UC Berkeley
Research Reports

Title

Counteracting the Bullwhip Effect with Decentralized Negotiations and Advance Demand Information

Permalink
https://escholarship.org/uc/item/0489z4qg

Authors
Ouyang, Yanfeng
Daganzo, Carlos

Publication Date
2005-10-01

eScholarship.org Powered by the California Digital Library
University of California

Institute of Transportation Studies
University of California at Berkeley

Counteracting the Bullwhip Effect with Decentralized Negotiations and Advance Demand Information

Yanfeng Ouyang, Department of Civil and Environmental Engineering University of Illinois at Urbana-Champaign and Carlos Daganzao

RESEARCH REPORT
UCB-ITS-RR-2005-6

October 2005
ISSN 0192 4095

Counteracting the Bullwhip Effect with Decentralized Negotiations and Advance Demand Information

Yanfeng Ouyang
Department of Civil and Environmental Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61801

Carlos Daganzo
Institute of Transportation Studies and
Department of Civil and Environmental Engineering
University of California
Berkeley, CA 94720

October 7, 2005

Abstract

This paper shows how to reduce the bullwhip effect by introducing advance demand information (ADI) into the ordering schemes of supply chains. It quantifies the potential costs and benefits of ADI, and demonstrates that they are not evenly distributed across the chain. Therefore, market-based strategies to re-distribute wealth without penalizing any supplier are presented. The paper shows that if a centralized operation can eliminate the bullwhip effect and reduce total cost, then some of this reduction can also be achieved with decentralized negotiation schemes. Their performance is evaluated under different modes of probabilistic supplier behavior. For some forms of behavior the optimum is reached. But if suppliers are greedy and impatient the expected gain in wealth is relatively small. This is a case of economic "market failure".

1 Background and Definitions

In the supply chain literature, the term "bullwhip effect" refers to a phenomenon where the fluctuations in order sequence are usually greater upstream than downstream of a chain. This phenomenon is repeatedly shown in industry operations [7, 8, 15, 16], macroeconomic data [1, 2, 10, 11, 17, 20], and simulations such as "beer games" [9, 12, 21]. The bullwhip effect results in huge extra operation costs for suppliers; in some cases reported to be as much as 25% [3, 13, 14]. The rest of this section explains ADI and its effect on, both, the bullwhip effect and the cost of operation; the key ideas are taken from [19].

1.1 Supply chain operation with ADI

Consider a multi-echelon chain with $i = 1, 2, \cdots, I + 1$ suppliers and one final customer (treated as supplier $i = 0$), as shown in Figure 1. Every supplier ($i = 0, 1, 2, \cdots, I$) orders $u_i(t)$ items from its upstream neighbor at discrete times $t = \cdots, -2, -1, 0, 1, 2, \cdots$, and receives the items after a constant lead time $l_i = 0, 1, 2, \cdots$. The conservation equations for the supplier's *inventory position* at time t (cumulative orders placed minus orders received, $x_i(t)$) and for the *in-stock inventory* at time t (cumulative items received minus orders received, $y_i(t)$) are:

$$x_i(t + 1) = x_i(t) + u_i(t) - u_{i-1}(t), \forall i = 1, 2, \cdots, \tag{1}$$

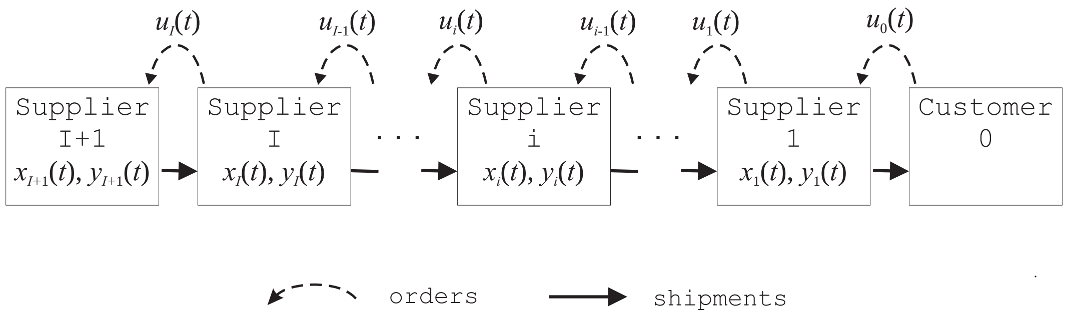

Figure 1: A representation of a supply chain.

and

$$y_i(t+1) = y_i(t) + u_i(t - l_i) - u_{i-1}(t), \forall i = 1, 2, \cdots. \tag{2}$$

The supply chain is decentralized if suppliers act independently, placing orders based on private information; i.e., on all the inventory records and the histories of orders received and placed. Since it turns out that the history of orders placed is redundant information [19], all the information available to supplier i at time t is encapsulated in the following information set:

$$\begin{aligned}\mathcal{I}_i(t) \quad := \quad & \{\, x_i(t), x_i(t-1), \cdots, x_i(-\infty); y_i(t), y_i(t-1), \cdots, y_i(-\infty); \\ & u_{i-1}(t-1), u_{i-1}(t-2), \cdots, u_{i-1}(-\infty)\}.\end{aligned}$$

Therefore, the most general linear and time-invariant (LTI) ordering policy only uses elements of this set as inputs, and can be written as follows:

$$u_i(t) = \gamma_i + A_i(P)x_i(t) + B_i(P)y_i(t) + C_i(P)u_{i-1}(t-1), i = 1, 2, \cdots, \tag{3}$$

where parameter γ_i is a real number, and $A_i(\cdot), B_i(\cdot), C_i(\cdot)$ are polynomials with real coefficients. The symbol P is a backward lag operator; i.e., $P^k x(t) = x(t-k), \forall k = 0, 1, 2, \cdots$. The polynomials $A_i(P)$ and $B_i(P)$ indicate the influence of inventory history, and $C_i(P)$ the history of orders received. We shall abbreviate the policy of supplier i by $\mathcal{F}_i := \{\gamma_i, A_i, B_i, C_i\}, \forall i$.

