Upstream versus downstream information and its impact on the bullwhip effect

Rachel Croson[a] and Karen Donohue[b]*

Rachel Croson is Associate Professor of Operations and Information Management at the Wharton School of the University of Pennsylvania. She holds a PhD in Economics from Harvard University. Her research uses methods from experimental and behavioral economics to explore strategic interaction.

Karen Donohue is Associate Professor of Operations and Management Science at the Carlson School of Management at the University of Minnesota. She holds a PhD in Industrial Engineering from Northwestern University. Her research interests include coordination and behavioral issues in supply chain management.

Abstract

This paper reports the results of an experiment to examine whether giving supply chain partners access to downstream inventory information is more effective at reducing bullwhip behavior, and its associated costs, than similar access to upstream inventory information. Bullwhip behavior refers to the tendency of orders to increase in variation as they are passed upstream in a supply chain (i.e., away from the final consumer). We use a controlled version of the Beer Distribution Game as the setting for our experiment, and vary the amount and location of inventory information shared across treatments. We first independently test whether sharing upstream or downstream inventory information helps reduce bullwhip behavior, and find that only downstream information sharing leads to significantly lower order oscillations throughout the supply chain. We then compare the reduction in order oscillations experienced by supply chain level and find that upstream supply chain members benefit the most from downstream information sharing. Copyright © 2005 John Wiley & Sons, Ltd.

Syst. Dyn. Rev. **21**, 249–260, (2005)

Introduction

There is a large and growing literature on the bullwhip effect and its impact on supply chain performance. The bullwhip effect is a magnification of order oscillations as one moves up the supply chain, away from the final customer. This magnification is usually measured in terms of a change in the variance of orders placed at each supply chain level. For example, in a two-level supply chain consisting of a retailer and manufacturer, bullwhip behavior would imply that the variance of orders received by the manufacturer is higher than the variance of demand experienced by the retailer. This behavior is witnessed in a number of industries, ranging from consumer packaged goods to real estate (Sterman 2000). Its existence is attributed to both operational and behavioral factors.

Lee *et al.* (1997) were the first to identify four key operational factors that encourage bullwhip behavior. These factors include (1) fixed costs in production, ordering, or shipping, which encourage order batching, (2) shortage gaming, which encourages phantom orders, (3) price promotions, which encourage forward buying, and (4) errors in demand signaling, which encourage order adjustments. Behavioral factors, primarily attributable to cognitive limitations

[a] Department of Operations and Information Management, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104-6366, U.S.A.
[b] Department of Operations and Management Science, The Carlson School, University of Minnesota, Minneapolis, MN 55455-9940, U.S.A.
* Correspondence to: Karen Donohue. Email: kdonohue@csom.umn.edu

System Dynamics Review Vol. 21, No. 3, (Fall 2005): 249–260
Published online in Wiley InterScience
(www.interscience.wiley.com). DOI: 10.1002/sdr.320

Received April 2004
Accepted June 2005

of the decision makers, include the tendency to not fully account for one's supply line when making ordering decisions (Sterman 1989), and the tendency to mistrust, and thus develop counteracting strategies for, the performance of one's fellow supply chain members (Croson *et al.* 2004).

Information sharing, particularly sharing information on inventory levels, has been cited as a possible countermeasure to the bullwhip effect. From an operational perspective, inventory information can be used to update demand forecasts and lessen the impact of demand-signaling errors and delays. In fact, such information may even be helpful in supply chains where the demand distribution is known to all supply chain members and each member makes ordering decisions based on an order-up-to policy. For example, analytical research on inventory management in two-echelon supply chains with a single supplier and one or more retailers (e.g., Bourland *et al.* 1996; Lee *et al.* 1997; Cachon and Fisher 2000; Gavirneni *et al.* 1999) concludes that sharing inventory information can improve supply chain performance, with the upstream member (i.e., the supplier) enjoying most of the benefits. In these analytical models, inventory information provides the supplier with more timely and less distorted demand signals. These signals are then factored into the supplier's order decisions, resulting in lower safety stock and/or higher service levels compared to cases where no inventory information is shared. This improvement at the supplier level also translates into less need for safety stock at the retail site, although the inventory savings for the retailer turns out in most cases to be less than that of the supplier. In multi-echelon supply chains, inventory information also allows supply chain members to manage orders based on echelon inventory level rather than the order quantity placed by one's immediate customer, which is known to lead to better performance (Chen 1998). This prior research sheds light on how access to inventory information improves operational factors leading to better inventory management. It does not focus on the bullwhip effect *per se* since these models assume rational decision makers that operate according to order-up-to policies (which effectively passes through orders with no amplification).

