Optimum Policies for Dynamic Supply Chain Management: Competitive vs. Cooperative View

A. Waddell[1] & J. Hosang[1]

[1] Institute for Data Analysis and Process Design, Zurich University of Applied Sciences, Technopark, Jägerstrasse 2, 8401 Winterthur, Switzerland, juerg.hosang@zhwin.ch

Abstract

In this paper, we study a supplier-retailer supply chain under nonstationary stochastic demand. Depending on the logistics cost rates of the supply chain partners and their accounting schemes, there may be substantial differences between individually and collectively optimum inventory management policies. We propose transfer payments between the supply chain partners as a measure to eliminate incentives for departing from collectively optimum policies and thus make them incentive compatible. In case the supply chain faces changes in the demand pattern, demand forecasts are required for both stages for periodic updates of their strategy parameters. We propose a simple smoothing technique to obtain these forecasts. Both, transfer payments and the forecasting method combined, provide a method to manage a supply chain optimally with respect to three criteria. First, total operating costs are minimized, second, no member of the supply chain has incentives to deviate from collective optimality and third, the system is reactive to changes in demand patterns.

1. INTRODUCTION

The objective of a supply chain is to maximize profit, i.e., the difference between the revenues generated from customers and the costs incurred in supply chain operation. This objective can be best accomplished if the supply chain is well coordinated. Lack of coordination may lead to phenomena such as the well-known bullwhip effect, i.e., increased demand fluctuations at different stages in the upstream direction of the supply chain. The bullwhip effect can cause severe inefficiencies because maintaining acceptable service levels may require very high inventory levels.

Despite the above facts many supply chains are poorly coordinated. There are many obstacles to coordination, such as conflicting incentives of individual supply chain partners, lack of adequate information at the different stages or operational inadequacies, leading e.g. to large replenishment lead times.

This paper focuses on incentive obstacles for coordination in supply chains. If supply chain partners would care only for systemwide performance, they would behave in such a way as to maximize overall profit. Unfortunately, the behavior which maximizes overall profit may not maximize each partner's individual profit. E.g. suppliers typically do not care for consumer backorders and therefore tend to hold lower inventories than those required for an overall optimum. In practice, neither are supply chain partners altruistic nor is there any authority which can enforce decisions to be taken with only overall system performance in view. Transfer payments between supply chain members are a possibility to eliminate their incentives to deviate from collectively optimum behavior.

We study a one-product supply chain consisting of a supplier and a retailer under stochastic nonstationary demand. The supplier is fed by an outside source with infinite capacity. Further, there are constant replenishment lead times between this source and the supplier and between the supplier and the retailer. Costs are accounted for holding inventory and for consumer backorders, but there are no fixed ordering costs. We propose a method to manage such a system with three objectives in mind:

- Inventory replenishment policies for all stages are due to minimize total supply chain costs.
- The system is to be incentive compatible in the sense that no member of the supply chain can reduce his own costs by deviating from the behaviour required for minimum total supply chain costs.
- The supply chain is to be capable of responding to instationarities in demand.

2. DESCRIPTION OF THE SYSTEM AND MINIMUM COSTS

The system corresponds to that studied by [1] and [2] and we use the same notation. Consider a one-product supply chain with one supplier (Stage 2) and one retailer (Stage 1). Time is divided into an infinite number of periods of length T. There are constant lead times L_2 for shipments from the outside source to the supplier and L_1 for shipments from the supplier to the retailer. Each stage may order any nonnegative amount of the product in each period at the next higher stage. There are no fixed order costs and no quantity discounts. The supplier is charged holding cost h_2 per period for each unit in his stock or in transit to the retailer. The retailer's holding cost is $h_1 + h_2$ per period for each unit in his stock ($h_1 \geq 0, h_2 > 0$). Customer demand which cannot be met immediately by the retailer is backlogged and each unit charged with backorder cost p per period. All backorders are filled in the end.

In each period t, the following sequence of events occurs:

1. Shipments arrive at each stage.
2. Possible backorders are filled by the retailer if he has enough stock available.
3. Replenishment orders are placed at both stages.
4. Orders are shipped and will arrive in period $t + L_{1/2}$.
5. Consumer demand occurs and is obeyed if the retailer has enough stock available. Unmet demand is backlogged.
6. Inventory holding costs and backorder costs are charged.