We can also allow suppliers to incorporate advance demand information (ADI) into their policies [4, 5, 6, 18, 19]. We say that ADI is available to supplier $i+1$ if its downstream neighbor (i) announces at every t the orders it will place in $\bar{h}_{i,i+1} = 0, 1, \cdots$ future periods, and commits to these orders with a contract. We call $\bar{h}_{i,i+1}$ the commitment index between i and $i+1$, and the difference $h_i := \bar{h}_{i-1,i} - \bar{h}_{i,i+1}, i = 1, \cdots, I$ the ADI level of i; see Figure 2. [1]

If the ADI level of a supplier is positive ($h_i > 0$), the supplier (i) is a net consumer of ADI and can incorporate extra information into its policies. Its order $u_i(t)$ is determined and committed at time $t - \bar{h}_{i,i+1}$ based on the following information set:

$$\mathcal{I}'_i(t - \bar{h}_{i,i+1}) := \mathcal{I}_i(t) \cup \{\, u_{i-1}(t), u_{i-1}(t+1), \cdots, u_{i-1}(t + h_i - 1)\}, \text{ if } h_i > 0.$$

A general LTI policy based on this information set has the following form [19]:

$$u_i(t) = \gamma_i + A_i(P)x_i(t) + B_i(P)y_i(t) + C_i(P)u_{i-1}(t + h_i - 1), \; \forall i, t, \text{ if } h_i > 0. \tag{4}$$

Conversely, when $h_i < 0$, supplier i is a net provider of ADI who bears risks and has less flexibility. Its order size $u_i(t)$ is again determined at time $t - \bar{h}_{i,i+1}$ based on the following information set:

$$\mathcal{I}''_i(t - \bar{h}_{i,i+1}) := \mathcal{I}_i(t + h_i) \cup \{\, u_i(t + h_i), u_i(t + h_i + 1), \cdots, u_i(t-1)\}, \text{ if } h_i < 0.$$

[1] We assume that the customer and supplier $I+1$ do not make commitments; i.e. we let $\bar{h}_{0,1} = \bar{h}_{I+1,I+2} = 0$.

A general LTI policy based on this information set is:

$$\begin{aligned} u_i(t) &= \hat{\gamma}_i + A_i(P)x_i(t+h_i) + B_i(P)y_i(t+h_i) + C_i(P)u_{i-1}(t+h_i-1) \\ & +D_i(P)[A_i(P) + B_i(P)P^{l_i}]u_i(t-1), \ \forall i, t, \ \text{if } h_i < 0, \end{aligned} \tag{5}$$

where $D_i(P) = (1 + P + P^2 + \cdots + P^{|h_i|-1})$ and $\hat{\gamma}_i = \gamma + h_i[A_i(1) + B_i(1)]u^\infty$.

We have chosen a somewhat awkward way of expressing (5) because then policies (3), (4) and (5) have the same steady states $(u^\infty, x_i^\infty, y_i^\infty)$.[2] Note that the only ADI parameter on the right side of these equations is h_i. Thus, these equations define a suite of ADI policies characterized by $\mathcal{F}_i$, whose members are distinguished by $h_i = 0, > 0, < 0$.

We are interested in determining how different choices of $\{h_i\}$ affect the bullwhip effect. To this end, we shall assume from now on that the system is in a steady state prior to $t = 0$, and the evaluation is done for $t = 0, 1, \cdots, \infty$.

1.2 The Bullwhip effect

Let σ_0 be the root mean square error (RMSE) of the customer order sequence:

$$\sigma_0 = \lim_{T\to\infty} \left[\frac{1}{T} \sum_{t=0}^{T} (u_0(t) - u^\infty)^2 \right]^{\frac{1}{2}}.$$

The policies of a supply chain uniquely transform this input sequence into a set of order sequences $\{u_i(t)\}$ for each supplier. Because an analyst does not know or control $\{u_0(t)\}$, we define σ_i as the largest possible RMSE of $\{u_i(t)\}$ for all possible $\{u_0(t)\}$ with a given σ_0; i.e.,

$$\sigma_i := \sup_{\forall \{u_0(t)\} \text{ with RMSE } \sigma_0} \left\{ \lim_{T\to\infty} \left[\frac{1}{T} \sum_{t=0}^{T} (u_i(t) - u^\infty)^2 \right]^{\frac{1}{2}} \right\}. \tag{6}$$

We say that the bullwhip effect exists at the upstream end of the chain if:

$$W_I := \frac{\sigma_I}{\sigma_0} = \prod_{i=1}^{I} \frac{\sigma_i}{\sigma_{i-1}} > 1. \tag{7}$$

Recent studies [4, 5, 6, 18, 19] reveal that the amplification factor for stage i, $\mathcal{A}_i := \frac{\sigma_i}{\sigma_{i-1}}$, is given by a formula, $\mathcal{A}_i = \mathcal{A}(\mathcal{F}_i, h_i, l_i)$. A general expression for $\mathcal{A}$ is given in [19]. This formula establishes a quantitative relationship between the ADI level of a supplier h_i and its RMSE amplification (σ_i/σ_{i-1}), for a given policy family. It is shown in [18, 19] that if $\mathcal{F}_i$ amplifies RMSE without ADI (i.e., $\mathcal{A}(\mathcal{F}_i, 0, l_i) > 1$), then there is an $h_i > 0$ that reduces the amplification. For many policies ADI completely eliminates it. It can also be shown that for all $h_i < 0$, $\mathcal{A}(\mathcal{F}_i, h_i, l_i) = \mathcal{A}(\mathcal{F}_i, 0, l_i)$; see Appendix A.