From a behavioral perspective, inventory information can also provide a means to affect behavior and, as a result, increase trust (or at least understanding) throughout the supply chain. In an experimental setting based on the popular Beer Distribution Game, Croson and Donohue (2004) showed that human decision makers in a four-member, serial supply chain continued to exhibit bullwhip behavior in their ordering patterns even when all the operational causes of the bullwhip were removed. They further found that sharing everyone's inventory information throughout the entire supply chain significantly dampened order oscillations, although it did not eliminate the effect completely. However, consistent with previous analytical research, the benefit of information sharing was more significant for upstream players. This led to the conjecture that the critical part of an inventory-sharing information system is communicating the inventory position of downstream players to upstream players and not vice versa. If true, this implies that an inventory-sharing system may still be effective if upstream firms are reluctant or unable to share their inventory information with downstream firms. It is this conjecture that we test in this paper. This paper thus differs from previous work by disentangling the effects of upstream and downstream information sharing. In Croson and Donohue (2004) we compared no inventory information with full inventory information. Here we compare no inventory information with only upstream inventory information and with only downstream inventory information.

More specifically, we run three treatments to test the conjecture that giving supply chain partners access to downstream inventory information (i.e., inventory levels of one's customers and their customers) is more effective at reducing bullwhip behavior, and its associated costs, than similar access to upstream inventory information. Like Croson and Donohue (2004), we use a controlled version of the Beer Distribution Game as the setting for our experiment. The experiment consists of three separate treatments, varying in the amount and location of inventory information shared across a four-member, serial supply chain. Our results confirm that access to downstream inventory information does significantly reduce order oscillation throughout the supply chain, with the most significant improvement at upstream levels (i.e., the level of distributor and manufacturer). Access to upstream inventory information, on the other hand, provides no significant improvement in order oscillations for the supply chain as a whole or for any particular member. The paper's main contribution is to pinpoint the type of inventory information sharing that is most beneficial in dampening the bullwhip effect.

We continue in the next section with a description of the Beer Game and our experimental design. In the third section, we introduce hypotheses and report on the results of our three treatments. The paper concludes with a discussion of main results and future research in the final section.

Overview of experimental setting and treatments

The rules and mechanics of the Beer Distribution Game are well documented (see Sterman 1992 or Croson and Donohue 2004 for more details). Its rich history began with the work of Forrester (1958). The game consists of four players who take the role of inventory managers at one of four echelons within an integrated supply chain. Figure 1 illustrates the supply chain structure, with players taking on the role of retailer, wholesaler, distributor, and manufacturer. Within each role, the decision maker is responsible for placing orders to his upstream supplier and filling orders placed by his downstream customer. This decision is made repeatedly over a series of periods (referred to as weeks in the game). Within each period, events occur in the following order: (1) shipments arrive from one's upstream suppliers; (2) new orders arrive from downstream customers; (3) new orders are filled and shipped from inventory;

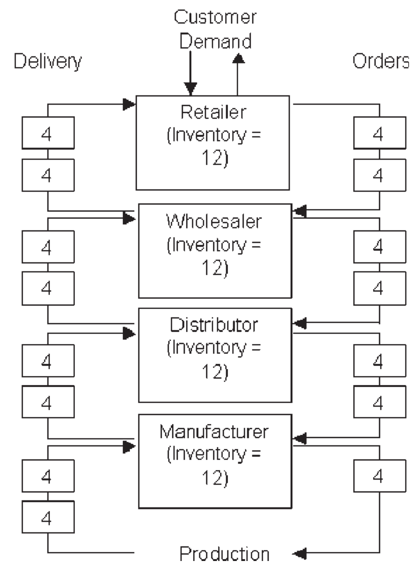

Fig. 1. Initial conditions for the Beer Distribution Game

however, if demand is higher than inventory on-hand, unfilled demand is placed in a backlog and filled once the inventory becomes available in a future period; and (4) each supply chain member places an order to his upstream supplier. This order decision is the single decision variable and the focus of our analysis over the multi-period game.