Before demand, we define for Stage i: echelon inventory level, IL_{it}, is all inventory at stage i or lower in the system minus consumer backorders; echelon inventory position, IP_{it}, is IL_{it} plus inventory currently in transit to stage i. D^t, total demand over t periods is a random variable with density and distribution $f^t(x)$ and $\Phi^t(x)$, respectively and mean value m^t. We assume $\Phi^1(x)$ is continuous and increasing for $x \geq 0$ and $\Phi^1(0) = 0$. Mathematical notation is as follows: $[x]^+ = \max(0, x)$, $[x]^- = \max(0, -x)$ and $E[x]$ is the expected value of x.

2.1 COLLECTIVE OPTIMALITY

Our objective is to identify ordering policies that minimize overall costs for systems of the above type. This involves – for the case of two stages and a finite time horizon – recursive solution of a two-dimensional functional equation. [3], [4] and [5] prove that this problem can be decomposed in two one-dimensional optimization problems which can be solved sequentially if an infinite time horizon is considered and total costs are – virtually – partitioned in a particular way between the retailer and the supplier. An additional result is that among periodic-review policies, an echelon base stock policy minimizes total average costs. In a an echelon base stock setting, each stage i orders – in each period – the amount required to raise his echelon inventory position to a specific echelon base stock level s_i.

We define $\hat{G}_1^0(IL_{1t} - D^1)$ as the retailer's costs in period t, where

$$\hat{G}_1^0(x) = h_1[x]^+ + (h_2 + p)[x]^- . \tag{1}$$

In the same period, the retailer's expected costs for period $t + L_1$, $G_1^0(IP_{1t})$, are

$$G_1^0(y) = E\left[\hat{G}_1^0\left(y - D^{L_1+1}\right)\right]. \tag{2}$$

The value of y for which $G_1^0(y)$ is minimal, is the retailers optimum echelon base stock level s_1^*. It is found by inverting

$$\Phi^{L_1+1}\left(s_1^*\right) = \frac{h_2 + p}{h_1 + h_2 + p} . \tag{3}$$

Next, we define the 'induced penalty function' $\underline{G}_1^0(y)$ that measures the costs arising at the retailer as a result of the supplier's inability to meet his demand

$$\underline{G}_1^0(y) = G_1^0\left(\min\left(s_1^*, y\right)\right) - G_1^0\left(s_1^*\right). \tag{4}$$

This penalty is charged to the supplier. So his total costs in period t amount to

$$\hat{G}_2^0(y) = h_2\left(y - D^1\right) + \underline{G}_1^0(y). \tag{5}$$

The supplier's expected costs in period t are $G_2^0(IP_{2t})$, where

$$G_2^0(y) = E\left[\hat{G}_2^0\left(y - D^{L_2}\right)\right]. \tag{6}$$

The supplier's optimum base stock level is s_2^*, the value of y that minimizes $G_2^0(y)$.

2.2 Individual Optimality

In the following, we give the costs actually incurred at both stages. The retailer is charged $h_1 + h_2$ for each unit stocked and p for each unit backordered in each period. To allow more flexibility, we introduce a backorder cost partitioning between the retailer and the supplier. The coefficient a and $1 - a$ describe the retailer's and the suppliers share, respectively. We define $\hat{G}_1(IL_{1t} - D^1)$ as the retailer's costs in period t,

$$\hat{G}_1(y) = (h_1 + h_2)[y]^+ + a\,p[y]^- . \tag{7}$$

The retailer's expected costs in period $t + L_1$, $G_1(IP_{1t})$, are

$$\begin{aligned} G_1(y) &= E\left[\hat{G}_1\left(y - D^{L_1+1}\right)\right] \\ &= (h_1 + h_2)\left(y - m^{L_1+1}\right) + (h_1 + h_2 + a\,p)\int_y^\infty (x - y) f^{L_1+1}(x)\,dx \end{aligned} . \tag{8}$$

The retailer's cost not only depends on his own echelon base stock level, s_1, but also on the supplier's, s_2. If s_2 is low, there is a chance that the supplier cannot completely fill the retailer's order. Therefore, the retailer's cost, as a function of the strategy parameters are

$$\begin{aligned} H_1(s_1, s_2) &= E\left[G_1\left(\min\left(s_2 - D^{L_2}, s_1\right)\right)\right] \\ &= \Phi^{L_2}(s_2 - s_1)G_1(s_1) + \int_{s_1 - s_2}^\infty G_1(s_2 - x) f^{L_2}(x)\,dx \end{aligned} . \tag{9}$$