1.3 Supplier costs

A supplier's operating costs (including production, storage, handling, inventory, and transportation) depend on the character of the supplier's input and its ordering policy. In our analysis, we shall hold constant the demand trend, σ_0, and $\mathcal{F}_i$, and only focus on the interplay between the h_i and the order variability. Thus, we express the operating costs estimated by supplier i as:

$$\mathcal{C}_i\left(\sigma_{i-1}, h_i\right), \ i = 1, 2, \cdots, I+1. \tag{8}$$

[2] To see this, simply verify that in the steady state (3), (4) and (5) reduce to:

$$u^\infty = \gamma_i + A_i(1)x_i^\infty + B_i(1)y_i^\infty + C_i(1)u^\infty, i = 1, 2, \cdots.$$

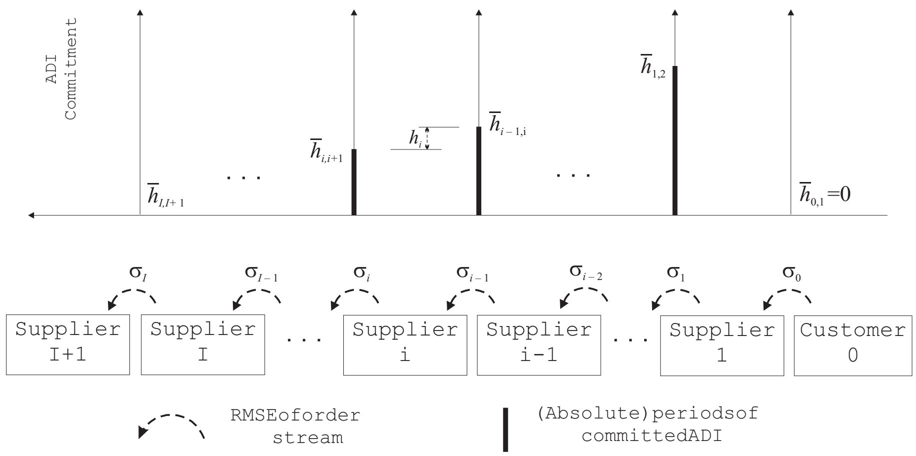

Figure 2: Supplier cost structure and ADI commitment index.

The function $\mathcal{C}_i$ should increase (or at least not decrease) with the variability of the orders received, σ_{i-1}, and decrease (or at least not increase) with the ADI level, h_i.

Properly implemented, ADI would extract commitments from downstream suppliers (imposing a cost) to reduce order variability for upstream suppliers (providing a benefit). Therefore ADI would transfer wealth from the downstream end of the chain to the upstream end. This unbalance needs to be addressed in an ADI implementation.

This paper proposes negotiation and contracting options for the selection of $\bar{h}_{i,i+1}$ that improve cost efficiency while maintaining balance. Section 2 studies an idealized cooperative scheme for the selection, which yields a lower bound for the optimum system cost. Section 3 then proposes a decentralized, non-cooperative market-based procedure and evaluates its performance.

2 Optimum Performance in Cooperative Chains

We compare here the supply chain costs with and without ADI for a given policy, $\mathcal{F}_i$. At stage i, the non-ADI benchmark for comparison is $\mathcal{C}_i(\sigma^0_{i-1}, 0)$. The benchmark cost for the whole chain is $\sum_{i=1}^{I+1} \mathcal{C}_i(\sigma^0_{i-1}, 0)$. Note that the RMSEs in this expression are linked and fixed by the amplification relations $\sigma^0_i = \sigma^0_{i-1} \cdot \mathcal{A}(\mathcal{F}_i, 0, l_i)$, $\forall i$ and $\sigma^0_0 = \sigma_0$. Therefore, the benchmark cost is a constant. Because $\mathcal{F}_i$ and l_i are fixed, we shall write $\mathcal{A}_i(h_i)$ instead of $\mathcal{A}(\mathcal{F}_i, h_i, l_i)$ from now on.

If suppliers cooperate, an agent could find and set the commitment periods at each stage so as to minimize the total operating cost; i.e., as the solution of:

$$\text{Problem (MP):} \quad \min_{\bar{h}_{1,2},\cdots,\bar{h}_{I,I+1},\sigma_1,\cdots,\sigma_I} \sum_{i=1}^{I+1} \mathcal{C}_i\left(\sigma_{i-1}, h_i\right) \tag{9}$$

$$\text{s.t.} \quad \sigma_i = \sigma_{i-1} \cdot \mathcal{A}_i(h_i),\ \forall i = 1, 2, \cdots, I \tag{10}$$

$$h_i = \bar{h}_{i-1,i} - \bar{h}_{i,i+1}, \forall i \tag{11}$$

$$\bar{h}_{i,i+1} \in \{0, 1, 2, \cdots\},\ \forall i \tag{12}$$

The agent does not need to specify the RMSEs because, as we see from (10), knowledge of the commitment levels (h_i's) implies knowledge of the RMSEs (σ_i's). The solution of (9) is the theoretical best one can do with ADI. The benchmark (non-ADI) cost is obtained by setting $h_i = 0, \forall i$ in (9). Since $h_i = 0, \bar{h}_{i,i+1} = 0, \forall i$ is a feasible solution of problem MP, its cost must match or exceed the ADI optimum. Clearly then, ADI can only improve on the benchmark. The magnitude of the improvement generally depends on the cost and amplification functions, $(\mathcal{C}_i, \mathcal{A}_i)$, which vary across industries.