As shown in Figure 1, ordering and shipment activities within the supply chain are complicated by the presence of delays. Two periods are required to process orders at the wholesaler and distributor levels (e.g., an order placed at the retail level in period 1 will be acted on by the wholesaler in period 3). Similarly, two periods are required to ship orders from the distributor and wholesaler. Finally, the manufacturer experiences a three-period delay between placing a production order and completing this production. These delays complicate ordering decisions by reducing supply chain responsiveness.

The game was programmed in Visual Basic to run off a client/server platform where each participant within a team works off a separate computer. See Figure 2 for a sample screen. The conditions of all three treatments were identical except for the amount of inventory information displayed. Inventory information was displayed in a bar chart with each bar representing the inventory position of a different supply chain member within the team (negative for backlogs). The chart was updated automatically at the beginning of each period. In treatment one (our baseline case) participants only saw their own inventory levels (i.e., one bar in the chart). This setting is identical to the setup used in the baseline study of Croson and Donohue (2004). In treatment two (upstream only) participants saw their own inventory level as well as the

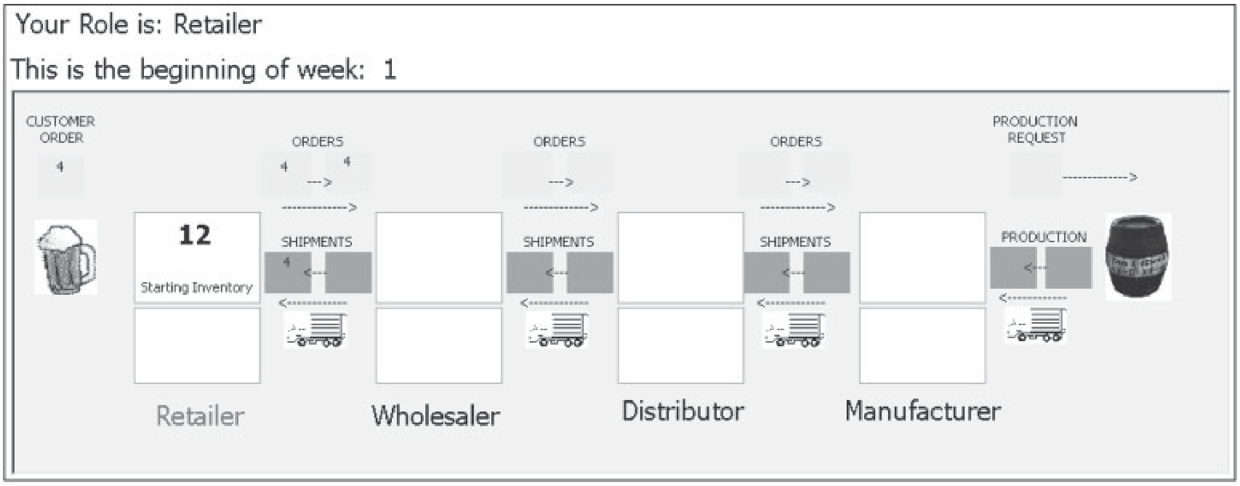

Fig. 2. Sample screen of the Beer Distribution Game

inventory levels of all upstream members (e.g., the distributor saw inventory levels for himself and the manufacturer). In treatment three (downstream only) participants had access to their own inventory level as well as the inventory levels of all downstream customers (e.g., the distributor now sees his own as well as the retailer and wholesaler inventory levels). In each treatment, our primary dependent measure is the variance of the orders placed by a given individual over the course of the multi-period game. We compare these variances for the different positions of the supply chain and conclude that the bullwhip effect exists if higher-echelon members exhibit higher variances.

Participants were drawn from Carlson MBA students enrolled in an Introduction to Operations Management course during the fall of 2003. Ninety-six students participated in total, with 28 students (seven teams) taking part in treatment 1, 36 students (nine teams) taking part in treatment 2, and 32 students (eight teams) taking part in treatment 3. The treatments ran back to back on the same day, and students were cautioned not to discuss the exercise with other participants.

Participants arrived in the computer lab at a predetermined time and were randomly assigned to a computer terminal, which determined their role and team assignment. Once seated, participants were oriented to the rules and objectives of the game. They were instructed that each role would incur unit holding costs of $0.50, unit backlog costs of $1, and unit revenues of $2 per period. They were also told that retail demand was uniformly distributed between 0 and 8 cases per period and independently drawn between periods.

