coupling* arguments to show that $Y^L \geq^{st} Y^H$, where $\geq^{st}$ denotes stochastic ordering. This will imply, in particular, that $P(Y^L(\infty) \geq 2B) \geq P(Y^H(\infty) \geq 2B)$. The sample path coupling argument works as follows: Note that $Y^L(\cdot) \leq 2M$ and $L^H(\cdot) \leq M$. We can construct versions of $Y^L(\cdot)$ and $Y^H(\cdot)$ (which for notational simplicity are denoted the same as the original processes), such that $Y^L(\cdot) \geq Y^H(\cdot)$ with probability 1. Specifically, we let $v = 2\lambda + \mu^H + 2M\eta$ and assume that $Y^L(0) \geq Y^H(0)$. Then, by coupling the transitions of both chains and using induction on n, it follows that $Y^L(n) \geq Y^H(n)$. □

Proposition 4 in the body of the paper remains unchanged for the FIFO allocation policy. For completeness, we include it here as well.

Proposition A.4. *Let $\widehat{\eta} = \lim_{T\to\infty} \widehat{\eta}(T)$ be the limit of the reneging rate ML estimator as the time horizon grows to infinity. Then $\widehat{\eta} - \eta$ is increasing in μ at $\mu = 0$.*

Proof: It is easy to see that when $\mu = 0$, $\widehat{\eta} - \eta = 0$. However $\widehat{\eta} - \eta > 0$ for all $\mu > 0$. □

A.4 The Manufacturer's "Optimal" Capacity Investment

Suppose the manufacturer chooses capacity according to

$$\min_{\mu} \left[c\eta E(X_1^- + X_2^- - D) + k\mu \right], \tag{A.1}$$

where the decision variable μ is the total manufacturer's capacity, c is the manufacturer's contribution per unit sold, so the first term in the objective function is the expected cost of lost sales; the second term is the cost of capacity. (Without loss of generality, we will assume that $c = 1$.) In practice, a distributors' inventory policy depends on the delivery leadtime and hence upon the manufacturer's capacity. However, in solving for the optimal capacity in (A.1) we disregard strategic interaction, and assume that the base stock level B per distributor is fixed. The solution to (A.1) can be interpreted as the manufacturer's best response to the distributors' inventory policies.

Consider a system in which $\alpha = 1$ (every customer that must wait for the product will place a duplicate order) but the manufacturer believes that $\alpha = 0$. As was observed above, for this system, the total number of outstanding orders at the manufacturer's queue (counting double orders only once) follows a birth and death process with steady state probability distribution:

$$P(Y(\infty)=n)=P(Y(\infty)=0)\cdot\frac{(2\lambda)^n}{\mu^n\prod_{m=2B}^{n}\eta(m-2B)},\quad n=0,1,...$$

This closed form expression allows us to compute the limiting maximum likelihood estimators and the cost function exactly.

When $\alpha=0$ things are more intricate. Here we have a FIFO system with two customer classes, with exponential reneging whenever the number of customers in the queue from one particular class exceeds B. We are not aware of a closed form solution for the steady state distribution of this two dimensional process. However, in the pure loss system (i.e. $\eta=\infty$), the steady-state distribution follows a product form, with the number of customers in queue from each class restricted to be at most B. Under these assumptions, the state descriptor $C=(c(1),c(2),...,c(n))$ (with n the total number of outstanding orders, and $c(k)\in\{1,2\}$ is the distributor who made the order that is currently in position k in line ($k=1,...,n$)) is a continuous time Markov chain. Let $Y_i(C)=\sum_{k=1}^{n}1_{\{c(k)=i\}}$, $i=1,2$, be the number of orders from distributor i currently in queue; $Y_i(C)$ is constrained to the set $\{0,...,B\}$. Without this constraint the steady state of the system has a product form, and hence, it has this form on the constrained state space as well. In summary, the steady state distribution is as follows:

$$\pi(C)=\frac{b}{A}\left(\frac{\rho}{2}\right)^n,$$

where $\rho=\frac{2\lambda}{\mu}$, $b=\frac{1-\rho}{1-\rho^{2B+1}}$ if $\rho\neq1$, and $b=\frac{1}{2B+1}$ when $\rho=1$. Also, $A=1-b\sum_{n=B+1}^{2B}M_n\left(\frac{\rho}{2}\right)^n$, with $M_n=2\sum_{k=B+1}^{n}\binom{n}{k}$. The loss rate in this system is $\lambda P(Y_1(C)=B)+\lambda P(Y_2(C)=B)=2\lambda P(Y_1(C)=B)=2\lambda\sum_{n=B}^{2B}\binom{n}{B}\frac{b}{A}\left(\frac{\rho}{2}\right)^n$.