Section 3 will show that much of this improvement can be achieved in a non-cooperative, market-based setting. But before this is done, some asymptotic results are presented for the cooperative case.

Assume that supplier 1 increases its commitment index with 2, $\bar{h}_{1,2}$, allowing 2 to increase its ADI level, h_2, without any change in $h_3, h_4, \cdots$. Consideration of (10) shows that the increase in h_2 reduces σ_2, and this in turn reduces $\sigma_3, \sigma_4, \cdots$ by the same factor. Thus, the single sacrifice made by supplier 1 benefits suppliers $2, 3, \cdots, I+1$. Clearly, the greater I the larger the benefit. And, if the benefit derived by the individual suppliers is bounded from below by a positive constant, total cost would decline for sufficiently large I. This shows that ADI will reduce cost for any realistic $\mathcal{C}_i$ if a chain is sufficiently long. That is why it is so appealing.

This argument should be true even if value-adding activities are distributed across the chain, strongly hinting that ADI policies have much potential for reducing cost. We now introduce a market mechanism that allows all suppliers to capture some of this benefit without a central agent.

3 Non-cooperative Chains: A Market for ADI

So far, our analysis has assumed that suppliers cooperate with a central agent who knows all the supplier cost and amplification functions ($\mathcal{C}_i$ and $\mathcal{A}_i$). A truly powerful agent could enforce cooperation by compensating or penalizing each supplier for the difference between the supplier's new and benchmark operating costs, and keep all the profits. In cases when one party in the chain possesses dominating negotiation power, the chain may operate in ways that resemble this idealization.

In most practical cases, however, suppliers are reluctant to share cost information, except for information that is indirectly released through bilateral negotiations in which transaction prices are set. Free market mechanisms would set these prices somewhere between (i) the maximum price a supplier is willing to pay to its downstream neighbor for a reduction in RMSE or an increase in the commitment index, and (ii) the minimum price this neighbor is willing to take to adjust its ordering procedures (ADI and RMSE commitments). The gap between these two bounds is the potential profit. We will assume that downstream suppliers make offers and upstream suppliers accept or decline them.

3.1 Negotiation rounds

Suppliers i and $i+1$ negotiate the values of state variables $\{\bar{h}_{i,i+1}, \sigma_i\}$, where σ_i now represents an RMSE guarantee provided by supplier i. Of course, if this supplier is smart, the guarantee should equal its worst-case RMSEs prediction (10). This is illustrated in Figure 3. To formalize these ideas we now introduce a function $\texttt{price}_{i,i+1}(\bar{h}_{i,i+1}, \sigma_i)$ to denote the minimum compensation (taken as a discount to the item price) supplier i would be willing to receive from $i+1$ in exchange for providing $\bar{h}_{i,i+1}$ periods of ADI and an RMSE bound of σ_i. This may include a profit mark-up known only to supplier i.

During a negotiation, a supplier proposes a "price function" to its upstream neighbor. The upstream supplier can then either decline the offer and stay with the status quo, or else choose the arguments of the price function and agree to a new contract. The contract only defines the price function. The upstream supplier retains the right to change the arguments of the function with sufficient notice. The detailed mechanism of a bilateral negotiation is as follows.

Bilateral negotiation mechanism

Assume that supplier $i = 1, 2, \cdots$ decides to evaluate its ADI policy options and change $\texttt{price}_{i,i+1}$. If this does not directly follow a recent update of $\texttt{price}_{i-1,i}$ from supplier $i-1$, we say that supplier i has initiated a new negotiation round. For any desired $\texttt{profit}_i$ and given the price function from its downstream neighbor,

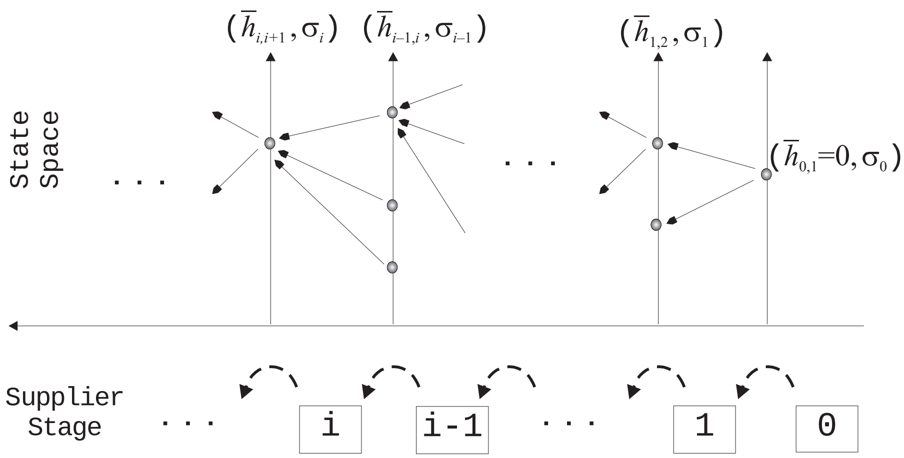

Figure 3: Stage-wise negotiations.