Each echelon began with an initial inventory level of 12, outstanding orders of 4 for the last two periods, and an incoming shipment of 4 in the next two periods (see Figure 1). Participants were not informed how many periods the experiment would run to avoid end-of-game behavior that might trigger over- or under-ordering. The actual number of periods was 48 for all experiments. All experiments also used the same random number seed to generate demand. This allowed us to isolate variations due to ordering behavior from variations due to different demand streams.

To incentivize the participants to choose orders in a manner that maximizes their team's cumulative profit, we adopted the payment scheme of Croson and Donohue (2004). This scheme was announced before the start of the game. Each participant was given a base compensation of $5 with a possible bonus of up to $20. The bonus was computed as follows:

$$b^g = \$20 \frac{\pi^g - \min_g(\pi^g)}{\max_g(\pi^g) - \min_g(\pi^g)}$$

where π^g denotes total supply chain profit for team g (i.e., sum of profits for the four players) at the end of the game. Maximum and minimum profit levels were computed and compared separately within each treatment.

As noted by Croson and Donohue (2004), this experimental setting controls for all four of the operational causes of the bullwhip effect. Order batching is avoided since the setting has no fixed costs for ordering; thus there is no incentive to hold orders and submit them in one large chunk. Shortage gaming is not an issue since the manufacturer has no capacity constraint; thus there is no reason to over-order to guarantee you will receive a shipment. Price promotions also do not come into play since the price is fixed at $2 throughout the game; thus there is no incentive to over-order when the price is low or restrain ordering when the price is high. Finally, errors in demand signaling are avoided since the consumer demand distribution is commonly known by all players through its announcement at the beginning of the game, alleviating the need for order adjustments as you would find if observed demand conveyed information about future demand. Controlling for these operational factors allows us to focus our analysis on the behavioral impact of sharing inventory information.

Hypotheses and experimental results

Although Croson and Donohue (2004) speculate that downstream inventory information may be more beneficial than upstream, their paper does not provide evidence of this. One might conjecture that both types of inventory information (i.e., upstream and downstream) could be helpful in dampening the bullwhip effect. For upstream inventory, access to such information may provide a forewarning of when suppliers are running short of inventory and thus lessen a decision maker's tendency to overreact when the order he receives from his supplier falls short of his original order request. In this way, upstream information may increase the decision maker's connection between cause and effect by allowing him to learn more about his supplier's inventory management process. This leads to our first hypothesis.

Hypothesis 1: Sharing *upstream* inventory information within the supply chain will decrease the level of order oscillations throughout the supply chain relative to the baseline treatment.

On the other hand, downstream inventory information may offer other important benefits. First, having access to the retailer's inventory information gives all upstream decision makers a means to compute the quantity of consumer demand in each period.[1] This information, in theory, could be used to better gauge the inventory needs of the supply chain and react more quickly to random demand spikes and troughs. Second, by viewing all downstream inventory levels, a decision maker could base his order decisions on total echelon inventory (rather than local inventory position), which would lead to a more stable ordering pattern. The following hypothesis summarizes the implication of these claims.

Hypothesis 2: Sharing *downstream* inventory information within the supply chain will decrease the level of order oscillations throughout the supply chain relative to baseline treatment.

Croson and Donohue (2004) observed that upstream supply chain members (i.e., manufacturers and wholesalers) enjoy significantly more reduction in order oscillations than downstream members (i.e., retailers and wholesalers) when all inventory information (i.e., both upstream and downstream) is shared across the supply chain. It is interesting to consider whether upstream members continue to enjoy the lion's share of benefit when inventory information is shared in only one direction. This leads to our second set of contrasting hypotheses.

Hypothesis 3: Sharing *upstream* inventory information across the supply chain will lead to a greater reduction in order oscillations for manufacturers and distributors than for retailers and wholesalers relative to the baseline treatment.

Hypothesis 4: Sharing *downstream* inventory information across the supply chain will lead to a greater reduction in order oscillations for manufacturers and distributors than for retailers and wholesalers relative to the baseline treatment.