The supplier's period t backorder costs, $\hat{G}_2(IL_{1t} - D^1)$, are

$$\hat{G}_2(y) = (1 - a)\,p[y]^- \tag{10}$$

and his expected backorder costs in period $t + L_1$, $G_2(IP_{1t})$

$$G_2(y) = E\left[\hat{G}_2\left(y - D^{L_1+1}\right)\right]. \tag{11}$$

After inventory holding costs are accounted for, the supplier's expected costs can also be written as a function of s_1 and s_2

$$\begin{aligned} H_2(s_1,s_2) &= h_2 m^{L_1} + h_2 \int_0^{s_2-s_1} (s_2 - s_1 - x) f^{L_2}(x)\,dx + \Phi^{L_2}(s_2-s_1)\,G_2(s_1) \\ &= + \int_{s_2-s_1}^{\infty} G_2(s_2-x) f^{L_2}(x)\,dx \end{aligned} . \qquad (12)$$

As (9) and (12) show, costs for Stage i, $i \in \{1,2\}$, $H_i(s_1,s_2)$, depend on Stage i's own strategic decision as well as that of the other Stage j, $j \neq i$, i.e. on s_1 and s_2. In order for the retailer to react optimally on the supplier's choice s_2, he has to choose the s_1 that minimizes his costs given the supplier's s_2. We denote Stage i's best reply mapping by $r_i(s_j)$,

$$\begin{aligned} r_1(s_2) &= \left\{ s_1 \in \mathbb{R}^+ \,\middle|\, H_1(s_1,s_2) = \min_{x \in \mathbb{R}^+} H_1(x,s_2) \right\} \\ r_2(s_1) &= \left\{ s_2 \in \mathbb{R}^+ \,\middle|\, H_2(s_1,s_2) = \min_{x \in \mathbb{R}^+} H_2(s_1,x) \right\} \end{aligned} . \qquad (13)$$

At any point (s_1^{**}, s_2^{**}) with $s_1^{**} \in r_1(s_2^{**})$ and $s_2^{**} \in r_2(s_1^{**})$, each stage chooses a best reply to the other stage's optimum base stock level. Such points are called Nash equilibria and are self-inforcing in the sense that none of the stages has an incentive to deviate unilaterally from it.

2.3 COMPARISON OF INDIVIDUAL AND COLLECTIVE OPTIMA

We use an example to illustrate some important properties of individual and collective optima. Parameters are given in Tab. 1. Further, we assume one-period demand to be normally distributed with mean 1 and standard deviation ¼. This implies that only a fraction of about $3 \cdot 10^{-5}$ of the total demand is <0. So the above assumption that demand is positive holds reasonably well.