it should be able to calculate:

$$\begin{aligned}
\texttt{price}_{i,i+1}(\bar{h}_{i,i+1}, \sigma_i) \quad &:= \quad \min_{\bar{h}_{i-1,i}, \sigma_{i-1}} \left[\mathcal{C}_i(\sigma_{i-1}, h_i) + \texttt{price}_{i-1,i}(\bar{h}_{i-1,i}, \sigma_{i-1}) \right] + \texttt{profit}_i \qquad (13) \\
\text{s.t.} \quad & \quad \sigma_i \geq \sigma_{i-1} \cdot \mathcal{A}_i\left(h_i\right), \\
& \quad h_i = \bar{h}_{i-1,i} - \bar{h}_{i,i+1}.
\end{aligned}$$

Note that the first constraint is an inequality because supplier i can always try a σ_{i-1} smaller than strictly necessary. If its offer is accepted, supplier i obtains a profit, $\texttt{profit}_i$, regardless of the choice made by $i+1$. Supplier i may have to change its state variables with $i-1$ as a result of this choice, though.

Supplier $i+1$ will accept the offer if there exists arguments:

$$\begin{aligned}
\{\bar{h}^*_{i,i+1}, \sigma^*_i\} \quad &:= \quad \arg \min_{\bar{h}_{i,i+1}, \sigma_i} \left[\mathcal{C}_{i+1}(\sigma_i, h_{i+1}) + \texttt{price}_{i,i+1}(\bar{h}_{i,i+1}, \sigma_i) \right] \qquad (14) \\
\text{s.t.} \quad & \quad \sigma_{i+1} \geq \sigma_i \cdot \mathcal{A}_{i+1}\left(h_{i+1}\right), \\
& \quad h_{i+1} = \bar{h}_{i,i+1} - \bar{h}_{i+1,i+2},
\end{aligned}$$

which increase its profit, $\texttt{profit}_{i+1} = \texttt{price}_{i+1,i+2} - (\mathcal{C}_{i+1} + \texttt{price}_{i,i+1})$, from the current value without changing $\texttt{price}_{i+1,i+2}$; i.e., if $\mathcal{C}_{i+1}(\sigma^*_i, h^*_{i+1}) + \texttt{price}_{i,i+1}(\bar{h}^*_{i,i+1}, \sigma^*_i)$ decreases.

Multilateral negotiation mechanism

Multilateral negotiations happen when supplier $i+1$, upon receiving a price function offer from i, $\texttt{price}_{i,i+1}$, updates its price function offer to supplier $i+2$, $\texttt{price}_{i+1,i+2}$, and waits for its response before replying to i. The updating process uses (13), and may be repeated at suppliers $i+2, i+3, \cdots, j$ until supplier $j+1$ finally accepts or declines the offer, using (14).

If the offer is accepted, supplier $j+1$ then pays $\texttt{price}_{j,j+1}(\bar{h}^*_{j,j+1}, \sigma^*_j)$ to supplier j in exchange for $\bar{h}^*_{j,j+1}$ and σ^*_j. Supplier j will then integrate the choice of $j+1$ (even if the offer is declined) into (14) to make its own choice. This process continues toward the downstream of the chain until reaching supplier i, who initiated the negotiation round. At that point, all the suppliers involved in the multilateral transaction will have revised their contracts and adjusted their state variables (i.e., the commitments).

Initially, the price function has to be consistent with the equilibrium that exists before the introduction of ADI. Therefore, we initialize the system with $h_i = 0$ and price functions

$$\mathtt{price}_{i,i+1}(\bar{h}_{i,i+1}, \sigma_i) = \begin{cases} 0, & \text{if } \sigma_i \geq \sigma_i^0 \text{ and } \bar{h}_{i,i+1} = 0 \\ \infty, & \text{otherwise} \end{cases} .$$

These price functions preclude greater variances than those obtained without ADI. The non-ADI equilibrium is broken when suppliers realize that they can increase profit by changing this function and starting their negotiations.

3.2 Behavioral Scenarios and Analytic Results

In decentralized chains, every supplier may independently initiate a negotiation round, set its own profit level, and choose among accepting, transmitting or declining an offer. Some of these choices (e.g., the selection of a profit level and the decision to transmit an offer) are random, and the probabilistic associated with the choices can affect the course of the negotiations. We now show how much.

Systematic multi-lateral negotiations with negligible profit

First we consider an idealized process where only supplier 1 initiates a round, and all offers are transmitted upstream with probability $p = 1$. We also assume that every supplier charges a negligible profit. Thus, the only beneficiary of the cost reduction is supplier $I + 1$.

It should be obvious that such negotiation exactly solves problem MP in one round. This is because the procedure is just the stage-wise optimization algorithm that solves dynamic-programming-type problems of type (9).

Random negotiations with negligible profit

We still assume that $\mathtt{profit}_i = 0, \forall i$ but now allow all suppliers to initiate negotiations randomly, and upon receiving offers to transmit them upstream with probability $p_i < 1$. Since $\mathtt{profit}_i = 0, \forall i$, suppliers have total flexibility to choose any price at all time. Clearly then, an infinite number of negotiation rounds would have a subsequence that updates $\mathtt{price}_i$ from $i = 1$ to I. This ensures asymptotic convergence — i.e., with probability 1 the optimal solution of (9) is reached in finite time.