Impact on overall supply chain

Figure 3 displays the variance of orders placed for participants in each of the three treatments. It appears that order oscillations continue to be high under each

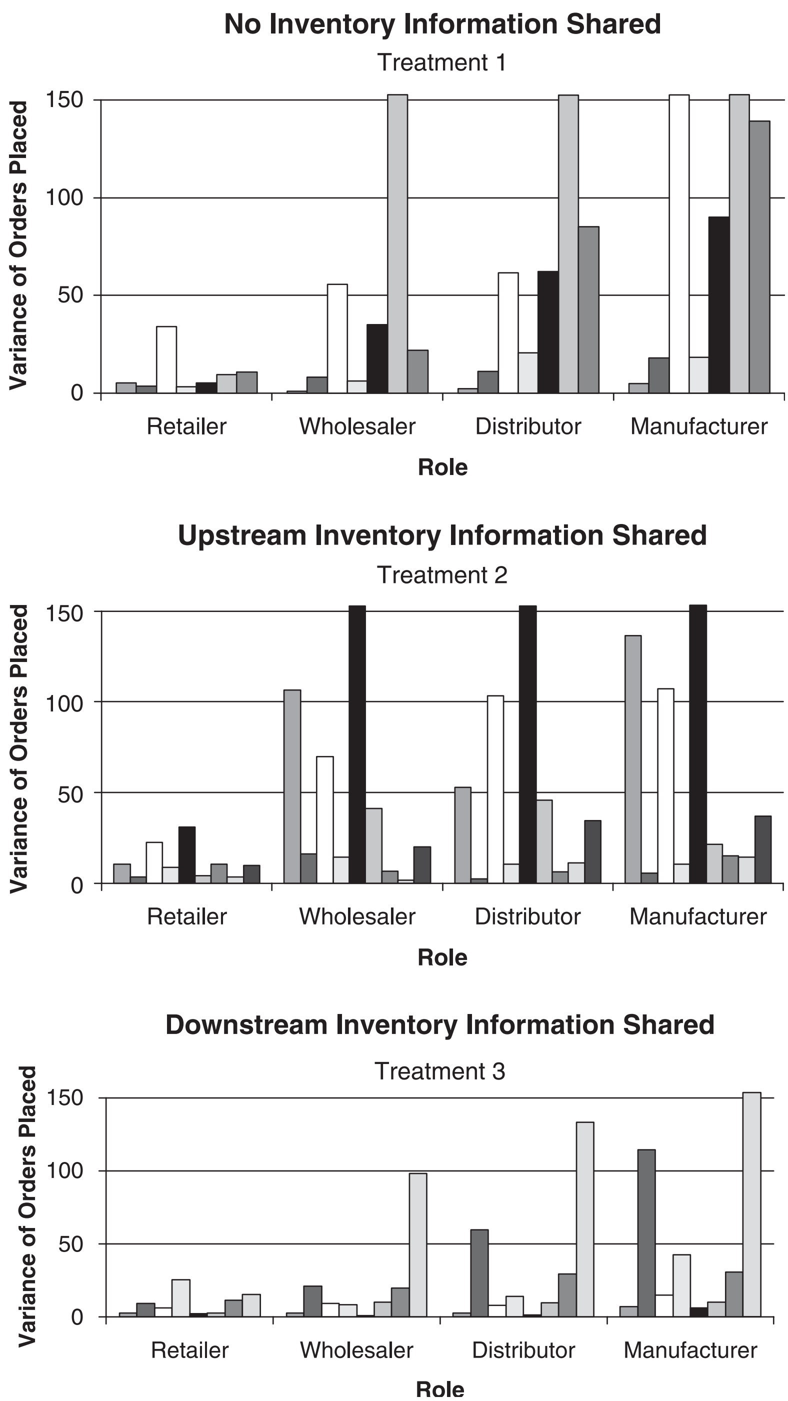

Fig. 3. Dependent measure: variance of orders placed over 48 periods

treatment, and that the magnitude of oscillations increases as one moves further upstream. In other words, the bullwhip effect appears alive and well. To test whether the order variance amplification is statistically significant, we use a simple non-parametric sign test (see Seigel 1965, p. 68). We code an increase in the variance of orders placed between each role (for each group) as a success and a decrease as a failure. The null hypothesis of no amplification suggests only a 50% success rate. All our treatments reject this hypothesis. Comparing immediate neighbors (i.e., retailer versus wholesaler, wholesaler versus distributor, etc.), we find a 90.5% success rate for the baseline treatment which is significantly different from 50% ($p = 0.0001$). Our second treatment (upstream information) yields a 66.7% success rate, which is also significantly different from 50% ($p = 0.0349$). Finally, our third treatment (downstream information) also illustrates the persistence of the bullwhip effect at 83.3% ($p = 0.0006$).[2]