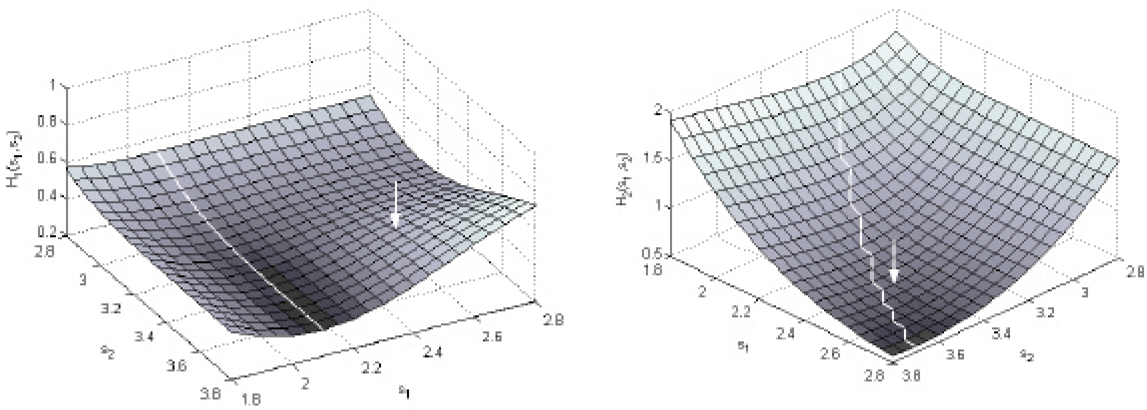

Figure 1. The retailer's (left) and supplier's (right) costs as a function of s_1 and s_2. The white curves in both figures represent best reply curves for both stages. For the retailers' best reply, the s_2-axis is the abscissa, so irrespective of the supplier's choice he chooses $s_1 = 2.1$ as his optimum base stock level. As for the supplier, his abscissa is the s_1-axis. His optimum base stock level increases with that of the retailer. There is only one intersection point of the two reply curves, $(s_1^{**}, s_2^{**}) = (2.1, 3.3)$ so this is the only Nash equilibrium. The collective optimum (white arrows) is far from the retailers best reply curve but relatively close to that of the retailer.

Collectively optimum base stock levels s_1^* and s_2^* are obtained from (3) and (6) as 2.5 and 3.5, respectively. Fig. 1 shows costs for both stages as calculated with (9) and (12) along with best reply mappings. The only individual optimum (Nash equilibrium), (s_1^{**}, s_2^{**}), is the intersection point of the two reply curves at (2.1,3.3). It is obvious that there is a discrepancy between both equilibria. The

Parameter	Value
a	0.3
p	5
h_1	0.5
h_2	0.5
L_1	1
L_2	1

Table 1. Parameters for the example considered

average per-period costs amount to 1.13 and 1.36 at the collective and the individual optimum, respectively, so the difference is of the order of 20%. Note that in our example, the supplier is charged 70% of the system backorder costs. With $a = 1$ the retailer would have to pay all backorder costs and there would be no incentive for the supplier to hold any stock. Moreover, s_1 and s_2 both refer to echelon inventory levels. Therefore, if $s_2 < s_1$ the supplier chooses an echelon inventory level that implies that the sum of the amount of material in his local stock and in transit to it be negative. Thus, we require that $s_2 \geq s_1$.

3. MANAGING THE SYSTEM

In this section we propose a method to manage the system described above. The method consists of two parts. First, the collective optimum described above is made incentive compatible via transfer-payments between the two stages. Second, both the supplier and the retailer have to rely on observed demand data to choose their base optimum base stocks. We describe an approach to dynamically estimate the two strategy parameters from potentially nonstationary demand data.

3.1 TRANSFER PAYMENTS

To achieve optimum systemwide performance, the retailer and supplier have to choose s_1^* and s_2^* as their echelon base stock levels. Thus, they could contract to use (s_1^*, s_1^*) in order to achieve optimum performance. The problem with such a contract is that both sides have an incentive to deviate from this optimum because it is not a Nash equilibrium. Therefore the contract would also have to specify penalties for deviations which clearly is not beneficial to developing the level of trust required for successful supply chain partnerships. As an alternative approach, the supplier and the retailer could agree on a transfer payment $T(s_1, s_2)$ from the supplier to the retailer which eliminates incentives to diverge from the optimal solution. Formally, we can describe this payment as

$$\begin{aligned} H_1^c(s_1,s_2) &= H_1(s_1,s_2) - T(s_1,s_2) \\ H_2^c(s_1,s_2) &= H_2(s_1,s_2) + T(s_1,s_2) \end{aligned}, \quad (14)$$

where $H_i^c(s_1, s_2)$ is Stage i's costs after the payment. [1] propose to parameterizes the payment as a linear combination of Stage 1's on-hand stock I_{1t} and both stages' backorders, B_{it}, all accounted for at the end of period t,

$$i_1 I_{1t} + b_1 B_{1t} + b_2 B_{2t}. \quad (15)$$

We define $T_1(IP_1)$ as the expected transfer payment at the end of period $t + L_1$ as

$$\begin{aligned} T_1(y) &= E\left[i_1\left[y - D^{L_1+1}\right]^+\right] + E\left[b_1\left[y - D^{L_1+1}\right]^-\right] \\ &= i_1\left(y - m^{L_1+1}\right) + \left(i_1 + b_1\right)\int_y^\infty (x-y) f^{L_1+1}(x)\,dx \end{aligned} \quad (16)$$

and $T(s_1, s_2)$ as the expected payment from the supplier to the retailer,

$$\begin{aligned} T(s_1,s_2) &= E\left[b_2\left[s_2 - s_1 - D^{L_2}\right]^-\right] + T_1\left(s_1 + \min\left(0, s_2 - s_1 - D^{L_2}\right)\right) \\ &= b_2 \int_{s_2-s_1}^\infty (x - s_2 + s_1) f^{L_2}(x)\,dx + \Phi^{L_2}(s_2 - s_1) T_1(s_1) \\ &\quad + \int_{s_2-s_1}^\infty T_1(s_2 - x) f^{L_2}(x)\,dx \end{aligned}. \quad (17)$$