We can analytically calculate the number of negotiation rounds needed to achieve the optimum. Consider, for example, a homogeneous chain with $I + 1$ suppliers where $p_i = p$, and every supplier is equally likely to initiate a negotiation round. It can be easily shown that the expected number of negotiation rounds, $\mathrm{E}\, X_I$, can be calculated from the following recursion (see Appendix B):

$$\mathrm{E}\, X_i = \frac{I + (p^i - p^{I+1}) \sum_{j=1}^{i} p^{-j} \mathrm{E}\, X_{j-1}}{1 + (p - p^{I-i+1})/(1-p)}, \forall i = 1, 2, \cdots, I, \text{ with } \mathrm{E}\, X_0 = 0. \tag{15}$$

Greedy suppliers

In reality, suppliers try to achieve some non-negligible profit and use negotiations as a means to increase it. Nobody's profit should decrease after a successful negotiation round. This behavior adds considerable complexity to the process, and influences the possibility of reaching the optimum. We investigate this issue with simulations.

3.3 Simulations

Consider a homogeneous chain with $I + 1 = 7$, $W = 3$, and $h = 2$. Each supplier has cost function $\mathcal{C}(\sigma_{i-1}, h_i) = \frac{\sigma_{i-1}}{2}(2)^{\frac{1}{2}}$ for $h_i > 0$, and $\frac{\sigma_{i-1}}{2}(1 + W^2)^{\frac{1}{2}} \cdot \max(1, 1 - h_i)$ for $h_i \leq 0$. When a supplier enters a new negotiation round, it tries to increase its profit by a random quantity that is uniformly distributed

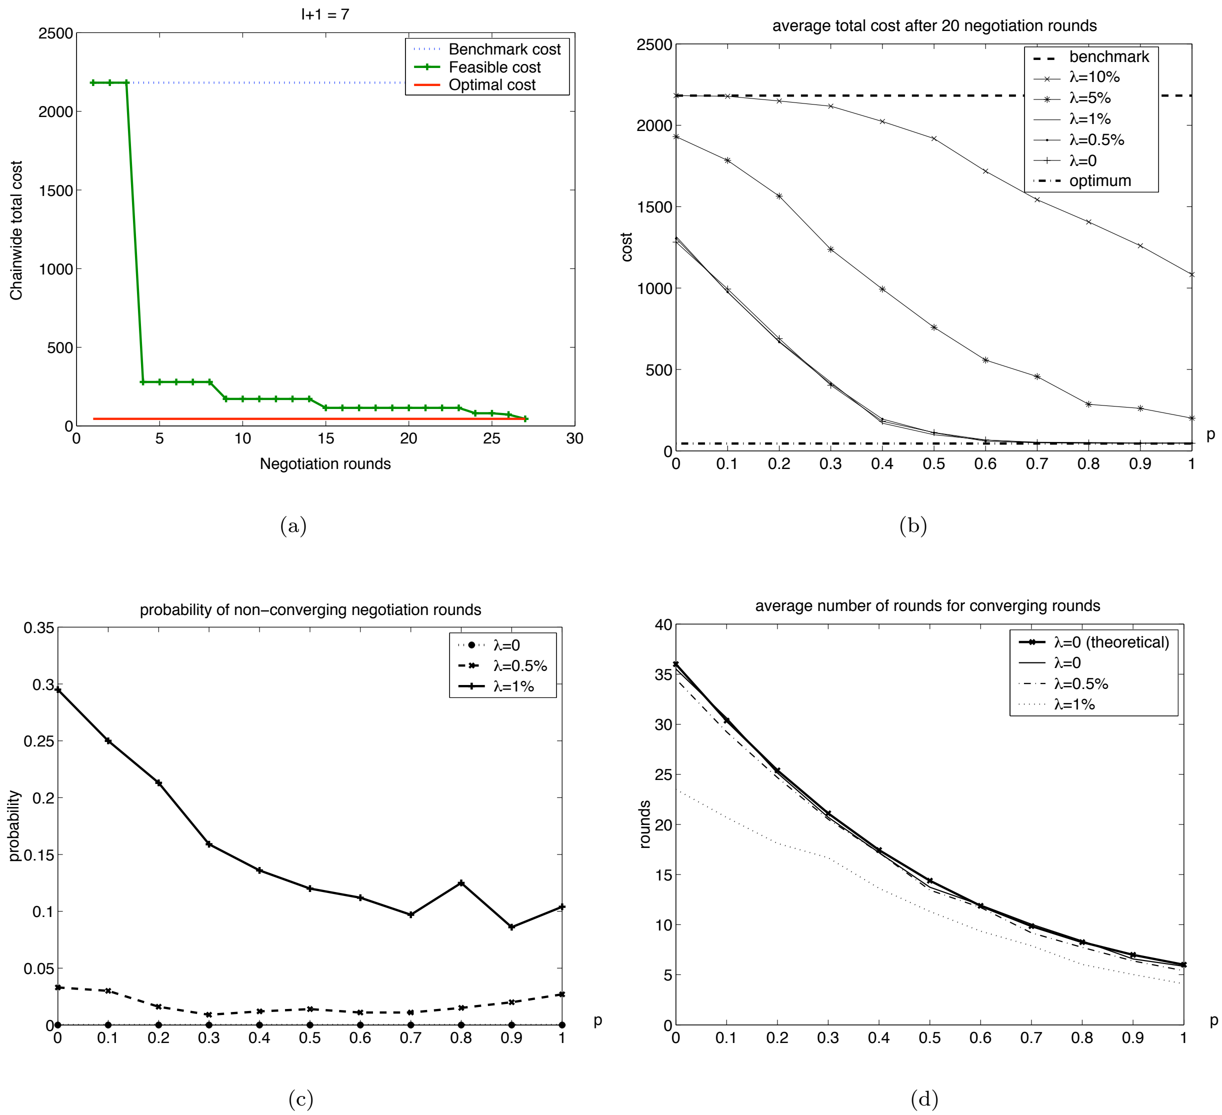

Figure 4: Simulations: (a) A typical realization ($p = 0.5, \texttt{profit}_i = 0, \forall i$); (b) Average cost after 20 negotiations; (c) probability of non-converging rounds; (d) Expected number of rounds among those that converged.

in $[U/2, U]$, where the limit U is set as a small percentage, λ, of a supplier's share of the total possible cost savings; i.e., as a percentage of (benchmark cost − optimal cost)$/(I+1)$. We then simulated many random realizations of negotiation processes, assuming that all suppliers are equally likely to initiate a negotiation round.