While sharing inventory information does not eliminate the bullwhip effect, comparing the results of the treatments in Figure 1 it seems that inventory information may lead to some reduction in order oscillations. To test for this reduction, we use a non-parametric Wilcoxon test comparing the variance of orders placed over time. Since providing information should improve performance, we use one-tailed tests to compare the results of the baseline with each of the two treatments. The variances of orders placed in the baseline treatment have an average of 80.66, compared with 55.10 in the upstream information treatment. Information does appear to help in absolute terms; however, this improvement is not statistically significant ($n = 28$, $m = 36$, $U = 49$, $z = 0.81$, $p = 0.234$). This leads us to reject Hypothesis 1 and conclude that sharing only upstream inventory information offers little benefit in reducing bullwhip behavior.

In contrast, the variances of orders placed are significantly different between the first and third treatments. The average variance of orders placed when downstream inventory information is available is 30.10, which is statistically different from the 80.66 average of the baseline ($n = 28$, $m = 32$, $U = 357$, $z = 1.348$, $p = 0.044$). This implies that sharing only downstream inventory information offers significant benefit, in the form of decreased order oscillations. In fact, when comparing the second and third treatments we also find that sharing downstream inventory leads to significantly lower order oscillations than sharing upstream inventory information ($n = 36$, $m = 32$, $U = 444$, $z = 1.622$, $p = 0.026$). These results support Hypothesis 2.

Impact on upstream versus downstream members

We now turn to comparing the benefit that inventory information offers upstream versus downstream supply chain members. To perform these comparisons, we once again use a Wilcoxon test but now divide retailer/wholesaler and distributor/manufacturer pairs into two groups and compare each group separately across treatments.

Since our results in the previous section showed that upstream inventory information led to no significant improvement in order oscillations overall, it is not surprising that the retailer/wholesaler and distributor/manufacturing groups also experienced no significant improvement from upstream information. More specifically, in comparing treatments 1 and 2, we find no significant reduction in the variance of orders for the upstream ($n = 14$, $m = 18$, $z = 0.722$, $p = 0.147$) or downstream ($n = 14$, $m = 18$, $z = 0.714$, $p = 0.151$) players. This leads us to reject Hypothesis 3 since there is no detectable difference in impact between the two groups.

In contrast, the benefit of sharing downstream inventory information does vary by location. The manufacturer/distributor group enjoys a significant reduction in order oscillations ($n = 14$, $m = 16$, $z = 1.921$, $p = 0.026$), while the retailer/wholesaler group shows no significant difference ($n = 14$, $m = 16$, $z = 0.540$, $p = 0.167$). This confirms Hypothesis 4. The fact that upstream suppliers enjoy larger benefits is not surprising given that the difference in information available to upstream players is much greater across the two treatments (e.g., the retailer sees no additional information between treatments 3 and 1, while the manufacturing sees three new inventory levels). It is more surprising that downstream members do not modify their ordering behavior in reaction to their suppliers' more stable production patterns.

Discussion and conclusions

Our data clearly reveals that the impact of sharing inventory information varies by where such information resides. The conjecture stated in Croson and Donohue (2004) holds true. Sharing downstream inventory information is more effective at reducing bullwhip behavior than sharing similar upstream information. Furthermore, sharing only upstream information offers no significant performance improvement in our setting.

Table 1 summarizes the average variance, by role, for each study as well as the percentage improvement offered by the two information-sharing schemes. In column 5 we see that sharing upstream information actually increases average order oscillations for the retailer and wholesaler, relative to the baseline case. Although this performance decrease is not statistically significant, it does raise questions about how these downstream members are using this information. In informal post-game discussions with student participants, we found that some downstream members felt frustrated by not being able to counteract what they saw as poor decisions being made by their suppliers. They did not feel they could adjust their own ordering patterns to accommodate the suppliers' inventory levels (e.g., they did not want to decrease orders when their suppliers were out of stock since they felt this would only frustrate their own inventory position). A similar frustration was voiced by some of the manufacturers in treatment 3. In this case, the manufacturer has a bird's-eye

Table 1. Average variance of orders placed over 48 periods

Role	Average variance of orders			Improvement (%)	
	Treatment 1 No. info.	Treatment 2 upstream info.	Treatment 3 downstream info.	Impact of upstream info. (2 versus 1)	Impact of downstream info. (3 versus 1)
1. Retailer	10.19	11.54	9.42	–13.3%	7.6%
2. Wholesaler	43.37	59.30	21.35	–36.7%	50.8%
3. Distributor	87.42	56.17	32.31	35.8%	63.0%*
4. Manufacturer	181.66	93.38	57.32	48.6%	68.5%*

* $p < 0.05$.

view of inventory levels throughout the supply chain but could only help by adjusting their orders to counteract spikes in demand. They felt powerless to do anything about excess inventory in the system. One participant suggested he be allowed to institute Vendor Managed Inventory (VMI) so he could correct the "errors" he was seeing.