We have to determine the parameters i_1, b_1 and b_2 such that the collective optimum (s_1^*, s_2^*) is a Nash equilibrium for the cost functions $H_i^c(s_1, s_2)$. To do so, we first note that $H_i^c(s_1, s_2)$ is strictly convex in s_i, given that the supplier chooses s_j^*, $j \neq i$ ([1]). Thus, to find minimum costs for Stage i, given s_j, we require the first partial derivative of $H_i^c(s_1, s_2)$ with respect to s_i to vanish and its second partial derivative with respect to s_i to be positive.

Several parameter sets qualify as solutions to the above problem. Among those, we choose $i_1 = (1-a)(h_1 + h_2)$, $b_1 = 0$ and $b_2 = g/(1-g)\, a\, h_2$, where $g = \Phi^{L_2}(s_2^* - s_1^*)$. Fig. 2 shows that transfer payments based on these parameters do actually eliminate incentives to deviate from the collective optimum for both stages.

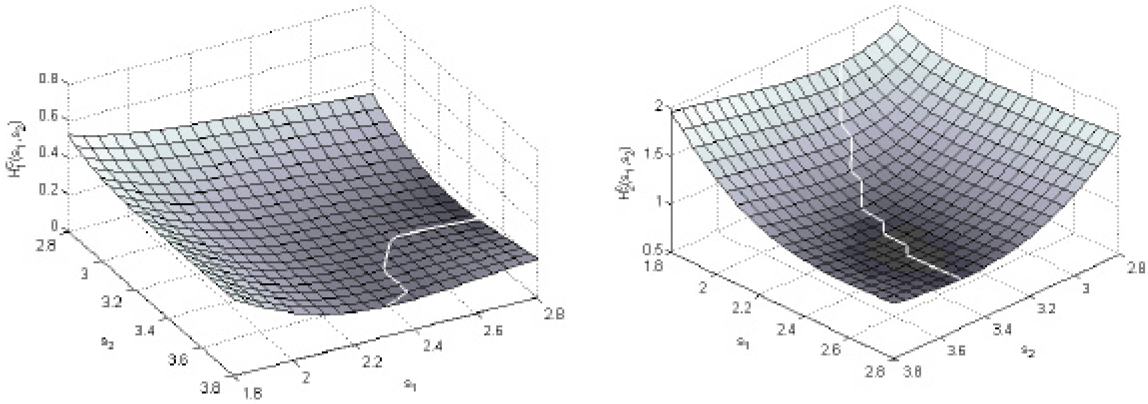

Figure 2. The retailer's (left) and supplier's (right) new cost functions after accounting for transfer payments. The white curves in both figures represent best reply curves for both stages. Note that the only intersection point of the curves is now at the collective optimum ($(s_1^*, s_2^*) = (2.5, 3.5)$). Thus, the transfer payments eliminate both stages' incentives to depart from the collective optimum.

3.2 DYNAMIC ESTIMATION OF s_1^* AND s_2^*

We assume that the supply chain is subjected to the some potentially time-dependent demand process of which the actual parameters are unknown and have to be estimated from observed demand data. Now consider the situation for the supplier and the retailer in period t. The data at hand are the individual orders observed in the past, i.e., tuples $(t_i, d_i), t_i \leq t$, where t_i and d_i denote the time of occurrence and the size of order i, respectively. To minimize the operating costs of the supply chain, the probability distribution of demand and hence s_1^* and s_2^* have to be estimated for period $t+1$. We assume that the distribution of aggregated demand is estimated for each period t' up to t from the set of tuples with $t' \leq t_i < t'+1$. Exponential smoothing is a simple and versatile method to obtain one-period forecasts from these estimates. If $y(t)$ is the variable for which a forecast is required,

$$\begin{aligned} \tilde{y}(t) &= q\, y(t) + (1-q)\, \tilde{y}(t-1) \\ \hat{y}(t+1) &= \tilde{y}(t) \end{aligned}, \qquad (18)$$

where $\tilde{y}$ and $\hat{y}$ are smoothed and estimated values of y, respectively. The smoothing constant q determines to what extent past values are considered in the forecast. For $q = 1$, past values are not at all accounted for, which corresponds to a random walk model without drift, whereas for $q = 0$, (18) becomes a mere mean model.