Section 3.2 showed that when the suppliers are not greedy ($\lambda = 0$), asymptotic convergence is ensured. Figure 4(a) confirms this result. It shows a typical realization with $p = 0.5$, where the total cost converges to the optimum in 27 rounds. Note the rapid improvement after just a few rounds.

We are interested in how negotiations, especially the first few rounds, reduce total system cost in the (realistic) greedy case. Figure 4(b) shows the average total system cost achieved after 20 rounds (i.e., each supplier on average initiates 3 rounds) for different values of p and λ. We see that for small λ, almost 100% of the cost savings can be achieved if suppliers are enterprising ($p > 0.6$), and that 45% can be achieved even if suppliers are totally passive ($p = 0$). When λ increases (i.e., suppliers try to increase profit more rapidly), the total cost savings decreases. The problem is that if one of the suppliers succeeds in extracting too large a profit early and if this is suboptimal, then the process will fail to get the optimum, since the supplier in question will be unwilling to lower its profit in future negotiation rounds. Therefore, the suppliers' rush for profit (even if moderate) hinders convergence toward optimality and increases cost. This is an interesting example of "market failure".

We also study how parameters λ and p influence the rate of convergence to the optimum. Figure 4(c) confirms that with zero profit ($\lambda = 0$), the optimum is asymptotically achieved. But when λ increases, the convergence process is also more likely to fail. On the other hand, among those rounds that do converge, the average number of rounds for convergence decreases with λ and p. This is shown in Figure 4(d), which also includes theoretical result (15).

4 Conclusion

This paper built on previous findings on the bullwhip effect and ADI [4, 18, 19], by proposing new ordering and contracting schemes that distribute wealth across the chain. It is found that significant benefits can be achieved through a free market mechanism, but greed and impatience reduce the total benefit.

Our result assumed a specific negotiation process, in which downstream suppliers make offers to their upstream neighbors, and the latter accept or decline the offers. But in reality, upstream suppliers can also initiate offers and the real process is likely to be more disorganized. Symmetry considerations, however, suggest that our qualitative result would not change.

This paper also assumed that suppliers follow up with their commitments, and thus does not discuss a mechanism for enforcing compliance. A mechanism could be based on the price function itself, and could include clauses with extra penalties for large deviations. This would require the contracts to be evaluated over extended periods of time, and would preclude rapid changes.

Future research can be extended in several directions. First, our negotiation framework is based on worst-case RMSE bounds. This may be changed to expected RMSEs, which can better match empirical observations. We should also allow suppliers to change $\mathcal{F}_i$ when h_i changes. This paper also assumed that all suppliers know their own cost functions and are able to make optimal decisions in a very structured way. In reality suppliers may not be perfectly competent or have enough information to be so precise. Thus, one may wish to include an adaptive learning mechanism in the game. Maybe convergence can be improved with more frequent negotiations and this learning mechanism. Real experiments with people (similar to the beer game) could be a good way to verify these conjectures.

References

[1] Baganha, M.P., Cohen, M.A. (1998) The stabilizing effect of inventory in supply chains. *Operations Research*, 46(3): 572 - 583.

[2] Blinder, A.S. (1986) Can the production smoothing model of inventory behavior be saved? *Quarterly Journal of Economics*, 101: 431 - 454.

[3] Cooke, J.A. (1993) The $30 Billion Promise. *Traffic Management*, 32, December 1993: 57-59.

[4] Daganzo, C.F. (2001) *A theory of supply chains*. Institute of Transportation Studies Research Report, UCB-ITS-RR-2001-7, University of California, Berkeley, CA.

[5] Daganzo, C.F. (2003) *A Theory of Supply Chains*. Springer, Heidelberg, Germany.

[6] Daganzo, C.F. (2004) On the stability of supply chains. *Operations Research* 52(6): 909-921.

[7] Forrester, J. (1958) Industrial dynamics, a major breakthrough for decision makers. *Harvard Business Review*, July - August, 36: 37 -66.

[8] Forrester, J. (1961) *Industrial Dynamics*. MIT Press, Cambridge MA.

[9] Goodwin, J., Franklin, S. (1994) The beer distribution game: using simulation to teach systems thinking, *Journal of Management Development*, 13(8): 7-15.

[10] Holt, C.C., Modigliani, F., Muth, J. and Simon, H.A. (1960) *Planning Production, Inventories and Work Force*, Prentice Hall, Englewood Cliffs, NY.

[11] Kahn, J. (1987) Inventories and the volatility of production. *American Economic Review*, 77: 667 -679.

[12] Kaminsky, P. and Simchi-Levi, D. (1998) The Computerized Beer Game: Teaching the Value of Integrated Supply Chain Management, *Supply Chain and Technology Management*, (Hau Lee and Shu Ming Ng, eds.), The Production and Operations Management Society.

[13] Lee, H.L., Padmanabhan, V., Whang, S. (1997a) The Bullwhip effect in supply chains. *Sloan Management Review*, Spring 38 (3): 93 - 102.