The expected rate of lost sales is strictly lower in the system with $\alpha=1$ than in the system with $\alpha=0$ because, in choosing to double order, each customer increases the likelihood that she will obtain the product before reneging. Effectively, inventory and capacity are pooled in the system with $\alpha=1$. Furthermore, as illustrated by Figure A.1, when the capacity μ is very small, the marginal value of capacity is greater in the system with $\alpha=1$ than in the system with $\alpha=0$ (i.e., increasing μ does more to reduce lost sales when $\alpha=1$ than when $\alpha=0$). However, if the capacity μ is sufficiently large, additional capacity is more beneficial when $\alpha=0$ than when $\alpha=1$. Therefore, if the manufacturer knows the true demand rate and reneging rate, but incorrectly assumes that $\alpha=0$, he will underinvest when the cost of capacity k is large and overinvest when the cost of capacity k is small.

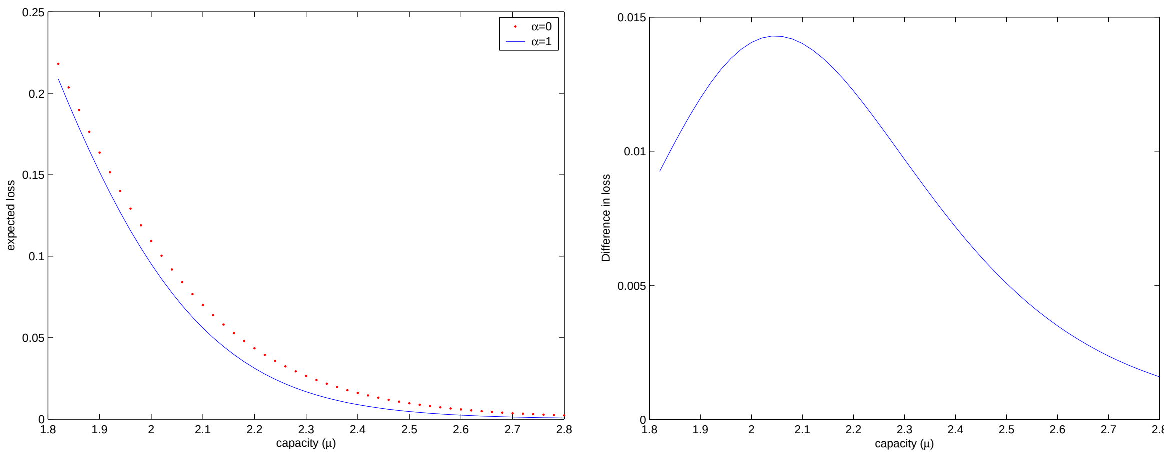

Figure A.1: The difference in the expected rate of lost sales in the case $\alpha = 0$ and the case $\alpha = 1$, for a system with $\lambda = 1$, $\eta = \infty$, and $B = 10$.

Suppose that the manufacturer has been operating the system at some fixed initial level of capacity μ, and uses the maximum likelihood estimators $\widehat{\lambda}$ and $\widehat{\eta}$ (which depend on the level of capacity μ) to compute his "optimal" capacity investment. Figure A.2 shows that when the initial capacity level is larger than the demand rate and the cost of capacity is relatively large, this "optimal" capacity investment will be strictly smaller than the true optimal capacity. That is, the manufacturer will underinvest in capacity.

Now, let us suppose that the manufacturer *repeatedly* runs the system for long enough to compute the estimators $\widehat{\lambda}$ and $\widehat{\eta}$, and then adjusts capacity to the "optimal" level. In all of our numerical experiments, the capacity converges to an equilibrium which appears to be "optimal" if the manufacturer assumes that $\widehat{\lambda}$ and $\widehat{\eta}$ (evaluated at the current capacity level) are the true demand rate and reneging rate. Figure A.3 shows that in equilibrium, the manufacturer overinvests in capacity.

A.5 Maximum Likelihood Estimation when $\alpha > 0$

If the manufacturer is aware of the potential for double orders and observes the system continuously, she can recognize a double order whenever both distributors order simultaneously, or cancel an order simultaneously. One may argue that in reality no two events will occur at exactly the same time. As a practical alternative, if the manufacturer has visibility of end customers' identities, she can pair two orders made by the same customer at approximately the same time, and recognize a double order. In this section, we spell out the maximum likelihood estimators of λ, η and α in the case