3.3 NUMERICAL STUDY

We conduct a number of simulations to analyze the costs actually incurred when the supply chain is operated with the methods described above. Assume that actual demands correspond to a lognormal compound Poisson process, i.e. the intervals between two consecutive orders are exponentially distributed and the individual amounts ordered follow a lognormal distribution. Under this assumption, the demand process can be described completely by the Poisson rate l and the parameters of the lognormal distribution $m_{\log}$ and $s_{\log}$. Although costs and optimum values for s_1 and s_2 can be calculated for any distribution of demand, we use demand processes that are not too sporadic. Under this condition, i.e., for $lT >> 1$, [6] show that period-aggregated demand D^1 becomes asymptotically normally distributed with

$$E\left[D^1\right] = \exp\left(m_{\log} + \frac{1}{2} s_{\log}^2\right) lT \tag{19}$$

and

$$Var\left[D^1\right] = \left[E\left[D^1\right] \exp\left(\frac{1}{2} s_{\log}^2\right) \sqrt{\frac{1}{lT}}\right]^2 . \tag{20}$$

During the simulations, l, $m_{\log}$ and $s_{\log}$ are estimated from the individual orders observed during each period t and then used to obtain one-period forecasts $\hat{l}(t+1)$, $\hat{m}_{\log}(t+1)$ and $\hat{s}_{\log}(t+1)$. From these estimates, the mean and variance of period $t+1$ demand are estimated by means of (19) and (20).

Parameter	1. Experiment	2. Experiment Before Step	2. Experiment After Step
l	20	20	20
$m_{\log}$	-3.1073	-3.1073	-2.4142
$s_{\log}$	0.4724	0.4724	0.4724

Table 2. Parameters sets used for the simulations. The step change in the second experiment occurs after 50 time units.

We conduct two numerical experiments. In the first study, the supply chain is subjected to stationary demand, i.e. a demand process with constant parameters l, $m_{\log}$ and $s_{\log}$, whereas in the second, the demand process takes a step change from one parameter set to another at $t = 50$ (Tab. 2). The mean and standard deviation of demand represented by the first parameter set (1. Experiment, 2. Experiment before the step) are 1 and ¼, respectively. For the second parameter set (2. Experiment after the step), the mean and standard deviation amount to 2 and ½, respectively. All other parameter values are given in Tab. 1. In the first experiment, minimum costs would be incurred if the true stationary demand process were known and constant values of s_1^* and s_2^* were chosen such as to cope with the true process. In practice, actual costs will be higher, as the true demand process is not know but has to be estimated from an observed realization. In the second experiment, there are two distinct time periods with stationary demand, respectively. So there are still costs due to imperfect knowledge of the process. However, in the analysis of this experiment, we will merely concentrate on estimating the costs involved in adapting the supply chain to the new level of demand that results from the step change.

In the first experiment, the forecasted mean and variance of demand are highly dependent on the smoothing parameter q. This parameter was varied between the simulations to estimate the cost

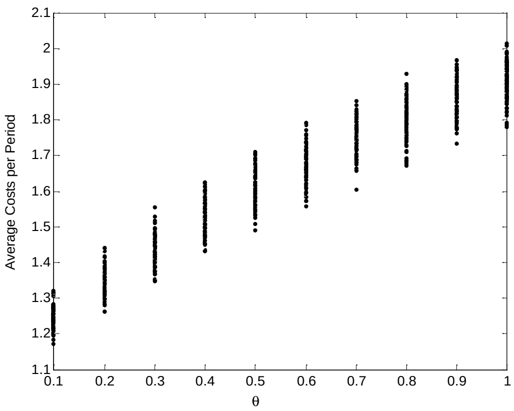

Figure 3. Average per period costs as a function of the smoothing constant. Each dot represents the average costs calculated from one simulation over 500 time steps for a given q.

effects involved. 100 simulations were conducted for each value of q and each simulation covered 500 time units. Fig. 3 shows that average per-period costs increase with q . This can be explained by the fact that with high values of the smoothing parameter the system becomes too reactive to high-frequency fluctuations in demand which leads to inventory levels that are either too low or too high. With lower values of q , the system becomes less and less reactive to high-frequency fluctuations. In the limit case $q = 0$ (not shown), the system does no longer react to ‘new’ data at all but chooses s_1^* and s_2^* such as to match some demand assumed on an a-priori basis or the first value observed. Note that the average costs shown in Fig. 3 converge to the 1.13 calculated above for (s_1^*, s_2^*) .