[14] Lee, H.L., Padmanabhan, V., Whang, S. (1997b) Information Distortion in a Supply Chain: The Bullwhip Effect. *Management Science*, 43(4): 546 -558.

[15] Magee, J.F. (1956) Guides to inventory control (Part II), *Harvard Business Review*, March-April 1956: 106-116.

[16] Magee, J.F. and Boodman, D. (1967) *Production Planning and Inventory Control (2nd edition)*, McGraw-Hill, NY.

[17] Naish, H.F. (1994) Production smoothing in the linear quadratic inventory model. *The Economic Journal*, 104(425): 864-875.

[18] Ouyang, Y. (2005) *System-level Stability and Optimality of Decentralized Supply Chains*, Ph.D. Dissertation, University of California, Berkeley.

[19] Ouyang, Y. and Daganzo C.F. (2005) Some properties of decentralized supply chains; *Institute of Transportation Studies Research Report*, UCB-ITS-RR-2005-3.

[20] Ramey, V.A. (1991) Nonconvex costs and the behavior of inventories. *Journal of Political Economy*, 99: 306-334.

[21] Sterman, J. (1989) Modelling managerial behaviour: Misperceptions of feedback in a dynamic decision making experiment. *Management Science*, 35(3):321 -339.

A The amplification factor of (5) for $h_i < 0$

This appendix shows that policy (5) with $h_i < 0$ yields the same RMSE amplification as (3) with $h_i = 0$. To do this, let us first express policy (5) into a form similar to (3) by redefining the inventory and order variables. Since $D_i(P) := 1 + P + \cdots + P^{-h_i-1}$, the total number of advance orders committed by supplier i at time $t + h_i$ equals $D_i(P)u_i(t-1) = u_{i-1}(t-1) + \cdots + u_{i-1}(t+h_i)$. Then, define the following variables:

$$\begin{aligned}
&u_i'(t) := u_i(t),\ \forall t,\\
&u_{i-1}'(t) := u_{i-1}(t+h_i),\ \forall t,\\
&x_i'(t) := x_i(t+h_i) + D_i(P)u_i(t-1),\ \forall t,\\
&y_i'(t) := y_i(t+h_i) + D_i(P)u_i(t-l_i-1),\ \forall t.
\end{aligned}$$

By substituting these new variables into (5), we find:

$$\begin{aligned}
u_i'(t) = u_i(t) \quad &= \quad \hat{\gamma}_i + A_i(P)[x_i'(t) - D_i(P)u_i(t-1)] \ + B_i(P)[y_i'(t) - D_i(P)u_i(t-l_i-1)]\\
&\quad\ + C_i(P)u_{i-1}'(t-1) + D_i(P)[A_i(P) + B_i(P)P^{l_i}]u_i(t-1)\\
&= \quad \hat{\gamma}_i + A_i(P)x_i'(t) + B_i(P)y_i'(t) + C_i(P)u_{i-1}'(t-1)
\end{aligned} \tag{16}$$

Note that policy (16) is based on an information set that does not include terms from the future:

$$\begin{aligned}
\mathcal{I}_i'''(t) \quad := \quad &\{\, x_i'(t), x_i'(t-1), \cdots, x_i'(-\infty); y_i'(t), y_i'(t-1), \cdots, y_i'(-\infty);\\
&u_{i-1}'(t-1), \cdots, u_{i-1}(-\infty)\}.
\end{aligned}$$

Note too that (17) has the same function form as policy (3), except for a different constant $\hat{\gamma}_i$. Since constants do not influence RMSE, it follows that the worst-case RMSE amplification is the same for both policies.

B Proof for (15)

Let $\mathrm{E}\,X_i$ denote the expected number of negotiation rounds needed to reach convergence when $i = 0, 1, \cdots, I$ upstream suppliers (i.e., $j = I, \cdots, I-i+1$) have not yet found their optimal $\texttt{price}_{j,j+1}$. Obviously $\mathrm{E}\,X_0 = 0$. To prove (15), consider the realization of one additional round, which may be initiated by supplier k and concluded by supplier $s(> k)$. By conditional expectation,

$$\mathrm{E}\,X_i = 1 + \sum_{\substack{k,\,s\\ 1 \le k < s \le I+1}} \left\{ \begin{array}{l} \text{(expected rounds after this realization)}\\ \cdot \Pr\{\text{initiated by } k,\ \text{concluded by } s\} \end{array} \right\}. \tag{17}$$

We examine all possible combinations of k, s and i, and (17) becomes

$$\begin{aligned}
\mathrm{E}\,X_i \quad &= \quad 1 + \sum_{s=1}^{i} \left\{ \mathrm{E}\,X_{s-1} \cdot \sum_{k=i}^{I} [\frac{1}{I} p^{k-s}(1-p)] \right\} + \sum_{s=i+1}^{I} \left\{ \mathrm{E}\,X_i \cdot \sum_{k=s}^{I} \left[\frac{1}{I} p^{k-s}(1-p) \right] \right\}\\
&\quad + \frac{i-1}{I} \mathrm{E}\,X_i + \frac{1}{I} \sum_{k=i}^{I} \mathrm{E}\,X_0\\
&= \quad 1 + \frac{p^i - p^{I+1}}{I} \sum_{s=1}^{i} \left[\mathrm{E}\,X_{s-1} p^{-s} \right] + \left(\frac{I-1}{I} - \frac{p - p^{I-i+1}}{I(1-p)} \right) \mathrm{E}\,X_i.
\end{aligned}$$

Solving $\mathrm{E}\,X_i$ from this equation yields (15).