These informal observations suggest that upstream and/or downstream inventory information sharing would be more powerful if it came bundled with authorization. In other words, the transfer of decision rights may be necessary for the supply chain to fully exploit the benefits of inventory information sharing. This suggests a number of questions worth exploring in future research. For example, in a system where downstream information is shared, does moving toward VMI (where the manufacturer makes all order decisions) eliminate the bullwhip effect? Could the same improvement be obtained by simply allowing the supply chain members to communicate as a group (perhaps to discuss conflicting ordering strategies) at key points in time? Also, assuming authorization for order decisions will be passed to one supply chain member, who is in the best position to take on this responsibility (e.g., the retailer or manufacturer)? For now, our results suggest that companies are better served by investing in information systems which track inventory at the retail and wholesale levels. That said, the incentives in place in supply chains may make this difficult. In particular upstream members stand to gain the most but the information that needs to be shared resides downstream. This raises issues about how upstream members can incentivize downstream members to share their information in order to realize efficiency gains.

Our research also highlights how information can affect the decision making of individuals in dynamic situations. Information sharing can lead to frustration and somewhat worse performance in settings where one cannot fully act on the information. For example, the retailers' and wholesalers' order variation in the second (upstream only) treatment appears higher than in the baseline. Thus it is not the information *per se* but the interaction between the information and the decision setting that has the potential to improve performance in dynamic tasks. We look forward to seeing more research on these issues in the coming years.

Notes

1. Previous work has compared outcomes when information about the realization of demand (point-of-sale data) is directly shared with other members of the supply chain (e.g., Croson and Donohue 2003).
2. We also performed a sign test to detect order amplification between a player and all his downstream customers (e.g., manufacturer versus distributor, wholesaler and retailer). All comparisons were significantly different from 50% here as well. More specifically, we found 90.5% ($p = 0.0001$) for the baseline, 74.1% ($p = 0.0001$) for upstream inventory information, and 87.5% ($p = 0.0001$) for downstream inventory information.

References

Bourland K, Powell S, Pyke D. 1996. Exploiting timely demand information to reduce inventories. *European Journal of Operational Research* **92**: 239–253.

Cachon G, Fisher M. 2000. Supply chain inventory management and the value of shared information. *Management Science* **46**(8): 1032–1048.

Chen F. 1998. Echelon reorder points, installation reorder points, and the value of centralized demand information. *Management Science* **44**: S221–S234.

Croson R, Donohue K. 2003. The impact of POS data sharing on supply chain management: an experimental study. *Production and Operations Management* **12**: 1–11.

——, ——. 2004. Behavioral causes of the bullwhip and the observed value of inventory information. Working paper, Carlson School, University of Minnesota (forthcoming in *Management Science*).

Croson R, Donohue K, Katok E, Sterman J. 2004. Order stability in supply chains: coordination risk and the role of coordination stock. Working paper, Smeal College of Business, Penn State University.

Forrester J. 1958. Industrial dynamics: a major breakthrough for decision makers. *Harvard Business Review* **36**: 37–66.

Gavirneni S, Kapuscinski R, Tayur S. 1999. Value of information in capacitated supply chains. *Management Science* **45**: 16–24.

Lee H, Padmanabhan P, Whang S. 1997. Information distortion in a supply chain: the bullwhip effect. *Management Science* **43**: 546–558.

Seigel S. 1965. *Nonparametric Statistics for the Behavioral Sciences*. Wiley: New York.

Sterman JD. 1989. Modeling managerial behavior: misperceptions of feedback in a dynamic decision making experiment. *Management Science* **35**: 321–339.

——. 1992. Teaching takes off: flight simulators for management education. *OR/MS Today* **19**(5): 40–44.

——. 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Irwin–McGraw-Hill: New York.