Fig. 4 shows the results of a single simulation within the second experiment. As a result of the step change from one parameter set to the other at $t = 50$, average period-aggregated demand changes from 1 to 2 units at the same relative standard deviation (25%). The smoothed $m_{\log}$ takes about 50 time units to cope with the new demand level as a result of the low value of the smoothing constant ($q = 0.1$). This phenomenon is well reflected in the dynamics of the optimum echelon base stock levels and the operating costs of a system. Costs are constantly high between $t = 50$ and $t = 100$ and decrease subsequently to fluctuate around their new mean value.

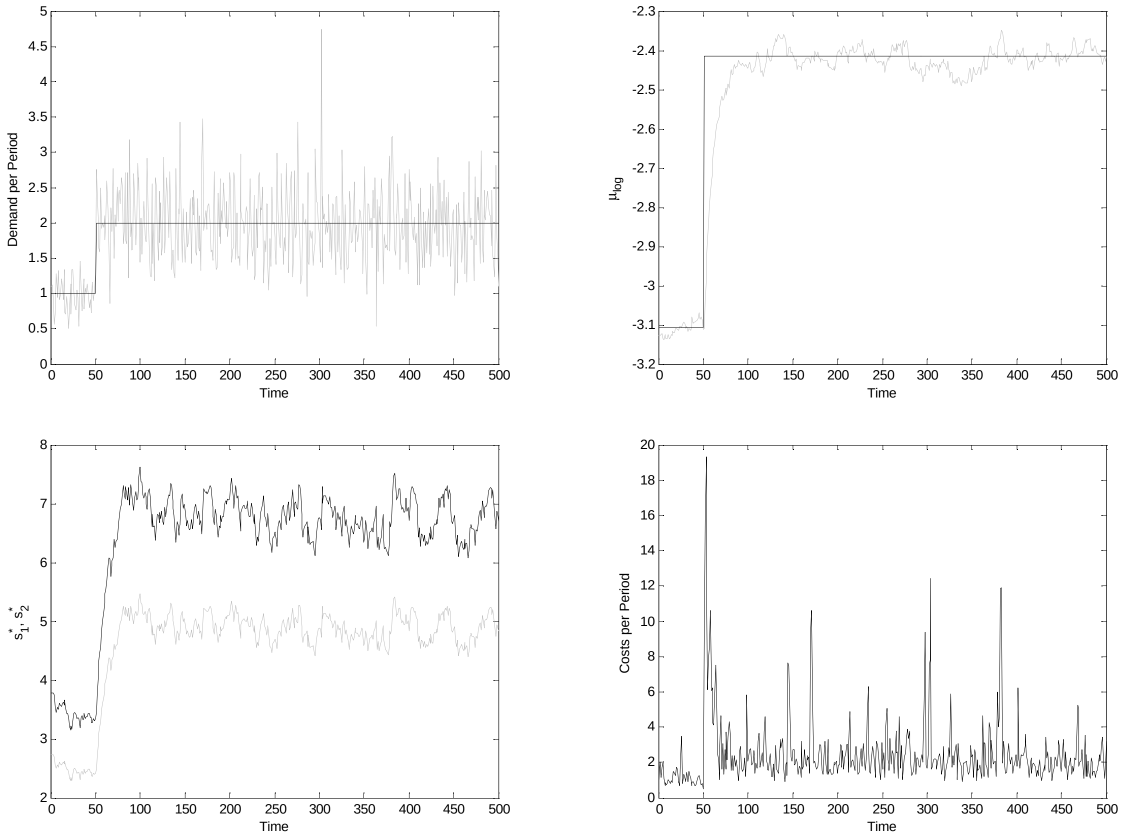

Figure 4. Dynamics of the key variables in a single simulation, $q = 0.1$. Above left: actual (grey) and average theoretical (black) per-period demand. Above right: the smoothed $m_{\log}$ (grey) lags about 50 time units behind its true value (black). Below left: s_1^* (gray) and s_2^* (black) both increase to cope with the new demand but also lag behind their optimum values by about 50 time units. Below right: Too low values for s_1^* and s_2^* (below left) lead to very high costs between $t = 50$ and $t = 100$.