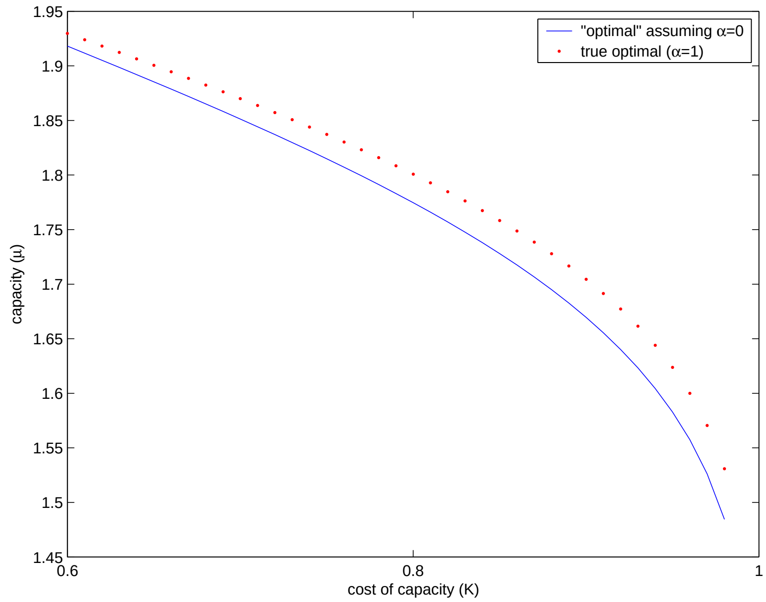

Figure A.2: True optimal capacity and the "optimal" capacity investment for a manufacturer who assumes that $\alpha = 0$, $\lambda = \widehat{\lambda}$ and $\eta = \infty$, for the system with $\alpha = 1$, $\eta = \infty$, $\lambda = 1, B = 10$ and initial capacity $\mu_1 = 2.5$.

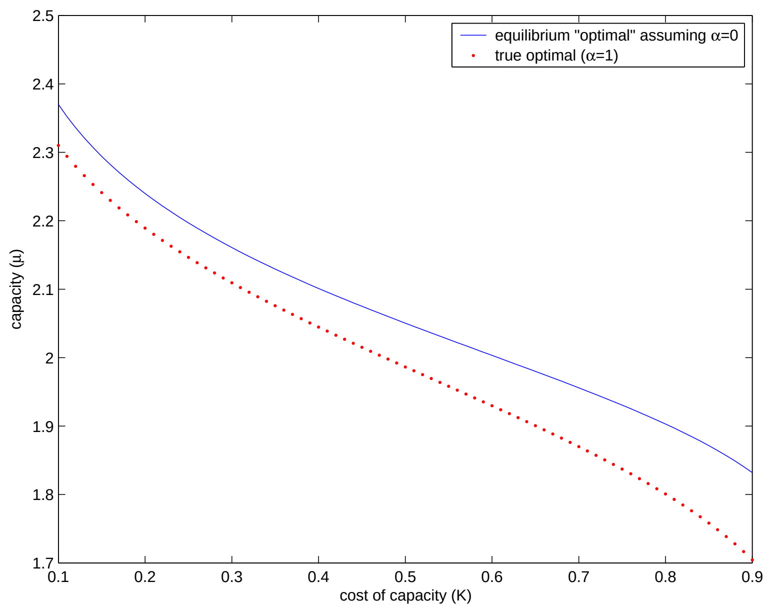

Figure A.3: True optimal capacity and equilibrium "optimal" capacity investment for a manufacturer who assumes that $\alpha = 0$, $\lambda = \widehat{\lambda}$ and $\eta = \infty$, for the system with $\alpha = 1$, $\eta = \infty$, $\lambda = 1$ and $B = 10$.

of full information (continuous observation or visibility of customers' identities). These estimators are valid for a general shipment schedule from the manufacturer to the distributors.

In order to write down the maximum likelihood estimators for these three parameters, we need to introduce some additional notation. Let $N(T)$ denote the total number of orders made by both distributors in the period $[0, T]$ (accounting only once for those orders that are immediately switched from an out-of-stock distributor to a distributor with item in inventory, but twice for double orders). Also, let $Z(T)$ correspond to the total number of order cancellations from both distributors in the same time interval. Let $D_{in}(T)$ be the total number of duplicate orders made between time 0 and time T (counting only those duplicate orders that occur while both distributors are out of stock), and let $D_{out}(T)$ be the total number of double orders (counted in $D_{in}(T)$) that have both been cancelled by time T. Finally, let $Sw(T)$ be the number of customers who switch from an out-of stock distributor to one with positive inventory in the time interval $[0, T]$, and denote by $N_{-}(T)$ the total number of arriving customers who find the first distributor they turn to being out-of-stock. As in the case with $\alpha = 0$ analyzed in Section A.3, maximizing the likelihood function given continuous time transition information yields the maximum likelihood estimators described in the following proposition. The proof of Proposition A.5. is very similar to that of Proposition A.1., hence omitted.

Proposition A.5. *The maximum likelihood estimators of* λ, η *and* α *are given by:*

$$\widetilde{\lambda}(T) = \frac{N(T) - D_{in}(T)}{2T},$$

$$\widetilde{\eta}(T) = \frac{Z(T) - D_{out}(T)}{\int_0^T [X_1^-(t) + X_2^-(t) - D(t)]dt},$$

$$\widetilde{\alpha}(T) = \frac{D_{in}(T) + Sw(T)}{N_{-}(T)}.$$

These estimators are consistent: $(\widetilde{\lambda}(T), \widetilde{\eta}(T), \widetilde{\alpha}(T)) \to (\lambda, \eta, \alpha)$ *as* $T \to \infty$*.*