As in the first experiment, we consider the cost effects of various values of q . As mentioned above, we are now interested in estimating the costs involved in reacting to a systematic change in demand. To estimate these, we first estimate average per-period costs after the system has fully adapted to the new demand level and subtract these from the costs actually observed in the simulation after the step change. These are the costs caused by the step change. These costs typically have high values at the beginning ($t = 50$) and subsequently decrease to fluctuate around zero. To account for the total of these costs, we cumulate them up to the end of the simulations ($t = 500$). Fig. 5 shows that for $q \le 0.7$ the total adaptation costs decrease with increasing q. This can be explained by the fact that with low values for q, our base stock levels lag behind their optimum values. This results in too low inventory levels and high backorder penalties for a certain time after the step change.

From the analysis of both experiments, we can conclude that with high values of q, we are overly reactive to short-term demand fluctuations whereas with low q, we are at risk of laging too far behind any systematic changes of demand. So, the choice of the optimum level of smoothing must be made such as to avoid both problems as far as possible. This can only be based on data or a-priori assumptions about the demand process.

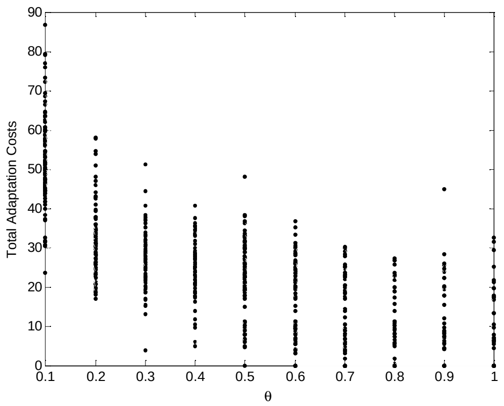

Figure 5. Total costs involved in the adaptation of the supply chain to a new demand level as a function of the smoothing constant. Each dot represents total costs calculated from a simulation over 500 time steps for a given q.

4. DISCUSSION AND CONCLUSIONS

We propose a method to optimally manage a supply chain. The method relies in the first place on a one-period forecast of the distribution of demand. We use exponential smoothing of the parameters of our process but smoothing the parameters of period-aggregated demand would work as well. On the basis of predicted demand, transfer payments can be determined that compensate both the retailer and the supplier for their costs incurred in not choosing individually optimum base stock levels but a collective optimum instead. This eliminates incentives to deviate from collective optimality and ensures that overall costs are actually minimized.

To implement our method in a supply chain, an essential prerequisite is that point-of-sale information is available at both stages. Moreover, both stages have to agree to the cost accounting schemes involved and be ready to disclose their logistics cost rates. Especially the last point may be critical. Our work shows that cooperation in a supply chain may reduce costs substantially. This is a strong argument towards more cooperation in a supply chain.

REFERENCES

[1] G.P. Cachon and P.H. Zipkin (1999): Competitive and Cooperative Inventory Policies. Management Science, 45, 936-953.

[2] A. Waddell (2005): Individuell und kollektiv optimale Lagerstrategien in einer Supply Chain. Projektarbeit im Rahmen des Studiengangs ‚Datenanalyse und Prozessdesign’, Zürcher Hochschule Winterthur.

[3] A. Clark and H. Scarf (1960): Optimal Policies for a Multi-Echelon Inventory Problem. Management Science, 6, 475-490.

[4] A. Federgruen and P. Zipkin (1984): Computational Issues in an Infinite-Horizon, Multi-Echelon Inventory Model. Operations Research, 32, 818-836.

[5] F. Chen and Y. Zheng (1994): Lower Bounds for Multi-Echelon Stochastic Inventory Systems. Management Science, 40, 1426-1443.

[6] J. Hosang, Chr. Heitz and S. Mönkeberg (2002): Demand Forecasting for Inventory Control and Production Planning. Proceedings of the IPLnet 2002 Workshop, September 10-11, 2002, Saas-Fee, Switzerland, 5.3, 1-10.